Production Function Estimation in a Network*

Preliminary

Andrew B. Bernard
Tuck@Dartmouth
CEPR, CESiFo & NBER

Glenn Magerman
ECARES, Université libre de Bruxelles
CEPR & CESiFo

August 28, 2026

Abstract

A firm in a production network buys its inputs from other firms and sells its output to other firms. Its supplier base affects its marginal cost, and its customer base scales the demand it faces. Standard production-function estimators assume a single unobservable state, inputs bought at common prices, an input choice that inverts to reveal productivity, and survival governed by productivity alone. We consider a two-sided networked firm with general returns to scale in the variable inputs and characterize in closed form what estimation recovers when the data come from such a firm. The supplier network enters cost exactly where technology does, so measured productivity is technology compounded with supplier breadth and composition. The cost advantages of a firm's suppliers pass through prices into its own measured productivity. The customer network contaminates the investment proxy, measured persistence, and exit selection. These biases are general properties of networked production and the same firm-to-firm data can be used to correct them and the central correction needs no prices.

Keywords: production networks, customer capital, supplier capital, sourcing technology, productivity measurement

JEL: D24, D25, D85, E22, L11, L14.

*We thank ... for helpful comments. Monte Carlo results are produced by accompanying scripts and verify the closed-form identities to machine precision.

1 Introduction

The estimation of firm-level production functions and total factor productivity is one of the most consequential measurement exercises in economics, disciplining models of misallocation and reallocation, firm dynamics, knowledge spillovers, trade, and growth, among others. The estimated output elasticities directly deliver returns to scale and markups. Modern production-function and markup estimation rests on a compact view of the firm as an isolated producer buying and selling in anonymous markets. A scalar productivity term is the only unobservable, inputs and output are homogeneous and traded at common prices, a flexible input or investment inverts to reveal productivity, and exit, when modeled, is governed by productivity and capital.

This view underpins the proxy and control-function estimators [Olley and Pakes, 1996, Levinsohn and Petrin, 2003, Ackerberg et al., 2015], the gross-output approach [Gandhi et al., 2020], and the production approach to markups [De Loecker and Warzynski, 2012]. We call the compact view of the firm with these four assumptions – a single latent state variable, common prices, an invertible proxy, and survival on productivity and capital – the *standard* estimation problem.

These assumptions describe a firm that meets its suppliers and customers in undifferentiated anonymous markets. They fail, by construction, for a firm embedded in a production network. Such a firm carries two margins that are outside the standard problem: its supplier base, the number and identity of its input suppliers, and its customer base, the number and identity of its downstream customers. Both margins are accumulated relationship stocks, built with labor, slow to depreciate and on the same footing as physical capital. Both relationship stocks result from firm choices that are correlated with the firm's unobserved productivity, further complicating its identification. The margins cannot be treated as exogenous to the productivity the estimator sets out to recover. Moreover, unlike the conventional inputs the estimator nets out through their elasticities, they enter the productivity residual itself. This paper asks what standard estimation recovers when the data come from such a firm, based on the two-sided networked firm of Bernard [2026], and provides answers in closed form.

These stocks enter the firm's problem through separate channels: the supplier base enters through marginal cost and the customer base through demand. On the cost side, a firm's suppliers are firms charging specific prices, so the price of its input bundle is firm-specific. The input price is lower for firms with more suppliers and for firms with cheaper suppliers. A supplier's own cost advantage passes through, in turn, to its customers, and further down the production chain. On the demand side, each customer is a firm whose orders scale

the demand the seller faces. With decreasing returns in the variable inputs, an increase in customers raises scale and marginal cost, affecting revenue and the investment proxy. Each bias from the standard estimator is an explicit function of observable network objects, making the corresponding correction exact and constructible from the same network data. The mechanisms behind the biases, however, are more general than the CES structure that delivers these exact expressions.

The answers to our question start from one observation. A firm's supplier base lowers its marginal cost in exactly the same place that its own technology does. Cheaper or more numerous inputs and better technology are indistinguishable to an estimator that sees only quantities and common deflators. What the residual from the standard estimator actually measures is not technology but what we call 'effective productivity', technology compounded with the breadth and composition of the supplier base. A firm is mechanically measured as more productive simply for buying from more or cheaper suppliers.

Two example firms fix ideas and appear throughout the paper. The *artisan* is technologically excellent and sources only from a handful of suppliers. The *networker* is technologically ordinary and sources from many. The two have identical marginal cost and identical measured productivity. At conventional parameter values, a firm with twice the suppliers of an otherwise identical rival is measured as about twelve percent more productive with no difference in technology. Every result below can be read as an answer to one question: which of the two does the estimator call productive, and why?

We do not abandon the standard toolkit; instead, we generalize the structure of the underlying firm. The networked firm still holds physical capital, predetermined at production and accumulated by investment, so the entire proxy literature applies here without amendment. We add two relationship stocks and general returns to scale in the variable inputs, and we then ask the estimator's own questions of this richer firm. Doing so connects three literatures that have largely run in parallel: the productivity-estimation literature built on the compact firm [Olley and Pakes, 1996, Levinsohn and Petrin, 2003, Ackerberg et al., 2015, Gandhi et al., 2020]; the prices-and-markups critique, which showed that revenue-based residuals confound technology with output prices and markups [Klette and Griliches, 1996, De Loecker and Warzynski, 2012, De Loecker et al., 2016], with its structural treatment of unobserved output prices [De Loecker, 2011a] and its markup application [De Loecker et al., 2020]; and the production-network literature on how firm-level shocks propagate through the input–output structure [Hulten, 1978, Acemoglu et al., 2012, Carvalho, 2014, Baqaee and Farhi, 2019, 2020, Atalay, 2017]. The network is what ties them together. It is why input prices are firm-specific, and it is the channel along which one firm's costs become another firm's measured productivity.

Seen this way, the biases are not a list of unrelated problems but two families with a common origin. The static family is present for any networked firm before any dynamics. Because a firm's input price is firm-specific set by the number and identity of its suppliers, measured productivity absorbs its own supplier variety on the diagonal and its suppliers' cost advantages on the off-diagonal, carried downstream through a damped Leontief inverse exactly as own productivity is in Hulten [1978]. The same firm-specific price contaminates the intermediate coefficient as an errors-in-variables problem, attenuating returns to scale. The buyer-seller relationships are themselves heterogeneous, born small and deepening with tenure apart from buyer or seller characteristics. Link age needs to be correctly controlled for both in measures of firm prices and in measures of the relationship stocks.[1] Similarly, total labor at the firm contains both production labor and labor used for other activities that do not scale with output. Pooling both types of labor biases the labor elasticity. Revenue-based measures recover the customer base and markups rather than technology Klette and Griliches [1996] and value-added inflates every one of these biases by the intermediate multiplier.

The dynamic family follows from the stocks being state variables omitted by the standard estimator. Optimal investment depends on the marginal value of capital, which rises with scale, which in turn depends on both relationship stocks. Investment now prices three assets at once, so inverting it for a scalar productivity fails. Measured persistence is pulled away from the rate of technological change and toward the survival rate of relationships, because the residual inherits the smoothness of two slow-moving stocks. Exit depends on the two relationship stocks as well as physical capital, so survivors with thin networks are positively selected on productivity.

The constructive half of the paper is that the same firm-to-firm data that expose these biases also correct them, and the central correction needs no price data. The buyer–supplier panel already contains the network matrix from expenditure shares, the two relationship stocks as in- and out-degree, and the value of every transaction. From the values alone, a two-way seller–buyer decomposition of bilateral sales, the network analogue of the matched employer–employee methods of Abowd et al. [1999] implemented in Bernard et al. [2022], recovers each firm's relative output price from the seller effect and its exact input price index from the buyer effect. Deflation, which the literature has treated as a maintained assumption, becomes instead the first step of estimation.

Investment is therefore a function of the two relationship stocks alongside capital and productivity, so the proxy inversion of Olley and Pakes [1996] and its descendants is not

[1] An Appendix discusses the three capabilities (physical productivity φ, customer-acquisition efficiency χ, and supplier-acquisition efficiency ξ) that are drawn jointly at entry from a fixed distribution G and evolve exogenously around the drawn type.

one-to-one. Conditioning on the two observed stocks then restores the proxy inversion and the survival correction. In short, the contaminating prices and states are not unobservables to be modeled away; they are data, and the structural content migrates from the estimator to measurement.

The estimation guide is designed to match the likely available information set. We assume the researcher has a linked panel containing the firm-to-firm network plus firm-level data from a Census or balance sheets, so labor, capital, total intermediates, and total sales sit alongside the bilateral transactions. We outline two identification paths to deliver the parameters. Path 1 is a staged pipeline building the firm-specific prices from the network data alone, yielding effective productivity, and then moving to parameter estimation. Path 2 stacks the same information into one joint GMM system, asymptotically efficient, with the surplus moments delivered as specification tests. Both paths deliver the three technology elasticities. However, parameter scale must come from outside information on the substitution elasticity.

We simulate the network firm economy and ask what our estimation strategy recovers and where it does well and poorly over a range of parameter values. The labor and intermediate share parameters are recovered essentially exactly, the genuinely estimated parameters of the capital share and the productivity persistence are centered at their true values, and the recovered productivity series are strongly correlated with the true series. However, the network estimator only provides a diagnostic for the elasticity of substitution and when an incorrect external value is provided, the input elasticities move with it. Two standard estimators perform less well on the same data, attaining lower productivity series correlations, and generating negative capital coefficients and inflating the labor coefficients precisely where the network states carry the most variance.

The paper connects the productivity-estimation literature, the prices-and-markups critique, and the production-network literature on propagation and aggregation to the micro-evidence on heterogeneous buyer–seller relationships [Bernard et al., 2022, Oberfield, 2018, Demir et al., 2024, Dhyne et al., 2025, 2021]. Atalay [2014] documents that differences in the materials prices plants pay pass directly into measured productivity. The static layer in our framework gives that bias a network origin as a function of the supplier count and propagated supplier prices, as well as a correction.

Unobserved input-price and input-quality heterogeneity is a long-standing obstacle to estimation, [Katayama et al., 2009, Kugler and Verhoogen, 2012, De Loecker et al., 2016], met either with structure, first-order conditions that back the missing prices out [Grieco et al., 2016], or with auxiliary data, as in the imported-inputs estimates of Halpern et al. [2015]. The network panel makes it directly observable, so the missing input prices are recovered rather than modeled.

Doraszelski and Jaumandreu [2013] and Aw et al. [2011] argue that productivity is itself a controlled state, breaking the exogenous-Markov assumption. De Loecker [2013] makes the point with a demand-side control, export participation. The relationship stocks here are observed, demand- and cost-side counterparts. The customer base as a costly, sluggish, demand-shifting asset comes from Gourio and Rudanko [2014], Drozd and Nosal [2012], Foster et al. [2016], and Piveteau [2021]. Somaini and Einav [2013] price a customer stock and obtain a markup that moves with it. The acquisition technology that builds it here, costly reach with diminishing returns, follows the market-penetration costs of Arkolakis [2010]. Customer-facing labor is measured by He et al. [2024]. The distortions that omitted intangible capital imparts to measured productivity and markups are studied by Crouzet and Eberly [2019] and De Ridder [2024].

A related paper is Iyoha [2023] who incorporates productivity spillovers in a network into a control-function framework. We address a distinct but complementary question: whether network interactions can themselves be exploited to identify latent productivity in a fully interconnected production network.

On the estimator side, Norris Keiller et al. [2024] replace the proxy inversion with subjective expectations data, Utamaru [2026] develops identification that is robust to misspecified productivity dynamics, and Collard-Wexler and De Loecker [2016] treat measurement error in capital. Each relaxes a maintained assumption that this paper shows the network violates for reasons of its own.

The value-only identification adapts the matched-panel methods of Bernard et al. [2022], descended from the matched employer–employee literature [Abowd et al., 1999, Andrews et al., 2008, Kline et al., 2020]. The data-generating process throughout is the two-sided networked firm of Bernard [2026], which develops the firm's problem, prices the two relationship networks as accumulated capital, and closes the model in general equilibrium. This paper works with firm-specific (realized-network) prices, in which a firm's input price index loads on which suppliers it holds and on their propagated effective productivity, and asks what standard estimation recovers when the data come from such a firm.

The paper proceeds as follows. Section 2 sets out the two-sided networked firm and isolates the features that break estimation. Section 3 states what the econometrician assumes and defines the two information sets. Sections 4 and 5 develop the static and dynamic layers respectively, while Section 6 combines them. The implementable estimator, in its two identification paths, is laid out in Section 7. Monte Carlo results are presented in Section 8. Section 9 [TBD] presents results from production network data. Section 10 concludes.

2 The two-sided networked firm

The data-generating process is the firm of Bernard [2026], presented compactly. There are N firms, indexed i, observed over periods t (the time index is suppressed where it is not needed). Firm i has Hicks-neutral productivity φ_i ($\omega_i \equiv \ln \varphi_i$), physical capital K_i, a stock of customer capital n_i equal to its number of active downstream buyer links, and a stock of supplier capital v_i equal to its number of active upstream supplier links. Factor prices (the wage w and the investment-goods price p_I) and the substitution elasticity $\sigma > 1$ are common. A demand-capacity index $\bar{\psi}$, the demand level per customer, held common across firms at the baseline (Section 2.2 gives the general, firm-specific object), and the prices its suppliers charge summarize the network environment and are taken as given by the firm.

2.1 Technology, sourcing, and effective productivity

Gross output uses physical capital, production labor, and a CES intermediate bundle,

$$Q_i = \varphi_i \, K_i^{\alpha_k} (L_i^p)^{\alpha_l} M_i^{\alpha_m}, \qquad s \equiv \alpha_l + \alpha_m \leq 1, \qquad M_i = \Big[\textstyle\int_{\mathcal{S}_i} m_{ki}^{(\sigma-1)/\sigma} \, dk \Big]^{\frac{\sigma}{\sigma-1}}, \tag{1}$$

where L_i^p is production labor, $\mathcal{S}_i$ is the firm's supplier set, of measure v_i, and m_{ki} is the quantity i sources from supplier k at the price p_{ki}^M.[2] The bundle elasticity equals the demand elasticity mechanically: the firm's customers are themselves firms, and each buyer's derived demand for its supplier's variety carries the elasticity of its own bundle, so one σ governs both sides.

Capital and the two relationship stocks are predetermined at production. Production labor and the bundle are the variable inputs chosen at production, so their returns to scale s govern the short-run cost function, while $\eta \equiv \alpha_k + s$ is the total returns to scale across all three inputs.[3] Non-production labor produces no current output; it builds and maintains the two relationship stocks through the laws of motion (7)–(8) below, which raise output only through those stocks, v through inputs and n through demand. Cost minimization over the bundle gives the firm-specific input price index $P_i^M = (\int_{\mathcal{S}_i} (p_{ki}^M)^{1-\sigma} dk)^{1/(1-\sigma)}$ and the cost-minimizing expenditure share of supplier k,

$$a_{ki} = \frac{(p_{ki}^M)^{1-\sigma}}{\int_{\mathcal{S}_i} (p_{k'i}^M)^{1-\sigma} dk'}, \qquad \int a_{ki} \, dk = 1, \tag{2}$$

[2] The Cobb–Douglas upper tier keeps the output elasticities constant, so the returns to scale are parameters rather than functions of the input mix.

[3] With $s < 1$ marginal cost rises with output, $mc_i \propto Q_i^{(1-s)/s}$, and profit is concave in the customer stock. In the estimating equation the elasticities are unrestricted and η denotes their sum.

which defines the column-stochastic network matrix $A = [a_{ki}]$, built from expenditures. By Shephard's lemma the share a_{ki} is the elasticity of firm i's input price index with respect to supplier k's price, Hulten's theorem at the firm-pair level. Read backwards, observed shares reveal relative supplier prices from the structure, which is why the estimator of Section 7 never needs price data. The shares are the weights by which upstream costs propagate downstream. The input price index decomposes exactly into a variety (extensive) margin and a composition (intensive) margin. Let $\bar{p}_i^M \equiv \left(\frac{1}{v_i}\int_{\mathcal{S}_i}(p_{ki}^M)^{1-\sigma}dk\right)^{1/(1-\sigma)}$ denote the firm's average per-variety price, the power mean of order $1-\sigma<0$ of its realized supplier prices. Then

$$P_i^M = v_i^{1/(1-\sigma)}\,\bar{p}_i^M, \qquad \frac{\partial \ln P_i^M}{\partial \ln v_i}\Big|_{\bar{p}_i^M} = \frac{1}{1-\sigma} < 0. \tag{3}$$

The variety margin falls with the supplier count. The composition margin $\bar{p}_i^M$ is firm-specific because suppliers charge heterogeneous prices, so the index loads on which suppliers the firm holds. Holding the average price fixed, a firm whose suppliers charge more dispersed prices has a lower $\bar{p}_i^M$: substitution converts dispersion into savings, since any mean-preserving spread of prices raises $\mathbb{E}[p^{1-\sigma}]$ by Jensen's inequality. A common sectoral deflator removes neither margin. In log deviations the composition margin is, to first order, the expenditure-share-weighted average of the suppliers' relative prices (Section 4). The variety elasticity in (3), taken at fixed composition, does not depend on the returns to scale. Minimizing variable cost over (L_i^p, M_i) conditional on the predetermined K_i gives a short-run marginal cost that carries a scale term when $s<1$,

$$mc_i = \kappa_s\, Q_i^{(1-s)/s}\,\frac{w^{\alpha_l/s}(P_i^M)^{\alpha_m/s}}{(\varphi_i\, K_i^{\alpha_k})^{1/s}}, \qquad \kappa_s \equiv \alpha_l^{-\alpha_l/s}\alpha_m^{-\alpha_m/s}. \tag{4}$$

with κ_s the usual Cobb–Douglas cost constant. Substituting the decomposition (3), the supplier stock enters marginal cost in the same place as productivity, for all s. Writing

$$\boxed{\hat{\varphi}_i \equiv \varphi_i\, v_i^{\zeta}, \qquad \zeta \equiv \frac{\alpha_m}{\sigma-1} > 0,} \qquad mc_i = \kappa_s\, Q_i^{(1-s)/s}\,\frac{w^{\alpha_l/s}(\bar{p}_i^M)^{\alpha_m/s}}{(\hat{\varphi}_i\, K_i^{\alpha_k})^{1/s}}, \tag{5}$$

the cost shifter is $(\hat{\varphi}_i\, K_i^{\alpha_k})^{1/s}$, and the composition margin $\bar{p}_i^M$ enters where a factor price does, firm-specific through the identity of the suppliers.[4] Marginal cost moves with capital, with technology, with the variety of the supplier base, and with its composition, and the four shifters differ in observability. Capital is observed in standard balance-sheet data. The supplier count is observed in the network panel. The unobserved pair is technology

[4]The exponent ζ is the intermediate share times the sourcing elasticity. The CES sourcing technology fixes the sourcing elasticity at the constant $1/(\sigma-1)$. In Bernard [2026], the gain from one more supplier, the sourcing elasticity $\nu(v)$, may vary with the stock and $\zeta(v) = \alpha_m \nu(v)$.

and the composition margin: what estimation must disentangle is effective productivity $\hat{\varphi}_i$, technology compounded with the variety-deflated supplier stock, operating alongside the firm-specific $\bar{p}_i^M$, the average price its realized suppliers charge. The variety part of the conflation is the easier part, removable with the observed count. We show how to recover a firm-specific input price from a seller–buyer decomposition in Section 4.5. Both φ and v are predetermined when the firm produces, φ as production productivity (a firm-level capability) and v as an accumulated stock of sourcing relationships. v is distinctive in that it lowers cost through the extensive margin of sourcing, the number of input relationships.

Conditional input demands follow: $L_i^p = \alpha_l mc_i Q_i/w$ and $M_i = \alpha_m mc_i Q_i/P_i^M$. Because mc_i carries the scale term when $s < 1$, in reduced form the variable inputs scale as $Q_i^{1/s}$ rather than Q_i. Three features of (4) organize the paper: the scale term when $s < 1$, whose consequences are carried by the dampening scalar of the next subsection; the firm-specific input price index, the static layer of Section 4; and the cost shifter $\varphi_i K_i^{\alpha_k}$, the information-set boundary of Section 3.2.

A running example. The two firms of the introduction can now be made precise. They are identical on the demand side, the same customers and the same sales, facing the same average supplier price $\bar{p}^M$, and differ only in how they produce their cost advantage. The artisan is technologically excellent, with high φ, and sources from a handful of suppliers. The networker is technologically ordinary and sources from many. Its cost advantage is v^ζ, bought link by link with sourcing labor. The two have identical marginal cost, identical effective productivity $\hat{\varphi} = \varphi v^\zeta$, and identical measured productivity. Holding $\bar{p}^M$ equal across the two shuts off the supplier composition by construction. Every result below can be read as an answer to one question: which of the two does the estimator call productive, and why?

2.2 Demand, customers, the two relationship stocks, and the scale-dampening scalar

Each active customer's derived demand has constant elasticity σ. Aggregating over the firm's n_i buyers, quantity demanded is $Q_i^d = n_i \bar{\psi} p_i^{-\sigma}$, with $\bar{\psi}$ the demand-capacity index introduced at the opening of this section, so the customer stock shifts demand multiplicatively.[5] The firm sets the constant markup $\mu = p_i/mc_i \equiv \sigma/(\sigma-1)$.

[5]Customer j's capacity is $\vartheta_{ij}(P_j^M)^{\sigma-1}X_j$, set by its match quality, its own input price index, and its spending. The firm-specific index $\bar{\psi}_i \equiv n_i^{-1}\sum_j \vartheta_{ij}(P_j^M)^{\sigma-1}X_j$ is the average capacity of the firm's realized customers.

Start with $s = 1$. Marginal cost is then constant in output, the price is a closed-form function of fundamentals, the customer stock drops out of the price, per-customer operating profit π_i^* is constant, and $\Pi_i = \pi_i^* n_i$, exactly linear in the customer stock.

With $s < 1$ the scale term in (4) reconnects price to demand: mc_i depends on Q_i and Q_i on p_i, so $p_i = \mu\, mc_i$ becomes a fixed point. Its solution (Appendix E) rescales every exponent of the $s = 1$ case by the same factor, carried by the single scale-dampening scalar $D \equiv s + \sigma(1-s) = \sigma - (\sigma-1)s \geq 1$, with $D = 1$ exactly at $s = 1$ and $D - 1 = (\sigma-1)(1-s)$. $1/D$ is the dampening of the firm's short-run marginal-cost shifters in its price and the elasticity of revenue and operating profit with respect to the customer stock. Solving the fixed point also loads the firm's own customer stock onto its price, since more customers raise scale, raise marginal cost, and raise price, with elasticity $(1-s)/D > 0$. In reduced form, operating profit is a power function of the firm's states in which the demand side stands apart from the cost side,

$$\Pi_i = \Big(1 - \tfrac{s}{\mu}\Big) R_i = \bar{B} \Big[\underbrace{\bar{\psi}\, n_i}_{\text{demand side}} \cdot \big(\underbrace{\varphi_i v_i^{\zeta}\, K_i^{\alpha_k}\, (\bar{p}_i^{M})^{-\alpha_m}}_{\text{cost side}} \big)^{\sigma-1} \Big]^{1/D}, \tag{6}$$

with $\bar{B}$ a constant in $(w, \mu, \alpha_l, \alpha_m)$. The demand-side is the firm's market scale, its customer count times the per-customer demand capacity. The cost-side is the cost shifter $\hat{\varphi}_i\, K_i^{\alpha_k}$ of (5) together with the composition margin. This includes technology, compounded with capital, with the variety of the supplier base through v_i^{ζ}, and with the prices its realized suppliers charge through $(\bar{p}_i^{M})^{-\alpha_m}$. The cost side is levered by $\sigma - 1$, because a lower marginal cost earns a higher margin on every unit sold.

Every estimation result below sorts into the same two channels: the supplier stock and the supplier prices contaminate measurement through cost, the customer stock through demand. At $D = 1$ profit is linear in the customer stock, $\Pi_i = \pi_i^* n_i$ with the per-customer profit $\pi_i^* \propto \bar{\psi} \left(\varphi_i v_i^{\zeta} K_i^{\alpha_k} (\bar{p}_i^{M})^{-\alpha_m}\right)^{\sigma-1}$ set entirely by the cost side; for $D > 1$ profit is concave in customers, because a larger customer base also raises marginal cost through scale.

Returns to scale, short run and long run. The passage from decreasing variable returns to rising marginal cost uses only cost minimization over the variable inputs at given input prices with a homothetic technology. The variable cost function is then homogeneous of degree $1/s$ in output and $mc_i \propto Q_i^{(1-s)/s}$, a statement a multi-product cost function would not deliver without further restrictions on joint costs. Everything is short run: capital is predetermined, so s, marginal cost, and D are conditional-on-K objects. The markup is constant, and a cost change passes into price completely after accounting for the induced output adjustment. If $1/D < 1$, the output expansion raises marginal cost and absorbs part

of any cost shifter.

The two relationship stocks are capital. The problem is dynamic: the firm invests today in stocks that produce tomorrow, and the full forward-looking program is stated in Section 2.3. A link fails for reasons outside the firm's control at an exogenous rate, and the firm invests to replace and grow the stock. They evolve through depreciation and acquisition,

$$n_{i,t+1} = (1-\delta)\, n_{it} + \chi_{it}(L^n_{it})^{\gamma_n}, \tag{7}$$

$$v_{i,t+1} = (1-\delta)\, v_{it} + \xi_{it}(L^v_{it})^{\gamma_v}, \tag{8}$$

with acquisition labor L^n, L^v, acquisition efficiencies χ, ξ, a common exogenous link-failure rate $\delta \in (0,1)$, and diminishing acquisition returns $0 < \gamma_n, \gamma_v < 1$. The firm does not choose how long a relationship lasts, only how many new ones to form. The failure rate is common to the two sides because a link's failure is one event, removing a customer from its seller and a supplier from its buyer. Because relationships are heterogeneous, a firm's two stocks can churn at different rates through the composition of their link sizes. Servicing the installed base is a technological requirement. Maintaining a customer (supplier) link costs a fixed amount of labor c_n (c_v) per period, so servicing labor is proportional to the stocks. Total firm labor splits into three uses,

$$L_i = \underbrace{L^p_i}_{\text{production}} + \underbrace{L^n_i + c_n\, n_i}_{\text{customer-facing}} + \underbrace{L^v_i + c_v\, v_i}_{\text{sourcing}}, \tag{9}$$

of which only L^p_i enters the production function (1). Within non-production, the servicing labor component tracks the slow-moving stocks of links and the acquisition labor component moves with the new link flows.[6] Physical capital accumulates with investment subject to the convex adjustment cost $\Theta(I,K)$, $K_{i,t+1} = (1-\delta_K)K_{it} + I_{it}$, with $\delta_K \in (0,1)$ the capital depreciation rate.

Relationships within a stock are heterogeneous. A link is born small, deepens with tenure, and fails at a size-dependent rate, so a large flow of small links churns around a durable, valuable core [Martin et al., 2026, Fontaine et al., 2023, Huneeus, 2018]. Nothing below restricts the tenure profile to be monotone.[7] The same process generates a cross-sectional distribution of link sizes, right-skewed with a thin tail of old, deep links. Calculating the correct input price index (3) requires a correction for the age of each link. The input price

[6] Real firms also employ workers in other activities, e.g. administration, finance, and accounting, which may be time and scale-invariant. Overhead labor changes no first-order condition. As the fixed cost of carrying the firm it opens the exit margin of Section 5.2, and Section 7.2 adds it to the estimation guide.

[7] Link size z_{ik} is the value of the transaction. A link is born at the match size $\bar{z}$ and follows the tenure profile thereafter.

index falls with the number of supplier varieties, and the marginal profit of a customer or a supplier in (12) is the value of one more link.

Link heterogeneity affects the firm's value. A relationship contributes to demand in proportion to its size. The firm's value, and through it the investment policy and the exit margin of Section 5.2, therefore load on the total size of a portfolio rather than the number of links. The value-relevant state is the *value-weighted stock* $n_i^\star = \sum_k z_{ik}$, the within-firm sum of link sizes, and symmetrically $v_i^\star$ upstream, read straight off the panel as the firm's sales downstream and its intermediate purchases upstream. The count n_i is a noisy proxy: projecting the portfolio on the count gives $n_i^\star = \beta_\star n_i + u_i^\star$, with $\beta_\star$ the mean link size and $u_i^\star$ the within-firm size dispersion, and because small links fail first $n^\star$ is stickier than the count. The investment proxy and the exit correction therefore condition on $(n^\star, v^\star)$; the cost-side variety term and the per-link marginal profits stay in counts.

2.3 The firm's problem and the estimation-relevant policies

The firm carries three investable states (K, n, v) and the exogenous capabilities (φ, χ, ξ). Each period it decides whether to continue, and a continuing firm chooses investment and the two acquisition labor inputs, pricing and the flexible inputs being already folded into the operating profit Π of (6), discounting at the factor $\beta \in (0,1)$,

$$V(K,n,v;\varphi,\chi,\xi) = \max\Big\{\underline{V},\ \max_{I,L^n,L^v}\Big[\,\Pi - w(c_n\,n + c_v\,v) - w(L^n+L^v) - p_I I - \Theta(I,K) \\ + \beta\,\mathbb{E}\big[V(K',n',v';\varphi',\chi',\xi')\big]\Big]\Big\}, \tag{10}$$

subject to (7)–(8) and the capital law of motion, with $\underline{V}$ the value of exiting.[8] The exit option binds only once the firm pays a strictly positive carrying cost. The fixed overhead cost that opens the margin, and the exogenous death hazard δ_f, are introduced in Section 5.2, and for a continuing firm every first-order condition below is unchanged.

The capabilities are exogenous to this problem. Each firm draws a permanent capability type $(\bar\varphi, \bar\chi, \bar\xi)$ once, at entry, from a fixed distribution G, and its realized capabilities evolve as exogenous AR(1) processes in logs around that type, exactly as productivity is treated in the canonical production function. The acquisition capabilities belong to the state vector because the firm knows (χ_t, ξ_t) when it hires acquisition labor. The costates are customer lifetime value $\lambda^n = V_n \equiv \partial V/\partial n$, supplier lifetime value $\lambda^v = V_v \equiv \partial V/\partial v$, and Tobin's

[8] The environment the firm takes as given, the demand index $\bar\psi$ and the supplier prices behind P^M, is stationary from the firm's viewpoint and is suppressed from the state vector. Appendix C adds the firm-specific input price to the states where matching makes it payoff-relevant.

$q = V_K \equiv \partial V/\partial K$. The first-order conditions for the three kinds of investment are

$$\begin{aligned} (customers) \qquad & w = \beta\, \chi_t \gamma_n (L_t^n)^{\gamma_n - 1}\, \mathbb{E}_t\big[\lambda_{t+1}^n\big] \\ (suppliers) \qquad & w = \beta\, \xi_t \gamma_v (L_t^v)^{\gamma_v - 1}\, \mathbb{E}_t\big[\lambda_{t+1}^v\big] \\ (physical\ capital) \qquad & p_I + \Theta_I(I_t, K_t) = \beta\, \mathbb{E}_t\big[q_{t+1}\big]. \end{aligned} \tag{11}$$

Each equates the marginal cost of a unit of next-period capital, a link or a machine, to its discounted shadow value. Differentiating the reduced-form profit (6), holding composition fixed so that the marginal supplier is priced at the firm's average, gives the marginal profits in closed form,

$$\Pi_K \equiv \frac{\partial \Pi}{\partial K} = \frac{\alpha_k(\sigma - 1)}{D}\, \frac{\Pi_i}{K_i}, \qquad \Pi_n \equiv \frac{\partial \Pi}{\partial n} = \frac{1}{D}\, \frac{\Pi_i}{n_i}, \qquad \Pi_v \equiv \frac{\partial \Pi}{\partial v} = \frac{\alpha_m}{D}\, \frac{\Pi_i}{v_i}. \tag{12}$$

One more customer scales the demand side, with the exponent $1/D$, while capital and the supplier stock work through the cost side, their exponents levered by $\sigma - 1$. The cross-partials keep their sign and are damped by $1/D^2$, $\Pi_{Kn} = \frac{\alpha_k(\sigma-1)}{D^2}\Pi_i/(K_i n_i) > 0$ and $\Pi_{Kv} = \frac{\alpha_k \alpha_m(\sigma-1)}{D^2}\Pi_i/(K_i v_i) > 0$.

Applying the envelope theorem to (10) with $\partial n'/\partial n = 1 - \delta$ gives $\lambda_t^n = \Pi_{n,t} - w c_n + \beta(1-\delta)\, \mathbb{E}_t[\lambda_{t+1}^n]$, which iterates to

$$\lambda_t^n = \mathbb{E}_t \sum_{j \geq 0} \big[\beta(1-\delta)\big]^j \big(\Pi_{n,t+j} - w c_n\big), \tag{13}$$

the expected discounted stream of marginal profits from one more customer, net of servicing, over the link's stochastic lifetime, with the survival-adjusted discount factor $\beta(1-\delta)$. This is the customer-lifetime-value object of the marketing literature and of the customer-capital models of Gourio and Rudanko [2014] and Drozd and Nosal [2012], and symmetrically λ^v is supplier lifetime value.

Two features of the solution break standard estimation. The first is the labor composite. Total labor (9) mixes production, customer-facing, and sourcing uses, and only production labor enters the technology (1). The second is the investment policy, which prices three assets at once. Solving (10) delivers the policies that matter for estimation. Production labor and the intermediate bundle are proportional to output and to marginal cost (4)–(5). Acquisition labor is set by the two lifetime values and the wage through (11), and servicing labor is proportional to the stocks. Investment depends on every state the firm carries,

$$I_i = \mathcal{I}(K_i, n_i, v_i, \varphi_i). \tag{14}$$

Optimal investment equates the marginal cost of capital to its shadow value $\beta\mathbb{E}_t[q_{t+1}]$. By $\Pi_{Kn} > 0$, that shadow value rises with the customer stock (physical and customer capital are complements). Through effective productivity (5) and $\Pi_{Kv} > 0$, the shadow value rises with the supplier stock (more suppliers lower cost and raise the scale at which capital operates). The acquisition capabilities enter (14) only through the continuation values and are held at their types here. Equation (9) makes observed labor a composite of three uses, and equation (14) makes investment carry three states besides K: the two relationship stocks (n, v) and productivity φ.[9]

3 The econometrician's problem

This section sets up the econometrician's problem twice: first the standard one, with standard firm-level data and the maintained assumptions under which the estimators cited above operate, and then the networked one with the additional firm-to-firm network information.

3.1 The standard setup

A researcher using standard (balance sheet) data observes, for each firm and period, revenue R_i, a labor quantity L_i and its wage bill, capital K_i, investment I_i, and intermediate expenditure $E_i^M = P_i^M M_i$.[10] Nominal magnitudes are deflated by common industry indices, $\bar{P}$ for output and $\bar{P}^M$ for intermediates. The researcher works with

$$y_i \equiv \ln \frac{R_i}{\bar{P}} = q_i + \tilde{p}_i, \qquad x_i \equiv \ln \frac{E_i^M}{\bar{P}^M} = m_i + \widetilde{P^M}_i, \tag{15}$$

where $q_i = \ln Q_i$, $m_i = \ln M_i$ are the true (unobserved) physical log quantities and $\tilde{p}_i, \widetilde{P^M}_i$ are the firm-specific relative output and input prices (deviations from the common deflators). A tilde denotes a log deviation from the cross-sectional mean throughout, the unweighted arithmetic mean of logs within the year, taken within industry wherever the deflators are industry-specific. The labor regressor is $\ell_i \equiv \ln L_i$, and $k_i \equiv \ln K_i$ is log capital. The maintained model is a scalar-productivity gross-output function

$$y_i = \beta_k k_i + \beta_l \ell_i + \beta_m x_i + \hat{\omega}_i + \varepsilon_i, \tag{16}$$

[9] To be precise, the baseline policy (14) suppresses the variation of the acquisition capabilities around their types. Letting (χ_t, ξ_t) vary would make it five states besides K, and the footnote to Proposition 5 handles that case. The composition margin $\bar{p}_i^M$ is likewise payoff-relevant through (6); the baseline holds it at its cross-sectional mean, and Appendix C adds its firm-specific variaation to the conditioning sets.

[10] Data sources might include Census or administrative firm panels, balance sheets, or annual accounts.

(ε_i i.i.d. measurement error) estimated by any standard method (OLS, a proxy/control function for $\hat{\omega}$, the gross-output share approach, or an instrumented revenue function) under Assumption 1. $\hat{\omega}_i$ is the econometrician's recovered residual, while the model's effective productivity $\hat{\varphi}_i$ of (5) is a primitive of the data-generating process. When, and in what form, the first recovers the second is the subject of Section 4.

Assumption 1 (Standard maintained assumptions)**.**

(i) Scalar unobservable: productivity ω_i is the only unobserved state besides observed capital k_i, and it enters Hicks-neutrally.

(ii) Exogenous Markov: ω_t follows an exogenous first-order Markov process, $\omega_t = \mathbb{E}[\omega_t \mid \omega_{t-1}] + \epsilon_t^{\omega}$.[11]

(iii) Scalar-unobservable invertibility: a proxy (investment $i_t = i(k_t, \omega_t)$ or intermediate demand $m_t = m(k_t, \omega_t)$) is monotone in ω_t given k_t, so $\omega_t = h(k_t, \text{proxy}_t)$.

(iv) Homogeneous input at a common price: the intermediate input is homogeneous and priced at $\bar{P}^M$, so x_i measures m_i.

(v) Output measurement: output is physical, or y_i measures q_i up to a common deflator.

(vi) Labor: ℓ_i is the production labor of (1).

The two-sided firm violates (i), (ii), and (iii) through the two relationship stocks. It violates (iv) and (v) mechanically through firm-specific prices, the input-price heterogeneity familiar from De Loecker et al. [2016], though here the prices carry the supplier stock and the network rather than unobserved quality. It violates (vi) because the observed labor quantity is, in the data-generating process, the three-use composite (9) rather than the production input L_i^p.[12] Section 4 signs the static (price) consequences and Section 5 the dynamic (state) consequences.

Assumption 1(i) also builds in Hicks neutrality: one scalar shifts the whole isoquant map. The networked firm suggests where that fails, since the two acquisition capabilities are factor-specific productivities for the two relationship technologies. Demirer [2026] identifies factor-augmenting technology directly and finds that imposing Hicks neutrality distorts estimated markups. The corrections below stay within a scalar production state.

[11] The innovation is written ϵ_t^{ω}, keeping ξ for the supplier-acquisition capability of (8) and ε for measurement error.

[12] In the data-generating process productivity itself is exogenous Markov. Each capability evolves exogenously around a type the firm sets at entry. Assumption (ii) fails not for ω but for the residual that estimation recovers, which mixes in the two controlled, slow-moving relationship stocks. The violations are entirely the network's.

3.2 Two information sets

The corrections in this paper are constructed from data rather than from added structure. What can be constructed depends on which datasets are linked, and we consider two nested information sets. The first is the *network-only* set $\mathcal{I}_1$, the firm-to-firm panel alone, including the bilateral transaction values, the matrix A, the two link counts, link-level sizes and tenures, and network intermediate expenditure. Linking a standard firm panel of annual accounts adds labor (employment and the wage bill), capital K_i, investment I_i, total intermediate purchases and total sales including final demand, giving the *balance-sheet-linked* set $\mathcal{I}_2$. Labor by occupation, the split of total labor into production, customer-facing, and sourcing uses, is in the *occupational* set $\mathcal{I}_3$, less readily available.

The default throughout the paper is $\mathcal{I}_2$. The guide assumes the researcher holds the network panel and the balance sheets, and it does not assume occupational labor. The assignment matters because the production-function input is production labor L_i^p, not total labor. Under $\mathcal{I}_3$, L_i^p is observed. Under $\mathcal{I}_2$ it is constructed by the structural split of Section 7, and using total labor generates the mismeasurement analyzed in Section 5.4. The labor revenue share, the labor elasticity α_l, the returns to scale $s = \alpha_l + \alpha_m$, and the dampening scalar D are all $\mathcal{I}_2$ objects.

Observed capital is Assumption 1(i) itself and without it capital cannot be separated from technology, so the supply-side object recoverable under $\mathcal{I}_1$ is the *capital-inclusive cost fundamental*

$$\omega_i^c \equiv \omega_i + \alpha_k \tilde{k}_i \tag{17}$$

(log capital in deviations, ω_i read in the same deviations), technology compounded with capital intensity.

3.3 Generality

The results below are derived for one data-generating process in order to give closed form to a general set of estimation issues for networked firms. Here we briefly discuss the generality of the biases.

None of these biases is a property of the constant-elasticity sourcing technology. Five biases follow from the fact that the firm buys from and sells to other firms: input-price heterogeneity, propagation along the input–output matrix, the generated regressor in the deflated intermediate, the extra states in the proxy, and selection on more than productivity. The two margins enter asymmetrically, the supplier base through cost and the customer base through demand, but both are omitted states, so measured cost, the investment proxy, and exit each depend on network states that an anonymous-firm estimator omits. Any technology

in which the supplier base lowers the input price index carries the first two failures, input-price heterogeneity and propagation [Atalay, 2014, Antras et al., 2017, Halpern et al., 2015, Hulten, 1978, Baqaee and Farhi, 2019]. The remaining three, the generated regressor, the extra proxy states, and exit selection, are about the count of payoff-relevant states and the exit margin, not the functional form. The logic also extends to any accumulated input bought at a firm-specific price: heterogeneous capital types acquired at different prices, as under a financing pecking order, enter measured productivity through the same channel the supplier stock does, accumulated cost advantage indistinguishable from technology without the corresponding price data.

The constant-elasticity structure given below buys the exact expression for each bias. The variety term in the input price index, the single scalar D that carries the returns-to-scale generality, the propagation operator $(\alpha_m/D)A^\top$, the two-way fixed-effects deflation that recovers the input price index exactly, and the share identities that read the elasticities off revenue shares are all closed forms of the constant-elasticity model.

A different model of the networked firm would keep the mechanisms and the direction of the biases. Under a non-CES or translog technology the biases keep their sign and lose their exact form. The input price index is no longer variety plus composition, and the two-way decomposition of bilateral sales recovers a first-order approximation to the firm's input price rather than the exact index, so the share identities hold locally rather than globally.

When relationships follow different dynamics, the sign of the persistence bias, measured persistence pulled toward the relationship-survival rate, survives as long as the stocks are slow-moving. The value-weighted-stock point of Section 2 is the constant-hazard instance of that general point.

4 The static layer: variety, prices, and propagation

We first hold the network and the two stocks fixed and study the cross-section of prices. This section isolates the price layer assuming that log production labor $\ell_i^p \equiv \ln L_i^p$ is the labor regressor, or equivalently that production labor can be isolated from total labor. Substituting the truth $q_i = \beta_k k_i + \beta_l \ell_i^p + \alpha_m m_i + \omega_i$ and (15) into (16), the residual the researcher recovers at true slope coefficients is

$$\hat{\omega}_i = \omega_i + \underbrace{\tilde{p}_i}_{\text{output-price}} - \underbrace{\alpha_m \widetilde{P^M}_i}_{\text{input-price}}. \tag{18}$$

This accounting identity holds whatever the returns to scale. What changes with s is the content of the two price wedges, which we now make explicit.

The paper's focus is the *realized network*, in which it matters who a firm's suppliers are. A firm that happens to buy from cheap, productive suppliers, suppliers priced below the industry-year mean, has cheap inputs, and its suppliers' advantages pass through prices to it, and onward downstream.

Normalizations and what is identified. All price and productivity objects in this section are read in log deviations from cross-sectional year means, within industry where deflators are industry-specific. The deviations absorb the aggregate price normalization, the common deflators, and every constant parameter, so no separate numeraire needs to be chosen, and stacked expressions such as $A^\top \tilde{\boldsymbol{p}}$ are linear maps of the deviation vectors. The recovery is correspondingly relative. Rankings, dispersion, and comovement along the network are identified; the level of aggregate TFP and its growth must come from the deflators themselves, which the corrections take as given. Capital rides along throughout: with predetermined capital every formula below holds with the cost fundamental $\omega_i^c = \omega_i + \alpha_k \tilde{k}_i$ of (17) in place of ω_i.

4.1 The price system

Taking logs of the markup and marginal cost (4) and solving the pricing fixed point (Appendix E) gives the relative output price

$$\tilde{p}_i = \frac{1}{D}\big(\alpha_m \widetilde{P^M}_i - \omega_i^c\big) + \frac{1-s}{D}\widetilde{\ln n}_i, \tag{19}$$

with ω_i^c the cost fundamental of (17). Cost shifters enter price dampened by the scale feedback, $1/D < 1$, and the last term loads the firm's own customer stock onto its price through scale. Focusing on unobservables, this subsection sets $\alpha_k = 0$, so $\omega_i^c = \omega_i$. Every formula holds in general with ω read as ω^c.[13] The input price index has the two margins of (3) and is, to first order in log deviations,

$$\widetilde{P^M}_i = \underbrace{-\tfrac{1}{\sigma-1}\widetilde{\ln v}_i}_{\text{variety (own } v_i)} + \underbrace{(A^\top \tilde{\boldsymbol{p}})_i}_{\text{composition (suppliers' prices)}}. \tag{20}$$

The first term is the own-variety gain from (3) (more suppliers, lower index) and the second is the composition margin $\bar{p}_i^M$ in deviations, to first order the expenditure-share-weighted average of the realized suppliers' relative prices. Both are unchanged by the returns to scale

[13] Capital is observed but its elasticity α_k is not. It is estimated from the capital timing moment of Section 7, and until then the capital term stays bundled inside the cost fundamental.

since the bundle is unaffected by s. The composition margin is where propagation lives: in equilibrium the suppliers' relative prices carry their own effective productivities, so every off-diagonal term below is the composition margin written out. Substituting (20) into (19) and stacking,

$$\Big(\mathbb{I}-\tfrac{\alpha_m}{D}A^\top\Big)\tilde{\boldsymbol{p}} = -\tfrac{1}{D}\big(\boldsymbol{\omega}+\zeta\,\widetilde{\ln \boldsymbol{v}}-(1-s)\widetilde{\ln \boldsymbol{n}}\big), \qquad \tilde{\boldsymbol{p}} = -\tfrac{1}{D}\,\Psi_s\big(\boldsymbol{\omega}+\zeta\,\widetilde{\ln \boldsymbol{v}}-(1-s)\widetilde{\ln \boldsymbol{n}}\big), \tag{21}$$

using $\alpha_m \cdot \frac{1}{\sigma-1} = \zeta$ to collapse the variety term, where $\Psi_s \equiv (\mathbb{I} - \frac{\alpha_m}{D}A^\top)^{-1} = \sum_{t\geq 0}(\frac{\alpha_m}{D}A^\top)^t$ is the upstream Leontief inverse with the propagation coefficient damped from α_m to α_m/D (well defined since $\alpha_m/D < 1$ and A is column-stochastic). At $D = 1$ this is the standard network price system, the dual (price) form of Hulten's theorem [Hulten, 1978, Baqaee and Farhi, 2019], with own productivity $\boldsymbol{\omega}$ replaced everywhere by effective productivity $\boldsymbol{\omega}+\zeta\widetilde{\ln \boldsymbol{v}}$. With $D > 1$, the propagation is damped by the scale feedback and the firm's own customer stock enters with coefficient $(1 - s)/D$.

Capital and the two information sets. With the predetermined capital of Section 2 in marginal cost (4), the price system (21) holds with $\boldsymbol{\omega}$ replaced by the cost fundamental $\boldsymbol{\omega}^{\boldsymbol{c}} = \boldsymbol{\omega} + \alpha_k\tilde{\boldsymbol{k}}$ of (17). In the realized-network case measured productivity gains one further component, upstream capital intensity: $\hat{\boldsymbol{\omega}} = \Psi_s(\boldsymbol{\omega} + \zeta\widetilde{\ln \boldsymbol{v}}) - (1 - s)(\Psi_s - \mathbb{I})\widetilde{\ln \boldsymbol{n}} + (\Psi_s - \mathbb{I})\,\alpha_k\tilde{\boldsymbol{k}}$, so a firm is also measured as productive when its suppliers are capital-deep. The two information sets of Section 3.2 then split exactly here. Under $\mathcal{I}_2$ the capital quantity is data and its elasticity is an estimate, $\hat{\alpha}_k$ from the timing moment of Section 7, so each correction subtracts the measured term $\hat{\alpha}_k\tilde{k}_i$, isolating technology ω. Under $\mathcal{I}_1$ the same variety and propagation inversions are feasible, with ω^c everywhere in place of ω. They deliver the firm's supply-side fundamental, technology compounded with capital intensity, which cannot be split further without a balance sheet. In the revenue case the cancellation of Proposition 2 leaves $-(1-1/D)\alpha_k\tilde{k}_i$, absorbed by the included capital regressor under $\mathcal{I}_2$, so the technology statement is unchanged. The propositions below keep the $\alpha_k = 0$ convention; under $\mathcal{I}_2$ they are exact after the capital subtraction, and under $\mathcal{I}_1$ they hold with ω^c in place of ω.

4.2 The effective-productivity bias, nested

A firm's measured productivity absorbs the inputs its network delivers. It picks up the firm's own supplier variety on the diagonal, and it picks up the costs of the firms upstream of it, carried down the network on the off-diagonal. Suppose first that the econometrician observes perfect quantity data, output physically measured at the firm level; the assumption is a benchmark, relaxed immediately below.

Proposition 1 (Network propagation of measured effective productivity)**.** *When the data-generating process is the networked firm of Section 2, intermediate expenditure is deflated by the common index $\bar{P}^M$, and output is physically measured, the recovered productivity is the network-propagated effective productivity,*

$$\hat{\boldsymbol{\omega}} = \Psi_s \left(\boldsymbol{\omega} + \zeta \widetilde{\ln \boldsymbol{v}}\right) - (1-s)\left(\Psi_s - \mathbb{I}\right) \widetilde{\ln \boldsymbol{n}}, \qquad \Psi_s = \left(\mathbb{I} - \tfrac{\alpha_m}{D} A^\top\right)^{-1}. \tag{22}$$

(Realized-network case.) *With heterogeneous realized supplier prices the damped operator carries the effective productivity of a firm's upstream suppliers downstream to it. A supplier's cost advantage passes through to its customers' input prices, so measured productivity loads on the effective productivity of all firms upstream of the firm. It is additionally measured as productive when its suppliers have thin customer bases (and so price low), the $-(1-s)(\Psi_s - \mathbb{I})\widetilde{\ln \boldsymbol{n}}$ term, which is purely off-diagonal. The customer-stock contamination of measured productivity under physical output is therefore a network-propagation phenomenon, requiring both $s < 1$ and off-diagonal structure in A. With physical output and $s = 1$ the customer stock drops out entirely: the output-price channel is absent by measurement and the cost-side channel is absent by constant variable returns.*

Physically measured output is itself a strong requirement. Output prices and quantities are not available in typical production-network data, and often not in other firm-level datasets. Section 4.5 shows that bilateral transaction values substitute for them, and Proposition 2 characterizes what deflated revenue delivers instead.

This is Hulten [1978]'s theorem at the level of estimation, now with contributions to the measured residual from the firm's own supplier stock through variety (the diagonal, $\zeta\widetilde{\ln v_i}$), its suppliers' propagated effective productivity (the off-diagonal, $\Psi_s - \mathbb{I}$), and, when $s < 1$, its suppliers' customer stocks. A firm is mechanically measured as productive if it is technologically good, or sources from many varieties, or sources from productive (low-price) suppliers, or its suppliers themselves source from many or productive suppliers, and so on up the chain. In the running example, the artisan and the networker are measured as equally productive. Cheap inputs are indistinguishable from technology in common-deflated data, and only the variety subtraction, using the networker's observed supplier count, re-ranks them.

Magnitudes. At conventional values, $\sigma = 4$, $\alpha_m = 0.5$, $s = 0.9$, the variety elasticity is $\zeta = \alpha_m/(\sigma-1) \approx 0.17$ and the propagation coefficient is $\alpha_m/D \approx 0.38$. Doubling the supplier base raises measured productivity by $2^\zeta - 1 \approx 12$ percent with no change in technology. A firm's direct suppliers' effective productivity enters its measured productivity with share-

weighted coefficient 0.38, the second tier with 0.15, and the total upstream loading, the row sum of $\Psi_s - \mathbb{I}$, is $\alpha_m/(D - \alpha_m) \approx 0.63$, two thirds of the weight on the firm's own fundamentals. The off-diagonal is not a marginal contributor. At realistic parameters a large share of measured cross-sectional variation can originate upstream.

Corollary 1 (Dispersion inflation and network correlation)**.** *Let $\omega_i^e = \omega_i + \zeta\widetilde{\ln v_i}$ be effective log productivity. If ω^e is cross-sectionally i.i.d., the average variance of measured productivity $\hat{\boldsymbol{\omega}} = \Psi_s\,\boldsymbol{\omega}^e$ weakly exceeds that of ω^e, strictly when $\alpha_m > 0$ and the network is non-trivial. If, in addition,* $\mathrm{Cov}(\omega, \ln v) \geq 0$*, which is the case when more-productive firms hold more suppliers, then*

$$\mathrm{Var}(\hat{\omega}) \;\geq\; \mathrm{Var}(\omega^e) \;\geq\; \mathrm{Var}(\omega),$$

so measured productivity is at least as dispersed as technology. Measured productivity is positively correlated across linked firms even when $\boldsymbol{\omega}$ is i.i.d., and the firm ranking by measured productivity over-ranks firms with large or productive supplier bases.

The source of the inflation differs from Asker et al. [2014], where costly adjustment of a measured dynamic input inflates the dispersion of revenue products. Here the dynamic inputs are omitted by the estimator and enter the residual itself.

Corollary 2 (Exact correction)**.** *If the network matrix A, the elasticity α_m, the substitution elasticity σ, the returns to scale s, and the two stocks (v, n) are observed, true productivity is recovered by inverting the damped operator and subtracting the variety and customer terms,*

$$\boldsymbol{\omega} = \Big(\mathbb{I} - \tfrac{\alpha_m}{D}A^\top\Big)\,\hat{\boldsymbol{\omega}} - \zeta\,\widetilde{\ln \boldsymbol{v}} + (1-s)\Big(\mathbb{I} - \tfrac{\alpha_m}{D}A^\top\Big)(\Psi_s - \mathbb{I})\widetilde{\ln \boldsymbol{n}}. \tag{23}$$

The correction is feasible precisely in firm-to-firm data, where A is observed from downstream purchases, α_m from the intermediate cost share, and v_i, n_i from the counts of active suppliers and buyers. All ingredients are read directly from the firm-to-firm panel.

Implication: measured productivity spillovers. Corollary 1 has a direct consequence for the literature that estimates productivity spillovers from the comovement of measured TFP across linked firms, whether by regressing pre-estimated TFP on partner TFP [Serpa and Krishnan, 2018] or by embedding partners' average contemporaneous productivity in the law of motion [Iyoha, 2023]. Measured residuals of connected firms are correlated through propagated input prices and the supplier margin, even when true technologies are independent. In Section 9, we show that these properties of measured productivity hold for the Belgian manufacturing data.

Because the relationship stocks accumulate, partners' lagged measured effective productivity predicts own measured effective productivity, so timing-based instruments such as the network's lagged TFP do not separate a technological externality from price-based spillovers. An estimated spillover parameter therefore loads on both. The decomposition implies an order of operations. Apply the static corrections first, and the network correlation that survives in the corrected residual is the candidate technological spillover, estimable with the instruments of Iyoha [2023]. The exercise of separating spillover channels parallels Bloom et al. [2013].

4.3 Output-price, coefficients, markups, and matching

Proposition 2 (Revenue productivity and the returns to scale)**.** *Under constant-elasticity demand and the constant markup at the true coefficients, deflated revenue recovers*

$$\hat{\omega}_i^{\text{rev}} = \Big(1 - \tfrac{1}{D}\Big)\big(\omega_i^c - \alpha_m \widetilde{P^M}_i\big) + \frac{1-s}{D}\,\widetilde{\ln n}_i, \qquad 1 - \tfrac{1}{D} = \frac{(\sigma-1)(1-s)}{D}. \tag{24}$$

At constant variable returns ($D = 1$) deflated revenue contains no residual technology component, the output-price wedge cancels the cost fundamental ω_i^c exactly, and the variation orthogonal to inputs is demand variation scaled by the common markup (the customer stock n_i, the buyer set, the demand index), not technology. With decreasing returns ($D > 1$) revenue recovers a fraction $1 - 1/D$ of the physical-output residual and additionally loads on the customer stock through the scale feedback. In either case, with a valid output-side instrument the recoverable elasticities are the true ones attenuated by the markup, $\hat{\beta}_j = \alpha_j/\mu$ for each input $j \in \{k, l, m\}$, and $\widehat{RTS} = \eta/\mu$ with $\eta = \alpha_k + s$.[14]

This is the network form of the revenue-elasticity problem. Output elasticities are not recovered from revenue data alone [Bond et al., 2021, De Loecker, 2021], and Ackerberg and De Loecker [2024] give the conditions for identification under imperfect competition. The constant-elasticity demand and constant markup of Section 2 are the restriction that reduces the attenuation here to the single scalar $1/\mu$.

Constant returns and TFPR. At constant variable returns, revenue productivity contains no technology at all: TFPR is customer capital. Constant variable returns is not constant returns at the firm level. Under constant total returns, $\eta = 1$ and $s = 1 - \alpha_k$, the dampening scalar is $D = 1 + (\sigma - 1)\alpha_k$, and deflated revenue retains the fraction $(D-1)/D = (\sigma-1)\alpha_k/\big(1 + (\sigma-1)\alpha_k\big)$ of the physical-output residual, 0.47 at $\sigma = 4$

[14] No step of the guide in Section 7 relies on an output-side instrument.

and $\alpha_k = 0.1$. Roughly a quarter of a CRS firm's revenue-productivity variation is network-contaminated technology and the rest is customer capital. Because the customer stock is dispersed, persistent, and predictive of growth and survival, a residual that is half demand capital passes every standard validation exercise. The reading matches the direct evidence: where output prices are observed, the demand component of revenue productivity is large, persistent, and dominant for survival [Foster et al., 2008, 2016], which is the behavior of n_i here.

Coefficient and returns-to-scale bias. The intermediate regressor $x_i = m_i + \widetilde{P^M}_i$ contains the firm-specific input price (20) leading to an errors-in-variables problem with a correlated error. Cheaper sourcing (lower $\widetilde{P^M}_i$, from a larger v_i or productive suppliers) induces more intermediate purchases, so $\mathrm{Cov}(m_i, \widetilde{P^M}_i) < 0$. The probability limit of the intermediate coefficient is

$$\operatorname{plim} \hat{\beta}_m = \alpha_m \frac{\mathrm{Var}(m) + \mathrm{Cov}(m, \widetilde{P^M})}{\mathrm{Var}(m) + 2\mathrm{Cov}(m, \widetilde{P^M}) + \mathrm{Var}(\widetilde{P^M})}, \tag{25}$$

which is below α_m if and only if $\mathrm{Var}(\widetilde{P^M}) + \mathrm{Cov}(m, \widetilde{P^M}) > 0$. The wedge dispersion $\mathrm{Var}(\widetilde{P^M})$, by (21), rises with the dispersion and covariance structure of the supplier stock, of supplier prices, and of upstream effective productivity. With $s < 1$ there are two offsetting effects as the propagation operator is damped (smaller composition variance) and the suppliers' customer-stock dispersion enters the wedge (additional variance). Returns to scale, $\eta = \alpha_k + s$, inherit the bias.

The regressor is expenditure deflated by a common index, but the firm's true input price differs from that index, precisely because of its supplier network. The input quantity is measured with an error that is larger and more correlated with behavior, the more the network matters. The estimated intermediate elasticity is dragged down, and returns to scale with it. The expression characterizes the bias at the true coefficients rather than the plim of a feasible two-step procedure, since the productivity term it partials out is unobserved. Appendix E gives the derivation.

Markups, matching, and buyer power. The intermediate revenue share is $s_i^M = \alpha_m/\mu$ for any s (Appendix E), so the gross-output share identity of Gandhi et al. [2020] and the production-approach markup $\hat{\mu}_i = \theta_i^m / s_i^M$ of De Loecker and Warzynski [2012], stated compactly in De Loecker [2011b], with θ_i^m the output elasticity of the intermediate input, here the constant α_m, are unaffected by the returns to scale. Returns to scale do not enter a single flexible-input first-order condition. This is the one-flexible-input logic of De Loecker and Warzynski [2012]: any flexible input prices the markup in principle, and equivalence

across inputs is itself a testable restriction. Through the lens of the model, markup estimates that differ with the variable input used [Raval, 2023] point to firm-specific input prices in one of the inputs, here the network's intermediate price, while production labor is contaminated instead by the three-use composite of Section 5.4, so the intermediate first-order condition applied to properly deflated expenditure is the clean choice. De Loecker and Scott [2026] consider the production approach against demand-side markup estimates in the same industry. If market shares vary with network position [Bernard et al., 2022] and markups are variable, the elasticity attenuation is firm-specific, and $\hat{\mu}_i$ inherits the network's share variation. With variable markups the propagation of Proposition 1 would also no longer be summarized by $(\alpha_m/D)A^\top$ alone: prices carry markup terms, sales and cost-based weights separate, and an allocative-efficiency term appears in aggregation [Baqaee and Farhi, 2020]. The constant-markup benchmark is what buys the closed forms. The leading case it rules out is the customer-market pricing of Somaini and Einav [2013]. Endogenous matching makes intermediate demand a function of own productivity and the realized supplier set, violating scalar-unobservable invertibility (Assumption 1(iii)). Buyer power drives a position-dependent wedge $s_i^M = \alpha_m/(\mu\varpi_i^B)$ between the intermediate revenue share and its output elasticity, violating the share identity. Here $\varpi_i^B \leq 1$ is the ratio of marginal to average expenditure on the input, below one because an extra unit lowers the price paid on inframarginal units. Demirer and Rubens [2025] identify buyer and seller power jointly from production data.

4.4 Value added

Most applied work estimates value-added specifications in the tradition of Olley and Pakes [1996], Levinsohn and Petrin [2003], and Ackerberg et al. [2015] rather than gross output. Value added amplifies the static biases. The subtraction of intermediates is exact only under fixed proportions and a common intermediate price, and the sourcing technology (1) excludes both, the first because $\sigma > 1$ makes the bundle substitutable, the second because the input price index is firm-specific.

Proposition 3 (Value-added estimation)**.** *Let value added be constructed by double deflation at the observed intermediate revenue share, $\widetilde{va}_i = (\tilde{y}_i - s^M\tilde{x}_i)/(1-s^M)$ to first order, and let the researcher estimate the value-added function $\widetilde{va}_i = \beta_k\tilde{k}_i + \beta_l\tilde{\ell}_i + \omega_i^{VA}$. At the value-added coefficients $\beta_j = \alpha_j/(1-s^M)$ the recovered residual is*

$$\omega_i^{VA} = \frac{1}{1-s^M}\Big[\hat{\omega}_i + \big(\alpha_m - s^M\big)\,\tilde{x}_i\Big], \qquad \alpha_m - s^M = \alpha_m\Big(1 - \tfrac{1}{\mu}\Big) > 0, \tag{26}$$

with $\hat{\omega}_i$ the gross-output residual of (18). (i) *Every static bias, variety, propagation, the customer term, and the input-price error, carries over scaled by $1/(1-s^M)$, roughly a factor of two at intermediate shares near one half.* (ii) *A new term appears: with a markup, the revenue share understates the output elasticity, so deflated intermediate expenditure $\tilde{x}_i$ survives in the residual as an omitted flexible input, correlated with output and the network states, the value-added markup bias of Klette and Griliches [1996].* (iii) *At $\mu = 1$ the extra term vanishes and $\omega_i^{VA} = \hat{\omega}_i/(1-\alpha_m)$, so the sourcing advantage enters measured value-added productivity with weight $\zeta/(1-\alpha_m)$ rather than ζ.* (Proof: Appendix E.)

Value-added estimation adds to the gross-output biases, and the corrections below apply to it with every static term multiplied by the intermediate multiplier $1/(1-s^M)$. The clean route is to correct at the deflation stage and estimate gross output, or, where value added is required, to construct it from the firm-specific deflators of Section 4.5 before subtracting.

4.5 Identification without prices

In domestic firm-to-firm data such as VAT or e-invoice networks, transaction prices are rarely observed. What is observed is the value of each link, $r_{ij} = p_i q_{ij}$, the link size z of Section 2 written here to emphasize its role as a bilateral revenue, with firm totals. Pairwise values nonetheless identify relative productivity, by exploiting variation across the buyers a seller serves. Taking logs of the bilateral value $r_{ij} = \vartheta_{ij}(p_i/P_j^M)^{1-\sigma}X_j$, with ϑ_{ij} the match-specific component and X_j buyer j's total intermediate spending, shown below to equal the observed E_j^M, yields the seller–buyer decomposition. The match-specific component carries the deepening of Section 2. A link is born at the match size and grows along the tenure profile, so $\vartheta_{ij} = h_{a_{ij}}\,\bar{\vartheta}_{ij}$, with a_{ij} the link's age, h_a the common tenure profile, and $\bar{\vartheta}_{ij}$ the idiosyncratic match quality. In logs,

$$\ln r_{ij} = \underbrace{(1-\sigma)\ln p_i}_{\psi_i} + \underbrace{(\sigma-1)\ln P_j^M + \ln X_j}_{\theta_j} + \underbrace{\ln h_{a_{ij}}}_{g_{a_{ij}}} + \underbrace{\ln \bar{\vartheta}_{ij}}_{\varepsilon_{ij}}. \tag{27}$$

This corresponds to a two-way fixed-effects model with seller effect ψ_i, buyer effect θ_j, and match residual ε_{ij}, identified up to normalization from a connected network under exogenous mobility, as in Bernard et al. [2022] (*seller–buyer decomposition*). The tenure term g_a is common to all links of age a, so it is absorbed by link-age dummies added to the two-way regression, and the match residual ε_{ij} is what remains of match quality after age. Omitting the dummies pushes each firm's link-age composition into ψ_i and θ_j, and both deflators inherit it. The regression is run year by year, so the effects are within-year objects; match quality may drift over time, and the time index is suppressed.

Both fixed effects are structural measurements. The seller effect is the output price, $\psi_i = (1-\sigma)\ln p_i$. The buyer effect is the input price index. Writing $\ln S_j \equiv \ln \sum_i \vartheta_{ij}\, p_i^{1-\sigma}$ for the log CES input price aggregator of buyer j (an aggregator of prices, not the quantity bundle M_j), the accounting identity $X_j = E_j^M$ and the price definition $(P_j^M)^{1-\sigma} = S_j$ turn θ_j of (27) into $\theta_j = x_j - \ln S_j$, with $x_j = \ln E_j^M$ observed.[15]

Lemma 1 (The buyer effect is the exact input price index)**.** *Under the CES demand of Section 2, the buyer effect satisfies $\theta_j = x_j - \ln S_j$, so the log input aggregator is recovered from observed expenditure and the estimated effect, $\ln S_j = x_j - \theta_j$, and the firm-specific input price index is exact in deviations,*

$$\widetilde{P^M}_j = -\frac{1}{\sigma-1}\big(x_j - \theta_j\big). \tag{28}$$

The value-only construction (30) *approximates $\ln S_j$ by the supplier count and the share-weighted mean seller effect. It drops a Jensen composition term in the dispersion of match quality that* (28) *carries. The two agree at $D = 1$ with degenerate match quality and differ by the composition term otherwise. The identity is exact in population; the estimated effects carry sampling noise, taken up in the generated-deflator discussion of Section 7.*

Proposition 4 (Value-only recovery of productivity)**.** *The estimated seller fixed effect equals the network-propagated effective productivity (net of the customer term), scaled by the demand elasticity and the scale dampening,*

$$\tilde{\psi}_i = \frac{\sigma-1}{D}\left[\Psi_s\big(\boldsymbol{\omega} + \zeta\widetilde{\ln \boldsymbol{v}} - (1-s)\widetilde{\ln \boldsymbol{n}}\big)\right]_i, \tag{29}$$

so with A, α_m, σ, s, and the two stocks recovered from values, expenditure shares, and link counts, true productivity is identified by inverting the damped operator and subtracting the variety and customer terms,

$$\boldsymbol{\omega} = \frac{D}{\sigma-1}\left(\mathbb{I} - \tfrac{\alpha_m}{D}A^\top\right)\tilde{\boldsymbol{\psi}} - \zeta\,\widetilde{\ln \boldsymbol{v}} + (1-s)\,\widetilde{\ln \boldsymbol{n}}. \tag{30}$$

No price or quantity data are used. Cross-buyer variation in a seller's sales substitutes for the missing output price, the expenditure-share matrix for the missing input-price composition, the supplier count for the variety term, and the customer count for the scale term. At $D = 1$ this is the constant-returns recovery $\boldsymbol{\omega} = \frac{1}{\sigma-1}(\mathbb{I} - \alpha_m A^\top)\tilde{\boldsymbol{\psi}} - \zeta\widetilde{\ln \boldsymbol{v}}$.

Comparing how much different sellers sell to the same buyer nets out the buyer, so the

[15] Expenditure here is nominal, while the intermediate regressor of (15) is deflated by $\bar{P}^M$. The common deflator is a year constant and drops from every deviation, so the same symbol serves.

variation that remains in a seller's sales across its buyers is its price competitiveness. Under constant markups price competitiveness is effective marginal cost. No price data are needed. The seller fixed effect of Bernard et al. [2022] is not a measure of own productivity but of network-propagated effective productivity. A firm scores high because it is productive, because it buys from many varieties, or because it buys from productive suppliers. The same network inversion that corrects the price-based estimate corrects the value-based one, and the supplier and customer counts needed for the subtractions are supplied directly by the firm-to-firm panel.[16]

Deflation as the first step. Both effects are more than diagnostics. The seller effect delivers each firm's relative output price, $\tilde{p}_i = -\tilde{\psi}_i/(\sigma-1)$. The buyer effect delivers the firm-specific input price index exactly, $\widetilde{P^M}_i = -(x_i - \theta_i)/(\sigma-1)$ by Lemma 1, and we build the input deflator from it. The count-plus-mean form $\widetilde{P^M}_i = -\frac{1}{\sigma-1}\widetilde{\ln v}_i + (A^\top \tilde{\boldsymbol{p}})_i$ is the same object with the Jensen composition term dropped, and it is the fallback where the buyer effect is poorly measured. Deflating revenue and intermediate expenditure by these constructed prices converts the data available to the econometrician into the physically measured data the propositions assume. The fixed-effects regression can serve as a universal first step, removing the output-price, variety, propagation, and errors-in-variables biases at the deflation stage. The dynamic corrections of Section 5 then operate on properly deflated quantities.[17] Read against the older literature, this is a solution by measurement to the unobserved input-price problem [Katayama et al., 2009, Grieco et al., 2016], and to the output-price counterpart treated structurally by De Loecker [2011a]. The network supplies every firm's input price index, variety from counts, composition from shares and recovered supplier prices. Supplier quality differences, being priced, travel through the same composition term.

Observed prices are valuable where they exist. The construction inherits the assumptions of Proposition 4, constant markups and separable bilateral demand, so the recovery is cardinal only under CES; it needs a connected network and the parameters (α_m, σ, s); and the estimated effects carry sampling noise, which passes into the constructed deflators as measurement error in the regressors. The construction is also first order, since with heterogeneous match quality the exact CES index adds a Jensen term in the dispersion of ϑ_{ij} that the log-linear deflator omits. Within the model, however, the answer to whether output prices are needed at all is no. Bilateral values substitute for them everywhere.

[16]The maintained assumption is exogenous mobility. If high-productivity sellers systematically match buyers with whom match quality is unusually high, $\tilde{\psi}_i$ is biased, the network analogue of sorting bias in matched employer–employee models [Andrews et al., 2008, Kline et al., 2020], and the standard exogenous-mobility diagnostics of Bernard et al. [2022] apply.

[17]It also makes the dynamic revenue problem moot. What deflated revenue contains when the stocks accumulate, and why partial measures load on the customer base, is recorded in Appendix A.

The corrections are valid under any matching technology measurable in firm-level states, random, assortative, or search-frictional, as long as the one pairwise restriction of exogenous mobility holds, i.e., there is no selection on the pairwise match component ϑ_{ij}. Appendix C gives the full statement, including the single conditioning variable that non-representative matching adds to the dynamic corrections.

5 The dynamic layer: customer and supplier capital

We let the two stocks accumulate and study the dynamic estimation. The standard procedure follows the proxy recipe: (i) recover expected output $\hat{\phi}_t$ from the first stage; (ii) invert investment, given capital, for the scalar unobservable; (iii) impose the exogenous Markov process of Assumption 1(ii) to form the innovation; (iv) estimate the elasticities from the innovation's orthogonality to capital and lagged inputs; and (v), where exit matters, add a survival propensity in capital and investment.

The networked firm breaks this recipe at three points, the inversion, the Markov step, and the survival correction, and contaminates its ingredients at two more, the labor regressor (9) and deflated revenue, the latter treated in Appendix A because the deflation step of Section 4.5 resolves it. This layer assumes availability of the linked panel $\mathcal{I}_2$.

5.1 The proxy inversion and contaminated dynamics

If the researcher inverts investment for productivity, a higher draw means more investment (given capital), so investment is asked to reveal the shadow value of one asset. The networked firm prices three assets at once, so investment reveals none of them alone.

Under the policy (14), $i_t = \mathcal{I}(k_t, n_t, v_t, \omega_t)$, and the correspondence from (k_t, i_t) to ω_t is one-to-many:

$$i_t = \mathcal{I}(k_t, n_t, v_t, \omega_t) \implies \omega_t \neq h(k_t, i_t) \text{ for any function } h. \tag{31}$$

The control-function step $h(k_t, i_t)$ returns a productivity proxy that is a mixture of ω_t, n_t, and v_t. There are three payoff-relevant states where the method assumes one, so the first stage of Olley and Pakes [1996] is misspecified, as are those of Levinsohn and Petrin [2003] and Ackerberg et al. [2015] with $m_t = m(k_t, n_t, v_t, \omega_t)$, and the elasticity and productivity estimates are inconsistent. Investment is not invertible for productivity given capital alone, and it becomes invertible only conditional on (n_t, v_t).

The deepest problem is dynamic. Relationship stocks move slowly, a customer base survives a bad productivity draw, so a residual contaminated by the stocks inherits their

smoothness. The econometrician mistakes the firm's connections for its ability, and connections outlive bad luck. In the running example, the networker's measured productivity is the steadier series even if the artisan's technology is the more persistent.

Write the recovered proxy as $\hat{\omega}_t = \omega_t + b_n\tilde{n}_t + b_v\tilde{v}_t$, with $\tilde{n}_t, \tilde{v}_t$ the components of the two stocks orthogonal to ω_t, persistent at the common link-survival rate $1-\delta$. This is a stylized log-linear approximation that attributes the stocks' persistence to mechanical survival. The acquisition response to persistent profit adds to it, strengthening persistence of the stocks. The estimated first-order autocorrelation is the variance-weighted average

$$\widehat{\rho} = \frac{\rho_\varphi \mathrm{Var}(\omega) + (1-\delta)\big(b_n^2\mathrm{Var}(\tilde{n}) + b_v^2\mathrm{Var}(\tilde{v})\big)}{\mathrm{Var}(\omega) + b_n^2\mathrm{Var}(\tilde{n}) + b_v^2\mathrm{Var}(\tilde{v})} \tag{32}$$

of the technology persistence ρ_φ and the link-survival rate $1-\delta$. The loadings b_n, b_v depend on the returns to scale (they run partly through the price channel (19) when $s < 1$). The variance-weighted average and the direction of the bias do not. Estimated productivity is therefore "too persistent" precisely when relationships are sticky relative to technology.

The composition of the bias also drifts with age. Young firms hold thin networks, so their measured productivity is mostly technology. Old firms carry large accumulated stocks, so theirs is increasingly relationship capital. Measured productivity that grows more persistent with firm age is, through this lens, the accumulation of relationship capital rather than a change in the technology process, and the corrected residual provides the test.

5.2 Exit selection

Exit depends on the same stocks. A firm carries itself through a bad year on its installed base of customers and suppliers, so a survivor without a dense network must have had an unusually good productivity draw. The artisan survives on its technology, the networker on its connections.

The firm exits when the value of continuing falls below the cost of carrying the firm, alongside the exogenous death hazard δ_f. The cost of carrying the firm is the servicing bill plus a fixed overhead labor cost $w\bar{\ell}_i$, the overhead stock introduced with the labor split in Section 7.2.[18] Because the value function is increasing in all of (K, n, v, ω) for any $s \leq 1$, the exit threshold is a surface $\underline{\omega}(K, n, v)$, decreasing in each stock. Among survivors, selection induces a negative correlation between ω and the two stocks conditional on fundamentals, concentrated among young, small-network firms. The survival correction is restored by conditioning the propensity on the observed stocks, $P_t = P(k_t, i_t, n_t^\star, v_t^\star)$.

[18]The overhead cost makes the carrying cost strictly positive and opens the exit margin. It changes no first-order condition and none of the closed forms above.

The omitted threshold variation loads on the stocks exactly as the proxy contamination does. The testable signature is direct. Given recovered productivity, the exit hazard should fall with both observed stocks, and the firm hazard should fall with age through the accumulating portfolio, the firm-level face of the duration dependence developed in Section 2.

5.3 The correction

The core failures are repaired by measurement rather than structure, because the contaminating states are visible in the firm-to-firm panel.

Proposition 5 (Failure and repair of the proxy method)**.** *Let the data come from the firm of Section 2, with the two stocks n_t and v_t observed as the active link counts under $\mathcal{I}_2$.*

(i) Inversion. *The investment policy* (14) *makes the correspondence from (k_t, i_t) to ω_t one-to-many, equation* (31)*, and likewise for an intermediate-demand proxy, so the recovered control is a mixture of ω_t, n_t, and v_t.*

(ii) Dynamics. *The estimated persistence of the recovered proxy is the variance-weighted average* (32)*, pulled from the technology persistence ρ_φ toward the stocks' persistence.*

(iii) Exit. *The exit threshold is a surface $\underline{\omega}(K, n, v)$ decreasing in each stock, so a survival propensity conditioned on (k_t, i_t) alone is misspecified, and surviving thin-network firms are positively selected on productivity.*

(iv) Repair. *Augmenting the control function to $\omega_t = h(k_t, i_t, n_t, v_t)$ restores the inversion, investment being strictly monotone in ω_t conditional on (k_t, n_t, v_t), a monotonicity which holds where the firm's problem is concave, the pointwise sub-critical region of Bernard [2026].*[19] *Conditioning the survival propensity on (k_t, i_t, n_t, v_t) restores the exit correction. With heterogeneous relationships the conditioning stocks are the value-weighted $(n^\star, v^\star)$ (Section 2). Neither the failures nor the repair depends on the returns to scale. The state count and the conditioning set are the same for any s.* (Proof: Appendix E.)

The contrast with the latent-state literature is the methodological payoff and it now applies twice. Doraszelski and Jaumandreu [2013] handle a latent controlled productivity state by modeling its law of motion and instrumenting, as does De Loecker [2013] where the control is a participation decision. Here the two controlled states, customer capital on the demand side and supplier capital on the cost side, are both observed as the firm's portfolio of links. Each enters the control function as a conditioning variable rather than a latent process to

[19] Strictly, with the acquisition capabilities χ_t, ξ_t varying over time around their entry-set types, investment carries them as further states. They are firm constants up to that AR(1) variation, so firm fixed effects absorb the types. Under $\mathcal{I}_2$ the AR(1) innovations enter the first-stage error, attenuating rather than overturning the correction.The statement in the text is exact when χ, ξ sit at their types.

be inverted. Survey-based expectations offer a complementary route that dispenses with the inversion altogether [Norris Keiller et al., 2024], and identification schemes designed to be invariant to the productivity process [Utamaru, 2026] are robust, by construction, to the persistence contamination of Proposition 5(ii), though not to the static price layer.

5.4 A measurement problem: labor

One further bias is a problem of measurement rather than inversion. The researcher's labor regressor is total employment or the wage bill. In the data-generating process most of a networked firm's non-production employees carry its relationships, servicing the installed bases and prospecting for new links, so the wage bill is not a measure of the production input.

Using total labor $L_i = L_i^p + (L_i^n + c_n\, n_i) + (L_i^v + c_v\, v_i)$ in place of production labor L_i^p biases the labor elasticity and returns to scale. The customer-facing and sourcing components rise with the two stocks, the stocks rise with output, so L_i is a contaminated measure of the production input with two non-classical error components, and $\operatorname{plim} \hat{\beta}_l \neq \alpha_l$ (Appendix E).

Neither non-production use responds to current output. The servicing components are proportional to the slow-moving stocks, and the acquisition components respond to the relevant lifetime value and the wage. Of the biases in this paper, this is the one least repaired by data. The occupational split that removes it is rarely recorded. A model-based substitute is developed in Section 7.

6 The combined model

The static and dynamic layers act on the same measured object. A firm's relationships are its links in the network and its accumulated relationship capital. The customer stock n_i is its out-degree and its demand-side capital, the supplier stock v_i its in-degree and its cost-side capital. Throughout this section ω denotes the firm's technology state of Sections 4–5, $\omega = \ln \varphi$ under the $\alpha_k = 0$ convention. With predetermined capital, read the cost fundamental ω^c of (17) in its place (under $\mathcal{I}_2$ the capital terms are then subtracted as in the capital paragraph of Section 4). Collecting Proposition 1 and Proposition 5, the productivity

recovered from networked, dynamic data is, to first order,

$$\hat{\omega}_i = \underbrace{\omega_i}_{\text{technology}} + \underbrace{\zeta\widetilde{\ln v_i}}_{\substack{\text{supplier capital}\\ \text{(variety, cost side)}}} + \underbrace{\big[\Psi_s - \mathbb{I}\big](\boldsymbol{\omega} + \zeta\widetilde{\ln \boldsymbol{v}})\big|_i}_{\text{upstream propagation}} - \underbrace{(1-s)\big[\Psi_s - \mathbb{I}\big]\widetilde{\ln \boldsymbol{n}}\big|_i}_{\substack{\text{customer capital, cost side}\\ \text{(only if } s<1\text{)}}}$$
$$+ \underbrace{b\big(\tilde{n}_i^\star, \tilde{v}_i^\star\big)}_{\substack{\text{relationship capital}\\ \text{(through the proxy)}}} + \underbrace{\tilde{\bar{g}}_i}_{\substack{\text{portfolio age}\\ \text{(deepening composition)}}} + \text{(price, markup, labor terms)}. \tag{33}$$

The static terms are levels from the cross-section of prices and partners. The proxy term, $b(\cdot)$, is the dynamic omitted-state contamination the naive control function attributes to productivity. It is a function of the value-weighted stocks that is $b_n\tilde{n}_i^\star + b_v\tilde{v}_i^\star$ to first order and the polynomial control function $h(k, i, n^\star, v^\star)$ of Proposition 5 in estimation. The portfolio-age term $\tilde{\bar{g}}_i$ is the firm's average link tenure.[20] The components are corrected at distinct stages of the estimation pipeline, and none of the corrections removes another's bias. With capital in marginal cost, own capital $\alpha_k\tilde{k}_i$ sits inside the technology term and upstream capital inside the propagation term (Section 4).

In a production network with accumulated customer and supplier capital, measured productivity is a composite of own technology, the firm's own supplier capital (through input variety), the network-propagated effective productivity of upstream suppliers, the firm's own customer capital (through propagated prices when $s < 1$), the tenure composition of its relationships when the decomposition omits the deepening term, and capital intensity. Its cross-sectional dispersion overstates technological dispersion, and its estimated persistence is biased toward the stocks' persistence. Consistency requires correcting at both the deflation/fixed-effects stage (variety, propagation, and the customer scale term) and the proxy/Markov stage (supplier and customer capital). The corrections are jointly feasible because the buyer-supplier firm-to-firm panel observes the network matrix A, both relationship stocks (n, v), and the firm-specific transaction values.

The two stocks are symmetric in the model but asymmetric in how each contaminates the measured residual, the asymmetry of the two brackets of (6). The supplier stock sits in the cost-side bracket, so it contaminates the productivity term directly and is stripped at the deflation/network-inversion stage. The customer stock sits in the demand-side bracket, so it contaminates the investment proxy (and, with decreasing returns, prices), and is stripped by conditioning the control function on the observed link count and by the network inversion. A researcher who corrects only one leaves the other's bias intact. Table 1 collects the full

[20] This appears whenever the seller–buyer decomposition omits the tenure term of (27): the two firm effects then absorb each firm's link-age composition, and both deflators inherit it. It is removed at the decomposition stage by estimating g_a jointly. No age term survives in the corrected recovery.

Table 1: The biases at a glance.

Bias	Mechanism	Symptom	Fix (data required)
Effective productivity (Prop. 1)	more suppliers ⇒ cheaper inputs ⇒ lower measured cost	residual correlated with supplier count; inflated dispersion	subtract the variety term $\zeta\widetilde{\ln v}$ (supplier counts; ζ estimated in Sec. 7)
Upstream propagation (Prop. 1)	suppliers' productivity passes through input prices	residual correlated along A	invert $(\mathbb{I} - \frac{\alpha_m}{D}A^\top)$ (expenditure shares; α_m, D estimated in Sec. 7)
Revenue residual (Prop. 2; App. A)	the markup cancellation makes revenue productivity load on demand; at $s < 1$ it recovers a fraction $1 - 1/D$ of technology	residual orthogonal to technology at $D = 1$; loads on the customer base	physical output or value-only FE; condition on n (link counts)
Value added (Prop. 3)	subtraction of intermediates exact only under fixed proportions and common prices	every static bias inflated by $1/(1 - s^M)$; omitted-intermediate markup term	estimate gross output, or subtract using the firm-specific deflators (Sec. 4.5)
Coefficients (eq. 25)	firm-specific input price inside deflated intermediates	attenuated $\hat{\beta}_m$ and returns to scale	firm-specific deflator $v^{1/(1-\sigma)}\bar{p}^M$ (supplier counts and recovered supplier prices)
Proxy inversion, persistence (Prop. 5)	investment reflects three assets, not one	persistence pulled toward the link-survival rate	condition the control function on $(n^\star, v^\star)$ (link panel)
Exit selection (Prop. 5(iii))	participation threshold falls with both stocks	surviving thin-network firms look too productive	survival propensity on $(k, i, n^\star, v^\star)$ (value-weighted stocks)
Labor (Sec. 5.4)	servicing and acquisition labor scale with the stocks, not output	labor elasticity falls when total labor is used	production labor (occupations, or the structural split of Sec. 7)
No balance sheet ($\mathcal{I}_1$, Sec. 3.2)	link panel contains no capital	residual mixes technology and capital intensity	static corrections recover $\omega^c = \omega + \alpha_k\tilde{k}$; linked accounts split it

set.

What the biases do and do not contaminate. The contamination is cross-sectional and dynamic, not aggregate-on-impact. With the relationship stocks predetermined, the first-order response of real aggregate output to any profile of firm shocks is the Domar-weighted sum, so static growth accounting at the aggregate level is correct even while every firm-level residual is mismeasured. The network's signature appears instead in the firm-level cross-section (dispersion, ranking, and correlation along the network), in measured persistence, and in the time series of aggregates. Because idiosyncratic shocks are amplified over time as relationship stocks accumulate, aggregate TFP shows delayed growth after dispersed shocks, an endogenous-network echo that a static decomposition in the spirit of Baqaee and Farhi [2019] attributes to changing Domar weights.

7 The estimation strategy

The corrections are operations on one object, the buyer–supplier firm-to-firm panel linked to firm accounts where every estimated object names the moment that identifies it. We give two identification paths that deliver the same parameters. Path 1 is a staged pipeline whose one fixed point, a scalar crossing in σ, we characterize with an existence and uniqueness argument. Path 2 stacks the same information into a single joint GMM system. Both identify $(\sigma, \alpha_k, \alpha_l, \alpha_m)$ and, from them, the derived objects (s, D, ζ, μ). The default information set is the linked panel $\mathcal{I}_2$. Appendix B gives the moment system for Path 2.

7.1 Data, the ledger, and the identifying assumptions

The model maps to a linked firm-to-firm panel linked to annual accounts. Table 2 lists each object, its data counterpart or identifying moment, and its status. One group of objects needs no model parameters at all and is read, or estimated as a sample object, directly from the link panel. This includes the network matrix and the two counts; the gross formation flows $F^n_{it} \equiv n_{i,t+1} - (1 - \hat{\bar{\delta}})n_{it}$ and symmetrically F^v_{it}, with $\hat{\bar{\delta}}$ the estimated mean link-failure rate and $\hat{\delta}(\cdot)$ the size-dependent failure schedule from the link-survival hazard, $\delta(z) = \bar{\delta}\, z^{-\eta_\delta}$ with z the link size and η_δ the failure size-elasticity; the deepening profile h_a, link size as a function of tenure a; and the value-weighted stocks $n^\star, v^\star$ as within-firm sums of link sizes, a link's size being its sales value downstream and its purchase value upstream.

The recipe requires persistent link identifiers over time, since tenure, the survival hazard, and the formation flows all track a link across years. This is the one genuinely network-

Table 2: The ledger: model objects, data counterparts or moments, and status.

Object	Data counterpart / identifying moment	Status
network matrix A; stocks n_i, v_i; flows F^n, F^v; link sizes, tenures	bilateral values r_{ki}; shares r_{ki}/E_i^M; link counts and their panel dynamics	observed ($\mathcal{I}_1$)
$n^\star, v^\star$; $\delta(\cdot)$, h_a, η_δ	sums of link sizes; link-survival hazard; size by tenure	estimated from the link panel, no parameters ($\mathcal{I}_1$)
R_i, E_i^M, s_i^M, L_i, wage bill, K_i, I_i	annual accounts	observed ($\mathcal{I}_2$)
seller, buyer effects $\tilde{\psi}_i, \theta_j$	seller–buyer decomposition of $\ln r_{ij}$, eq. (27)	estimated; Assumption 2(i)
c_n, c_v, d_n, d_v; $b_R = \alpha_l/\mu$	split regression (34)	estimated ($\mathcal{I}_2$)
σ	the pass-through crossing (37); Prop. D.1	estimated ($\mathcal{I}_2$); external value only as robustness
α_m	share identity $s^M = \alpha_m(\sigma-1)/\sigma$	estimated ($\mathcal{I}_2$)
α_k, ρ_φ	timing moments (35) on σ-deflated data	estimated ($\mathcal{I}_2$)
α_l; s, D, ζ, μ	$\mu(\hat{\sigma})\hat{b}_R$; functions of $(\sigma, \alpha_l, \alpha_m)$	derived; delta method
technology ω	corrected residual; Markov innovation	estimated
survival propensity	probit on $(k, i, n^\star, v^\star)$	estimated ($\mathcal{I}_2$)
γ_n, γ_v; χ, ξ	formation regressions	extension ($\mathcal{I}_3$)

specific requirement. It needs a firm-level wage measure, the wage bill and employment. It needs a firm exit indicator for the survival propensity and standard deflators for capital and investment levels. The last three are standard in linked accounts.

Most VAT and e-invoice panels impose reporting thresholds, so the observed network truncates small links. The counts are then lower bounds while the value-weighted stocks are nearly complete, small links carrying little value, and robustness can be assessed by imposing harsher artificial thresholds on the observed panel and re-running the pipeline. And domestic panels miss imports, so the supplier set of an importing firm is incomplete. Where customs data exist they complete it, in the spirit of the imported-inputs estimates of Halpern et al. [2015], since imported varieties enter the same input price index.

Assumption 2 (Identification conditions)**.** .

(i) Connectedness and exogenous mobility: the network is connected, and link formation does not select on the pairwise match component ϑ_{ij} (Appendix C).

(ii) CES and constant markups: demand on both sides is CES with common σ, so prices

carry the constant markup $\mu = \sigma/(\sigma - 1)$.[21]

(iii) Timing and the law of motion: K, n and v are predetermined at t; the value-weighted stocks $n^\star$ and $v^\star$ are predetermined only at a lag.[22] Technology follows the AR(1) law $\omega_{it} = (1-\rho_\varphi)\ln\bar{\varphi}_i + \rho_\varphi\,\omega_{i,t-1} + \epsilon^\omega_{it}$ around a permanent type $\ln\bar{\varphi}_i$ drawn at entry, with the innovation ϵ^ω_{it} orthogonal to the firm's $t-1$ information. The permanent type is a firm fixed effect in the law of motion, so the innovation is formed by a within-firm transformation. Endogenous exit selects on ϵ^ω_{it}, so the transformation is one whose inclusion condition is predetermined (Section 7).

(iv) Servicing and acquisition technology: servicing labor is linear in the stocks, and acquisition labor is, to first order, proportional to the gross formation flows.

(v) Aggregation: portfolio heterogeneity enters value and the policies only through $(n^\star, v^\star)$ (Section 2).

(vi) Measurement: the network objects are measured without error. Classical error in them attenuates the corrections, so the estimated biases are conservative.

Each stage below names the parts of Assumption 2 it uses. Where a condition fails, the stage that leans on it fails alone, and the ledger says which.

7.2 The labor split

We suggest three ways to obtain the required production labor input. The best would be to have access to detailed occupation data to measure the production labor input directly. The model supplies another option from objects in the panel that can be estimated. Finally, production labor can be assumed to be a fraction of total observed labor.

From the structure of the model, production labor satisfies $wL^p_{it} = (\alpha_l/\mu)R_{it}$ exactly, for any s. The revenue coefficient $b_R \equiv \alpha_l/\mu$ defined by this identity is the object the split estimates. Servicing labor is $c_n\,n_{it} + c_v\,v_{it}$ by technology. Acquisition labor is, to first order, proportional to the gross formation flows, $L^n_{it} \approx d_n F^n_{it}$ and $L^v_{it} \approx d_v F^v_{it}$, which the link panel observes [Assumption 2(iv)]. The split regression is

$$L_{it} = b_R\,\frac{R_{it}}{w} + c_n\,n_{it} + c_v\,v_{it} + d_n F^n_{it} + d_v F^v_{it} + e_{it}, \qquad b_R = \frac{\alpha_l}{\mu}, \tag{34}$$

where w is the firm's average wage, its wage bill divided by its employment, both $\mathcal{I}_2$ objects.

[21] The deflators and the share identities are exact only under this structure.

[22] The counts are fixed by acquisition effort at $t-1$ and are predetermined at t, but the value of each relationship at t is a contemporaneous choice that moves with ω_{it}, so $n^\star_{it} = \sum_j r_{ijt}$ inherits the current innovation. The estimator uses $n^\star_{t-1}, v^\star_{t-1}$ in the timing moments and in the survival propensity. $n^\star_t$ and $v^\star_t$ appear only in the first-stage inversion as arguments of the investment policy.

The model's common numeraire wage corresponds to the economy-wide average, so using the firm's own wage lets the revenue leg absorb firm-level wage variation. Production labor is then constructed by subtraction, $\hat{L^p}_{it} = L_{it} - \hat{c}_n\, n_{it} - \hat{c}_v\, v_{it} - \hat{d}_n F^n_{it} - \hat{d}_v F^v_{it}$, so the constructed series keeps the residual variation e_{it}.[23]

α_l is identified by the share moment $b_R = \alpha_l/\mu$, not by a production-function moment. $L^p \propto R$ holds exactly in the model, so no quantity-side moment can separate α_l from the output scale, and the production-function moments below identify only α_k and α_m.

The split is estimated by ordinary least squares, with no instrument. Consistency rests on Assumption 2(iv). The servicing legs are technological and the formation flows proxy acquisition labor to first order, so the residual is measurement. Identification comes from the artisan-versus-networker variation, firms with the same revenue and different link counts. Acquisition effort that scales with revenue without producing links would violate it.

Real firms also employ workers in none of the three activities, e.g., administration, finance, and accounting. We model them as a time-invariant overhead stock $\bar{\ell}_i$, drawn at entry, so total labor adds a fourth component and the firm pays the fixed cost $w\bar{\ell}_i$ every period. Overhead changes no first-order condition. The split regression (34) gains a firm fixed effect and is estimated within firm and the recovered fixed effect $\hat{\bar{\ell}}_i$ is itself an estimate of overhead. The exit threshold gains a firm-specific fixed cost, so the survival propensity conditions on $\hat{\bar{\ell}}_i$ alongside the stocks.[24]

7.3 Path 1: the staged pipeline

The pipeline has five stages. Stages A through C need no parameters. Stage D estimates every parameter, and its only fixed point is a scalar equation in σ whose properties we state as a proposition. Stage E applies the corrections. The sequence is acyclic. No stage consumes an output produced by a later stage.

In the estimation literature the elasticity of substitution σ is usually a nuisance, calibrated or set aside, and taking it from a demand system is standard practice. Here it is a load-bearing parameter. Every bias correction is a known function of it, so its value propagates through the entire pipeline. The two factor shares are $\alpha_l = \mu(\sigma)\hat{b}_R$ and $\alpha_m = \mu(\sigma)\bar{s}^M$, with $\bar{s}^M$ the cross-sectional mean of the intermediate revenue share s^M_i, so both output elasticities, the returns to scale $s = \alpha_l + \alpha_m$, the dampening scalar D, and the markup $\mu = \sigma/(\sigma - 1)$ are all functions of σ. Without it the network delivers the deflated data

[23] In finite samples the subtraction can turn negative. Such rows have to be trimmed from the quantity moments, and their share in the data is itself a diagnostic of the split.

[24] The slope coefficients of (34) are identified from within-firm variation. The recovered fixed effect is then average labor net of the fitted components at those slopes, so it estimates the level of overhead under one restriction the model supplies, that overhead is the only time-invariant component of the split.

and every revenue share but not a single elasticity. Everything the estimator reports about technology rests on pinning σ down, and the network cannot read it off any one firm's own accounts (Proposition D.2). The pipeline's one network-internal route to σ is the pass-through crossing below. That crossing pins σ down only under decreasing variable returns and only in a large panel. In practice we take σ from outside, a demand system or trade elasticities, and read the crossing as a check on it.

Stage A: network objects (uses $\mathcal{I}_1$; no parameters). Build A, the counts (n, v), the flows (F^n, F^v), the failure schedule $\hat{\delta}(\cdot)$ with (h_a, η_δ), and the value-weighted stocks $(n^\star, v^\star)$, as in the ledger.

Stage B: the seller–buyer decomposition (uses $\mathcal{I}_1$; Assumption 2(i)). Estimate the two-way fixed-effects regression (27) on log bilateral values over the largest connected component including link-age dummies. Somaini and Wolak [2016] give an algorithm for two high-dimensional effects at panel scale. Recover the seller effects $\tilde{\psi}_i$ and the buyer effects θ_j with the seller effects normalized to mean zero without loss of generality. The seller effect is the output price and the buyer effect is the exact input price index of Lemma 1, and the input deflator is built from the buyer effect. The estimated effects carry sampling noise that is serially correlated because links persist, so it enters the constructed deflators as a correlated measurement error. We remove the first-order channel by cross-fitting [Chernozhukov et al., 2018]. Buyers are assigned to K folds ($K = 10$ in our implementation), the decomposition is estimated once per excluded fold, and every supplier effect entering a buyer's deflator is taken from the regression that saw none of that buyer's links, so the covariance between a buyer's deflator and its own residuals is zero exactly and by construction. The own effect $\tilde{\psi}_i$ and the buyer effect θ_j come from the full sample, whose own links are their identifying data. The cost is one fold's share of the links per estimation, about a five percent increase in the standard deviation of an effect at $K = 10$. The role of the remaining noise, and why the cross-fit suffices for the headline parameters, is taken up at the close of the section.

Stage C: the labor split (uses $\mathcal{I}_2$; Assumption 2(ii), (iv)). Estimate (34), obtaining $(\hat{b}_R, \hat{c}_n, \hat{c}_v, \hat{d}_n, \hat{d}_v)$ and the constructed $\hat{L^p}_{it}$.

Stage D: the parameters (uses $\mathcal{I}_2$; Assumption 2(ii), (iii)). For a candidate σ, form the deflators and deflated quantities

$$\tilde{p}_i(\sigma) = -\frac{\tilde{\psi}_i}{\sigma - 1}, \qquad \widetilde{P^M}_i(\sigma) = -\frac{x_i - \theta_i}{\sigma - 1},$$
$$y^d(\sigma) = y - \tilde{p}(\sigma), \qquad m^d(\sigma) = x - \widetilde{P^M}(\sigma),$$

with the input index the exact form of Lemma 1. Impose the two share identities $\alpha_l(\sigma) = \mu(\sigma)\hat{b}_R$ and $\alpha_m(\sigma) = \mu(\sigma)\bar{s}^M$, with $\mu(\sigma) = \sigma/(\sigma - 1)$. The intermediate share is $s^M = \alpha_m/\mu$ exactly under the maintained CES structure (Section 4.3), so α_m is a share object and the quantity-side moments carry no separate information about it (Proposition D.2 below). The remaining parameters (α_k, ρ_φ) and the elasticity σ are then identified together by the pass-through system below, and the timing moments become over-identifying content.

The technology residual at the candidate is $\omega_{it}(\alpha_k; \sigma) = y^d_{it}(\sigma) - \alpha_l(\sigma) \ln \hat{L}^p{}_{it} - \alpha_k k_{it} - \alpha_m(\sigma) m^d_{it}(\sigma)$. Its law of motion is AR(1) around the entry-drawn type [Assumption 2(iii)], so the quasi-difference $a_{it} = \omega_{it} - \rho_\varphi \omega_{i,t-1} = (1 - \rho_\varphi) \ln \bar{\varphi}_i + \epsilon^\omega_{it}$ still carries the firm fixed effect.

To remove it, first form the innovation by the forward pairwise difference $\Delta a_{it} = a_{it} - a_{i,t+1} = \epsilon^\omega_{it} - \epsilon^\omega_{i,t+1}$, which nets out the type and whose two-period window is deterministic. Second, trim own-rows whose ex-ante survival propensity, the logit of continuation on the $t - 1$ state including lagged revenue, is in the bottom decile, so the rows that remain have inclusion probability near one and contribute no selection. Selection bias needs rows whose inclusion was in doubt, so dropping the rows with low ex-ante survival propensity removes the correlation between survival and the innovation. The timing moments are then

$$\mathbb{E}\big[\Delta a_{it}(\alpha_k; \sigma) \left(k_{it},\ \ln v_{i,t-1},\ i_{i,t-1}\right)'\big] = 0, \tag{35}$$

on the trimmed sample. These hold at the truth and are reported as specification tests; they are not the estimating equations for (α_k, ρ_φ), for a reason we return to at the end of the stage. The instruments are predetermined raw data, capital and the firm's own lagged supplier count and investment, not the constructed lagged deflator m^d_{t-1}, which would carry the serially correlated first-step sampling noise of the estimated effects (discussed at the close of the section). The first stage purges measurement error by a flexible regression of y^d on $(k, \ln \hat{L}^p, m^d, i, \ln n^\star, \ln v^\star)$, where i is log investment, the proxy, and the tuple is the augmented state of Proposition 5(iv), with a full second-order polynomial and year effects as the default basis. The lagged stocks $(n^\star_{t-1}, v^\star_{t-1})$ enter as additional orthogonality conditions whose sample violations are specification tests, not as states in the law of motion. Under the

DGP the technology process is exogenous, and a nonzero stock coefficient is evidence that the inversion failed [Assumption 2(iii)].

Elasticity σ is identified by the pass-through of the input price index into the output price, a price-to-price relation distinct from the scale dampening $1/D$. The pricing identity of the model, written in seller-effect units, is the exact within-year relation

$$\tilde{\psi}_i = \alpha_m \ln S_i + (1-\sigma)(1-s)\, y_i - (1-\sigma)\big(\omega_i + \alpha_k k_i\big), \qquad \ln S_i = x_i - \theta_i, \tag{36}$$

in which the coefficient on the observable index $\ln S_i$ is α_m alone, the scale denominator D having cancelled against the demand feedback. Form the candidate residual $u_i(\sigma) = \tilde{\psi}_i - \alpha_m(\sigma) \ln S_i - \gamma_y(\sigma)\, y_i - \gamma_k k_i$, with $\gamma_y(\sigma) = (1-\sigma)(1-s(\sigma))$ imposed by the same share logic that fixes $\alpha_m(\sigma)$, and $\gamma_k = -(1-\sigma)\alpha_k$ a free coefficient estimated inside this block. At the truth $u_i = (\sigma-1)\omega_i$ up to sampling noise, so u_i inherits the same AR(1)-around-the-type structure as ω and takes the same forward-difference-and-trim transform.

The instrument must be supplier-side, since own-size instruments cannot see the α_m direction (Proposition D.2). We use $\bar{k}^S_{i,t-1}$, the mean over firm i's current supplier set of the suppliers' lagged log capital, which is other firms' predetermined data and is orthogonal to firm i's own sampling errors and innovations by construction. Two features of its construction carry real weight. The mean is *expenditure-share weighted*: $\ln S_i$ is a share-weighted aggregate of supplier prices, so a supplier accounting for forty percent of i's input bill and one accounting for half a percent are not equally informative about it, and an unweighted average diversifies away most of the supplier cost variation as the network thickens. The shares are ratios within the buyer, so i's own scale and productivity cancel and only supplier-side composition survives. The aggregate must also be *re-centered within year*: averaging a year-demeaned firm variable over a firm's suppliers is not mean zero, because being a supplier is a selected state, and the residual mean multiplies the mean of Δu and adds a level term to (37) with no economic content. Close the stage with the scalar equation

$$M(\sigma) \equiv \mathbb{E}\big[\Delta u_{it}(\sigma)\, \bar{k}^S_{i,t-1}\big] = 0, \tag{37}$$

where $\Delta u_{it}(\sigma) = \big[u_{it}(\sigma) - \rho_\varphi u_{i,t-1}(\sigma)\big] - \big[u_{i,t+1}(\sigma) - \rho_\varphi u_{it}(\sigma)\big]$, the same forward pairwise difference as in the timing moments (35), on the trimmed sample. $\hat{\sigma}$ solves (37). Then $\hat{\alpha}_l = \mu(\hat{\sigma})\hat{b}_R$, $\hat{\alpha}_m = \mu(\hat{\sigma})\bar{s}^M$, $\hat{s} = \hat{\alpha}_l + \hat{\alpha}_m$, $\hat{D} = \hat{\sigma} - (\hat{\sigma}-1)\hat{s}$, and $\hat{\zeta} = \hat{\alpha}_m/(\hat{\sigma}-1)$, together with $(\hat{\alpha}_k, \hat{\rho}_\varphi)$ from the pass-through system at $\hat{\sigma}$, so the two paths deliver the same parameter vector.

Both the pass-through system and the timing moments can deliver (α_k, ρ_φ), and we take them from the former. The two differ in what sits on the left-hand side. The pass-through

residual u_i is built from the seller effect, an estimated *price*, and inherits only the sampling error of that effect. The timing residual ω_{it} is built from recovered productivity, which additionally carries the labor split $\hat{L^p}$ and the input price index, each estimated and each with serially correlated error. Since the AR term instruments productivity with its own lag, that error enters the moment on both sides and the quasi-difference reads part of it as mean reversion: the classical errors-in-variables attenuation, arriving through the objects the timing residual adds and the price residual does not. The timing moments remain in the system as over-identifying restrictions, and Stage E still uses ω_{it} as the productivity series.

An incorrect σ enters the whole pipeline through a single scalar, and the crossing $M(\sigma)$ has the truth as a root, locally unique when $M'(\sigma_0) \neq 0$, with $\hat{\sigma}$ consistent; a supplier-side instrument is forced because every within-firm σ contrast is vacuous. These properties, the rescaling that carries misdeflation through the pipeline, and the degeneracy, are collected in Appendix D.

The effects $\tilde{\psi}$ and θ are estimated, so the deflators and $\ln S$ carry first-step sampling noise. We recommend a bootstrap that holds the network fixed, resamples the match residuals of (27), and re-runs Stages B through E in each draw. Stacking the fixed-effects normal equations into the moment system is the formal alternative [Kline et al., 2020]. An external estimate of σ, from a demand system or trade elasticities [e.g. Hottman et al., 2016], replaces the crossing (37) and shortens the pipeline by one equation. It is the default: the crossing identifies σ only under decreasing variable returns and only in a large panel, neither of which the data certify, so we anchor σ externally and use the crossing as a within-network check, comparing the two and reporting both sets of corrected objects if they differ materially.

Stage E: corrections and cross-validation (uses Stage D). Apply the static correction at the estimates,

$$\hat{\boldsymbol{\omega}}^{\,\mathrm{S}} = \frac{\hat{D}}{\hat{\sigma}-1}\Big(\mathbb{I} - \tfrac{\hat{\alpha}_m}{\hat{D}} A^{\top}\Big)\tilde{\boldsymbol{\psi}} - \hat{\zeta}\,\widetilde{\ln \boldsymbol{v}} + (1-\hat{s})\,\widetilde{\ln \boldsymbol{n}} - \hat{\alpha}_k \tilde{\boldsymbol{k}}, \tag{38}$$

the value-only recovery of Proposition 4 with the capital term subtracted under $\mathcal{I}_2$. The Stage-D residual $\hat{\omega}_{it}$ is the dynamic recovery of the same object from disjoint moments, and the two should coincide firm by firm up to sampling error. Their divergence is the pipeline's summary specification test (Section 7.6). Standard errors for every derived object, $(\hat{s}, \hat{D}, \hat{\zeta}, \hat{\mu})$ and the corrections, follow by the delta method within the bootstrap.

7.4 Path 2: one joint GMM system

Path 2 estimates the parameter vector $\boldsymbol{\theta} = (\sigma, \alpha_k, \alpha_l, \alpha_m, \rho_\varphi, c_n, c_v, d_n, d_v)$, with true value $\boldsymbol{\theta}_0$, the bold $\boldsymbol{\theta}$ distinct from the buyer effect θ_j, in a single step from the stacked moment vector

$$\mathbb{E}\left[\begin{array}{l} \left(L_{it} - \frac{\alpha_l}{\mu(\sigma)}\frac{R_{it}}{w} - c_n\, n_{it} - c_v\, v_{it} - d_n F^n_{it} - d_v F^v_{it}\right) \cdot \left(\frac{R_{it}}{w},\, n_{it},\, v_{it},\, F^n_{it},\, F^v_{it}\right)' \\ s^M_{it} - \alpha_m \dfrac{\sigma - 1}{\sigma} \\ \Delta a_{it}(\boldsymbol{\theta}) \cdot \left(k_{it},\ i_{i,t-1},\ \ln n^\star_{i,t-1},\ \ln v^\star_{i,t-1}\right)' \\ \Delta u_{it}(\boldsymbol{\theta}) \cdot \bar{k}^S_{i,t-1} \end{array}\right] = \mathbf{0}, \qquad (39)$$

where $a_{it}(\boldsymbol{\theta}) = \omega_{it}(\boldsymbol{\theta}) - \rho_\varphi \omega_{i,t-1}(\boldsymbol{\theta})$ is the quasi-difference of $\omega_{it}(\boldsymbol{\theta}) = y^d_{it}(\sigma) - \alpha_k k_{it} - \alpha_l \ln \hat{L}^p{}_{it}(\boldsymbol{\theta}) - \alpha_m m^d_{it}(\sigma)$, Δ is the forward pairwise difference on the exit-trimmed sample of Stage D, $u_{it}(\boldsymbol{\theta})$ is the pass-through residual of (36), $\bar{k}^S_{i,t-1}$ is the supplier-capital instrument, and $\hat{L}^p(\boldsymbol{\theta})$ is the subtraction split evaluated at the candidate (c_n, c_v, d_n, d_v). Eleven moments identify nine parameters. The two surplus conditions, the stock orthogonalities, are exactly the specification tests of Path 1, now delivered as a J statistic. The deflators are the same functions of $(\sigma, \tilde{\psi}, \theta)$ as in Path 1, so the system is smooth in $\boldsymbol{\theta}$ through the two scalars of Lemma D.1.

Proposition 6 (Identification of the joint system)**.** *Under Assumption 2, $\boldsymbol{\theta}_0$ solves the population system* (39)*. Given σ, the remaining blocks are recursively identified. The split moments identify $(\alpha_l/\mu(\sigma), c_n, c_v, d_n, d_v)$ under a full-rank condition on the split regressors. The share moment then reads off α_m. The timing moments identify (α_k, ρ_φ). Elasticity σ is identified by the pass-through moment, which supplies the one cross-firm restriction the within-firm blocks cannot (Proposition D.2). Local identification of the full vector therefore reduces to the scalar condition $M'(\sigma_0) \neq 0$ of Proposition D.1. The population Jacobian of* (39) *is block-triangular apart from the σ column, and its determinant equals $M'(\sigma_0)$ times a product of nonzero block determinants (Appendix B).* (Proof: Appendix E.)

The joint system stacks both blocks, so the choice of which block reports (α_k, ρ_φ) in Path 1 does not arise here: Path 2 uses all of the moments at once, and the errors-in-variables argument of Stage D shows up instead as the relative weight the efficient metric places on the timing block. The two paths are the same information arranged differently. Path 1 is Path 2 under a block-recursive weighting, so both are consistent under the same assumptions; Path 2 is asymptotically efficient and delivers the specification tests as one J statistic, while Path 1 is transparent, diagnosable stage by stage, and cheaper to compute. We recommend running Path 1 first and Path 2 as the efficient one-step check. As in Path 1, an external σ turns

the pass-through moment into an over-identifying check and (39) estimates the remaining eight parameters. Inference in both treats $\tilde{\psi}$ as generated: the bootstrap of Stage D, or the stacked normal equations [Kline et al., 2020].

7.5 What $\mathcal{I}_2$ delivers, and the $\mathcal{I}_1$ boundary

Under $\mathcal{I}_2$, with the substitution elasticity σ supplied from outside, the pipeline delivers the technology elasticities $(\alpha_k, \alpha_l, \alpha_m)$ and the derived (s, D, ζ, μ); technology ω purged of variety, propagation, the customer channels, and exit selection; the accumulation primitives $\delta(\cdot), h_a, \eta_\delta$ and the value-weighted stocks; the labor-split parameters (c_n, c_v, d_n, d_v); and the survival propensity.

The network panel alone ($\mathcal{I}_1$) is not an estimator. It runs Stages A and B and every accumulation object, and it supports the variety and propagation structure of the corrections. However, it does not identify any parameter scale. The intermediate revenue share needs total sales, the quantity moments need capital and investment, and the labor split needs employment, all $\mathcal{I}_2$ objects, so σ is not identified within $\mathcal{I}_1$ and the crossing (37) cannot run. Given (σ, α_m, s) from elsewhere, $\mathcal{I}_1$ delivers the propagated, variety-corrected cost fundamental $\omega^c = \omega + \alpha_k \tilde{k}$ of (17), the supply-side object that governs prices, costs, and survival, and a partial exit correction conditioning on the stocks and age but not capital.

7.6 Over-identifying tests

The framework over-identifies, and each maintained assumption maps to a test on the same panel. *(i)* The uncorrected residual should be correlated along A and with the supplier count. The corrected residual should not (Corollary 1). *(ii)* Estimated persistence should move toward the technology rate when the control function conditions on the stocks. *(iii)* The Stage-E value-only recovery and the Stage-D control-function residual estimate the same object from disjoint moments and should coincide firm by firm. In Path 2 the same content arrives as the J test. *(iv)* The exit hazard should fall with both observed stocks given recovered productivity (Proposition 5(iii)). *(v)* The downstream and upstream link-survival hazards should estimate the same failure schedule, since a link's failure is one event. *(vi)* Where occupational labor exists, the split validates leg by leg. Where observed input prices exist, the constructed deflators can be compared with them directly, a test of the CES variety elasticity $1/(1-\sigma)$ of (3), which a richer aggregator would relax. One caveat applies throughout. The corrections treat the network objects as measured without error [Assumption 2(vi)]. Classical error attenuates them, so the estimated biases, and the gaps the tests detect, are conservative. The parallel problem for capital is treated by Collard-Wexler

and De Loecker [2016].

Generated-deflator noise. The seller and buyer effects are estimated, so the constructed deflators carry sampling noise. Two features make it more than variance inflation. The noise is serially correlated, because a firm's links persist, so it survives the innovation transform. And it is correlated with the data, because the same match components sit in the effects, in revenue, and in the shares. The default construction is the K-fold cross-fit over buyers of Stage B, the sample-splitting counterpart of the leave-out estimator of Kline et al. [2020] and the cross-fitting of Chernozhukov et al. [2018]. Where the leave-out algebra removes the own-observation term to first order through leverage corrections, the cross-fit removes it exactly and by construction, with no leverage computations, and the removal is robust to heteroskedasticity and to the serial dependence within a match that pair-level algebra treats only approximately. The exact pair-level leave-out remains implemented and serves as the validation at moderate scale, where its dense linear algebra is feasible.

The design keeps this problem small where it matters. The headline parameters barely touch the constructed deflators. α_m comes from the share identity, σ from the pass-through moment built on $\tilde{\psi}$, θ, and raw data, and α_l from the labor split. The deflators enter only the (α_k, ρ_φ) timing step. A bootstrap correction of the timing moments is available for the researcher who needs that step clean. It rebuilds the effects in a synthetic network that preserves the link structure, propagates the induced noise into revenue, stocks, and shares, and subtracts the resulting bias. We keep it as an option, not a default, because the cross-fit effects suffice for the headline parameters. One data-requirements point closes the account. The noise variance scales with the average network degree, not the sample size, so larger panels do not remove it. Denser networks, longer panels, and lower match dispersion do.

8 Numerical evidence

In this section we generate simulated data from the networked firm of Section 2 and ask what our estimation strategy from Section 7 recovers. We also run the data through two standard estimators ignoring the network and ask what they recover on the same data. The section closes by moving the parameters of the economy one at a time to evaluate which outcomes are sensitive to parameter choices.

Table 3: The data-generating process: parameters.

		value
elasticity of substitution	σ	4
capital elasticity	α_k	0.10
labor elasticity	α_l	0.40
intermediate elasticity	α_m	0.50
persistence of productivity	ρ_φ	0.70
innovation standard deviation	$\mathrm{sd}(\nu)$	0.15
imposed labor split	λ	0.60
implied		
variable returns	$s = \alpha_l + \alpha_m$	0.90
scale dampening	$D = \sigma - (\sigma - 1)s$	1.30
markup	$\mu = \sigma/(\sigma - 1)$	1.33
variety elasticity	$\zeta = \alpha_m/(\sigma - 1)$	0.17
relationships and matching		
Pareto shape of match quality	k	2.5
birth depth of a link	h_0	0.30
deepening rate	ρ_h	0.60
failure hazard, scale	$\bar{\delta}$	0.30
failure hazard, size elasticity	η_δ	0.50
exogenous firm exit hazard	δ_f	0.02
panel		
firms	N	2,000
years	T	15

Notes: the implied rows are functions of the parameters above and are not set separately. The imposed labor split λ is the estimator's stage C setting rather than a primitive of the generator. The relationships panel lists the distributional choices the model leaves free, with the deepening profile $h_a = 1 - (1 - h_0)\rho_h^a$ at tenure a.

8.1 The economy we simulate

The data-generating process is the networked-firm economy of Section 2, simulated under the linked information set $\mathcal{I}_2$. The capability types and their dynamics, the two relationship stocks, link deepening, size-dependent failure, and the exit margin are implemented as in Section 2. The simulation therefore pins down only what the model leaves free. Match quality is assumed tyo be Pareto; a link is born at a fraction h_0 of its match size and approaches it geometrically in tenure; and the failure hazard is $\delta(z) = \bar{\delta}\, z^{-\eta_\delta}$, truncated to lie between 0.02 and 0.95 so that the smallest links do not fail with certainty and the largest never become immortal. Table 3 collects these choices together with the technology parameters and the

Table 4: The simulated economy.

	value
sd(ln revenue)	1.10
sd(ln employment)	0.83
sd(ln capital)	0.91
mean suppliers per buyer	20.3
sd(ln suppliers)	0.78
mean share of the largest supplier	0.28
effective number of suppliers (1/HHI)	5.9
mean relationship length (years)	4.9
annual exit rate	0.083
top one percent share of revenue	0.10

Notes: $N = 2{,}000$ firms over $T = 15$ years, seed 7. The effective number of suppliers is the inverse Herfindahl of within-buyer expenditure shares, averaged over buyers.

objects they imply.

The econometrician sees this economy through data that might contain measurement error. Reported revenue carries five percent log measurement error, reported employment and capital are measured with eight percent error, and log wages are dispersed with standard deviation 0.30. The measurement error is a modeling requirement. With exact accounts the gross-output revenue regressions are degenerate, as the intermediate first-order condition makes expenditure proportional to revenue, and the regression collapses onto the share identity.

Two components of the data generator are deliberate approximations. First, capital follows a partial-adjustment rule that is monotone in productivity rather than the solution of the full dynamic program. Monotonicity given the other states is the invertibility property the proxy estimators rely on, and the rule preserves it exactly. When a proxy estimator fails on this data, it is therefore not because the generator broke the assumption its inversion needs. Second, acquisition labor is chosen against the current marginal value of a relationship with idiosyncratic cost noise, a static stand-in for the forward-looking lifetime value. Chosen labor therefore loads on the acquisition capability, so the data pose the same transmission problem the estimation of the formation technology faces. The formation technology itself holds exactly, and the cost noise is what lets a marginal-value instrument shift labor orthogonally to the capability.

Table 4 reports the moments of the resulting economy. It is disciplined to empirically reasonable magnitudes. The standard deviation of log revenue is 1.10. A buyer holds twenty suppliers on average, the largest accounting for 28 percent of the input bill, with an effective number (inverse Herfindahl) of about six. Relationships last five years on average, the annual

exit rate is eight percent, and the top one percent of firms hold ten percent of revenue.

8.2 Implementation

Each period the generator solves the within-period price fixed point on the realized network. The solver iterates on the full price vector with damping, computes every buyer's input price index as the exact CES aggregate over its realized suppliers, and runs a nested revenue fixed point inside each pass. The outer loop converges at a maximum log-price change of 10^{-11} and the inner loop at a relative revenue change of 10^{-12}. Every price, expenditure share, and transaction value is therefore an equilibrium object.[25] The generator holds the exact nonlinear price system where the estimator's deflators and propagation operator are first-order approximations, so the Monte Carlo prices the estimator's approximation error along with measurement and sampling noise.

A Monte Carlo cell consists of fifty independent draws. A draw is a complete economy, with new capability draws, a new realized network, and new measurement error. Cross-draw means and standard deviations are reported throughout. Single-draw diagnostics use a chosen single seed's panel, the reference economy of Table 4.

The estimator runs on every draw in a fixed configuration. The seller–buyer decomposition carries the tenure profile. The deflators use the $K = 10$ buyer-fold cross-fit of Section 7.3. The intermediate share is the expenditure-weighted aggregate. Firm-years with nonpositive reported value added are not in the moment sample. Production labor is imposed as a fraction λ of the median wage-bill share, with the estimated split recorded beside it. The timing and pass-through moments use the forward pairwise difference with a ten percent ex-ante exit-propensity trim. The elasticity of substitution σ is supplied externally at its true value, and the internal crossing is computed and recorded as a diagnostic on every draw.[26]

Table 5: The estimator in the baseline economy.

	mean	sd across draws	truth
α_m (share identity)	0.4995	0.0004	0.500
α_l (imposed λ)	0.3935	0.0018	0.400
α_k	0.111	0.023	0.100
ρ_φ	0.712	0.030	0.700
corr$(\hat{\omega}, \omega)$, Stage D	0.941	0.004	1
corr$(\hat{\omega}, \omega)$, value-only	0.913	0.010	1
share of rows with $\ell^p \leq 0$	0.002	0.000	0
over-identification t at the anchor	0.88	0.55	0
recorded diagnostics			
internal crossing	3.28	0.49	4.00
stage C split b_R	0.303	0.003	–

Notes: the economy of Table 4, estimator in the configuration of record. (α_k, ρ_φ) are the pass-through values, the ones the estimator reports. α_l comes through the imposed-λ stage C. The last row is the estimated split coefficient recorded beside the imposition.

8.3 What the estimator recovers

Table 5 reports what the estimator recovers on the baseline economy. We first discuss the elasticity estimates and then the recovered productivity series. Two of the three technology elasticities are measurements rather than estimates. Given the anchored σ, the share identities deliver $\alpha_m = \mu(\sigma)\,\bar{s}^M$ from the expenditure share and $\alpha_l = \mu(\sigma)\,b_R$ from the labor split directly and both are precisely measured at their true values. The capital elasticity α_k and persistence ρ_φ are the two parameters the moments genuinely estimate, and again both are centered near their true values. On revenue data this succeeds because the network correctly supplies the missing prices. The constructed deflators turn deflated revenue into the quantity object that the timing and pass-through moments assume, and the standard dynamic machinery then applies unchanged.

The recovered productivity series correlates 0.94 with the underlying true productivities through the stage-D residual and 0.91 through the value-only recovery. The two recoveries

[25]The equilibrium converges and is unique. Each buyer's log input price index is a weighted average of its suppliers' log prices with weights summing to one, and it enters the price update with coefficient $\alpha_m/D < 1$, the damped propagation coefficient of the price system. The map is therefore a contraction, the within-period equilibrium exists and is unique, and the iteration reaches it from any starting point. The nested revenue system is linear with spectral radius below $\theta_M < 1$. The tolerances then state how exactly the equilibrium conditions hold at the reported solution.

[26]This version reports point estimates with their cross-draw dispersion. Per-draw sampling uncertainty will be added in a later version of the paper, using the match-residual bootstrap that holds the network fixed and, for the leave-out variance components at scale, a random-projection (Johnson–Lindenstrauss) approximation of the leverage terms in the spirit of Kline et al. [2020].

Table 6: What the pass-through moment says about σ, by network density.

mean degree	crossing	95% set	width	α_k (pass-through)
10.0	4.22	[3.82, 4.77]	0.95	0.082
16.2	3.80	[3.02, 5.42]	2.40	0.096
20.3	3.31	[2.53, 5.39]	2.86	0.102

Notes: the same economy at three values of the baseline relationship-failure rate, which is what governs equilibrium degree. Everything else is held fixed, and the last row is the panel of Table 4. The set $\mathcal{S}$ is defined in the text, with the endpoints interpolated between grid points.

are built from disjoint moments, so their agreement firm by firm is the cross-validation of Section 7.3, here passed on every draw. The sources of the correlation below one are various sources of measurement error. Regenerating the same economy with clean accounts and then replacing estimated objects with their true counterparts one at a time helps isolate the sources. With everything true, the identity returns a correct correlation of 1. Adding the reporting noise alone takes the correlation down to 0.979. That is a ceiling no estimator can beat, because the noise is in the data the econometrician is handed. Adding the estimated stage-B deflators on otherwise clean accounts takes it to 0.953, and the two error sources are almost orthogonal. The small remaining error is due to the noise in the labor split and parameter estimation error.

The over-identification t at the anchor does not reject on any draw, so the external anchor is consistent with the internal moments when the anchor is right. This raises the questions of what the network can say about the elasticity on its own and what the costs are of an incorrect anchor. The pass-through moment $M(\sigma)$ of Section 7.3 has a root at the truth, so the internal crossing is in principle an estimator. However its performance is governed by how close the sample moment sits to zero at the truth and by how steeply the profile passes through it. Neither is a property of the estimator alone. Both depend on the economy, and network density is the feature that moves them most. The profile can be shallow which suggests that a point estimate is the wrong object to report. We therefore report the set of elasticities the moment cannot reject, $\mathcal{S} = \{\sigma : |M(\sigma)/\widehat{\mathrm{se}}(\sigma)| < 1.96\}$, which is valid whether or not the moment identifies σ strongly.

Table 6 reports the crossing and the set across three network densities. Density does two opposite things. It sharpens the seller and buyer effects, which improves every deflator in the estimator, and it dilutes the supplier-cost instrument, which flattens the pass-through profile. For the elasticity, dilution wins. Across the draws of the baseline economy the crossing has mean 3.28 and standard deviation 0.49, and the set covers the true value in 97.5 percent of draws. On a network this dense, the crossing rules out a great deal without

Table 7: Estimates when the external anchor is wrong (truth $\sigma = 4$).

σ supplied	2.5	3.0	3.5	4.0	4.5	5.0	6.0
error in the anchor	−38%	−25%	−12%	0	+12%	+25%	+50%
exposed to the anchor							
markup $\hat{\mu}$	1.67	1.50	1.40	1.33	1.29	1.25	1.20
$\hat{\alpha}_l$	0.510	0.459	0.428	0.408	0.393	0.382	0.367
$\hat{\alpha}_m$	0.626	0.563	0.526	0.501	0.483	0.469	0.451
$\hat{D}$	0.80	0.96	1.12	1.27	1.43	1.59	1.91
$\hat{\alpha}_k$	0.218	0.152	0.117	0.102	0.082	0.073	0.060
invariant to the anchor							
$\hat{\gamma}_k = (\sigma - 1)\hat{\alpha}_k$	0.33	0.30	0.29	0.31	0.29	0.29	0.30
$\hat{\rho}_\varphi$	0.719	0.712	0.707	0.703	0.698	0.694	0.688
corr$(\hat{\omega}, \omega)$	0.77	0.89	0.88	0.85	0.81	0.79	0.74
over-identification t	−2.27	−0.79	0.16	0.95	1.28	1.64	2.13

Notes: the panel of Table 4 is unchanged. Only the value of σ handed to the estimator moves. True values are $\mu = 1.33$, $\alpha_l = 0.40$, $\alpha_m = 0.50$, $D = 1.30$, $\alpha_k = 0.10$ and $\rho_\varphi = 0.70$. α_k and ρ_φ are the pass-through values, the ones the estimator reports. The productivity correlation is the better of the Stage-D and value-only series. The last row is the t statistic of the pass-through moment evaluated at the supplied σ, the over-identification check the external policy prescribes. A value of $|t| > 1.96$ is a rejection. Two further draws at $\sigma \in \{3, 4, 5\}$ reproduce the invariance of $\hat{\gamma}_k$ and $\hat{\rho}_\varphi$ and the location of the productivity-correlation peak below the truth.

pinning the elasticity down. It is a diagnostic and is close to a reasonable substitute for an external anchor only on sparse networks.

Anchoring externally helps only if the external anchor is right. In Table 7 we vary the supplied anchors from 2.5 to 6 against a true value of 4. The exercise yields two clear results. The anchor sets the level of the elasticities, the markup and the dampening scalar, each through a closed-form expression. All of these comove with the supplied anchor. However the persistence and the productivity series barely move.

The split has a practical reading. The over-identification t catches anchor errors of roughly forty percent and larger. Because the exposure is analytic, the useful convention is to report the anchor-free objects such as the productivity series and its law of motion in addition to the elasticities, $\hat{\gamma}_k$, $\bar{b}_R$ and $\bar{s}^M$.

8.4 Sensitivity analysis on parameter values

We now vary the structural parameters, one at a time, rerun our estimator and report results from the GNR (gross output) and ACF-VA estimators.Here there is only one draw per cell, anchored at each economy's own true σ. Table 8 reports the parameter estimates and Table

Table 8: Both estimators across twenty economies: parameters.

		this paper				GNR (gross output)			ACF-VA (materials)	
what moves	value	α_l	α_m	α_k	ρ_φ	β_k	β_l	β_m	β_k	β_l
baseline	–	0.408	0.500	0.102	0.705	0.13	0.54	0.38	0.19	0.84
σ	2.5	0.416	0.500	−0.036	0.785	0.21	0.50	0.30	0.19	0.83
	6.0	0.405	0.500	0.221	0.752	0.04	0.69	0.42	–	–
s	1.0	0.499	0.499	0.516	0.837	−0.06	0.90	0.38	−0.11	1.45
(α_l, α_m)	$(0.3, 0.6)$	0.305	0.599	0.108	0.755	0.14	0.47	0.45	0.24	0.82
	$(0.5, 0.4)$	0.501	0.400	0.172	0.713	0.10	0.64	0.30	0.10	0.93
α_k	0.05	0.403	0.499	0.008	0.814	0.08	0.58	0.38	0.13	0.87
	0.20	0.403	0.500	0.666	0.829	0.15	0.59	0.38	0.16	1.01
ρ_φ	0.40	0.405	0.500	0.072	0.401	–	–	–	−0.31	1.64
	0.90	0.392	0.499	0.295	0.970	−0.03	0.83	0.38	−0.12	1.45
sd(ν)	0.08	0.399	0.499	0.095	0.714	0.15	0.51	0.38	−0.03	1.14
	0.25	0.406	0.499	0.134	0.689	0.08	0.63	0.38	0.16	0.92
sd of χ	0.35	0.403	0.499	0.121	0.717	0.11	0.55	0.38	0.15	0.89
	0.60	0.404	0.500	0.149	0.749	0.10	0.57	0.38	0.12	0.92
sd of φ	0.30	0.394	0.500	0.255	0.774	−0.17	1.10	0.38	−0.41	2.00
sd of ξ	0.35	0.405	0.499	0.112	0.693	0.11	0.57	0.38	0.14	0.93
all three	$(0.3, 0.8, 0.8)$	0.396	0.499	0.325	0.796	−0.39	1.49	0.38	−0.79	2.70
corr(φ, χ)	0.0	0.406	0.499	0.100	0.717	0.12	0.55	0.38	0.17	0.89
	−0.8	0.403	0.500	0.109	0.694	0.15	0.52	0.38	0.22	0.79
	+0.3	0.398	0.499	0.105	0.726	0.12	0.57	0.38	0.15	0.90
truth (baseline)		0.400	0.500	0.100	0.700	0.10	0.40	0.50	0.20	0.80

Notes: One economy per row, one parameter moved from the baseline of Table 4, one draw each, estimator in the configuration of record anchored at each economy's true σ. α_l is the split-based value. α_k and ρ_φ are the pass-through values. Dashes mark standard estimator fits that did not converge. Convergence is recorded, not waited out. For GNR the gross-output truth is the baseline row's. For ACF value added the truth is $\beta_k = \alpha_k/(1-\alpha_m)$ and $\beta_l = \alpha_l/(1-\alpha_m)$, which is $(0.20, 0.80)$ except where the factor split moves.

9 reports the properties of the recovered productivity series.

For parameters, our estimator splits into a robust half and an exposed half. The share parameters are measurement and read from identities, so they survive every economy. The dynamic pair is the exposed half. The persistence tracks its truth from 0.40 through 0.97 across its own rows. The capital elasticity is right in the interior and fails at the corners. It overshoots where its true value is large, 0.67 at a truth of 0.20, and where variable returns are constant, 0.52 at $s = 1$. At $s = 1$, the scale feedback that gives the pass-through system its identifying variation is absent and revenue carries no technology, so the capital moment has little to work with. It also drifts upward near the unit root and where the capability types are widely dispersed, 0.30 at $\rho_\varphi = 0.90$ and 0.33 when all three capability dispersions are raised at once (all three). In both rows the persistence estimate errs in the same direction, which makes sense since the pair is estimated jointly from the same quasi-differenced residual.

The estimated coefficients can be reasonable while the recovered productivity series can

Table 9: The recovered productivity series.

		ours, stage D			ours, value-only	GNR		ACF-VA (materials)		
what moves	value	corr	sd ratio	AR1	corr	corr	sd ratio	corr	sd ratio	AR1
baseline	–	0.94	1.06	0.76	0.91	0.67	0.94	0.70	1.55	0.74
σ	2.5	0.88	1.18	0.82	0.87	0.62	0.66	0.63	0.93	0.69
	6.0	0.41	1.02	0.78	0.44	0.61	1.31	–	–	–
s	1.0	−0.26	3.92	0.96	−0.27	0.41	1.87	0.41	2.98	0.82
(α_l, α_m)	$(0.3, 0.6)$	0.92	0.97	0.72	0.92	0.63	0.99	0.66	1.87	0.74
	$(0.5, 0.4)$	0.95	1.06	0.71	0.92	0.69	0.91	0.70	1.28	0.75
α_k	0.05	0.94	1.14	0.79	0.90	0.68	0.83	0.73	1.35	0.73
	0.20	−0.14	3.65	0.96	−0.15	0.57	1.61	0.55	2.61	0.83
ρ_φ	0.40	0.93	1.14	0.67	0.88	–	–	0.11	1.74	0.80
	0.90	0.52	0.92	0.81	0.57	0.48	0.90	0.40	1.41	0.83
sd(ν)	0.08	0.91	1.14	0.83	0.85	0.57	1.03	0.47	1.70	0.80
	0.25	0.93	0.91	0.68	0.95	0.70	0.87	0.75	1.46	0.73
sd of χ	0.35	0.95	1.04	0.74	0.92	0.66	0.94	0.68	1.52	0.74
	0.60	0.94	1.01	0.73	0.92	0.67	0.93	0.67	1.51	0.76
sd of φ	0.30	0.65	0.82	0.77	0.70	0.25	0.91	0.05	1.56	0.83
sd of ξ	0.35	0.94	1.04	0.74	0.92	0.65	0.93	0.66	1.48	0.75
all three	$(0.3, 0.8, 0.8)$	0.43	0.93	0.81	0.45	−0.12	1.27	−0.24	2.41	0.84
corr(φ, χ)	0.0	0.94	1.05	0.76	0.91	0.67	0.93	0.69	1.50	0.74
	−0.8	0.94	1.06	0.76	0.91	0.70	0.95	0.73	1.57	0.73
	+0.3	0.94	1.04	0.75	0.92	0.66	0.95	0.68	1.54	0.75

Notes: same economies and draws as Table 8. Here corr is the correlation of the within-year-demeaned series with the sealed true ω. The sd ratio is the ratio of the series' within-year standard deviation to the truth's, so 1 is unbiased dispersion. AR1 is the series' own first-order autocorrelation, which mixes the persistent type with the AR(1) around it and therefore exceeds ρ_φ even for the truth. Dashes mark non-converged island fits.

be wrong in dispersion and persistence. In Table 9 we report the correlation of the productivity series as well as its AR(1) coefficient (persistence) and the ratio of the within-year standard deviation to the true productivity series. The stage-D series correlation with the true productivity is at or above 0.88in fourteen of the twenty economies, and its dispersion ratio sits between 0.9 and 1.2. The series is close to unbiased in spread, not only in ranks. The six failures are exactly the economies where the capital elasticity is badly estimated, the $s = 1$ and $\alpha_k = 0.20$ corners, $\sigma = 6$, $\rho_\varphi = 0.90$, the high-dispersion productivity economy and the overall high dispersion economy (all three).

The static correction subtracts $\hat{\alpha}_k k$, so a bad capital elasticity propagates straight into the series, and in the two worst corners it also inflates the dispersion ratio to 3.9 and 3.7. Where α_k is right the series is right. The value-only column tracks the stage-D column within a few hundredths in every economy, which is a cross-validation across all the economies.

The standard estimators ignoring the network data present a clear contrast.[27] We report

[27]Results for other standard estimators are available from the authors.

results for GNR (gross output) estimator and the ACF-VA estimator with materials proxy on the baseline economy and for all the parameter variations. As seen in Table 8, GNR recovers the intermediate elasticity at 0.38 for the baseline economy. It sits at 0.38 wherever the baseline shares and elasticity hold and moves one for one with α_m/μ when either moves, 0.30 at $\sigma = 2.5$ and 0.45 when the true intermediate share rises. It responds to nothing else in the table because expenditure over revenue does not. The capital and labor coefficients absorb what the intermediate share leaves, and they degrade exactly where the network states are most dispersed. The capital coefficient turns negative at $s = 1$, in the high-dispersion productivity economy and in the all-three dispersion economy, reaching -0.39, while the labor coefficient climbs to 1.49. The estimator does not converge in the low-persistence economy, where its AR(1) stage has little signal to work with.

ACF on value added has the widest range of the three. At the baseline and in the factor-split and capability-correlation rows, it sits near its true value-added values. It collapses precisely where the network does the most work. The capital coefficient falls to -0.31 at low persistence, -0.41 in the high-dispersion productivity economy and -0.79 in the all-three economy, while the labor coefficient climbs to 1.64, 2.00 and 2.70 in the same rows, and the $\sigma = 6$ economy does not converge. The value-added subtraction inflates every static bias by the intermediate multiplier, and the proxy inversion omits the two relationship stocks, so where the stocks and the propagated supplier prices carry more variance, the omitted states load onto labor and push capital negative.

In the moments from productivity series of Table 9, the standard estimators never reach the level of our estimator. GNR tops out at a correlation of 0.70 and turns negative in the all-three dispersion economy. ACF on value added collapses in precisely the economies where the network does the most work, reaching -0.24 with a dispersion ratio of 2.4 in the all-three economy. Its dispersion ratio exceeds 1.3 in most rows. The productivity series from the ACF estimator overstates dispersion by thirty percent to a factor of three even where its correlation looks tolerable.

Persistence completes the picture. The autocorrelation of the standard estimator series barely responds to the truth, sitting at 0.80 when the true persistence is 0.40 and 0.83 when it is 0.90. That is the persistence bias of Proposition 5 visible in series form. Our network estimator moves in the right direction, 0.67 in the low-persistence economy and 0.81 in the high one. One negative result is worth recording. Moving the correlation between the productivity and customer-acquisition capabilities from -0.8 to $+0.3$ changes essentially nothing for either estimator. The dispersion of the capabilities is what matters, not their correlation.

The takeways from these simulations are that results from identities are robust every-

where and what is estimated from dynamic moments is as good as its identifying variation. A mispecified proxy fails in proportion to how much the network matters.

9 Application to production network data

[TBD]

10 Conclusion

Standard productivity estimation describes a firm that buys inputs at common prices and sells into anonymous demand. A firm in a production network is not that firm. It has a supplier base and a customer base both of which are the result of firm decisions, built with labor, and slow to depreciate. We characterize in closed form what standard estimation recovers from the data generated by networked firms. The organizing object is effective productivity, technology compounded with supplier breadth. Because the supplier stock lowers marginal cost in exactly the place technology does, common-deflated estimation cannot separate them, and what propagates downstream is effective productivity rather than technology. The customer stock contaminates a second set of objects, the investment proxy, measured persistence, and exit selection, because it is a state the standard estimator omits.

The constructive half of the paper is that the firm-to-firm panel that exposes these biases also corrects them, and the central correction needs no price data. A two-way decomposition of bilateral sales delivers each firm's relative output price from the seller effect and its exact input price index from the buyer effect, so deflation becomes the first step of estimation rather than a maintained assumption. We give two identification paths that deliver the same parameters from a linked panel, and Monte Carlo evidence that the pipeline recovers technology where standard estimators degrade in proportion to how much the network matters.

Our results are of interest beyond the estimation of the production function itself. They apply anywhere measured residual productivity is an object of interest. Research on misallocation and productivity dispersion often links the spread of measured residuals to the variation of technology. Work on supply-chain spillovers reads the comovement of residuals across linked firms as a technological externality. Here, prices propagated into marginal cost can produce the same relationship when technologies are independent. The production approach to markups inherits the network's variation in shares and input prices, so estimates of market power built on it are exposed to the same biases. In trade, the gains from importing and from input variety currently arrive in measured productivity, which the variety term

makes explicit and removes. Our results on the dynamics suggest revenue-based measures load on the accumulated customer base, so growth and survival regressions that treat them as technology may actually be viewing demand capital instead, and event studies of shocks that move the network will attribute relationship reallocation to technical change unless they control for the relationship stocks. In each case the requirement is the same, a buyer-supplier panel with persistent link identifiers, and the structural content migrates from the estimator to measurement.

References

John M. Abowd, Francis Kramarz, and David N. Margolis. High wage workers and high wage firms. *Econometrica*, 67(2):251–333, 1999.

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

Daniel A. Ackerberg and Jan De Loecker. Production function identification under imperfect competition. CEPR Discussion Paper 19640, Centre for Economic Policy Research, 2024.

Daniel A. Ackerberg, Kevin Caves, and Garth Frazer. Identification properties of recent production function estimators. *Econometrica*, 83(6):2411–2451, 2015.

Martyn J. Andrews, Leonard Gill, Thorsten Schank, and Richard Upward. High wage workers and low wage firms: Negative assortative matching or limited mobility bias? *Journal of the Royal Statistical Society: Series A*, 171(3):673–697, 2008.

Pol Antras, Teresa C Fort, and Felix Tintelnot. The margins of global sourcing: Theory and evidence from us firms. *American Economic Review*, 107(9):2514–64, 2017.

Costas Arkolakis. Market penetration costs and the new consumers margin in international trade. *Journal of Political Economy*, 118(6):1151–1199, 2010.

John Asker, Allan Collard-Wexler, and Jan De Loecker. Dynamic inputs and resource (mis)allocation. *Journal of Political Economy*, 122(5):1013–1063, 2014.

Enghin Atalay. Materials prices and productivity. *Journal of the European Economic Association*, 12(3):575–611, 2014.

Enghin Atalay. How important are sectoral shocks? *American Economic Journal: Macroeconomics*, 9(4):254–280, 2017.

Bee Yan Aw, Mark J. Roberts, and Daniel Yi Xu. R&d investment, exporting, and productivity dynamics. *American Economic Review*, 101(4):1312–1344, 2011.

David Rezza Baqaee and Emmanuel Farhi. The macroeconomic impact of microeconomic shocks: Beyond hulten's theorem. *Econometrica*, 87(4):1155–1203, 2019.

David Rezza Baqaee and Emmanuel Farhi. Productivity and misallocation in general equilibrium. *Quarterly Journal of Economics*, 135(1):105–163, 2020.

Andrew B. Bernard. Networks as capital: Two-sided accumulation and firm size. Working paper, Tuck School of Business at Dartmouth, 2026.

Andrew B. Bernard, Emmanuel Dhyne, Glenn Magerman, Kalina Manova, and Andreas Moxnes. The origins of firm heterogeneity: A production network approach. *Journal of Political Economy*, 130(7):1765–1804, 2022.

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

Steve Bond, Arshia Hashemi, Greg Kaplan, and Piotr Zoch. Some unpleasant markup arithmetic: Production function elasticities and their estimation from production data. *Journal of Monetary Economics*, 121:1–14, 2021.

Vasco M. Carvalho. From micro to macro via production networks. *Journal of Economic Perspectives*, 28(4):23–48, 2014.

Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 2018.

Allan Collard-Wexler and Jan De Loecker. Production function estimation and capital measurement error. NBER Working Paper 22437, National Bureau of Economic Research, 2016.

Nicolas Crouzet and Janice C. Eberly. Understanding weak capital investment: The role of market concentration and intangibles. NBER Working Paper 25869, National Bureau of Economic Research, 2019.

Jan De Loecker. Product differentiation, multiproduct firms, and estimating the impact of trade liberalization on productivity. *Econometrica*, 79(5):1407–1451, 2011a.

Jan De Loecker. Recovering markups from production data. *International Journal of Industrial Organization*, 29(3):350–355, 2011b.

Jan De Loecker. Detecting learning by exporting. *American Economic Journal: Microeconomics*, 5(3):1–21, 2013.

Jan De Loecker. Comment on (un)pleasant ... by bond et al (2020). *Journal of Monetary Economics*, 121:15–18, 2021.

Jan De Loecker and Paul T. Scott. Markup estimation using production and demand data: An application to the US brewing industry. *Review of Economic Studies*, 2026. Forthcoming; earlier version circulated as NBER Working Paper 22957.

Jan De Loecker and Frederic Warzynski. Markups and firm-level export status. *American Economic Review*, 102(6):2437–2471, 2012.

Jan De Loecker, Pinelopi K. Goldberg, Amit K. Khandelwal, and Nina Pavcnik. Prices, markups, and trade reform. *Econometrica*, 84(2):445–510, 2016.

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

Maarten De Ridder. Market power and innovation in the intangible economy. *American Economic Review*, 114(1):199–251, 2024.

Banu Demir, Ana Cecília Fieler, Daniel Yi Xu, and Kelly Kaili Yang. O-ring production networks. *Journal of Political Economy*, 132(1), 2024. Turkish firm-to-firm (VAT) data; page numbers to verify.

Mert Demirer. Production function estimation with factor-augmenting technology: An application to markups. *Econometrica*, 2026. Forthcoming.

Mert Demirer and Michael Rubens. Welfare effects of buyer and seller power. NBER Working Paper 33371, National Bureau of Economic Research, 2025.

Emmanuel Dhyne, Ayumu Ken Kikkawa, Magne Mogstad, and Felix Tintelnot. Trade and domestic production networks. *Review of Economic Studies*, 88(2):643–668, 2021.

Emmanuel Dhyne, Ayumu Ken Kikkawa, Toshiaki Komatsu, Magne Mogstad, and Felix Tintelnot. Firm responses and wage effects of foreign demand shocks with fixed labor costs and monopsony. *American Economic Review*, 115(12):4328–68, December 2025. Working paper, May 2025. PDF in Literature folder: DKKMT2025.pdf.

Ulrich Doraszelski and Jordi Jaumandreu. R&d and productivity: Estimating endogenous productivity. *Review of Economic Studies*, 80(4):1338–1383, 2013.

Lukasz A. Drozd and Jaromir B. Nosal. Understanding international prices: Customers as capital. *American Economic Review*, 102(1):364–395, 2012.

Francois Fontaine, Julien Martin, and Isabelle Mejean. Frictions and adjustments in firm-to-firm trade. Discussion Paper 18110, CEPR, 2023.

Lucia Foster, John Haltiwanger, and Chad Syverson. Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? *The American Economic Review*, 98 (1):394–425, 2008.

Lucia Foster, John Haltiwanger, and Chad Syverson. The slow growth of new plants: Learning about demand? *Economica*, 83(329):91–129, 2016. Ungated WP: NBER 17853.

Amit Gandhi, Salvador Navarro, and David A. Rivers. On the identification of gross output production functions. *Journal of Political Economy*, 128(8):2973–3016, 2020.

François Gourio and Leena Rudanko. Customer capital. *Review of Economic Studies*, 81(3): 1102–1136, 2014. Ungated WP: NBER 17191.

Paul L. E. Grieco, Shengyu Li, and Hongsong Zhang. Production function estimation with unobserved input price dispersion. *International Economic Review*, 57(2):665–690, 2016.

László Halpern, Miklós Koren, and Adam Szeidl. Imported inputs and productivity. *American Economic Review*, 105(12):3660–3703, December 2015.

Bianca He, Lauren Mostrom, and Amir Sufi. Investing in customer capital. NBER Working Paper 33171, National Bureau of Economic Research, 2024.

Colin J. Hottman, Stephen J. Redding, and David E. Weinstein. Quantifying the sources of firm heterogeneity. *Quarterly Journal of Economics*, 131(3):1291–1364, 2016.

Charles R. Hulten. Growth accounting with intermediate inputs. *Review of Economic Studies*, 45(3):511–518, 1978.

Federico Huneeus. Production network dynamics and the propagation of shocks. Phd thesis, Princeton University, 2018. Working paper.

Ebehi Iyoha. Estimating productivity in the presence of spillovers: Firm-level evidence from the US production network. Working Paper 24-033, Harvard Business School, 2023.

Hajime Katayama, Shihua Lu, and James R. Tybout. Firm-level productivity studies: Illusions and a solution. *International Journal of Industrial Organization*, 27(3):403–413, 2009.

Tor Jakob Klette and Zvi Griliches. The inconsistency of common scale estimators when output prices are unobserved and endogenous. *Journal of Applied Econometrics*, 11(4): 343–361, 1996.

Patrick Kline, Raffaele Saggio, and Mikkel Sølvsten. Leave-out estimation of variance components. *Econometrica*, 88(5):1859–1898, 2020.

Maurice Kugler and Eric Verhoogen. Prices, plant size, and product quality. *Review of Economic Studies*, 79(1):307–339, 2012.

James Levinsohn and Amil Petrin. Estimating production functions using inputs to control for unobservables. *Review of Economic Studies*, 70(2):317–341, 2003.

Julien Martin, Isabelle Méjean, and Mathieu Parenti. Relationship stickiness, international trade, and economic uncertainty. *Review of Economics and Statistics*, 108(1):179–193, 2026.

Agnes Norris Keiller, Aureo de Paula, and John Van Reenen. Production function estimation using subjective expectations data. NBER Working Paper 32725, National Bureau of Economic Research, 2024.

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

G. Steven Olley and Ariel Pakes. The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, 64(6):1263–1297, 1996.

Paul Piveteau. An empirical dynamic model of trade with consumer accumulation. *American Economic Journal: Microeconomics*, 13(4):23–63, 2021.

Devesh Raval. Testing the production approach to markup estimation. *Review of Economic Studies*, 90(5):2592–2611, 2023.

Juan Camilo Serpa and Harish Krishnan. The impact of supply chains on firm-level productivity. *Management Science*, 64(2):511–532, 2018.

Paulo Somaini and Liran Einav. A model of market power in customer markets. *Journal of Industrial Economics*, 61(4):938–986, 2013.

Paulo Somaini and Frank A. Wolak. An algorithm to estimate the two-way fixed effects model. *Journal of Econometric Methods*, 5(1):143–152, 2016.

Rentaro Utamaru. Nonparametric identification and estimation of production functions invariant to productivity dynamics. Working paper, Waseda University, 2026.

A Revenue confounding in the dynamic model

Deflated revenue is a measurement problem that the deflation-first step of Section 4.5 resolves wherever the link panel is available; this appendix records what revenue-based residuals contain when it is not. At the true coefficients the recovered residual is the static expression (24), $\hat{\omega}_i^{\text{rev}} = (1-1/D)(\omega_i^c - \alpha_m \widetilde{P^M}_i) + \frac{1-s}{D}\widetilde{\ln n_i}$. At constant returns ($D = 1$) it is degenerate: the constant markup cancels technology (Proposition 2) and the customer stock, which drives the input demands, is absorbed by conditioning on the full input vector, so what revenue retains is markups and the accumulated customer base. With decreasing returns it carries a fraction $1 - 1/D$ of technology, genuine information, plus an additive customer-stock term that conditioning on the observed n_t removes. The dynamics make the bias bite: n_t is a slow state, so the customer-base content of partial measures such as TFPR and value added per worker inherits its persistence. Where the deflation step is unavailable, the correction is to condition on n_t and the demand index.

B Path 2: the moment system in detail

Parameters and moments. $\boldsymbol{\theta} = (\sigma, \alpha_k, \alpha_l, \alpha_m, \rho_\varphi, c_n, c_v, d_n, d_v)$, nine parameters. The moment vector (39) stacks four blocks. The split block (five moments; Assumption 2(ii), (iv)) is the set of normal equations of (34), with the revenue coefficient written $\alpha_l/\mu(\sigma)$ so the block ties α_l to σ. The share block (one moment; Assumption 2(ii)) is $\mathbb{E}[s^M - \alpha_m(\sigma - 1)/\sigma] = 0$, the intermediate first-order condition of Section 4.3. The timing block (four moments; Assumption 2(iii)) sets the forward-differenced innovation $\Delta a_{it}(\boldsymbol{\theta})$, formed on the exit-trimmed sample, orthogonal to $(k_t,\ i_{t-1},\ \ln n^\star_{t-1},\ \ln v^\star_{t-1})$. The pass-through block (one moment; Assumption 2(ii), (iii)) sets $\Delta u_{it}(\boldsymbol{\theta})$ orthogonal to the supplier-capital instrument $\bar{k}^S_{i,t-1}$, and it is the block that identifies σ (Proposition D.2). The deflated quantities, $\ln S$, and $\hat{L}^p(\boldsymbol{\theta})$ are the Path-1 constructions evaluated at candidate parameters. Eleven moments, nine parameters. The two surplus conditions are the stock orthogonalities, which hold under the DGP only if the augmented inversion has removed the relationship stocks from the recovered process, so the J statistic is an economically named specification test rather than generic over-identification.

The Jacobian. Order $\boldsymbol{\theta} = (\sigma\,;\ \alpha_l, c_n, c_v, d_n, d_v\,;\ \alpha_m\,;\ \alpha_k, \rho_\varphi)$. Given σ, the split block involves only the second group, the share block only α_m, and the timing block only the fourth group once the others are substituted in, so the population Jacobian is block lower-triangular apart from the σ column. Its diagonal blocks are $\mathbb{E}[WW']$ (split), with $W =$

$(R/w, n, v, F^n, F^v)$ the split regressors, $(\sigma-1)/\sigma$ (share), and the timing-moment Jacobian, each nonsingular under the rank conditions stated in Appendix E. The σ column is anchored by the pass-through moment, whose derivative $M'(\sigma_0)$ of Proposition D.1 is nonzero. Expanding along the σ column, the determinant equals the product of the block determinants times $M'(\sigma_0)$. The joint system is locally identified exactly when the pass-through crossing is regular. The two paths are the same information. Path 1 solves this system by block-recursive substitution.

Implementation and inference. Estimate by two-step GMM: a first step at the block-recursive (Path-1) solution, then efficient weighting with a firm-clustered covariance. The seller effects $\tilde{\psi}$ are generated regressors from the fixed-effects step. We recommend the network-fixed bootstrap of Section 7.3, which holds the network fixed and so preserves connectedness, redoing the whole system in each draw. Stacking the fixed-effects normal equations into (39) is the formal alternative, with leave-out corrections where the network is sparse [Kline et al., 2020]. Unbalanced panels enter through the observed (i,t) cells, and where exit is non-trivial the production block conditions on the survival propensity exactly as in Path 1. Report the J statistic together with its components: a violation of the stock orthogonalities localizes the failure to the dynamic layer, a violation of the share moment points to markups or buyer power (Section 4.3), and a violation of the split moments points to the servicing technology, each naming the assumption to relax.

C Matching technologies and the corrections

The corrections in this paper are stated for the realized network, so they do not restrict how firms meet. This appendix states the matching class under which they remain valid, and the one place where richer matching adds a conditioning variable.

The matching class. Let links form by any process measurable in firm-level states: random matching, assortative matching on productivity, size, prices, or the relationship stocks, search frictions with firm-level intensities, or relationship stickiness. Three restrictions remain. *(i) Timing*: formation decisions use information up to $t-1$, so the period-t stocks are predetermined and matching does not condition on the current innovation ϵ_t^{ω}. *(ii) Exogenous mobility*: formation does not select on the pairwise match component ϑ_{ij}. *(iii)* The demand and technology structure of Sections 2–3 is unchanged. What is relaxed is who meets whom.

Static corrections. Every static proposition is stated conditional on the realized network. Over the realized supplier set $P_i^M = v_i^{1/(1-\sigma)}\bar{p}_i$ with $\bar{p}_i$ the power mean of the realized supplier prices, so $\ln P_i^M$ decomposes exactly into a variety term and a share-weighted composition term, whatever process selected the suppliers; selective matching correlates composition with the count, but both are constructed from data and subtracted separately, so the correlation is harmless, and propagation conditions on the observed A.

Dynamic corrections. The input price index a firm faces, $\widetilde{P^M}_i$, varies across firms and enters its value function. A firm matched with cheap suppliers earns a higher margin on every customer. It is therefore a payoff-relevant firm-level state, and the investment policy and exit threshold become $I = \mathcal{I}(K, n, v, \varphi, \widetilde{P^M})$ and $\underline{\omega}(K, n, v, \widetilde{P^M})$. The corrections extend by adding $\widetilde{P^M}_i$ to the conditioning sets of the control function and the survival propensity. It is constructible from the deflators of Section 4.5, so the extension costs no new data. Conditional on the enlarged set, investment remains strictly monotone in ω_t and the inversion logic of Proposition 5 goes through unchanged. The functional-form requirements, log-separable cost, a homothetic input aggregator, and constant markups for exactness, are separate maintained assumptions, untouched by the matching class.

D The network-internal σ crossing

Two results discipline the crossing. The first says that a wrong σ enters the entire pipeline through one scalar.

Lemma D.1 (Misdeflation is a rescaling)**.** *Within the log-linear price system* (21)*, the seller effect satisfies* $\tilde{\boldsymbol{\psi}} = (1-\sigma_0)\tilde{\boldsymbol{p}}$ *and the true input-price wedge satisfies* $\widetilde{\boldsymbol{P^M}} = -\frac{1}{\sigma_0 - 1}\big[\widetilde{\ln \boldsymbol{v}} + A^\top \tilde{\boldsymbol{\psi}}\big]$*, where* σ_0 *is the true elasticity. Hence for any candidate* $\sigma > 1$ *the constructed deflators are exact rescalings,* $\tilde{p}(\sigma) = \kappa(\sigma)\,\tilde{p}$ *and* $\widetilde{P^M}(\sigma) = \kappa(\sigma)\,\widetilde{P^M}$ *with* $\kappa(\sigma) = (\sigma_0 - 1)/(\sigma - 1)$*, and the deflated data satisfy*

$$y^d(\sigma) = q + \big(1 - \kappa(\sigma)\big)\,\tilde{p}, \qquad m^d(\sigma) = m + \big(1 - \kappa(\sigma)\big)\,\widetilde{P^M},$$

with $1 - \kappa(\sigma) = (\sigma - \sigma_0)/(\sigma - 1)$ *strictly increasing in* σ *and zero at* σ_0*. Misdeflation therefore perturbs the estimating equations only through the two monotone scalars* $\kappa(\sigma)$ *and* $\mu(\sigma)$*.* (Proof: Appendix E.)

Proposition D.1 (The pass-through σ crossing: existence, uniqueness, consistency)**.** *Let* $M(\sigma)$ *be the population version of* (37) *under Assumption 2.* (i, the truth is a root) *At* $\sigma = \sigma_0$

the deflators and the imposed coefficients $(\alpha_m(\sigma_0), \gamma_y(\sigma_0), \gamma_k)$ *are exact, so* $u_i(\sigma_0) = (\sigma_0 - 1)\omega_i$ *up to sampling noise. The forward pairwise difference nets out the entry type and the exit-propensity trim makes each retained row's inclusion predetermined, so* $\Delta u_{it}(\sigma_0)$ *is orthogonal to the predetermined supplier-side instrument* $\bar{k}^S_{i,t-1}$ *and* $M(\sigma_0) = 0$*. Without the selection treatment* $M(\sigma_0)$ *is displaced by the exit-selection channel of Section 7, amplified here by the factor* $(\sigma_0 - 1)$ *that multiplies* ω *in* u*.* (ii, existence) M *is continuous on* $(1, \infty)$*. The imposed* $\alpha_m(\sigma) = \mu(\sigma)\bar{s}^M$ *is strictly decreasing, so the loading of the observable index* $\ln S$ *in* $u(\sigma)$ *moves monotonically and* M *changes sign across a bracket that the researcher verifies from the grid.* (iii, uniqueness) *The root at* σ_0 *is locally unique if and only if* $M'(\sigma_0) \neq 0$*. The slope collects the derivative of the imposed* $\alpha_m(\sigma)$ *times the index loading and the* $(\sigma - 1)$ *scaling of the productivity term, given in closed form below. Global uniqueness holds when* M *is strictly monotone on the bracket, a property the researcher verifies directly by evaluating* M *on a grid. We recommend reporting the plot.* (iv, consistency) *Under (ii), (iii), and uniform convergence of the sample* M *over a compact bracket,* $\hat{\sigma} \xrightarrow{p} \sigma_0$*, and the remaining parameters inherit consistency by continuity.* (Proof: Appendix E.)

The minimal iteration is bisection on the scalar M, which converges under (ii) and monotonicity. The crossing compares two different firms' Stage-B objects, a seller's price response against its suppliers' predetermined costs, so it survives the degeneracy of Proposition D.2.

Proposition D.2 (Degeneracy of within-firm σ contrasts)**.** *Impose the share identities* $\alpha_m(\sigma) = \mu(\sigma)\bar{s}^M$ *and* $\alpha_l(\sigma) = \mu(\sigma)\hat{b}_R$*. Let* $\omega_D(\sigma)$ *be the production-side recovery from the deflated timing equation and* $\omega_S(\sigma)$ *the value-only recovery of Proposition 4, both rebuilt at a candidate.* (i) *With the exact index of Lemma 1, the contrast* $\omega_D(\sigma) - \omega_S(\sigma)$ *is identically zero for every* σ*, not only at* σ_0*.* (ii) *If access is measured by the customer count instead, the contrast equals* $(1 - s(\sigma))(\mathcal{N}_i - \ln n_i)$*, with* $\mathcal{N}_i$ *the exact log demand access defined in the proof, a fixed observable-versus-latent wedge whose root does not involve* σ_0*.* (iii) *Under the share identity* $\ln S_i = y_i - \theta_i$ *up to a constant, so any instrument built from firm* i*'s own size is orthogonal to the* α_m *direction, and the pass-through moment* (37) *must use a supplier-side instrument.* (Proof: Appendix E.)

The reason is a scale the network does not separate from σ. Stage B measures $(1-\sigma)\ln p$, never $\ln p$, so every linear recovery built from $\{y, x, \tilde{\psi}, \theta\}$ inherits that product. Once the share identity is imposed the three model identities make any two such recoveries algebraically dependent, and their contrasts recombine into observables whose candidate-dependence cancels. This is why the quantity-side timing moments cannot identify σ, and why cross-recovery correlations peak off the truth. The pass-through restriction (36) escapes

the trap because it links two different firms' Stage-B objects, a seller's own price response and its suppliers' costs, rather than differencing one firm's own identities.

The sensitivity behind Proposition D.1(iii). The pass-through residual $u_i(\sigma) = \tilde{\psi}_i - \alpha_m(\sigma)\ln S_i - \gamma_y(\sigma)y_i - \gamma_k k_i$ depends on σ through the two imposed coefficients $\alpha_m(\sigma)$ and $\gamma_y(\sigma)$. Differentiating the moment (37) at the root and using $\mathbb{E}[\Delta u_{it}(\sigma_0)\,\bar{k}^S_{i,t-1}] = 0$,

$$M'(\sigma_0) = -\alpha'_m(\sigma_0)\,\mathbb{E}\big[\Delta(\ln S_i)\,\bar{k}^S_{i,t-1}\big] - \gamma'_y(\sigma_0)\,\mathbb{E}\big[\Delta y_i\,\bar{k}^S_{i,t-1}\big], \qquad \alpha'_m(\sigma_0) = -\frac{\bar{s}^M}{(\sigma_0-1)^2},$$

with $\gamma'_y(\sigma_0)$ from $\gamma_y(\sigma) = (1-\sigma)(1-s(\sigma))$. The supplier-capital instrument $\bar{k}^S_{i,t-1}$ moves the input index $\ln S_i$ through the suppliers' costs and is orthogonal to firm i's own innovations by construction, so the first term is nonzero and carries the identification. The own-size loading $\mathbb{E}[\Delta y_i\,\bar{k}^S_{i,t-1}]$ is second order here, because $\bar{k}^S$ is supplier-side while y_i is own-size, the orthogonality of Proposition D.2(iii). The slope is therefore dominated by $-\alpha'_m(\sigma_0)\,\mathbb{E}[\Delta(\ln S_i)\,\bar{k}^S_{i,t-1}]$, nonzero with sign set by the covariance of the differenced index with the supplier-cost instrument. It is plug-in verifiable at the estimates, and the grid plot of Step 7 is the global check.

E Proofs

Throughout the proofs $\alpha_k = 0$, as in the text. With predetermined capital, replace ω by the cost fundamental $\omega^c = \omega + \alpha_k\tilde{k}$ of (17) (Section 4, capital paragraph), every added term being observable under $\mathcal{I}_2$.

Reduced-form profit and costates, eqs. (19), (6), (12). This derives the pricing fixed point (19), the reduced-form profit (6), and the costates (12) under general variable returns $s \le 1$, with $D = \sigma - (\sigma-1)s$.

Variable cost minimization over (L^p_i, M_i) for the degree-s technology gives marginal cost $mc_i = \Lambda_i Q_i^{(1-s)/s}$ with scale-free part $\Lambda_i = \varphi_{K,i}^{-1/s}(w/\alpha_l)^{\alpha_l/s}(P^M_i/\alpha_m)^{\alpha_m/s}$ and $\varphi_{K,i} = \varphi_i K_i^{\alpha_k}$, and $VC_i = s\,mc_i Q_i$ (Euler, the cost function being homogeneous of degree $1/s$ in Q). Constant-markup pricing $p_i = \mu\,mc_i$ and CES demand $Q_i = n_i\bar{\psi}p_i^{-\sigma}$ form a fixed point. Since $p_i = \mu\Lambda_i(n_i\bar{\psi}p_i^{-\sigma})^{(1-s)/s}$ and $1 + \sigma(1-s)/s = D/s$,

$$p_i = (\mu\Lambda_i)^{s/D}(n_i\bar{\psi})^{(1-s)/D}, \quad Q_i = (n_i\bar{\psi})^{s/D}(\mu\Lambda_i)^{-\sigma s/D}, \quad R_i = (n_i\bar{\psi})^{1/D}(\mu\Lambda_i)^{-s(\sigma-1)/D}.$$

Taking logs of p_i in deviations gives (19). Operating profit is $\Pi_i = R_i - VC_i = (1 - s/\mu)R_i$. Using $\Lambda_i^{-s(\sigma-1)/D} \propto \varphi_{K,i}^{(\sigma-1)/D}(P^M_i)^{-\alpha_m(\sigma-1)/D}$ and the exact decomposition (3), $P^M_i =$

$v_i^{1/(1-\sigma)}\bar{p}_i^M$, so that $(P_i^M)^{-\alpha_m(\sigma-1)/D} = v_i^{\alpha_m/D}(\bar{p}_i^M)^{-\alpha_m(\sigma-1)/D}$, and collecting the demand terms as $(\bar{\psi}n_i)^{1/D}$ and the cost terms under the common exponent $(\sigma-1)/D$, gives (6). Differentiating (6) at fixed composition delivers the costates (12) and the cross-partials. □

□

Proof of Proposition 1. With physically measured output, $y_i = q_i$, (18) gives $\hat{\boldsymbol{\omega}} = \boldsymbol{\omega} - \alpha_m\widetilde{\boldsymbol{P^M}}$. From (20), $-\alpha_m\widetilde{\boldsymbol{P^M}} = \zeta\widetilde{\ln \boldsymbol{v}} - \alpha_m A^\top\tilde{\boldsymbol{p}}$. By the price system (21), $\tilde{\boldsymbol{p}} = -\frac{1}{D}\Psi_s(\boldsymbol{\omega} + \zeta\widetilde{\ln \boldsymbol{v}} - (1-s)\widetilde{\ln \boldsymbol{n}})$, so $-\alpha_m A^\top\tilde{\boldsymbol{p}} = \frac{\alpha_m}{D}A^\top\Psi_s(\cdots) = (\Psi_s - \mathbb{I})(\cdots)$ by the Neumann identity $\frac{\alpha_m}{D}A^\top\Psi_s = \Psi_s - \mathbb{I}$. Therefore

$$\hat{\boldsymbol{\omega}} = \boldsymbol{\omega} + \zeta\widetilde{\ln \boldsymbol{v}} + (\Psi_s - \mathbb{I})\big(\boldsymbol{\omega} + \zeta\widetilde{\ln \boldsymbol{v}} - (1-s)\widetilde{\ln \boldsymbol{n}}\big) = \Psi_s(\boldsymbol{\omega} + \zeta\widetilde{\ln \boldsymbol{v}}) - (1-s)(\Psi_s - \mathbb{I})\widetilde{\ln \boldsymbol{n}},$$

which is (22). The own customer stock appears only off-diagonally through $\Psi_s - \mathbb{I}$. At $D = 1$ the expression is the constant-returns $\Psi(\boldsymbol{\omega} + \zeta\widetilde{\ln \boldsymbol{v}})$. □

Proof of Corollary 1. The object $\omega_i^e = \omega_i + \zeta\widetilde{\ln v_i}$ has variance $\mathrm{Var}(\omega) + \zeta^2\mathrm{Var}(\ln v) + 2\zeta\mathrm{Cov}(\omega, \ln v) \geq \mathrm{Var}(\omega)$ whenever $\mathrm{Cov}(\omega, \ln v) \geq 0$, which holds because more-productive firms hold more suppliers. With propagation, $\hat{\boldsymbol{\omega}} = \Psi_s\,\boldsymbol{\omega}^e$ (suppressing the off-diagonal customer term, which only adds dispersion). For i.i.d. ω^e with variance s_e^2, the average measured variance is $\frac{s_e^2}{N}\|\Psi_s\|_F^2 \geq s_e^2$ since $\Psi_s = \mathbb{I} + \frac{\alpha_m}{D}A^\top\Psi_s$ has diagonal entries at least one and non-negative entries, with equality iff $\alpha_m = 0$. Chaining gives $\mathrm{Var}(\hat{\omega}) \geq \mathrm{Var}(\omega^e) \geq \mathrm{Var}(\omega)$. For $i \neq i'$, $\mathrm{Cov}(\hat{\omega}_i, \hat{\omega}_{i'}) = s_e^2(\Psi_s\Psi_s^\top)_{ii'} \geq 0$, strictly positive when i, i' share an upstream path, and $\hat{\omega}_i$ is increasing in v_i through the diagonal $\zeta\widetilde{\ln v_i}$. □

Proof of Corollary 2. By Proposition 1, $\hat{\boldsymbol{\omega}} = \Psi_s(\boldsymbol{\omega} + \zeta\widetilde{\ln \boldsymbol{v}}) - (1-s)(\Psi_s - \mathbb{I})\widetilde{\ln \boldsymbol{n}}$. Premultiplying by $\Psi_s^{-1} = \mathbb{I} - \frac{\alpha_m}{D}A^\top$ gives $(\mathbb{I} - \frac{\alpha_m}{D}A^\top)\hat{\boldsymbol{\omega}} = \boldsymbol{\omega} + \zeta\widetilde{\ln \boldsymbol{v}} - (1-s)(\mathbb{I} - \frac{\alpha_m}{D}A^\top)(\Psi_s - \mathbb{I})\widetilde{\ln \boldsymbol{n}}$, hence (23). □

Proof of Proposition 2. With deflated revenue, $y_i = q_i + \tilde{p}_i$, so $\hat{\omega}_i = \omega_i + \tilde{p}_i - \alpha_m\widetilde{P^M}_i$. By (19), $\tilde{p}_i = \frac{1}{D}(\alpha_m\widetilde{P^M}_i - \omega_i) + \frac{1-s}{D}\widetilde{\ln n_i}$ (with $\omega_i^c = \omega_i$ at $\alpha_k = 0$), so

$$\hat{\omega}_i = \omega_i + \tfrac{1}{D}(\alpha_m\widetilde{P^M}_i - \omega_i) + \tfrac{1-s}{D}\widetilde{\ln n_i} - \alpha_m\widetilde{P^M}_i = \big(1 - \tfrac{1}{D}\big)(\omega_i - \alpha_m\widetilde{P^M}_i) + \tfrac{1-s}{D}\widetilde{\ln n_i},$$

which is (24), using $1 - 1/D = (D-1)/D = (\sigma-1)(1-s)/D$. At $D = 1$ both coefficients vanish and $\hat{\omega}_i = 0$. For the attenuation, CES demand gives $R_i \propto Q_i^{(\sigma-1)/\sigma}$ up to a demand shifter, so $\ln R_i = \frac{\sigma-1}{\sigma}(\sum_j \alpha_j x_{ji} + \omega_i) + \text{shifter}$. With an instrument for the shifter the revenue function identifies $\hat{\beta}_j = \alpha_j/\mu$ and $\widehat{RTS} = \eta/\mu$ with $\eta = \alpha_k + s$. □

Proof of Proposition 3. Double deflation at the share s^M gives, to first order in log deviations, $\widetilde{va} = (\tilde{y} - s^M\tilde{x})/(1 - s^M)$. Substitute $\tilde{y} = \tilde{q} + \tilde{p}$ and $\tilde{x} = \tilde{m} + \widetilde{P^M}$ from (15), the

technology $\tilde{q} = \alpha_k\tilde{k}+\alpha_l\tilde{\ell}+\alpha_m\tilde{m}+\omega$, and evaluate the residual at the value-added coefficients $\beta_j = \alpha_j/(1-s^M)$:

$$(1-s^M)\,\omega^{VA} = \alpha_m\tilde{m}+\omega+\tilde{p}-s^M\big(\tilde{m}+\widetilde{P^M}\big) = \big[\omega+\tilde{p}-\alpha_m\widetilde{P^M}\big]+\big(\alpha_m-s^M\big)\big(\tilde{m}+\widetilde{P^M}\big) = \hat{\omega}+\big(\alpha_m-s^M\big)\tilde{x},$$

using (18), which is (26). The share identity $s^M = \alpha_m/\mu$ (Section 4.3) gives $\alpha_m - s^M = \alpha_m(1-1/\mu) > 0$ for $\mu > 1$. At $\mu = 1$ the extra term vanishes and $\omega^{VA} = \hat{\omega}/(1-\alpha_m)$, so every additive component of $\hat{\omega}$, including the variety term $\zeta\widetilde{\ln \boldsymbol{v}}$ and the propagation terms, is scaled by $1/(1-\alpha_m)$. Exactness of the subtraction would require the $\tilde{m}$ and price terms to vanish identically, which needs $\alpha_m = s^M$, fixed proportions, and a common intermediate price, conditions the sourcing technology (1) excludes for $\sigma > 1$. □

Derivation of the errors-in-variables limit (25). Partial capital, labor, and the productivity term out of y_i and x_i, isolating the input-price measurement error. Let m and $x = m + \widetilde{P^M}$ denote the residualized intermediate quantity and its observed counterpart, with residualized $y = \alpha_m m$. Then $\operatorname{plim}\hat{\beta}_m = \operatorname{Cov}(y,x)/\operatorname{Var}(x) = \alpha_m\operatorname{Cov}(m, m+\widetilde{P^M})/\operatorname{Var}(m+\widetilde{P^M})$, which is (25). Cost-minimizing intermediate demand decreases in its own price, so $\operatorname{Cov}(m, \widetilde{P^M}) < 0$. The ratio is below one iff $\operatorname{Var}(\widetilde{P^M}) + \operatorname{Cov}(m, \widetilde{P^M}) > 0$. By the wedge (20) and the price system (21), $\widetilde{P^M}_i = -\frac{1}{\sigma-1}\widetilde{\ln v}_i - \frac{1}{D}(A^\top\Psi_s(\boldsymbol{\omega} + \zeta\widetilde{\ln \boldsymbol{v}} - (1-s)\widetilde{\ln \boldsymbol{n}}))_i$, so $\operatorname{Var}(\widetilde{P^M})$ loads on the dispersion of the supplier stock and of upstream effective productivity (damped by $1/D$) and, when $s < 1$, on the suppliers' customer-stock dispersion. Returns to scale, summing the input coefficients, inherit the bias. □

Derivation of the share identity and its failures (Section 4.3). The intermediate first-order condition $P_i^M = mc_i\,\partial Q_i/\partial M_i$ gives $P_i^M M_i = mc_i\theta_i^m Q_i$ with $\theta_i^m = \alpha_m$, so $s_i^M = P_i^M M_i/(p_i Q_i) = \theta_i^m/\mu = \alpha_m/\mu$ for any s (returns to scale enter the level of cost, not a single input's share). (i) With a share-dependent demand elasticity, $p_i = \mu_i mc_i$ and applying Proposition 2 firm-by-firm makes the attenuation firm-specific; $\hat{\mu}_i = \theta_i^m/s_i^M$ inherits the network share variation. (ii) Buyer power scales the share to $s_i^M = \alpha_m/(\mu\varpi_i^B)$, so $\hat{\alpha}_m = \mu s_i^M = \alpha_m/\varpi_i^B$. (iii) With endogenous matching, $m_i = m(k_i, \omega_i, \mathcal{P}_i)$ where the realized partner set $\mathcal{P}_i$ determines $\widetilde{P^M}_i$ independently of ω_i, so $(k_i, m_i) \mapsto \omega_i$ at common prices is not well defined. □

Proof of Proposition 4. In deviations the seller fixed effect in (27) is $\tilde{\psi}_i = (1-\sigma)\tilde{p}_i$. By the price system (21), $\tilde{p}_i = -\frac{1}{D}(\Psi_s(\boldsymbol{\omega} + \zeta\widetilde{\ln \boldsymbol{v}} - (1-s)\widetilde{\ln \boldsymbol{n}}))_i$, so $\tilde{\psi}_i = \frac{\sigma-1}{D}(\Psi_s(\boldsymbol{\omega} + \zeta\widetilde{\ln \boldsymbol{v}} - (1-s)\widetilde{\ln \boldsymbol{n}}))_i$, which is (29). Multiplying by $D/(\sigma-1)$ and premultiplying by $\Psi_s^{-1} = \mathbb{I} - \frac{\alpha_m}{D}A^\top$ gives $\boldsymbol{\omega} + \zeta\widetilde{\ln \boldsymbol{v}} - (1-s)\widetilde{\ln \boldsymbol{n}} = \frac{D}{\sigma-1}(\mathbb{I} - \frac{\alpha_m}{D}A^\top)\tilde{\boldsymbol{\psi}}$, and rearranging gives (30). Only $\ln r_{ij}$ enters the regression and only $A, \alpha_m, \sigma, s, v, n$ enter the inversion, so no prices are used. At $D = 1$ this is the constant-returns recovery. □

Derivation of the labor mismeasurement (Section 5.4). Write $\ln L_i = \ln L_i^p + \zeta_i^L$ with $\zeta_i^L = \ln(1 + (L_i^n + c_n\, n_i + L_i^v + c_v\, v_i)/L_i^p)$ the log gross-up by the non-production labor share. The servicing components are proportional to the two stocks by technology, and the acquisition components are governed by the two lifetime values and the wage, so ζ_i^L is increasing in both n_i and v_i. Since both stocks scale demand and hence output, ζ_i^L is correlated with y_i and with the other inputs. Replacing $\ln L_i^p$ by $\ln L_i$ in (16) introduces a non-classical regressor error with two components, one per stock, so $\operatorname{plim} \hat{\beta}_l \neq \alpha_l$, and because ζ_i^L co-moves with the remaining inputs the bias spreads. □

Proof of Proposition 5. (i) By (14), $i_t = \mathcal{I}(k_t, n_t, v_t, \omega_t)$ with $\partial\mathcal{I}/\partial n_t > 0$ and $\partial\mathcal{I}/\partial v_t > 0$ (the shadow value of capital rises with scale, $\Pi_{Kn} > 0$, and with effective productivity, $\Pi_{Kv} > 0$, both signs holding for any s by (12)). Fix (k_t, i_t). The equation $\mathcal{I}(k_t, n_t, v_t, \omega_t) = i_t$ defines a two-dimensional locus along which ω_t varies with (n_t, v_t), so ω_t is not a function of (k_t, i_t) alone and $h(k_t, i_t) = \mathbb{E}[\omega_t \mid k_t, i_t]$ is a mixture of ω_t, n_t, v_t. The same holds for an intermediate proxy. *(ii)* Write $\hat{\omega}_t = \omega_t + b_n\tilde{n}_t + b_v\tilde{v}_t$ by (i), with $\tilde{n}_t, \tilde{v}_t$ the components of the stocks not spanned by ω_t, mutually uncorrelated, AR(1) with persistence $1-\delta$, and ω_t AR(1) with persistence ρ_φ. Then $\widehat{\rho} = \operatorname{Cov}(\hat{\omega}_t, \hat{\omega}_{t-1})/\operatorname{Var}(\hat{\omega}_t)$ equals (32), a convex combination of ρ_φ and $1-\delta$ with weights proportional to the variance contributions, moving toward $1-\delta$ as the contamination weights grow. The loadings b_n, b_v depend on s but the combination form does not. *(iii)* For an active firm the value function is strictly increasing in ω and, through $\lambda^n = V_n > 0$, $\lambda^v = V_v > 0$, and $q = V_K > 0$, in each stock, for any $s \leq 1$. The participation condition $V(K, n, v, \omega) \geq \underline{V}$, with $\underline{V}$ the fixed value of exiting, which binds at a finite threshold because the firm pays the fixed carrying cost $w\bar{\ell}_i$ of Section 5.2, defines a threshold $\underline{\omega}(K, n, v)$ decreasing in each argument. Survival truncates the productivity distribution below at $\underline{\omega}(K, n, v)$, and the truncated mean is increasing in the threshold, hence decreasing in the stocks. Among survivors, ω and (n, v) are negatively correlated conditional on fundamentals, most strongly for young thin-network firms. The Olley and Pakes [1996] propensity $P(k_t, i_t)$ omits the stocks, and $P(k_t, i_t, n_t, v_t)$ spans the threshold. *(iv)* Conditioning on the observed (n_t, v_t) leaves ω_t as the only unobservable in (14). Since $\mathcal{I}_\omega \neq 0$, $\mathcal{I}(k_t, n_t, v_t, \cdot)$ is strictly monotone and invertible, so $\omega_t = h(k_t, i_t, n_t, v_t)$ restores the first-stage purge, and conditioning the survival propensity on the stocks restores the exit-selection correction by (iii). Hence the elasticities and the productivity process are consistently estimated, and the conditioning set does not depend on s. With heterogeneity the value function loads on the value-weighted stocks (Section 2). □

Derivation of the composite (33). Log-linearize the measurement equations. The static layer (Proposition 1) contributes $\Psi_s(\boldsymbol{\omega} + \zeta\widetilde{\ln \boldsymbol{v}}) - (1-s)(\Psi_s - \mathbb{I})\widetilde{\ln \boldsymbol{n}}$, the first four terms of (33). The

dynamic layer (Proposition 5(i)–(ii)) contributes the additive contamination $b_n\tilde{n}$, the proxy's loading on the ω-orthogonal customer capital. To first order the layers enter additively, giving (33). The dispersion statement is Corollary 1. The persistence statement is (32). Each correction acts on its own term, and none removes another's, so consistency requires all. □

Proof of Lemma D.1. From the proof of Proposition 4, $\tilde{\psi}_i = (1-\sigma_0)\tilde{p}_i$ in deviations, so $\tilde{p}(\sigma) = -\tilde{\psi}/(\sigma-1) = \kappa(\sigma)\tilde{p}$ with $\kappa(\sigma) = (\sigma_0-1)/(\sigma-1)$. The true wedge (20) is $\widetilde{P^M} = -\frac{1}{\sigma_0-1}\widetilde{\ln v} + A^\top\tilde{\boldsymbol{p}} = -\frac{1}{\sigma_0-1}\big[\widetilde{\ln v} + A^\top\tilde{\boldsymbol{\psi}}\big]$, and the constructed index replaces $\sigma_0 - 1$ by $\sigma - 1$, so $\widetilde{P^M}(\sigma) = \kappa(\sigma)\widetilde{P^M}$. Then $y^d(\sigma) = y - \tilde{p}(\sigma) = q + (1-\kappa)\tilde{p}$ and $m^d(\sigma) = x - \widetilde{P^M}(\sigma) = m + (1-\kappa)\widetilde{P^M}$ by (15). Finally $1-\kappa(\sigma) = (\sigma-\sigma_0)/(\sigma-1)$, with derivative $(\sigma_0-1)/(\sigma-1)^2 > 0$, zero at σ_0, limit $-\infty$ as $\sigma \downarrow 1$ and limit 1 as $\sigma \to \infty$. □

Proof of Lemma 1. In (27) the buyer effect is $\theta_j = (\sigma-1)\ln P_j^M + \ln X_j$. Market clearing on buyer j's purchases, $\sum_i r_{ij} = E_j^M$, together with the CES definition $(P_j^M)^{1-\sigma} = S_j$ and the demand system $r_{ij} = \vartheta_{ij}(p_i/P_j^M)^{1-\sigma}X_j$, forces $X_j = E_j^M$ and $(P_j^M)^{\sigma-1}S_j = 1$. Hence $\ln X_j = x_j$ and $(\sigma-1)\ln P_j^M = -\ln S_j$, so $\theta_j = x_j - \ln S_j$ and $\ln S_j = x_j - \theta_j$. The input price index in deviations is $\widetilde{P^M}_j = \ln P_j^M - \overline{\ln P^M} = -\ln S_j/(\sigma-1)$ up to the year constant, which is (28). The value-only construction (30) replaces $\ln S_j$ by $\ln v_j + (A^\top\tilde{\boldsymbol{p}})_j$, the log count plus the share-weighted mean of the recovered supplier prices, which equals $\ln S_j$ only when the log of a sum equals the sum of a mean, that is, at degenerate match quality, and otherwise omits the Jensen dispersion term. □

Proof of Proposition D.1. (i) At σ_0 the deflators are exact (Lemma D.1, $\kappa = 1$) and the imposed coefficients equal their true values, $\alpha_m(\sigma_0) = \alpha_m$, $\gamma_y(\sigma_0) = (1-\sigma_0)(1-s_0)$, and $\gamma_k = -(1-\sigma_0)\alpha_k$ from the timing step. Substituting the pricing identity (36) into $u_i(\sigma_0) = \tilde{\psi}_i - \alpha_m \ln S_i - \gamma_y y_i - \gamma_k k_i$ leaves $u_i(\sigma_0) = (\sigma_0-1)\omega_i$ up to the Stage-B sampling errors $\varepsilon_i^\psi - \alpha_m\varepsilon_i^\theta$. The productivity term is AR(1) around the entry type [Assumption 2(iii)], so the quasi-difference $u - \rho_\varphi u_{-1}$ carries the type as a firm constant, and the forward pairwise difference Δ removes it, leaving $(\sigma_0-1)(\epsilon_{it}^\omega - \epsilon_{i,t+1}^\omega)$ plus differenced noise. The instrument $\bar{k}_{i,t-1}^S$ is the suppliers' lagged capital, predetermined and drawn from other firms, hence orthogonal to ϵ_{it}^ω, $\epsilon_{i,t+1}^\omega$, and firm i's own sampling errors. The exit-propensity trim makes each retained row's inclusion predetermined, removing the selection through the endogenous threshold that would otherwise correlate the retained innovations with the instrument. Hence $M(\sigma_0) = \mathbb{E}[\Delta u_{it}(\sigma_0)\,\bar{k}_{i,t-1}^S] = 0$. Without the trim and the forward difference the selection channel displaces $M(\sigma_0)$, scaled by (σ_0-1) through the productivity loading in u. *(ii)* $u_i(\sigma)$ depends on σ through the continuous imposed coefficients $\alpha_m(\sigma) = \mu(\sigma)\bar{s}^M$ and $\gamma_y(\sigma) = (1-\sigma)(1-s(\sigma))$, so $M(\sigma)$ is continuous. Since $\alpha_m(\sigma)$ is strictly decreasing, the loading

of $\ln S$ in u moves monotonically in σ, so M changes sign across a bracket the researcher verifies from the grid. *(iii)* Differentiating gives $M'(\sigma_0)$, the implicit-function expression below, and the root is locally unique iff $M'(\sigma_0) \neq 0$. *(iv)* With a compact bracket on which the population M has the unique zero σ_0 and the sample counterpart converges uniformly, the standard consistency argument for scalar Z-estimators applies, and the remaining parameters are continuous transformations of $\hat{\sigma}$ and sample moments. □

Proof of Proposition D.2. Work in within-year log deviations with the year constants dropped. Write $D(\sigma) = \sigma - (\sigma - 1)s(\sigma)$ for the candidate dampening scalar, $D_0 = D(\sigma_0)$, and $\kappa = (\sigma_0 - 1)/(\sigma - 1)$ of Lemma D.1. *(i)* With the exact index of Lemma 1 and share-identity labor, the production-side recovery is $\omega_D(\sigma) = (1-\alpha_l(\sigma))y + \tilde{\psi}/(\sigma-1) - \alpha_k k - \alpha_m(\sigma)[x - \widehat{P}(\sigma)]$ and the value-only recovery rebuilt at a candidate is $\omega_S(\sigma) = D(\sigma)\,\tilde{\psi}/(\sigma-1) + \alpha_m(\sigma)\widehat{P}(\sigma) + (1 - s(\sigma))(y - \tilde{\psi}) - \alpha_k k$, with $\widehat{P}(\sigma) = (x - \theta)/(1 - \sigma)$. Substituting $m = x - P$ and the share-identity constancy of $y - x$, the difference reduces to $\omega_D(\sigma) - \omega_S(\sigma) = \Phi(\sigma)\,p$ with $\Phi(\sigma) = \kappa(D(\sigma)-1) - (D_0 - 1) + (s_0 - s(\sigma))(1-\sigma_0)$. Since $\kappa(D(\sigma)-1) = (\sigma_0 - 1)(1 - s(\sigma))$, the three terms cancel and $\Phi(\sigma) \equiv 0$ for all σ. *(ii)* Replacing exact access $y - \tilde{\psi}$ by the customer count $\ln n$ in ω_S leaves the residual wedge $(1 - s(\sigma))(\mathcal{N} - \ln n)$, where $\mathcal{N}$ is the exact log demand access, the log of the customer-weighted market capacity facing the firm, not the network matrix A. Its orthogonality condition $\mathbb{E}[(\mathcal{N} - \ln n)\,\bar{k}^S] = 0$ fixes a relation among observables that does not contain σ_0. *(iii)* Under the share identity intermediate expenditure is proportional to revenue, so $x = y$ up to a constant, and by Lemma 1 $\ln S = x - \theta = y - \theta$ up to a constant. Any instrument that is a function of firm i's own size moves $\ln S$ only through the y component, which is the γ_y direction, not the α_m direction carried by θ. The pass-through moment (37) therefore requires a supplier-side instrument, orthogonal to own size, to load on α_m. □

Proof of Proposition 6. That $\boldsymbol{\theta}_0$ solves the population system is Assumption 2 plus the text derivations. The split moments hold by (9) and Assumption 2(iv), the share moment by the intermediate first-order condition (Section 4.3), the timing moments by the AR(1)-around-the-type law with the forward difference and trim of Stage D, and the pass-through moment by part (i) of Proposition D.1. For identification, fix σ. The split block is a linear projection on $W = (R/w, n, v, F^n, F^v)$, and $\mathbb{E}[WW']$ nonsingular identifies its five coefficients, hence $(\alpha_l/\mu(\sigma), c_n, c_v, d_n, d_v)$. The share moment is linear in α_m with nonzero coefficient $(\sigma-1)/\sigma$. Given these, the timing block determines (α_k, ρ_φ) under the timing rank conditions, and the two stock orthogonalities remain as over-identifying content. Elasticity σ is pinned by the pass-through moment, the one restriction that is not degenerate in σ (Proposition D.2). Ordering the parameters $(\sigma; \alpha_l, c_n, c_v, d_n, d_v; \alpha_m; \alpha_k, \rho_\varphi)$, the population Jacobian is block

lower-triangular apart from the σ column, its diagonal blocks are nonsingular by the rank conditions just stated, and expanding along the σ column shows its determinant equals $M'(\sigma_0)$ times the product of the block determinants (Appendix B). Local identification therefore holds iff $M'(\sigma_0) \neq 0$, the condition of Proposition D.1(iii). □