Markup and Spillover Attribution in Demand–Production Networks: An AI Exposure Application

Zhongji Wei*

June 5, 2026

Abstract

What do firm-level markups and spillovers measure when firms are connected through rivalry, joint use, and production links at once? I develop a multilayer attribution framework that separates final-use substitution, final-use complementarity, and production dependence, and embeds them in a local single-price general equilibrium. Firms post one price to final and intermediate buyers, so final-use demand, downstream intermediate demand, and upstream cost feedback discipline the same markup, and final-use shocks are signed by the demand system before production links propagate them. I implement the framework for U.S. public firms. The 2022 complementarity layer is sparse, covering 0.336 percent of possible firm pairs, and is largely distinct from text-based rivalry and strong input–output exposure. Under pass-through discipline, final-use pricing pressure is 54 percent of the downstream intermediate-demand term, while cost feedback lowers the no-cost-feedback Lerner margin from 0.338 to 0.280. The one-price benchmark is consequential: markup movements that originate on the final-demand side create cross-outlet input-cost externalities that vanish under segmented pricing. Applied to AISC strict AI input-use records, one unit of direct downstream exposure generates 2.38 units of gross firm-level incidence in the 2023 pre-network, with 66.7 percent landing outside directly exposed firms. I then use the same outlet discipline to separate this production-heavy input-use incidence from customer-pull supplier demand and BEA final-use AI demand.

Keywords: demand networks, production networks, markups, complementarity, AI input use, technology diffusion, pass-through.

JEL: D21, D43, D57, D85, L13, L16, O33.

*Universitat Pompeu Fabra, zhongji.wei@upf.edu. I thank Dimin Liu, Manchun Wang, Shengxing Zhang, and the audience at the UPF Macro Reading Group for helpful comments and discussion. All errors are my own.

1 Introduction

Firm-level markups and sales spillovers are observed after several economic relations have already been compressed into one price or output response. What does a firm's markup measure when the firm sells to both households and other firms? A high markup can reflect weak substitution from final consumers, joint-use complementarity with other products, inelastic demand from downstream buyers, or upstream cost feedback through input-output links. The same observed sales response can also come from different structural layers. The paper's object is therefore not only how large a markup or spillover is, but which economic relation generates it.

The paper's leading application is a salient recent shock: firm-level AI exposure. It is a measured-exposure application, not a single AI shock. Direct AISC records identify downstream AI input users and generate broad, production-heavy incidence. The framework then separates this input-use object from two demand-facing AI objects: customer-pull supplier demand from AI-adopting customers and BEA final-use AI demand. The distinction is the point: the same AI adoption episode can enter as input use, customer demand, or final-use demand, and each object belongs to a different operator and markup margin. AI thus motivates the attribution problem, but the framework is general—the same objects attribute any markup or spillover once firms are linked by rivalry, complementarity, and production at the same time.

That framework starts from a measurement fact: product-market proximity, final-use complementarity, and production dependence are different primitives. Treating them as one network can assign the right exposure to the wrong margin, or even to the wrong sign. A production network can propagate activity, but it cannot say whether nearby final-use firms are complements or substitutes. A product-market network can measure proximity, but it cannot separate final-use signs from production dependence. A markup table can report price-cost margins, but not which demand or cost margin disciplines them.

The framework separates three primitive relations: final-use rivalry, final-use complementarity, and production dependence. I denote the two final-demand layers by $\boldsymbol{\Sigma}^S$ and $\boldsymbol{C}$. The production transaction layer is used in two normalizations: one for upstream cost feedback and one for downstream derived demand and output propagation. Firms post one price for a product that can be sold both to final users and to intermediate buyers. In that benchmark, the local Lerner restriction combines the final-use outlet share, the final-use elasticity, the downstream intermediate-demand elasticity, and the upstream cost-feedback elasticity. This is a pricing restriction from the firm's local FOC, not an observed revenue identity. Same markup, different constraining margin.

The second object attributes output spillovers across layers. In a fixed-price, fixed-input-coefficient benchmark, write $\boldsymbol{R}^M = (\boldsymbol{I} - \boldsymbol{M})^{-1}$ for production propagation and $\boldsymbol{H}^F = (\boldsymbol{I} + \boldsymbol{\Sigma}^S - \boldsymbol{C})^{-1}$ for signed final-demand propagation. A final-demand shifter satisfies

$$\mathrm{d}\boldsymbol{y} = \boldsymbol{R}^M \boldsymbol{H}^F \mathrm{d}\boldsymbol{b}. \tag{1}$$

Observed output responses are therefore final-demand movements after production propagation.

This is an ordered propagation restriction: signed final demand selects which quantities move before input-output links propagate them. A supplier can gain from a shock to a customer's complement or lose from a shock to a customer's substitute even when the input-output link is unchanged. Same output response, different structural source.

A central measurement caveat follows. Demand alone reveals the reduced-form final-demand operator $\boldsymbol{G}^F = \boldsymbol{\Sigma}^S - \boldsymbol{C}$, not the two primitive layers separately. Appendix A formalizes this reduced-form invariance. The analysis is therefore conditional on primitive-layer measurement: rivalry, joint-use complementarity, and production dependence must be measured or validated as distinct economic relations before the markup and propagation decompositions are computed. The empirical burden is not deconvolution from a single demand matrix; it is separate measurement of primitives plus equilibrium attribution.

The same local objects close in a benchmark general equilibrium, and that benchmark is where the quantitative work is done. Markup movements are not macro-relevant only because a firm is large or has a high markup; their aggregate price and welfare incidence depends on final-use expenditure weights, household income, profit rebates, and production-network cost pass-through. The benchmark keeps an active household budget with labor income and profit rebates together with input-output price propagation, and I evaluate explicit local comparative statics along a regular branch: the one-price cross-outlet cost externality of a markup movement, the welfare incidence of a located cost reduction, and the signed cross-sectional reallocation of a located demand pull. These are model-implied local-equilibrium counterfactuals, not causal aggregate effects, and the fixed-price composite operator $\boldsymbol{R}^M\boldsymbol{H}^F$ becomes the quantity sub-block diagnostic within that equilibrium rather than a separate device.

The quantitative implementation follows the same sequence. It first measures the primitive-layer system, then decomposes markups by elasticity and cost-feedback margin, and finally uses the same primitives to attribute spillovers. It validates the new complementarity layer as a sparse final-use primitive, anchors the calibration with external demand and pass-through evidence, and compares the composite propagation operator $\boldsymbol{R}^M\boldsymbol{H}^F$ with one-layer alternatives. Section 7 is the demand-facing evidence block: signed final-demand measurement changes markup attribution and spillover signs in product-facing diagnostics. Section 8 is the measured-exposure application. It asks a different firm-level question: once AI input use is observed at a limited set of provider-user edges, which firms are incident to that exposure, through which economic margins, and how does that incidence relate to markup-relevant outlet structure?

The paper makes three contributions, each with a quantitative counterpart. First, it decomposes firm-level markups by constraining margin—final-use demand, downstream intermediate demand, and upstream cost feedback—rather than as a single-network statistic; under pass-through discipline the final-use pricing term is 54 percent of the downstream intermediate-demand term, while cost feedback lowers the no-cost-feedback Lerner from 0.338 to 0.280. Second, it decomposes output spillovers by ordered propagation $\boldsymbol{R}^M\boldsymbol{H}^F$, with final-demand signs

measured before input-output links propagate them; the measured 2022 complementarity layer that supplies those signs is sparse, 33,085 firm pairs or 0.336 percent of all pairs, much thinner than text-based rivalry. Third, it applies the same attribution system to AI exposure. It begins with AISC strict downstream input-use records, which produce broad but production-heavy network incidence—2.38 units of gross incidence per unit of direct exposure, with 66.7 percent landing outside directly exposed firms. It then separates that object from two demand-facing AI vectors: customer-pull supplier demand from AI-adopting customers and BEA final-use AI demand. The distinction is the point: the framework tells whether an AI exposure vector belongs to input use, customer demand, or final-use demand before assigning signs, propagation operators, and markup margins.

The empirical blocks are deliberately separated by estimand, so that the quantitative payoff is read at the right altitude. The signed-demand evidence is carried by the C-measurement, markup-attribution, and demand-facing diagnostic blocks; the AI input-use block is a measured-exposure application whose lesson is an undercount result, not a large signed-demand amplification.

The main text states the primitive layers, the markup attribution object, the propagation object, and the AI input-use application. A full technical appendix contains the supporting derivations, robustness tables, disclosure audits, and proofs; the NBER submission version keeps the main PDF focused on the ten paper-facing tables.

Related literature and contribution. The closest demand-network benchmark is product-market rivalry. Hoberg and Phillips (2016) construct scalable firm-level product-market networks from product text, and Pellegrino (2025) shows how product-market centrality disciplines markups in a final-demand oligopoly model. This paper keeps the scalable network ambition but changes the object being attributed: a close product-market relation need not be only rivalry. It can be final-use substitution, final-use complementarity, or a relation that is better classified as production dependence. The contribution is therefore not a new centrality statistic, but a primitive-layer decomposition of the pricing and propagation objects.

Relative to Pellegrino (2025), the paper does not add another centrality index; it separates product proximity into rivalry and joint use. Relative to Wu (2025), it keeps the single-price production-network markup logic but replaces a collapsed final-demand object with signed firm-pair final-use layers.

The second strand is differentiated demand and product complementarity. Berry et al. (1995) and Nevo (2001) provide the canonical link between demand elasticities and pricing conduct in differentiated-product markets. Gentzkow (2007) is especially close to the sign problem here: the question is whether a new product steals demand from an existing product or raises demand for it. This paper uses that sign logic in a firm-network setting: complementarity is a final-demand primitive, separate from rivalry and production dependence.

The third strand is production-network propagation (Atalay et al., 2011; Barrot and Sauvagnat,

2016; Bigio and La'O, 2020; Baqaee and Farhi, 2020). These papers show how shocks move through input-output links. Here production links are one layer of a composite propagation system. A final-demand shifter is first reallocated by the signed demand operator $\boldsymbol{H}^F$ and only then propagated by $\boldsymbol{R}^M$. Relative to recent multilayer evidence such as Huremović et al. (2026), the first layer here is signed final-demand exposure rather than financial exposure.

Finally, the paper speaks to markup measurement and network Lerner restrictions (Hall, 1988; De Loecker and Warzynski, 2012; De Loecker et al., 2020; Wu, 2025). The object here is not an observed revenue identity. It is a layer-attribution restriction: final-use elasticity, downstream intermediate-demand elasticity, and upstream cost-feedback elasticity constrain the same markup through outlet structure. Relative to this literature, the contribution is to ask which primitive layer constrains the observed markup, rather than to summarize all market power in a single market-power statistic.

The rest of the paper proceeds as follows. Section 2 defines the primitive layers and the local demand system, including the reduced-form caveat for separating substitution from complementarity. Section 3 isolates the final-demand sign logic. Section 4 derives the benchmark markup attribution object. Section 5 composes signed demand responses with production propagation. Section 6 derives macro incidence for welfare, aggregate prices, and downstream costs. Section 7 maps the theory into measurable decompositions and validation exercises. Section 8 applies the incidence operator to measured AI input use, and Section 9 concludes.

Figure 1 illustrates the basic measurement problem. A close firm pair is not a scalar proximity statistic: it can reflect rivalry, final-use joint use, production dependence, or several relations at once.

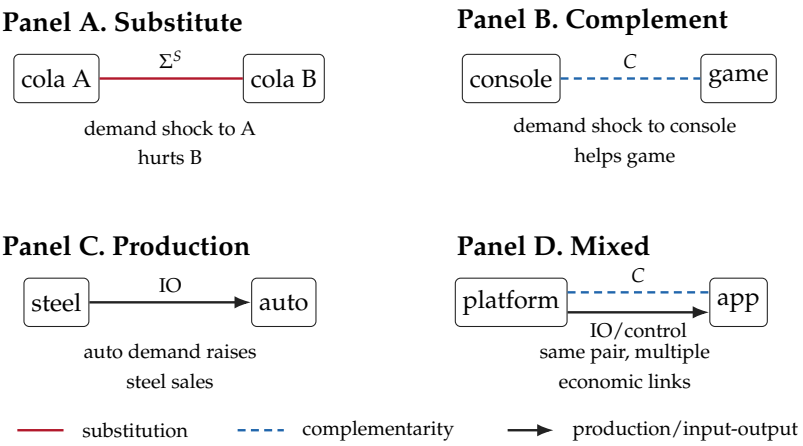


Figure 1: Proximity is not a primitive. The same close firm pair can represent rivalry, final-use complementarity, production dependence, or a mixture. The paper measures these primitive relations separately before composing them into markup and spillover objects.

The main takeaway is that product-market proximity is an input to measurement, not the

economic primitive being propagated. This matters because the paper's markup and spillover objects are composed from separately measured primitives: signed final-use demand first, then production and cost feedback where those margins belong.

2 Primitive Layers and Local Demand

The paper starts from three primitive relations rather than from one generic proximity matrix. Final-demand substitution $\mathbf{\Sigma}^S$ records rivalry for final users, final-demand complementarity $\boldsymbol{C}$ records joint use or compatibility, and the production transaction layer records intermediate input dependence. These layers may coexist for the same firm pair, but they enter different economic margins and therefore must be measured separately. The production transaction layer is later used through $\widetilde{\mathbf{\Omega}}$ for cost feedback and $\boldsymbol{M}$ for output propagation.

2.1 Economic Boundaries of the Layers

Table 1: Primitive-layer measurement discipline

Layer	Economic content	Positive evidence	Main exclusion
$\mathbf{\Sigma}^S$	Horizontal product-market rivalry	competitor mentions, product overlap, TNIC-style similarity	complements, suppliers, broad topical similarity
$\boldsymbol{C}$	Final-use joint demand	compatibility, bundles, co-usage, app or platform ecosystems, "works with" relations	rivalry, vertical buyer-supplier dependence, shared customer labels alone
$\widetilde{\mathbf{\Omega}}, \boldsymbol{M}$	Production and input-output dependence	input-output tables, buyer-supplier links, sectoral use tables	consumer joint use and horizontal rivalry

The first boundary is between product-market rivalry and all other forms of similarity. Two cola producers are close in product space because a buyer who chooses one cola is less likely to choose the other. That relation belongs in $\mathbf{\Sigma}^S$. A broad text similarity score may also put a console maker close to a game publisher, or a cloud platform close to an enterprise software producer, but proximity alone does not say whether the pair competes for the same final purchase, is used jointly by customers, or is linked through production. The substitution layer is therefore not every close pair. It is the part of proximity that acts like final-use rivalry.

The second boundary is between final-use complementarity and vertical dependence. A game and a console are complements because demand for one makes the other more useful to final consumers. A mobile operating system and an app can have the same economic content when the

relation is final use or ecosystem compatibility. By contrast, steel and autos are linked primarily through production: steel demand rises when auto production expands, but that is not a consumer joint-use relation between steel and autos. That link belongs in the production transaction layer, not in $\boldsymbol{C}$. The distinction matters because a final-use complement softens residual demand in the markup theorem, while a production link changes downstream derived demand or upstream input costs.

The same firm pair can appear in multiple layers. A platform and an app may be final-use complements because users value them jointly; the platform may also exercise control over distribution, fees, or access; and firms in related sectors may be connected through input-output tables. The model does not force exclusive labels. It forces each measured relation to enter the margin where it has economic content: final-use rivalry in $\boldsymbol{\Sigma}^S$, final-use joint demand in $\boldsymbol{C}$, and production dependence through $\widetilde{\boldsymbol{\Omega}}$ or $\boldsymbol{M}$.

Terminology: "final use" versus household consumption. The term "final use" carries three related but distinct meanings, and the implementation keeps them separate. In the markup and incidence decompositions, "final use" means non-intermediate or end-use demand, not necessarily household consumption: the two-outlet model partitions output into intermediate use and non-intermediate final use, so the broad final-use outlet includes personal consumption, private fixed investment, government final demand, exports, and inventory changes. In the complementarity layer, $\boldsymbol{C}$ is an end-user joint-use relation, which may involve a household or an enterprise end user but never a current production input. In the welfare closure, statements that rely on a representative household are read as applying most directly to the household/PCE subcomponent or to a reduced-form final-user aggregator. The broad final-use share is the outlet-share object for markup and incidence attribution; the PCE-only share is reported as the closer household-demand robustness.

2.2 Demand Block Intuition

The two final-demand layers collapse into one inverse-demand operator,

$$\boldsymbol{G}^F = \boldsymbol{\Sigma}^S - \boldsymbol{C}, \qquad \boldsymbol{H}^F = (\boldsymbol{I} + \boldsymbol{G}^F)^{-1}. \tag{2}$$

The sign convention is easiest to see in inverse demand. In the local linear-quadratic demand block,

$$p_i = b_i - q_i - \sum_{j \neq i} \Sigma^S_{ij} q_j + \sum_{j \neq i} C_{ij} q_j. \tag{3}$$

More rival quantity lowers firm i's willingness to pay; more complement quantity raises it. The collapsed curvature $G^F_{ij} = \Sigma^S_{ij} - C_{ij}$ is the object inverted in local equilibrium, but the two components have opposite economic interpretations. A positive entry in Σ^S is competitive pressure in final demand. A positive entry in C is joint-use value in final demand.

I use the local inverse-demand system as the main text object. It provides the curvature needed for the sign logic, the final-use elasticity in the markup theorem, and the first stage of the propagation operator. Appendix A gives a representative-consumer construction that rationalizes this block in the symmetric benchmark, together with a latent-projection implementation for the two measured demand layers. The pricing and incidence results below require only the local demand curvature, the two primitive final-demand matrices, and the regularity of the collapsed inverse-demand operator.

2.3 Measurement Discipline and Reduced-Form Caveat

The collapsed object in (2) is useful for equilibrium calculations, but it is not an identification result. Demand alone identifies G^F, not $\mathbf{\Sigma}^S$ and C separately. The empirical strategy is therefore primitive-layer measurement, not deconvolution from a single cross-demand matrix. Appendix A formalizes this reduced-form invariance and the separation bound behind the measurement discipline.

In practice, the measured rivalry layer $\widehat{\mathbf{\Sigma}}^S$ can come from TNIC-style product similarity, product-overlap text, and close competitor mentions. The measured complementarity layer $\widehat{C}$ requires different evidence: compatibility, bundles, co-usage, app ecosystems, accessories, product integrations, or "works with" language. The production layer $\widehat{\mathbf{\Omega}}$ comes from input-output tables, sectoral use tables, or buyer-supplier links; this hat notation denotes the raw empirical transaction layer before it is normalized into $\widehat{\widehat{\mathbf{\Omega}}}$ and $\widehat{M}$. A common corpus can feed the construction of more than one layer, but the same cosine-similarity matrix cannot be used as both $\widehat{\mathbf{\Sigma}}^S$ and $\widehat{C}$. Otherwise the paper would only rename proximity rather than separate economic margins.

This caveat changes the empirical burden of the paper. I do not claim that a single text model can recover substitution and complementarity from the same embedding space without additional structure. The measurement task is instead to assign different pieces of evidence to different primitives. Competitor mentions and product overlap support final-use rivalry. Compatibility, bundling, co-use, and ecosystem language support final-use complementarity. Buyer-supplier information and input-output accounts support production dependence. Once the primitive layers are constructed, the model tells us how they enter markups and spillovers. The empirical object is therefore a validated layer system, not a purely statistical decomposition.

The payoff from this discipline is that the same firm pair can be interpreted without forcing a single label. Consider a platform firm and an app developer. They may be complements for users, competitors in some services, and linked through distribution or fees. A one-layer proximity measure has to compress these relations into one number. The multilayer construction does not. It allows the pair to load on final-use complementarity when users value the two products jointly, on substitution when they compete for the same final purchase, and on production or distribution dependence when one side affects the other's costs or access. This is the sense in which the paper is about attribution: the economic sign of a spillover and the margin constraining a markup depend

on which layer is active.

The reduced-form caveat also explains why the paper is conservative about embedding-based expansion. Embeddings are good at finding semantic neighbors, but semantic neighbors are not automatically substitutes or complements. Without a mechanism label, a close pair may be a rival, a supplier, a customer, an infrastructure relation, a shared technology cluster, or a generic topical association. The quantitative implementation therefore treats embeddings as a search technology and validates candidate links before using them as C. This choice lowers recall, but it keeps the sign of the final-demand operator tied to economics rather than to text distance alone.

3 Signed Final-Demand Responses

The first payoff is a sign result. Holding posted prices fixed, a final-demand shifter first moves final-use quantities through the signed demand operator before the production network is introduced. This is not yet a statement about a full event-time response. An observed event may also move costs, capacity, inventories, or input-output coefficients. The proposition isolates the final-demand component so that the sign carried into later incidence calculations is well defined.

Proposition 1 (Final-demand sign logic)**.** *Fix prices and suppose $\boldsymbol{I} + \boldsymbol{G}^F$ is nonsingular. For a local final-demand shifter* $\mathrm{d}\boldsymbol{b}$*,*

$$\mathrm{d}\boldsymbol{q}^F = \boldsymbol{H}^F \mathrm{d}\boldsymbol{b}. \tag{4}$$

If $\|\boldsymbol{G}^F\|$ is small and $G^F_{ii} = 0$, then for $i \neq j$,

$$\frac{\partial q_i^F}{\partial b_j} = -\Sigma^S_{ij} + C_{ij} + O(\|\boldsymbol{G}^F\|^2). \tag{5}$$

A positive demand expansion to firm j lowers final demand for net substitutes and raises final demand for net complements.

Demand contractions reverse these signs. The proposition therefore applies to the demand-intercept component of an event. Cost, capacity, or policy shocks must first be separated from demand and production-side components.

The result is simple, but it is doing important work. Many empirical spillover exercises start from a nonnegative exposure measure: two firms are close in product space, connected through an input-output link, or located in the same technological neighborhood. Such objects are useful for measuring exposure, but they do not by themselves say whether the first-stage demand effect is positive or negative. In this model, that sign comes from the difference between substitution and complementarity. A shock to a substitute is a negative final-use spillover; a shock to a complement is a positive final-use spillover. The two can be close in the same text space, but they have opposite signs in the demand system.

This sign logic is also why the paper does not interpret C as "negative competition" in a loose sense. C enters with the opposite sign from Σ^S in inverse demand because the economic relation

is different. More quantity of a complement raises the willingness to pay for the product. More quantity of a substitute lowers it. The model lets these effects net out in G^F, but the primitive layers are still needed to understand why the net sign is positive or negative. A collapsed matrix can reproduce the local demand response, but it cannot tell the reader whether that response came from competition, joint use, or a mixture of the two.

The proposition should therefore be read as a measurement test as much as a mathematical statement. If the measured C layer contains mostly rivals, then the sign predictions will be wrong. If the measured S layer contains many complements, the same problem appears in the opposite direction. The high-confidence examples and the 2022 C-validation table below are not decorative; they are what make the sign logic credible. The model supplies the formula, but the empirical layer construction supplies the sign content.

Box 1. Two-firm demand-layer arithmetic

Write $I + G^F = \begin{pmatrix} 1 & g \\ g & 1 \end{pmatrix}$, where $g = \Sigma^S_{12} - C_{12}$. A positive demand shifter to firm 1 gives

$$\frac{\partial q_2^F}{\partial b_1} = -\frac{g}{1-g^2}.$$

Case	g	$\partial q_2^F / \partial b_1$
Net substitute	$+0.3$	-0.330
Net complement	-0.3	$+0.330$

The box fixes the sign convention for the final-demand component. Appendix C uses the same arithmetic as a collapsed-proximity diagnostic. The main text uses it only to identify the quantity vector that later enters markup attribution and production-network propagation.

4 Markup Attribution Across Outlets and Production Margins

The main pricing result studies a firm that posts one price to final consumers and intermediate buyers. The same price loads on three distinct margins: final-use elasticity, downstream intermediate demand, and upstream cost feedback.

There are two demand outlets. In the final-use outlet, households buy differentiated products and the relevant cross-firm objects are $\mathbf{\Sigma}^S$ and $\boldsymbol{C}$. In the intermediate outlet, downstream producers buy inputs, and the relevant object is the transaction layer induced by input-output accounts. Following the normalization discipline in Wu (2025), I use two normalizations of this transaction layer. The cost-share matrix $\widetilde{\mathbf{\Omega}}$ governs which upstream input prices enter a sector's unit cost. The intermediate-demand coefficient matrix $\boldsymbol{M}$, equivalently the revenue/output normalization embedded in downstream purchases and firm allocation weights, governs derived demand and output propagation. One posted price forces the firm to price against all of these margins at once.

The entry for $\boldsymbol{M}$ refers to the direct downstream intermediate-demand coefficient. The full technical appendix also tracks the indirect cost-feedback component inside downstream buyers' expenditure, which is governed by $\widetilde{\boldsymbol{\Omega}}$ for those buyers.

The theorem is stated at a regular single-price equilibrium. The firm problem is stated only at the level needed for the one-price FOC: one posted price, final and intermediate demand cleared jointly, and a local FOC system that can be differentiated. Appendix B gives the full clearing equations and derivative algebra.

Definition 1 (Regular single-price equilibrium)**.** Given $(\boldsymbol{\Sigma}^S, \boldsymbol{C}, \widetilde{\boldsymbol{\Omega}}, \boldsymbol{M}, \boldsymbol{b})$, a regular single-price equilibrium is an interior allocation and price vector such that final demand, intermediate demand, sectoral prices, unit costs, and goods clearing are mutually consistent; each firm chooses one posted price for both final and intermediate buyers; the input-output fixed point is well defined; and the local first-order-condition system is nonsingular. Firm i solves

$$p_i^* \in \arg\max_{p_i} \left[p_i - \kappa_i(p_i, \boldsymbol{p}_{-i}^*)\right] Y_i(p_i, \boldsymbol{p}_{-i}^*), \tag{6}$$

where Y_i is total residual demand from final-use and intermediate outlets after the network-clearing mapping is applied.

Let $\varphi_i = q_i^F / y_i$ be the final-use outlet share, E_i^F the final-use own-price elasticity, E_i^M the intermediate-demand own-price elasticity, and η_i^κ firm i's unit-cost elasticity with respect to its own posted price along the unilateral pricing margin. These are local unilateral pricing elasticities: the derivative is with respect to firm i's own posted price, holding rival posted prices fixed while the model-implied final-demand, intermediate-demand, and cost objects adjust.

On the cost side, the main text uses only a local dual-cost representation. For firm i in sector $s(i)$,

$$\mathrm{d}\log\kappa_i = \beta_{s(i)}^L \mathrm{d}\log w + \sum_r \widetilde{\Omega}_{s(i)r} \mathrm{d}\log P_r, \qquad \eta_i^\kappa \equiv \mathcal{D}_i \log\kappa_i. \tag{7}$$

Here P_r is the sectoral input-price index, and $\widetilde{\Omega}_{s(i)r}$ is the cost-share exposure of firm i's sector to input sector r. Sectoral input prices are local fixed-share aggregates of firm posted prices, and $\mathcal{D}_i$ denotes the local derivative induced by changing firm i's posted price while holding rival posted prices fixed. Thus η_i^κ is not a free markup residual: it is the local upstream cost-feedback elasticity implied by the cost-share normalization along firm i's unilateral pricing margin. Appendix B gives a Cobb-Douglas microfoundation for this local representation.

Theorem 1 (Local markup attribution)**.** *At any regular single-price equilibrium in the LQ/local-cost-share implementation, the local Lerner index satisfies*

$$L_i = \frac{1 - \eta_i^\kappa}{\varphi_i E_i^F + (1 - \varphi_i) E_i^M - \eta_i^\kappa}. \tag{8}$$

A finite, positive markup requires the regularity condition $1 - \eta_i^\kappa > 0$ and $\varphi_i E_i^F + (1 - \varphi_i) E_i^M > \eta_i^\kappa$, so the perceived total elasticity exceeds the cost-feedback elasticity; the full technical appendix states the remaining

equilibrium regularity conditions. The final-use component is governed by the two signed final-demand layers. In the weak-network region,

$$E_i^F \approx \frac{p_i}{a_i}\left[1+\sum_{j\neq i}(\Sigma_{ij}^S - C_{ij})\frac{a_j}{a_i}\right], \qquad a_i = b_i - p_i. \tag{9}$$

The proof is the one-price FOC plus an outlet-share decomposition. Along the unilateral pricing margin, the finite-firm FOC can be written

$$L_i = \frac{1-\eta_i^\kappa}{\varepsilon_i^J - \eta_i^\kappa}, \qquad \varepsilon_i^J = \varphi_i E_i^F + (1-\varphi_i)E_i^M.$$

Substituting the second identity into the first gives (8). The only role of the LQ demand block in the main text is to make the final-use elasticity transparent: substitution exposure increases the final-use elasticity in the weak-network expansion, whereas complementarity exposure decreases it. Appendix B derives the general local wedge, the residual-demand decomposition, and the exact perturbation formula used in quantitative work.

Equation (8) is the paper's central market-power attribution object. It separates the observed firm-level markup into three constraining margins: final-use demand, downstream intermediate demand, and upstream cost feedback. The production transaction layer enters through two normalizations rather than through a single centrality: $\widetilde{\boldsymbol{\Omega}}$ supplies cost shares for upstream cost exposure, while $\boldsymbol{M}$ supplies intermediate-demand coefficients for downstream derived demand. Within the final-use margin, substitution and complementarity determine the local final-use elasticity. Their role is to discipline E_i^F, one component of the markup denominator.

Several implications follow. First, the observed markup is not a sufficient statistic for final-use market power. A high firm-level markup can arise because final users face weak substitution, because downstream intermediate buyers are inelastic, or because cost feedback is small. Conversely, a firm can have strong final-use market power but a modest observed markup if downstream buyers are price sensitive or if the one-price cost-feedback term is large. The theorem therefore shifts the question from "what is the markup?" to "which margin constrains the markup?" This is the economic object that the quantitative section estimates.

Second, final-use complementarity affects markups through elasticity, not through an ad hoc conduct term. In the weak-network expansion, a stronger substitution edge raises E_i^F and lowers the final-use component of the Lerner index. A stronger complementarity edge lowers E_i^F and raises the final-use component. The same C link can also have a welfare value in the GE closure—in the symmetric integrable benchmark—because it represents joint-use utility. This gives complementarity a dual role: it can be a source of real value and a source of residual-demand softening. The markup theorem isolates the second force; the welfare section returns to the first.

Third, the two production normalizations answer different questions. The cost-share normalization $\widetilde{\boldsymbol{\Omega}}$ answers an upstream question: if firm i raises its posted price, how much does that

movement feed back into input costs through the sectors it buys from or affects? The intermediate-demand normalization $\boldsymbol{M}$ answers a downstream question: how much does demand from buyer sectors discipline firm i's posted price and how do output requirements propagate? Collapsing these into one production centrality would hide the distinction between cost feedback and derived demand. This is why the one-price Lerner denominator has both E_i^M and η_i^κ.

Fourth, outlet shares determine when final-use demand matters for observed firm markups. The final-use elasticity is multiplied by φ_i. A firm with strong complementarity exposure but little final-use outlet share may have limited final-use markup attribution in the pooled one-price object. The same complementarity exposure matters more for a firm that sells heavily to households or final users. This is a useful empirical prediction: the importance of C for markups should be largest where the final-use outlet is large enough for final-demand curvature to discipline the pooled price.

Remark: segmented-pricing benchmark. The single-price formula is the benchmark for a firm-level markup when final-use and intermediate-use prices are not separately observed. If the two outlets can be priced separately, the attribution changes: each outlet has its own Lerner condition and firm-level markups are aggregated after outlet-specific pricing. In the cost-taking comparison,

$$L_i^{1p} = \frac{1}{\varphi_i E_i^F + (1-\varphi_i) E_i^M}, \qquad \bar{L}_i^{seg} = s_i^F \frac{1}{E_i^F} + s_i^M \frac{1}{E_i^M}.$$

Here s_i^F and s_i^M are outlet revenue shares; when comparing to the one-price benchmark on the same outlet mix, set $s_i^F = \varphi_i$ and $s_i^M = 1 - \varphi_i$. The segmented benchmark therefore clarifies what is being attributed in Theorem 1: not a universal pricing rule, but the constraining margins behind a single observed firm-level markup. Appendix B gives the full segmented-pricing benchmark.

5 Demand-Production Composition

Markup attribution explains the pricing margin. Output spillovers require a second composition. A demand shock first moves final-use demand through the signed demand operator and is then propagated through input-output links. Proposition 2 is an incidence object, not a pricing object. It fixes prices and input coefficients to isolate how a final-demand shifter travels through layers. The purpose is to separate the sign chosen by final-demand measurement from the amplification or attenuation created by production propagation.

Proposition 2 (Composite demand-production incidence)**.** *Fix prices, marginal costs, and the intermediate-demand coefficient matrix $\boldsymbol{M}$ at a regular allocation with $\rho(\boldsymbol{M}) < 1$. Let $\boldsymbol{R}^M = (\boldsymbol{I} - \boldsymbol{M})^{-1}$. For a local final-demand shifter,*

$$\mathrm{d}\boldsymbol{y} = \boldsymbol{R}^M \boldsymbol{H}^F \mathrm{d}\boldsymbol{b}. \tag{10}$$

In the weak-final-demand-network region,

$$\frac{\partial y_i}{\partial b_j} \approx R^M_{ij} - \sum_k R^M_{ik}\Sigma^S_{kj} + \sum_k R^M_{ik}C_{kj}. \tag{11}$$

The proposition says that neither layer alone is the observed output response. The final-demand operator determines which final-use quantities move, and the production inverse determines how those movements become total output and intermediate-input demand. Figure 2 summarizes this ordered propagation mechanism: the final-demand shifter first becomes a signed quantity vector, and production links only then propagate that vector.

Corollary 1 (Final-use and customer-demand incidence)**.** *Fix the same prices, marginal costs, and input coefficients as in Proposition 2. If the economy receives a final-use demand shifter* $\mathrm{d}\boldsymbol{b}^F$ *and an externally measured customer-demand or intermediate-demand shifter* $\mathrm{d}\boldsymbol{x}^M$*, then the fixed-price incidence object is*

$$\mathrm{d}\boldsymbol{y} = \boldsymbol{R}^M\left(\boldsymbol{H}^F\mathrm{d}\boldsymbol{b}^F + \mathrm{d}\boldsymbol{x}^M\right). \tag{12}$$

The corollary is an outlet discipline. Final-use shocks enter the signed final-demand system before production propagation. Customer-demand or intermediate-demand shocks enter the goods-demand vector directly and are then propagated by $\boldsymbol{R}^M$. When an externally measured activity vector can be split into non-intermediate and intermediate outlets, only the non-intermediate branch belongs inside $\boldsymbol{H}^F$. This corollary is the outlet discipline used in the AI application of Section 8. A measured final-use AI demand vector enters $\boldsymbol{H}^F$ before production propagation, so it cannot be replaced by a customer-pull or intermediate-demand vector, which enters goods demand directly and is then propagated by $\boldsymbol{R}^M$. A customer-pull vector is therefore demand-side but not automatically a final-use object.

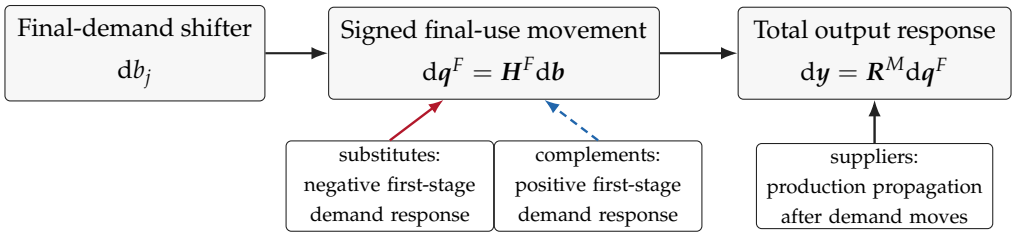


Figure 2: Ordered propagation: signed demand first, production second.

The main takeaway is that output spillovers are not demand-only and not production-only. Substitution and complementarity choose the sign of the first-stage final-use movement; $\boldsymbol{R}^M$ then maps that movement into total output and intermediate-input demand. For customer-demand shifters, the same production propagation applies without first treating the customer-pull vector as a final-use shifter. This matters because the same observed sales response can come from different ordered paths, so a direct IO benchmark, a customer-demand benchmark, and the composite object $\boldsymbol{R}^M\boldsymbol{H}^F$ answer different questions.

Empirical diagnostics. Proposition 2 implies three empirical cautions. First, production-only exposure is incomplete because total output responses include final-demand reallocations before IO propagation. Second, total sales are not final-demand sales: observed output combines final-use demand and intermediate demand. Third, one-layer proximity measures are useful benchmarks, but the main object is the separated demand-production composition $\boldsymbol{R}^M\boldsymbol{H}^F$. Appendix C states the diagnostic conditions formally without adding additional main-text corollaries.

The ordering also helps interpret sales regressions. Suppose a firm gains sales after a positive shock to a nearby firm. A production-only reading would say that the firm is upstream or downstream from the shocked firm. A demand-only reading would say that the firm is a complement. The composite object allows both. The shock may raise final-use demand for complements, and that final-use movement may then require output from suppliers. Or the shock may reduce final-use demand for substitutes while still raising output for some upstream firms. The same observed sales response can therefore be generated by different ordered paths.

This is especially important when the observed outcome is total sales. Total sales mix final-use sales and intermediate sales. A firm may lose final-use customers because a substitute expands, while gaining intermediate demand because a downstream sector expands. Without outlet information or a model that separates layers, the sign of the final-use response cannot be read from total output alone. The paper's decomposition is meant to discipline exactly this ambiguity. It does not claim that every observed event is a clean demand shock; it gives the map needed once the demand component has been isolated.

The production inverse also changes magnitudes. A small final-use movement at a central downstream firm can require large upstream output when the input network is deep. Conversely, a large final-use movement at a firm with few input-output links may have little aggregate production propagation. The spillover object therefore has two pieces of economic content: the sign chosen by the demand layers and the amplification chosen by the production layer. This is why the quantitative section reports both signed-demand exposure and direct IO benchmarks rather than replacing one with the other.

6 Benchmark Local Quantitative GE

The preceding theorem is a local pricing theorem. To interpret welfare and aggregate price incidence, the model adds a local GE closure with an active household budget, labor clearing, profit rebates, production costs, input-output clearing, and firm first-order conditions. There are therefore two levels of equilibrium in the paper. Theorem 1 is the local pricing restriction that defines which margins constrain one posted price. The GE closure does not replace that theorem; it gives price-index, welfare, and cost-incidence meanings to local markup movements. All comparative statics in this section are evaluated along a regular local branch. The model does not claim global uniqueness of the equilibrium correspondence.

The pricing theorem uses the quasi-linear perceived-demand margin, so E_i^F is the local final-use elasticity along firm i's own-price direction. If firms internalize household income feedback at the pricing margin, E_i^F is replaced by the active-budget GE elasticity $E_{ii}^{F,GE}$ defined in the full technical appendix; the benchmark below keeps the quasi-linear margin and uses the closure only to attach welfare and price-incidence meanings to the same local allocation.

This section states the benchmark economy of the paper. Local markup and demand movements are evaluated as comparative statics of a regular active-budget single-price general equilibrium: the same nonsingular local system maps primitive shocks into prices, quantities, income, and money-metric welfare along a regular comparative-static branch. The objects reported below are sub-blocks of that branch rather than a separate accounting device. The welfare envelope gives the household or final-user interpretation of the demand block. The broad final-use outlet share used in the implementation is an outlet-share proxy for non-intermediate demand; PCE-only shares are reported as the closer household-demand robustness.

The household maximization problem is kept in Appendix D; the main text uses only the closure definition, the welfare envelope, and the price-incidence equation.

Definition 2 (Regular active-budget GE closure)**.** A regular active-budget GE extends the regular single-price equilibrium by requiring the household budget, labor market, profit rebates, production costs, input-output clearing, and firm FOCs to hold jointly. The equilibrium is regular if the local system of household FOCs, goods clearing, labor clearing, income closure, unit-cost equations, and firm FOCs has a nonsingular Jacobian. Counterfactuals below are evaluated along this regular local branch, not as a global uniqueness claim.

Proposition 3 (Local GE welfare and price incidence)**.** *At a regular active-budget GE, the local money-metric welfare envelope can be written as*

$$\mathrm{d}\mathcal{W} = \mathrm{d}Y - \boldsymbol{q}^{F\prime}\mathrm{d}\boldsymbol{p} + \frac{1}{\lambda^H}\left[\boldsymbol{q}^{F\prime}\mathrm{d}\boldsymbol{b} - \frac{1}{2}\boldsymbol{q}^{F\prime}\mathrm{d}\boldsymbol{B}^F\boldsymbol{q}^F\right], \qquad \boldsymbol{B}^F = \boldsymbol{I} + \boldsymbol{\Sigma}^S - \boldsymbol{C}. \tag{13}$$

Local markup movements pass through production costs and sectoral price indices according to

$$\mathrm{d}\log\boldsymbol{p} = \boldsymbol{A}_{\widetilde{\Omega}}\left[\mathrm{d}\log\boldsymbol{\mu} + \boldsymbol{\beta}^L\mathrm{d}\log w\right], \qquad \boldsymbol{A}_{\widetilde{\Omega}} = (\boldsymbol{I} - \boldsymbol{B}_{\widetilde{\Omega}}\boldsymbol{\Lambda})^{-1}. \tag{14}$$

The same nonsingular local system delivers a regular comparative-static branch for primitive shocks to $(\boldsymbol{b}, \boldsymbol{A}, \boldsymbol{\Sigma}^S, \boldsymbol{C}, \widetilde{\boldsymbol{\Omega}}, \boldsymbol{M})$.

Direct welfare value of complementarity. The last term in (13) is the direct demand-layer effect. When the perturbation is an increase in complementarity and other primitives are held fixed, the term contains $q^{F\prime}\mathrm{d}Cq^F/(2\lambda^H)$. This is positive for an entrywise increase in a symmetric complementarity layer. Complementarity is therefore not merely "less substitution." It can directly raise joint-use value even before any markup or price response is considered.

Quantitatively the measured complementarity layer is sparse, so this direct value and the markup softening it induces are modest. Table 2 removes the measured complementarity layer in the benchmark equilibrium. At the preferred ratio $r = \alpha_C/\alpha_S = 0.12$, removing C lowers sales-weighted markups by 0.056 percent and the final-use price index by 0.076 percent, and moves money-metric welfare by -0.029 percent of model income once markups respond, or -0.024 percent at constant markups. The effect scales with r over the range 0.10 to 0.30 and is smaller for the strict-core layer. Complementarity refines the demand side; it is not a first-order welfare force in this sample.

Table 2: Removing the measured complementarity layer in the benchmark equilibrium, 2022

Counterfactual (remove measured C)	$d\log\mu$ (%)	$d\log P^F$ (%)	direct/Y (%)	dW/Y const-μ (%)	dW/Y endog-μ (%)
Main C, $r = 0.10$ (range low)	-0.0466	-0.0632	-0.0320	-0.0198	-0.0241
Main C, $r = 0.12$ (preferred)	-0.0561	-0.0760	-0.0384	-0.0237	-0.0289
Main C, $r = 0.30$ (range high)	-0.1428	-0.1936	-0.0961	-0.0584	-0.0719
Strict-core C, $r = 0.12$	-0.0406	-0.0245	-0.0171	-0.0096	-0.0458

Notes: Comparative static $\mathrm{d}C = -\alpha_C\hat{C}$ along the regular branch. The constant-markup column fixes μ; the endogenous-markup column lets markups respond through the exact network perturbation $\mathrm{d}H^F = -H^F\,\mathrm{d}G^F\,H^F$ of the final-use elasticity. The direct term is $q^{F\prime}\mathrm{d}C\,q^F/(2\lambda^H Y)$.

Markup pass-through. The matrix $A^{\widetilde{\Omega}}$ in (14) is the local production-network pass-through operator built from the cost-share normalization. A markup increase at one firm raises its posted price; that price enters sectoral input price indices; and those indices enter the unit costs of downstream firms. Here Λ maps firm posted prices into sectoral input-price indices, and $B^{\widetilde{\Omega}}$ maps sectoral input prices into firm unit costs; the full technical appendix gives the fixed-share definitions and derivation. The inverse in $A^{\widetilde{\Omega}} = (I - B_{\widetilde{\Omega}}\Lambda)^{-1}$ sums those cost-price feedbacks. Thus the aggregate final-price impact of a markup movement depends not only on the firm's own markup but also on where the firm sits in production price propagation and in final-use expenditure weights.

One-price cross-outlet cost externality. Because one posted price is charged to both final users and intermediate buyers, a markup movement that originates in the final-demand side can raise input costs elsewhere. A complementarity shock can lower a firm's final-use elasticity and raise its markup; the higher posted price then applies to downstream firms that use the product as an input. Under full segmentation, the final-use price need not enter the intermediate input price index. This cross-outlet cost externality is therefore a one-price channel.

The size of this externality is not negligible. The pass-through calculation evaluates $\mathrm{d}\log\boldsymbol{\kappa} = \boldsymbol{B}_{\widetilde{\Omega}}\boldsymbol{\Lambda}\boldsymbol{A}_{\widetilde{\Omega}}\,\mathrm{d}\log\boldsymbol{\mu}$ for final-demand-driven markup movements. A one percent markup increase at the top-decile complementarity-central firms raises the final-use price index by 0.051 percent di-

rectly and by a further 0.088 percent through the input-cost externality, so the production network nearly triples the price impact, and the externality reaches 4,289 firms that did not change their own markup. Under segmentation the externality is identically zero. The one-price assumption is thus quantitatively consequential, not a harmless normalization.

The closure supports several located exercises, each holding a different set of objects fixed. In the submission version, the distinction is stated in the exercise text: the complementarity and cost-cut exercises are local-GE objects, the AI incidence table is a fixed-network incidence index, and the demand-pull exercise is a partial-equilibrium sign-composition diagnostic.

7 Primitive-Layer Measurement and Quantitative Attribution

This section reports the quantitative objects that support the paper's multilayer attribution claim. The purpose is not to fit markup levels with every available primitive. I instead discipline different margins with different evidence: substitution is treated as a final-use rivalry layer, complementarity is validated as the missing final-use primitive, production objects discipline downstream demand and upstream cost feedback, demand scales are fixed by external demand and markup-share moments, and pass-through evidence anchors η_i^κ. Markups are then used for attribution and diagnostics.

The section follows the order of the model. It first lists the primitive-layer system, then validates the signed final-demand measurement, documents the production and outlet-share objects, states the calibration hierarchy, reports markup attribution, and finally compares signed final-demand propagation with direct input-output exposure. It also states the event-validation design that turns the signed-layer prediction into a short-window outcome test. Welfare, regularity, embedding, collapsed proximity, placebo C, and reduced-form fit checks are reported as diagnostics rather than calibration targets.

7.1 Empirical Primitive-Layer System

The minimal empirical package has four objects. First, the substitution layer $\widehat{\boldsymbol{\Sigma}}^S$ measures final-use rivalry. Its natural inputs are TNIC-style product descriptions, product-overlap text, and close competitor mentions. Second, the complementarity layer $\widehat{\boldsymbol{C}}$ measures final-use joint demand. Its inputs are compatibility language, bundles, co-usage, accessories, app ecosystems, product integrations, and "works with" relations. Third, the production layer $\widehat{\boldsymbol{\Omega}}$ measures input-output exposure using sectoral use tables or buyer-supplier links; it is then normalized into $\widehat{\widehat{\boldsymbol{\Omega}}}$ for cost feedback and $\widehat{\boldsymbol{M}}$ for output propagation. Fourth, the implementation matches these layers to firm outcomes: sales, final-use proxies where available, COGS or gross margins, sector assignment, event shocks, and outlet-share proxies.

The key discipline is that $\widehat{\boldsymbol{C}}$ must be a final-use object. A cloud provider and an enterprise application may be close in text because the application runs on the cloud; that fact alone is vertical

or infrastructure dependence, not final-use complementarity. A console and a game, a phone and an app, or a payment network and a merchant-facing checkout service are closer to the complementarity concept because the final user values the products jointly. The implementation therefore records the evidence behind high-C pairs rather than treating a similarity score as self-explanatory.

C receives more validation than the other primitives because it is the least standard measured object. Rivalry layers build on the TNIC/product-market tradition, and production layers build on input-output accounting. There is no off-the-shelf firm-level complementarity matrix. The empirical question is therefore not whether the whole paper is a C-matrix paper; it is whether the missing final-use primitive is credible enough to complete signed final demand.

The production and outlet-share objects receive parallel empirical discipline: the cost-share normalization governs the pass-through wedge η_i^K, the M normalization governs downstream intermediate-demand discipline and production propagation, and φ_i determines how strongly final-use demand enters a firm's one-price markup denominator. The currently completed historical C layers cover 2000, 2005, 2010, 2015, 2020, and 2022, so I report the outlet-share and markup-attribution objects as a short panel rather than as a single-year cross-section. The detailed production and outlet-share diagnostics remain in the full technical appendix.

7.2 Outlet Shares and IO-Imputed Mixed Exposure

The outlet share is measured before markups are decomposed. I do not infer φ_i from the markup equation. The baseline object is an IO-imputed final-use, or non-intermediate, outlet share:

$$\widehat{\varphi}_{i,t}^{FU} = \widehat{\varphi}_{s(i,t),t}^{FU}.$$

At the BEA sector level, the numerator is broad final use and the denominator is broad final use plus intermediate use. Broad final use includes personal consumption expenditures, private fixed investment, government final demand, exports, and inventory changes as recorded in the use table. This is the right object for the two-outlet model because the model partitions output into non-intermediate and intermediate uses, not into household consumption and everything else. A PCE-only share is reported as a narrower household-demand robustness, but it is not the baseline φ_i.

Firm-level outlet shares are therefore firm-indexed but partly sector-imputed. Without firm-level channel or customer-type data, I do not claim point identification of the true firm outlet mix. The empirical claim is more modest: BEA use tables discipline the sectoral split between final use and intermediate use, and the firm-year implementation assigns that split through sector mappings. The panel then asks whether the assigned outlet exposure is economically heterogeneous and whether firms receive non-degenerate exposure to both final and intermediate uses under the IO mapping.

Mixed exposure under the IO sector mapping is not a knife-edge assumption. Across the six completed C-layer years, the sales-weighted broad final-use share is between 0.52 and 0.56,

while the PCE-only share of total use is lower, between 0.31 and 0.36. The difference is expected: investment goods, government purchases, exports, and inventories are final uses in the IO sense but are not household consumption. Using the baseline assigned mixed-exposure definition $0.1 < \widehat{\varphi}_{i,t} < 0.9$, about 86–92 percent of firms, and 86–88 percent of sales, lie away from the two pure-outlet extremes. Figure 3 makes the measurement granularity visible: the mass points reflect sector-level use-table assignment, not direct firm-level channel data.

This same use-table distinction becomes central in the AI application. Industry-buyer columns classify current intermediate demand and build $\boldsymbol{M}$ and $\boldsymbol{R}^M$; final-use columns classify non-intermediate demand and build the outlet share φ_s and $\mathrm{d}\boldsymbol{b}^F$-type final-use AI exposure vectors; the PCE subset is reported only as a narrower household-use robustness. The production matrix excludes the final-use columns by construction, so those columns never enter the input–output operator; they are retained precisely to construct φ_s and the BEA final-use AI demand shifter of Section 8.3. A firm-customer relation therefore need not be treated mechanically as current intermediate demand: business purchases of capital equipment, exports, government purchases, inventories, and household purchases are final uses in the IO sense, while current services and materials purchased as production inputs are intermediate uses.

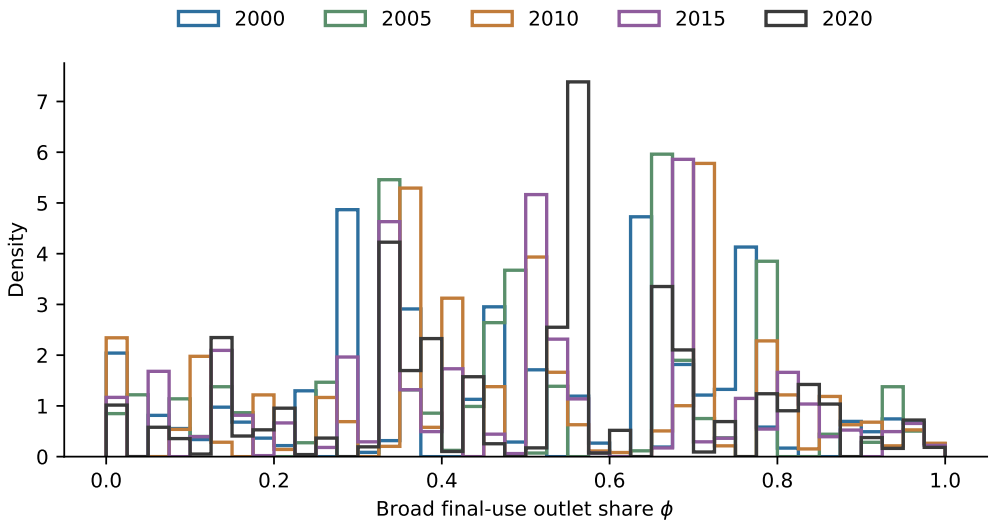


Figure 3: Distribution of broad final-use outlet shares. The figure reports the IO-imputed broad final-use share $\widehat{\varphi}_{i,t}^{FU}$ for the years with completed C layers. The mass points reflect sector-level use-table assignment. PCE-only shares are lower and are reported as a robustness object rather than as the baseline two-outlet share.

7.3 Validating the Signed Final-Demand Layer

The empirical foundation for signed final demand is the separate measurement of rivalry and joint-use relations. The new object is a final-use complementarity layer that is sparse, mechanism-based, and distinct from product-market rivalry. Table 3 separates the construction rules from their paper roles. The baseline expanded layer, labeled local balanced, contains 33,085 links, or 0.336 percent of the 4,437-firm pair universe. A high-precision lower bound contains 10,574 links. The broad candidate universe is kept as an audit pool, not as a paper-facing C matrix. Figure 4 then reports the empirical density and overlap checks against TNIC-style rivalry and IO exposure.

The embedding exercise is deliberately weaker than matrix construction. It ranks candidates outside the local-balanced layer and then asks whether a separate mechanism judge confirms final-use complementarity. Only 131 of the top 1,000 embedding-ranked candidates are confirmed. I therefore use embedding as a recall audit and candidate-discovery device, not as a standalone complementarity matrix.

Table 3: C-layer construction and validation roles, 2022

Layer	Construction rule	Pairs	Paper role	Main risk
Strict core	Prior judge plus strict deterministic rules	10,574	High-precision lower bound	Low recall
Local high precision	Deterministic rules with tighter local evidence	16,184	Conservative robustness layer	Misses weaker complements
Local balanced	Deterministic expanded rules	33,085	Baseline C layer	Generic links
Broad candidate universe	Search pool before final C screening	105,797	Audit only	Too broad for structural C
Embedding top-1,000 confirmed	Embedding recall audit plus mechanism judge	131	Candidate discovery	Not a standalone matrix

Notes: The denominator is all unordered pairs among 4,437 2022 firms. The embedding row reports confirmed additions among the top 1,000 embedding-ranked candidates outside the local-balanced C layer; it is an LLM-confirmation recall audit, not a hand-audited precision estimate. Figure 4 reports the corresponding density and overlap diagnostics against rivalry and IO layers.

Figure 4 reports whether the measured complementarity layer is just TNIC rivalry or strong IO exposure under another name. It puts the main density and overlap diagnostics in one place because this is the reader's first empirical check on the new primitive.

The main takeaway is that local-balanced C is sparse and mostly outside the two most likely confounds: only 11.5 percent of its links are also TNIC rivalry links, 5.2 percent are top-IO links, and about 85 percent are neither. This matters because the later sign and markup exercises require C to carry joint-use content rather than generic product-market proximity or production depen-

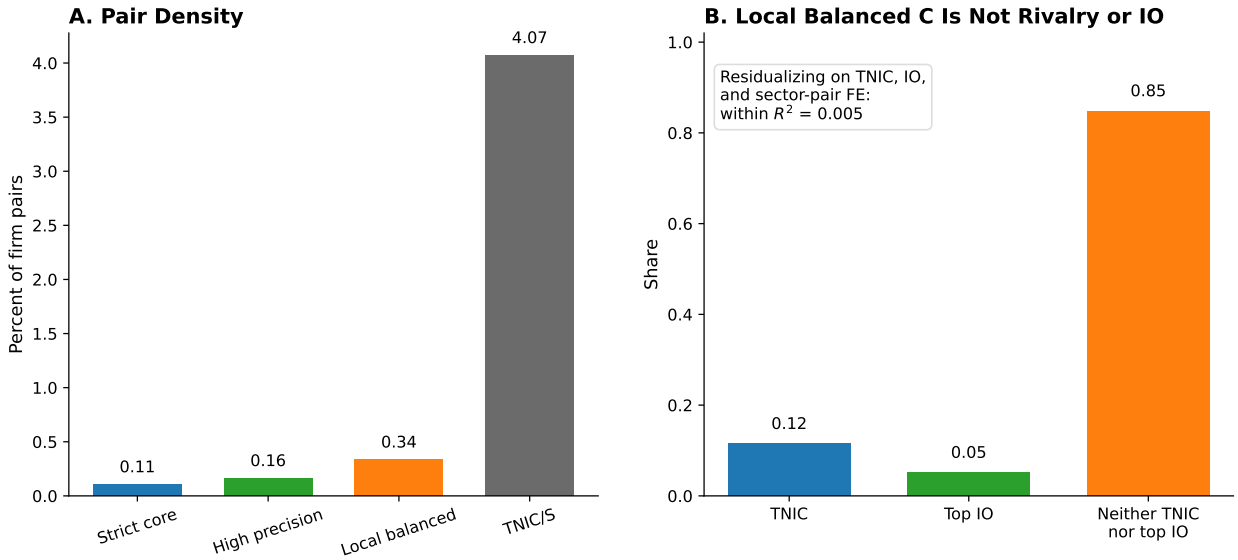


Figure 4: C-layer validation dashboard, 2022. Panel A compares the density of the measured complementarity layer with TNIC-style rivalry. Panel B reports the overlap of the baseline local-balanced C layer with rivalry and strong IO exposure. The residual within-sector-pair R^2 is reported as an annotation because it is a residual diagnostic, not an overlap share.

dence.

The table and figure have one main role: they show that C is a thin, mechanism-confirmed joint-use layer, not a generic similarity matrix. The accepted links require evidence that one product raises the value or use of the other for the same final user. Shared product categories, common technology words, customer overlap, and supply-chain dependence are not enough because those relations belong to other layers of the model.

This convention disciplines the later quantitative exercises. In markup attribution, C can lower residual-demand elasticity only where there is a joint-use relation. In spillover exercises, C creates positive final-use propagation while substitution creates negative final-use propagation. The embedding audit is used only as a recall check; the full technical appendix reports the detailed positive-control, boundary-screen, and residual-C diagnostics behind the dashboard.

7.4 Calibration Strategy

The measured layers must be placed on a common demand scale before they enter G^F. Write

$$\mathbf{\Sigma}^S = \alpha_S \widehat{\mathbf{\Sigma}}^S, \qquad \boldsymbol{C} = \alpha_C \widehat{\boldsymbol{C}}, \qquad r = \alpha_C / \alpha_S.$$

The hatted matrices are measurement objects. Their scale must be fixed by an explicit normalization, such as row-sum or operator-norm normalization, before the scale parameters receive a quantitative interpretation. A regular implementation also requires $I + \alpha_S \widehat{\mathbf{\Sigma}}^S - \alpha_C \widehat{\boldsymbol{C}}$ to be invertible on the relevant local branch. This condition prevents the complementarity layer from making

the local demand system explosive. It is a regularity screen, not a substitute for validating the economic content of the layer.

The paper-facing calibration hierarchy is simple. The substitution scale α_S is fixed at 0.12 as an external final-demand scale, in the spirit of product-market demand-network calibrations such as Pellegrino (2025). The complement scale is reported as the ratio $r = \alpha_C/\alpha_S$. I treat $r = 0.12$ as the preferred center and $[0.10, 0.30]$ as the maintained quantitative range, disciplined by a weak markup-share slope check rather than by full markup-level fit. Cost feedback is disciplined separately by pass-through evidence. The observed markup identifies the joint denominator $\varphi_i E_i^F + (1-\varphi_i)E_i^M - \eta_i^\kappa$; it does not identify each primitive margin separately.

Observed markups enter in two limited ways. First, I use the Compustat-based translog gross markup `markup_tl_kx`, constructed with COGS as the variable input, as the observed markup object. I convert it firm by firm into a price-cost margin, $L_i = 1 - 1/\mu_i^{TL,KX}$, and report sales-weighted moments in the markup attribution table. Second, the maintained range for $r = \alpha_C/\alpha_S$ is disciplined by a weak markup-share slope check using $\log \mu_i$. Neither use fits markup levels to choose the primitive layers or the cost-feedback wedge. In particular, observed markups are not used to estimate $\widehat{\boldsymbol{\Sigma}}^S$, $\widehat{\boldsymbol{C}}$, $\widehat{\boldsymbol{M}}$, $\widehat{\boldsymbol{\Omega}}^e$, φ_i, α_S, or η_i^κ.

The hierarchy is important because the same observed markup can be generated by different primitive margins. I therefore do not choose α_S, α_C, outlet shares, and η_i^κ jointly to match markup levels. The substitution scale is fixed by an external demand anchor, $r = \alpha_C/\alpha_S$ is reported as a maintained range that preserves regularity and the markup-share slope check, and cost feedback is disciplined by pass-through evidence. Observed markups then become attribution objects: they show how separately disciplined margins combine in the one-price Lerner denominator.

7.5 External Discipline for Cost Feedback

The cost-feedback term is disciplined by pass-through evidence, not by fitting observed markups. The markup equation shows why. A markup identifies the joint denominator $\varphi_i E_i^F + (1-\varphi_i)E_i^M - \eta_i^\kappa$; it does not by itself separate final-use demand, intermediate demand, and upstream cost feedback. The empirical target for η_i^κ is therefore an output-price response to an upstream cost shock,

$$\rho_{MC}^{data} = \frac{\mathrm{d}\log p^{out}}{\mathrm{d}\log \kappa^{cost}}.$$

I use the U.S. manufacturing evidence in Ganapati et al. (2020) as the main external anchor. Their energy-cost design implies a short- to medium-run marginal-cost pass-through of about 0.70. The baseline supply-side calibration therefore chooses a common or sector-scaled cost-feedback parameter so that the model-implied pass-through matches

$$\rho_{MC}^{model}(\eta; \alpha_S, r, \widehat{\widehat{\boldsymbol{\Omega}}}, \widehat{\boldsymbol{M}}, \varphi) = 0.70,$$

holding the demand scale α_S and the complement-to-substitute ratio $r = \alpha_C/\alpha_S$ fixed at their demand-side values. In the first implementation,

$$\eta_i^\kappa = \eta\, x_i^\Omega,$$

where x_i^Ω is a normalized cost-feedback exposure constructed from $\widehat{\widehat{\mathbf{\Omega}}}$; the common-$\eta$ case sets $x_i^\Omega = 1$. I report robustness to pass-through targets 0.60 and 0.80. As a public-data validation, I also estimate the same object using BLS satellite input-to-industry price indexes and PPI industry output indexes (U.S. Bureau of Labor Statistics, 2026). A 2018–2026 panel of matched three-digit NAICS industries gives a twelve-month pass-through coefficient of 0.63 with industry and calendar-month fixed effects. The full technical appendix reports the coefficient, robust standard error, fixed effects, sample, and period. This public-data estimate is a validation moment, not a causal replacement for the input-cost-shock evidence. The cost-taking case $\eta_i^\kappa = 0$ remains a transparent benchmark, but it is no longer the maintained supply-side calibration. Evidence on variable markups and incomplete cost pass-through, such as Amiti et al. (2019), provides a secondary robustness benchmark rather than the main U.S. target.

7.6 Outlet Shares and Markup Decomposition

The empirical counterpart of Theorem 1 is the denominator of the one-price Lerner index. A complete implementation reports the three constraining components

$$\varphi_i E_i^F, \qquad (1-\varphi_i)E_i^M, \qquad \eta_i^\kappa.$$

These objects answer different questions. The first asks how much the final-use market constrains the firm. The second asks how much downstream intermediate buyers constrain the firm. The third asks how much a price change feeds back into the firm's own unit cost through upstream input prices. A table can report the distribution by firm and sector, the firms with high complementarity exposure, and the firms for which complementarity matters most because the final-use outlet share is large.

This decomposition also clarifies a non-identification point. An observed markup identifies the joint denominator $\varphi_i E_i^F + (1-\varphi_i)E_i^M - \eta_i^\kappa$, not the demand-side primitive scale parameters alone. A numerical implementation must discipline final-use demand, intermediate demand, and cost feedback separately. The pass-through target above provides the external discipline for η_i^κ; markups are then used for attribution and validation rather than for calibrating every primitive at once.

Table 4 reports the pass-through anchored markup attribution. The cost-taking benchmark provides the pre-pass-through reference, but it does not discipline cost feedback. Under the literature pass-through target $\rho_{MC} = 0.70$, the sales-weighted observed Lerner index is 0.280, while the no-cost-feedback counterfactual is 0.338. Cost feedback lowers the no-cost-feedback Lerner by 0.059. Under the BLS validation target, $\rho_{MC} = 0.6267$, the no-cost-feedback Lerner is 0.358 and the

cost-feedback effect is -0.078. The reading is not that the model mechanically fits markup levels. It is that the final-use pricing term is quantitatively meaningful relative to intermediate-demand discipline, and cost feedback materially lowers the no-cost-feedback Lerner.

Here and below, "price pressure" means an effective denominator contribution in the one-price Lerner restriction, not a claim about price levels or inflation.

Table 4: Calibration-Conditional Markup Attribution: Effective Denominator Terms, 2022

Target	Obs. L	No-cost-fb. L	Cost-fb. effect	Final-use pressure	Intermed. pressure	Cost-fb. pressure
Literature, $\rho_{MC} = 0.70$	0.280	0.338	-0.059	0.395	0.733	-0.128
BLS, $\rho_{MC} = 0.6267$	0.280	0.358	-0.078	0.444	0.740	-0.184

Notes: Reported moments are sales-weighted. The no-cost-feedback Lerner sets the pass-through/cost-feedback term to zero while holding the calibrated demand components fixed. Observed L is constructed firm by firm as $L_i = 1 - 1/\mu_i^{TL,KX}$ using the translog COGS-based gross markup and then sales-weighted in the pass-through sample. The final-use, intermediate-demand, and cost-feedback columns are normalized terms in the one-price pricing condition; they are not revenue shares and need not sum to one when the cost-feedback term is negative.

The final three columns should be read as price-pressure attribution, not as revenue shares. The final-use term is the margin that changes when substitution or complementarity changes residual demand; the intermediate term is discipline from downstream buyers; and the negative cost-feedback term captures the one-price channel through upstream input costs. The main message is not that final-use demand explains all markups. Under the baseline pass-through target, the final-use pricing term is 0.395, or 54 percent of the downstream intermediate-demand term of 0.733, while cost feedback lowers the no-cost-feedback Lerner by 0.059. The BLS validation target gives the same qualitative reading. This is why markups are used for attribution after the primitive margins have been separately disciplined.

These are not causal estimates of markup changes. They are the accounting attribution implied by separately measured primitives, BEA outlet shares, and externally disciplined pass-through.

The same accounting attribution is available for the short historical C-layer panel. Across 2000, 2005, 2010, 2015, 2020, and 2022, demand-only Lerner indexes remain systematically above observed Lerner indexes once cost feedback is removed, and the final-use term accounts for roughly 40–48 percent of the effective denominator. The full technical appendix keeps the panel decomposition, sector, and segmented-pricing tables because their role is to check stability, not to introduce a separate main-text result.

7.7 Demand-Side Attribution Benchmarks

The spillover exercise compares one-layer and composed exposure objects. For a vector of final-demand shifters $d\boldsymbol{b}$, let $\Delta\boldsymbol{q}^0$ denote the direct final-use quantity movement used by the production-only benchmark. Under the own-curvature normalization used in the quantitative exercise, $\Delta\boldsymbol{q}^0 =$

d$\boldsymbol{b}$. The three benchmarks are

$$\Delta \boldsymbol{y}^{IO} = \boldsymbol{R}^{M} \Delta \boldsymbol{q}^{0}, \qquad \Delta \boldsymbol{q}^{D} = \boldsymbol{H}^{F} \mathrm{d}\boldsymbol{b}, \qquad \Delta \boldsymbol{y}^{full} = \boldsymbol{R}^{M} \boldsymbol{H}^{F} \mathrm{d}\boldsymbol{b}. \tag{15}$$

The first object is a production-only benchmark. The second isolates signed final-demand reallocation. The third is the paper's composite output-response object. Comparing them shows whether observed output responses are driven by direct input-output exposure, displacement or amplification inside final demand, or production propagation of the signed final-demand movement.

When the outcome is total sales or output rather than final-use quantity, the regression must include production-network controls or use the model-based decomposition. A positive sales spillover can reflect final-use complementarity, downstream derived demand, or both. The cleanest sign tests use final-use quantities or events for which the demand-shifter component is separated from supply, capacity, and inventory shocks.

The current sign-disagreement diagnostic asks whether direct IO and the separated signed-demand system assign the same sign to predicted spillovers. Direct IO propagation is nonnegative in these exercises, while signed final-demand propagation can be negative because substitution can dominate complementarity. At the preferred calibration $r = 0.12$, signed final demand disagrees with direct IO for 37.6 percent of broad final-use sales responses. For a more product-facing app/platform complementarity exposure, the disagreement rises to 63.8 percent. Table 5 summarizes the two main demand-side diagnostics. The comparison is useful because the objects answer different questions: direct IO asks how output needs propagate after demand has moved, while signed demand asks which final-use quantities move in the first place. The composite object $\boldsymbol{R}^{M}\boldsymbol{H}^{F}$ keeps both pieces.

Table 5: Demand-facing signed-layer diagnostics

Shock or exposure	Signed vs. direct IO sign disagreement	Signed vs. collapsed sign change	Interpretation
Broad final-use sales	0.376	0.093	Signed demand changes broad spillover signs
App/platform C exposure	0.638	0.078	Signed demand matters more for ecosystem/product exposure

Notes: Values use the preferred $r = 0.12$ calibration and the 2022 layer system. The collapsed benchmark codes complements as rivalry; direct IO is nonnegative production propagation. Detailed C-definition and placebo-C diagnostics are reported in the full technical appendix.

These sign diagnostics are not driven by the C pairs that also load on rivalry or input-output links. Rebuilding the signed operator on the roughly 85 percent of C pairs that are neither TNIC nor top-IO leaves the broad final-use and app/platform disagreement shares essentially unchanged.

As an external sign check, complement-linked peers earn positive short-window abnormal returns around positive focal-firm adoption, fleet-order, and charging-access news. The preferred $[-1, +1]$ specification gives a complement coefficient of 0.0054 with a wild-cluster $p = 0.010$, while randomization inference is suggestive rather than decisive ($p = 0.117$); the $[-2, +2]$ window gives 0.0070 ($p = 0.030$ wild, $p = 0.066$ randomization). I therefore treat the exercise as sign validation, not as a powered causal event study.

7.8 Robustness, Diagnostics, and Measured Exposure Vectors

The remaining quantitative objects are diagnostics rather than main calibration targets. C-definition robustness varies the deterministic acceptance rule while holding demand scale, production objects, and pass-through fixed. Collapsed-proximity robustness incorrectly codes complements as rivalry; in the preferred 2022 calibration, this wrong coding changes the signed demand sign for 9.3 percent of firms under the broad final-use sales shock. Placebo-C exercises show that sparsity and degree sequences alone are not enough to validate C. The full technical appendix reports these detailed robustness, positive-control, and residual-C diagnostics.

The complementarity layer also has a direct demand-layer welfare interpretation. Under the baseline max-row normalization, the direct complementarity value relative to the direct substitution cost is 0.456 at $r = 0.10$, 0.547 at the preferred $r = 0.12$, and 1.368 at $r = 0.30$. I report these as direct final-use network implications, not as full GE welfare counterfactuals, because total welfare also depends on price, income, markup pass-through, and production incidence. The maintained range satisfies the local regularity screen: the minimum eigenvalue of $I + \alpha_S \widehat{\Sigma}^S - \alpha_C \widehat{C}$ is above 0.996 under the max-row normalization and above 0.964 under the spectral normalization. The full technical appendix collects this screen, the production and pass-through spectral radii, and the finite-markup conditions in a regularity dashboard.

Finally, the same incidence operator can receive externally measured exposure vectors. This is the case needed for the AI application, where AISC records measure downstream input-use exposure rather than a structural final-demand primitive.

Definition 3 (Measured exposure incidence)**.** For an externally measured firm exposure vector x, define

$$\begin{aligned} \mathcal{I}^{direct}(x) &= x, & \mathcal{I}^{IO}(x) &= \boldsymbol{R}^M x, \\ \mathcal{I}^{full}(x) &= \boldsymbol{R}^M \boldsymbol{H}^F x, & \mathcal{I}^{SD}(x) &= \boldsymbol{R}^M (\boldsymbol{H}^F - \boldsymbol{I}) x. \end{aligned} \tag{16}$$

When $x = \mathrm{d}b$, $\mathcal{I}^{full}(x)$ is the fixed-price structural response in Proposition 2. When x is an externally measured activity vector, such as AISC AI input use, the same object is an incidence index. I report net incidence $1'z$, gross incidence $1'|z|$, outside-direct incidence $\sum_{i:x_i=0} |z_i| / \sum_i |z_i|$, concentration, and comparisons with one-layer benchmarks.

Definition 3 is appropriate when the measured vector is treated as a firm-level activity or exposure index. When the measured object is explicitly customer-demand or can be split by outlet,

Corollary 1 supplies the entry rule: non-intermediate components pass through $\boldsymbol{H}^F$, while intermediate or customer-demand components enter before $\boldsymbol{R}^M$ without first being treated as final-use demand.

8 AI Exposure: Input Use, Customer Pull, and Final-Use Demand

AI exposure is not a single shock. AISC strict records measure downstream input use. Customer links measure the demand exposure of suppliers to AI-adopting customers. BEA final-use columns measure non-intermediate final-use demand. These objects enter different operators. The input-use vector is a measured incidence index; customer pull enters goods demand through the customer/intermediate branch; final-use AI demand enters the signed final-demand system before production propagation.

This section therefore separates three exposure objects rather than forcing one AI shock into one operator. Section 8.1 reports the completed AISC input-use incidence; Sections 8.2 and 8.3 define the customer-pull and BEA final-use demand objects as measurement boundaries; Section 8.4 collects the three vectors in one comparison; Section 8.5 locates incidence in markup-margin space; and Section 8.6 reports located GE illustrations at the AISC set. The demand-facing rows are object definitions, included so the quantified AISC input-use exercise is not misread as an aggregate AI demand estimate.

8.1 Direct Downstream AI Input-Use Incidence

The AISC vector measures downstream input use, not final-use AI demand. In this application $\boldsymbol{x}^{AI,input}$ is an externally measured downstream input-use vector, not a final-use demand shifter, so I read $\boldsymbol{R}^M\boldsymbol{H}^F\boldsymbol{x}$ as a normalized incidence and attribution index rather than a structural output response. This boundary makes the positive object clear: direct AI records, IO-only propagation, and product-market proximity each answer one part of the incidence question; the multilayer object answers how those parts compose.

The shock comes from AISC strict input relationships in 2024 (Cen et al., 2025). I direct an AISC strict-input edge from the AI provider to the downstream public firm and assign the shock to the downstream user. The baseline weight is AISC relationship-record multiplicity: each dated strict-input record contributes one unit to the downstream public firm. I do not weight records by source type. The baseline vector is therefore $x_i^{AI,input}$, a downstream AI input-use vector measured by strict-input record counts.

The direct, production-only, signed-demand-plus-production, and incremental signed-demand objects are

$$\begin{aligned} \mathcal{I}^{AI}_{direct} &= \boldsymbol{x}^{AI,input}, & \mathcal{I}^{AI}_{IO} &= \boldsymbol{R}^M\boldsymbol{x}^{AI,input}, \\ \mathcal{I}^{AI}_{SD+IO} &= \boldsymbol{R}^M\boldsymbol{H}^F\boldsymbol{x}^{AI,input}, & \mathcal{I}^{AI}_{SD} &= \boldsymbol{R}^M(\boldsymbol{H}^F-\boldsymbol{I})\boldsymbol{x}^{AI,input}. \end{aligned} \tag{17}$$

If $\boldsymbol{x}^{AI,input}$ were a final-demand shifter, $\mathcal{I}^{AI}_{SD+IO}$ would coincide with the fixed-price structural response. In the AISC application, it is an externally measured downstream input-use exposure vector, so I interpret (17) as incidence indices rather than structural output responses.

The main timing design uses the 2024 shock with the 2023 pre-network: official HP/TNIC-derived substitution, the 2023 product-catalog complementarity layer, and the BEA production network. I use the observed 2024 network only as robustness because its substitution layer is a project-built HP-style Item 1 proxy, not official TNIC.

Table 6: AISC strict AI input use and pre-network timing

Panel A. Shock series

Year	Main markup rows	Strict-input exposed	AISC records	Role
2021	2,919	19	37	Sparse placebo / pre-period
2022	2,878	51	145	Pre-period
2023	2,762	97	446	Post-ChatGPT expansion
2024	2,609	130	743	Main shock year

Panel B. Baseline propagation timing

Shock	Network used	S source	C source	Paper role
2024 AISC strict-input multiplicity	2023 pre-network	Official HP/TNIC-derived	2023 product catalog	Main incidence design
2024 AISC strict-input multiplicity	2024 observed network	Project-built HP-style Item 1 proxy	2024 product catalog	Observed-network robustness

Notes: AISC strict input links are directed from AI providers to downstream public firms. Panel A reports the main US markup sample; all-row counts are larger because the raw Compustat roster includes non-US and non-markup rows. The shock is assigned to the downstream public firm and is not constructed from the 10-K LLM role classifier. Propagation tables below use network-year rosters after alignment: the 2023 pre-network main design has 127 positive exposed roster firms and 750 record-count mass, while the 2024 observed-network robustness has 151 positive exposed roster firms because the network roster is not the same denominator as Panel A.

The direct vector is not a one-off anecdotal exposure measure. Within the main markup sample, strict-input records rise from 37 in 2021 to 743 in 2024, while directly exposed downstream firms rise from 19 to 130. The record margin therefore expands by a factor of 20.1 and the exposed-firm margin by a factor of 6.8. The intensive margin also rises: records per exposed firm increase from 1.9 in 2021 to 5.7 in 2024. I therefore use 2024 as the main shock year and the 2023 pre-network as the baseline propagation network.

Direct-only undercount and outside-direct incidence. Table 7 reports the normalized incidence objects. The main scale-free object is incidence reach: how much gross exposure is generated per unit of direct record mass, and how much of that exposure reaches firms with no direct AISC record. In the 2023 pre-network design, one unit of direct AI input use becomes 2.38 units of gross firm-level incidence, and 66.7 percent of that gross incidence is received by firms with no direct

AISC strict input-use record. Direct exposure alone therefore captures only about 42 percent of SD+IO gross incidence. The 2024 observed-network robustness gives a similar gross incidence of 2.44 and an outside-exposed share of 66.8 percent.

Table 7: Normalized AI incidence decomposition

Shock	Network used	Object	Net sum	Gross sum	Outside-exposed share
2024 AISC	2023 pre-network	x	1.000	1.000	0.000
2024 AISC	2023 pre-network	$H^F x$	1.000	1.003	0.003
2024 AISC	2023 pre-network	$R^M x$	2.378	2.378	0.667
2024 AISC	2023 pre-network	$R^M H^F x$	2.377	2.377	0.667
2024 AISC	2023 pre-network	$R^M(H^F - I)x$	-0.001	0.005	0.874
2024 AISC	2024 observed	x	1.000	1.000	0.000
2024 AISC	2024 observed	$H^F x$	0.999	1.004	0.004
2024 AISC	2024 observed	$R^M x$	2.443	2.443	0.668
2024 AISC	2024 observed	$R^M H^F x$	2.438	2.438	0.668
2024 AISC	2024 observed	$R^M(H^F - I)x$	-0.006	0.006	0.876

Notes: All rows use the 2024 AISC strict input-use multiplicity shock; the 2023 and 2024 labels refer to the network layers used for propagation. Columns report normalized model incidence per unit of direct AI input use. "Outside-exposed" firms have no direct AISC strict input-use record. $R^M x$ is the IO-only benchmark, and $R^M(H^F - I)x$ is the incremental signed-demand correction beyond IO-only propagation.

The table also clarifies why one-layer objects are insufficient. Direct AISC records identify AI-using firms but miss outside-direct incidence. IO-only propagation nearly matches gross reach in this input-use exercise, but gross reach is not attribution. IO-only propagation cannot say whether a non-user is reached as a supplier, a complement, a substitute, or a mixed-path receiver. For this downstream input-use vector, signed demand primarily attributes the IO-reached firms rather than changing the aggregate multiplier. Its role is to supply the sign, path, and markup-margin interpretation around the IO benchmark.

Supporting reach-concentration and benchmark-error diagnostics show that the IO-only ranking is almost identical to the full ranking for this downstream input-use vector. I read that as a feature of this production-heavy application and its small signed-demand correction, not as a general equivalence between IO-only and signed-demand-plus-production propagation.

The gross multiplier is a sector-composition fact. The 2.38 gross-incidence number is benchmarked against placebo vectors with the same number of directly exposed firms and the same record mass. A sector-matched placebo almost recovers the AISC multiplier, so the gross multiplier is largely a sector-composition fact for this input-use vector rather than evidence that the specific AI users have unusually high within-sector centrality; the outside-direct share is lower for the actual vector than for matched placebos because the direct AI users are themselves large receivers of the incidence. In leaveout diagnostics, dropping the top AISC firm leaves gross incidence at 2.40, while dropping the top five firms or the Mag7 raises the outside-direct share above

85 percent.

Why the input-use vector is production-heavy. The first-order decomposition makes the same point numerically:

$$R^M H^F x \simeq R^M x - R^M \alpha_S \Sigma^S x + R^M \alpha_C C x + \text{higher-order loops.}$$

In the AISC input-use exercise, production propagation accounts for most gross incidence. The signed-demand correction is small in the aggregate because substitution drag and complementarity boosts partly offset, and because this shock is assigned to downstream input users rather than to AI product/platform providers. The framework separates the production channel from final-use attribution and shows which channel matters for this measured vector.

The gross signed-demand correction is 0.005 in the 2023 pre-network design, about 0.2 percent of SD+IO gross incidence; in the 2024 robustness design it is 0.006, about 0.3 percent. I therefore do not interpret the AISC input-use vector as evidence of large aggregate signed-demand amplification. The small correction is not evidence that AI has no demand-side effects; it shows that the strict AISC downstream input-use vector is not a final-use AI demand vector, which is exactly why the customer-pull and BEA final-use objects below are needed. The result is sharper: AI input-use incidence is large and indirect because it propagates through production, while signed demand tells us how the IO benchmark is economically attributed.

The outside-direct share is not a residual category. The framework classifies non-users by the path through which they are reached. Some are reached only because directly exposed AI users require intermediate inputs; others are reached because they are final-use complements or substitutes of directly exposed users; and many are mixed-path receivers, reached through both production and signed final-use channels. A direct-record design cannot construct this taxonomy, and an IO-only design cannot sign the final-use part of it.

Robustness and disclosure scope. The incidence reach and outside-direct share are stable across binary and log record weightings, so the headline numbers are not an artifact of record-count multiplicity. A separate 10-K audit shows that AI-related firm roles are broader than the strict AISC input-use vector: the audit confirms that AISC is clean but conservative, and it remains a disclosure audit rather than a separately validated exposure vector.

8.2 Customer-Pull Supplier-Demand Exposure

Customer pull is demand-side, but it is not automatically final-use demand. The direct AISC vector answers who uses AI inputs; a distinct demand-side object asks who sells to AI-adopting customers. Let B_{ij}^{pre} denote the pre-AI share of supplier i's sales accounted for by customer j:

$$B_{ij}^{pre} = \frac{Sales_{ij}^{pre}}{\sum_k Sales_{ik}^{pre}}. \tag{18}$$

Define the resulting customer-pull exposure as

$$x_i^{AI,cp} = \sum_j B_{ij}^{pre} x_j^{AI,input}. \tag{19}$$

AI adoption is measured at the customer side, but demand-pull exposure is assigned to the pre-existing suppliers. This separates the origin of AI adoption from the receiver of AI-related demand.

Because customer pull is demand-side but not automatically final-use demand, I split it by the BEA outlet share:

$$x_i^{AI,cp,F} = \varphi_{s(i)}^{FU,pre} x_i^{AI,cp}, \qquad x_i^{AI,cp,M} = \left(1 - \varphi_{s(i)}^{FU,pre}\right) x_i^{AI,cp}. \tag{20}$$

The corresponding incidence object follows the outlet discipline of Corollary 1:

$$\mathrm{d}\boldsymbol{y}^{AI,cp} = \boldsymbol{R}^M \left(\boldsymbol{H}^F \boldsymbol{x}^{AI,cp,F} + \boldsymbol{x}^{AI,cp,M}\right). \tag{21}$$

The final-use branch enters the signed final-demand system; the intermediate branch enters goods demand directly and then propagates through production links. A GPU illustrates the outlet distinction: the same physical product may be bought by households for gaming, by firms as data-center capital equipment, or as cloud compute purchased as a current production input. Without a clean outlet split, the customer-pull vector is the customer/intermediate-demand object $\boldsymbol{R}^M \boldsymbol{x}^{AI,cp}$ rather than a final-use demand shifter.

The clean implementation requires pre-period firm-customer links or segment major-customer links. I do not substitute a sector-level IO proxy for B_{ij}^{pre}, because that proxy is mechanically close to the first-step Leontief object $\boldsymbol{M}\boldsymbol{x}^{AI,input}$. The direct input-use exercise should therefore be read as production-heavy measured incidence, not as evidence that AI has no demand-side component.

8.3 BEA Final-Use AI Demand Exposure

The customer-pull object is supplier demand; a third object is the non-intermediate final-use AI demand that belongs inside $\boldsymbol{H}^F$. It is built from the BEA final-use columns and the AI-relevant sectors, not by attaching a final-demand shock to the AISC input-user set. For an AI-relevant sector s, let $g_s^{F,2023\to2024}$ denote the growth of its BEA final-use column over the AI-boom window, and define the sectoral final-use AI demand shifter

$$x_s^{AI,F} = AIRelevant_s \cdot g_s^{F,2023\to2024}, \tag{22}$$

allocated to firms by within-sector final-use weights $\lambda_{i,s(i)}$,

$$x_i^{AI,F} = \lambda_{i,s(i)}\, x_{s(i)}^{AI,F}. \tag{23}$$

Because this object is a non-intermediate final-use demand shifter, it enters the signed final-demand system before production propagation, exactly as in Corollary 1:

$$\mathrm{d}\boldsymbol{y}^{AI,F} = \boldsymbol{R}^M \boldsymbol{H}^F \boldsymbol{x}^{AI,F}. \tag{24}$$

This is the object for which signed demand does first-order work: a positive final-use AI demand shifter displaces the substitutes of the AI-relevant sellers and pulls up their complements before $\boldsymbol{R}^M$ propagates the movement.

Two boundaries keep this object distinct. First, $\boldsymbol{x}^{AI,F}$ originates in BEA final-use columns and AI-relevant sectors, so it is not the AISC input-user set with a stipulated final-demand shock attached. Second, the located demand pull of Section 8.6 is a controlled counterfactual at the AISC set, not this measured final-use object; the two are reported separately and are not added together.

8.4 Comparing the Three AI Exposure Vectors

Table 8 collects the three exposure objects in one place. The input-use row is quantified from the incidence decomposition above; the customer-pull and BEA final-use rows report the corresponding demand-side entry rules and empirical anchors. The point of the table is not to force a large signed-demand multiplier on every AI vector, but to record which operator and markup margin each object loads on.

Table 8: Three AI exposure vectors: object, operator, and where incidence lands

Vector	Economic object	Origin → receiver	Model entry	Empirical content here
$\boldsymbol{x}^{AI,input}$	Downstream input use	AISC user → same user	Incidence index $\boldsymbol{R}^M\boldsymbol{H}^F\boldsymbol{x}$ (signed corr. 0.005)	Quantified: gross incidence 2.38, outside-direct share 0.667
$\boldsymbol{x}^{AI,cp}$	Customer-pull supplier demand	AI customer → supplier	$\boldsymbol{R}^M\boldsymbol{x}$ or outlet split $\boldsymbol{R}^M(\boldsymbol{H}^F\boldsymbol{x}^F + \boldsymbol{x}^M)$	Formal object requires pre-period firm-customer links; 2023 SEC/procurement seller anchors below
$\boldsymbol{x}^{AI,F}$	BEA final-use AI demand	Final-use sector → seller	$\boldsymbol{R}^M\boldsymbol{H}^F\boldsymbol{x}$	BEA final-use sector proxy and 2023 seller-side micro evidence below

Notes: The input-use row reproduces the 2023 pre-network numbers of Table 7: gross SD+IO incidence per unit of direct exposure, outside-direct share, and (in parentheses) the gross signed-demand correction $\boldsymbol{R}^M(\boldsymbol{H}^F - \boldsymbol{I})\boldsymbol{x}$. The customer-pull and BEA final-use rows are demand-facing measurement objects. The formal customer-pull object still requires pre-period firm-customer sales links; the micro evidence below instead uses observed 2023 seller-side AI demand records as location anchors. The table classifies how alternative AI exposure vectors enter the model; it is not a causal estimate of aggregate AI demand.

The 2023 SEC/procurement exercise gives a narrow measured demand-side anchor for where AI demand is sold, not a complete outlet split. I collect SEC exhibit windows classified as AI-

related final-demand or procurement evidence in which the public firm is the seller, and federal AI procurement awards matched to Compustat vendors. I normalize each seller-side channel to unit mass and pass it through the same ordered operator used for final-demand incidence, $\boldsymbol{R}^M\boldsymbol{H}^F\boldsymbol{x}$, on the 2023 network. Table 9 shows the result. The SEC channel is the broader micro anchor, with 329 positive seller evidence rows mapped to 139 public firms; federal procurement is deliberately narrower, with three matched 2023 vendors. In both cases, production propagation carries most gross reach after the demand location is set: the combined SEC-plus-procurement vector generates 3.30 units of gross SD+IO incidence per unit of direct seller-side demand mass, with 95.5 percent landing outside the directly anchored sellers. This "gross incidence" is an absolute receiver-mass statistic: the direct seller-side shock is normalized to one, and the table sums the absolute firm-level responses after signed demand and production propagation. It is not an aggregate output or GDP multiplier.

Table 9: 2023 seller-side AI demand evidence from SEC contracts and federal procurement

Micro demand evidence	Object	Active firms	Abs. mass	Outside-direct	Top-10
SEC seller contracts	x	139	1.000	0.0	21.8
SEC seller contracts	$H^F x$	3261	1.003	0.3	21.7
SEC seller contracts	$R^M x$	3811	2.434	94.9	19.2
SEC seller contracts	$R^M H^F x$	3811	2.419	94.9	19.2
SEC seller contracts	$R^M(H^F - I)x$	3811	0.014	98.4	16.2
Federal procurement, amount	x	3	1.000	0.0	100.0
Federal procurement, amount	$H^F x$	1058	1.000	0.0	100.0
Federal procurement, amount	$R^M x$	3811	4.182	99.0	16.1
Federal procurement, amount	$R^M H^F x$	3811	4.181	99.0	16.1
Federal procurement, amount	$R^M(H^F - I)x$	3811	0.001	98.7	22.3
SEC + procurement, equal channel	x	142	1.000	0.0	59.1
SEC + procurement, equal channel	$H^F x$	3199	1.001	0.2	59.0
SEC + procurement, equal channel	$R^M x$	3811	3.308	95.5	14.5
SEC + procurement, equal channel	$R^M H^F x$	3811	3.300	95.5	14.6
SEC + procurement, equal channel	$R^M(H^F - I)x$	3811	0.007	98.2	15.7

Notes: The direct vector x is normalized to unit positive mass within each channel. SEC seller contracts count positive seller-side evidence rows; federal procurement weights matched vendors by award amount; the combined row gives equal channel weight to SEC and procurement. $H^F x$ is signed demand propagation, $R^M x$ is production propagation, $R^M H^F x$ is ordered demand-then-production incidence, and $R^M(H^F - I)x$ is the signed-demand correction after production propagation. Abs. mass is $\sum_i |z_i|$ for each reported incidence vector z; it can exceed one because direct normalized seller-side demand is redistributed through signed demand and production links. Outside-direct is the share of absolute incidence outside the directly anchored sellers; Top-10 is the top-ten absolute-incidence share.

8.5 AI Incidence in Markup-Margin Space

The markup-space decomposition is central because incidence and market power are different objects. Direct AI records identify users, and $R^M H^F x$ identifies receivers, but Theorem 1 identifies which pricing margin characterizes those receivers. The question is not whether AI has changed markups; it is whether observed AI input-use incidence lands on firms whose prices are mainly disciplined by final-use demand, intermediate demand, or cost feedback. The same template applies to each of the three exposure vectors above—input use, customer pull, and final-use demand—and asks which markup margin receives the incidence, the question that separates this exercise from a generic AI-spillover regression.

Sorting firms by predicted SD+IO AI incidence shows that exposure does not simply land on the highest-markup firms. In the 2023 design, the top SD+IO-incidence quintile receives 85.8 percent of SD+IO gross incidence, but its sales-weighted markup is 1.61, below the lowest-incidence bin. The high-incidence bin is instead relatively intermediate-heavy in the one-price attribution. This is exactly the distinction between who is exposed and which margin constrains the exposed firm's markup.

This table is a structural location exercise, not an independent validation test. The incidence object and the markup attribution both use $\boldsymbol{H}^F$, while the incidence ranking also depends on $\boldsymbol{R}^M$ and the AISC shock vector. The shared operator is the reason the table is useful for attribution, and also the reason I do not interpret the association as a causal relationship between AI exposure and markups.

Table 10: AI incidence in markup space

Network	Bin	Direct-exposed	Incidence mass	Markup	Final-use pressure	Intermed. pressure	Cost-fb. pressure
2023	Q1	0.008	0.001	2.559	0.647	0.588	-0.235
2023	Q2	0.008	0.007	1.844	0.533	0.629	-0.163
2023	Q3	0.034	0.030	1.527	0.424	0.705	-0.129
2023	Q4	0.072	0.090	1.556	0.407	0.727	-0.135
2023	Q5	0.182	0.858	1.609	0.434	0.702	-0.135
2024	Q1	0.003	0.001	3.189	0.636	0.612	-0.248
2024	Q2	0.016	0.005	2.176	0.579	0.623	-0.202
2024	Q3	0.033	0.026	1.797	0.452	0.702	-0.154
2024	Q4	0.059	0.086	1.604	0.426	0.711	-0.138
2024	Q5	0.186	0.902	1.700	0.429	0.710	-0.138

Notes: Bins are quintiles of predicted SD+IO AI incidence $R^M H^F x$ within the network-year design. Markup and the three pricing-pressure columns are sales-weighted means. The pressure columns are the effective-denominator terms of the one-price Lerner restriction—$\varphi_i E_i^F$, $(1-\varphi_i)E_i^M$, and $-\eta_i^\kappa$, each divided by the denominator $\varphi_i E_i^F + (1-\varphi_i)E_i^M - \eta_i^\kappa$—under the literature pass-through anchor $\rho_{MC} = 0.70$, the same convention as Table 4; the three sum to one. A larger final-use pressure means the bin's markup is more disciplined by final-use demand; the negative cost-feedback column shrinks toward zero in the high-incidence bins. Incidence mass is the bin's share of SD+IO gross incidence. The table describes where incidence lands in markup space; it is not evidence that AI changed markups.

The same incidence can be located by receiver class rather than by incidence quintile. Crossing the five-way receiver taxonomy with the cost-feedback markup decomposition shows that mixed signed-path non-users receive about half of gross incidence and are intermediate-heavy in the pricing attribution, the production-only non-users are also intermediate-leaning, and the smaller substitute- and complement-linked classes sit at different points of the final-use versus intermediate margin split. This is the sense in which direct records identify users, the incidence operator identifies receivers, and the markup theorem identifies the pricing margin of those receivers. The markup columns cover only receivers with a 2022 markup, so the thin substitute- and complement-linked cells should be read as indicative.

8.6 Located GE Illustrations at the AISC Set

These exercises place stipulated shocks at the AISC exposure set and read where they would land in the calibrated local equilibrium of Section 6. They are useful for reading the GE machinery, but they are distinct from the measured input-use, customer-pull, and BEA final-use AI vectors above: they are located counterfactuals at the AISC set, not measured AI demand, AI productivity, or a general-equilibrium estimate of AI's aggregate effect. I report three objects, all holding the numeraire wage and constant markups fixed on the calibrated networks; the shock sizes are normalizations, so the objects are read for incidence shape, not for the level magnitude of any actual AI rollout.

A located cost reduction. Treat AI adoption as a one percent unit cost reduction at the exposed users and let it travel through the production network. The welfare gain is modest in level, $dW/Y = +0.196$ percent for a one percent cost cut on this narrow set, about twice the $+0.087$ percent of a sector-matched placebo. About three quarters of the gross-output incidence lands on firms with no direct AISC record (outside-direct share 0.77); that share is *below* the placebo's 0.91 because the exposed AI users are themselves large receivers of the incidence rather than peripheral to it. For this cost shock the incidence is governed by the production network: shutting the demand network ($B^F = I$) leaves both the welfare number and the cross-section essentially unchanged.

The input requirement and break-even productivity. Adoption is not free: the user buys AI inputs worth a share χ of its cost from marked-up providers. The price-level break-even is clean—the user passes the full marked-up input price into its own cost, so the final-use price level is unchanged exactly when the productivity gain equals the input requirement, $\gamma_P^* = \chi$. The welfare break-even lies below it, because in the closed economy the provider markup is a transfer back to the household, so the input's resource cost is below the price the user pays. The schedule gives $\gamma_W^* \approx 0.6\text{–}0.8\,\chi$, so adoption that merely holds the price level flat already raises welfare. The

schedule is illustrative over the assumed input requirement χ; the robust object is the ordering $\gamma_W^* < \gamma_P^* = \chi$, not the level of AI productivity.

A located demand pull. Place a one percent final-demand pull at the same firms instead of a cost reduction. The operative operator is now the signed demand system $H^F = (B^F)^{-1}$ rather than the production network. Because B^F is a small, sparse perturbation of the identity, the demand network moves only about 0.79 percent of the direct pull onto non-exposed firms in the preferred 2023-network specification; but it sets the *sign* of that spillover, displacing the substitutes of the pulled firms and pulling up their complements, with the complementarity share rising in r. The contrast with the cost-shock incidence is the division of labor the paper has emphasized throughout: the production network carries the magnitude of incidence, while the demand network sets its sign and composition. This located demand pull is a controlled counterfactual on the AISC set; it is distinct from the downstream customer-pull object of Section 8.2 and the BEA final-use object of Section 8.3, as well as the seller-side SEC/procurement evidence in Table 9.

9 Conclusion

The paper develops a multilayer attribution framework for markups and spillovers in demand and production networks. It separates final-use substitution, final-use complementarity, and production dependence, then composes them into markup and incidence objects. Final-use demand disciplines markups only through the outlet share; downstream intermediate demand and upstream cost feedback enter the same one-price Lerner restriction. A final-demand shock first passes through the signed demand system and only then through the production network, so production-only exposure is an incomplete economic attribution even when it captures much of the gross reach.

The AI input-use application makes this attribution problem visible in a measured input-use vector. Direct AISC records identify downstream AI use, not the full incidence of AI input use. In the 2023 pre-network design, the 2024 AISC strict input-use vector generates 2.38 units of gross firm-level incidence per unit of direct AI input use, and 66.7 percent of that incidence is received by firms with no direct AISC input-use record. IO links explain most of the gross reach. The signed final-use layers explain how the IO benchmark is decomposed into substitution drag, complementarity boost, and higher-order residual terms. The result is not that AI has no final-demand component; it is that observed strict input-use exposure is production-heavy.

The quantitative results make those distinctions concrete. The 2022 complementarity layer is sparse relative to product-market rivalry, so it is not just a low-threshold similarity matrix. Under pass-through discipline, final-use demand accounts for a meaningful share of markup-denominator pressure, intermediate-demand discipline remains large, and cost feedback lowers the no-cost-feedback Lerner. In the demand-facing spillover exercises, measuring signed final-

demand responses before production propagation changes signs for economically large shares of firms. In the AI application, direct records, IO-only propagation, and product-market proximity each answer only one part of the same question, and the framework specifies how demand-side seller evidence from SEC contracts, procurement records, or pre-period firm-customer links adds another object. The paper's contribution is to organize those parts into measurable layers and to compose them in the order required by the economics.

The AI application separates three exposure objects rather than forcing one AI shock into one operator. Direct AISC records identify downstream input use and generate broad, production-dominant incidence. Customer-pull exposure assigns AI adopters' demand to their pre-existing suppliers and therefore belongs to the customer/intermediate-demand margin. BEA final-use exposure isolates the non-intermediate demand margin that belongs in the signed final-use system. The framework's contribution is not that signed demand is large for every AI vector; it is that the framework tells which AI exposure is input use, customer demand, or final-use demand, and which markup margin receives the incidence. Removing any one layer changes the estimand: direct records alone answer who uses AI inputs, not who is incident to them; IO alone answers reach, not final-use sign or markup margin; and product-market proximity alone answers closeness, not whether closeness is rivalry, joint use, or production dependence.

References

Amiti, M., Itskhoki, O., and Konings, J. (2019). International shocks, variable markups, and domestic prices. *Review of Economic Studies*, 86(6):2356–2402.

Atalay, E., Hortaçsu, A., Roberts, J., and Syverson, C. (2011). Network structure of production. *Proceedings of the National Academy of Sciences*, 108(13):5199–5202.

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

Barrot, J.-N. and Sauvagnat, J. (2016). Input specificity and the propagation of idiosyncratic shocks in production networks. *Quarterly Journal of Economics*, 131(3):1543–1592.

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890.

Bigio, S. and La'O, J. (2020). Distortions in production networks. *Quarterly Journal of Economics*, 135(4):2187–2253.

Cen, S. H., Gailmard, L., Bommasani, R., Ho, D. E., and Liang, P. (2025). AI Supply Chain Mapping: An Analysis of the Complex Relationships in the AI Ecosystem. Forthcoming.

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

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

Ganapati, S., Shapiro, J. S., and Walker, R. (2020). Energy cost pass-through in us manufacturing: Estimates and implications for carbon taxes. *American Economic Journal: Applied Economics*, 12(2):303–342.

Gentzkow, M. (2007). Valuing new goods in a model with complementarity: Online newspapers. *American Economic Review*, 97(3):713–744.

Hall, R. E. (1988). The relation between price and marginal cost in u.s. industry. *Journal of Political Economy*, 96(5):921–947.

Hoberg, G. and Phillips, G. (2016). Text-based network industries and endogenous product differentiation. *Journal of Political Economy*, 124(5):1423–1465.

Huremović, K., Jiménez, G., Moral-Benito, E., Peydró, J.-L., and Vega-Redondo, F. (2026). Production and financial networks in interplay. *American Economic Review*, 116(5):1611–1647.

Nevo, A. (2001). Measuring market power in the ready-to-eat cereal industry. *Econometrica*, 69(2):307–342.

Pellegrino, B. (2025). Product differentiation and oligopoly: A network approach. *American Economic Review*, 115(4):1170–1225.

U.S. Bureau of Labor Statistics (2026). Producer price indexes: Satellite inputs to industry indexes. https://www.bls.gov/ppi/input-indexes/. Accessed May 27, 2026.

Wu, L. (2025). Network lerner index: Demand and distortions across industries. Conditionally accepted, Review of Economic Studies.