

When Dense Networks Diffuse Sparsely: A Stochastic Propagation Operator for Weighted Directed Networks

Pablo Reyes Granados

`pablo.reyes4@est.uexternado.edu.co`

Universidad Externado de Colombia <https://orcid.org/0009-0002-2906-4182>

German Combariza

National University of Colombia

Research Article

Keywords: Weighted directed networks, Stochastic diffusion, Resolvent centrality, Sparse activation, Production networks, Input-output networks

Posted Date: October 16th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11382078/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

When Dense Networks Diffuse Sparsely: A Stochastic Propagation Operator for Weighted Directed Networks

Pablo Reyes¹ and Germán Combariza²

¹Universidad Externado de Colombia, pablo.reyes4@est.externado.edu.co

²Universidad Nacional de Colombia, gcombariza@unal.edu.co

Abstract

Classical resolvent centralities on weighted networks aggregate propagation over all observed links. We develop a stochastic propagation operator for weighted directed networks in which only a random subset of links is active at each round, impulses may persist, and receiving nodes may face finite capacity. In the linear regime, expected propagation is governed exactly by an effective operator that converges monotonically to the dense resolvent limit and factorizes into emitter-level retention factors, bounded by the inverse effective number of destinations, and a mass-preserving reallocation within each emitter's links. Applied to the 2021 Colombian input–output network, which is topologically dense but concentrated in a few effective trading relationships, sparse activation preserves the broad supply-side ordering while substantially reordering the most central sectors. The reordering is not an attenuation effect: a uniformly attenuated benchmark at the same distance to criticality reproduces 3% of it, whereas retention alone reproduces 58% of the rank changes and 97% of the variation in log scores. Nor is it backbone extraction: dense resolvents on deterministic backbones reproduce at most 29%. The effective operator predicts simulated cascade reach better than Katz, PageRank, eigenvector, communicability, and Viral Centrality computed on network weights, except under strongly binding capacity. Across 79 OECD input–output tables, retention reproduces most of the reordering in every economy (median 63%), attenuation and backbones reproduce little, and the industries promoted by sparse activation are the same across countries. Dense connectivity need not imply dense effective diffusion.

Keywords: weighted directed networks; stochastic diffusion; resolvent centrality; backbone extraction; production networks; sparse activation; input–output networks.

1 Introduction

A central problem in network science is to rank nodes by their ability to transmit influence through a weighted directed graph. Katz–Bonacich centrality, communicability, and related resolvent measures do so by aggregating weighted walks of increasing length (Katz, 1953; Bonacich, 1987; Estrada and Hatano, 2008; Newman, 2010). These approaches provide a natural way to summarize how local connections combine into multistep patterns of influence across a network.

This perspective is particularly relevant for economic production networks, where sectors are connected through flows of intermediate goods and services. An input–output matrix provides a direct representation of this structure: sectors can be viewed as nodes and intermediate sales as weighted directed links. Shocks may then propagate through successive buyer–supplier relationships, with each round of transmission corresponding to an additional step through the production network. Classical input–output multipliers summarize this propagation. The Leontief inverse aggregates the direct and indirect upstream requirements

generated by a demand shock, while the Ghosh inverse captures the corresponding downstream allocation effects of a supply shock (Leontief, 1936; Ghosh, 1958; Miller and Blair, 2009). From a network perspective, both can be interpreted as orientation-specific resolvent measures.

These classical resolvent representations imply a dense propagation mechanism in which every observed productive link contributes at every step, so all weighted paths enter simultaneously into measured influence. Economically, this amounts to treating all observed productive relationships as active transmission channels, despite potentially large heterogeneity in their relevance for actual propagation (Pichler et al., 2022). In weighted production networks, however, a graph may appear dense because sectors are connected to many trading partners while the bulk of their transactions is concentrated in only a small subset of them. The network can therefore be dense in observed connectivity but sparse in economically relevant relationships, so a shock need not propagate through every observed link at each round.

Several strands of the literature have studied these mechanisms separately. Production-network models examine how shocks propagate through input–output linkages (Acemoglu et al., 2012; Carvalho, 2014; Carvalho and Tahbaz-Salehi, 2019), while dynamic input–output models show how production constraints and inventories can shape the propagation process (Pichler et al., 2022). Cascade and influence models study how network structure determines spreading and node importance (Watts, 2002; Kempe et al., 2003; Gleeson et al., 2025), and centrality measures have been designed specifically to approximate cascade reach on weighted directed networks (Fink et al., 2023). Temporal-network models consider settings in which links are not continuously available (Holme and Saramäki, 2012; Perra et al., 2012), and dynamic communicability extends walk-based centrality to evolving networks (Grindrod et al., 2011). A separate literature addresses the gap between topological and weighted density by filtering: backbone-extraction methods such as the disparity filter retain only the links whose weight is statistically significant relative to a uniform null (Serrano et al., 2009; Radicchi et al., 2011), building on measures of weight disparity at the node level (Barthélemy et al., 2005). These ideas have also begun to intersect with production networks: cascade-based diffusion has been applied directly to empirical input–output systems (Alatryste Contreras and Fagiolo, 2014), while firm-level evidence shows that disruptions propagate through both direct and indirect supplier–customer relationships (Barrot and Sauvagnat, 2016; Carvalho et al., 2021).

To bridge this gap, we develop an operator-based extension of dense resolvent propagation to stochastically activated weighted directed networks. The framework allows sectors to transmit through a random subset of outgoing links, reflecting the idea that observed relationships need not all become active transmission channels at each round. Activation depends on transaction weights and productive proximity, while a separate base matrix determines transmission magnitudes conditional on activation. The model also allows part of the impulse to persist across rounds and receiving nodes to face finite absorption capacity, so shocks may propagate gradually and become nonlinear once operational bottlenecks bind.

Replacing random edge states by their expected occupation is closely related to annealed network representations. Our setting also differs from cascade-based centralities for weighted directed networks, such as Viral Centrality (Fink et al., 2023), which approximate the expected reach of an independent-cascade process in which link weights are themselves transmission probabilities. Here, activation probabilities and transmitted magnitudes are distinct objects, impulses are divisible and may persist across rounds, and the expected multiround operator is characterized exactly. Our contribution is to characterize this effective operator and to connect it analytically to its dense full-activation resolvent limit. We further show that, in the linear regime, persistence and attenuation enter infinite-horizon scores only through a single effective attenuation, and that the effective operator factorizes into an emitter-level retention factor and a within-row reallocation of transmission. The retention factor is bounded by the inverse effective number of destinations of each emitter, which links the reordering produced by sparse activation to the concentration of outgoing transactions. This characterization establishes how heterogeneous activation can alter influence rankings even when expected total propagation is attenuated. Finite receiving capacity then provides a separate nonlinear extension that breaks scale invariance once bottlenecks bind. In the input–output application, the resulting

framework provides a supply-side extension of the classical Ghosh benchmark (Ghosh, 1958).

In the linear no-capacity regime, this characterization takes a particularly simple form. Expected propagation is governed exactly by

$$\bar{B} = \Omega(G \odot \Pi)^\top,$$

where G is the base transmission matrix and Π contains link-activation probabilities. Since $0 \leq \bar{B} \leq \Omega G^\top$, the dense full-activation operator is an upper envelope of expected sparse diffusion. Yet heterogeneous activation reweights links non-uniformly, so influence rankings can differ systematically even when expected total propagation is attenuated.

Applying the framework to the Colombian production network, we find that sparse diffusion preserves the broad Ghosh-oriented ordering but substantially reorders the sectors that are most central under dense propagation, while capacity adds further reordering only once bottlenecks bind under stress.

We then identify what drives this reordering. It is not an attenuation effect: a uniformly attenuated operator at the same distance to criticality reproduces almost none of it, whereas the retention channel of the factorization reproduces most of the rank changes and nearly all the variation in log scores. Nor is it sparsification: dense resolvents on deterministic backbones reproduce at most 29%. Activation information also predicts simulated cascades better than centralities computed from network weights alone, outside the strongly capacity-constrained regime. Finally, the mechanism generalizes. In synthetic dense networks, the dispersion of retention factors governs the magnitude of reordering, and across 79 OECD input–output tables retention is the dominant channel in every economy, with the same industries gaining and losing influence across countries.

The remainder of the paper is organized as follows. Section 2 formally defines the stochastic propagation operator. Section 3 establishes the core theoretical results. Section 4 introduces the Colombian input–output data. Section 5 applies the framework empirically, detailing the calibration strategy and documenting shifts in sectoral influence across density and capacity regimes. Section 6 identifies the mechanism behind the reordering, compares it with backbone extraction and standard centralities, and maps it across synthetic networks and 79 OECD input–output tables. Section 7 summarizes the robustness exercises, which are reported in full in Appendix B, and Section 8 concludes.

2 The propagation operator

Throughout this section, uppercase letters denote matrices and lowercase letters scalars or vectors, unless otherwise stated. A tilde denotes row normalization, $\mathbf{1} \in \mathbb{R}^n$ is the vector of ones, and $e_s \in \mathbb{R}^n$ is the s th standard basis vector.

2.1 Primitives

The economy is a weighted directed graph $\mathcal{N} = (V, E)$ with $n = |V|$ sectors. Let $W \in \mathbb{R}_+^{n \times n}$ denote the matrix of intermediate transactions, with W_{ji} the flow of inputs sold by sector j to sector i , and let $x_j > 0$ denote the gross output of sector j . Since the propagation network represents transmission across distinct sectors, define its off-diagonal transaction matrix as

$$W^\circ := W - \text{diag}(W).$$

Thus, diagonal within-sector transactions remain part of the classical input–output benchmarks but are not treated as propagation links. We distinguish the matrix that controls the magnitude transmitted through an active edge from the kernel that controls whether that edge is active.

- **Base transmission matrix** $G \in \mathbb{R}_+^{n \times n}$, with zero diagonal and support contained in E . Conditional on activation, G_{ji} determines the amount of carried impulse transmitted from emitter j to receiver i . The

theory is written for a generic nonnegative G ; in the input–output application,

$$G = B_G^\circ, \quad (B_G)_{ji} = \frac{W_{ji}}{x_j}, \quad B_G^\circ := B_G - \text{diag}(B_G),$$

so the stochastic operator uses the off-diagonal part of the Ghosh allocation matrix. The full B_G is retained separately for the classical Ghosh benchmark.

- **Row-normalized sales weights** $\tilde{\omega} \in \mathbb{R}_+^{n \times n}$, with

$$\tilde{\omega}_{ji} = \frac{(W_{ji}^\circ)^\gamma}{\sum_k (W_{jk}^\circ)^\gamma}, \quad \tilde{\omega}_{jj} = 0.$$

Each positive row is a probability distribution over intersectoral destinations. The curvature $\gamma > 0$ tunes weight concentration; the baseline fixes $\gamma = 1$, so that weights are proportional to observed intersectoral sales. The support of row j is $V_j = \{i : \tilde{\omega}_{ji} > 0\}$. Unlike G , $\tilde{\omega}$ does not set transmitted magnitudes; it summarizes the relative composition and concentration of outgoing sales and enters the activation mechanism.

- **Productive proximity** $\tilde{\phi} \in \mathbb{R}_+^{n \times n}$, an auxiliary relatedness matrix constructed from the intersectoral production network. Let $p_{ji} = W_{ji}^\circ / \sum_k W_{jk}^\circ$ denote the row-normalized sales share of sector j going to sector i . The motivation is that shock transmission need not depend only on the strength of an observed bilateral link: sectors that are more closely related in the production structure may also be more likely to participate in the same propagation process. Following the bilateral proximity logic of the product-space literature (Hidalgo et al., 2007), we define raw bilateral proximity as $\phi_{ji} = \min\{p_{ji}, p_{ij}\}$, set the diagonal to zero, optionally apply a curvature transformation, and row-normalize to obtain $\tilde{\phi}$.
- **Selection kernel** $B = \lambda \tilde{\omega} + (1 - \lambda) \tilde{\phi}$, with $\lambda \in [0, 1]$, set to zero on the diagonal and row-normalized. It governs which intersectoral partners a sector activates. Thus G and B have distinct roles: G determines transmission conditional on activation, whereas B determines link-inclusion probabilities.
- **Global transmission attenuation** $\Omega \in (0, 1]$: a common pass-through factor multiplying the base transmission matrix in each round. It is a common normalization and does not replace the cross-sector allocation shares already embedded in $G = B_G^\circ$.
- **Capacities** $C_i > 0$: the maximum impulse sector i can absorb per round, e.g. $C_i = \kappa \cdot \text{output}_i$.
- **Persistence** $\delta \in [0, 1]$: the share of carried impulse retained across rounds.

We track two state variables: $P_i(t)$, the impulse *absorbed* by sector i in round t , and $S_i(t)$, the impulse *carried* by sector i into round t . The carried state combines retained impulse from the previous round with newly absorbed impulse; its formal update is introduced below.

2.2 Stochastic activation

The economic content of the operator is that, each round, an emitter transmits through only a sparse subset of its links. Activation captures whether an existing relationship effectively participates in the transmission of a shock in a given round, while the underlying network structure remains fixed. We encode activation by a random binary matrix $I(t) \in \{0, 1\}^{n \times n}$, where $I_{ji}(t) = 1$ if the channel $j \rightarrow i$ fires in round t , with $I_{jj}(t) = 0$ and $I_{ji}(t) = 0$ whenever $(j, i) \notin E$. We assume the $\{I(t)\}_{t \geq 1}$ are independent across rounds and independent of the initial shock. All theoretical results below depend on the activation scheme *only* through its inclusion-probability matrix

$$\Pi = \mathbb{E}[I(t)], \quad \pi_{ji} = \mathbb{P}(I_{ji}(t) = 1) \in [0, 1], \quad \pi_{ji} = 0 \text{ if } (j, i) \notin E.$$

For each emitter j , define its effective diversification as

$$n_j^{\text{eff}} = \left(\sum_i \tilde{\omega}_{ji}^2 \right)^{-1},$$

with $n_j^{\text{eff}} = 0$ for an empty row. It is the inverse of the node-level weight disparity $Y_j = \sum_i \tilde{\omega}_{ji}^2$ (Barthélemy et al., 2005). Let

$$\hat{s}_j = \frac{\sum_i W_{ji}^\circ}{\max_k \sum_i W_{ki}^\circ}, \quad \hat{d}_j = \frac{n_j^{\text{eff}}}{\max_k n_k^{\text{eff}}},$$

so that both measures lie in $[0, 1]$. We construct the baseline intensity profile as

$$\tilde{\alpha}_j = \alpha_{\max} [\eta \hat{s}_j + (1 - \eta) \hat{d}_j], \quad \alpha_{\max} \in [0, 1], \quad \eta \in [0, 1],$$

where $\alpha_{\max}$ controls the scale of the baseline profile and η determines the relative weight placed on intersectoral sales scale versus effective diversification. The implemented intensity is $\alpha_j = \alpha_{\text{scale}} \tilde{\alpha}_j$, where $\alpha_{\text{scale}} \geq 0$ is the common activation scale calibrated in Section 5.

For the empirical exercises, Π is generated by a marked-Poisson activation mechanism. Each candidate link activates independently with

$$\pi_{ji} = 1 - e^{-\alpha_j B_{ji}}. \quad (1)$$

Equivalently, a rate- α_j Poisson process of “transmission attempts” is split across destinations in proportions B_{ji} , and a link fires if it receives at least one attempt. This specification yields closed-form inclusion probabilities, makes links within a row mutually independent, and satisfies the monotonicity and limiting properties used below ($\pi_{ji} \uparrow \mathbf{1}(j, i) \in E$ as $\alpha_j \rightarrow \infty$, and $\pi_{ji} \rightarrow 0$ as $\alpha_j \rightarrow 0$). Within-row independence is not required for any first-moment result in Section 3, which uses only independence of $I(t)$ across rounds; it matters only for higher moments (Remark 5).

2.3 The round map and its operator form

Fix a round t . Given the carried state $S(t-1)$ and the realized activation $I(t)$:

(1) *Potential inflows.* Each active emitter j transmits through the links that fire in that round. Conditional on activation, the magnitude on link $j \rightarrow i$ is set by the base transmission matrix G and the common attenuation factor Ω :

$$\hat{\Delta}_{j \rightarrow i}(t) = I_{ji}(t) \Omega S_j(t-1) G_{ji}, \quad \hat{P}_i(t) = \sum_j \hat{\Delta}_{j \rightarrow i}(t).$$

Collecting terms, $\hat{P}(t) = A(t)S(t-1)$ with the random linear operator

$$A(t) = \Omega(G \odot I(t))^\top \in \mathbb{R}_+^{n \times n}, \quad A(t)_{ij} = \Omega G_{ji} I_{ji}(t), \quad (2)$$

where $\odot$ is the Hadamard product. The row-normalized matrices $\tilde{\omega}$ and $\tilde{\phi}$ affect $A(t)$ only indirectly through the selection kernel B and hence through the realization of $I(t)$.

(2) *Capacity rationing.* Potential inflows are scaled down so that no sector absorbs more than its capacity. We allow any *admissible saturation map* $\sigma_C : \mathbb{R}_+^n \rightarrow \mathbb{R}_+^n$, applied componentwise, satisfying

$$\sigma_C(0) = 0, \quad 0 \leq \sigma_C(x)_i \leq \min\{x_i, C_i\}, \quad \sigma_C \text{ is nondecreasing, concave, and 1-Lipschitz.} \quad (3)$$

The two leading instances are the hard cap

$$\sigma_C^{\text{hard}}(x)_i = \min\{x_i, C_i\},$$

and the smooth saturation

$$\sigma_C^{\text{soft}}(x)_i = C_i \left(1 - e^{-x_i/C_i}\right).$$

Absorbed impulse is $P(t) = \sigma_C(\hat{P}(t))$.

(3) *State update.*

$$S(t) = \delta S(t-1) + P(t).$$

The dynamics are therefore the random nonlinear recursion

$$S(t) = \delta S(t-1) + \sigma_C(A(t)S(t-1)), \quad S(0) = s_0 \geq 0. \quad (4)$$

In the no-capacity case $C_i \rightarrow \infty$ (so $\sigma_C = \text{id}$), this collapses to the random *linear* recursion

$$S(t) = (\delta I + A(t))S(t-1).$$

2.4 Sectoral importance

Definition 1 (Importance score). Given a unit source at sector s , so that $S(0) = e_s$, define the realized cumulative absorption as

$$T_s = 1 + \sum_{t \geq 1} \mathbf{1}^\top P(t),$$

where the leading one represents the initial unit impulse. The *importance score* of sector s is

$$m_s = \mathbb{E}[T_s],$$

and sectors are ranked in decreasing order of m_s . For stochastic specifications, m_s is estimated by averaging T_s across Monte Carlo realizations.

Definition 1 provides the common ranking criterion used throughout Section 5. Its connection to the classical Leontief and Ghosh multipliers is established in Proposition 1 and Corollary 1.

3 Theoretical results

Throughout, inequalities between matrices or vectors are entrywise, and a matrix is *nonnegative* if all entries are. For a nonnegative base transmission matrix $G \in \mathbb{R}_+^{n \times n}$ with zero diagonal, write the associated dense full-activation operator as $A_G^\star := \Omega G^\top$, and let $A^\star := A_G^\star$ when the base matrix is understood. The empirical propagation operator uses $G = B_G^\circ$, the off-diagonal Ghosh allocation matrix, whereas $\tilde{\omega}$ and $\tilde{\phi}$ enter only through the activation-selection kernel. The full matrices A and B_G are retained separately as the classical Leontief and Ghosh benchmarks. Let $\mathbf{1}_E$ denote the indicator of the support of G .

To distinguish the finite simulation horizon from the infinite-horizon resolvent, for any nonnegative matrix M define

$$\mathcal{R}_H(M) := \sum_{t=1}^H M(\delta I + M)^{t-1}, \quad H < \infty.$$

Whenever $\rho(M) < 1 - \delta$, its infinite-horizon counterpart is $\mathcal{R}_\infty(M) = M((1 - \delta)I - M)^{-1}$. The empirical propagation models use the common finite horizon $H = T$, while infinite-horizon expressions are used for the classical accounting benchmarks and theoretical limiting results. Because Definition 1 includes the initial unit impulse, the score associated with a deterministic propagation operator M at horizon H is $m_s^{(H)} = 1 + \mathbf{1}^\top \mathcal{R}_H(M) e_s$.

Recall also that $0 \leq M \leq M'$ implies $\rho(M) \leq \rho(M')$ for nonnegative matrices (Horn and Johnson, 2013). Whenever the marked-Poisson dense-activation limit is invoked, we additionally require that every transmissible edge can be selected, i.e. $\alpha_j B_{ji} > 0$ whenever $G_{ji} > 0$.

3.1 P1: Dense resolvent limit and classical input–output orientations

Run the operator with the generic base transmission matrix G defined in Section 2. The next result establishes its exact dense full-activation limit. In the input–output application, the orientation of this limit can then be compared with the classical Leontief and Ghosh accounting benchmarks.

Proposition 1 (Full-activation limit). *Under full deterministic activation $I(t) \equiv \mathbf{1}_E$ and no capacity ($\sigma_C = \text{id}$), cumulative absorption over any finite horizon H is*

$$\sum_{t=1}^H P(t) = \mathcal{R}_H(A_G^\star) s_0. \quad (5)$$

If in addition $\rho(A_G^\star) < 1 - \delta$, letting $H \rightarrow \infty$ gives

$$\sum_{t \geq 1} P(t) = A_G^\star ((1 - \delta)I - A_G^\star)^{-1} s_0. \quad (6)$$

For $\delta = 0$ and $\Omega = 1$, this becomes $((I - G^\top)^{-1} - I)s_0$. After including the initial unit impulse as in Definition 1, the importance score of sector s is the column- s sum of $(I - G^\top)^{-1}$, equivalently the row- s sum of $(I - G)^{-1}$. (Proof: Appendix A.1.)

Corollary 1 (Leontief and Ghosh orientations). *Let*

$$A^\circ := A - \text{diag}(A), \quad B_G^\circ := B_G - \text{diag}(B_G).$$

For $\delta = 0$, $\Omega = 1$, and $H = \infty$, whenever the corresponding resolvent exists:

- (i) (Leontief orientation.) *Setting $G = (A^\circ)^\top$ gives the off-diagonal demand-side resolvent $(I - A^\circ)^{-1}$.*
- (ii) (Ghosh orientation.) *Setting $G = B_G^\circ$ gives the off-diagonal supply-side resolvent $(I - B_G^\circ)^{-1}$.*

These limits have the same directional interpretation as the corresponding classical input–output systems, while the standard benchmarks $M0_L$ and $M0_G$ retain the full matrices A and B_G , including their diagonal entries. In the empirical application, $G = B_G^\circ$, so the stochastic models propagate exclusively along intersectoral sales links and $M0_G$ serves as the natural classical supply-side reference. (Proof: Appendix A.2.)

Remark 1 (Dense endpoint versus classical benchmark). The distinction between $M2b$ and $M0_G$ is deliberate. $M2b$ is the exact full-activation endpoint of the stochastic operator: it uses B_G° together with the calibrated Ω , δ , and finite horizon T . $M0_G$ is instead the standard infinite-horizon Ghosh multiplier constructed from the full allocation matrix B_G . Thus $M0_G$ identifies the classical supply-side orientation, whereas $M2b$ is the appropriate dense benchmark for isolating the effect of sparse activation.

3.2 P2: Mean-field linearization and operator domination

Proposition 2 (Effective operator). *In the no-capacity regime ($\sigma_C = \text{id}$), the expected state obeys*

$$\mathbb{E}[S(t)] = (\delta I + \bar{B})\mathbb{E}[S(t-1)], \quad \bar{B} := \mathbb{E}[A(t)] = \Omega(G \odot \Pi)^\top. \quad (7)$$

For every finite horizon H , expected cumulative absorption is $\mathcal{R}_H(\bar{B})s_0$, and the expected importance score of sector s is $1 + \mathbf{1}^\top \mathcal{R}_H(\bar{B})e_s$. If $\rho(\bar{B}) < 1 - \delta$, the corresponding infinite-horizon operator is $\bar{B}((1 - \delta)I - \bar{B})^{-1}$. Moreover:

- (i) **Domination.** $0 \leq \bar{B} \leq A^\star$ entrywise, hence $\rho(\bar{B}) \leq \rho(A^\star)$.

- (ii) **Finite- and infinite-horizon envelope.** For every finite H , $\mathcal{R}_H(\bar{B}) \leq \mathcal{R}_H(A^\star)$ entrywise. If $\rho(A^\star) < 1 - \delta$, the same ordering holds at infinite horizon:

$$\bar{B}((1 - \delta)I - \bar{B})^{-1} \leq A^\star((1 - \delta)I - A^\star)^{-1}.$$

Consequently, at a matched horizon and calibration, expected sparse propagation is bounded above by its dense full-activation counterpart.

- (iii) **Structural divergence.** Holding G , Ω , δ , and H fixed, the expected score vector is a deterministic functional of the inclusion-probability matrix Π . Hence systematic ranking differences between sparse and dense propagation are properties of the effective operator $\bar{B}$, not artifacts of Monte Carlo sampling. (Proof: Appendix A.3.)

Remark 2 (Relation to annealed network representations). Proposition 2 is closely related to annealed network formulations, where a network ensemble is represented through an annealed adjacency matrix (Guerra and Gómez-Gardeñes, 2010), and to bond-percolation representations in which random edge states are summarized by their occupation probabilities (Newman, 2002). In our setting, the annealed representation is exact under the stated independence assumptions: $\bar{B} = \Omega(G \odot \Pi)^\top$ governs expected multiround propagation. Proposition 2 establishes that $\bar{B}$ is entrywise dominated by $A^\star$ for fixed Π ; Proposition 4 shows that along the activation-intensity path, this bound is achieved in the limit, $\bar{B}(\theta) \uparrow A^\star$ as $\theta \rightarrow \infty$. Proposition 5 shows that this linear representation no longer characterizes the dynamics once finite absorption capacity binds, breaking scale invariance through a separate nonlinear channel.

Remark 3. Part (ii) separates attenuation in the *amount* of propagation from reordering in its *distribution*: sparse activation cannot increase expected cumulative propagation relative to full activation, yet the heterogeneous reweighting $G \mapsto G \odot \Pi$ can change relative sectoral scores. Part (iii) therefore gives a closed-form population target against which the Monte Carlo implementation can be validated. Importantly, $\bar{B}$ is not introduced as an ad hoc reweighted adjacency matrix: it is the exact population operator induced by the underlying stochastic propagation process under the independence assumptions of Proposition 2. These assumptions concern only independence across rounds: arbitrary dependence among links activated within the same round leaves $\bar{B}$ unchanged.

3.3 P2': Effective attenuation and the anatomy of sparse reweighting

The next two results decompose the passage from $A^\star$ to $\bar{B}$. The first shows that persistence and attenuation are not separately relevant at infinite horizon. The second shows that sparse activation acts through two channels: an emitter-specific attenuation and a reallocation of transmission within each emitter's outgoing links.

Corollary 2 (Effective attenuation). *Let $M \geq 0$.*

(i) $\rho(\delta I + M) = \delta + \rho(M)$.

(ii) *If $\rho(M) < 1 - \delta$, then*

$$\mathcal{R}_\infty(M) = K(I - K)^{-1}, \quad K := \frac{M}{1 - \delta}.$$

In particular, for $M = \Omega(G \odot \Pi)^\top$, including the dense case $\Pi = \mathbf{1}_E$, infinite-horizon scores depend on (Ω, δ) only through the effective attenuation $\Omega_{\text{eff}} := \Omega/(1 - \delta)$, and $\rho(K) = \rho(M)/(1 - \delta) \in [0, 1)$ measures the distance of the propagation process to criticality. (Proof: Appendix A.4.)

At finite horizon H , $\mathcal{R}_H(M)$ is not an exact function of K , but the discrepancy vanishes as $(\delta + \rho(M))^H \rightarrow 0$. Under the baseline calibration, $\Omega_{\text{eff}} \approx 2.31$, and the dense and sparse operators lie at distances $\rho(A^\star)/(1 - \delta) \approx 0.86$ and $\rho(\bar{B})/(1 - \delta) \approx 0.60$ from criticality (Appendix F).

Proposition 3 (Row retention and within-row reallocation). *For each emitter j with $g_j := \sum_i G_{ji} > 0$, define its retention factor*

$$r_j := \frac{1}{g_j} \sum_i G_{ji} \pi_{ji} \in [0, 1].$$

- (i) **Exact factorization.** *Let $\widehat{G}_{ji} := G_{ji} \pi_{ji} / r_j$ if $r_j > 0$, and $\widehat{G}_j := G_j$ otherwise. Then $\sum_i \widehat{G}_{ji} = g_j$ for every j , and*

$$G \odot \Pi = D_r \widehat{G}, \quad \bar{B} = \Omega \widehat{G}^\top D_r, \quad D_r := \text{diag}(r).$$

Relative to $A^\star = \Omega G^\top$, sparse activation therefore (a) scales each emitter's total outgoing transmission by r_j , and (b) reallocates that transmission across the emitter's links without changing its total.

- (ii) **Concentration bound.** *Suppose $G_{ji} = g_j \bar{\omega}_{ji}$ and $\pi_{ji} = 1 - e^{-\alpha_j \bar{\omega}_{ji}}$, i.e. marked-Poisson activation with $\lambda = 1$. Writing $\bar{\omega}_j := \max_i \bar{\omega}_{ji}$,*

$$\frac{\alpha_j}{n_j^{\text{eff}}} \left(1 - \frac{\alpha_j \bar{\omega}_j}{2} \right) \leq r_j \leq 1 - \exp\left(-\frac{\alpha_j}{n_j^{\text{eff}}}\right) \leq \frac{\alpha_j}{n_j^{\text{eff}}}.$$

Hence $r_j = (\alpha_j / n_j^{\text{eff}}) (1 + O(\alpha_j \bar{\omega}_j))$ at low intensity, and r_j is nondecreasing in α_j .

- (iii) **Concentration ordering.** *Under the conditions of (ii), let $\alpha_j = \alpha_k = \alpha$. If $\bar{\omega}_j$ majorizes $\bar{\omega}_k$ and $\alpha \bar{\omega}_j \leq 2$, then $r_j \geq r_k$: at equal intensity, the more concentrated emitter retains more of its transmission.*

(Proof: Appendix A.5.)

In the input–output application with $\gamma = 1$, the proportionality condition in (ii) holds exactly, since $(B_G^\circ)_{ji} = W_{ji}^\circ / x_j = (\sum_k W_{jk}^\circ / x_j) \bar{\omega}_{ji}$, so g_j is the intersectoral sales share of sector j 's output. Parts (ii)–(iii) are stated for $\lambda = 1$, which corresponds to the weights-only kernel in Table 11; part (i) holds for any activation kernel.

Remark 4 (Attenuation versus selection). Corollary 2 and Proposition 3 imply that the comparison between $A^\star$ and $\bar{B}$ combines two distinct effects. Even uniform activation $\Pi = p \mathbf{1}_E$, which reweights no link relative to another, gives $\bar{B} = p A^\star$ and generally changes the resolvent ranking when $p < 1$, because resolvent rankings depend on the distance to criticality: as $\rho(K) \downarrow 0$ the score ordering approaches the ordering of outgoing transmission g_s , and as $\rho(K) \uparrow 1$ it approaches the Perron-vector ordering of G . The component attributable to *selection* is isolated by comparing $\bar{B}$ with the spectrally matched uniform benchmark $p^\star A^\star$, $p^\star := \rho(\bar{B}) / \rho(A^\star)$, which shares the distance to criticality of $\bar{B}$ but applies no differential reweighting. By Proposition 3(i), any remaining difference between $p^\star A^\star$ and $\bar{B}$ arises entirely from the dispersion of the retention factors r_j around $p^\star$ and from the within-row reallocation $\widehat{G}$. Under the baseline calibration, $p^\star \approx 0.69$, and Exercise E4 shows that the uniform benchmark $p^\star A^\star$ reproduces only 3% of the observed rank reordering, while the retention channel reproduces 58%.

Remark 5 (Second moments). Proposition 2 depends on the activation law only through Π , but higher moments do not. In the no-capacity regime, let $Q(t) := \mathbb{E}[S(t)S(t)^\top]$ and $M(t) := \delta I + A(t)$. Since $\text{vec}(M(t)S(t-1)S(t-1)^\top M(t)^\top) = (M(t) \otimes M(t)) \text{vec}(S(t-1)S(t-1)^\top)$ and $M(t)$ is independent of $S(t-1)$,

$$\text{vec } Q(t) = \mathbb{E}[M(t) \otimes M(t)] \text{vec } Q(t-1).$$

The operator $\mathbb{E}[M(t) \otimes M(t)]$ depends on the joint inclusion probabilities $\mathbb{E}[I_{ji}(t)I_{lk}(t)]$, not only on Π . Two activation schemes with identical Π therefore produce identical expected scores but, in general, different dispersion of realized cascades. Under the marked-Poisson kernel, links are independent, so $\mathbb{E}[I_{ji}I_{lk}] = \pi_{ji}\pi_{lk}$ for distinct links and $\mathbb{E}[I_{ji}^2] = \pi_{ji}$.

3.4 P3: Monotone interpolation between sparse propagation and the dense benchmark

Proposition 4 (Activation-intensity interpolation). *Index activation by a scalar $\theta \geq 0$ through $\alpha_j \mapsto \theta \alpha_j$, and suppose each $\pi_{ji}(\theta)$ is continuous and nondecreasing, with $\pi_{ji}(0) = 0$ and $\pi_{ji}(\theta) \rightarrow \mathbf{1}\{G_{ji} > 0\}$ as $\theta \rightarrow \infty$. The marked-Poisson kernel $\pi_{ji}(\theta) = 1 - e^{-\theta \alpha_j B_{ji}}$ satisfies these conditions whenever $\alpha_j B_{ji} > 0$ on the support of G . Then, in the no-capacity regime, $\bar{B}(\theta) = \Omega(G \odot \Pi(\theta))^\top$ satisfies:*

- (i) $\theta \mapsto \bar{B}(\theta)$ is entrywise continuous and nondecreasing, with $\bar{B}(0) = 0$ and $\bar{B}(\theta) \uparrow A^*$ as $\theta \rightarrow \infty$;
- (ii) $\theta \mapsto \rho(\bar{B}(\theta))$ is continuous and nondecreasing;
- (iii) for every finite horizon H , $\mathcal{R}_H(\bar{B}(\theta))$ is entrywise continuous and nondecreasing in θ , and $\mathcal{R}_H(\bar{B}(\theta)) \rightarrow \mathcal{R}_H(A^*)$ as $\theta \rightarrow \infty$. If $\rho(A^*) < 1 - \delta$, the same convergence holds for the infinite-horizon resolvent. Hence the expected score vector varies continuously with activation intensity and converges to the dense full-activation score vector. The induced ranking can change only when sectoral score paths cross; if the dense endpoint has no ties, the ranking eventually coincides with the dense ranking for sufficiently large θ . (Proof: Appendix A.6.)

Remark 6. Proposition 4 is the theoretical counterpart of the activation-intensity exercise in Appendix E3. At the empirical horizon $H = T$, increasing θ traces the transition from sparse mean-field propagation toward *M2b*. Rank-similarity statistics need not themselves vary continuously—rankings are discrete objects—but the underlying sectoral score vector does.

3.5 P4: Capacity, boundedness, and loss of scale invariance

Proposition 5 (Nonlinearity from capacity). *Let σ_C be any admissible saturation map in (3).*

- (i) **Global boundedness.** *For every round and realization, $0 \leq P_i(t) \leq C_i$. If $\delta < 1$, then $\limsup_{t \rightarrow \infty} S_i(t) \leq C_i / (1 - \delta)$, independently of $\rho(A^*)$. Thus capacity bounds the nonlinear state even when the corresponding dense linear system is unstable.*
- (ii) **Loss of scale invariance.** *Fix $v \geq 0$, $v \neq 0$, and let $s_0 = c v$, $c \geq 0$. For every finite horizon H , define total absorption including the initial impulse by*

$$\mathcal{T}_H(c) := c \mathbf{1}^\top v + \sum_{t=1}^H \mathbf{1}^\top P(t; c).$$

Pathwise, $\mathcal{T}_H(c)$ is nondecreasing and concave in c , with $\mathcal{T}_H(0) = 0$. Therefore $\mathcal{T}_H(c) / (c \mathbf{1}^\top v)$ is nonincreasing for $c > 0$. The same properties hold after taking expectations. Strict decline occurs whenever the induced pathwise total is strictly concave over the relevant scale range, i.e. whenever saturation produces a strict departure from proportional propagation along a positive-flow path.

- (iii) **Normalized stress intensity.** *For the hard-cap and smooth-exponential maps used in the paper, let $C_i = \kappa x_i$ and $s_0 = qv$. Since both maps satisfy $\sigma_{\kappa x}(\kappa y) = \kappa \sigma_x(y)$, the normalized finite-horizon response depends on (q, κ) only through $\zeta := q/\kappa$:*

$$\frac{\mathcal{T}_H(q, \kappa)}{\|s_0\|_1} = \Psi_H(\zeta; v, x)$$

for a function Ψ_H determined by the network and the remaining model parameters. This applies both to sector-specific shocks $v = x_s e_s$ and to aggregate output-proportional shocks $v = x$.

If the corresponding infinite-horizon cumulative absorption is finite, the conclusions in part (ii) extend to $H = \infty$. (Proof: Appendix A.7.)

Remark 7 (Economic reading). Capacity changes two properties of the linear diffusion benchmark. First, it guarantees global boundedness of the carried state even when unrestricted dense propagation would be

unstable. Second, it breaks proportionality with respect to shock size: once saturation becomes quantitatively relevant, cumulative absorption per unit of initial impulse weakly declines. With capacities $C_i = \kappa x_i$, the natural stress coordinate is therefore $\zeta = q/\kappa$, exactly the shock-to-capacity ratio used in the empirical capacity exercise.

Remark 8 (From operator properties to ranking divergence). Proposition 2 through Proposition 5 establish properties of the propagation operator, not the magnitude of empirical rank changes. Sparse activation can reorder sectors because Π reweights bilateral transmission non-uniformly, while capacity adds a separate path-dependent nonlinear compression once bottlenecks become relevant. Whether these channels generate economically meaningful reordering is therefore a quantitative question taken to the data in Section 5, and their relative contributions are quantified in Section 6.

4 Data


We use the 2021 symmetric input–output matrix for Colombia published by the Departamento Administrativo Nacional de Estadística (DANE), under the activity-by-activity convention, base 2015, and classified according to CIIU Rev. 4 adapted for Colombia (Departamento Administrativo Nacional de Estadística (DANE), 2025). The matrix is the most recent input–output table currently published by DANE. Following international recommendations, DANE evaluates and/or compiles the input–output matrix at approximately five-year intervals, or earlier when structural changes or statistical considerations warrant an update.

After aligning sector labels and removing non-productive aggregates, we work with $n = 67$ sectors. Gross output by sector (x_j , the “Producción total” row) is taken from the full table. The complete input–output matrix is used to construct the Leontief technical-coefficient matrix A and the Ghosh allocation matrix B_G . For the stochastic propagation network, we remove within-sector self-links and use $G = B_G^\circ := B_G - \text{diag}(B_G)$ as the transmission matrix. The same intersectoral convention is used when constructing the activation weights and productive-proximity kernel. Sectoral capacities are parameterized as $C_i = \kappa x_i$.

Three stylized facts motivate the exercise. First, the intermediate-transactions graph is *nearly complete*: of the $n(n - 1) = 4,422$ possible directed links, 4,142 carry positive flow, implying a density of 0.94 and a mean out-degree of about 62. Second, this high topological density masks substantial concentration in the distribution of transaction weights. The inverse-Herfindahl effective degree n_j^{eff} , defined in Section 2, has a median of approximately 6.3 and a mean of 7.1. Thus, a typical sector trades with around sixty destinations but concentrates most of its sales on the equivalent of only six or seven equally weighted partners. Third, the technical-coefficient matrix has spectral radius $\rho(A) \approx 0.54$, ensuring that the Leontief inverse is well defined and that the frictionless benchmark is non-degenerate.

Figure 1 illustrates the distinction between topological and effective connectivity. Panel (a) displays the complete positive-flow network, whereas panel (b) retains, for each emitting sector j , its $\text{round}(n_j^{\text{eff}})$ largest outgoing links. Node size and color are held fixed across panels and represent total weighted strength, defined as the sum of weighted in-strength and weighted out-strength. The common layout makes the comparison purely about the edge structure.

Figure 1: Topological and effective connectivity in the Colombian production network



Notes: Panel (a) shows the full directed input–output network, including every positive intermediate transaction between the 67 sectors. Panel (b) shows an effective network obtained by retaining, for each sector j , the $\text{round}(n_j^{\text{eff}})$ largest outgoing links, where $n_j^{\text{eff}} = (\sum_i p_{ji}^2)^{-1}$ and $p_{ji} = W_{ji}^\circ / \sum_k W_{jk}^\circ$. Node positions are identical across panels. Node size and color report total weighted strength, equal to weighted in-strength plus weighted out-strength, and use the same scale in both panels. Edge opacity is used only for visualization and does not alter the underlying transaction weights.

The contrast highlights the distinction between topological and effective connectivity: although almost all sector pairs are linked, the economically relevant structure is concentrated in a much smaller set of relationships. This motivates modeling propagation through a sparse subset of active links. The proximity matrix $\tilde{\phi}$ is constructed from the same input–output network, as defined in Section 2.

5 Empirical strategy

The empirical analysis applies the propagation framework to Colombia’s 2021 input–output network and examines how sectoral influence rankings change across alternative propagation mechanisms. The exercises compare a family of diffusion operators on the same weighted directed graph, so the underlying network is held fixed throughout the analysis. All propagation models from *M2b* onward use the finite-horizon analogue of the importance criterion in Definition 1: cumulative network-wide absorption generated by a unit sectoral impulse. The classical Leontief and Ghosh benchmarks retain their standard infinite-horizon resolvent definitions.

A central empirical issue is orientation. The Leontief inverse is demand-oriented: it ranks sectors by the total upstream requirements induced by a final-demand shock. The Ghosh inverse is supply-oriented: it ranks sectors by the total downstream allocation effects of an output shock. Since the stochastic operator developed in this paper propagates forward along directed sales links, the empirical design reports both classical benchmarks and separates orientation effects from sparse-activation effects.

5.1 Model ladder

The empirical comparison includes two classical input–output benchmarks and a sequence of supply-side propagation models. Let W_{ji} denote the flow from sector j to sector i , and let x_j be gross output of sector j .

The classical Leontief technical-coefficient matrix is

$$A_{ji} = \frac{W_{ji}}{x_i},$$

while the classical Ghosh allocation matrix is

$$(B_G)_{ji} = \frac{W_{ji}}{x_j}.$$

Both retain the diagonal entries of the input–output table. The stochastic propagation network instead excludes within-sector self-links and uses

$$B_G^\circ := B_G - \text{diag}(B_G)$$

as its supply-side transmission matrix. This preserves the forward sales orientation of Ghosh while restricting stochastic propagation to intersectoral links.

$M0_L$ — Leontief benchmark.

The classical demand-side benchmark. Scores are the column sums of

$$(I - A)^{-1},$$

corresponding to standard backward-linkage multipliers.

$M0_G$ — Ghosh benchmark.

The classical supply-side benchmark. Scores are the row sums of

$$(I - B_G)^{-1},$$

which measure cumulative downstream allocation effects from a sectoral output shock. It is retained as the standard accounting benchmark for the supply-side orientation.

$M2b$ — Dense supply-side operator.

The full-activation counterpart of the stochastic propagation model. It uses B_G° as the transmission matrix and the same finite horizon T , attenuation parameter Ω , and persistence parameter δ used in the sparse models. It is therefore the exact dense endpoint of the empirical propagation family.

$M3_{mf}$ — Mean-field sparse operator.

The closed-form expected ranking of the sparse stochastic model without capacity. Its effective propagation matrix is

$$\bar{B} = \Omega(B_G^\circ \odot \Pi)^\top,$$

where Π_{ji} is the activation probability of link $j \rightarrow i$. Relative to $M2b$, the only change is that intersectoral links are activated probabilistically rather than fully.

$M3$ — Stochastic sparse diffusion, no capacity.

The Monte Carlo implementation of $M3_{mf}$. Each active sector transmits through a random subset of outgoing intersectoral links, with transmission magnitudes determined by B_G° and activation probabilities determined by the selection kernel.

$M4$ — Stochastic sparse diffusion with capacity.

The capacity-constrained version of $M3$, with sectoral absorption limits

$$C_i = \kappa x_i$$

and saturation map σ_C . Relative to $M3$, the only additional mechanism is the nonlinear capacity constraint studied in Proposition 5.

The classical benchmarks $M0_L$ and $M0_G$ establish the demand- and supply-side accounting orientations. Conditional on the supply-side propagation design, the nested empirical sequence is

$$M2b \longrightarrow M3_{mf} \longrightarrow M3 \longrightarrow M4.$$

The comparison of $M2b$ with $M0_L$ and $M0_G$ establishes the orientation of the dense propagation operator rather than an exact algebraic nesting relation. Within the propagation family, $M2b$ versus $M3_{mf}$ isolates sparse activation, $M3_{mf}$ versus $M3$ validates the Monte Carlo implementation, and $M3$ versus $M4$ isolates capacity. All models from $M2b$ onward use the same off-diagonal transmission matrix, finite horizon T , and score definition. Infinite-horizon resolvents are used only for the classical Leontief and Ghosh benchmarks.

5.2 Calibration

The calibration uses observed features of the network to set the main scale parameters. Other parameters are fixed in the baseline and varied in the robustness exercises. The baseline is chosen to place the sparse diffusion process on an interpretable scale and to provide a common reference point for the exercises that follow.

The common transmission attenuation is normalized using the aggregate intermediate-output share of the economy,

$$\Omega = \frac{\sum_{j,i} W_{ji}}{\sum_j x_j} = 0.474.$$

Because the base transmission matrix is $G = B_G^\circ$, sector-specific intermediate-output shares are already contained in the row sums of G . The role of Ω is therefore to provide a common attenuation factor for the propagation process while preserving the cross-sector heterogeneity embedded in B_G° .

The activation scale is chosen so that the mean expected number of active links across emitters matches the median inverse-Herfindahl effective degree of the sales network, equal to 6.29 in the Colombian matrix. This uses the observed concentration of intersectoral sales to set the baseline degree of sparsity. The activation-intensity exercise in Appendix E3 varies this scale directly, and Appendix B considers alternative cross-sector and bilateral activation specifications with comparable levels of sparsity.

The persistence parameter is chosen so that the mean-field process has a diffusion half-life of eight rounds,

$$\rho(\delta I + \bar{B}) = 2^{-1/8} = 0.917,$$

which implies $\delta = 0.795$ given the calibrated effective operator. Under the same parameters, the dense full-activation operator satisfies

$$\rho(\delta I + A^\star) = 0.971 < 1,$$

so both the sparse and dense propagation processes remain stable.

The baseline horizon is $T = 20$, the activation kernel is marked Poisson, and the selection kernel uses $\lambda = 0.8$, placing most of the weight on observed sales while still allowing productive proximity to affect link activation. Capacity is parameterized as

$$C_i = \kappa x_i,$$

with baseline $\kappa = 0.5$ and smooth saturation. Baseline ranking exercises use a unit sectoral impulse, consistent with Definition 1. Under this calibration, capacity is largely inactive in the baseline ranking exercise and is examined separately through the saturation analysis, where larger shock-to-capacity ratios generate binding bottlenecks.

The full calibration table is reported in Appendix F.

5.3 Exercises

5.3.1 E1 — Orientation validation and localized rank reordering

The first exercise establishes the relevant benchmark for the sparse propagation models and then documents where the sparse ranking departs from the classical supply-side ordering. We compute the Leontief benchmark $M0_L$, the Ghosh benchmark $M0_G$, the dense finite-horizon supply-side benchmark $M2b$, the closed-form sparse operator $M3_{mf}$, and the Monte Carlo sparse ranking $M3$. To separate algorithmic variation from cross-sector stability, Monte Carlo rankings are averaged across independent seed streams, while confidence intervals are computed by nonparametric bootstrap resampling of sectors.

Table 1: Orientation validation and nested propagation diagnostics


Model	Reference	Spearman		Kendall		Top-10	Top-20	Mean $ \Delta r $
		Coeff.	[95% CI]	Coeff.	[95% CI]			
$M2b$	$M0_L$	0.010	[−0.255, 0.277]	0.015	[−0.177, 0.206]	0.30	0.25	20.99
$M2b$	$M0_G$	0.932	[0.871, 0.965]	0.794	[0.713, 0.864]	0.60	0.70	5.10
$M3_{mf}$	$M2b$	0.945	[0.884, 0.975]	0.825	[0.743, 0.890]	0.90	0.70	4.45
$M3$	$M3_{mf}$	1.000	[0.999, 1.000]	0.999	[0.994, 1.000]	1.00	1.00	0.03

Notes: $M0_L$ and $M0_G$ are the classical Leontief and Ghosh benchmarks. $M2b$ is the dense finite-horizon supply-side operator aligned with the stochastic model. $M3_{mf}$ is its sparse mean-field counterpart, so the comparison between $M2b$ and $M3_{mf}$ isolates the effect of probabilistic link activation. $M3$ is the Monte Carlo implementation of $M3_{mf}$. Top- k reports the share of common sectors in the top k of both rankings. Mean $|\Delta r|$ is the average absolute rank difference. Confidence intervals are percentile intervals from 5,000 nonparametric bootstrap resamples of sectors. They should be interpreted as resampling-based measures of cross-sector stability.

The benchmark comparisons in Table 1 confirm the supply-side orientation of the empirical operator. Both the dense full-activation benchmark $M2b$ and the sparse rankings $M3_{mf}$ and $M3$ are strongly correlated with $M0_G$ and essentially uncorrelated with $M0_L$. This pattern is consistent with the construction of the operator, which propagates shocks forward through intersectoral sales links using B_G° as its transmission matrix. Because $M0_G$ is computed from the full classical Ghosh matrix whereas the propagation models exclude self-links, the $M0_G$ – $M2b$ comparison should be interpreted as an orientation benchmark rather than an exact nesting comparison. The bootstrap intervals reinforce this result: correlations with $M0_G$ remain high under sector resampling, whereas those with $M0_L$ remain centered near zero. The last row additionally confirms that the Monte Carlo implementation is numerically indistinguishable from the closed-form mean-field ranking; Appendix E2 documents its convergence across simulation budgets and independent seed streams.

Having established the relevant supply-side orientation, we next examine where the sparse ranking departs from the classical Ghosh ordering. Although the sparse ranking remains highly correlated with Ghosh in the aggregate, this global statistic masks substantial heterogeneity across the ranking distribution. To characterize where this departure occurs, we compute the mean absolute rank change between the Ghosh and $M3$ rankings within a moving window of 15 sectors ordered by their Ghosh position, and plot this quantity against the center of the window. Because $M0_G$ retains self-links and is evaluated at infinite horizon, this profile describes departures from the classical supply-side ordering and combines sparse activation with these benchmark differences; the component attributable to sparse activation alone is isolated by the $M2b$ – $M3_{mf}$ comparison in E2. Since the $M3$ and $M3_{mf}$ rankings are numerically indistinguishable (last row of Table 1), using either produces the same profile.

Figure 2: Localized rank reordering along the Ghosh ranking



Notes: The solid line reports the mean absolute difference between the global Ghosh and $M3$ ranks in overlapping moving windows of 15 sectors, ordered by the Ghosh benchmark. The horizontal axis indicates the center of each window. Shaded areas are 95% percentile intervals from 5,000 nonparametric bootstrap resamples of sectors within each window and summarize cross-sector stability. Because adjacent windows overlap, the estimates and confidence intervals are not statistically independent.

Figure 2 reveals a clear gradient in rank reordering. Mean absolute rank changes are high at the top of the Ghosh ranking, peak in the upper-middle part of the distribution, and then decline steadily toward the lower tail. In the first half of the ranking, average rank differences remain around 8–10 positions, whereas near the bottom they fall to roughly 2 positions. Across window centers, the Spearman correlation between Ghosh position and mean absolute rank change is approximately -0.91 , indicating a very strong decline in reordering as one moves down the supply-side ranking.

This pattern shows that sparse activation does not replace the Ghosh benchmark with an unrelated ordering. Instead, it preserves the broad supply-side structure while substantially reshuffling the sectors that are most central under dense propagation. The lower tail is comparatively stable, whereas the ordering of the most influential supply-side sectors is much more sensitive to the activation mechanism.

The moving-window evidence does not depend on arbitrary rank cutoffs: the divergence between dense and sparse propagation is concentrated among the sectors that matter most in the dense supply-side benchmark.

This localized reordering is the first empirical indication that dense resolvent centrality and sparse realized diffusion encode different notions of influence on the same production network. The rank-change decomposition in E2 isolates the corresponding sparse-activation component by comparing the aligned dense operator $M2b$ with its mean-field sparse counterpart $M3_{mf}$, and examines which network features account for those movements.

5.3.2 E2 — Accounting for rank reordering under sparse activation

The previous exercises show that sparse activation preserves the broad supply-side ordering while producing substantial changes in the relative position of individual sectors. This exercise examines whether those changes are systematically related to observable features of the production network.

The outcome is the change in sector i 's position between the dense, full-activation supply-side benchmark

and the sparse mean-field operator,

$$\Delta r_i = r_i^{M2b} - r_i^{M3mf}, \quad (8)$$

where lower numerical ranks indicate greater importance. Hence, $\Delta r_i > 0$ means that sector i moves upward under sparse activation, while $\Delta r_i < 0$ means that it loses positions relative to the dense benchmark. We use $M3_{mf}$ because it is the closed-form expected ranking of the no-capacity stochastic model under the marked-Poisson activation kernel.

Our baseline specification is

$$\Delta r_i = \beta_0 + \beta_1 \log(\text{OutStrength}_i) + \beta_2 n_i^{\text{eff}} + \beta_3 \text{DownstreamReach}_i + \varepsilon_i, \quad (9)$$

where OutStrength_i is the total value of intermediate sales, n_i^{eff} is the inverse-Herfindahl effective number of outgoing destinations, and $\text{DownstreamReach}_i = \sum_j [(I - B_G)^{-1}]_{ij}$ is the row sum of the Ghosh inverse, capturing cumulative downstream reach through the production network. All covariates are standardized, so coefficients are measured in rank positions associated with a one-standard-deviation change in the corresponding network characteristic.

We then consider two complementary specifications. The propagation-depth specification replaces DownstreamReach_i with DenseMeanRound_i , the absorption-weighted mean propagation round under the aligned dense operator $M2b$. This variable captures how strongly a sector's influence depends on later stages of propagation. The extended specification retains DownstreamReach_i and adds topological out-degree and proximity in-strength, allowing the baseline correlates to be assessed alongside broader measures of connectivity and productive relatedness.

Because the baseline activation intensity depends partly on sector size and effective diversification, these regressions are descriptive. Proposition 3 also implies that part of their sign pattern is structural: at a given intensity, emitters with a larger effective number of destinations retain less of their transmission, while the baseline profile assigns higher intensity to larger emitters. The regressions therefore quantify how strongly these mechanisms operate in the Colombian network, not whether they exist. The coefficients summarize network characteristics associated with rank reordering under the baseline calibration and should not be interpreted as causal effects. Inference is based on heteroskedasticity-robust HC1 standard errors and is conditional on the 67-sector input-output system studied here.

Table 2: Network correlates of rank reordering under sparse activation

	(1) Baseline	(2) Propagation depth	(3) Extended
Log out-strength	3.87*** (0.74)	2.46*** (0.62)	3.95*** (0.90)
Effective degree, n^{eff}	-2.44*** (0.69)	-1.55** (0.68)	-2.60*** (0.80)
DownstreamReach	-2.54*** (0.70)		-2.51*** (0.73)
DenseMeanRound		-2.10*** (0.76)	
Out-degree			-0.82** (0.36)
Proximity in-strength			0.52 (0.80)
Observations	67	67	67
R^2	0.332	0.319	0.349
Adjusted R^2	0.300	0.287	0.296
Maximum VIF	1.43	1.11	1.98

Notes: The dependent variable is $\Delta r_i = r_i^{M2b} - r_i^{M3mf}$, so positive coefficients indicate characteristics associated with upward movement under sparse activation. Column (2) replaces DownstreamReach with DenseMeanRound, the absorption-weighted mean propagation round under $M2b$. Regressors are standardized. HC1 robust standard errors are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

The results provide empirical evidence of a consistent role for propagation depth. In the baseline specification, greater cumulative downstream reach is associated with losses in rank under sparse activation. Column (2) shows that the same empirical pattern survives when this resolvent-based measure is replaced by a direct measure of propagation depth: sectors whose dense influence is concentrated in later rounds also tend to lose positions when activation becomes sparse. Column (3) shows that these associations remain similar after adding broader measures of network connectivity and productive relatedness.

To connect these average relationships with the sector-level reordering, Table 3 reports the seven largest upward and downward movements between $M2b$ and $M3_{mf}$.

Table 3: Largest sectoral rank changes under sparse activation

Sector	r^{M2b}	r^{M3mf}	Δr
Largest upward movements			
Coffee	42	23	19
Livestock	35	16	19
Grain milling and processed foods	33	17	16
Non-metallic mineral products	26	11	15
Oils and fats	23	14	9
Specialized construction activities	21	12	9
Agriculture excluding coffee	34	26	8
Largest downward movements			
Other mining and quarrying	15	27	-12
Machinery and equipment	10	22	-12
Gas production and distribution	18	30	-12
Textiles	22	33	-11
Vehicle maintenance and repair	14	24	-10
Forestry and logging	19	28	-9
Air transport	30	38	-8

Notes: The table reports the seven largest upward and downward movements according to $\Delta r_i = r_i^{M2b} - r_i^{M3mf}$. Positive values indicate upward movement under sparse activation, and lower numerical ranks indicate greater importance. The complete comparison is reported in Appendix E4.


The reordering is not concentrated in a single broad industry group. The largest gains include primary activities, such as coffee and livestock, but also food-processing, mineral-products, and construction-related activities. What these sectors appear to share is not their economic classification, but a transmission structure concentrated in a relatively small set of relevant destinations. Sparse activation therefore gives greater weight to sectors whose influence can be sustained through a limited number of strong channels.

The largest declines include machinery, gas distribution, textiles, transport, and vehicle-maintenance activities. Several of these sectors are highly connected under the dense benchmark, but their influence is spread across a broader collection of outgoing relationships. When only a subset of links is activated, this broad connectivity no longer guarantees that the paths supporting their dense centrality remain simultaneously available. Sparse activation thus shifts influence away from sectors whose ranking relies more heavily on the joint contribution of many links.

These patterns are consistent with the regression results: greater outgoing scale is associated with upward movement, whereas effective diversification and longer downstream dependence are associated with rank losses. They also match the retention channel in Proposition 3: a sector that concentrates its sales on a few strong links keeps most of its outgoing transmission when activation is sparse, whereas a diversified sector loses a larger share, with r_j bounded above by $1 - \exp(-\alpha_j/n_j^{\text{eff}})$. Nevertheless, the magnitude of some sectoral changes remains larger than predicted by these aggregate characteristics. This is consistent with the link-specific structure of the effective operator, $\bar{B} = \Omega (B_G^\circ \odot \Pi)^\top$, under which reordering depends on the bilateral alignment between transmission weights and activation probabilities.

Figure 3 illustrates this distinction by plotting the observed rank change against the fitted value from the parsimonious specification. Sectors close to the 45-degree line are well summarized by the broad network characteristics, whereas large residuals identify cases in which bilateral activation structure plays an additional role.

Figure 3: Observed and predicted rank changes under sparse activation



Notes: The horizontal axis reports the fitted rank change from the parsimonious specification in equation (9); the vertical axis reports the observed change $\Delta r_i = r_i^{M2b} - r_i^{M3mf}$. Selected sectors with the largest absolute residuals are labeled.

Overall, sparse activation produces a systematic pattern of rank reordering that broad network characteristics only partly summarize. The regressions are descriptive, and part of their sign pattern is mechanical. Section 6 replaces them with an exact decomposition of the effective operator: the emitter-level retention channel of Proposition 3 alone accounts for 97% of the cross-sector variation in log score ratios, and the remaining rank changes reflect its interaction with the within-row reallocation of transmission along multi-step paths.

5.3.3 E3 — Capacity-binding and saturation regime

To study when the capacity mechanism becomes quantitatively relevant, we instead scale the initial shock with the output of the seed sector while normalizing the resulting score by the initial impulse. We therefore treat capacity as a nonlinear stress-regime channel and increase the initial impulse relative to the capacity scale.

Let q denote the initial shock as a fraction of the output of the seed sector, so that $S_s(0) = qx_s$. Because both the initial condition and capacities scale with sectoral output, the relevant normalized stress intensity is

$$\zeta = \frac{q}{\kappa}.$$

Holding the calibrated value of κ fixed, we trace the separation between the capacity-constrained model $M4$ and the no-capacity model $M3$ as ζ increases. Ranking scores are normalized by the initial shock, so that departures from $M3$ reflect the nonlinear capacity mechanism rather than the mechanical scale of the impulse.

Table 4: E3: Transition into the capacity-constrained regime

ζ	Spearman	Kendall	Mean $ \Delta r $	Max $ \Delta r $	Top-20	Mean score loss	Binding nodes
1	1.000	1.000	0.00	0	1.00	0.000	0
2	1.000	1.000	0.00	0	1.00	0.000	0
4	1.000	0.997	0.09	3	1.00	0.001	1
8	0.995	0.981	0.60	13	0.95	0.006	5
20	0.976	0.933	1.91	22	0.90	0.030	27
40	0.938	0.851	4.21	30	0.85	0.066	49
100	0.822	0.693	8.15	37	0.65	0.135	64

Notes: $\zeta = q/\kappa$ is the initial impulse relative to the capacity scale. Rank statistics compare $M4$ with the no-capacity ranking $M3$. Mean score loss is the cross-sector average of $1 - \text{score}_i^{M4}/\text{score}_i^{M3}$. Binding nodes counts receptors whose conditional binding frequency exceeds 0.1% of the rounds in which they receive positive flow. The table reports selected points from the full stress grid $\zeta \in \{0.10, 0.25, 0.50, 1, 2, 4, 8, 10, 20, 40, 80, 100\}$. The calibrated baseline, $q = 0.01$ and $\kappa = 0.5$, corresponds to $\zeta = 0.02$ and lies in the capacity-inactive region. Higher values of ζ are stress-regime counterfactuals used to characterize the nonlinear capacity mechanism rather than calibrated sectoral shock magnitudes.


Table 4 reveals the regime transition implied by the capacity mechanism in Proposition 5. For $\zeta \leq 2$, $M4$ is numerically indistinguishable from $M3$: rank correlations and top- k overlaps equal one, no sector changes rank, and no receiver exhibits recurrent binding. Capacity first becomes detectable around $\zeta = 4$, where the departure remains highly localized. At $\zeta = 20$, the Spearman correlation remains high at 0.976, but the mean absolute rank change rises to 1.91 positions and individual sectors move by as many as 22 positions. At $\zeta = 40$, the departure becomes more widespread: Spearman falls to 0.938, the mean absolute rank change exceeds four positions, and 49 of the 67 sectors bind in more than 0.1% of their active receiving events.

The reordering is highly uneven across sources. At $\zeta = 20$, coffee falls from rank 23 to 45 and livestock from 16 to 36, with normalized score losses of approximately 40 and 31 percent. Grain milling, agriculture excluding coffee, petroleum extraction, textiles, and basic metals also lose positions. Under stronger stress, these movements intensify, while pharmaceuticals, information and publishing activities, gas distribution, machinery, water transport, and vehicle maintenance gain positions because their propagation scores are less compressed than those of the most capacity-exposed sources.

These rank changes need not coincide with a sector's own binding frequency as a receiver. A source loses influence when the cascade it generates encounters capacity-constrained nodes anywhere along its downstream paths. Consistent with Proposition 5, capacity is therefore inactive in the low-stress regime but, once binding occurs, breaks scale invariance and compresses propagation non-uniformly. In the production network, this nonlinear compression also translates into selective reordering according to the bottlenecks encountered by each cascade.

Figure 4 shows the corresponding nonlinear response in aggregate propagation. At low stress intensities, $M4$ coincides with the no-capacity benchmark. Once capacity begins to bind, absorption per unit of initial impulse declines and the share of propagation lost to capacity rises.

Figure 4: Capacity-induced nonlinearity across stress regimes



Notes: The horizontal axis reports normalized stress intensity $\zeta = q/\kappa$. Panel (a) reports cumulative absorption normalized by the initial impulse. Panel (b) reports the percentage of cumulative absorption lost relative to the no-capacity benchmark. Shaded areas are 95% percentile intervals across Monte Carlo replications. The dotted line marks the onset of detectable capacity effects.

Taken together, the ranking and aggregate-response results show that capacity is inactive in the baseline regime but becomes increasingly important as the system moves into a high-stress regime.

6 What drives sparse reordering

The exercises in Section 5 establish that sparse activation reorders sectoral influence. This section asks four questions that any sparse-diffusion centrality must answer before the reordering can be attributed to the activation mechanism itself. First, is the reordering a mere attenuation effect, generated by moving the operator away from criticality? Second, is sparse activation equivalent to extracting a deterministic backbone and computing a dense resolvent on it? Third, does the effective operator carry predictive information about realized cascades beyond standard centralities? Fourth, which network properties govern the magnitude of the reordering, and does the mechanism hold beyond the Colombian table? E4 to E6 use the baseline calibration of Section 5, and E7 recalibrates the same rules for each of 79 OECD economies. E4, E6, and E7 are closed-form; E5 uses Monte Carlo ground truths.

6.1 E4 — Attenuation, retention, and reallocation

Proposition 3 and Remark 4 decompose the passage from $A^* = \Omega G^\top$ to $\bar{B} = \Omega(G \odot \Pi)^\top$ into separate channels. We compute the score vector of five operators that share $G = B_G^\circ$, Ω , δ , and T :

$$\underbrace{G}_{M2b} \longrightarrow \underbrace{p^* G}_{\text{uniform attenuation}} \longrightarrow \underbrace{D_r G}_{\text{retention only}} \longrightarrow \underbrace{c \hat{G}}_{\text{reallocation only}} \longrightarrow \underbrace{G \odot \Pi = D_r \hat{G}}_{M3_{mf}}, \quad (10)$$

where $p^* = \rho(G \odot \Pi)/\rho(G) = 0.694$ places the uniformly attenuated benchmark at the same distance to criticality as $\bar{B}$, $D_r = \text{diag}(r)$ scales each emitter's total transmission by its retention factor without reallocating it across destinations, and $c = \rho(G \odot \Pi)/\rho(\hat{G})$ rescales the mass-preserving reallocation $\hat{G}$ to the same spectral radius. For each candidate operator X we report its rank distance to $M3_{mf}$ and the share of the dense-to-sparse reordering it reproduces,

$$\text{Gap closed}(X) = 1 - \frac{d(X, M3_{mf})}{d(M2b, M3_{mf})},$$

with d the mean absolute rank change. Because rank statistics are sensitive to the local density of scores, we also report a score-based statistic: the share of the centered cross-sector variance of $\log(m_s^{M3mf}/m_s^{M2b})$ reproduced by $\log(m_s^X/m_s^{M2b})$ (log-ratio R^2).

Table 5: E4: Attenuation, retention, reallocation, and deterministic sparsification

Operator	Links	$\rho/(1-\delta)$	Spearman	Mean $ \Delta r $	Gap closed	Log-ratio R^2
<i>Panel A. Channels of the effective operator</i>						
Dense, $M2b$	61.8	0.859	0.945	4.45	0.00	0.00
Uniform attenuation, p^*A^*	61.8	0.596	0.952	4.30	0.03	0.75
Retention only, $D_r G$	61.8	0.581	0.988	1.88	0.58	0.97
Reallocation only, $c\hat{G}$	61.8	0.596	0.960	3.79	0.15	0.79
Sparse, $M3mf$	6.29 ^a	0.596	1.000	0.00	1.00	1.00
<i>Panel B. Dense resolvent on deterministic backbones</i>						
Top- k ($k_j = \text{round}(n_j^{\text{eff}})$), raw	7.09	0.650	0.962	3.58	0.20	0.76
Top- k ($k_j = \text{round}(n_j^{\text{eff}})$), ρ -matched	7.09	0.596	0.964	3.52	0.21	0.83
Disparity filter, matched sparsity	6.28	0.596	0.945	4.15	0.07	0.66
Disparity filter ($\alpha = 0.30$), ρ -matched	8.70	0.596	0.969	3.16	0.29	0.82
Global threshold, matched, ρ -matched	6.28	0.596	0.950	4.16	0.06	0.58

Notes: All operators use $G = B_G^\circ$, the calibrated Ω and δ , and $T = 20$. Spearman and mean $|\Delta r|$ are computed against $M3mf$. Gap closed is $1 - d(\cdot, M3mf)/d(M2b, M3mf)$ with d the mean absolute rank change; negative values indicate operators farther from $M3mf$ than $M2b$. Log-ratio R^2 is the share of the centered variance of $\log(m^{M3mf}/m^{M2b})$ reproduced by the candidate. ρ -matched operators are rescaled to $\rho(\bar{B})$, so they share the distance to criticality of $M3mf$. The disparity filter (Serrano et al., 2009) is applied to outgoing links; “matched sparsity” selects the significance level ($\alpha = 0.154$) whose backbone retains, on average, the baseline expected number of active links; $\alpha = 0.30$ is the best-performing level on the grid $\{0.01, 0.05, 0.1, 0.2, 0.3, 0.5\}$. Links are retained (or expected active) links per emitter. The raw and ρ -matched versions of the matched disparity backbone coincide to two decimals. ^aExpected number of active links per emitter.

Panel A of Table 5 answers the first question. Attenuation alone does not produce the reordering. The uniformly attenuated benchmark shares the distance to criticality of $M3mf$, yet it reproduces only 3% of the rank reordering. Its own rank changes relative to $M2b$ average 1.55 positions but are nearly unrelated to the observed ones (rank-change correlation 0.21). In score space it does account for a sizable share of the log-ratio variation ($R^2 = 0.75$), as expected from Remark 4: moving away from criticality reduces the weight of long walks for every emitter. This compression, however, rarely alters relative positions. The retention channel behaves differently. Scaling each emitter by r_j , without reallocating transmission across its destinations, reproduces 58% of the rank reordering, 97% of the log-ratio variation, and its rank changes correlate at 0.88 with those of $M3mf$. The within-row reallocation $\hat{G}$ alone reproduces 15%. The remaining rank gap therefore reflects the interaction between retention and reallocation along multi-step paths, which matters for ranks even though it accounts for only 3% of the variance in log scores.


The sector-level pattern matches the mechanism in Proposition 3. Livestock ($n^{\text{eff}} = 1.9$, $r = 0.91$) rises from rank 35 to 16, and coffee ($n^{\text{eff}} = 1.1$, $r = 0.92$) from 42 to 23; under retention alone they already reach ranks 15 and 21. Machinery ($n^{\text{eff}} = 10.6$, $r = 0.51$) falls from 10 to 22, and gas distribution ($n^{\text{eff}} = 10.6$, $r = 0.53$) from 18 to 30; under retention alone they fall to 24 and 32. Across sectors, the Spearman correlation between r_j and n_j^{eff} is -0.48 . The correlation is far from perfect because r_j also depends on the activation intensity α_j , which the baseline profile ties to sector size. Numerically, the bounds in Proposition 3(ii) hold for every sector under the weights-only kernel.

Panel B answers the second question. A natural alternative reading of Figure 1 is that sparse activation is simply a dense resolvent computed on the network’s backbone. We test this against three deterministic filters: the row-wise top- k backbone of Figure 1(b), the disparity filter (Serrano et al., 2009; Radicchi et al., 2011), and a global weight threshold. Each is computed with and without rescaling to the distance to criticality of

$M3_{mf}$. Even the best of them, a disparity backbone retaining 8.7 links per emitter, reproduces 29% of the reordering, half of what retention alone achieves. When the disparity filter is calibrated to the same sparsity as the baseline, it reproduces 7%. The global threshold, which isolates 10% of emitters, performs similarly. The difference is structural. A backbone decides deterministically which links survive and transmits the full dense weight on them, so a diversified emitter keeps most of its transmission as long as its largest links pass the filter. Under stochastic activation, by contrast, every emitter’s total transmission is scaled by a retention factor that penalizes diversification in proportion (Proposition 3). Sparse activation is therefore not equivalent to sparsification.

A natural objection is that the mechanism is built into the calibration: the baseline intensity profile α_j depends on sector size and on n_j^{eff} , so finding that diversification is penalized could be mechanical. Table 15 in Appendix E1 repeats the decomposition under alternative activation rules, each with the baseline sparsity. The evidence points the other way. With homogeneous intensity, so that α_j carries no information on size or diversification, reordering is larger (mean $|\Delta r| = 7.02$), the retention channel reproduces 74% of it, and the correlation between r_j and n_j^{eff} rises from -0.48 to -0.95 , as Proposition 3(ii) predicts. The baseline profile therefore dilutes the mechanism, because it assigns higher intensity to large and diversified emitters. Retention also dominates under size-only intensity (82%) and the weights-only kernel (66%). The one exception is intensity proportional to n_j^{eff} alone, which offsets the retention penalty by construction: there reordering is small (mean $|\Delta r| = 2.33$), and retention (33%) and the top- k backbone (50%) share the explanation.

Figure 5: Retention accounts for the reordering in scores



Notes: Each point is a sector. The vertical axis reports the observed centered log-ratio $\log(m_s^{M3_{mf}}/m_s^{M2b})$; the horizontal axis reports the centered log-ratio implied by uniform attenuation $p * A *$ (left) and by retention only $D_r G$ (right). The dotted line is the 45-degree line.

6.2 E5 — Predictive content relative to standard centralities

The third question is whether the effective operator carries information about realized propagation that standard centralities do not. We compare scores computed from the network alone against ground truths obtained by simulation. The thirteen candidates are out-strength and out-degree; PageRank (Brin and Page, 1998) and its reversed-network counterpart, CheiRank (Ermann et al., 2015); the right Perron vector of G ; Katz centrality on G with attenuation $0.85/\rho(G)$; communicability $\exp(G/\rho(G))$ (Estrada and Hatano, 2008); Viral Centrality (Fink et al., 2023), computed both on network weights alone and on a matrix that includes Π ; the Ghosh benchmark $M0_G$; the dense operator $M2b$; the first-order retention score $1 + \Omega_{\text{eff}} r_j g_j$;

and $M3_{mf}$. All scores are oriented as source influence except PageRank, which measures the receiving side and is reported only for completeness.

The ground truths avoid circularity. The Monte Carlo mean of $M3$ is excluded, since $M3_{mf}$ is its exact expectation by Proposition 2; that comparison is only a sanity check (Spearman 1.000). Instead we use the following outcomes:

- (i) the 90th percentile of realized cascade size in $M3$, a tail statistic that a first-moment operator does not target;
- (ii) normalized absorption in $M4$ at stress intensities $\zeta = 20$ and $\zeta = 40$, where capacity binds and the dynamics are nonlinear (E3);
- (iii) mean reach of a binary independent cascade (Kempe et al., 2003), in which each newly activated sector attempts once to activate each downstream partner $j \rightarrow i$ with probability $P_{ji} = \min\{1, \Omega_{\text{eff}} G_{ji} \pi_{ji}\}$.

The independent cascade is a different dynamic, binary and without divisibility or persistence, and is the setting for which Viral Centrality was designed. With $P = \Pi$ alone, the cascade saturates the dense network from every seed and provides no ranking information, which is why P combines activation and magnitude. Ground truths use 500 Monte Carlo replications per seed (2,000 for the independent cascade), and confidence intervals come from 2,000 paired bootstrap resamples of sectors.

Table 6: E5: Which network-only score predicts simulated cascade reach?

Score	Spearman correlation with simulated ground truth			
	$M3$, 90th pct.	$M4$, $\zeta = 20$	$M4$, $\zeta = 40$	Independent cascade
Sparse, $M3_{mf}$	0.997	0.976	0.937	0.991
Viral Centrality, with Π	0.991	0.989	0.966	0.998
Dense, $M2b$	0.953	0.971	0.970	0.963
Communicability (out)	0.959	0.958	0.952	0.963
Katz (out)	0.914	0.951	0.957	0.927
Viral Centrality, network weights	0.904	0.939	0.954	0.924
Ghosh, $M0_G$	0.920	0.925	0.909	0.921
First-order retention, $r_j g_j$	0.934	0.892	0.850	0.921
Eigenvector (out)	0.858	0.910	0.922	0.876
CheiRank	0.629	0.659	0.626	0.646
Out-strength	0.643	0.626	0.566	0.639
Out-degree	0.012	0.023	-0.001	0.006
PageRank	-0.157	-0.155	-0.204	-0.150
<i>$M3_{mf}$ minus comparator, 95% bootstrap interval</i>				
Best score without Π	0.037	0.018	-0.020	0.027
	[0.019, 0.077]	[-0.008, 0.050]	[-0.087, 0.033]	[0.009, 0.063]
Viral Centrality, with Π	0.006	-0.013	-0.028	-0.007
	[-0.001, 0.020]	[-0.042, 0.004]	[-0.070, -0.001]	[-0.023, 0.002]

Notes: Each column is a simulated ground truth; each row is a score computed from the network alone. Bold marks the best predictor. The independent cascade uses transmission probabilities $P_{ji} = \min\{1, \Omega_{\text{eff}} G_{ji} \pi_{ji}\}$. Viral Centrality is computed by node-level probability propagation, $x_i \leftarrow 1 - \prod_j (1 - P_{ji} x_j)$ with the seed fixed at one, either on this P (with Π) or on $P_{ji} = \min\{1, \Omega_{\text{eff}} G_{ji}\}$ (network weights). “Best score without Π ” is the best of the activation-free centralities other than the paper’s benchmarks $M2b$ and $M0_G$: communicability in all columns except $\zeta = 40$, where it is Katz. The circular comparison with the Monte Carlo mean of $M3$ is omitted.


Table 6 yields four results. First, in the linear and moderate-capacity regimes, no activation-free centrality matches $M3_{mf}$. On the tail outcome, $M3_{mf}$ reaches 0.997, against 0.959 for communicability, 0.953 for $M2b$, and 0.914 for Katz, with the bootstrap interval of the difference from the best activation-free score excluding zero. The advantage also holds for the independent cascade, a distinct simulated dynamic not governed by our linear recursion. Second, the advantage comes from activation information not from a

particular propagation formula. Viral Centrality computed on network weights performs at the level of Katz (0.904–0.954), whereas incorporating Π makes it statistically indistinguishable from $M3_{mf}$ in three of four outcomes and better under strongly binding capacity. The predictive gain therefore comes from the activation structure, which improves even alternative centrality formulations; what the effective operator adds is an exact, closed-form characterization of that structure, which is what makes the decomposition in E4 and the cross-network evidence in E6 and E7 possible. Third, under strongly binding capacity ($\zeta = 40$), the ranking reverses: $M2b$ performs best and $M3_{mf}$ falls to 0.937. Consistent with E3, high-retention sectors such as coffee and livestock experience greater compression of their concentrated cascades when bottlenecks bind. The linear effective operator is therefore most informative under unconstrained or moderately constrained propagation. Fourth, the first-order retention score $r_j g_j$, requiring no resolvent, already performs comparably to the Ghosh benchmark on linear outcomes (0.921–0.934), underscoring the importance of emitter-level retention.

6.3 E6 — What governs the magnitude of reordering

The fourth question concerns generality. The synthetic nulls in Table 12 (Appendix B) show that reordering does not require the Colombian topology, but they do not identify which property of a network controls its magnitude. We therefore generate dense networks in which two candidate properties are varied independently. Each synthetic network preserves, for every emitter, its empirical intersectoral sales and gross output, so the row sums g_j of $G = B_G^\circ$ are identical to those of the Colombian table. Destinations are drawn as a Bernoulli graph with the empirical density 0.937, and within-row weights are drawn from a log-normal distribution with dispersion $\sigma \in \{0.25, 0.5, 1, 1.5, 2, 2.5, 3\}$, which controls within-row concentration. Independently, δ is set so that the dense operator lies at distance to criticality $c = \rho(A^*)/(1 - \delta) \in \{0.30, 0.50, 0.70, 0.86, 0.95\}$, holding Ω fixed. By Corollary 2, this spans the relevant range of effective attenuation. In every network the activation scale is recalibrated with the baseline rule, so that the mean expected number of active links equals that network’s median n^{eff} . Each (σ, c) cell averages 20 independent networks, and every network is evaluated at all five values of c .

Figure 6: Phase diagram of sparse–dense reordering



Notes: Left: mean absolute rank change between $M2b$ and $M3_{mf}$. Right: share of that reordering reproduced by the retention-only operator $D_r G$. Columns vary within-row weight dispersion σ (realized mean n_j^{eff}/d_j in parentheses); rows vary the distance to criticality of the dense operator. Each cell averages 20 synthetic networks that preserve the empirical row sums of B_G° and gross output. The star locates the Colombian input–output network.

Table 7: E6: Single network statistics as predictors of reordering

Statistic	OLS R^2	Spearman	Slope
Dispersion of $\log r_j$ across emitters	0.920	0.949	19.26
Coefficient of variation of r_j	0.912	0.949	17.93
Median n_j^{eff}	0.860	-0.935	-0.042
Mean concentration n_j^{eff}/d_j	0.856	-0.930	-2.63
Spectrally matched scale p^*	0.755	-0.806	-6.49
Distance to criticality, sparse operator	0.135	-0.330	-1.75
Distance to criticality, dense operator	0.000	-0.011	-0.01
Dispersion of $\log r_j \times$ dense criticality (with interaction)	0.920	—	—

Notes: Univariate OLS of the mean absolute rank change between $M2b$ and $M3_{mf}$ on each statistic, pooling the 700 synthetic network–criticality evaluations. The last row adds the dense distance to criticality and its interaction with the dispersion of $\log r_j$.

Figure 6 and Table 7 deliver a sharp answer. Reordering rises monotonically with within-row concentration, from about 0.5 positions at $\sigma = 0.25$ to about 3 positions at $\sigma = 3$, and is essentially flat along the criticality axis: at $\sigma = 2$ it ranges between 2.30 and 2.41 across all five distances to criticality. A single statistic, the cross-emitter dispersion of $\log r_j$, explains 92% of the variation in reordering across all synthetic networks. The dense distance to criticality explains none of it, and adding it with an interaction does not improve the fit. This is the network-level counterpart of the E4 result: the magnitude of sparse reordering is governed by how unevenly activation retains transmission across emitters, not by how far the operator lies from criticality. The right panel adds a qualification. The share of the reordering that retention alone reproduces is close to 0.9 when weights are nearly uniform, at any distance to criticality. With concentrated weights it declines as c approaches one, from 0.76 to 0.39 at $\sigma = 2$, because long paths give more room to the within-row reallocation $\widehat{G}$.

Colombia sits outside this synthetic family in an informative way. Its dispersion of $\log r_j$ is 0.178, well inside the range of synthetic draws, for which the fitted relationship predicts a mean rank change of 3.68. The observed value is 4.45, above every one of the 700 synthetic evaluations, whose maximum is 4.39. Since the synthetic networks preserve row sums and within-row concentration but assign weights to destinations at random, the excess reflects the empirical matching between large transactions and particular downstream buyers. This is consistent with the weight-shuffle null in Table 12, which breaks exactly that matching and lowers the mean rank change to 3.45. Within-row concentration sets the baseline magnitude of sparse reordering, and the alignment of strong links with specific downstream paths amplifies it. E7 turns to empirical tables, where weights are not assigned at random.

6.4 E7 — Cross-country evidence

The Colombian table is a single network. To test whether the mechanism is a regularity of production networks rather than a feature of one table, we apply the identical pipeline to the 2021 tables of the OECD Input–Output Tables, 2025 edition, for the 79 economies available other than Colombia (OECD, 2025). The tables share a common classification of 50 industries; industries with zero output or no intermediate transactions are dropped, leaving 49 in every economy. Each table is recalibrated with the rules of Section 5: Ω equals the aggregate intermediate share, δ implies a diffusion half-life of eight rounds, and the activation scale matches the median n^{eff} of that economy. The main specification uses total intermediate flows, domestic plus imported inputs, which is the convention of the Colombian table. Allocation coefficients are computed over total supply (output plus imports), so that no row of B_G exceeds one. Results with domestic flows only are reported in Appendix B.8. Because the number of sectors may differ across tables, reordering is normalized as mean

$|\Delta r|/n$. The OECD tables share the stylized facts of the Colombian one: they are nearly complete (median density 0.996) but concentrated, with a median n^{eff} of 8.8 among 48 possible destinations.


Table 8: Sparse reordering and its channels across OECD input–output tables

Median across economies	All ($N = 79$)	Core sample ($N = 20$)	Advanced ($N = 37$)	Emerging ($N = 42$)
Sectors, n	49	49	49	49
Density	1.00	1.00	0.98	1.00
Median n^{eff}	8.8	9.2	8.9	8.5
Spearman $M2b$ vs. $M0_G$	0.95	0.92	0.94	0.95
Spearman $M2b$ vs. $M3_{mf}$	0.981	0.977	0.979	0.983
Mean $ \Delta r /n$	0.039	0.042	0.041	0.035
Gap closed, p^*A^*	0.05	0.07	0.06	0.05
Gap closed, D_rG	0.63	0.59	0.62	0.66
Gap closed, best backbone	0.15	0.15	0.09	0.24
Log-ratio R^2 , D_rG	0.98	0.98	0.98	0.98
Gap closed, D_rG , homogeneous α	0.82	0.81	0.82	0.82
Spearman(r_j, n_j^{eff})	−0.13	−0.18	−0.05	−0.19

Notes: Intermediate flows from the OECD Input–Output Tables (2025 edition): domestic plus imported inputs (total flows), as in the Colombian table, with allocation coefficients computed over total supply (output plus imports) unless stated otherwise. Each country is calibrated with the rules of Section 5. Medians across countries. “Core sample” is a representative subset of 20 economies (USA, DEU, JPN, GBR, FRA, KOR, AUS, ESP, BRA, MEX, ARG, CHL, PER, CRI, CHN, IND, IDN, VNM, ZAF, TUR); “Emerging” includes emerging and developing economies (IMF classification). Gap closed is defined as in Table 5; the best backbone is the larger of the ρ -matched top- k and matched-sparsity disparity backbones.

Table 8 and Figure 7 show that the mechanism is the same in every economy, even though its magnitude varies. Sparse activation reorders the OECD tables less than the Colombian one. The median Spearman correlation between $M2b$ and $M3_{mf}$ is 0.981 (interquartile range 0.974–0.986), and the median normalized rank change is 0.039, against 0.066 for Colombia, which lies above every OECD economy (maximum 0.064). The Colombian table is more disaggregated (67 against 49 sectors), and its dispersion of retention factors (0.178) exceeds that of every OECD table (maximum 0.167). The decomposition, however, is remarkably stable. The retention channel reproduces a median 63% of the reordering (interquartile range 55–71%), more than half in 89% of economies, and a median 98% of the variation in log scores (minimum 93%). Uniform attenuation reproduces a median 5%, and the best ρ -matched backbone 15%. Retention exceeds both alternatives in all 79 economies. Under the baseline profile, the median correlation between r_j and n_j^{eff} is only −0.13, because α_j rises with sector size and effective diversification and partly offsets the retention penalty, as in the Colombian table (−0.48). With homogeneous intensity, this correlation reaches −0.94 and retention reproduces a median 82% of the reordering, so the cross-country evidence also rules out that the mechanism is an artifact of the baseline intensity profile.

Figure 7: Channels of sparse reordering across 79 OECD input–output tables



Notes: Each dot is an economy; the vertical bar marks the median. The horizontal axis reports the share of the $M2b \rightarrow M3_{mf}$ reordering reproduced by each operator, as defined in Table 5. Total intermediate flows over total supply, 2021. The last row recomputes $M3_{mf}$ under homogeneous activation intensity.

The magnitude of reordering across economies is only partly predictable from single network statistics. The dispersion of $\log r_j$ explains 13% of the cross-country variation in normalized reordering and the dense distance to criticality 19%; since the two are uncorrelated (-0.04), jointly they explain 33%. This qualifies the synthetic result of E6, where criticality played no role. When weights are assigned to destinations at random, within-row concentration is the only systematic source of variation; in empirical tables, differences in the calibrated attenuation and in the structure of long paths also matter. The channel through which sparse activation reorders influence is therefore universal, whereas its magnitude depends on more than one network property.

Table 9: E7: Industries that move systematically under sparse activation across countries

Code	Industry	Mean $\Delta r/n$	Share up	Share down	Median n^{eff}
<i>Largest average gains</i>					
A01	Crop and animal production	0.137	0.96	0.01	1.5
G	Wholesale and retail trade	0.066	0.90	0.01	14.5
A03	Fishing and aquaculture	0.062	0.71	0.09	1.8
C23	Other non-metallic mineral products	0.051	0.70	0.20	2.1
C10T12	Food, beverages, and tobacco	0.048	0.73	0.00	3.7
C24A	Basic iron and steel	0.039	0.59	0.16	4.9
<i>Largest average losses</i>					
H53	Postal and courier activities	-0.049	0.10	0.67	7.7
C16	Wood and wood products	-0.041	0.22	0.63	4.1
H51	Air transport	-0.037	0.13	0.68	15.6
A02	Forestry and logging	-0.033	0.20	0.59	2.9
H50	Water transport	-0.032	0.16	0.61	9.9
E	Water supply, sewerage, and waste	-0.032	0.16	0.73	12.3

Notes: $\Delta r = r^{M2b} - r^{M3mf}$ normalized by the number of sectors in each country; positive values indicate upward movement under sparse activation. Shares are computed over the countries in which the industry is present (79 for most codes). Total flows over total supply.

Because the OECD tables share a classification, we can also ask which industries move consistently. Table 9 shows a sharp pattern. Crop and animal production rises under sparse activation in 96% of economies and falls in only 1%; trade, fishing, food products, and non-metallic minerals also gain in most economies. Water, air, and postal transport, wood products, forestry, and water supply and waste lose positions in roughly two thirds of economies. This mirrors the Colombian results, where livestock, coffee, and non-metallic minerals rise while air and water transport and forestry fall. Most gains concentrate in primary and processing industries with few dominant buyers (median n^{eff} between 1.5 and 3.7), which keep most of their transmission when activation is sparse, while the losses include transport and utility services whose sales are spread across many users (median n^{eff} between 7.7 and 15.6), as the retention mechanism of Proposition 3 predicts. Two patterns depart from this reading. Wholesale and retail trade rises despite diversified sales (median $n^{\text{eff}} = 14.5$) because its scale assigns it a high activation intensity under the baseline profile, which offsets the diversification penalty in r_j . Forestry and wood products lose positions despite concentrated sales, which suggests that their dense centrality runs through the downstream paths of their buyers rather than through their own retention. The complete country-level statistics are reported in Appendix H.

7 Robustness

Appendix B reports the full set of robustness exercises; we summarize them here. The sparse–dense reordering and the retention mechanism survive every variation we consider.

- (i) *Activation intensity* (Table 10). Reordering persists, and grows, when α_j is homogeneous across emitters (mean $|\Delta r| = 7.01$) or depends only on size (7.58), so it is not generated by the baseline intensity profile.
- (ii) *Activation kernel* (Table 11). Holding each emitter’s expected number of active links fixed, weighted kernels give Spearman correlations with $M2b$ between 0.916 and 0.964; uniform and proximity-only activation reorder more.
- (iii) *Network structure* (Table 12). Directed Erdős–Rényi nulls with matched density and sparsity produce reordering of the same magnitude (about 4.6 positions), while shuffling weights within rows lowers it to 3.45.
- (iv) *Leave one sector out* (Table 13). The mean rank change ranges between 3.42 and 4.94 across all 67

exclusions.

- (v) *Persistence and horizon* (Table 14). Spearman correlations remain between 0.922 and 0.951 across half-lives of 4 to 12 rounds and horizons of 10 to 100.
- (vi) *Bounded allocation coefficients* (Appendix B.6). Bounding the four rows of B_G° whose intermediate sales exceed domestic output leaves the results unchanged: Spearman 0.939, and retention reproduces 62% of the reordering.
- (vii) *Decomposition under alternative activation rules* (Table 15). Retention remains the dominant channel, and becomes stronger with homogeneous intensity.
- (viii) *Domestic flows in the OECD tables* (Table 16). With domestic intermediate flows only, retention reproduces a median 67% of the reordering and exceeds attenuation and every backbone in all 79 economies.

Appendix E2 and Appendix E3 report the convergence of the Monte Carlo implementation and the activation-intensity transition predicted by Proposition 4.

8 Conclusion

This paper develops a stochastic propagation operator for weighted directed networks in which transmission occurs through a random subset of available links, shocks may persist across rounds, and receiving nodes may face finite absorption capacity. The framework has an exact dense full-activation limit and, in the linear no-capacity regime, expected diffusion is governed by a closed-form effective operator. This operator factorizes into emitter-level retention factors, bounded by the inverse effective number of destinations, and a mass-preserving reallocation of transmission within each emitter's links. This makes it possible to compare sparse and dense propagation analytically while separating the effects of heterogeneous link activation from those generated by nonlinear capacity constraints.

In the Colombian input–output application, sparse activation preserves the broad supply-side orientation while substantially reorganizing influence among the sectors that are most central under dense propagation. The reordering is systematic: sectors with greater intermediate-sales scale tend to move upward, whereas broader effective diversification and greater dependence on longer downstream paths are associated with rank losses. The decomposition of the effective operator identifies the mechanism. Emitter-level retention, which penalizes diversified emitters in proportion to their effective number of destinations, reproduces most of the reordering, whereas uniform attenuation to the same distance to criticality reproduces almost none, and dense resolvents on deterministic backbones reproduce at most 29%. The effective operator also predicts simulated cascade reach better than standard resolvent, spectral, and random-walk centralities outside the strongly capacity-constrained regime. Across synthetic dense networks, the dispersion of retention factors governs the magnitude of reordering, while the empirical alignment of strong links with specific downstream buyers amplifies it in the Colombian table. Across 79 OECD input–output tables, the decomposition is the same in every economy: retention reproduces a median 63% of the reordering, attenuation and backbones reproduce little, and the industries promoted and demoted by sparse activation coincide across countries. The magnitude of reordering varies across economies and is only partly predictable from single network statistics, whereas the channel that produces it does not vary. The sparse–dense divergence is also robust to alternative activation kernels, synthetic network nulls, and sector exclusions. Capacity is largely inactive around the baseline regime; the stress exercise shows that, once normalized shocks become sufficiently large relative to receiving capacity for downstream bottlenecks to bind, propagation is compressed non-uniformly and sectoral rankings begin to separate from their no-capacity counterpart.

These results are especially relevant for networks that are topologically dense but effectively concentrated. Dense resolvent measures treat all observed channels as continuously available, whereas sparse activation allows realized influence to depend on which links are actually engaged in propagation. More generally, when

node importance is defined by the ability to generate downstream influence, the resulting ranking depends not only on network structure, but also on the propagation mechanism operating on that structure. In topologically dense but effectively concentrated networks, retention factors, which require no simulation, identify the nodes whose influence dense centralities overstate.

9 Limitations

The framework is a reduced-form model of propagation on weighted directed networks. It does not include prices, substitution, inventories, optimizing agents, market clearing, or welfare, so the scores measure network influence rather than structural output losses or policy multipliers. The activation process and the capacity rule are calibrated, not estimated, and the same rules are applied to every economy; the high-stress capacity exercise therefore characterizes the nonlinear regime of the model rather than empirically plausible shock magnitudes.

The empirical evidence covers a single year (2021) and a single class of networks. The Colombian and OECD tables also differ in aggregation (67 sectors against 49 industries), which limits direct comparison of reordering magnitudes between them. Although the theoretical results hold for any nonnegative weighted directed network, whether retention also dominates sparse reordering in other dense but concentrated networks, such as trade, financial, or mobility networks, remains open.

The phase diagram assigns synthetic weights to destinations at random, so it characterizes concentration but not the alignment of strong links with particular buyers, and single network statistics explain only part of the cross-country variation in reordering. In the centrality benchmark, all ground truths are simulated and the independent cascade includes Π , so agreement with it supports the relevance of the activation structure but is not validation against observed shock propagation. Finally, dense measures outperform the effective operator once capacity binds strongly, so it should be read as a centrality for the unconstrained and moderately constrained regimes.

Data Availability

The data used in this study are publicly available from the Departamento Administrativo Nacional de Estadística (DANE), Colombia. Specifically, the analysis uses the 2021 symmetric input–output matrix available through DANE’s National Accounts complementary matrices. The cross-country analysis uses the OECD Input–Output Tables, 2025 edition (domestic output and imports tables), publicly available from the OECD (OECD, 2025).

A Proofs

This appendix collects the proofs of the results in Section 3. Throughout, matrix and vector inequalities are entrywise, $A_G^\star = \Omega G^\top$, and $\mathcal{R}_H(M) = \sum_{t=1}^H M(\delta I + M)^{t-1}$.

A.1 Proof of Proposition 1

Under full deterministic activation $I(t) \equiv \mathbf{1}_E$, the support convention for G implies $G \odot \mathbf{1}_E = G$, and therefore $A(t) = \Omega G^\top = A_G^\star$. With $\sigma_C = \text{id}$,

$$S(t) = (\delta I + A_G^\star)S(t-1), \quad P(t) = A_G^\star S(t-1).$$

Hence $S(t) = (\delta I + A_G^\star)^t s_0$ and $P(t) = A_G^\star (\delta I + A_G^\star)^{t-1} s_0$. Summing over a finite horizon gives

$$\sum_{t=1}^H P(t) = \sum_{t=1}^H A_G^\star (\delta I + A_G^\star)^{t-1} s_0 = \mathcal{R}_H(A_G^\star) s_0,$$

which proves (5) without requiring a stability condition.

If $\rho(A_G^\star) < 1 - \delta$, then $\rho(\delta I + A_G^\star) = \delta + \rho(A_G^\star) < 1$ because $A_G^\star$ is nonnegative. The corresponding Neumann series converges, yielding

$$\sum_{t \geq 1} P(t) = A_G^\star \sum_{k \geq 0} (\delta I + A_G^\star)^k s_0 = A_G^\star ((1 - \delta)I - A_G^\star)^{-1} s_0,$$

which is (6). For $\delta = 0$ and $\Omega = 1$, $A_G^\star = G^\top$ and $G^\top(I - G^\top)^{-1} = (I - G^\top)^{-1} - I$. Thus cumulative spillovers are $((I - G^\top)^{-1} - I)s_0$. Adding the initial unit impulse gives the column- s sum of $(I - G^\top)^{-1}$; since $(I - G^\top)^{-1} = ((I - G)^{-1})^\top$, this equals the row- s sum of $(I - G)^{-1}$. $\square$

A.2 Proof of Corollary 1

Set $\delta = 0$, $\Omega = 1$, and take the infinite-horizon limit. By Proposition 1, the total score associated with a base matrix G is the row- s sum of $(I - G)^{-1}$.

(i) *Leontief orientation.* Take $G = (A^\circ)^\top$. Then

$$(I - G)^{-1} = (I - (A^\circ)^\top)^{-1} = ((I - A^\circ)^{-1})^\top,$$

so the resulting score is the column- s sum of $(I - A^\circ)^{-1}$. This is the off-diagonal demand-side resolvent associated with the Leontief orientation.

(ii) *Ghosh orientation.* Take $G = B_G^\circ$. The resulting score is the row- s sum of $(I - B_G^\circ)^{-1}$, the off-diagonal supply-side resolvent associated with the Ghosh orientation.

The standard Leontief and Ghosh benchmarks used empirically are instead constructed from the full matrices A and B_G . Hence the corollary establishes orientation and the exact dense endpoint of the propagation family, without requiring the stochastic network to retain within-sector self-links. $\square$

A.3 Proof of Proposition 2

Linearity of the mean. Without capacity, $P(t) = A(t)S(t-1)$ and $S(t) = (\delta I + A(t))S(t-1)$. Since $S(t-1)$ depends only on $s_0, I(1), \dots, I(t-1)$ and $I(t)$ is independent across rounds, $A(t)$ is independent of $S(t-1)$. Therefore

$$\mathbb{E}[A(t)S(t-1)] = \mathbb{E}[A(t)] \mathbb{E}[S(t-1)] = \bar{B} \mathbb{E}[S(t-1)], \quad \bar{B} = \Omega(G \odot \Pi)^\top.$$

This yields (7). Iterating, $\mathbb{E}[S(t)] = (\delta I + \bar{B})^t s_0$ and $\mathbb{E}[P(t)] = \bar{B}(\delta I + \bar{B})^{t-1} s_0$, so for finite H , $\sum_{t=1}^H \mathbb{E}[P(t)] = \mathcal{R}_H(\bar{B}) s_0$. If $\rho(\bar{B}) < 1 - \delta$, the same Neumann-series argument used above gives $\mathcal{R}_\infty(\bar{B}) = \bar{B}((1 - \delta)I - \bar{B})^{-1}$.

(i) *Domination*. Since $0 \leq \pi_{ji} \leq 1$ on the support of G , $\bar{B} = \Omega(G \odot \Pi)^\top \leq \Omega G^\top = A^\star$. Both matrices are nonnegative, hence Perron–Frobenius monotonicity gives $\rho(\bar{B}) \leq \rho(A^\star)$.

(ii) *Envelope*. From $0 \leq \bar{B} \leq A^\star$, induction gives $\bar{B}^r \leq (A^\star)^r$ for every integer $r \geq 0$. Moreover,

$$M(\delta I + M)^{t-1} = \sum_{r=1}^t \binom{t-1}{r-1} \delta^{t-r} M^r,$$

which is a nonnegative polynomial in M . Applying the power inequality term by term therefore gives $\bar{B}(\delta I + \bar{B})^{t-1} \leq A^\star(\delta I + A^\star)^{t-1}$ for every t , and summing to H yields $\mathcal{R}_H(\bar{B}) \leq \mathcal{R}_H(A^\star)$. If $\rho(A^\star) < 1 - \delta$, then part (i) implies $\rho(\bar{B}) < 1 - \delta$ as well, and taking $H \rightarrow \infty$ preserves the entrywise inequality, producing the stated resolvent envelope. Adding the common initial unit impulse preserves the score ordering.

(iii) *Structural divergence*. For fixed G , Ω , δ , and H , the finite-horizon expected score vector is $1 + \mathbf{1}^\top \mathcal{R}_H(\Omega(G \odot \Pi)^\top)$ and is therefore a deterministic function of Π . The infinite-horizon expression has the same property whenever it exists. Monte Carlo affects only estimation of this fixed population object. $\square$

A.4 Proof of Corollary 2

(i) The eigenvalues of $\delta I + M$ are $\delta + \mu$, with μ ranging over the eigenvalues of M . Hence $|\delta + \mu| \leq \delta + |\mu| \leq \delta + \rho(M)$. By the Perron–Frobenius theorem, $\rho(M)$ is itself an eigenvalue of $M \geq 0$, so the bound is attained and $\rho(\delta I + M) = \delta + \rho(M)$.

(ii) If $\rho(M) < 1 - \delta$, then $\rho(K) < 1$, so $I - K$ is invertible and

$$M((1 - \delta)I - M)^{-1} = (1 - \delta)K((1 - \delta)(I - K))^{-1} = K(I - K)^{-1}.$$

For $M = \Omega(G \odot \Pi)^\top$, $K = \Omega_{\text{eff}}(G \odot \Pi)^\top$, so the infinite-horizon operator, and hence the score vector, depends on (Ω, δ) only through Ω_{eff} . $\square$

A.5 Proof of Proposition 3

(i) If $r_j > 0$, then $\sum_i \hat{G}_{ji} = r_j^{-1} \sum_i G_{ji} \pi_{ji} = g_j$ and $r_j \hat{G}_{ji} = G_{ji} \pi_{ji}$. If $r_j = 0$, then $G_{ji} \pi_{ji} = 0$ for every i , so row j of $G \odot \Pi$ vanishes and equals row j of $D_r \hat{G}$. Thus $G \odot \Pi = D_r \hat{G}$, and transposing gives $\bar{B} = \Omega(D_r \hat{G})^\top = \Omega \hat{G}^\top D_r$.

(ii) Under the stated conditions, $r_j = \sum_i \tilde{\omega}_{ji} f(\tilde{\omega}_{ji})$ with $f(x) := 1 - e^{-\alpha_j x}$. Let X be a random variable taking the value $\tilde{\omega}_{ji}$ with probability $\tilde{\omega}_{ji}$. Then $r_j = \mathbb{E}[f(X)]$ and $\mathbb{E}[X] = \sum_i \tilde{\omega}_{ji}^2 = 1/n_j^{\text{eff}}$. Since f is concave, Jensen's inequality gives $r_j \leq f(1/n_j^{\text{eff}}) = 1 - e^{-\alpha_j/n_j^{\text{eff}}}$, and $1 - e^{-y} \leq y$ gives the last inequality. For the lower bound, $1 - e^{-y} \geq y - y^2/2$ for $y \geq 0$ (the difference vanishes at $y = 0$ and has derivative $e^{-y} - 1 + y \geq 0$). Hence

$$r_j \geq \alpha_j \sum_i \tilde{\omega}_{ji}^2 - \frac{\alpha_j^2}{2} \sum_i \tilde{\omega}_{ji}^3 \geq \frac{\alpha_j}{n_j^{\text{eff}}} - \frac{\alpha_j^2 \bar{\omega}_j}{2n_j^{\text{eff}}},$$

using $\sum_i \tilde{\omega}_{ji}^3 \leq \bar{\omega}_j \sum_i \tilde{\omega}_{ji}^2$. Monotonicity in α_j holds termwise.

(iii) Write $r = \sum_i h(\tilde{\omega}_i)$ with $h(x) := x(1 - e^{-\alpha x})$. Then $h''(x) = \alpha e^{-\alpha x}(2 - \alpha x) \geq 0$ on $[0, 2/\alpha]$, so h is convex on this interval. If $\tilde{\omega}_j$ majorizes $\tilde{\omega}_k$, then $\bar{\omega}_k \leq \bar{\omega}_j \leq 2/\alpha$, so both vectors lie in $[0, 2/\alpha]^n$. By Karamata's inequality, $\sum_i h(\tilde{\omega}_{ji}) \geq \sum_i h(\tilde{\omega}_{ki})$, i.e. $r_j \geq r_k$. $\square$

A.6 Proof of Proposition 4

(i) Each entry $\bar{B}(\theta)_{ij} = \Omega G_{ji} \pi_{ji}(\theta)$ is a nonnegative constant times a continuous nondecreasing function of θ . Hence $\bar{B}(\theta)$ is entrywise continuous and nondecreasing. The boundary conditions give $\bar{B}(0) = 0$, while $\pi_{ji}(\theta) \rightarrow 1$ on the support of G implies $\bar{B}(\theta) \uparrow \Omega G^\top = A^\star$.

(ii) If $\theta \leq \theta'$, then $0 \leq \bar{B}(\theta) \leq \bar{B}(\theta')$, so Perron–Frobenius monotonicity gives $\rho(\bar{B}(\theta)) \leq \rho(\bar{B}(\theta'))$. Continuity follows from continuity of the spectral radius as a function of the matrix entries.

(iii) For finite H , every entry of $\mathcal{R}_H(\bar{B}(\theta))$ is a finite polynomial in the entries of $\bar{B}(\theta)$ with nonnegative coefficients. The same power-ordering argument as in Proposition 2 therefore implies that $\mathcal{R}_H(\bar{B}(\theta))$ is entrywise nondecreasing in θ , while continuity follows from the finite polynomial representation. Since $\bar{B}(\theta) \rightarrow A^\star$, continuity also gives $\mathcal{R}_H(\bar{B}(\theta)) \rightarrow \mathcal{R}_H(A^\star)$.

If $\rho(A^\star) < 1 - \delta$, domination implies $\rho(\bar{B}(\theta)) < 1 - \delta$ for every finite θ , and continuity of the resolvent on this stable set yields

$$\bar{B}(\theta)((1 - \delta)I - \bar{B}(\theta))^{-1} \longrightarrow A^\star((1 - \delta)I - A^\star)^{-1}.$$

Thus the expected score vector converges continuously to its dense endpoint. Rankings themselves are discrete: they can change only when two score paths cross. If the dense endpoint has distinct sectoral scores, its minimum pairwise score gap is positive, so continuity implies that for all sufficiently large θ the ordering coincides with the dense ranking. $\square$

A.7 Proof of Proposition 5

The hard cap $\sigma_C(x)_i = \min\{x_i, C_i\}$ and the smooth saturation $\sigma_C(x)_i = C_i(1 - e^{-x_i/C_i})$ are nondecreasing, concave, 1-Lipschitz, satisfy $\sigma_C(0) = 0$, and obey $0 \leq \sigma_C(x)_i \leq \min\{x_i, C_i\}$. The proof of parts (i)–(ii) uses only the admissibility properties in (3).

(i) *Global boundedness.* Admissibility gives $0 \leq P_i(t) = \sigma_C(\hat{P}(t))_i \leq C_i$. Hence $S_i(t) = \delta S_i(t-1) + P_i(t) \leq \delta S_i(t-1) + C_i$, and iteration yields

$$S_i(t) \leq \delta^t S_i(0) + C_i \sum_{r=0}^{t-1} \delta^r, \quad \limsup_{t \rightarrow \infty} S_i(t) \leq \frac{C_i}{1 - \delta} \quad (\delta < 1).$$

No spectral property of $A^\star$ is used, so this bound remains valid even when the corresponding dense linear recursion is unstable.

(ii) *Loss of scale invariance.* Fix a realization of $\{I(t)\}$, a direction $v \geq 0$, and set $S(0; c) = cv$. We show by induction that every $c \mapsto S_i(t; c)$ is nondecreasing and concave. The claim is immediate at $t = 0$. If it holds at $t - 1$, then $\hat{P}_i(t; c) = \sum_j A(t)_{ij} S_j(t-1; c)$ is a nonnegative linear combination of nondecreasing concave functions and therefore has the same properties. Since $\sigma_C(\cdot)_i$ is itself nondecreasing and concave, $P_i(t; c) = \sigma_C(\hat{P}(t; c))_i$ is nondecreasing and concave, and so is $S_i(t; c) = \delta S_i(t-1; c) + P_i(t; c)$. This closes the induction.

For any finite H , both $\sum_{t=1}^H \mathbf{1}^\top P(t; c)$ and $\mathcal{T}_H(c) = c \mathbf{1}^\top v + \sum_{t=1}^H \mathbf{1}^\top P(t; c)$ are therefore nondecreasing and concave, with $\mathcal{T}_H(0) = 0$. For every nonnegative concave function f with $f(0) = 0$, the secant slope $f(c)/c$ is nonincreasing on $c > 0$; hence $\mathcal{T}_H(c)/(c \mathbf{1}^\top v)$ is nonincreasing. Strict decline follows whenever the pathwise total is strictly concave over the relevant interval, which requires saturation to generate a strict departure from proportional propagation. Expectations preserve monotonicity and concavity. If the infinite-horizon cumulative total is finite, the same conclusions follow by taking the monotone limit $H \rightarrow \infty$.

(iii) *Normalized stress intensity.* For either implemented capacity map,

$$\sigma_{\kappa x}(\kappa y) = \kappa \sigma_x(y).$$

Let $C = \kappa x$, $S(0) = qv$, and define $\tilde{S}(t) := S(t)/\kappa$ and $\tilde{P}(t) := P(t)/\kappa$. Writing $\zeta = q/\kappa$, the initial state becomes $\tilde{S}(0) = \zeta v$. Since $\tilde{P}(t) = A(t)S(t-1) = \kappa A(t)\tilde{S}(t-1)$, joint homogeneity of the capacity map gives

$$\tilde{P}(t) = \sigma_x \left(A(t)\tilde{S}(t-1) \right), \quad \tilde{S}(t) = \delta \tilde{S}(t-1) + \tilde{P}(t).$$

Thus the normalized recursion is independent of q and κ separately and depends on them only through ζ . Consequently $\mathcal{T}_H(q, \kappa) = \kappa \tilde{\mathcal{T}}_H(\zeta)$, while $\|s_0\|_1 = q\|v\|_1 = \kappa\zeta\|v\|_1$, so

$$\frac{\mathcal{T}_H(q, \kappa)}{\|s_0\|_1} = \frac{\tilde{\mathcal{T}}_H(\zeta)}{\zeta\|v\|_1} =: \Psi_H(\zeta; v, x).$$

This proves that the shock-to-capacity ratio $\zeta = q/\kappa$ is the sufficient scale index for the capacity experiment, holding the network and all remaining parameters fixed. $\square$

B Robustness

This appendix reports the robustness exercises summarized in Section 7. Unless stated otherwise, each exercise holds the baseline calibration fixed and compares the sparse mean-field operator $M3_{mf}$ with its dense counterpart $M2b$.

B.1 Activation-intensity heterogeneity

The rank-change accounting exercise in Exercise E2 uses the baseline construction of α_j , which varies across emitters with sector size and effective diversification. To assess whether the resulting reordering is mechanically driven by this cross-sector heterogeneity, we recompute the sparse operator under alternative activation-intensity rules. In each case, the mean expected number of active links per emitter is held fixed, so the exercise changes only how activation intensity varies across sectors.

Table 10: Reordering under alternative activation-intensity rules

Activation rule	Spearman vs. $M2b$	Kendall vs. $M2b$	Mean $ \Delta r $	Max $ \Delta r $
Baseline	0.945	0.825	4.45	19
Homogeneous intensity	0.885	0.719	7.01	26
Size only	0.870	0.698	7.58	30
Effective degree only	0.982	0.907	2.33	14

Notes: Baseline uses the calibrated intensity profile in which α_j depends on both emitter size and effective diversification. Homogeneous intensity sets the same α_j for every emitter. Size only lets α_j vary only with total intermediate sales, while Effective degree only lets it vary only with n_j^{eff} , the inverse-Herfindahl effective number of outgoing destinations. All variants are rescaled to match the baseline mean expected number of active links, so they differ only in the cross-sector distribution of activation intensity.

Substantial reordering persists under homogeneous activation and becomes stronger when intensity depends only on sector size. Hence, the departure from $M2b$ is not generated solely by the baseline construction of α_j . It also reflects the bilateral interaction between transmission weights and activation probabilities in $\tilde{B} = \Omega(B_G^\circ \odot \Pi)^\top$.

B.2 Alternative activation kernels

The previous exercise varies emitter-level heterogeneity in α_j . Here, we hold each emitter's expected number of active links fixed and vary how activation probability is distributed across its outgoing links, that is,

the within-row structure of Π_{ji} . This directly tests whether the sparse–dense reordering depends on the marked-Poisson activation rule used in the baseline.

The underlying network, transmission matrix, Ω , δ , and horizon T are kept fixed. Every alternative also preserves, separately for each emitting sector, the same expected number of active outgoing links as the baseline. The exercise therefore changes the *allocation* of activation probabilities across links without changing the degree of sparsity.

The alternatives modify how the selection kernel B_{ji} is mapped into activation probabilities. Bernoulli linear uses probabilities proportional to B_{ji} , with row-specific scaling and clipping at one. Power specifications replace B_{ji} by $B_{ji}^{0.5}$ or B_{ji}^2 before calibration: the former flattens differences across links, whereas the latter concentrates activation more strongly on high- B_{ji} links. Uniform assigns the same activation probability to every available outgoing link of an emitter. Finally, weights only and proximity only construct activation using exclusively the transaction-weight component $\tilde{\omega}$ or the productive-proximity component $\tilde{\phi}$ of the baseline selection kernel.

Table 11: Robustness to the activation kernel

Kernel	Spearman	Kendall	Mean $ \Delta r $	Top-20
Marked Poisson	0.945	0.825	4.45	0.70
Bernoulli linear	0.955	0.838	4.06	0.75
Power 0.5	0.916	0.779	5.13	0.75
Power 2	0.964	0.856	3.61	0.80
Weights only	0.941	0.814	4.54	0.70
Uniform	0.732	0.569	9.97	0.65
Proximity only	0.645	0.537	10.72	0.60

Notes: Statistics compare $M3_{mf}$ with the aligned dense benchmark $M2b$. All specifications preserve the baseline expected number of active links for each emitter.

The baseline pattern is stable across the weighted alternatives. Bernoulli linear, both power transformations, and weight-only activation generate rank correlations and mean rank changes close to the marked-Poisson benchmark. Flattening activation completely through the uniform rule does not eliminate reordering; it substantially increases it. Proposition 3 clarifies why: with $\pi_{ji} = k_j/d_j$ constant within row j , the retention factor becomes $r_j = k_j/d_j$, which depends on topological out-degree d_j rather than on the concentration of transactions, so the emitter-level attenuation profile differs sharply from the weighted kernels. The proximity-only specification produces the largest departure, with a Spearman correlation of 0.645 and a mean absolute rank change of 10.72. When activation is driven by a qualitatively different bilateral signal, rankings therefore become substantially more sensitive, consistent with the dependence of the effective operator on the structure of Π .

B.3 Alternative network structures

We next test whether the sparse–dense reordering depends on the particular structure of the Colombian input–output network. For each counterfactual network, the activation scale is recalibrated to match the empirical mean expected number of active links, so differences are not driven mechanically by different sparsity levels. We consider two directed adaptations of the classical random-graph nulls. In $G(n, p)$ (Gilbert, 1959), each possible directed link is included independently with probability equal to the empirical density; in $G(n, m)$ (Erdős and Rényi, 1959), exactly the empirical number of directed links is drawn uniformly across ordered sector pairs. In both cases, each emitter preserves its empirical total intermediate sales, while the identity of its destinations and the bilateral allocation of those sales are randomized.

We also construct a more restrictive MIP weight-shuffle null that preserves the observed network topology.

Each sector keeps exactly the same destinations, out-degree, total intermediate sales, and set of positive bilateral weights as in the empirical matrix. The only change is that those weights are randomly reassigned across the sector’s existing destinations. This preserves essentially all row-level characteristics while breaking the empirical matching between transaction intensity and the identity of the downstream buyer.

Table 12: Robustness to network structure

Network	Spearman	Kendall	Mean $ \Delta r $	Top-20
Empirical MIP	0.945	0.825	4.45	0.70
$G(n, p)$	0.945	0.816	4.62	0.82
$G(n, m)$	0.943	0.815	4.64	0.82
MIP weight shuffle	0.966	0.863	3.45	0.88

Notes: Synthetic-network statistics are averages across 250 realizations. $G(n, p)$ matches the empirical directed density and $G(n, m)$ the empirical number of directed links. Activation intensity is recalibrated in every realization to match the empirical mean expected number of active links.

The two Erdős–Rényi nulls generate almost the same sparse–dense rank divergence as the empirical network, with mean absolute changes of about 4.6 positions versus 4.45 in Colombia. The phenomenon therefore does not require the particular empirical topology. By contrast, preserving the Colombian topology but randomizing the bilateral assignment of weights reduces the mean rank change to 3.45 positions and increases rank concordance. This suggests that the empirical matching between link intensity and downstream destination contributes to the magnitude of the observed reordering. Exercise E6 generalizes these nulls into a phase diagram over within-row concentration and distance to criticality.

B.4 Leave-one-sector-out stability

To assess whether the sparse–dense reordering is driven by a small number of influential sectors, we perform a leave-one-sector-out (LOSO) exercise. For each sector, we remove it as both an emitter and a receiver and reconstruct the dense and sparse propagation operators on the remaining network. Gross output and all baseline propagation parameters are kept fixed, while the activation scale is rescaled only to preserve the baseline mean expected number of active links. Transactions involving the omitted sector are not redistributed across the remaining network.

Table 13: Leave-one-sector-out stability of sparse–dense reordering

	Spearman	Kendall	Mean $ \Delta r $	Top-20
Full network	0.945	0.825	4.45	0.70
LOSO median	0.944	0.824	4.39	0.70
LOSO 5th pct.	0.939	0.812	4.07	0.65
LOSO 95th pct.	0.953	0.838	4.65	0.75
LOSO minimum	0.927	0.800	3.42	0.65
LOSO maximum	0.967	0.868	4.94	0.80

The results are highly stable across sector exclusions. The median LOSO statistics are nearly identical to the full-network estimates, and even the most influential omission leaves substantial sparse–dense reordering. Thus, the baseline result is not driven by any single sector.

B.5 Persistence and horizon sensitivity

We examine whether the sparse–dense reordering depends on the dynamic parameters governing shock persistence and propagation depth. We vary the realized diffusion half-life over $h \in \{4, 8, 12\}$ rounds and the finite propagation horizon over $T \in \{10, 20, 40, 100\}$, holding all remaining calibration parameters fixed. For each half-life, δ follows the baseline calibration rule,

$$\rho(\delta I + \bar{B}) = 2^{-1/h}.$$

Table 14: Sensitivity to persistence and propagation horizon

Half-life h	T	Spearman	Kendall	Mean $ \Delta r $
4	10	0.951	0.839	4.15
4	20	0.946	0.834	4.39
4	40	0.943	0.824	4.51
4	100	0.942	0.822	4.54
8	10	0.951	0.838	4.18
8	20	0.945	0.825	4.45
8	40	0.935	0.814	4.90
8	100	0.922	0.796	5.34
12	10	0.951	0.837	4.18
12	20	0.942	0.822	4.51
12	40	0.933	0.807	5.10
12	100	0.929	0.804	5.31

The sparse–dense divergence remains stable throughout the expanded grid. Rank correlations remain high in every specification, while mean absolute rank changes remain economically meaningful. At a fixed horizon, the half-life barely moves the reordering (mean absolute rank changes between 4.39 and 4.51 at $T = 20$), even though, by Corollary 2, a longer half-life raises the effective attenuation $\Omega/(1 - \delta)$ and brings both operators closer to criticality. This matches the phase diagram in Exercise E6, where distance to criticality explains none of the variation in reordering. Longer horizons, by contrast, amplify reordering moderately, consistent with heterogeneous retention compounding over successive propagation rounds.

B.6 Intermediate uses exceeding domestic output

The Colombian table records total intermediate flows, including imported inputs, while the Ghosh coefficients divide by domestic gross output. For four sectors (wood products, chemicals, rubber and plastics, and electrical and electronic equipment), intermediate sales exceed domestic output by between 9% and 20%, so their rows of B_G° sum to slightly more than one. Because $\Omega = 0.47$ scales every row, their effective transmission remains well below one (0.51–0.57), and the operators remain stable. To verify that these rows do not drive the results, we bound all allocation coefficients by replacing x_j with $\max\{x_j, \sum_i W_{ji}\}$, recalibrate with the same rules, and repeat the decomposition of Table 5. The sparse–dense Spearman correlation moves from 0.945 to 0.939 and the mean absolute rank change from 4.45 to 4.66. The retention channel reproduces 62% of the reordering (against 58%) with a log-ratio R^2 of 0.97, while uniform attenuation reproduces 2% and the best backbone 24%. The conclusions are unchanged and, if anything, slightly strengthened.

B.7 The decomposition under alternative activation rules

Table 15: Appendix E1: The retention mechanism under alternative activation intensities and kernels

Activation rule	$M2b$ vs. $M3_{mf}$		Gap closed			Spearman(r_j, n_j^{eff})
	Spearman	Mean $ \Delta r $	p^*A^*	D_rG	Best backbone	
Baseline intensity	0.945	4.45	0.03	0.58	0.21	−0.48
Homogeneous intensity	0.885	7.01	0.10	0.74	0.37	−0.95
Size-only intensity	0.870	7.58	0.05	0.82	0.14	−0.39
Effective-degree-only intensity	0.982	2.33	0.18	0.33	0.50	−0.41
Weights-only kernel ($\lambda = 1$)	0.941	4.54	0.03	0.66	0.26	−0.57
Uniform kernel	0.732	9.97	−0.10	1.00	−0.14	0.71

Notes: Each row recomputes $M3_{mf}$ under an alternative activation rule and repeats the decomposition of Table 5. Intensity rules are rescaled to the baseline mean expected number of active links; kernel rules preserve each emitter’s baseline expected number of active links. Best backbone is the largest gap closed among the ρ -matched top- k and matched-sparsity disparity backbones. Under the uniform kernel π_{ji} is constant within each row, so $\widehat{G} = G$ and retention accounts for the reordering by construction.

B.8 Domestic intermediate flows in the OECD tables

Exercise E7 uses total intermediate flows, the convention of the Colombian table. Restricting propagation to domestic intermediate flows, with coefficients computed over domestic output, leaves the cross-country results essentially unchanged (Table 16). The median Spearman correlation between $M2b$ and $M3_{mf}$ is 0.983, retention reproduces a median 67% of the reordering (interquartile range 59–75%), uniform attenuation 4%, and the best backbone 12%. Retention again exceeds both alternatives in all 79 economies.

Table 16: Cross-country evidence with domestic intermediate flows

Median across economies	All ($N = 79$)	Core sample ($N = 20$)	Advanced ($N = 37$)	Emerging ($N = 42$)
Sectors, n	49	49	49	49
Density	0.99	0.99	0.98	0.99
Median n^{eff}	9.0	9.0	9.2	8.6
Spearman $M2b$ vs. $M0_G$	0.95	0.92	0.95	0.95
Spearman $M2b$ vs. $M3_{mf}$	0.983	0.977	0.981	0.985
Mean $ \Delta r /n$	0.036	0.042	0.037	0.034
Gap closed, p^*A^*	0.04	0.05	0.03	0.05
Gap closed, D_rG	0.67	0.64	0.68	0.67
Gap closed, best backbone	0.12	0.12	0.03	0.22
Log-ratio R^2 , D_rG	0.99	0.98	0.99	0.98
Gap closed, D_rG , homogeneous α	0.84	0.82	0.84	0.83
Spearman(r_j, n_j^{eff})	−0.08	−0.12	−0.02	−0.14

Notes: Intermediate flows from the OECD Input–Output Tables (2025 edition): domestic intermediate flows only, with allocation coefficients computed over domestic output. Each country is calibrated with the rules of Section 5. Medians across countries. “Core sample” is a representative subset of 20 economies (USA, DEU, JPN, GBR, FRA, KOR, AUS, ESP, BRA, MEX, ARG, CHL, PER, CRI, CHN, IND, IDN, VNM, ZAF, TUR); “Emerging” includes emerging and developing economies (IMF classification). Gap closed is defined as in Table 5; the best backbone is the larger of the ρ -matched top- k and matched-sparsity disparity backbones.

C Monte Carlo convergence and ranking stability

The second exercise verifies that the stochastic rankings are numerically reproducible. We re-estimate $M3$ and $M4$ over independent Monte Carlo seed streams and compare the finite-budget $M3$ rankings with the closed-form mean-field ranking $M3_{mf}$.

Table 17: Appendix E2: Monte Carlo convergence toward the mean-field ranking

n_{mc}	Mean Spearman	[2.5%, 97.5%]	Mean $ \Delta r $	[2.5%, 97.5%]
50	0.9995	[0.9993, 0.9998]	0.256	[0.157, 0.372]
100	0.9997	[0.9995, 0.9998]	0.206	[0.098, 0.269]
200	0.9997	[0.9996, 0.9999]	0.179	[0.076, 0.260]
300	0.9998	[0.9996, 0.9999]	0.109	[0.038, 0.222]
500	0.9999	[0.9998, 1.0000]	0.080	[0.030, 0.141]
800	0.9999	[0.9998, 1.0000]	0.065	[0.000, 0.111]

Notes: Each row summarizes 12 independent Monte Carlo seed streams. The intervals report the 2.5th and 97.5th percentiles across streams. Mean $|\Delta r|$ is the average absolute rank difference relative to $M3_{mf}$.

Table 18: Appendix E2: Ranking reproducibility across independent seed streams

Statistic	Value
Spearman between averaged $M3$ and $M3_{mf}$	1.0000
Mean pairwise Spearman across $M3$ streams	0.9997
Minimum pairwise Spearman across $M3$ streams	0.9994
Mean pairwise Spearman across $M4$ streams	0.9997
Minimum pairwise Spearman across $M4$ streams	0.9994

Notes: Pairwise correlations are computed across independent Monte Carlo seed streams using sectoral rankings. $M4$ includes the capacity channel.

The Monte Carlo implementation is effectively converged. Across independent seed streams, mean rank correlation with $M3_{mf}$ increases monotonically with the simulation budget, while the mean absolute rank error falls from 0.256 positions at 50 replications to 0.065 at 800. Rankings are also nearly identical across independent streams for both $M3$ and $M4$.

These results support Proposition 2(iii): sparse stochastic propagation estimates a deterministic effective-network centrality. The divergence documented in E1 is therefore structural rather than a Monte Carlo artifact.

D Activation-intensity transition

This exercise traces the activation-intensity path predicted by Proposition 4. For each value of θ , we compute the closed-form sparse mean-field ranking $M3_{mf}(\theta)$ and compare it with three benchmarks: the aligned dense supply-side limit $M2b$, the classical Ghosh benchmark $M0_G$, and the classical Leontief benchmark $M0_L$. Since θ scales the activation vector α , low values describe sparse activation, while high values approach the full-activation limit.

Figure 8 shows that the ranking generated by $M3_{mf}(\theta)$ approaches $M2b$ progressively as activation becomes dense. The Spearman correlation with $M2b$ rises from 0.824 at $\theta = 0.05$ to 0.999 at $\theta = 200$, while the Kendall correlation rises from 0.661 to 0.993. The same path remains strongly aligned with Ghosh and

essentially unrelated to Leontief. This confirms that the transition is an activation-intensity transition within the supply-side orientation, not a change in the direction of propagation.




Figure 8: Activation-intensity transition. The figure reports rank similarity between the closed-form sparse mean-field ranking $M3_{mf}(\theta)$ and three benchmarks: the aligned dense supply-side limit $M2b$, the classical Ghosh benchmark $M0_G$, and the classical Leontief benchmark $M0_L$. As θ increases, the sparse operator approaches its full-activation counterpart $M2b$. The persistent alignment with $M0_G$ and the near-zero correlation with $M0_L$ confirm that the operator remains supply-oriented along the transition path.

The numerical transition is summarized in Table 19. This compact representation shows three facts. First, the spectral radius of the effective operator $\rho(\bar{B})$ increases with θ , consistent with the monotone expansion of the effective propagation network. Second, both Spearman and Kendall correlations with $M2b$ rise sharply and approach one, confirming convergence toward the full-activation supply-side limit. Third, the ranking remains highly correlated with Ghosh and nearly orthogonal to Leontief throughout the path, reinforcing that Appendix E3 isolates activation intensity rather than orientation.

Table 19: Appendix E3: Activation-intensity transition

θ	$\rho(\bar{B})$	Spearman correlation			Kendall correlation		Top-10 vs. $M2b$
		vs. $M2b$	vs. $M0_G$	vs. $M0_L$	vs. $M2b$	vs. $M0_G$	
0.0500	0.0272	0.8241	0.8317	-0.0060	0.6608	0.6536	0.40
0.4801	0.1044	0.8863	0.8709	-0.0067	0.7422	0.7006	0.60
1.0205	0.1228	0.9442	0.9128	-0.0022	0.8227	0.7576	0.90
9.7992	0.1612	0.9957	0.9313	0.0085	0.9575	0.7910	0.90
200.00	0.1755	0.9996	0.9320	0.0086	0.9928	0.7919	1.00

Notes: The table reports selected points from the activation-intensity path. $\bar{B}$ is the closed-form effective sparse operator. $M2b$ is the dense finite-horizon supply-side benchmark aligned with the stochastic operator; $M0_G$ is the classical Ghosh benchmark; and $M0_L$ is the classical Leontief benchmark. Top-10 reports the share of sectors common to the top ten of $M3_{mf}(\theta)$ and $M2b$. Kendall correlations with $M0_L$ are omitted from the table because they remain close to zero throughout the path, matching the Spearman evidence and confirming the orientation gap.

E Algorithm

Algorithm 1 One Monte-Carlo run of the propagation operator (seed sector s)

Require: transmission matrix G ; selection kernel B ; intensities $\{\alpha_j\}$; Ω, δ ; capacities $\{C_i\}$ or ∞ ; horizon T ; seed s ; initial shock x_0

```

1: Initialize  $S \leftarrow x_0 e_s$ ; total  $\leftarrow x_0$ 
2: for  $t = 1, \dots, T$  do
3:    $\hat{P} \leftarrow \mathbf{0} \in \mathbb{R}^n$ 
4:   for each emitter  $j$  with  $S_j > 0$  and  $\alpha_j > 0$  do
5:     Activate links: for each  $i \in V_j$ , set

$$I_{ji} \leftarrow \mathbf{1}\{U_{ji} < 1 - e^{-\alpha_j B_{ji}}\}, \quad U_{ji} \sim \text{Unif}(0, 1)$$

6:     for each  $i$  with  $I_{ji} = 1$  do
7:        $\hat{P}_i \leftarrow \hat{P}_i + \Omega S_j G_{ji}$ 
8:     end for
9:   end for
10:   $P \leftarrow \sigma_C(\hat{P})$  ▷  $P = \hat{P}$  without capacity
11:  total  $\leftarrow$  total +  $\mathbf{1}^\top P$ 
12:   $S \leftarrow \delta S + P$ 
13: end for
14: return  $T_s^{(T)} =$  total

```

In the empirical propagation models, $G = B_G^\circ$ and the baseline ranking exercises use a unit seed, $x_0 = 1$. Alternative shock scales enter the capacity-stress exercise.

Algorithm 2 Finite-horizon mean-field operator and ranking (no capacity)

Require: transmission matrix G ; selection kernel B ; intensities $\{\alpha_j\}$; Ω, δ ; horizon T

```

1:  $\Pi_{ji} \leftarrow 1 - e^{-\alpha_j B_{ji}}$  for  $i \in V_j$ ; set  $\Pi_{ji} = 0$  otherwise ▷ marked-Poisson activation
2:  $\bar{B} \leftarrow \Omega(G \odot \Pi)^\top$ 
3:  $\mathcal{R}_T(\bar{B}) \leftarrow \sum_{t=1}^T \bar{B}(\delta I + \bar{B})^{t-1}$ 
4:  $\text{score}_s \leftarrow 1 + \mathbf{1}^\top \mathcal{R}_T(\bar{B}) e_s$  for each  $s$ 
5: Rank sectors by decreasing  $\text{score}_s$ 
6: return scores and ranking

```

Under marked-Poisson activation and no capacity, Algorithm 2 gives the exact expected finite-horizon score of Algorithm 1. Monte Carlo averages therefore converge to these scores as the number of replications increases. If $\rho(\bar{B}) < 1 - \delta$, the corresponding infinite-horizon propagation operator is $\bar{B}((1 - \delta)I - \bar{B})^{-1}$.

F Calibration details

Table 20 reports the full baseline calibration used in the empirical exercises. Parameters are calibrated by network moments and should not be interpreted as structural estimates.

Table 20: Full baseline calibration

Object	Value
Transmission rate Ω	0.4743
Persistence δ	0.7946
Activation scale α_{scale}	81.5405
Capacity scale κ	0.5000
Selection weight λ	0.8000
Weight curvature γ	1.0000
Simulation horizon T	20
Pure-persistence half-life implied by δ	3.0152
Aggregate intermediate-transactions share of gross output	0.4743
$\rho(\delta I + A^*)$	0.9709
$\rho(\bar{B})$	0.1224
Infinite-horizon dense operator converges	Yes
Mean expected active links	6.2891
Median expected active links	5.5789
Mean effective degree n^{eff}	7.1145
Median effective degree n^{eff}	6.2891
Target active links	6.2891
$\rho(B_G^\circ)$, intersectoral Ghosh allocation matrix	0.3717
$\rho(A^*) = \Omega \rho(B_G^\circ)$	0.1763
Effective attenuation $\Omega/(1 - \delta)$	2.3095
Distance to criticality, dense: $\rho(A^*)/(1 - \delta)$	0.8585
Distance to criticality, sparse: $\rho(\bar{B})/(1 - \delta)$	0.5959
Spectrally matched uniform scale $p^* = \rho(\bar{B})/\rho(A^*)$	0.6941
$\rho(A)$, Leontief technical-coefficient matrix	0.5413
$\rho(\delta I + \bar{B})$	0.9170
Realized diffusion half-life	8.0000
Ω calibration mode	Intermediate-transactions share
δ calibration mode	Diffusion half-life

Notes: $A^* = \Omega(B_G^\circ)^\top$ is the empirical dense full-activation operator and $\bar{B} = \Omega(B_G^\circ \odot \Pi)^\top$ is the closed-form effective sparse operator. The realized half-life is calibrated from $\rho(\delta I + \bar{B}) = 2^{-1/8} = 0.9170$. The pure-persistence half-life reports the half-life implied by δ alone, before accounting for propagation through $\bar{B}$. Spectral radii are computed directly from the calibrated operators and satisfy $\rho(\delta I + A^*) = \delta + \rho(A^*)$ (Corollary 2(i)). The attenuation parameter Ω is a common normalization anchored to the aggregate intermediate-transactions share of gross output, while B_G° preserves the cross-sector heterogeneity in transmission magnitudes. Expected active links are computed under the marked-Poisson activation kernel.

G Complete sectoral reordering (E4)

Table 21 reports the complete comparison between the dense finite-horizon supply-side ranking $M2b$ and the sparse mean-field ranking $M3_{mf}$. Positive values of Δr indicate that a sector becomes more central under sparse activation, whereas negative values indicate a loss of rank.

Table 21: E4: Complete sectoral rank changes under sparse activation

Sector	r^{M2b}	$r^{M3_{mf}}$	Δr
Livestock	35	16	19
Coffee	42	23	19
Grain milling and processed foods	33	17	16
Non-metallic mineral products	26	11	15
Specialized construction activities	21	12	9
Oils and fats	23	14	9
Agriculture excluding coffee	34	26	8
Crude petroleum and natural gas extraction	32	25	7
Mining-support activities	8	2	6
Petroleum refining	24	18	6
Basic metals and fabricated metal products	12	8	4
Real-estate activities	49	45	4
Storage and transport-support activities	7	4	3
Administrative and support services	9	6	3
Telecommunications	45	42	3
Electricity generation and distribution	46	43	3
Land and pipeline transport	17	15	2
Wholesale and retail trade	37	35	2
Food products, n.e.c.	39	37	2
Dairy products	52	50	2
Beverages and tobacco	54	52	2
Fishing and aquaculture	57	55	2
Water collection and distribution	50	49	1
Cocoa, chocolate, and confectionery	59	58	1
Coffee processing	63	62	1
Metal-ore extraction	1	1	0
Professional, scientific, and technical activities	20	20	0
Pharmaceutical products	31	31	0
Environmental sanitation services	40	40	0
Coal and lignite extraction	53	53	0
Meat and fish processing	56	56	0
Arts, entertainment, and recreation	60	60	0
Food and beverage services	61	61	0
Residential and non-residential construction	64	64	0
Public administration and defence	65	65	0
Civil-engineering construction	66	66	0
Human-health and social-work activities	67	67	0

Table 21 – *continued*

Sector	r^{M2b}	r^{M3mf}	Δr
Chemicals	2	3	-1
Rubber and plastic products	4	5	-1
Sugar and panela	38	39	-1
Insurance and pension funds	43	44	-1
Leather goods and footwear	47	48	-1
Sewerage	58	59	-1
Education	62	63	-1
Auxiliary financial activities	11	13	-2
Other manufacturing	44	46	-2
Apparel	55	57	-2
Waste collection and treatment	48	51	-3
Furniture	51	54	-3
Materials recovery	3	7	-4
Postal and courier activities	5	9	-4
Electrical, electronic, and optical equipment	6	10	-4
Financial intermediation	25	29	-4
Motor vehicles and other transport equipment	28	32	-4
Publishing, audiovisual, and information services	16	21	-5
Accommodation	36	41	-5
Wood and cork products	13	19	-6
Other service activities	41	47	-6
Water transport	27	34	-7
Paper, printing, and recorded media	29	36	-7
Air transport	30	38	-8
Forestry and logging	19	28	-9
Vehicle maintenance and repair	14	24	-10
Textiles	22	33	-11
Machinery and equipment	10	22	-12
Other mining and quarrying	15	27	-12
Gas production and distribution	18	30	-12

Notes: $\Delta r_i = r_i^{M2b} - r_i^{M3mf}$. Positive values indicate upward movement under sparse activation. Lower numerical ranks indicate greater importance. Rows are ordered by Δr , from the largest upward to the largest downward movement.

H Country-level results (E7)

Table 22 reports the main statistics of Exercise E7 for every economy, with total intermediate flows over total supply.

Table 22: Sparse reordering by country (OECD input–output tables)

Country	n	Spearman	Mean $ \Delta r /n$	Gap p^*A^*	Gap $D_r G$	Gap backbone	$sd(\log r_j)$
AGO (2021)	49	0.990	0.027	0.06	0.69	0.22	0.116
ARE (2021)	49	0.983	0.038	0.02	0.63	0.46	0.126
ARG (2021)	49	0.983	0.037	0.07	0.49	0.16	0.101
AUS (2021)	49	0.969	0.051	−0.05	0.57	0.25	0.105
AUT (2021)	49	0.983	0.037	0.11	0.78	0.07	0.110
BEL (2021)	49	0.979	0.033	0.05	0.68	0.12	0.090
BGD (2021)	49	0.974	0.040	0.12	0.46	0.29	0.117
BGR (2021)	49	0.981	0.039	0.13	0.66	0.09	0.110
BLR (2021)	49	0.987	0.031	−0.05	0.68	0.05	0.105
BRA (2021)	49	0.977	0.045	−0.09	0.52	−0.13	0.098
BRN (2021)	49	0.993	0.021	−0.04	0.60	0.00	0.097
CAN (2021)	49	0.968	0.055	0.09	0.70	0.12	0.097
CHE (2021)	49	0.974	0.040	0.02	0.58	−0.25	0.087
CHL (2021)	49	0.980	0.039	0.00	0.66	0.02	0.145
CHN (2021)	49	0.969	0.055	0.18	0.55	0.33	0.148
CIV (2021)	49	0.991	0.026	0.03	0.71	0.23	0.129
CMR (2021)	49	0.982	0.040	0.04	0.67	0.19	0.127
COD (2021)	49	0.986	0.037	0.00	0.89	0.30	0.148
CRI (2021)	49	0.987	0.032	0.11	0.76	0.29	0.089
CYP (2021)	49	0.976	0.047	0.12	0.75	0.07	0.127
CZE (2021)	49	0.981	0.039	0.13	0.53	0.13	0.112
DEU (2021)	49	0.985	0.032	0.08	0.61	−0.08	0.110
DNK (2021)	49	0.980	0.041	0.08	0.51	0.02	0.096
EGY (2021)	49	0.991	0.024	0.00	0.86	0.34	0.101
ESP (2021)	49	0.977	0.042	0.00	0.66	0.18	0.094
EST (2021)	49	0.982	0.040	0.12	0.65	0.15	0.096
FIN (2021)	49	0.984	0.037	−0.14	0.52	−0.05	0.127
FRA (2021)	49	0.972	0.046	0.00	0.58	0.02	0.108
GBR (2021)	49	0.974	0.046	0.11	0.62	0.13	0.113
GRC (2021)	49	0.980	0.040	0.06	0.71	−0.10	0.122
HKG (2021)	49	0.993	0.025	−0.07	0.77	−0.07	0.102
HRV (2021)	49	0.991	0.027	0.00	0.53	0.03	0.094
HUN (2021)	49	0.969	0.046	0.11	0.69	0.07	0.085
IDN (2021)	49	0.983	0.034	0.20	0.68	0.24	0.126
IND (2021)	49	0.971	0.045	0.02	0.67	0.02	0.145
IRL (2021)	49	0.987	0.030	0.14	0.75	−0.08	0.088
ISL (2021)	49	0.952	0.055	0.00	0.52	0.15	0.098
ISR (2021)	49	0.972	0.046	−0.07	0.60	0.13	0.123
ITA (2021)	49	0.979	0.043	0.12	0.63	0.25	0.097
JOR (2021)	49	0.989	0.028	0.18	0.59	0.29	0.122
JPN (2021)	49	0.973	0.047	0.09	0.51	0.05	0.153
KAZ (2021)	49	0.990	0.029	−0.14	0.63	0.26	0.097
KHM (2021)	49	0.994	0.021	0.04	0.72	0.20	0.089
KOR (2021)	49	0.955	0.054	−0.09	0.49	0.14	0.140
LAO (2021)	49	0.990	0.031	−0.03	0.68	0.35	0.128
LTU (2021)	49	0.973	0.042	0.18	0.71	−0.02	0.101
LUX (2021)	49	0.989	0.028	−0.06	0.62	0.06	0.076
LVA (2021)	49	0.985	0.035	0.02	0.76	0.17	0.129
MAR (2021)	49	0.987	0.028	0.03	0.71	0.26	0.104

Country	n	Spearman	Mean $ \Delta r /n$	Gap p^*A^*	Gap $D_r G$	Gap backbone	$sd(\log r_j)$
MEX (2021)	49	0.986	0.030	-0.03	0.39	0.25	0.112
MLT (2021)	49	0.978	0.044	0.00	0.70	0.02	0.114
MMR (2021)	49	0.984	0.030	0.11	0.39	0.25	0.083
MYS (2021)	49	0.985	0.029	-0.14	0.34	0.14	0.101
NGA (2021)	49	0.984	0.031	-0.11	0.57	0.24	0.117
NLD (2021)	49	0.978	0.040	0.17	0.71	0.10	0.094
NOR (2021)	49	0.964	0.056	-0.01	0.69	0.21	0.123
NZL (2021)	49	0.978	0.045	0.07	0.56	0.17	0.099
PAK (2021)	49	0.987	0.024	-0.14	0.45	0.28	0.086
PER (2021)	49	0.984	0.037	0.27	0.70	0.18	0.140
PHL (2021)	49	0.979	0.041	0.08	0.49	0.31	0.099
POL (2021)	49	0.966	0.046	0.16	0.64	0.13	0.105
PRT (2021)	49	0.979	0.036	-0.02	0.51	0.02	0.113
ROU (2021)	49	0.978	0.036	0.26	0.56	0.30	0.096
RUS (2021)	49	0.953	0.064	0.05	0.75	0.06	0.131
SAU (2021)	49	0.953	0.057	0.07	0.86	0.19	0.159
SEN (2021)	49	0.985	0.035	0.17	0.79	0.29	0.167
SGP (2021)	49	0.973	0.047	0.11	0.61	0.19	0.123
STP (2021)	49	0.994	0.021	-0.08	0.68	0.28	0.108
SVK (2021)	49	0.981	0.036	0.02	0.51	0.09	0.117
SVN (2021)	49	0.979	0.039	0.09	0.79	0.30	0.118
SWE (2021)	49	0.986	0.030	-0.11	0.61	-0.06	0.099
THA (2021)	49	0.973	0.043	0.10	0.73	0.23	0.131
TUN (2021)	49	0.983	0.028	-0.03	0.71	0.26	0.125
TUR (2021)	49	0.974	0.043	0.15	0.50	0.19	0.148
TWN (2021)	49	0.970	0.049	0.14	0.75	0.36	0.123
UKR (2021)	49	0.989	0.026	0.16	0.52	0.06	0.096
USA (2021)	49	0.981	0.041	0.12	0.59	-0.14	0.123
VNM (2021)	49	0.973	0.048	0.07	0.62	0.43	0.125
ZAF (2021)	49	0.978	0.041	-0.04	0.59	0.10	0.106

References

- Acemoglu, D., Carvalho, V. M., Ozdaglar, A., and Tahbaz-Salehi, A. (2012). The network origins of aggregate fluctuations. *Econometrica*, 80(5):1977–2016.
- Alatríste Contreras, M. G. and Fagiolo, G. (2014). Propagation of economic shocks in input-output networks: A cross-country analysis. *Physical Review E*, 90(6):062812.
- Barrot, J.-N. and Sauvagnat, J. (2016). Input specificity and the propagation of idiosyncratic shocks in production networks. *The Quarterly Journal of Economics*, 131(3):1543–1592.
- Barthélemy, M., Barrat, A., Pastor-Satorras, R., and Vespignani, A. (2005). Characterization and modeling of weighted networks. *Physica A: Statistical Mechanics and its Applications*, 346(1–2):34–43.
- Bonacich, P. (1987). Power and centrality: A family of measures. *American Journal of Sociology*, 92(5):1170–1182.
- Brin, S. and Page, L. (1998). The anatomy of a large-scale hypertextual web search engine. *Computer Networks and ISDN Systems*, 30(1–7):107–117.
- Carvalho, V. M. (2014). From micro to macro via production networks. *Journal of Economic Perspectives*, 28(4):23–48.
- Carvalho, V. M., Nirei, M., Saito, Y. U., and Tahbaz-Salehi, A. (2021). Supply chain disruptions: Evidence from the great east japan earthquake. *The Quarterly Journal of Economics*, 136(2):1255–1321.
- Carvalho, V. M. and Tahbaz-Salehi, A. (2019). Production networks: A primer. *Annual Review of Economics*, 11:635–663.
- Departamento Administrativo Nacional de Estadística (DANE) (2025). Matriz insumo-producto (mip) 2021. Dirección de Síntesis y Cuentas Nacionales, Cuentas Nacionales, Bogotá. Activity-by-activity input-output matrix, base 2015. Published September 29, 2025.
- Erdős, P. and Rényi, A. (1959). On random graphs i. *Publicationes Mathematicae Debrecen*, 6:290–297.
- Ermann, L., Frahm, K. M., and Shepelyansky, D. L. (2015). Google matrix analysis of directed networks. *Reviews of Modern Physics*, 87(4):1261–1310.
- Estrada, E. and Hatano, N. (2008). Communicability in complex networks. *Physical Review E*, 77(3):036111.
- Fink, C. G., Fullin, K., Gutierrez, G., Omodt, N., Zinnecker, S., Sprint, G., and McCulloch, S. (2023). A centrality measure for quantifying spread on weighted, directed networks. *Physica A: Statistical Mechanics and its Applications*, 626. arXiv:2303.09684.
- Ghosh, A. (1958). Input-output approach in an allocation system. *Economica*, 25(97):58–64.
- Gilbert, E. N. (1959). Random graphs. *The Annals of Mathematical Statistics*, 30(4):1141–1144.
- Gleeson, J. P., Cassidy, A., Giles, D., and Faqueh, A. (2025). Time-dependent influence metric for cascade dynamics on networks. *Physical Review E*, 111(5):054310.
- Grindrod, P., Parsons, M. C., Higham, D. J., and Estrada, E. (2011). Communicability across evolving networks. *Physical Review E*, 83(4):046120.

- Guerra, B. and Gómez-Gardeñes, J. (2010). Annealed and mean-field formulations of disease dynamics on static and adaptive networks. *Physical Review E*, 82(3):035101.
- Hidalgo, C. A., Klinger, B., Barabási, A.-L., and Hausmann, R. (2007). The product space conditions the development of nations. *Science*, 317(5837):482–487.
- Holme, P. and Saramäki, J. (2012). Temporal networks. *Physics Reports*, 519(3):97–125.
- Horn, R. A. and Johnson, C. R. (2013). *Matrix Analysis*. Cambridge University Press, Cambridge, 2 edition.
- Katz, L. (1953). A new status index derived from sociometric analysis. *Psychometrika*, 18(1):39–43.
- Kempe, D., Kleinberg, J., and Tardos, É. (2003). Maximizing the spread of influence through a social network. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '03, pages 137–146, New York, NY. Association for Computing Machinery.
- Leontief, W. W. (1936). Quantitative input and output relations in the economic system of the united states. *The Review of Economics and Statistics*, 18(3):105–125.
- Miller, R. E. and Blair, P. D. (2009). *Input–Output Analysis: Foundations and Extensions*. Cambridge University Press, Cambridge, 2 edition.
- Newman, M. E. J. (2002). Spread of epidemic disease on networks. *Physical Review E*, 66(1):016128.
- Newman, M. E. J. (2010). *Networks: An Introduction*. Oxford University Press, Oxford.
- OECD (2025). OECD input-output tables, 2025 edition: Domestic output and imports. https://webfs-sti.oecd.org/files/STI-PIE/IOTs/2025/IOTs_DOMIMP.zip. Accessed October 2026.
- Perra, N., Gonçalves, B., Pastor-Satorras, R., and Vespignani, A. (2012). Activity driven modeling of time varying networks. *Scientific Reports*, 2:469.
- Pichler, A., Pangallo, M., del Rio-Chanona, R. M., Lafond, F., and Farmer, J. D. (2022). Forecasting the propagation of pandemic shocks with a dynamic input-output model. *Journal of Economic Dynamics and Control*, 144:104527.
- Radicchi, F., Ramasco, J. J., and Fortunato, S. (2011). Information filtering in complex weighted networks. *Physical Review E*, 83:046101.
- Serrano, M. Á., Boguñá, M., and Vespignani, A. (2009). Extracting the multiscale backbone of complex weighted networks. *Proceedings of the National Academy of Sciences*, 106(16):6483–6488.
- Watts, D. J. (2002). A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences of the United States of America*, 99(9):5766–5771.