



City Research Online

City St George's, University of London

Citation: Amici, G., Fusai, G., Gambaro, A. M. & Marazzina, D. (2026). Navigating Supply Shocks: Sector Resilience and Production Prices Through Stochastic Input–Output Modeling. *Mathematical Finance*, mafi.70029. doi: 10.1111/mafi.70029

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/37297/>

Link to published version: <https://doi.org/10.1111/mafi.70029>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

Navigating Supply Shocks: Sector Resilience and Production Prices through Stochastic Input-Output Modeling

Giovanni Amici^{*}, Gianluca Fusai[†], Anna Maria Gambaro[‡] and Daniele Marazzina[§]

April 9, 2026

Short title: Navigating Supply Shocks.

Acknowledgment: This research has been funded by the European Union - Next Generation EU, Mission 4-Component 1-CUP D53D23017690001 - Project PRIN PNRR 2022 P20228CHNL "Measuring, managing and hedging indirect climate-transition risk"

^{*}Department of Industrial and Systems Engineering, North Carolina State University, 915 Partners Way, Raleigh, NC 27606, USA, gamici@ncsu.edu

[†]Faculty of Finance, Bayes Business School City St George's, University of London gianluca.fusai.1@city.ac.uk
Dipartimento di Studi per l'Economia e l'Impresa, Università del Piemonte Orientale, Via Perrone 18, 28100 Novara, Italy, gianluca.fusai@uniupo.it

[‡]Dipartimento di Studi per l'Economia e l'Impresa, Università del Piemonte Orientale, Via Perrone 18, 28100 Novara, Italy, annamaria.gambaro@uniupo.it

[§]Department of Mathematics, Politecnico di Milano, Piazza Leonardo da Vinci, 32, Milano, I-20133, Italy. daniele.marazzina@polimi.it

Abstract

This study develops a novel multivariate stochastic framework for assessing systemic risks, such as climate and nature-related shocks, within production or financial networks. By embedding a linear stochastic fluid network, interpretable as a generalized vector Ornstein-Uhlenbeck process, into the production network of interdependent industries, the model captures how physical shocks (e.g., extreme climate events or geopolitical disruptions) propagate through input-output (IO) linkages and affect sectoral price dynamics. The framework extends traditional IO models with advanced stochastic and dynamic features, enabling a quantification of both direct and indirect transmission channels of supply-cost shocks to production prices. Contributing to the literature on stochastic IO and Markovian networks, the model introduces the concept of divisible shocks, allowing for finer-grained simulation of adaptation responses and resilience across sectors. Empirical calibration leverages real-world economic data, including IO tables and historical industrial price indices. Sensitivity analyses are conducted using distributional risk measures, offering new tools for climate stress testing and medium to long-term risk assessment. Our findings support the optimal design of supply risk management strategies, including policy interventions and decentralized adaptation incentives for systemic stability under environmental stress.

Keywords: multivariate stochastic processes, production prices, divisible shocks, input-output tables, linear stochastic fluid network.

1 Introduction

Supply chain and global value chain risk management have become increasingly prominent fields of research, gaining substantial attention from both academia and industry. While disruptions in supply and value chains are not novel phenomena (Sodhi and Tang, 2012), their frequency, complexity, and severity have been significantly exacerbated by recent global events, most notably the COVID-19 pandemic, the intensification of physical risks driven by climate change, and a rise in geopolitical tensions. As a result, supply chain security has emerged as a critical concern, emphasizing the necessity for effective risk-sharing mechanisms and systemic resilience strategies.

Recent analyses, such as those by MSCI,¹ highlight that certain sectors, particularly energy, materials, utilities, and industrials, are especially vulnerable to climate-related risks. These industries face heightened exposure to physical disruptions caused by extreme weather events, including floods, heatwaves, and wildfires, which can severely impair production continuity and logistics. Consequently, integrating climate risk assessment into supply chain and value chain management has become essential to ensure operational stability and long-term sustainability.

In tackling these challenges, it is important to distinguish between direct risks, originating within a specific sector or firm, and indirect risks, which arise through exposure to interconnected supply networks. Quantifying losses from such disruptions remains inherently complex, given their extreme characteristics and the intricate web of interdependencies across sectors and countries. This complexity necessitates the use of dynamic and stochastic modeling frameworks that go beyond deterministic assumptions.

Our study contributes to this literature by exploring the dynamic resilience of economic sectors through input-output (IO) modeling, focusing on their response to stochastic supply-side

¹<https://www.msci.com/quick-take/which-sectors-are-most-affected/03827173296>

shocks. Methodologically, we position our work at the intersection of several strands of research, including stochastic shock propagation, uncertainty modeling, and network-based analyses. For example, Le Guenedal and Tankov (2025) investigate the propagation of carbon taxation shocks in interconnected economic systems, emphasizing scenario uncertainty and resilience metrics. Sopgoui (2024) examines the impacts of climate-related policy shocks on real estate markets using stochastic models built on Ornstein-Uhlenbeck (OU) dynamics. Approaches to modeling uncertainty, such as those in Cambou and Filipović (2017), and contagion mechanisms in financial networks, as explored by Veraart (2020), provide key insights for our framework on systemic risk. In parallel, climate transition risk is increasingly being studied using network methods, as in Pang and Shrimali (2024), underscoring the relevance and timeliness of network-based approaches to production and value chain disruptions.

To capture indirect risks, we utilize input-output tables, which describe the intersectoral flow of goods and services and can be naturally interpreted as weighted directed networks. These tools have become central to the rapid evaluation of economic disruptions. Our methodological foundation builds upon classical IO models, including the Dynamic Inoperability Input-Output Model (DIIM) proposed by Lian and Haines (2006) and its supply-driven variant developed by Xu et al. (2011). While these models represent important advances, they are deterministic and apply shocks only at the initial step, after which the system evolves in a fixed manner.

In contrast, a growing body of literature considers the propagation of stochastic shocks within IO networks. Blochl et al. (2011) and Mascaretti et al. (2022) use random walk-based indicators to study how shocks spread across sectors, while Kostoska et al. (2020) employs absorbing Markov chains to capture resilience in global trade networks. However, these models typically assume that shocks are indivisible and propagate sequentially across nodes, an assumption that may oversimplify real-world dynamics.

To address this point, we propose a novel stochastic framework based on the theory of linear stochastic fluid networks, as developed in Kella and Whitt (1999). Our model enables the representation of divisible shocks and allows for sector-specific responses and feedback. We adopt a two-factor structure inspired by Schwartz and Smith (2000) for commodity prices, in which production log-price dynamics are decomposed into a long-term equilibrium component and a short-term fluctuation component. The latter is modeled using a generalized vector OU process, which can be interpreted as a stochastic fluid network. This formulation captures both transient volatility and reversion to equilibrium, with the mean-reverting matrix expressed as the product of a Leontief IO matrix (encoding network structure) and a diagonal matrix of sectoral output rates.

While vector OU models have been extensively used for commodities (see Schwartz and Smith, 2000; Casassus and Collin-Dufresne, 2005; Cortazar et al., 2008; Chiang et al., 2015; Cortazar et al., 2019), to our knowledge, this is the first application of a generalized vector OU process to the dynamics of production prices in a networked economic system. This opens a path for future research into commodity-specific IO systems (see Bohlin and Widell, 2006; Weisz and Duchin, 2006), as well as broader macro-financial applications.

Our model addresses key research questions: (i) How do supply shocks, such as those caused by extreme weather events or geopolitical crises, propagate through production networks? (ii) What determines the recovery time of sectors following a disruption? (iii) How can policy

interventions mitigate long-term pricing effects?

Estimation is conducted using economic data, including IO tables and historical series of industrial production prices from the US. We estimate sector-specific resilience rates that characterize the speed of adjustment to equilibrium after shocks. We also study the long-term impact of persistent supply shocks on the stationary distribution of production prices, which can be derived semi-analytically in the case of Lévy-type shock processes. A sensitivity analysis is performed using risk measures of log-prices distributions, providing new tools for supply risk management. In particular, empirical analysis reveals that the tail risk associated with sectors such as manufacturing, utilities, and construction is primarily driven by fossil fuel prices, specifically, the production prices of the mining oil sector. Reducing the dependence of both the manufacturing and energy sectors on fossil fuels could therefore contribute to greater stability in production prices. Furthermore, the manufacturing production prices are also influenced by the tail risk of the agricultural production prices, that are heavily influenced by weather conditions.

In recent years, the increasing frequency and severity of exogenous shocks, particularly those driven by environmental and geopolitical factors have underscored the urgent need for economic models that can capture systemic vulnerabilities in interconnected production systems, see, e.g. Ciravegna et al. (2023) and references therein. Among these, climate-related risks have emerged as a particularly salient threat due to their global reach, structural persistence, and complex propagation patterns. Climate change gives rise to a novel class of systemic shocks, characterized by high uncertainty, heavy-tailed distributions, and non-linear amplification through production networks. Extreme weather events such as floods, heatwaves, and droughts can initiate localized supply disruptions, which, when transmitted via intersectoral dependencies, lead to cascading effects and potential macroeconomic instability.

Traditional input-output and macroeconomic models often lack the stochastic and dynamic components necessary to adequately capture these phenomena. The framework developed in this study fills this methodological gap by enabling the simulation and quantification of climate-induced supply shocks and their propagation across sectoral and geographical boundaries. By explicitly modeling both direct and indirect transmission channels, and by incorporating sector-specific resilience dynamics, our approach provides a versatile and analytically tractable tool for climate stress testing and scenario-based systemic risk assessment. These features make it particularly well-suited for informing adaptation strategies, guiding the allocation of public and private resources, and supporting the design of policy interventions aimed at enhancing the long-term resilience and sustainability of production systems under increasing climate uncertainty.

This work is organized as follows: after a short introduction on the literature of IO models in Section 2, Section 3 is devoted to the presentation of our model and to the main theoretical results of the article. Then, in Sections 4 and 5, we introduce the likelihood-based estimation procedure and estimate the model using historical data from several industrial sectors, analyzing the results. Section 6 presents risk management tools through a sensitivity analysis of price risk measures with respect to crucial inputs of the model, such as resilience rates and network links. Since the sensitivity analysis indicates that increasing a sector's resilience rate mitigates the impact of shocks on its price, Section 7 formulates the resource allocation problem as a multi-objective optimal control problem. Conclusions and future research are presented in Section 8.

2 Overview of the Input-Output equilibrium models

In this section, we provide a concise overview of the IO literature, with a particular emphasis on the interpretation of IO models as flow networks suitable for the analysis of indirect risks. We begin by introducing the foundational contributions of Leontief (1936) and Ghosh (1964), who developed the core static IO frameworks. These models describe the interdependencies among sectors in terms of either outputs (in the Leontief model) or inputs (in the Ghosh model), laying the groundwork for network-based representations of economic activity.

While the Leontief (1936) and Ghosh (1964) formulations focus on levels of production, we then introduce subsequent extensions to price and log-price dynamics, enabling the assessment of cost transmission mechanisms across sectors.

To incorporate temporal dynamics and post-shock adjustments, we finally consider the DIIM, introduced by Lian and Haines (2006). The DIIM extends the classical Leontief framework by introducing a dynamic adjustment process that governs how an economy gradually recovers from an initial supply or demand imbalance. A key feature of the DIIM is the notion of sector-specific production adjustment rates, which quantify the speed at which each sector returns to its pre-shock equilibrium. These rates play a central role in characterizing sectoral resilience and will be incorporated into our stochastic framework.

Leontief IO model IO models are a class of economic models used to describe how different industries or sectors within an economy interact. The Leontief (1936) IO model represents these interactions using a system of linear equations that describe how output from different sectors are interdependent. More precisely, it is based on the definition of the technical coefficient matrix A .

Definition 1. *Let an economy be divided into N productive sectors. For each pair of sectors (i, j) , let $A_{i,j}$ be the ratio of input from industry i to industry j , with respect to the total production of industry j . The technical coefficient matrix is given by $A = (A_{i,j})_{i,j=1,\dots,N}$.*

Remark 2. *As shown in Ronald E. Miller (2022)[Section 2.4.2], the above definition implies that all the elements of A are non negative and each of its columns sums to less than one. Thus A has spectral radius less than 1, $A^n \rightarrow 0$ as $n \rightarrow \infty$, $I - A$ is non-singular and $(I - A)^{-1} = \sum_{n=0}^{\infty} A^n$.*

In this framework, the Leontief IO model states that, at the equilibrium, the total production x and the final demand c are linked by

$$x_i = \sum_{j=1}^N A_{i,j} x_j + c_i, \quad i = 1, \dots, N,$$

or equivalently

$$x_i = \sum_{j=1}^N Z_{i,j} + c_i, \quad i = 1, \dots, N, \tag{2.1}$$

where Z is the IO matrix of an economic system or a supply chain: it represents the flow of products and services from each sector (or firm) to each of the other sector in monetary units.

The relation between the technical coefficient matrix A and the IO matrix Z is

$$A = ZD_x^{-1}, \quad (2.2)$$

with D_x the diagonal matrix of the vector x . This implies

$$x = (I - A)^{-1}c, \quad (2.3)$$

where I is the identity $N \times N$ matrix and $(I - A)^{-1}$ is called Leontief (inverse) matrix.

Notice that A , normalized by its row sums, is a stochastic matrix and can be interpreted as a generator matrix of a Markov chain, see Blochl et al. (2011) and Mascaretti et al. (2022). Moreover, in terms of graph representation, the technical coefficients matrix A represents an adjacency matrix, where the element $A_{i,j}$ represents the weight of the edge from node i to node j . Therefore, we can interpret the model as a directed (or flow) network.

Ghosh model Given the interpretation of the Leontief model as a flow network, we can apply the Flow Conservation Constraint (FCC), which states that the sum of the flow through edges directed toward a node plus that node's supply, equals the sum of the flow through edges directed away from that node plus that node's demand. Then, if we define the vector v as total value added to each node (total supply), we obtain the following algebraic representation of the FCC

$$\sum_{j=1}^N Z_{j,i} + v_i = \sum_{j=1}^N Z_{i,j} + c_i, \text{ for } i = 1, 2, \dots, N,$$

In the Ghosh model equation (2.1) becomes

$$x_i = \sum_{j=1}^N Z_{j,i} + v_i. \quad (2.4)$$

Therefore it holds that

$$x = (I - B)^{-1}v, \quad (2.5)$$

where

$$B = Z'D_x^{-1} = D_x A' D_x^{-1}, \quad (2.6)$$

Z' denotes the transpose of Z , and $\mathbf{1}$ is the N -dimensional vector of ones. The element $B'_{i,j} = B_{j,i}$ represents the ratio of output from i to j with respect to the total production of i , and B is called allocation coefficients matrix.

This second formulation of the IO equilibrium is called Ghosh (1964) model, while the matrix $(I - B)^{-1}$ is the Ghosh inverse matrix.

Notice that the Leontief model is a *demand side* model, while the Ghosh model is a *supply side* model. The Leontief model assumes fixed technical (input) coefficients, i.e., each sector fulfills its productive task by exploiting inputs in fixed proportions. The Ghosh model has fixed allocation (output) coefficients, implying “fixed output, or sales, distribution across sectors”

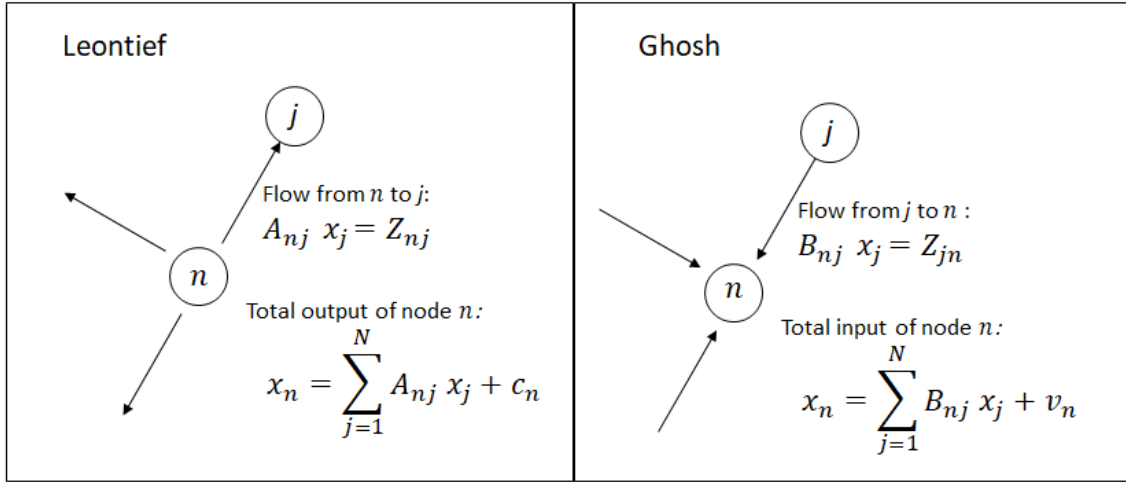


Figure 1: Graphical representation of the two IO models as flow networks.

(Ghosh, 1964). Figure 1 graphically represents the two models as flow networks; we recall that by the FCC the total input and the total output of each node must coincide.

IO price model Another class of IO models, instead of focusing on output levels, examines price formation in the economy (see, for instance, Ronald E. Miller, 2022). In this framework, if we denote the base year index prices by $p = (p_1, \dots, p_N)$, and $v^c = D_x^{-1}v = (v_1/x_1, \dots, v_N/x_N)$ is the supply cost per unit of production, then the input-output price model is

$$p = A'p + v^c. \quad (2.7)$$

To better highlight the link between the IO price model (2.7) and the IO Leontief and Ghosh models, let us consider that there are two scenarios: a base scenario with total supply vector v^0 and a stressed scenario with v^1 . The allocation coefficient matrix B is supposed to be the same in both scenarios. Following equation (2.4), the two equilibrium equations read $x^m = Bx^m + v^m$, $m = 0, 1$. We now define the price ratio between the two production vector as $p = D_{x^0}^{-1}x^1 = (x_1^1/x_1^0, \dots, x_N^1/x_N^0)$. Exploiting (2.6), we obtain

$$p = D_{x^0}^{-1}(Bx^1 + v^1) = D_{x^0}^{-1}BD_{x^0}p + D_{x^0}^{-1}v^1 = (A^0)'p + D_{x^0}^{-1}v^1,$$

where A^0 is the matrix of technical coefficients in the base scenario, and $D_{x^0}^{-1}v^1$ is the supply price per unit of production in the stressed scenario. This equation clarifies the role of the transpose of the technical coefficient matrix A in determining the equilibrium production prices in the IO static model.

IO log-additive models A parallel stream of economic literature, rooted in Cobb-Douglas production technologies, proposes log-linear IO models for output production $x(t)$ and prices $p(t)$ (see, for instance, Acemoglu et al., 2012; Carvalho and Tahbaz-Salehi, 2019). In particular, Carvalho and Tahbaz-Salehi (2019) present an IO network for the logarithm of the price to wage

ratios, that is

$$\log\left(\frac{p}{w}\right) = A' \log\left(\frac{p}{w}\right) + v^c, \quad (2.8)$$

where v^c captures supply-side log-cost shocks per unit of production, and w denotes wages. This specification aligns with the classical economic principle that wages reflect productivity growth: workers are compensated in proportion to their marginal product. Accordingly, as in Carvalho and Tahbaz-Salehi (2019), the wage w is identified as an index of economic growth.

DIIM models Another extension of the traditional IO models is the DIIM, see Lian and Haimes (2006), designed to analyze the propagation of disruptions and failures across economic sectors. Unlike the standard Leontief or Ghosh IO models, which assume normal operational conditions, DIIM introduces dynamic inoperability states to model how economic sectors recover from shocks (e.g., natural disasters, cyberattacks, financial crises, or pandemics).

The DIIM describes the adjustment dynamic of the level of the economy's production following an initial imbalance in supply and/or demand. Consider an initial value for the total production vector at time zero, $x(0)$, then, the dynamic of the production vector $x(t)$ for $t \geq 0$ is assumed to be described by the following system of ordinary differential equations (ODE)

$$\frac{dx(t)}{dt} = -K(I - A)x(t) + Kc, \quad (2.9)$$

where A is the technical coefficient matrix and K is a positive diagonal matrix, which elements $K_{ii} > 0$ for $i = 1, \dots, N$ are the production adjustment rates. We denote by \hat{x} the equilibrium value, that is, the value such that if the system is initialized at equilibrium at time $t = 0$, it remains in equilibrium for all $t > 0$. The equilibrium is given by the solution of the static Leontief model (2.3), namely $\hat{x} = (I - A)^{-1}c$, and, substituting this expression into the dynamics, the system can be rewritten as

$$\frac{dx(t)}{dt} = -K(I - A)(x(t) - \hat{x}), \quad x(0) = x_0. \quad (2.10)$$

In the following, we discuss how the DIIM represents the flow speed between sectors in the network and explain why this representation is not fully consistent with a linear fluid network model. Suppose that we have two sectors, and that only the first component of $x(0)$ deviates from its equilibrium value, i.e., $x_1(0) \neq \hat{x}_1$. We examine the value of $x(t) - \hat{x}$ after a very short time interval $\Delta t \simeq 0$. Using the Euler discretizations of the ODE (2.10), we obtain

$$x_2(\Delta t) - \hat{x}_2 \simeq k_2 A_{2,1}(x_1(0) - \hat{x}_1)\Delta t.$$

This shows that the contribution to sector 2 is proportional to the production adjustment rate k_2 and not k_1 . However, if we interpret the IO framework as a flow network, k_1 represents the output rate of sector 1, that is, the speed at which a shock originating in node 1 propagates to the other nodes. In contrast, as is evident from the expression above, within the DIIM framework a shock to the first sector propagates to the second one with speed k_2 . In the next section, we modify and extend the DIIM by incorporating insights from the theory of linear stochastic fluid networks, exchanging the positions of the technical coefficient matrix A and the adjustment-rate

matrix K . Under this modification, the ODE (2.10) becomes

$$\frac{dx(t)}{dt} = -(I - A)K(x(t) - \hat{x}), \quad x(0) = x_0.$$

3 A generalized vector Ornstein-Uhlenbeck dynamics for log-prices in an interconnected production system

In this section, we extend the DIIM model by incorporating a linear stochastic fluid network that captures random external shocks affecting both demand and supply costs. These shocks can influence either the production vector x or the price vector p , with our analysis primarily focused on the latter.

Formally, a linear stochastic fluid network is characterized by random external inputs, while, analogous to fluid dynamics, the internal flows remain deterministic and continuous, occurring at fixed proportional rates conditional on the system state (Kella and Whitt, 1999). This framework can be rigorously interpreted as a generalized vector Ornstein-Uhlenbeck process, in which the system exhibits mean-reverting behavior in response to stochastic disturbances. The matrix of mean-reversion coefficients is represented as the product of a matrix capturing the network structure, consistent with the IO and DIIM frameworks and a diagonal matrix of sector-specific proportional output rates.

Let $p(t)$ be the production price vector of N firms or industrial sectors. We assume that the logarithm of prices, $y(t)$, is composed of two factors

$$y(t) := \log(p(t)) = \zeta(t) + \xi(t), \quad (3.1)$$

where, following Carvalho and Tahbaz-Salehi (2019), see equation (2.8), $\zeta(t)$ is a vector of relative log-prices, defined as the logarithms of production prices $p(t)$ normalized by indexes of economic growth $e^{\xi(t)}$, such as wage or real GDP. A similar model structure has been analyzed by Schwartz and Smith (2000) in the context of univariate commodity price. Their study identifies two components of the log-price dynamic: an equilibrium price level, and price deviations that may arise from extreme weather conditions or supply disruption.

Given that production prices $p(t)$ and the aggregate price index $e^{\xi(t)}$ are observable, their logarithms are known. Our goal is therefore to study the dynamics of the relative log-price vector

$$\zeta(t) := \log(p(t)) - \xi(t) = y(t) - \xi(t),$$

which captures deviations of sectoral prices from the aggregate trend. We assume that the variations of $\zeta(t)$ are generated from supply shocks that hit the different industries (i.e., network nodes) randomly over time; such shocks are modeled by an N -dimensional stochastic process $v^c(t)$ corresponding to the instantaneous stochastic supply log-cost per unit of production. Consider a $N \times N$ technical coefficient matrix A , see Definition 1. Finally, suppose that the output rate of each node n is proportional to the value of $\zeta_n(t)$, with proportionality constant $k_n > 0$. The process $\zeta_n(t)$ results from the difference between the cumulative inputs and outputs of node

n from zero to t , and its dynamic is the solution of the following integral equation

$$\zeta_n(t) = \zeta_n(0) + \int_0^t dv_n^c(s) + \sum_{j=1}^N \int_0^t A_{jn} k_j \zeta_j(s) ds - \int_0^t k_n \zeta_n(s) ds, \quad (3.2)$$

in matrix form

$$\zeta(t) = \zeta(0) - \int_0^t (I - A') K \zeta(s) ds + \int_0^t dv^c(s). \quad (3.3)$$

Equation (3.2) is therefore composed of the following terms: (i) the initial value $\zeta(0)$, (ii) the cumulative external supply log-cost shocks, that are represented by the integral of the instantaneous supply log-price per unit of production $\int_0^t dv_n^c(s)$, (iii) the internal inputs from other nodes $\sum_{j=1}^N \int_0^t A_{jn} k_j \zeta_j(s) ds$, minus (iv) the total output from the node $\int_0^t k_n \zeta_n(s) ds$. In this decomposition, term (ii) captures the direct effect of external supply shocks on each sector, whereas term (iii) represents the indirect propagation of shocks through the network, reflecting the impact on upstream suppliers connected to the sector.

Assumption 3. A is a technical coefficient matrix, as stated in Definition 1.

Remark 4. Definition 1 and Remark 2 imply that A' is a substochastic transition matrix, each of its rows sums to less than one. Then, $(I - A')^{-1}$ is the fundamental matrix of an absorbing Markov chain and $(I - A')$ is the subgenerator of a Markov process.

Suppose that $\zeta(0)$ is a measurable random variable and $\{v^c(t)\}_{t \geq 0}$ is a semimartingale with respect to the filtration generated by $\{v^c(t)\}_{t \geq 0}$ and $\zeta(0)$. The existence and pathwise uniqueness of the solution of equation (3.3) is proved in Protter (1990)[chapter V, section 3]. The unique strong solution of the integral equation (3.3) is given by

$$\zeta(t) = e^{-(I-A')Kt} \zeta(0) + \int_0^t e^{-(I-A')K(t-s)} dv^c(s). \quad (3.4)$$

The first term captures the exponential decay of initial deviations. Each sector gradually reverts to equilibrium at a rate determined by its own k_i , and modulated by the intersectoral network A . The second term accumulates the effects of random supply shocks over time. The exponential term discounts past shocks, so that more recent shocks have a stronger influence on current prices and the effect of each shock vanishes over time. Therefore, the presence of this second term implies that production log-prices $y(t)$ do not fully revert to their equilibrium level $\xi(t)$, that is, $\zeta(t)$ do not fully revert to zero, after an initial supply shock. Instead, long-term deviations persist due to the stochastic nature of cost fluctuations. We will discuss stationary assumptions and results for the relative log-price vector $\zeta(t)$ in Section 3.2.

3.1 Characteristic function and moments

From this point onward, we state and prove the main theoretical results of our model under the following assumption.

Assumption 5. $v^c(t)$ is an N -dimensional Lévy process, with Lévy measure ν such that

$$\int_{\|x\|>1} \log(\|x\|) \nu(dx) < \infty.$$

Hereafter, we provide the characteristic exponents of the relative logarithmic prices $\zeta(t)$.²

Proposition 6. Let Assumption 5 for v^c hold, then the characteristic exponent of the relative log-prices $\zeta(t)$ is

$$\psi_t^\zeta(u) = -u' e^{-(I-A')Kt} \zeta_0 + \int_0^t \psi \left(e^{-K(I-A)(t-s)} u \right) ds, \quad (3.5)$$

where $u \in \mathbb{C}^N$ and

$$\psi(u) = -\log \left(\mathbb{E} \left[e^{-u' v^c(1)} \right] \right) \quad (3.6)$$

is the characteristic exponent of $v^c(1)$.

Proof. To prove (3.5) we apply the multi-dimensional extension of Lemma 15.1 of Cont and Tankov (2004), that is, assuming that $f : [0, t] \rightarrow \mathbb{R}^N$ is (left) continuous, then

$$-\log \mathbb{E} \left[e^{-\int_0^t (f(s))' dv^c(s)} \right] = \int_0^t \psi(f(s)) ds,$$

with in this case $f(s) = e^{-K(I-A)(t-s)} u$ and thus its transpose is $(f(s))' = u' e^{-(I-A')K(t-s)}$. \square

From Proposition 6, we derive analytically the conditional mean and covariance of the relative log-price and the conditional mean of the relative price. Firstly, the conditional mean of the relative log-prices $\zeta(t)$ is given by

$$\mathbb{E}[\zeta(t) | \zeta_0 = y_0 - \xi_0] = e^{-(I-A')Kt} (y_0 - \xi_0) + \int_0^t e^{-(I-A')K(t-s)} \mathbb{E}[v^c(1)] ds. \quad (3.7)$$

Assume that $v^c(1)$ has finite second moments with covariance matrix $S = -\nabla^2 \psi(0)$, then the conditional covariance matrix is

$$\text{Cov}[\zeta(t) | \zeta_0 = y_0 - \xi_0] = \int_0^t e^{-(I-A')K(t-s)} S e^{-K(I-A)(t-s)} ds. \quad (3.8)$$

Finally, the conditional mean of the relative price is given by

$$\mathbb{E}[e^{\zeta_j(t)} | \zeta_0 = y_0 - \xi_0] = e^{-\psi_t^\zeta(-e^j)}, \quad (3.9)$$

where ψ_t^ζ is defined in (3.5) and e^j is a vector of the canonical basis of \mathbb{R}^N , such that $e_i^j = \delta_{\{j=i\}}$, $i, j = 1, \dots, N$.

Given the characteristic exponent in Proposition 6, exploiting the Cauchy integral formula, we can numerically compute the moments at the desired order by computing an integral in the

²We use the following notation: given a process $x(t)$, its characteristic exponent $\psi_t^x(u)$ satisfies $\mathbb{E} \left[e^{-ux(t)} \right] = e^{-\psi_t^x(u)}$, $u \in \mathbb{C}$.

complex plane, see Choudhury and Lucantoni (1996). This approach avoids expensive calculations of high order derivatives of the characteristic function.

3.2 Stationary distribution of relative log-prices

In this section, we analyze in detail the behavior of the relative log-prices process $\zeta(t)$ as $t \rightarrow \infty$.

Proposition 7. *Let Assumptions 3 and 5 hold. Then, the process of relative log-prices $\zeta(t)$, defined in (3.4), has a stationary distribution with characteristic exponent*

$$\psi^\zeta(u) = \lim_{t \rightarrow +\infty} \psi_t^\zeta(u) = \lim_{t \rightarrow +\infty} \int_0^t \psi \left(e^{-K(I-A)(t-s)} u \right) ds = \int_0^\infty \psi \left(e^{-K(I-A)s} u \right) ds, \quad u \in \mathbb{C}^N, \quad (3.10)$$

where $\psi(u)$ is the characteristic exponent of $v^c(1)$ as defined in equation (3.6).

Proof. See Sato and Yamazato (1984)[Theorem 4.1]. Assumption 3 and the fact that K is a positive diagonal matrix guarantee that $|e^{-(I-A')Kt}| \rightarrow 0$ as $t \rightarrow \infty$, see Kella and Whitt (1999)[Theorem 4.2]. Then the result can be obtained simply taking the limit of equation (3.5) \square

Corollary 8. *Consider the process $\zeta(t)$ satisfying the assumptions of Proposition 7. Assume that $v^c(1)$ has finite first moment given by*

$$\mathbb{E}[v_i^c(1)] = \alpha_i < \infty, \text{ for all } i = 1, \dots, N.$$

Then, the stationary mean of $\zeta(t)$ is given by

$$\tilde{\zeta} := \lim_{t \rightarrow +\infty} \mathbb{E}[\zeta(t)] = ((I - A')K)^{-1} \alpha. \quad (3.11)$$

Corollary 8 shows that supply cost shocks affect the long-run mean of production prices when their expected value is non-zero. Importantly, this result does not imply that individual shocks have a permanent effect. Each realization of the supply shock has a vanishing impact over time due to the mean-reverting structure of the system. However, when stochastic supply shocks occur repeatedly and have a non-zero mean, their cumulative effect shifts the stationary mean of the relative log-price process. In this sense, the impact of supply cost shocks is persistent in expectation, as reflected by the non-zero long-run mean in equation (3.11). By contrast, in the absence of stochastic supply shocks (reducing the model to a deterministic dynamic inoperability framework; see Section 2) production prices revert to their pre-shock equilibrium levels after experiencing a shock at time zero.

Corollary 9. *In addition to the conditions in Proposition 7, assume that $v^c(1)$ has finite second moments with covariance matrix $S = -\nabla^2 \psi(0)$, then the stationary covariance matrix of $\zeta(t)$ is given by*

$$\tilde{S} = \int_0^{+\infty} e^{-(I-A')Ks} S e^{-K(I-A)s} ds, \quad (3.12)$$

and it is the unique solution of

$$K(I - A)\tilde{S} + \tilde{S}(I - A')K = S.$$

In addition, the stationary mean of prices normalized by the equilibrium trend is given by

$$\lim_{t \rightarrow \infty} \mathbb{E} \left[\frac{p_j(t)}{e^{\xi_j(t)}} \right] = \lim_{t \rightarrow \infty} \mathbb{E} \left[e^{\zeta_j(t)} \right] = e^{-\psi^\zeta(-e^j)}, \quad j = 1, \dots, N,$$

see equation (3.9).

4 Likelihood-based inference

In this section, we develop likelihood-based inference for the relative log-price dynamics derived in Section 3. We first derive the exact likelihood of the discretely observed process under general assumptions on the Lévy-driven supply cost shocks, and then introduce a tractable approximate likelihood that is more suitable for empirical implementation.

We suppose that the log-prices $y(t)$ and the exogenous (log-)indexes of economic growth $\xi(t)$ of equations (3.1) are observed at equally spaced discrete times $0 = t_0 < t_1 < \dots < t_M = T$, with $t_{j+1} - t_j = \Delta$ for $j = 0, 1, \dots, M - 1$. Moreover, we assume that the technical coefficient matrix A is known as well. Therefore, we are interested in estimating the dynamics of the relative log-prices $\zeta(t) = y(t) - \xi(t)$ described by the parameter vector θ , which includes the elements of the diagonal matrix K , and the parameters of the Lévy process $v^c(t)$.

Given equation (3.4), we can write

$$\zeta_{j+1} = e^{-(I-A')K\Delta} \zeta_j + v^*(\Delta), \quad (4.1)$$

where using similar reasons as in Valdivieso et al. (2009)[Proposition 3.1],

$$v^*(\Delta) = \int_{t_j}^{t_{j+1}} e^{-(I-A')K(t_{j+1}-s)} dv^c(s) \stackrel{d}{=} \int_0^\Delta e^{-(I-A')K(\Delta-s)} dv^c(s) \quad (4.2)$$

and $\stackrel{d}{=}$ means equal in distribution.

Let $(v^c(t))_{t \geq 0}$ be an \mathbb{R}^N -valued Lévy process with Lévy-Khintchine triplet (γ, Σ, ν) . The law of increment $v^*(\Delta)$ is absolutely continuous with respect to the Lebesgue measure if the Gaussian component is nondegenerate, i.e., Σ is positive definite (or, more generally, $\text{rank}(\Sigma) \geq 1$), in which case $v^c(t)$ admits a smooth density for all $t > 0$ due to convolution with a normal distribution. In the pure jump case ($\Sigma = 0$), absolute continuity fails if the Lévy measure is finite, $\nu(\mathbb{R}) < \infty$, since the process is then compound Poisson and possesses atoms. When $\nu(\mathbb{R}) = \infty$, absolute continuity typically holds under additional non-degeneracy conditions, such as sufficient small jumps activity (e.g. $\int_{|x| < 1} |x| \nu(dx) = \infty$) or the absence of lattice support of ν . These criteria and related refinements are discussed in detail in Sato (1999)[ch. 27, Theorem 27.7 and Theorem 27.10]. Firstly, we present a result for the absolutely continuous case.

Proposition 10. *Let $\zeta_0, \zeta_1, \dots, \zeta_M$ be observations of the process ζ defined in (3.4) and $\zeta(0) = \zeta_0$. Assume that $v^*(\Delta)$ is absolutely continuous with respect to the Lebesgue measure. Then, the likelihood function is given by*

$$L(\theta) = \prod_{j=1}^M f_{v^*} \left(\zeta_j - e^{-(I-A')K\Delta} \zeta_{j-1} \right),$$

where f_{v^*} is the probability density function of $v^*(\Delta)$.

Proof. This follows from Lu (2022)[Lemma 2] and from Lu (2022)[Proposition 2 and Remark 8], setting

$$M_{\zeta_{j-1}}^{-1}(x) = x - e^{-(I-A')K\Delta}\zeta_{j-1},$$

where we define the function $M_{\zeta_{j-1}}$ as

$$\zeta_j = M_{\zeta_{j-1}}(v^*(\Delta)) = e^{-(I-A')K\Delta}\zeta_{j-1} + v^*(\Delta),$$

see (4.1). □

In the interest of application, we specify the process of supply cost per unit of production $v^c(t)$ to be a multivariate compound Poisson process with absolutely continuous jump size distribution. We now extend the previous result to this case.

Theorem 11. *Let $v^c(t)$ be a multivariate compound Poisson process defined as*

$$v^c(t) = \sum_{k=1}^{N(t)} \mathbf{J}_k,$$

where $\mathbf{J}_k = (J_{k,1}, \dots, J_{k,N})'$, $k \geq 1$, are independent and identically distributed random vectors with absolutely continuous density function, and $N(t)$ is a Poisson process with intensity λ . Then,

(i) $v^*(\Delta)$ defined in (4.2) has a mixed discrete-continuous law, that is, with probability $p_0 = e^{-\lambda\Delta}$ then $v^*(\Delta) = 0$, otherwise with probability $1 - p_0$ there is at least one jump and $v^*(\Delta)$ has an absolutely continuous density f_c on \mathbb{R}^N .

(i) The likelihood is

$$\begin{aligned} L(\theta) &= \prod_{j=1}^M f_{v^*(\Delta)}(\zeta_j - e^{-(I-A')K\Delta}\zeta_{j-1}) \\ &= \prod_{j=1}^M \left[p_0 \mathbb{1}_{\{0\}}(\zeta_j - e^{-(I-A')K\Delta}\zeta_{j-1}) \right. \\ &\quad \left. + (1 - p_0) f_c(\zeta_j - e^{-(I-A')K\Delta}\zeta_{j-1}) \mathbb{1}_{\mathbb{R}^N \setminus \{0\}}(\zeta_j - e^{-(I-A')K\Delta}\zeta_{j-1}) \right]. \end{aligned} \tag{4.3}$$

where the density function f_c is obtained by inverting the characteristic function related to the following characteristic exponent

$$\psi_{\Delta}^c(u) = \int_0^{\Delta} \psi_c \left(e^{-K(I-A)(\Delta-s)} u \right) ds,$$

and

$$\psi_c(u) = -\log \left(\frac{\exp(\lambda(\phi_J(u) - 1)) - e^{-\lambda}}{1 - e^{-\lambda}} \right), \quad u \in \mathbb{C}^N,$$

and $\phi_J(u)$ is the multivariate moment generating function of the jumps \mathbf{J}_k .

Proof. From (4.2), we have

$$v^*(\Delta) = e^{-(I-A')K\Delta} \sum_{k=1}^{N(\Delta)} e^{(I-A')KT_k} \mathbf{J}_k,$$

where, conditional on having $N(\Delta)$ jumps, the jump times $T_1, \dots, T_{N(\Delta)}$ are i.i.d $\text{Unif}([0, \Delta])$ random variables and \mathbf{J}_k is independent of $N(\Delta)$. Then, the mixed discrete-continuous law of $v^*(\Delta)$ follows.

Following Lu (2022)[Corollary 1], we introduce the dominating measure

$$\mu := \delta_0 + L^N, \quad (4.4)$$

where δ_0 is the Dirac measure at zero and L^N denotes the Lebesgue measure on \mathbb{R}^N . Then $v^*(\Delta)$ admits the Radon–Nikodym derivative

$$f_{v^*(\Delta)}(z; \theta) := \frac{d\mathbb{P}_{v^*(\Delta)}}{d\mu}(z) = p_0 \mathbb{1}_{\{0\}}(z) + (1 - p_0) f_c(z; \theta) \mathbb{1}_{\mathbb{R}^N \setminus \{0\}}(z), \quad z \in \mathbb{R}^N, \quad (4.5)$$

where f_c depends on θ through λ and the jump size distribution. As in the previous proof, we define

$$\zeta_j = M_{\zeta_{j-1}}(v^*(\Delta)) = e^{-(I-A')K\Delta} \zeta_{j-1} + v^*(\Delta),$$

see equation (4.1), and, noticing that the dominating measure μ is independent of the parameters θ , Lu (2022)[Lemma 2] gives

$$L(\theta) = \prod_{j=1}^M f_{v^*(\Delta)}(\zeta_j - e^{-(I-A')K\Delta} \zeta_{j-1}; \theta).$$

Using the definition of $f_{v^*(\Delta)}$, we get equation (4.3). \square

While the exact likelihood above defined provides a fully consistent description of the discretely observed process under general Lévy-driven supply shocks, its numerical implementation may be challenging in practice, especially for large networks. To obtain a tractable likelihood based on Euler discretization, we consider the discrete-time approximation of the dynamics in (3.3) over the observation grid $\{t_j = j\Delta\}_{j=0, \dots, M}$, which reads

$$\zeta_j - \zeta_{j-1} = -(I - A')K \zeta_{j-1} \Delta + \int_{t_{j-1}}^{t_j} dv^c(s), \quad j = 1, \dots, M. \quad (4.6)$$

We define the Euler residual

$$r_j := \zeta_j - \zeta_{j-1} + (I - A')K \zeta_{j-1} \Delta, \quad j = 1, \dots, M,$$

which serves as an approximation of the innovation over the interval $[t_{j-1}, t_j]$. The Euler discretization itself does not rely on any distributional assumption on the jump sizes. However, to obtain a closed-form expression for the likelihood, we assume that the jump sizes of the Lévy process $v^c(t)$ are mutually independent Gaussian random variables with mean η_i and standard deviation σ_i , $i = 1, \dots, N$. Under this assumption, conditional on observing k jumps over an

interval of length Δ , the innovation r_j is Gaussian with mean

$$m_k := k \eta$$

and covariance matrix

$$\Sigma_k := k \operatorname{diag}(\sigma_1^2, \dots, \sigma_N^2).$$

As a consequence, the Euler (approximated) likelihood function can be written as

$$L^E(\theta) = \prod_{j=1}^M \left[e^{-\lambda\Delta} \mathbb{1}_{\{0\}}(r_j) + \sum_{k=1}^{\infty} \frac{(\lambda\Delta)^k e^{-\lambda\Delta}}{k!} g(r_j; m_k, \Sigma_k) \right], \quad (4.7)$$

where $g(\cdot; m_k, \Sigma_k)$ denotes the density of a multivariate Gaussian distribution with mean m_k and covariance matrix Σ_k .

4.1 Numerical implementation and regularization

In both the exact likelihood (4.3) and the Euler likelihood associated with the Gaussian jump size specification (4.7), the innovation over each sampling interval follows a mixed discrete-continuous distribution, with a point mass at zero corresponding to the event of no jumps. From a practical standpoint, the presence of the indicator function $\mathbb{1}_{\{0\}}(\cdot)$ may lead to numerical issues, since exact zeros are rarely observed in finite-precision data or in the presence of measurement noise.

To address this problem, we adopt a common regularization strategy in which the Dirac mass at zero is replaced by a narrow Gaussian density with mean zero and standard deviation $\varepsilon > 0$. This regularization is applied uniformly to both likelihoods, and can be interpreted either as a numerical smoothing device or as the introduction of a small measurement error in the observed increments.

Specifically, we approximate the Dirac measure δ_0 by the density

$$\varphi_\varepsilon(z) := \frac{1}{(2\pi\varepsilon^2)^{N/2}} \exp\left(-\frac{\|z\|^2}{2\varepsilon^2}\right), \quad z \in \mathbb{R}^N,$$

where $\varepsilon \in \mathbb{R}_+$. Equivalently, this amounts to replacing the indicator function $\mathbb{1}_{\{0\}}(\cdot)$ in the Radon–Nikodym derivative with $\varphi_\varepsilon(\cdot)$. Replacing $\mathbb{1}_{\{0\}}(\cdot)$ with $\varphi_\varepsilon(\cdot)$ yields a regularized likelihood that is absolutely continuous with respect to the Lebesgue measure, regardless of whether the transition density is obtained from the exact continuous-time model or from its Euler discretization under Gaussian jump sizes. As $\varepsilon \rightarrow 0$, the regularized likelihood converges pointwise to the original mixed discrete-continuous likelihood. In practice, the choice of ε entails a numerical trade-off, that is discussed for instance in literature concerning the estimation of processes of jump-diffusion type, see Kiefer (1978), (Hamilton, 1994, page 689) and Ramezani and Zeng (2007). When ε is too small, the regularized likelihood becomes sharply peaked around zero and may lead to numerical issues, reintroducing the numerical ill-conditioning associated with the Dirac mass representation. Conversely, excessively large values of ε may induce systematic estimation distortions by oversmoothing the no-jump component, thereby biasing the parameter estimates. For this reason, we select ε by minimizing an accuracy error as discussed in the next section. The selected value represents an acceptable trade-off between numerical stability and

statistical accuracy, allowing for the construction of standard errors via the information matrix.

4.2 Simulation experiment: exact versus Euler likelihood in a two-sector model

We now assess the finite-sample performance of the proposed likelihood-based estimators by means of a Monte Carlo experiment in a stylized two-sector economy ($N = 2$). The technical coefficient matrix A is assumed to be known, while the diagonal mean-reversion matrix $K = \text{diag}(k_1, k_2)$ and the parameters of the supply-cost shock process are estimated. More precisely, we generate relative log-prices $\{\zeta(t)\}_{t \geq 0}$ from the Lévy-driven network Ornstein–Uhlenbeck dynamics in (3.3). Supply-cost shocks are specified as a bivariate compound Poisson process

$$v^c(t) = \sum_{m=1}^{N(t)} \mathbf{J}_m, \quad N(t) \sim \text{Poisson}(\lambda t), \quad (4.8)$$

where the jump sizes $\mathbf{J}_m = (J_{m,1}, J_{m,2})'$ are i.i.d. Gaussian random vectors with mean $\eta = (\eta_1, \eta_2)'$ and covariance matrix $\Sigma_J = \text{diag}(\sigma_1^2, \sigma_2^2)$.³ Given a parameter vector

$$\theta_0 = (k_1, k_2, \lambda, \eta_1, \eta_2, \sigma_1, \sigma_2)$$

we simulate the continuous-time process exactly by drawing jump times and jump sizes and updating $\zeta(t)$ between jumps according to the deterministic flow implied by the matrix $(I - A')K$.

For each replication, we observe the process at equally spaced times $t_j = j\Delta$, $j = 0, 1, \dots, M$, over a fixed horizon $T = M\Delta$. The resulting sample is $\{\zeta_j\}_{j=0}^M$ with $\zeta_j = \zeta(t_j)$. For each simulated dataset we estimate θ using the two likelihoods (4.3) and (4.7), both implemented with the regularization described in Section 4.1, that is, replacing the Dirac component with a narrow Gaussian kernel of variance ε .

Dealing with performance measure, let $\hat{\theta}^{(r)}$ denote the estimator obtained in replication $r = 1, \dots, R$. We summarize estimation accuracy using root mean squared error (RMSE), computed in levels for all parameters. In particular, for the mean-reversion matrix K we compute

$$e_K^{(r)} = \left\| \hat{k}^{(r)} - k_0 \right\|_2, \quad \hat{k}^{(r)} = (\hat{k}_1^{(r)}, \hat{k}_2^{(r)})'$$

For the jump intensity and jump size parameters we define

$$e_\lambda^{(r)} = |\hat{\lambda}^{(r)} - \lambda_0|, \quad e_\eta^{(r)} = \left\| \hat{\eta}^{(r)} - \eta_0 \right\|_2, \quad e_\sigma^{(r)} = \left\| \hat{\sigma}^{(r)} - \sigma_0 \right\|_2,$$

where $\hat{\eta}^{(r)} = (\hat{\eta}_1^{(r)}, \hat{\eta}_2^{(r)})'$ and $\hat{\sigma}^{(r)} = (\hat{\sigma}_1^{(r)}, \hat{\sigma}_2^{(r)})'$. Here the subscript 0 corresponds to the true parameters.

As a numerical test, we set $\Delta = 1$ and $M = 60$. The technical coefficient matrix is fixed and given by

$$A = \begin{pmatrix} 0.20 & 0.15 \\ 0.12 & 0.08 \end{pmatrix}.$$

³For simplicity we assume independent jump components across sectors; allowing for cross-sectional correlation in Σ_J is straightforward.




Figure 2: RMSEs of the estimated parameters as functions of the regularization parameter ε (log-scale, Euler likelihood).

To select the regularization parameter ε , we construct an aggregate accuracy error that combines the RMSEs across parameter groups. Specifically, letting

$$\text{RMSE}(x) = \sqrt{\frac{1}{R} \sum_{r=1}^R (e_x^{(r)})^2}, \quad x \in \{K, \lambda, \eta, \sigma\},$$

we define the normalized aggregate error

$$\text{AE}(\varepsilon) = \frac{\text{RMSE}(K)}{\|k_0\|_2} + \frac{\text{RMSE}(\lambda)}{|\lambda_0|} + \frac{\text{RMSE}(\eta)}{\|\eta_0\|_2} + \frac{\text{RMSE}(\sigma)}{\|\sigma_0\|_2}, \quad (4.9)$$

we then select ε by minimizing $\text{AE}(\varepsilon)$ using the Euler likelihood. Figure 2 reports the RMSEs as functions of ε . The selected value $\varepsilon = 0.01$ corresponds to the minimizer of the aggregate error defined in Equation (4.9).

In Table 1 we report the RMSEs and the average values, that is,

$$\text{average}(x) = \frac{1}{R} \sum_{r=1}^R x^{(r)}$$

for $x \in \{K, \lambda, \eta, \sigma\}$, separately for the exact and the Euler likelihood estimators, together with their average computation times, based on $R = 50$ replications. The table compares the finite-sample performance of the exact likelihood and the Euler likelihood in this two-dimensional setting. Overall, both estimators deliver accurate parameter estimates, with RMSEs of comparable magnitude across all parameter groups. The most striking difference concerns computational cost. Evaluating the exact likelihood requires the numerical inversion of a two-dimensional characteristic function in order to recover the transition density at each observation, which results in a substantially higher runtime. In contrast, the Euler likelihood relies on a closed-form mixture

Table 1: Average values and RMSEs (50 replications): exact vs Euler likelihood

		Average		RMSE	
	True parameters	Exact	Euler	Exact	Euler
K	(0.05, 0.10)	(0.0501, 0.1001)	(0.0499, 0.0998)	0.0014	0.0009
λ	2	1.9248	1.9606	0.2612	0.3878
η	(0.10, 0.07)	(0.1028, 0.0723)	(0.1018, 0.0721)	0.0169	0.0206
σ	(0.08, 0.05)	(0.0757, 0.0460)	(0.0780, 0.0509)	0.0217	0.0142
Avg. time per replication (s)				179.84	0.69

Table 2: Euler likelihood performance for different sample lengths

	Sample length T			
	$T = 30$	$T = 60$	$T = 90$	$T = 120$
RMSE(K)	0.0064	0.0009	0.0007	0.0006
RMSE(λ)	0.6118	0.3878	0.2742	0.2218
RMSE(η)	0.0309	0.0206	0.0163	0.0133
RMSE(σ)	0.0208	0.0142	0.0117	0.0097
Avg. time per replication (s)	0.39	0.69	1.09	1.67

representation and can be evaluated very efficiently. As a result, the average computation time per replication is roughly two orders of magnitude larger for the exact likelihood than for its Euler counterpart.

Table 2 reports the performance of the Euler likelihood estimator as a function of the sample length T , keeping Δ fixed. As expected, estimation accuracy improves monotonically with T for all parameters, reflecting the increase in the effective number of observations. At the same time, the computational cost increases with the sample length, remaining however well below two seconds per replication, highlighting the excellent scalability of the Euler likelihood.

As an additional robustness check, we evaluate the performance of the Euler likelihood in a setting that closely mimics the empirical application. Specifically, we set the true parameter values equal to the estimates reported in Table 4 in the following section and consider the technical coefficient matrix A corresponding to the year 2023. We generate $R = 50$ synthetic samples of length $T = 30$ with quarterly frequency and estimate the model using the Euler likelihood. In this real-case calibration experiment, the Euler likelihood yields RMSEs equal to 0.0751 for the resilience vector K , 0.6854 for the jump intensity λ , 0.0025 for the jump-size mean η , and 0.0088 for the jump-size dispersion σ . The average computation time per replication for this six-dimension setting is approximately 19 seconds, reflecting the increased dimensionality and complexity of the empirical network compared to the stylized two-sector setting. Overall, these results confirm the practical feasibility of the Euler likelihood in empirically calibrated production networks.

Taken together, these results indicate that the Euler likelihood achieves a favorable trade-off between statistical accuracy and computational efficiency. While the exact likelihood provides a useful benchmark in low dimensions, its reliance on repeated numerical inversions of multivariate characteristic functions makes it computationally prohibitive in higher-dimensional settings. For

Code	Description
11	Agriculture, forestry, fishing, and hunting
211	Oil and Gas Extraction
212	Mining (Except Oil and Gas)
22	Utilities
23	Construction
31G	Manufacturing

Table 3: Selected NAICS sectors for the analysis of Section 5.2.

this reason, and in view of empirical applications involving larger production networks, we rely on the Euler likelihood in the analysis of real data in the next section.

5 Empirical Analysis

In this section, we estimate a specification of the multivariate production prices model presented in Section 2 using economic data, including IO tables and historical series of industrial production prices. In particular, we assess sector-specific resilience rates, which capture the speed at which each sector returns to its equilibrium state following supply shocks. In addition, we examine the long-term effects of persistent supply shocks by analyzing the stationary distribution of production prices.

5.1 Data

As log-prices we consider the logarithms of the US chain-type price indexes for gross output by industry (Bureau of Economic Analysis, 2025b), while as indexes of economic growth we consider the historical real log-GDP (Bureau of Economic Analysis, 2025c), as standard in the literature see Carvalho and Tahbaz-Salehi (2019). We focus on the six NAICS sectors listed in Table 3, several of which have been extensively studied in the supply chain literature, beginning with the foundational work of Long Jr and Plosser (1983) and followed by subsequent contributions. In particular, in order to investigate the role of the supply shocks of fossil fuels we consider the oil and gas extraction sector separately from the other mining activities. The total value added of the sectors of Table 3 for year 2023 approximately contributes to the 18% of the national GDP, representing a relevant share of the total production. Historical relative log-prices, computed as log-prices minus real log-GDP, see equation (3.1), for period 2010-2024, are displayed in the left-hand panel of Figure 3. Data are expressed in quarterly terms and using 2017 as base year: that is, all values are scaled in such a way that the average of the 2017 values equals 1.

In addition, following equation (2.2), we use the input–output matrices Z and the total production vectors x for each year of our analysis to recover the technical coefficient matrices A (data sourced from Bureau of Economic Analysis (2025a)). Therefore, in our empirical analysis, the technical coefficient matrix is allowed to vary over time, and the production network underlying the model becomes time dependent. Allowing for a time-varying production network raises additional analytical and computational issues. While the continuous-time dynamics of relative log-prices remain well defined under mild regularity conditions, the closed-form solution available in the constant-coefficient case no longer applies. In particular, when $A(t)$ varies




Figure 3: Left-hand panel: US historical relative log-prices with base year 2017; period: 2010-Q1 to 2024-Q4 in quarterly steps. Relative log-prices computed as producer log-price indexes minus real log-GDP (sourced from Bureau of Economic Analysis, 2025b,c). Right-hand panel: technical coefficient functions estimated via Nadaraya-Watson kernel regression (Nadaraya, 1964; Watson, 1964) using yearly data on technical coefficients (denoted with circles) between sector 211 (oil and gas extraction) and the other sectors. Yearly technical coefficients computed as in equation (2.2), for the years 2010, \dots , 2023 (input-output data sourced from Bureau of Economic Analysis, 2025a). Sectors reported in Table 3.

over time, the solution of the stochastic integral equation (3.3) involves a time-ordered matrix exponential (Lam, 1998). As a consequence, the explicit representation of the solution given in equation (3.4), which holds when A is constant, does not extend trivially to the time-dependent case. Consequently, extending the exact likelihood (4.3) derived under constant A in Section 4 to the case of time-varying production networks is non-trivial and computationally demanding. In particular, the presence of the time-ordered exponential rules out a simple expression for the one-step transition density and would require substantially more complex numerical methods.

From a theoretical perspective, however, the continuous-time model remains well posed. When the technical coefficient functions $A_{i,j}(t)$ are Lipschitz continuous in time, the stochastic differential equation associated with (3.3) admits a unique strong solution. This follows from standard results for stochastic differential equations with time-dependent coefficients; see, for example, Protter (1990, Chapter V, Section 3). Importantly, the Euler likelihood (4.7) does not depend on the explicit form of the continuous-time solution and therefore remains applicable when the technical coefficient matrix varies over time.

To ensure the Lipschitz continuity condition, we estimate the smooth technical coefficient functions $A_{i,j}(t)$, $i, j = 1, \dots, N$, using kernel regression techniques (Nadaraya, 1964; Watson, 1964). Specifically, let $\hat{A}_{i,j}^l$, for $l = 1, \dots, L$, denote the yearly technical coefficients obtained from the BEA databases, where L is the number of years considered. We then compute the pairwise technical coefficient functions as

$$A_{i,j}(t) = A_{i,j}(t; h) = \frac{\sum_{l=1}^L g_h(t - t_l) \hat{A}_{i,j}^l}{\sum_{l=1}^L g_h(t - t_l)}, \quad i, j = 1, \dots, N,$$

where $g_h(t) = \frac{1}{h} g\left(\frac{t}{h}\right)$, $g(\cdot)$ is a kernel function, and h is a smoothing parameter. Following standard practice we select the Gaussian kernel $g(t) = e^{-t^2}$, and we set the bandwidth size

$h = 0.49$ as a result of a modified mean squared error minimization as illustrated in Härdle et al. (1988). Specifically, we compute

$$h = \underset{\tilde{h} > 0}{\operatorname{argmin}} \frac{1}{L} \Xi(L, \tilde{h}) \sum_{l=1}^L \left(A_{i,j}(t_l; \tilde{h}) - \hat{A}_{i,j}^l \right)^2,$$

where $\Xi(L, \tilde{h}) = \left(1 - \frac{g(0)}{L\tilde{h}}\right)^{-2}$ is a generalized cross-validation correction factor (see Wahba, 1975).

The right-hand panel of Figure 3 shows the yearly technical coefficients, and the corresponding technical coefficient smooth functions, between the oil and gas extraction sector and the other sectors for the considered period.

5.2 Estimation results

In the interest of application, following equation (4.8), we specify the process of supply cost per unit of production $v^c(t)$ to be a multivariate compound Poisson process with components

$$v_i^c(t) = \sum_{k=1}^{N(t)} J_{i,k}, \quad i = 1, 2, \dots, N,$$

where $(J_{1,k}, \dots, J_{N,k})'$, $k \geq 1$, are independent and identically distributed random vectors with independent components and Gaussian distribution, and $N(t)$ is a Poisson process with intensity λ . In this case, the stationary mean of relative log-prices $\zeta(t)$ given by equation (3.11) reads

$$\tilde{\zeta} = \lambda \left((I - A')K \right)^{-1} \eta, \quad (5.1)$$

η being the vector having as i -th component $\mathbb{E}[J_i]$. According to the above specification, the time of the shocks are the same across all sectors, while the sizes of the shocks are idiosyncratic. The assumption that shock times are common across sectors is made to reflect the idea of economy-wide supply disruptions that affect multiple sectors simultaneously. As documented by Ciravegna et al. (2023), the contemporary economic environment is characterized by the multiplicity and simultaneity of crises, such as pandemics, armed conflicts, energy and climate challenges and financial turmoil. Under this assumption, heterogeneity across sectors arises from differences in shock magnitudes, from network propagation through the IO structure, and from sector-specific resilience rates. We highlight that the choice of a compound Poisson process is due to its piecewise-constant nature, describing the arrival of shocks at discrete time in our continuous-time framework.

We suppose that the technical coefficient matrices $A(t_j)$, $j = 1, \dots, M-1$, are known as well. Therefore, we are interested in estimating the dynamics of the relative log-prices $\zeta(t) = y(t) - \xi(t)$ described by the parameter vector θ , which includes the parameters that determine the diagonal matrix K , the distribution of the jump sizes J , and the intensity λ .

Table 4 presents the estimated parameters, with their standard errors reported in parentheses. Each standard error is computed as the square root of the corresponding Cramér-Rao lower bound of the estimator variance. All optimizations are carried out using the FMINCON function in MATLAB. To enhance the likelihood of reaching the global optimum, we execute the

Sector	K	η	σ
11	0.0509 (0.0528)	0.0015 (0.0046)	0.0196 (0.006)
211	0.0817 (0.06)	0.0002 (0.0096)	0.0623 (0.0232)
212	0.0363 (0.0303)	0.0012 (0.0026)	0.0123 (0.0034)
22	0.2815 (0.0788)	0.0042 (0.003)	0.0183 (0.0052)
23	0.0191 (0.0369)	0.0009 (0.0016)	0.0069 (0.0017)
31G	0.0599 (0.0786)	-0.0001 (0.003)	0.0092 (0.003)
λ	3.1742 (2.1353)		

Table 4: Estimated parameters and their standard deviations (in brackets) of the relative log-price process (equation (3.4)); standard deviations computed as the square roots of the Cramér-Rao lower bounds of the respective parameters. Sectors reported in Table 3. Estimation period: 2010-Q1 to 2024-Q4 in quarterly steps.

FMINCON routine from multiple randomly generated initial points and retain the best solution among these runs. However, we find that the resulting optimal solutions are generally robust to the choice of initial parameter values. Moreover, we fix ε to the tuned value obtained from the simulation experiment (0.01), see Section 4.2, and we verified that results are robust for ε in a neighborhood of such a value.

The resilience (or production adjustment) rates, K_{jj} , $j = 1, \dots, N$, determine –together with cross-sectoral effects– the speed at which a sector returns to its equilibrium state following an external direct or internal indirect shock. In our analysis, the estimated resilience rates range approximately between 0.02 and 0.28 (see Table 4), which correspond to the case of the construction (code 23) and utilities sectors (22), respectively. In the case of the construction sector, Figure 3 shows that the only significant jump occurs around the onset of the war in Ukraine. Following this event, the relative log-price does not revert to the equilibrium value, which is justified by the very low resilience rate (0.0191). The other sectors have a much larger resilience rate, and we observe a stronger tendency to revert to the zero equilibrium level. For example, the utilities sector (22) displays a more pronounced mean-reverting behavior: the two positive shocks appearing in Figure 3 –at the beginning of the previous decade and during the war in Ukraine– are absorbed although at different paces. Similarly, despite the substantial fluctuations experienced by the oil and gas extraction industry (211), this sector exhibits strong mean-reverting behavior, as indicated by the estimated second-highest resilience rate (0.0817).

Concerning the supply shock process, recalling that time is measured in quarters, the estimated value of λ indicates that there are on average approximately three jumps per quarter. Given that the model does not include a continuous component, all variations of the process occur through jumps. As a consequence, the estimated jump intensity also captures small, frequent adjustments that would otherwise be absorbed by a diffusive term in a jump-diffusion specification. Moreover, while the average number of jumps per quarter is approximately three, the estimated expected jump size η is relatively small. Hence, the implied dynamics are not charac-




Figure 4: Visualization of the flow network structure.

terized by large, infrequent discontinuities, but rather by frequent jumps of modest magnitude, which together reproduce the observed variability of the data. Comparing the different sectors, the most notable differences are observed in the diffusion coefficients σ . The pronounced fluctuations in the oil and gas extraction sector, as shown in Figure 3, result in a significantly higher estimate of σ —approximately ten times greater than that of the construction industry, which has the lowest diffusion coefficient. Finally, the jump size means η are positive in most cases, highlighting that price oscillations are often due to supply shortages that generate increases in the price levels.

5.2.1 Analysis of shock propagation

Based on the estimation results, we investigate how the shock produced by an industry is transmitted to the network. In the Figure 4 we show a network in which the nodes correspond to the industrial sectors and the thickness of the edges quantifies the interconnections between sectors, given by the off-diagonal element of matrix $e^{-(I-A')K}$. For this analysis, we use a constant matrix A corresponding to the most recent (i.e., year 2023) technical coefficient matrix available from the Bureau of Economic Analysis (BEA) database.

Given the network structure, we mimic the DIIM model (see Section 2) by computing the shock propagation over time produced by an initial jump of the oil and gas extraction industry (211) without any further shocks thereafter; we show it in Figure 5. In particular, we fix the shock size equal to $\bar{v}_{211} := \lambda(\eta_{211} + z\sigma_{211})$, where λ and η_{211} are given in Table 4 and $z = 2.34$ is fixed equal to the 99% quantile of the standard normal distribution, and compute the deterministic time propagation of the shock across sectors as $e^{-(I-A')Kt} \bar{v}_{211}^0$, $t \geq 0$, where $\bar{v}_{211}^0 = (0, \dots, \bar{v}_{211}, \dots, 0)'$.

As shown by the thickness of the arrows connecting the sectors in the left panel of Figure 5, the oil and gas extraction sector (211) has a strong direct impact on the manufacturing industry (31G), while its influence on the construction sector (23) is relatively limited. However, the construction sector is notably affected by the manufacturing sector, suggesting an indirect transmission of the initial shock in the relative log-price of oil and gas extraction. This propagation effect is also evident in the right panel of Figure 5, where the manufacturing (31G) and utilities

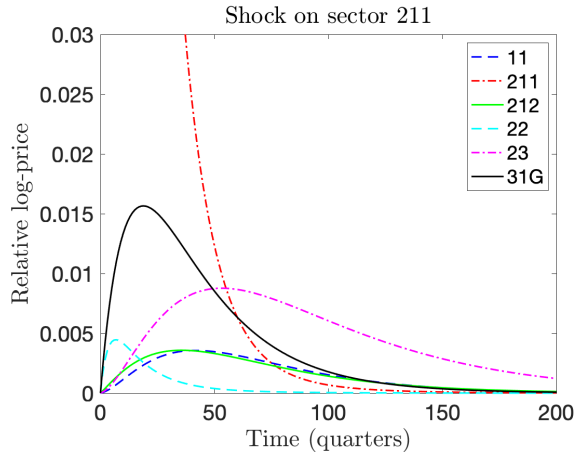


Figure 5: Deterministic shock propagation after an initial jump in the oil and gas extraction sector (211). Initial shock set as $\lambda(\eta_{211} + z\sigma_{211})$, where $z = 2.34$ is the 99% quantile of the standard normal distribution.

(22) sectors are the first to react to the shock, in line with the technical coefficients illustrated in Figure 3. The impact then spreads to the other sectors over time. Notably, the construction sector exhibits a slower and incomplete recovery compared to the others. The following section explores in greater detail the persistent effects of such supply shocks.

5.2.2 Stationary distribution

To assess the long run persistence of supply shocks on the industrial network, we further examine the stationary distribution of the relative log-prices on the basis of the results offered in Section 3.2. The computations are made assuming again that the 2023 technical coefficient matrix A remains stable in the future. In Table 5 we report the first four standardized marginal moments and the correlation matrix of the stationary distribution. The stationary mean vector $\tilde{\zeta}$ is computed as in (3.11), the standard deviation vector as $\tilde{\zeta}_s = (\sqrt{\tilde{S}_{11}}, \dots, \sqrt{\tilde{S}_{NN}})$, where \tilde{S} is recovered through numerical integration of equation (3.12), and the correlation matrix reads $\tilde{C} = \text{diag}(\tilde{\zeta}_s)^{-1} \tilde{S} \text{diag}(\tilde{\zeta}_s)^{-1}$. The skewness and excess kurtosis coefficients are computed by numerical integration of the density function, that is obtained through numerical inversion of the characteristic function, see equation (3.10) for the characteristic exponent, by using the COS method of Fang and Oosterlee (2009).

It can be seen from Table 5 that all the stationary means are positive. We note that even if the mean jump size of a particular sector can be negative –as it is for the manufacturing industry (code 31G, see Table 4)– the prevalence of positive shocks induced by the other sectors produces a positive stationary mean across all marginal distributions of the multivariate process. Moreover, it is interesting to observe that large long-term deviations are not necessarily given by sectors that exhibit large jumps. For example, the construction industry (23) has a relatively small jump mean and standard deviation (displayed in Table 4); however, due to its low resilience rate and its interconnections with the other industries, it largely embeds the long-term effects of shock propagation affecting the network.

In addition, given our assumption that the supply shocks are of Lévy type, the stationary distribution is infinitely divisible, heavy-tailed, with tails determined by the jumps distribution

\tilde{C}	11	211	212	22	23	31G
11						
211	0.0178					
212	0.0356	0.0303				
22	0.0199	0.0822	0.0325			
23	0.1031	0.0571	0.1091	0.029		
31G	0.1678	0.218	0.1105	0.0449	0.3549	
$\tilde{\zeta}$	0.1396	0.0137	0.1239	0.05	0.1605	0.0098
$\tilde{\zeta}_s$	0.1325	0.287	0.0862	0.0459	0.0779	0.0657
Skew	0.0227	0.0018	0.0282	0.1749	0.0196	0.0001
ExKrt	0.0327	0.0703	0.0305	0.2497	0.0117	0.0286

Table 5: Mean $\tilde{\zeta}$, standard deviation $\tilde{\zeta}_s$, skewness (Skew), excess kurtosis (ExKrt), and correlation matrix \tilde{C} of the stationary distribution of the relative log-price process $\zeta(t)$ (see equation (3.3)).




Figure 6: Stationary log-densities of the relative log-price dynamics of the manufacturing (code 31G) and utilities (22) industries, respectively, and corresponding Gaussian log-densities with the same means and standard deviations. Densities recovered using the COS method of Fang and Oosterlee (2009) for numerical inversion of the characteristic function (see exponent in (3.10)). Bounds of the x-axis corresponding to ten standard deviations away from the mean relative log-price of the respective sector.

and the arrival rate and asymmetric, unless the jump size mean η is zero. These last two points are indeed confirmed by the nonzero estimates of the skewness and excess kurtosis coefficients of Table 5, that however appear to have some variability across sectors.

To assess deviations from Gaussianity in the distribution of relative log-prices, Figure 6 presents illustrative examples of marginal log-densities for the manufacturing and utilities sectors—two industries directly impacted by oil and gas extraction (as indicated by the technical coefficients in Figure 3), where the most pronounced price fluctuations are observed. By comparing these empirical log-densities to Gaussian counterparts with matching means and standard deviations, we observe notable discrepancies in the tails for the utilities sector (22), suggesting heavier tails relative to the normal distribution. In contrast, the manufacturing sector (31G) shows a logarithmic density that closely aligns with the Gaussian benchmark, consistent with the very small skewness and excess kurtosis estimates in Table 5. These results imply that extreme price events are more probable in the utilities sector than in the manufacturing sector.

6 Risk management tools for direct and indirect supply shocks

In this section, we discuss various tools for managing direct and indirect supply cost risks within a production network. We demonstrate that our model can serve as a valuable tool for central controllers, such as governments or supranational entities, to implement measures aimed at mitigating the impact of supply cost shocks on production prices and, consequently, on final consumer prices. For instance, a central controller can allocate resources and efforts to enhance the resilience rate of a specific sector, denoted as K_{jj} for a given sector j , with the goal of reducing the long-term impact of shocks on prices, as outlined in Section 3.2. Alternatively, it could also intervene on the structure of the network, with policies, for example, that discourage dependence on the extraction/importation of fossil fuels. To evaluate these effects, we begin by analyzing the systemic risk inherent in the production network. Specifically, we adopt the framework proposed by Adrian and Brunnermeier (2016), computing the tail risk of each sector conditional on distress occurring in another sector. Subsequently, we conduct a sensitivity analysis to assess the influence of network linkages, represented by the matrix of technical coefficients A and resilience rates, captured by the matrix K , on the short and long-term distributions of production prices. For this, we follow the methodology of Tsanakas and Millossovich (2016), which involves computing the directional derivatives of risk measures with respect to model inputs, thereby quantifying the marginal impact of changes in A and K on log-price risk.

6.1 Analysis of the systemic tail risk

The Conditional Value at Risk (CoVaR), introduced by Adrian and Brunnermeier (2016), is a systemic risk measure used in finance. CoVaR measures the tail risk of a system (or of an industrial sector) conditional on a particular sector being in distress. We denote as ζ^∞ the random vector distributed according to the stationary law of the relative log-price $\zeta(t)$. We recall that, for each marginal density of the relative log-price ζ_i of the i -th sector, the Value at Risk (VaR) at level α is implicitly defined as the α -quantile, i.e.,

$$\text{VaR}_\alpha^i = \inf\{x \in \mathbb{R}, \mathbb{P}(\zeta_i^\infty \leq x) \geq \alpha\}. \quad (6.1)$$

The existence of the stationary distribution of the relative log-price $\zeta(t)$ is discussed in Section 3.2. The ΔVaR at level α is the difference between the VaR at level α and the median of the distribution, i.e.

$$\Delta\text{VaR}_\alpha^i = \text{VaR}_\alpha^i - \text{VaR}_{0.5}^i. \quad (6.2)$$

The $\text{CoVaR}_{\alpha,\beta}(i, j)$ is the VaR at level α of the sector i , conditional on a shock in the sector j (e.g., at its own VaR_β level). The CoVaR is defined as follows, see Adrian and Brunnermeier (2016).

Definition 12. *The $\text{CoVaR}_{\alpha,\beta}(i, j)$ is defined as the α -quantile of the conditional probability distribution*

$$\text{CoVaR}_{\alpha,\beta}(i, j) = \inf\left\{x \in \mathbb{R}, \mathbb{P}\left(\zeta_i^\infty \leq x \mid \zeta_j^\infty = \text{VaR}_\beta^j\right) \geq \alpha\right\}. \quad (6.3)$$

The $\Delta\text{CoVaR}_\alpha(i, j)$ is the difference between the $\text{CoVaR}_{\alpha,\alpha}(i, j)$ when sector j is in distress

and $\text{CoVaR}_{\alpha,0.5}(i, j)$ in normal conditions, that is, when the sector j is at median level,

$$\Delta\text{CoVaR}_{\alpha}(i, j) = \text{CoVaR}_{\alpha,\alpha}(i, j) - \text{CoVaR}_{\alpha,0.5}(i, j). \quad (6.4)$$

It is used to assess how much tail risk a single sector contributes to another. The diagonal elements of the ΔCoVaR matrix represent the marginal ΔVaR for each sector. The marginal VaRs and ΔVaR s are computed by numerically integrating the marginal density functions obtained via numerical inversion of the characteristic functions, applying the COS method of Fang and Oosterlee (2009). In particular, we use the characteristic exponents of ζ^{∞} given in equation (3.10). The ΔCoVaR matrix is estimated using Monte Carlo simulations of $\zeta(t)$ combined with quantile regression. Simulations are generated using the exact discretization scheme in equation (4.1), with

$$v^*(\Delta) = e^{-(I-A')K\Delta} \sum_{k=1}^{N(\Delta)} e^{(I-A')KT_k} \mathbf{J}_k,$$

where, conditional on $N(\Delta)$ jumps, the jump times $T_1, \dots, T_{N(\Delta)}$ are i.i.d. $\text{Unif}([0, \Delta])$ random variables, and the jump sizes \mathbf{J}_k are multivariate normal random vectors independent of $N(\Delta)$. The specification of the supply cost per unit of production $v^c(t)$ is detailed in Section 5.2. For each pair of simulated $(\zeta_j(t), \zeta_i(t))$, we perform a quantile regression. In linear models, quantile regression reduces to a linear programming problem and is a well-established methodology in economics and finance for estimating conditional risk measures such as VaR and CoVaR (see, e.g., Koenker, 2017, Barrera et al., 2026), without assuming any explicit parametric form for the joint distribution. Specifically, quantile regression is used to estimate the coefficient $\beta_{\alpha}^{i,j}$ in the linear relationship

$$\Delta\text{CoVaR}_{\alpha}(i, j) = \beta_{\alpha}^{i,j} \Delta\text{VaR}_{\alpha}^j.$$

Table 6 reports the ΔCoVaR matrix of joint stationary distribution of relative log-prices. From Table 6, we observe that the highest off-diagonal ΔCoVaR value is $\Delta\text{CoVaR}(31\text{G}, 211) = 0.1879$. This indicates that the tail risk of the manufacturing sector (31G) is significantly influenced by price shocks in the oil and gas extraction sector (211). Similarly, the tail risks of the utilities (22) and construction (23) sectors are also notably driven by shocks in the oil and gas extraction sector, with ΔCoVaR values of 0.1283 and 0.0768, respectively. Additionally, the tail risk of the manufacturing sector (31G) is affected by the agricultural sector (11), with $\Delta\text{CoVaR}(31\text{G}, 11) = 0.0530$. Moreover, the tail risk in the construction sector (23) appears to be influenced by the manufacturing sector (31G) and vice versa. These results are derived under the assumption that the technical coefficient matrix A is constant and set to its most recent historical value. The results are obtained using Monte Carlo simulations and quantile regression, as described above, based on 10^5 simulated paths and a time horizon of approximately 17 years. This horizon is assumed to be sufficient to reach stationarity, as the mean, covariance matrix, and marginal VaR stabilize after about 8 years.

These results highlight that the tail risk associated with sectors such as manufacturing, utilities, and construction is primarily driven by fossil fuel prices; specifically, the production prices of the mining oil sector (211). Reducing the dependence of both the manufacturing and energy sectors on fossil fuels could therefore contribute to greater stability in production prices. Furthermore, the manufacturing production prices are also influenced by the tail risk of the

ΔCoVaR	11	211	212	22	23	31G
11	0.3118	0.0327	0.0109	0.0062	0.0241	0.0295
211	0.0099	0.6929	0.0109	0.0128	0.0153	0.0443
212	0.0241	0.0430	0.2047	0.0068	0.0273	0.0237
22	0.0158	0.1283	0.0117	0.1164	0.0131	0.0157
23	0.0451	0.0768	0.0328	0.0081	0.1935	0.0623
31G	0.0530	0.1879	0.0275	0.0070	0.0792	0.1615

Table 6: ΔCoVaR_{ij} matrix of the stationary relative log-price ζ^∞ , see Section 3.2, that is the VaR of sector i (row index) conditional to sector j (column index).

agricultural production prices, that are consistently influenced by weather conditions.

6.2 Analysis of network links

To study the impact of network links on production prices, we calculate the change in relative log-price risk measures when we selectively deactivate certain input-output links, that is, we set a row of the matrix A to zero, except for the diagonal element. This sensitivity analysis is an extension, in a dynamic and stochastic case, of the classical static extraction analysis of economic IO literature; see, for instance, Dietzenbacher and Lahr (2013). For this analysis, we always assume a constant matrix A . Table 7 presents the percentage changes in the stationary ΔVaR of log-prices, when the output link of a particular sector is removed. Figure 7 shows two examples for the utilities (22) and manufacturing (31G) sectors of the evolution over time of the percentage change of ΔVaR of relative log-prices, when the output link of a particular sector is removed. We recall, that the marginal VaRs of $\zeta(t)$ can be defined similarly to the stationary ones as in equation (6.1), and they can be calculated inverting the characteristic function at finite time t , using the characteristic exponent defined in equation (3.5) and the COS method of Fang and Oosterlee (2009). First and foremost, the matrix A significantly influences the stationary (marginal) distributions, whereas the short-term distributions are primarily determined by the Lévy process parameters. However, despite the differences in magnitude, the trends observed in the sensitivity analyses are similar over time. The sensitivity analysis reveals a strong impact of the oil and gas extraction sector (211) on price distribution of other sectors, such as manufacturing (31G) and utilities (22). In fact, eliminating the output link of the oil and gas sector (211) consistently reduces the stationary ΔVaR of the other sectors: -1.14% for utilities (22), -6.13% for construction (23) and -10.95% for manufacturing (31G). This analysis reveals that the tail risks of utilities (22), construction (23) and manufacturing (31G) sectors are largely driven by the tail risk of fossil fuels production prices (sector 211). Since fossil fuels are a major contributor to greenhouse gas emissions, this dependency highlights a significant exposure of these sectors to climate transition risks, as shifts in climate policy, technological innovation, or market preferences could amplify volatility and reshape risk profiles.

6.3 Analysis of resilience rates

In this section, as in Tsanakas and Millossovich (2016), we estimate the derivatives of risk measures of the relative log-prices $\zeta(t)$ with respect to the resilience rates K . In particular, Table 8 reports the derivative of the 99% marginal ΔVaR in the resilience rates directions, $\frac{\partial \Delta \text{VaR}_{0.99}[\zeta_j(t)]}{\partial K_m}$, where the index j, m belongs to sectors and t is 21 months (or 7 quarters).

	$\Delta VaR_{0.99}[\zeta_{11}]$	$\Delta VaR_{0.99}[\zeta_{211}]$	$\Delta VaR_{0.99}[\zeta_{212}]$	$\Delta VaR_{0.99}[\zeta_{22}]$	$\Delta VaR_{0.99}[\zeta_{23}]$	$\Delta VaR_{0.99}[\zeta_{31G}]$
11	-0.72	-0.00	-0.10	-0.00	-1.26	-1.57
211	-0.24	-0.52	-0.82	-1.14	-6.13	-10.95
212	-0.01	-0.00	-0.09	-0.01	-0.35	-0.08
22	-0.06	-0.11	-0.40	-0.04	-0.42	-0.35
23	-0.01	-0.00	-0.15	-0.00	-0.12	-0.05
31G	-1.44	-0.46	-1.79	-0.14	-22.58	-1.86

Table 7: Percentage change in the stationary 99% ΔVaR of relative log-prices, when the output links of a particular sector are removed.

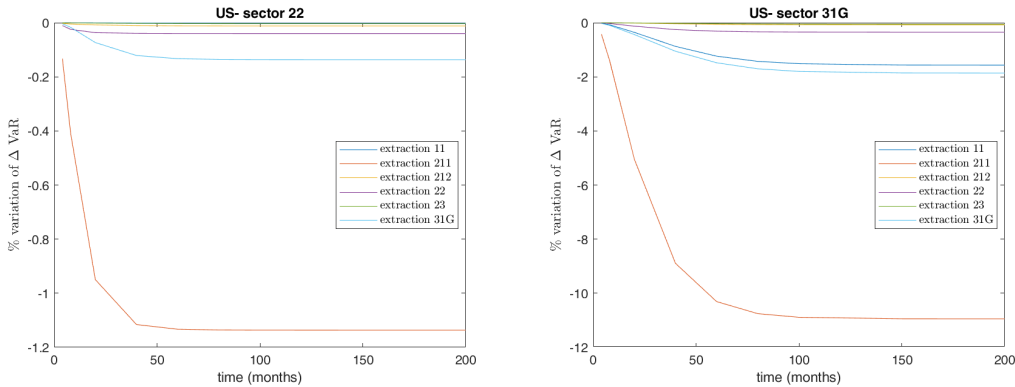


Figure 7: Evolution over time of the percentage change of ΔVaR of relative log-price for the utility (22) and manufacturing (31G) sectors, when the output links of a particular sector are removed.

The resilience (or production adjustment) rates determine both the stationary impact of supply shocks on prices and also the time required for a sector to return to its equilibrium state after experiencing an external direct or internal indirect shock. Allocating resources and efforts to enhance the resilience rate of a specific sector helps mitigate the impact of shocks on that sector's prices, in fact the derivative $\frac{\partial \Delta VaR_{0.99}[\zeta_j(t)]}{\partial K_m}$ are always negative if $j = m$. However, sensitivity analysis reveals that increasing the resilience rate of a specific sector may lead to unintended consequences. While it may benefit the sector itself, it can negatively impact other sectors by increasing the tail risk of price fluctuations, that is, the mixed derivatives (i.e., $\frac{\partial \Delta VaR_{0.99}[\zeta_j]}{\partial K_m}$ with $j \neq m$) can be positive. The cross sector results in Table 8 are in line with the $\Delta CoVaR$ results in Table 6; in fact, for example, the non negligible positive mixed derivatives are the ΔVaR of manufacturing (31G) with respect to production rates of the agricultural and oil and gas extraction sectors, K_{11} and K_{211} . The cross-sector impact is particularly relevant in the short period, while it mitigates in the long run, as shown in Figure 8 for the utilities (22) and manufacturing (31G) sectors as an example.

7 Formulation of multi-criteria optimal allocation problems and future research developments

Based on the sensitivity analysis presented in the preceding section, allocating resources to increase the resilience rate of a given sector helps mitigate the impact of shocks on that sector's prices. However, when resources are limited, determining how to allocate them efficiently across sectors is a nontrivial problem, particularly due to the network of input-output linkages that

$\partial \text{VaR}_{0.99}$	11	211	212	22	23	31G
K_{11}	-1.3685	0.0000	0.0002	0.0000	0.0008	0.0109
K_{211}	0.0006	-3.1620	0.0026	0.0131	0.0060	0.0856
K_{212}	0.0001	0.0001	-1.1492	0.0002	0.0008	0.0004
K_{22}	0.0001	0.0006	0.0003	-0.1896	0.0001	0.0002
K_{23}	0.0001	0.0000	0.0008	0.0001	-0.8967	0.0005
K_{31G}	0.0147	0.0208	0.0111	0.0012	0.1100	-0.5763

Table 8: Derivative of the VaR of the relative log-price $\zeta(t)$, see equation (3.1), with respect to each production rate K_j for $t = 21$ months (or 7 quarters).




Figure 8: Evolution over time of the derivative of ΔVaR of log-price for the utilities (22) and manufacturing (31G) sector in different resilience rates directions.

couples the dynamics of different sectors. Shocks and policy interventions in one sector can propagate through the network, affecting other sectors in indirect and potentially unintended ways.

A promising direction for future research is to formalize the resource allocation problem as a multi-objective optimal control problem, following the framework proposed by Ehrgott (2005). In this section, we present a preliminary formulation and illustrative examples to demonstrate this approach.

The resource allocation problem can then be expressed as the following vector optimization problem:

$$\min_{k \in \mathbb{R}_+^N} (\rho(w' \zeta^\infty; k), C(k)), \quad (7.1)$$

where $k = \text{diag}(K)$ is the vector of sectoral resilience rates, $\rho(\cdot; k)$ is a risk measure defined for any $k \in \mathbb{R}_+^N$, ζ^∞ denotes the (random) vector having the stationary distribution of relative log-prices, $w \in \mathbb{R}_+^N$ is a vector of weights such that $\sum_{j=1}^N w_j = 1$ and $C(k)$ is a cost function capturing the resources required to achieve a given resilience profile. From an economic perspective, the weighted sum $w' \zeta^\infty$ can be interpreted as a consumer price index basket. To analyze the solutions of this multicriteria optimization problem (MOP), standard scalarization techniques from the literature can be employed to reduce the MOP to a family of related single-objective problems. In particular, by applying the ϵ -constraint method, see (Ehrgott, 2005, Chapter 4), we obtain

the following two scalarized problems:

$$\begin{aligned} & \min_{k \in \mathbb{R}_+^N} \rho(w' \zeta^\infty; k) \\ & \text{subject to } C(k) \leq \bar{C}, \end{aligned} \quad (7.2)$$

and

$$\begin{aligned} & \min_{k \in \mathbb{R}_+^N} C(k) \\ & \text{subject to } \rho(w' \zeta^\infty; k) \leq \bar{\rho}, \end{aligned} \quad (7.3)$$

where \bar{C} is the total available budget and $\bar{\rho}$ is the target maximum risk value.

By analyzing the solutions of the two scalarized problems, we can study the Pareto efficiency of the solutions of the original MOP problem, as formalized in the two following propositions, see Ehrgott (2005, Chapter 4) for formal proofs.

Proposition 13. *If $k^* \in \mathbb{R}_+^N$ is the solution of at least one of the two scalarized problem (7.2) and (7.3), then k^* is a weakly Pareto efficient solution for the original MOP problem (7.1).*

Specifically, if k^* is a solution to (7.2), then there exists no $k \in \mathbb{R}_+^N$ such that $\rho(w' \zeta^\infty; k) < \rho(w' \zeta^\infty; k^*)$; however, there may exist a $k \neq k^*$ for which $C(k) < C(k^*)$. Conversely, if k^* is a solution to (7.3), then there exists no $k \in \mathbb{R}_+^N$ such that $C(k) < C(k^*)$, however, there may exist a $k \neq k^*$, such that $\rho(w' \zeta^\infty; k) < \rho(w' \zeta^\infty; k^*)$.

Proposition 14. *If there exist two values $(\bar{C}, \bar{\rho}) \in \mathbb{R}_+^2$, such that a vector $k^* \in \mathbb{R}_+^N$ is solution of both problems (7.2) and (7.3), then k^* is a strong Pareto efficient solution for the original MOP (7.1).*

If k^* is a strong Pareto efficient solution of (7.1), this means that for any $k \in \mathbb{R}_+^N$, $k \neq k^*$, $\rho(w' \zeta^\infty; k) > \rho(w' \zeta^\infty; k^*)$ or $C(k) > C(k^*)$.

The feasibility of problems (7.2) and (7.3), and consequently of the original problem (7.1), naturally depends on the properties assumed for the functions $\rho(w' \zeta^\infty; \cdot)$ and $C(\cdot)$, in particular their continuity, differentiability, and convexity. In the following, we present an illustrative example in which a unique strong Pareto-efficient solution exists in closed form.

An illustrative example of MOP for resource allocation

A simplified formulation arises when the risk measure $\rho(\cdot; k)$ is taken to be the expectation operator and the cost function $C(k)$ is linear. In this case, the MOP (7.1) problem becomes

$$\min_{k \in \mathbb{R}_+^N} \left(\mathbb{E}[w' \zeta^\infty], \sum_{j=1}^N c_j k_j \right), \quad (7.4)$$

$c_j > 0$ is the marginal cost of increasing the resilience rate $k_j = K_{jj}$ for sector j . Then, the two scalarized problems (7.2) and (7.3) reduces to:

$$\begin{aligned} & \min_{k \in \mathbb{R}_+^N} \mathbb{E}[w' \zeta^\infty] \\ & \text{subject to} \quad \sum_{j=1}^N c_j k_j \leq \bar{C}, \end{aligned} \tag{7.5}$$

and

$$\begin{aligned} & \min_{k \in \mathbb{R}_+^N} \sum_{j=1}^N c_j k_j \\ & \text{subject to} \quad \mathbb{E}[w' \zeta^\infty] \leq \bar{\rho}, \end{aligned} \tag{7.6}$$

Noticing that equation (3.11) implies

$$w' \mathbb{E}[\zeta^\infty] = w' \tilde{\zeta} = \sum_{j=1}^N \frac{w_j b_j}{k_j},$$

with $b = (I - A')^{-1} \alpha$, both the optimization problems (7.5) and (7.6) are convex and admit single closed-form solutions.

Remark 15. *We assume that the stationary means of relative log-prices are positive, and hence $b \in \mathbb{R}_+^N$. This is the case in our empirical analysis, as reported in Table 5. If a sector or industry has a negative value b_j , it can be excluded from the specific optimization problem, since in such cases there is no risk to be mitigated.*

Specifically, the optimal resilience rates are given by:

$$k_i^* = \frac{\bar{C}}{\sum_{j=1}^N \sqrt{w_j b_j / c_j}} \sqrt{w_i b_i / c_i}, \quad \text{for } i = 1, \dots, N, \tag{7.7}$$

for the constrained problem (7.5), and

$$k_i^* = \frac{\sum_{j=1}^N \sqrt{w_j b_j / c_j}}{\bar{\rho}} \sqrt{w_i b_i / c_i}, \quad \text{for } i = 1, \dots, N, \tag{7.8}$$

for the constrained problem (7.6).

Then, from Proposition 14, for any $\bar{C} \in \mathbb{R}_+ / \{0\}$ and $\bar{\rho} = \frac{(\sum_{j=1}^N \sqrt{w_j b_j / c_j})^2}{\bar{C}}$, k^* defined in (7.7) or equivalently in (7.8) is a strong efficient solution of the multi-objective problem (7.4). This result illustrates how an efficient allocation of resources can be achieved by balancing sector-specific stationary means b_j , weights w_j , and costs c_j .

Future research directions include analyzing the existence of efficient solutions for more general distributional risk measures and/or cost functions, as well as extending the formulation from the static stationary framework to a dynamic optimal control setting.

8 Conclusion

In this paper, we have developed a dynamic stochastic framework to study the propagation of supply shocks through industrial production networks, with a particular focus on climate-related and geopolitical disruptions. By introducing sector-specific resilience rates and deriving semi-analytical results for the long-term impact of persistent shocks on the stationary distribution of production prices, our model addresses key gaps left by traditional input-output approaches. Empirical analysis is conducted using US economic data, including input-output tables and historical series of industrial production prices, allowing us to quantify sector-specific adjustment dynamics. A sensitivity analysis based on risk measures of log-price distributions provides new tools for supply risk management. Our empirical findings highlight that the tail risk associated with sectors such as manufacturing, utilities, and construction is largely driven by the volatility of fossil fuel prices, particularly those of the mining oil sector. Reducing the dependence of manufacturing and energy sectors on fossil fuels could therefore enhance the stability of production prices.

Data availability statement

The data supporting the findings of this study are openly available in the Bureau of Economic Analysis (BEA) Interactive Data Application at <https://www.bea.gov/itable>, with reference number re3data ID: r3d100010187 .

References

- Acemoglu, D., Carvalho, V. M., Ozdaglar, A., and Tahbaz-Salehi, A. (2012). The network origins of aggregate fluctuations. *Econometrica*, 80(5):1977–2016.
- Adrian, T. and Brunnermeier, M. K. (2016). Covar. *American Economic Review*, 106(7):1705–41.
- Barrera, D., Crépey, S., Gobet, E., Nguyen, H. D., and Saadeddine, B. (2026). Statistical learning of value-at-risk and expected shortfall. *Mathematical Finance*, 36(1):156–179.
- BlochI, F., Theis, F., Vega-Redondo, F., and Fisher, E. (2011). Vertex centralities in input-output networks reveal the structure of modern economies. *Physical Review. E, Statistical, Nonlinear, and Soft Matter Physics*, 83:046127.
- Bohlin, L. and Widell, L. M. (2006). Estimation of commodity-by-commodity input-output matrices. *Economic Systems Research*, 18(2):205–215.
- Bureau of Economic Analysis (2025a). Annual input-output statistics. <https://www.bea.gov/itable/input-output>.
- Bureau of Economic Analysis (2025b). Chain-type price indexes for gross output by industry. <https://www.bea.gov/itable/gdp-by-industry>.

- Bureau of Economic Analysis (2025c). Real GDP. <https://www.bea.gov/itable/national-gdp-and-personal-income>.
- Cambou, M. and Filipović, D. (2017). Model uncertainty and scenario aggregation. Mathematical Finance, 27(2):534–567.
- Carvalho, V. M. and Tahbaz-Salehi, A. (2019). Production networks: A primer. Annual Review of Economics, 11(1):635–663.
- Casassus, J. and Collin-Dufresne, P. (2005). Stochastic convenience yield implied from commodity futures and interest rates. The Journal of Finance, 60(5):2283–2331.
- Chiang, I.-H. E., Huguen, W. K., and Sagi, J. S. (2015). Estimating oil risk factors using information from equity and derivatives markets. The Journal of Finance, 70(2):769–804.
- Choudhury, G. L. and Lucantoni, D. M. (1996). Numerical computation of the moments of a probability distribution from its transform. Operations Research, 44(2):368–381.
- Ciravegna, L., Ahlstrom, D., Michailova, S., Oh, C. H., and Gaur, A. (2023). Exogenous shocks and mnes: Learning from pandemics, conflicts, and other major disruptions. Journal of World Business, 58(6):101487.
- Cont, R. and Tankov, P. (2004). Financial Modelling with Jump Processes. Chapman and Hall/CRC Financial Mathematics Series, London, UK.
- Cortazar, G., Milla, C., and Severino, F. (2008). A multicommodity model of futures prices: Using futures prices of one commodity to estimate the stochastic process of another. Journal of Futures Markets, 28(6):537–560.
- Cortazar, G., Millard, C., Ortega, H., and Schwartz, E. S. (2019). Commodity price forecasts, futures prices, and pricing models. Management Science, 65(9):4141–4155.
- Dietzenbacher, E. and Lahr, M. (2013). Expanding extractions. Economic Systems Research, 25(3):341–360.
- Ehrgott, M. (2005). Multicriteria optimization., volume 491. Springer Science & Business Media.
- Fang, F. and Oosterlee, C. W. (2009). A novel pricing method for European options based on Fourier-cosine series expansions. SIAM Journal on Scientific Computing, 31(2):826–848.
- Ghosh, A. (1964). Experiments with input-output models. CUP Archive.
- Hamilton, J. (1994). Time Series Analysis. Princeton: Princeton University Press.
- Härdle, W., Hall, P., and Marron, J. S. (1988). How far are automatically chosen regression smoothing parameters from their optimum? Journal of the American Statistical Association, 83(401):86–95.
- Kella, O. and Whitt, W. (1999). Linear stochastic fluid networks. Journal of Applied Probability, 36(1):244–260.

- Kiefer, N. M. (1978). Discrete parameter variation: Efficient estimation of a switching regression model. Econometrica, 46(2):427–434.
- Koenker, R. (2017). Quantile regression: 40 years on. Annual Review of Economics, 9(Volume 9, 2017):155–176.
- Kostoska, O., Stojkoski, V., and Kocarev, L. (2020). On the structure of the world economy: An absorbing markov chain approach. Entropy, 22(4).
- Lam, C. (1998). Decomposition of time-ordered products and path-ordered exponentials. J. Math. Phys., 39:5543–5558.
- Le Guenedal, T. and Tankov, P. (2025). Corporate debt value under transition scenario uncertainty. Mathematical Finance, 35(1):40–73.
- Leontief, W. W. (1936). Quantitative input and output relations in the economic systems of the united states. The Review of Economic Statistics, pages 105–125.
- Lian, C. and Haimes, Y. Y. (2006). Managing the risk of terrorism to interdependent infrastructure systems through the dynamic inoperability input-output model. Systems Engineering, 9(3):241–258.
- Long Jr, J. B. and Plosser, C. I. (1983). Real business cycles. Journal of Political Economy, 91(1):39–69.
- Lu, K. W. (2022). Calibration for multivariate lévy-driven ornstein-uhlenbeck processes with applications to weak subordination. Statistical Inference for Stochastic Processes, 25(2):365–396.
- Mascaretti, A., Dell’Agostino, L., Arena, M., Flori, A., Menafoglio, A., and Vantini, S. (2022). Heterogeneity of technological structures between EU countries: An application of complex systems methods to input-output tables. Expert Systems with Applications, 206:117875.
- Nadaraya, E. A. (1964). On estimating regression. Theory of Probability & Its Applications, 9(1):141–142.
- Pang, R. K.-K. and Shrimali, G. (2024). Financial network valuation under climate transition risk. Available at SSRN 4905653.
- Protter, P. (1990). Stochastic Integration and Differential Equations. Springer-Verlag, Berlin, 2 edition.
- Ramezani, C. and Zeng, Y. (2007). Maximum likelihood estimation of the double exponential jump-diffusion process. Annals of Finance, 46(3):487–507.
- Ronald E. Miller, P. D. B. (2022). Input-Output Analysis: Foundations and Extensions. Cambridge University Press, 3 edition.
- Sato, K. (1999). Lévy Processes and Infinitely Divisible Distributions, volume 68 of Cambridge Studies in Advanced Mathematics. Cambridge University Press, Cambridge.

- Sato, K. and Yamazato, M. (1984). Operator-selfdecomposable distributions as limit distributions of processes of ornstein-uhlenbeck type. Stochastic Processes and their Applications, 17(1):73–100.
- Schwartz, E. and Smith, J. E. (2000). Short-term variations and long-term dynamics in commodity prices. Management Science, 46(7):893–911.
- Sodhi, M. S. and Tang, C. S. (2012). Strategic Approaches for Mitigating Supply Chain Risks, pages 95–108. Springer US, Boston, MA.
- Sopgoui, L. (2024). Modeling the impact of climate transition on real estate prices. arXiv preprint arXiv:2408.02339.
- Tsanakas, A. and Millossovich, P. (2016). Sensitivity analysis using risk measures. Risk Analysis, 36(1):30–48.
- Valdivieso, L., Schoutens, W., and Tuerlinckx, F. (2009). Maximum likelihood estimation in processes of ornstein-uhlenbeck type. Stat Inference Stoch Process, 12:1–19.
- Veraart, L. A. M. (2020). Distress and default contagion in financial networks. Mathematical Finance, 30(3):705–737.
- Wahba, G. (1975). Smoothing noisy data with spline functions. Numerische Mathematik, 24(5):383–393.
- Watson, G. S. (1964). Smooth regression analysis. Sankhyā: The Indian Journal of Statistics, Series A, pages 359–372.
- Weisz, H. and Duchin, F. (2006). Physical and monetary input–output analysis: what makes the difference? Ecological Economics, 57(3):534–541.
- Xu, W., Hong, L., He, L., Wang, S., and Chen, X. (2011). Supply-driven dynamic inoperability input-output price model for interdependent infrastructure systems. Journal of Infrastructure Systems, 17(4):151–162.