ECOOM–STORE
Vlamingenstraat 83/3550
B-3000 Leuven, België
https://www.ecoom.be

Beleidsrapport ECOOM–STORE–25–015

NeSESi-T, the Network Shock Estimation and Simulation Tool for a Firm-Level Production Network Approach

Yannick Bormans[◦,‡,1], Diederik Dujardin[◦,‡,2], & Jo Reynaerts[◦,‡,3]

[◦]*ECOOM–STORE*[*]

[‡]*Vlaams Instituut voor Economie & Samenleving (VIVES) , Faculteit Economie en Bedrijfswetenschappen , KU Leuven*

[1]*yannick.bormans@kuleuven.be*

[2]*diederik.dujardin@kuleuven.be*

[3]*jo.reynaerts@kuleuven.be*

22 december 2025

[*] De resultaten in dit rapport geven de mening van de auteur(s) weer en niet die van de Vlaamse Overheid: de Vlaamse Gemeenschap/het Vlaams Gewest is niet aansprakelijk voor het gebruik dat kan worden gemaakt van de in deze mededeling of bekendmaking opgenomen gegevens.

Abstract

Recent disturbances have exposed the fragility of domestic production networks to local shocks that can rapidly transmit across firms and sectors. This paper discusses the methodology behind NeSESi-T, an analytical tool that estimates and simulates the propagation and aggregation of shocks through production networks using detailed transaction-level data. Building on the framework of Carvalho *et al.* (2021), the model is calibrated through parameter estimation from a historical exogenous shock, after which it can be used for counterfactual simulations to assess potential network shocks and effective policy responses. It supports both *ex-ante* and *ex-post* analyses, guiding policymakers in the design of targeted interventions to mitigate network shocks. Although this version of NeSESi-T still has its limitations, future extensions will aim to expand its scope, evolving NeSESi-T into a comprehensive platform for analysing and managing dependency risks in domestic production networks, thereby offering a promising foundation for data-driven economic resilience strategies.

Contents

1. Introduction

The past few years have highlighted the vulnerability of global production networks. The COVID-19 pandemic, rising geopolitical tensions, and energy supply disruptions have shown how swiftly local shocks in one region can trigger widespread economic repercussions elsewhere. For governments and institutions, these events underscored the importance of strengthening strategic autonomy, i.e., the capacity to safeguard essential production capabilities and secure critical supply chains despite external disturbances. Policymakers therefore face an urgent need for effective tools to design shock-mitigation policies and network-resilience strategies that clarify which interventions can achieve desired outcomes under specific disruption scenarios.

A similar challenge arises within domestic economies. Whether geographically or sectorally concentrated, local shocks can propagate through supplier–customer relationships, transmitting the shock across firms and industries. This propagation threatens the stability of strategically important domestic supply chains. When a central firm — one highly interconnected with numerous suppliers and clients — experiences a severe disruption, the resulting cascade can undermine production throughout the entire sector, or even affect a large part of the wider economy. Such shocks may stem from factory closures, lockdowns, or regional disasters like earthquakes, blackouts, or nuclear incidents. Yet, most economic models fail to account for the complex interdependencies between firms. Existing frameworks, including those used by the European Commission, typically operate at the sectoral level, limiting their ability to generate detailed and targeted policy responses at the firm level. Policymakers are thus left without adequate tools to anticipate where and how shocks will spread.

This paper presents a first version of the Network Shock Estimation and Simulation Tool (NeSESi-T, pronounced as "necessity"), developed to largely address this gap. Following Carvalho *et al.* (2021), it enables policymakers and analysts to estimate how shocks propagate through production networks — from the micro to the macro level — and to identify vulnerabilities with precision. The model supports both preventive (*ex-ante*) and responsive (*ex-post*) analyses, helping to detect critical dependencies before crises occur and to simulate the impact of potential disruptions when urgent action is required. Although this paper restricts the discussion to a methodological elaboration of NeSESi-T, the framework should be capable of providing detailed network-based analyses, ideal for smart policy design, when it is combined with high-quality Belgian transaction-level business-to-business (B2B) data. An application of the NeSESi-T framework will be provided in a separate paper. While the current version of NeSESi-T focusses on short-term domestic shock propagation caused by negative shocks, future extensions will aim to expand its scope towards a more comprehensive resilience framework that is capable of capturing shock propagation over a longer time period in a production network with international linkages, incorporating both positive and negative shocks.

The remainder of this report is structured as follows: Section 2 explains the basic features and behaviour of production networks at the firm level according to the state-of-the-art literature. Section 3 elaborates the model and methodology of NeSESi-T, which heavily draws from Carvalho *et al.* (2021). In Section 4, the data needed to work with NeSESi-T are discussed, after which Section 5 illustrates how NeSESi-T works, and how it captures the production network features and behaviour explained in Section 2. Section 6 contains the discussion; Section 7 concludes.

2. The production network setting

A country's economy can be interpreted as a complex series of interlocking supply chains. This aggregate of supply chains forms a network of nodes (i.e., firms) and edges (i.e., transactions between firms). These edges have directions, as the transaction of a supplier to a client is not the same as a transaction from that client to the supplier. Furthermore, the edges are weighted by the transaction value: some transactions are more important than others, since the value of the transaction differs. Not all firms are equally important either, as the nodes in the production network are weighted as well. These node weights are usually expressed as the Domar weight, which equals the ratio of the firm's total sales over the country's GDP. These weighted nodes with directed, weighted edges between them form the domestic production network: it therefore contains all information on how the domestic economy is structured. It is the ideal setting to study how the microeconomic level - i.e., the separate nodes and edges - influences the macroeconomy - i.e., summary statistics of the entire network.

The Granularity Hypothesis of Gabaix (2011) explains how the behaviour of macroeconomic aggregates can be influenced by microeconomic events. So-called "incompressible grains" of economic activity - i.e., firms or disaggregated sectors - can result in macroeconomic fluctuations in the case of significant heterogeneity at the micro level, since heterogeneity reduces the extent to which a set of microeconomic shocks can cancel each other at the aggregate. Gabaix attributes this micro-level heterogeneity solely to significant differences in Domar weights, implying that idiosyncratic shocks can generate macroeconomic fluctuations even in the absence of specific network structure characteristics. However, the architecture of the production network itself also matters, since heterogeneity in sectors' positions within the network similarly influences aggregate outcomes.

Acemoglu *et al.* (2012) offer the insight that sufficient heterogeneity in different industries' roles as input-suppliers can significantly increase aggregate volatility in case of idiosyncratic shocks. The structure of the network therefore matters as well: microeconomic shocks impacting the aggregate output roughly symmetrically cancel out at the macro-level, while they do not if they impact the economy in an asymmetric way. In other words, if industries differ significantly in their roles as input-suppliers, shocks to industries that are considerably more important suppliers will propagate further in the network, affect consequently more firms in a more severe way, and hence do not cancel out with other idiosyncratic shocks. The same is true at the firm level.

At the heart of the matter lies the concept of a *focal* firm. A focal firm is an enterprise that because of its precise role, position and/or size in the production network exerts a disproportionally large influence on the other firms in the network. It may for example be a large firm at an early stage of the supply chain, holding the monopoly at its stage of the production process, and important for many firms as a crucial supplier for their production. The consequences of a negative shock on such a focal firm are very likely to be substantial, as many other firms may be severely affected, whilst this is much less the case for non-focal firms. The focality of a firm is linked to the concept of centrality: a firm i is more "central" in the network when it is a more important input-supplier and output-buyer to other firms. There exist many different centrality measures, such as the well-known Bonacich centrality (Bonacich, 1987). It is therefore not only a significant heterogeneity in Domar weights that can generate sizable aggregate fluctuations, a large disparity in centralities can have the same effect.

Production networks come in many forms and variants. Some networks will be strongly clustered with focal firms at the center of the clusters, playing an important role in the network as supplier and/or buyer of many other firms in the cluster. Other networks will consist of less clusters, being more symmetric with less firms having a more than proportional impact on other firms. Carvalho *et al.* (2019) use simplified examples very similar to the ones in Figure 2 to illustrate this. They explain that the star network, network 2a in the figure, exhibits maximal dispersion: one single firm (firm 1) is the only supplier of inputs to all other firms. A shock to this firm will propagate extensively through the network, and affect all other firms, leading to sizeable aggregate effects. At the other end of the network spectrum are networks 2b and 2c. Here, the economy is characterised by extremely symmetric input-supplier roles for all firms. Although microshocks will still propagate from firm to firm, the symmetric nature of this propagation will make shocks dissipate more, as they cancel each other out, with much smaller aggregate effects as a result for the same shock to firm 1 in the network.

Furthermore, Carvalho *et al.* (2019, Theorem 4) establishes how a pair of economies with identical Domar weights but a different level of interconnectivity behave differently. On one hand, the economy that has a network with many more connections between firms will be characterised by a higher average pairwise correlation of firms' output. In a highly connected economy, a shock hitting one firm will propagate through supply chain linkages to many other firms, although propagation is likely to be more attenuated thanks to the better diversification of firms, and the many substitution possibilities in the densely connected network. Given that more firms are (in)directly impacted, the output of firms will therefore show higher levels of co-movement, since their outputs move more in synchrony. On the other hand, the firms in the network with less linkages have a more volatile output, as they rely on

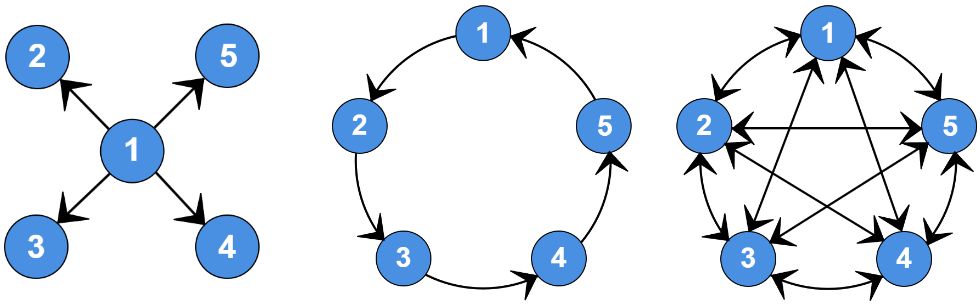

(a) (b) (c)

Figure 2: Three different types of stylised, fictive production networks, based on Figure 2 of Carvalho *et al.* (2019). Each blue node represent a firm. The arrows indicate a connection between firms: the arrow goes from supplier to buyer. Network (a) is a very centralised network with firm 1 as the focal firm. Network (b) is a circular network where each firm is the only supplier for the next firm in the chain. As the last firm is again the sole supplier to the first firm, the network forms a complete circle. Network (c) is a full, symmetric network, since all directed linkages are present in the network. This means that every firm is both supplier and client to any other firm.

less suppliers and clients. Firms in these lowly connected networks are less diversified with respect to supply chain shocks. They have less immediate substitution options when one of its suppliers/buyers is negatively impacted by a shock, impeding the mitigation of the negative shock impact for its own production. At the same time, it is more difficult for shocks to propagate to many other firms and sectors in the network, leading to a higher discrepancy between the fewer affected firms and the many non-affected. This difference is what distinguishes networks 2b and 2c in Figure 2. In the highly-connected economy of 2c, significant aggregate fluctuations will be accompanied by sizable comovement across a diverse range of firms, whilst aggregate fluctuations in the sparsely-connected economy of 2b are largely driven by high Domar weight firms being shocked and severely affecting its close suppliers and clients. In other words, there is a trade-off in economic network structures between interconnectivity and volatility.

To this point, the discussion has mainly focussed on the aggregate workings of production networks, and how macroeconomic indicators respond to a shock depending on the network structure. However, valuable insights can also be gained by analysing the impact at the firm level. Such a microeconomic perspective allows for an understanding of the underlying micro-level dynamics beyond the sheer macroeconomic impact. It highlights which firms experience the shock, in what magnitude, and the mechanisms through which these features arise.

As indicated before, a shock directly hitting a certain firm will propagate both downstream to the firm's clients, its clients' clients, etc., and upstream to the firm's suppliers, its suppliers' suppliers, and so on. The mechanism of these downstream and upstream propagation channels is explained in Carvalho *et al.* (2021) using a simplified

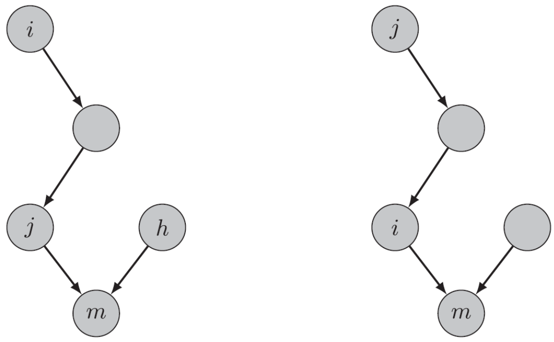

Figure 3: Figure X from Carvalho *et al.* (2021) explains the mechanism behind upstream and downstream shock propagation channels in the production network. Suppose that firm j is directly hit by a shock, than the shock will propagate upstream (left panel) or downstream (right panel) to firm i.

part of a network, visualised in Figure 3. Suppose that primary and intermediate inputs are complements, and that intermediate inputs among themselves are substitutes. In the left panel, firm i is purely upstream[1] to firm j. If firm j is hit by a negative shock, two effects will appear, resulting in a negative sales effect for firm i. Firstly, firm j reduces the expenditure share on intermediate inputs due to the negative shock. This drop in demand will propagate upstream through the supply chain, eventually lowering the sales of i. Secondly, the firm m downstream of j will substitute away from the production supply chain of firm j towards the supply chain of firm h, resulting in a second negative upstream effect affecting the firms upstream of firm j, including firm i. In the right panel, firm i is purely downstream[2] to firm j. If j is hit by a negative shock, the prices of good i will increase, making firm m that was using good i as an input to substitute away from i. This mechanism results accordingly in a negative downstream effect propagating through the network. Remark however that, as firm m can substitute away in inputs from the shocked supply chain, the impact on firm m will be more modest than for firms unable to substitute away. Substitution possibilities are consequently crucial for how severe the propagation effects are, and how far the propagation reaches. Moreover, it is much more difficult to quickly substitute away in the short term than in the medium to long term. The effect is therefore in general larger in the short term than in the medium to long term.

Interestingly, firms that are neither downstream nor upstream to a firm directly hit by a shock, may also experience shock propagation effects due to the reallocation effect. This effect is explained in Carvalho *et al.* (2019) using Figure 4. Imagine two separate parts of the production network, part (a) and part (b), where firms i and j supply firm k through a single supply chain. In (a), firms i and j indirectly supply

[1] Purely upstream means that there are only paths between firms i and j that connect the output side of firm i to the input side of firm j. There is no path, not even indirectly, from the output side of j to the input side of i.

[2] The notion of pure downstreamness is naturally derived from pure upstreamness.

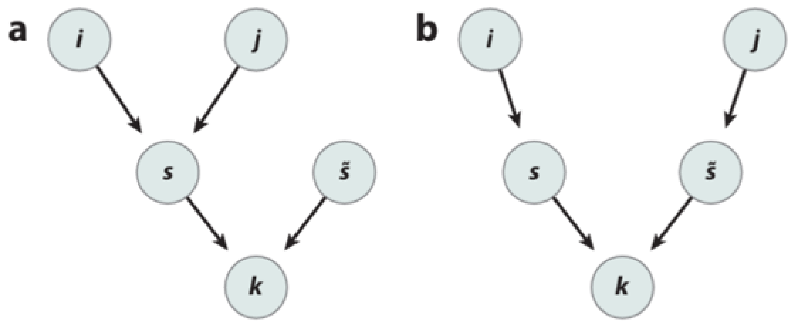

Figure 4: Figure 1 from Carvalho *et al.* (2019) illustrates how the reallocation effect translates a shock to firm j into an effect on firm i which is neither downstream nor upstream of j. Depending on the presence of overlap in the production chains, and the fact whether intermediate inputs are substitutes or complements, this effect will be negative or positive.

firm k through the same supplier s, while the production chains of i and j do not overlap each other in (b) until firm k. Assuming again that intermediate inputs are substitutes, and supposing firm j is hit by a shock, what will be the effect on firm i that is neither downstream nor upstream of j? In both network parts, the shock will propagate downstream from j to k. In order to attenuate the indirect network shock that k is experiencing, it will substitute away from the network branch of j. This will result in (a) in a drop in output of i, as firm k imports less from firm s and more from firm $\tilde{s}$. In (b) however, the output of i will increase, as k will still substitute away from $\tilde{s}$ to s. If intermediate inputs are however complements, firm k cannot substitute away from the supply chain of j, leading to the reverse conclusions: the output of i will increase in (a) and decrease in (b). This illustration of the reallocation effect indicates that the greater the overlap between the production chains of sectors i and j, the larger the adverse (favourable) impact on sector i when sector j is hit by a negative shock in an economy where intermediate inputs are gross substitutes (complements).

Unfortunately, the network is a complex combination of various overlapping supply chains, making it frequently severely opaque what the accumulated impact is of the many effects in the network. This is exactly where a tool as NeSESi-T should prove itself to be useful, by neatly combining the many channels through which network effects manifest themselves, and accumulating these effects into a clear measure in order to assess the impact of network shocks on firms and on the aggregate economy.

An excellent paper for more information on production networks, and their shock propagation behaviour is Carvalho *et al.* (2019). This paper is also a great starting point for people interested in the recent empirical evidence regarding network characteristics and shock propagation through networks, as it provides a clear overview of the literature.

3. Theoretical framework

In this section, we present the theoretical framework behind NeSESi-T, needed to perform the analysis leading to the estimation and simulation of production network shocks. This framework draws heavily from Carvalho *et al.* (2021); the same notation is used as in their paper.

3.1. Model

The model is based on a static economy that consists of n competitive firms denoted by $\{1, 2, \ldots, n\}$. Each firm produces a distinct product , which can be either destined for final consumption by households or used as an intermediate input by firms for production of other goods. Firms employ nested constant elasticity of substitution (CES) production technologies with constant returns, in which they transform labour, firm-specific capital, and intermediate goods into output. The output y_i of firm i is presented by the equation

$$y_i = \left[\chi(1-\mu)^{\frac{1}{\sigma}}\left((z_i k_i)^{\alpha} l_i^{1-\alpha}\right)^{\frac{\sigma-1}{\sigma}} + \mu^{\frac{1}{\sigma}} M_i^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}, \tag{1}$$

in which l_i is the firm's amount of employed labour, and k_i represents firm-specific capital. The capital-augmenting shock $z_i \leq 1$ to firm i has a normalised value of 1, meaning that a shock leads to the partial destruction of firm i's capital stock, which equals a post-shock endowment of $z_i k_i$. Also, μ captures the share of material inputs in the production technology of firms, α is the capital share in the primary factors bundle, and $\chi = (\alpha^{\alpha}(1-\alpha)^{1-\alpha})^{\frac{1}{\sigma}-1}$ is a normalisation constant. The elasticity of substitution between primary factors labour and capital, and the intermediate input bundle M_i is parametrised by σ. The intermediate input bundle itself is a CES aggregate of inputs from all firms:

$$M_i = \left[\sum_{j=1}^{n} a_{ij}^{\frac{1}{\xi}} x_{ij}^{\frac{\xi-1}{\xi}}\right]^{\frac{\xi}{\xi-1}}, \tag{2}$$

where x_{ij} is the quantity of good j used in the production process of i. The elasticity of substitution between different intermediate goods is ξ. The coefficient $a_{ij} \geq 0$ measures the importance of firm j as an intermediate input supplier for i. Ergo, the more important firm j is as a supplier for the production technology of firm i, the larger coefficient a_{ij} will be, whereas $a_{ij} = 0$ indicates that firm i does not have firm j as an input supplier. Throughout this paper, these coefficients are normalised by assuming that $\sum_{j=1}^{n} a_{ij} = 1$ for all i.

In addition to the production side, the economy also has a consumption side in the model. The unit mass of representative consumers, who populate the economy, inelastically supply to the economy L units of labour and $(k_1, \ldots, k_n)$ units of firm-specific capital. Their logarithmic preferences over the n goods are given by

$$u(c_1, \ldots, c_n) = \sum_{i=1}^{n} \beta_i \log\left(\frac{c_i}{\beta_i}\right), \tag{3}$$

in which c_i represents the consumed amount of good i, and constant coefficients $\beta_i \geq 0$ indicate the diverse goods' shares in the representative consumer's utility function. These coefficients are normalised as well, as $\sum_{i=1}^{n} \beta_i = 1$ is assumed.

The definition of the competitive equilibrium in the model's economy is as usual: the equilibrium of quantities and prices results in (i) maximal utility for the representative consumer, (ii) maximal profits for price-taking firms, and (iii) clearing of all markets.

Importantly, this model contains a couple of key concepts central to production network analysis. Firstly, there is the matrix $\boldsymbol{A} = [a_{ij}]$ summarising the interfirm input-output linkages; this matrix coincides with the steady-state input-output matrix. The Leontief inverse of the economy is then $\boldsymbol{L} = (\boldsymbol{I} - \mu \boldsymbol{A})^{-1}$, of which each (i,j) element denotes the importance of firm j as both a direct en indirect input supplier to firm i. I.e., a firm j does not supply directly nor indirectly via other firms to firm i if and only if $\ell_{ij} = 0$. Next, the Domar weight λ_i of a firm i is defined as its sales as a share of GDP: $\lambda_i = \frac{p_i y_i}{\text{GDP}}$, where p_i denotes the unit price of good i.

3.2. Methodology

With the model in hand, we can move towards a procedure for estimation and simulation of network shocks.

3.2.1. Estimation

Carvalho *et al.* (2021) explain how shocks propagate through the economy's network using the model elaborated in the previous subsection. In general, the model's equilibrium does not have a closed-form expression; hence why the authors make use of a first-order approximation around the point where elasticities σ and ξ are close to one. The net effect of a shock to firm j on the sales share of firm i is

$$\begin{aligned}\frac{d\log(\lambda_i)}{d\log(z_j)} = (\sigma - 1) \sum_{h=1}^{n} \alpha\mu(1-\mu)\frac{\lambda_h}{\lambda_i}\left[\left(\sum_{r=1}^{n} a_{hr}\ell_{rj}\right)\left(\sum_{s=1}^{n} a_{hs}\ell_{si}\right) - \ell_{hj}\sum_{r=1}^{n} a_{hr}\ell_{ri}\right] \\ + (\xi - 1)\sum_{h=1}^{n} \alpha\mu(1-\mu)\frac{\lambda_h}{\lambda_i}\left[\sum_{r=1}^{n} a_{hr}\ell_{ri}\ell_{rj} - \left(\sum_{r=1}^{n} a_{hr}\ell_{rj}\right)\left(\sum_{s=1}^{n} a_{hs}\ell_{si}\right)\right],\end{aligned} \tag{4}$$

with Leontief inverse $\boldsymbol{L} = [\ell_{ij}]$. As this closed-form expression for the first-order approximation is linear in both elasticities of substitution σ and ξ, it enables us to use linear regression in order to estimate these parameters from the data.

Equation (4) nicely illustrates the importance of the production network. Shocks to firm j induce its direct and indirect customers to change their expenditures composition on various inputs, leading to a change in demand expenditures on the good produced by firm i. The upper line at the right-hand side of Equation (4) captures the substitution between primary and intermediate inputs by any given firm h, whilst the lower line captures the substitution effect between the many intermediate inputs.

In Carvalho *et al.* (2021), Equation (4) eventually leads to an expression for the elasticities of substitution:

$$\Sigma = \alpha\mu(1-\mu)\boldsymbol{\Lambda}^{-1}\boldsymbol{L}'(\boldsymbol{A}'\boldsymbol{\Lambda}\boldsymbol{A} - \boldsymbol{A}'\boldsymbol{\Lambda})\boldsymbol{L}\ \Delta\log(z) \tag{5}$$

$$\Xi = \alpha\mu(1-\mu)\boldsymbol{\Lambda}^{-1}\boldsymbol{L}'(\mathrm{diag}(\boldsymbol{A}'\boldsymbol{\Lambda}\mathbf{1}) - \boldsymbol{A}'\boldsymbol{\Lambda}\boldsymbol{A})\boldsymbol{L}\ \Delta\log(z), \tag{6}$$

where Σ and Ξ are vectors of size n, and $\boldsymbol{\Lambda} = \mathrm{diag}(\lambda_1, \ldots, \lambda_n)$ is the diagonal matrix of the steady-state Domar weights. These equations result in the following specification that can be used in the estimation of σ and ξ:

$$\log(p_{it}y_{it}) = \gamma_i + \gamma_t + \beta_1(\Sigma_i \times \mathrm{year}_{t^*}) + \beta_2(\Xi_i \times \mathrm{year}_{t^*}) + \epsilon_{it}, \tag{7}$$

in which $\sigma = \beta_1 + 1$ and $\xi = \beta_2 + 1$. In Equation (7), t indicates the year, γ_i and γ_t are respectively firm and time fixed effects, and year_{t^*} is a time dummy for the post-shock year.

In order to estimate the elasticities of substitution Σ_i and Ξ_i, the authors explain how to construct the different terms in the equation. Firstly, note that the left-hand side of the equation can be derived from the matrix of Domar weights $\boldsymbol{\Lambda}$, which can be directly constructed from the data. Secondly, the parameters α and μ respectively coincide with the capital income shares and the aggregate intermediate input. Both values are to be calibrated using aggregate data, or from previous literature. Thirdly, the Leontief inverse $\mathbf{L} = (\boldsymbol{I} - \mu\boldsymbol{A})^{-1}$ is calculated from the matrix $\boldsymbol{A}$, which captures the production network and is directly derived from the data. Since it consists of a matrix inversion operation, the computation of the Leontief inverse can be computationally challenging, especially when the number of firms in the analysis is large. In Carvalho *et al.* (2021), the authors overcome this issue by approximating the inverse using the first 51 terms of its Neumann series. Finally, an estimate of the rate of capital destruction $\Delta\log(z)$ is needed. This estimate obviously depends on the specific network shock that is the object of examination, and needs to be based on the available information about the shock.

3.2.2. Aggregation

When sensible estimates of the elasticities of substitution σ and ξ have been found, the model can be employed to quantify the macroeconomic impact of the shock on the economy. As Carvalho *et al.* (2021) show, this aggregated impact on the economy's output can be approximated up to second order in shock size. This approximation is according to the model given by

$$\begin{aligned}\Delta\log(\text{GDP}) &= \alpha(1-\mu)\mathbf{1}'\left(\frac{\boldsymbol{\Lambda}+\boldsymbol{\Lambda}^*}{2}\right)\Delta\log(z) \\ &\quad + \frac{1}{2}\alpha^2\mu(1-\mu)(\sigma-1)\times\Delta\log(z)'\boldsymbol{\Lambda}(\boldsymbol{I}-\boldsymbol{A})\boldsymbol{L}\Delta\log(z),\end{aligned} \tag{8}$$

with diagonal matrices $\boldsymbol{\Lambda}$ and $\boldsymbol{\Lambda}^*$ of respectively firms' pre- and post-shock Domar weights. The authors explain in their paper that the advantage of a second-order approximation over a first-order is that it enables the capture of nonlinearities induced by nonunit elasticities, while simultaneously providing a tractable expression for the macroeconomic impact.

3.2.3. Simulation

The estimated model can eventually be employed to make general equilibrium forecasts for the aggregated impact of counterfactual shocks. For this counterfactual analysis, one needs three ingredients:

1. A clearly defined disaster-area indicates what firms are directly impacted by the shock. This disaster-area does not necessarily need to be a geographical area, but can also be an industrial sector, or any other classification that identifies impacted firms.

2. The capital destruction rate $\Delta\log(z)$ signifies the intensity of the counterfactual shock on the impacted firms. It informs us on the average capital destruction that impacted firms are confronted with.

3. The estimated model makes an estimate of the macroeconomic effect of the counterfactual shock possible through Equation (8).

Since the model is based on the estimates of the elasticities of substitution σ and ξ, these elasticities must be convincing in the counterfactual setting. Otherwise, the forecasts are likely to be nonsensical.

3.2.4. Multiplier

Instead of merely simulating different kinds of counterfactual shocks on the existing production network, one can also simulate shocks in a counterfactual version of the production network. If one removes all input-output linkages of directly hit firms - including the ones among the hit firms[3] - the shock can by construction no longer propagate through the production network. Other firms will consequently not be impacted, and the initial impact on directly hit firms is not amplified because of spill-overs from other hit firms: it is solely the initial impact that remains. Simulating the actual shock on such a counterfactual network, and comparing its effect on GDP to the GDP impact in the actual network, quantifies the contribution of shock propagation in the network to the aggregate impact. Dividing the actual GDP impact by the counterfactual GDP impact measures how much larger the accumulated shock effect is due to propagation. This value is called the network multiplier, as it indicates how hard the network features amplify the total shock impact.

[3]Here, we differ from Carvalho *et al.* (2021), as they keep the linkages between directly hit firms.

4. Data

The appropriate (type of) data is needed to work with the elaborated firm-level production network model NeSESi-T. The central dataset in the analysis is the B2B transaction-level dataset that informs us on the structure of the production network, which leads to matrix $\boldsymbol{A}$ in the analysis. If available for research, the national collection of VAT listings of enterprises can be used for this. VAT listings are particularly well suited, as they do not only tell what links there are between firms, but also how important these links are. If one does not know what the value of the transactions between firms are, then one needs to impute their value from elsewhere. In Carvalho *et al.* (2021), they use sectoral input-output tables to approximate the transaction values, which unfortunately results in a less detailed network structure.

In Belgium, the national VAT listings have been frequently used in the past for production network analysis. This dataset is of considerably high quality, as it does not only contain practically the entire universe of Belgian firms and their transactions, but it stretches over many years as well. Only Belgian enterprises completely exempt from VAT, or firms without a single domestic transaction above 250 euros are not included in the VAT listings. An extensive discussion on the Belgian VAT listings can be found in Dhyne *et al.* (2015). When considering a small, open economy in the analysis, firm-level import-export data is able to increase the completeness of the model, since the international linkages of the domestic firms examined in the analysis may play an important role in shock propagation and mitigation strategies.

Besides the B2B transactions dataset, other datasets may be useful in the analysis. Firm-level sales data are required to derive the Domar weights before and after the shock. Furthermore, an exhaustive list of all shock impacted firms is imperative for a good identification of struck enterprises, and is consequentially key for an accurate estimation. Information on the capital destruction rate of the impacted enterprises is essential as well for an analysis of high quality. Finally, data for the calibration of the capital income share and aggregate intermediate input are needed, if these cannot be calibrated using existing literature.

5. Illustrations

The elaborated NeSESi-T framework can be employed to estimate and simulate shocks in production networks at the firm level. In this section, the performance of NeSESi-T is illustrated using different types of network shocks in different kinds of fictitious, simplified networks. These illustrations are therefore not meant to simulate highly realistic production network shocks; they are however primarily intended to demonstrate how networks behave according to the NeSESi-T framework, reproducing the patterns extensively discussed in the literature. The illustrations are computed in reverse order in the way NeSESi-T is intended: firstly, the shocks are simulated, after which one can compute aggregate shock effects and network multipliers. Finally, the simulated scenarios are used to estimate the elasticities of substitution in order to check whether the used elasticity values can be recovered.

If all network parameters in the NeSESi-T framework are known, Equation (4) can be used to calculate the impact of a unit shock on the Domar weights of the firms in the network. The post-shock Domar weights λ_i^* for a shock $\boldsymbol{Z} = [z_i]$ can then easily be computed as follows:

$$\lambda_i^* = \lambda_i \cdot \exp\left(\sum_{j=1}^{n} \frac{d\, \log(\lambda_i)}{d\, \log(z_j)} \cdot \Delta\log(z_j) \right), \tag{9}$$

where λ_i is the pre-shock Domar weight. One can accordingly simulate what the effect of a fictional shock $\boldsymbol{Z}$ on the firms' Domar weight is, if one is able to realistically calibrate the NeSESi-T network parameters. These simulated post-shock Domar weights enable the computation of the counterfactual aggregate shock impact via Equation (8). Once the fictional scenarios are simulated, Equation (7) can be employed to estimate the elasticities of substitution used in the shock simulation.

The illustrations in this section will use different network types in order to simulate the effect of various network structures on shock propagation and aggregation. The networks in these simulations are based on the basic production networks presented in Figure 2. Put in a matrix format that captures the production network structure, the networks from Figure 2 are:

$$A_a = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \end{pmatrix} \; ; \quad A_b = \begin{pmatrix} 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{pmatrix} \; ; \quad A_c = \begin{pmatrix} 0.2 & 0.2 & 0.2 & 0.2 & 0.2 \\ 0.2 & 0.2 & 0.2 & 0.2 & 0.2 \\ 0.2 & 0.2 & 0.2 & 0.2 & 0.2 \\ 0.2 & 0.2 & 0.2 & 0.2 & 0.2 \\ 0.2 & 0.2 & 0.2 & 0.2 & 0.2 \end{pmatrix}.$$

Note that the values in matrix A_c have here been chosen such that all directed linkages are equally important. As long as all elements in matrix A_c would be different from zero, the production network would look the same as in the figure, although the weights of the directed links would of course be different.

The basic production networks from Figure 2 are expanded, such that they consist of exactly 100 firms. Network structure a is a very centralised network, henceforth called the central network, structured around a focal firm. All other firms in the network have the focal firm as their only supplier. The focal firm itself has no suppliers. Network b is a circular network, henceforth the circular network, in which each firm has the previous firm in matrix order as their only supplier. The first matrix firm has the last matrix firm as only supplier, in order to close the circle. Network c is a full network, as all linkages in the economy are present (albeit not necessarily equally important), hence this network is henceforth called the full network. These stylised networks are obviously not realistic production networks. However, they are simplified examples of production networks at the very extremes of possible network features, which is an ideal setting to study production network characteristics and shock propagation behaviour. Network a is a highly centralised network, network b is a sparsely-connected symmetric network, and network c is a densely-connected symmetric network. The matrices above are expanded in the natural way to size (100×100). These modest matrix sizes do not pose any computational problems for the matrix inversion needed to compute the Leontief inverse, while avoiding distortions caused by an excessively small network scale.

Unless specified otherwise, the transaction weights in the full economy in our simulations are randomly generated by drawing one sample of 100 components from a Dirichlet distribution with dimension 100. The central firm in network a gets the largest Domar weight of the entire economy in order to augment its focality even further. Otherwise, the heterogeneous Domar weights are randomly generated using a uniform distribution, after which they are all rescaled such that they sum up to two[4]. Although in reality these parameters are very unlikely to be uniformly distributed, this choice is sufficient to demonstrate the performance of NeSESi-T and the importance of heterogeneity in the network.

In order to simulate the worst case scenario in the different settings, the firm(s) that are the initial recipient of the shock are set to be either the focal firm in network a, or the most important firm(s) according to Domar weight in networks b and c. Random firms are shocked in b and c in the setting with homogeneous Domar weights. The z_i value of all other firms is equal to one: any shock damage at these firms solely

[4]The sum of all Domar weights should equal $\frac{1}{1-\mu}$

happens through shock propagation. In all simulations, unless stated otherwise, the parameters of the model are set according to the ones in Carvalho *et al.* (2021) in order to have realistic network behaviour: $\alpha = 0.36$, $\mu = 0.5$, $\sigma = 0.6$, and $\xi = 1.2$.

5.1. Simulating with NeSESi-T

NeSESi-T is capable of approximating both the aggregate shock impact as well as firm-specific ones. Before looking at the macroeconomic impact of different kind of shocks in different types of networks further below in this section, Appendix A contains a figure with a snippet of a fictive production network in which a network shock was induced. The figure indicates what the combined impact is of shock propagation and the reallocation effect for each single firm, which highlights NeSESi-T's ability to produce detailed firm-level outcomes, reflecting its strength in providing finely disaggregated policy recommendations.

As firm-level results are highly specific to a wide range of parameter settings - such as shock type, network type, network position, network distance from initial recipient, etc. - which would unnecessarily complicate the message of this section, the explanation below is focussed on the impact of shocks on the aggregate economy.

5.1.1. The relevance of network features

Following the literature discussed in Section 2, the aggregate economy in the centralised network is expected to be more susceptible to network shocks, than the circular and full economies, since shock propagation patterns will differ significantly between the different network types. Furthermore, heterogeneity in Domar weights is likely to lead to a worse post-shock macroeconomic outcome than a homogeneous Domar weight setting. The level of interconnectivity should impact shock propagation as well. Moreover, the easier it is to immediately substitute away from shock-affected supply chains, the lower the aggregated impact should be.

Network structure In the simulations with heterogeneous Domar weights - and heterogeneous transaction weights in case of the full network - the network slightly differs each round due to randomised value generation. Therefore, 100 Monte Carlo simulations are run to obtain an average effect in the network type. Each time the one firm with the highest Domar weight is shocked by destroying 75% of its capital[5],

[5]Remark that the equations in NeSESi-T are still approximations: the exact results are therefore most reliable for small shocks, and less for large ones. However, NeSESi-T is in first instance meant as a tool indicating which firms are worst impacted and in what order of magnitude, without requiring the exact value of the impact.

making the initial recipient of the shock being severely affected but not dropping out of the network. In Table 1, the results of the Monte Carlo simulations are presented for the various networks. As these simulated networks are highly stylised, and thus quite unnatural, these results are not to be interpreted as realistic representations of possible shocks in actual networks. They are rather meant to be interpreted relative to the other network types: what kind of network is worse affected, and in what order of magnitude?

The average simulation results in Table 1 are as expected. $\Delta\log(\text{GDP})$ indicates the approximate percentage drop in GDP, which is considerably larger in the central network, where the decrease is more than twice as large as in the two symmetric network types. In these symmetric networks, the shock aggregation is slightly larger in the circular network than in the full, highly-connected network. This corresponds to the discussion in Section 2, where it was stated that shocks attenuate more easily in networks with more connections. The simulated network multipliers corroborate these findings: while a shock to the focal firm that propagates through the centralised network leads to a macroeconomic impact of more than twice the size as if the shock were contained to the focal firm itself, shock propagation in the symmetric networks barely amplifies the network shock. In fact, in the fully diversified network the network multiplier appears to be on average even smaller than one, implying an aggregated shock effect that is even smaller when shocks can freely propagate through the network.

Table 1: Estimation of macroeconomic shock impact by network type

	Central network	Circular network	Full network
$\Delta\log(\text{GDP})$	-2.2478 (0.0009)	-1.0506 (0.0052)	-1.0370 (0.0064)
Network multiplier	2.1558 (0.0126)	1.0082 (0.0071)	0.9936 (0.0087)

This table reports the average results of the Monte Carlo simulations for various network types of dimension 100 with heterogeneous Domar weights and heterogeneous transaction weights. The initial recipient in the network is the firm with the highest Domar weight, which is the focal firm in the central network, or a random firm in the others. The shock destroys 75% of the initial recipient's capital. The first row of estimates indicates the approximate change in GDP in percentage, while the second row expresses how much larger the change in GDP is when the shock can propagate through the network relative to the setting where the shock cannot propagate to other firms.

Domar weight heterogeneity Table 2 presents the effect of Domar weight heterogeneity on shock aggregation. Looking at the worst case scenario, the table compares a heterogeneous Domar weight setting where the most important firm is directly hit with the homogeneous setting where a random firm is hit. In the central network, the focal firm is hit in both scenarios, only its weight differs. As expected, (more) heterogeneity in Domar weights increases the aggregated shock impact in the symmetric network types, since the drop in GDP is larger in a setting with heterogeneous weights compared to homogeneous weights. In the central network, the simulation results show however that the opposite is true. This unexpected finding for the central network is however not due to the fact that the network is a central one, but to the specific network structure in which there is only one focal firm and 99 client firms. Appendix B shows for example that expectations are confirmed when there are less client firms in the central network. These findings illustrate nicely how shock propagation and aggregation is influenced by a complex interplay of diverse network features. All of these features must be considered together to study realistic production network behaviour.

Connection density As distinct levels of interconnectivity lead to different shock propagation behaviour, simulations are also run for the full network type in which the percentage of linkages present in the network differs from 100%. However, in order to maintain the input-supplier and output-buyer symmetry among the firms, all enterprises in the network have the same amount of unique suppliers and unique clients. Using a homogeneous Domar weight distribution and homogeneous transaction weights, this simulation exercise purely looks at the effect of interconnectivity on shock aggregation, as the networks only differ in interconnectivity level from each other.

The simulation results are included in Table 3. For different levels of interconnectivity, the change in shock aggregation is at a more modest order of magnitude, albeit still statistically significant. The results indicate that the macroeconomic impact is the largest towards the two extreme levels of interconnectivity: it therefore appears that a trade-off is at play. This is consistent with the literature, as stated in Section 2. On the one hand, shock propagation in a more densely connected economy will affect more firms worse, as the network distance to the initial recipient is on average smaller. In other words, the shock has to travel less far in the network to impact more firms. On the other hand, a more highly connected network is a network where firms are more diversified, since there are more substitution options, dissipating propagating shocks faster.

Ease of substitution Shock propagation behaviour will naturally change if substitution in the production network gets easier. The easier it is for firms to substitute away from impacted inputs towards other primary or intermediate inputs that are viable substitutes, the less impact a certain shock will have on the network. This ease

Table 2: Estimation of macroeconomic shock impact by network type and Domar weight heterogeneity

	Central network		Circular network		Full network (100%)	
	Heterogeneous Domar weight	Homogeneous Domar weight	Heterogeneous Domar weight	Homogeneous Domar weight	Heterogeneous Domar weight	Homogeneous Domar weight
$\Delta\log(\text{GDP})$	−2.2478 (0.0009)	−3.2256	−1.0506 (0.0052)	−0.5324	−1.0370 (0.0064)	−0.5203
Network multiplier	2.1558 (0.0126)	6.1561	1.0082 (0.0071)	1.0161	0.9936 (0.0087)	0.9930

This table reports the average results of 100 Monte Carlo simulations for various network types of dimension 100 with heterogeneous Domar weights, as well as the results for the network types of the same dimension with homogeneous Domar weights. A homogeneous network setting also includes homogeneous transaction weights. As this homogeneous network setting does not include any randomisation, there are no error intervals. The initial shock recipient in the network is the firm with the highest Domar weight. In case the network has homogeneous Domar weights, the focal firm is shocked in the central network, and a random firm in the other networks. The shock destroys 75% of the initial recipient's capital. The first row of estimates indicates the approximate change in GDP in percentage, while the second row expresses how much larger the change in GDP is when the shock can propagate through the network relative to the setting where the shock cannot propagate to other firms.

Table 3: Estimation of macroeconomic shock impact by connection density

	Full network						
	(2%)	(5%)	(10%)	(25%)	(50%)	(75%)	(100%)
Δlog(GDP)	−0.52353	−0.51986	−0.51867	−0.51874	−0.51920	−0.52012	−0.52029
	(0.00024)	(0.00012)	(0.00013)	(0.00017)	(0.00020)	(0.00023)	
Network multiplier	0.99915	0.99215	0.98987	0.99001	0.99090	0.99265	0.99297
	(0.00047)	(0.00022)	(0.00025)	(0.00032)	(0.00037)	(0.00043)	

This table reports the average results of the Monte Carlo simulations for the full network type of dimension 100 with homogeneous Domar weights and homogeneous transaction weights. The interconnectivity varies from 2% to 100% of all possible network linkages. In order to have symmetric input-supplier roles in all scenarios, all firms had the same amount of suppliers, and same amount of clients. A random initial shock recipient in the network was chosen. The shock destroys 75% of the recipient's capital. The first row of estimates indicates the approximate change in GDP in percentage, while the second row expresses how much larger the change in GDP is when the shock can propagate through the network relative to the setting where the shock cannot propagate to other firms.

of substitution is measured by the elasticities of substitution: the higher their value, the easier it is to use alternative inputs.

In Table 4, the average result of 100 Monte Carlo simulations is shown for different elasticities of substitution in different network types. The results clearly demonstrate that higher elasticities reduce the shock aggregation. Changing the elasticity of substitution ξ does not adjust the macroeconomic impact of a shock in the central and circular network, since this elasticity indicates how easy it is to substitute between different intermediate goods, and in these kind of networks there is no possibility to switch to other intermediate inputs. The linkages are fixed, and firms in these networks have only one supplier: they have no opportunity to substitute towards other intermediate inputs if their supplier is impacted. Elasticity of substitution σ however measures the ease to substitute intermediate inputs with primary inputs such as capital or labour. As it is assumed that this is always possible in any kind of network, changing the value of σ does have an effect in all network types.

5.1.2. The relevance of shock features

Next to the network features are the characteristics of the shock itself hitting the initial recipient(s) of importance for shock propagation and aggregation. The shock can differ in two dimensions. Firstly, it can directly hit multiple firms. In that case, it does not only matter how many firms are impacted, but also what firms in particular. The latter issue was extensively investigated above. Secondly, the shock itself can change in intensity. A more intense shock damages ceteris paribus the initial recipient(s) worse, and results in stronger shock propagation waves through the network.

Shock prevalence If a larger percentage of firms in the network is directly hit by a shock, one expects the impact of the network shock to be larger, since the shock propagates from more starting points in the network. 100 Monte Carlo simulations are run in the completely full network type of dimension 100 with a shock destroying 75% of the capital of the top 1%, top 5%, and top 10% of firms according to Domar weight. The results in Table 5 demonstrate that the results indeed exhibit the expected kind of behaviour: the larger the shock prevalence, the higher the aggregated effect.

Shock intensity The stronger the shock hitting the initial recipients in the network, the worse one intuitively expects the macroeconomic impact to be, since a stronger shock signal is propagating through the network. Monte Carlo simulations are performed in a completely full network where a shock destroys 25%, 75%, 90%, and 99% of the capital of the top 10% of firms according to Domar weight. Again, the results in Table

Table 4: Estimation of macroeconomic shock impact by network type and ease of substitution

	Central network			Circular network			Full network (100%)		
	$\sigma = 0.6$ $\xi = 1.2$	$\sigma = 0.6$ $\xi = 1.7$	$\sigma = 0.9$ $\xi = 1.7$	$\sigma = 0.6$ $\xi = 1.2$	$\sigma = 0.6$ $\xi = 1.7$	$\sigma = 0.9$ $\xi = 1.7$	$\sigma = 0.6$ $\xi = 1.2$	$\sigma = 0.6$ $\xi = 1.7$	$\sigma = 0.9$ $\xi = 1.7$
Δlog(GDP)	−2.2478	−2.2478	−1.1926	−1.0506	−1.0483	−1.0082	−1.0370	−1.0281	−0.9975
	(0.0009)	(0.0010)	(0.0055)	(0.0052)	(0.0062)	(0.0061)	(0.0064)	(0.0058)	(0.0063)
Network multiplier	2.1558	2.1549	1.1759	1.0082	1.0076	1.0021	0.9936	0.9870	0.9876
	(0.0126)	(0.0127)	(0.0085)	(0.0071)	(0.0084)	(0.0085)	(0.0087)	(0.0077)	(0.0087)

This table reports the average results of 100 Monte Carlo simulations for various network types of dimension 100 with heterogeneous Domar weights. The ease of substitution differs between the simulations, as is indicated by the elasticities of substitution values in the table. The initial shock recipient in the network is the firm with the highest Domar weight. The shock destroys 75% of the initial recipient's capital. The first row of estimates indicates the approximate change in GDP in percentage, while the second row expresses how much larger the change in GDP is when the shock can propagate through the network relative to the setting where the shock cannot propagate to other firms.

Table 5: Estimation of macroeconomic shock impact by shock prevalence

	Percentage of initial recipients		
	1%	5%	10%
$\Delta\log(\text{GDP})$	-1.0370	-5.0680	-9.7790
	(0.0064)	(0.0269)	(0.0446)
Network multiplier	0.9936	0.9909	0.9875
	(0.0087)	(0.0074)	(0.0064)

This table reports the average results of 100 Monte Carlo simulations for a full network of dimension 100 with heterogeneous Domar weights and heterogeneous transaction weights. The initial shock recipients in the network are respectively either the top 1%, top 5%, or top 10% firms according to Domar weight. The shock destroys 75% of the initial recipient's capital. The first row of estimates indicates the approximate change in GDP in percentage, while the second row expresses how much larger the change in GDP is when the shock can propagate through the network relative to the setting where the shock cannot propagate to other firms.

6 display the expected behaviour, since the decrease in GDP is clearly larger for higher levels of capital destruction.

5.2. Estimating with NeSESi-T

Once a shock has been injected into the network, and its propagation has been simulated such that post-shock Domar weights are found, it is possible with NeSESi-T to estimate the elasticities of substitution that were used in the simulation. Table 7 presents the estimation results for a random simulation round in the full network type with 10% of all possible connections. These connections are randomly distributed in the network. It is very clear that the simulated elasticities of substitution are easily recovered. Estimating these elasticities for the network types in Table 1 leads to inaccurate results due to multicollinearity issues. In the central and circular network, the absence of substitution possibilities results in an erratic estimation of parameter ξ, whilst the estimation of σ is misleading in the completely full network, because of the absence of firms without substitution possibilities. This does however not pose any problems for the operationality of NeSESi-T, as these extremely stylised networks never appear as actual networks: real-life production networks are always complex combinations of these highly simplified networks.

Table 6: Estimation of macroeconomic shock impact by shock intensity

	Capital-augmented shock z_i			
	0.75	0.25	0.1	0.01
Δlog(GDP)	-1.9778	-9.7790	-16.6981	-35.2145
	(0.0105)	(0.0446)	(0.0883)	(0.1650)
Network multiplier	0.9973	0.9875	0.9800	0.9630
	(0.0075)	(0.0064)	(0.0073)	(0.0064)

This table reports the average results of the Monte Carlo simulations for a full network of dimension 100 with heterogeneous Domar weights and heterogeneous transaction weights. The initial shock recipients in the network are the top 10% firms according to Domar weight. The shock destroys initially respectively either 25%, 75%, 90%, or 99% of the recipients' capital. The first row of estimates indicates the approximate change in GDP in percentage, while the second row expresses how much larger the change in GDP is when the shock can propagate through the network relative to the setting where the shock cannot propagate to other firms.

Table 7: Estimation of elasticities of substitution

	Full network (10%)
σ	0.6
	(4.57×10^{-14})
ξ	1.2
	(7.76×10^{-15})
N	200
R^2	1.0

This table reports the estimation results of the elasticities of substitution used in a shock simulation. The network in the simulation has dimension 100, and has heterogeneous Domar weights and heterogeneous transaction weights. 10% of the possible linkages in the network are present. The initial recipient in the network is the firm with the highest Domar weight. The shock destroys 75% of the initial recipient's capital. Since there is no noice introduced in the simulations, the estimation errors are practically zero. The originally used elasticity values in the simulations were $\sigma = 0.6$ and $\xi = 1.2$.

6. Discussion

In this paper, a network shock estimation and simulation tool called NeSESi-T is derived from the model used in Carvalho *et al.* (2021). This simulation tool is capable of capturing network shock effects at the firm level, which facilitates detailed firm-level policy research. By quantifying both firm-specific losses as well as the aggregated economic loss, NeSESi-T enables policymakers to weigh micro-level vulnerabilities against macroeconomic stability, encouraging more balanced and targeted policy design.

The tool can be applied for both ex-post and ex-ante analyses. To conduct credible ex-ante simulations however, the framework must first be calibrated ex-post to ensure the accurate capture of network behavior. Calibration is central to the model's reliability, as the parameters embed how firms respond to shocks within the production network. These parameters can be estimated empirically, using data from exogenous shock events, or drawn from results in the existing literature. In any case, it is crucial to understand that the elasticity of substitution parameters in NeSESi-T are economy-wide elasticities. Estimates may however vary depending on shock type and most severely affected sector, since substitution behaviour differs from sector to sector. It is therefore key to know how realistic the obtained estimates are for the shock scenario under ex-ante investigation, in order to get convincing results.

Once credible elasticity estimates are obtained, they can inform counterfactual simulations. Performing simulation exercises with a plausible parameter range can help assess sensitivity and robustness of simulation results. Applying the framework to a series of historical shocks could further illuminate how elasticities evolve over time and across industries, enhancing the credibility and interpretability of simulation outcomes. These outcomes allow researchers and policymakers to quantify the potential effects of plausible network disruptions, identify vulnerable firms or sectors, and assess the likely effectiveness of various mitigation policies at a highly disaggregated level. In cases of unanticipated shocks, NeSESi-T can serve as a diagnostic instrument to approximate the expected consequences and explore effective intervention strategies.

Simulations have shown that NeSESi-T captures key mechanisms of shock propagation across production networks. Yet, several important limitations must be acknowledged. First, the model is static: the network structure is fixed and does not adjust endogenously to shocks. In the immediate time period after the network shock, a firm can only substitute away from the impacted supply chain via its current partners in the network. The tool therefore captures only short-term effects — roughly within a one-year horizon — before firms have time to adapt by forming new supply and demand linkages. New transaction partners increase the substitution options, leading to a considerably mitigated shock impact. In other words, substitution possibilities increase over time as firms reconfigure partnerships, mitigating the long-run impact of network shocks. Elasticities in the short term are consequently lower than in the medium to long horizon after a shock. Accordingly, the elasticities of substitution estimated by NeSESi-T should be interpreted as short-run elasticities, and both firm-level and aggregate impacts are expected to lessen in the medium and long term.

Second, the current model architecture incorporates only negative shocks. It cannot yet simulate positive disturbances such as subsidies or positive demand shocks, unless it is combined with a negative shock such that the net shock is still negative or zero[6]. The tool is also based on a model that uses first- and second-order approximations. Estimations and simulations therefore become less accurate when the elasticities deviate further from unity, and when the shock intensity increases. Combined with the economy-wide nature of the elasticities in the model, the ideal research setting of the model is thus a case where a large yet not devastating shock affects a very diverse range of sectors and firms.

Lastly, the model's accuracy depends critically on the comprehensiveness of input data. If the production network dataset excludes international linkages, the model overlooks potential cross-border propagation channels. For open economies, this restriction is likely to limit the accuracy and policy takeaway of results.

[6]For example, the setting in which a firm i is negatively impacted by a certain shock and is granted a shock-counteracting subsidy, is a combination of a negative and a positive shock. This setting can be studied using the current version of NeSESi-T as long as the net shock is still negative or zero, such that $z_i \leq 1$.

7. Conclusion

The first version of a production network shock estimation and simulation tool called NeSESi-T is developed in this paper, heavily based on the model of Carvalho *et al.* (2021). The overarching goal of NeSESi-T is to support ex-ante evaluation of production network shocks and government intervention policies, providing a data-driven foundation for firm-specific resilience strategies. Simulations show that the current version already succeeds in incorporating key mechanisms of shock propagation in production networks.

Future extensions will aim to expand both the theoretical and empirical scope. The first step is to demonstrate the present model's capabilities by providing illuminating applications. Subsequent developments will focus on accommodating positive shocks, incorporating international linkages, and ultimately introducing dynamic network formation. A dynamic extension would enrich understanding of how production networks reorganize and recover after disruptions, providing valuable insights into long-term resilience and adaptability. Improvements in the availability and granularity of firm-level international input–output data will further strengthen the model's analytical power. Extra applications on these future versions of NeSESi-T will provide further insights in network behaviour, and illustrate the tool's performance.

In the long run, NeSESi-T aspires to evolve into an integrated platform for production network analysis, capable of simulating both adverse and beneficial shocks, and informing evidence-based economic policy. By systematically applying and refining this framework, future research can advance understanding of how production networks propagate, absorb, and adapt to shocks, leading policy makers in the design of mitigation policies and resilience strategies. This may eventually establish the framework as a true necessity for effective and informed policy making.

References

Acemoglu, D., Carvalho, V.M., Ozdaglar, A. and Tahbaz-Salehi, A. (2012), "The Network Origins of Aggregate Fluctuations," *Econometrica*, 80(5), 1977–2016. [4]

Bonacich, P. (1987), "Power and Centrality: A Family of Measures," *The American Journal of Sociology*, 92(5), 1170–1182. [5]

Carvalho, V.M., Nirei, M., Saito, Y.U. and Tahbaz-Salehi, A. (2021), "Supply Chain Disruptions: Evidence from the Great East Japan Earthquake," *The Quarterly Journal of Economics*, 136(2), 1255–1321. [i, 2, 3, 6, 7, 9, 10, 11, 12, 13, 14, 17, 26, 28]

Carvalho, V.M., Tahbaz-Salehi, A., Aghion, P. and Rey, H. (2019), "Production Networks: A Primer," *Annual Review of Economics*, 11(1), 635–663. [5, 6, 7, 8]

Dhyne, E., Magerman, G. and Rubínová, S. (2015), "The Belgian Production Network 2002-2012," . [14]

Gabaix, X. (2011), "The Granular Origins of Aggregate Fluctuations," *Econometrica*, 79(3), 733–772. [4]

A. Illustration of shock propagation through the production network

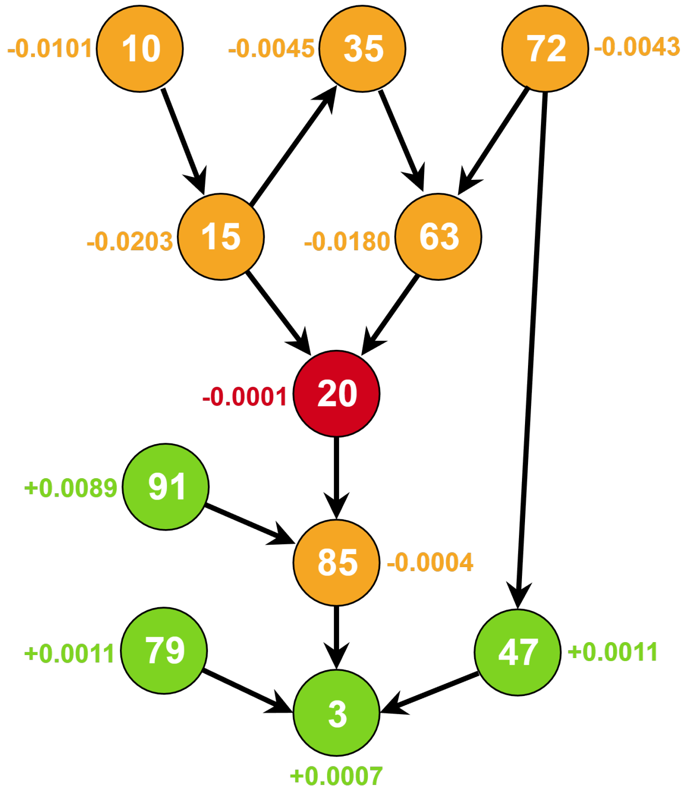

Figure 5: This figure presents a snippet of a production network generated in a random simulation round. The network has dimension 100, homogeneous Domar weights and transaction weights, and 2% of all possible linkages have materialised. These linkages are randomly distributed in the network. Various firms in the direct neigbourhood of Firm 20 are shown. In order to keep the figure readable, only linkages among the visualised firms are shown. The numbers next to the firms report the net impact on the firms' Domar weight caused by a unit shock directly hitting Firm 20.

All simulations in Section 5 compute first the firm-specific impact of a shock propagation through the network, before aggregating the effect to the macroeconomic level. Figure 5 presents a part of a production network generated in a random simulation round. The firms included in the figure are located in the network close to the firm that is the initial recipient of a shock. The values next to the firm nodes in the figure come from the matrix that is calculated in Equation 4: here, they indicate the net effect on firms' Domar weight of a unit shock that directly hits Firm 20. They show whether a specific firm is impacted, and how severely. These net values are the result of the complex interplay in the network between shock propagation and reallocation effects described in Section 2.

Figure 5 shows that the shock impact gets attenuated the further away a firm is from the initial recipient. The specific network structure gives some firms more opportunities to substitute away from impacted supply chains than others[7]. This reallocation response may result in a worse impact or a less bad one for a firm, depending on its own reaction and that of the other firms in the network[8]. Suppose for example that a firm is negatively impacted due to shock propagation, even though its reallocation response mitigates the original impact. If other firms in the network reallocate more towards that firm's supply chain than away, and that reallocation effect of the others is larger than the mentioned negative impact on the firm, the net effect for that specific firm can eventually even be positive. Hence, some firms may increase their Domar weight thanks to a certain network shock. Nevertheless, the overall impact on the economy will still be negative. Remark of course that the simulation here concerns only immediate substitution possibilities towards firms in the network that are already connected. In the medium to long term, firms can form new linkages with other firms that were previously unconnected, resulting in different shock impact patterns.

[7]Note however that linkages with firms outside the figure are not visualised.
[8]This explains why the net shock effect on Firm 20, which is the initial recipient, is modest.

B. Domar weight heterogeneity in a small central production network

Shock propagation and aggregation behaviour is very network specific. In Table 2 Domar weight heterogeneity leads counterintuitively to a better macroeconmic impact. However, running the same simulations in the same network type of dimension 10 - i.e., one focal firm being the only supplier of 9 client firms - gives the expected results, as Table 8 shows. However, the network multiplier is for both network dimensions larger in a homogeneous setting. This can be explained by the redistribution of Domar weights. Going from heterogeneous weights to homogeneous, the "excess" weight of one focal firm is spread over many client firms. The client firms therefore gain importance in the network, which increases the effect of shock propagation in the central network. The presence of a shock propagating in the homogeneous setting will thus affect more important firms, resulting in a worse aggregated impact, relative to the heterogeneous setting. Moreover, the initial recipient of the shock is in the homogeneous setting less of importance, lowering the shock impact on that firm. As both the impact on the initial recipient is less bad, and the impact on indirectly hit firms is larger, the difference between the scenario where a shock can propagate through the network and where it cannot is larger, giving a higher network multiplier value.

Table 8: Estimation of macroeconomic shock impact by Domar weight heterogeneity in a small central network

	Central network	
	Heterogeneous Domar weight	Homogeneous Domar weight
Δlog(GDP)	−9.5518 (0.1355)	−5.8680
Network multiplier	1.0610 (0.0221)	1.1199

This table reports the average results of 100 Monte Carlo simulations for a central network type of dimension 10 with heterogeneous Domar weights, as well as the results for the central network type of the same dimension with homogeneous Domar weights. A homogeneous network setting also includes homogeneous transaction weights. As this homogeneous network setting does not include any randomisation, there are no error intervals. The initial shock recipient in the network is the focal firm located central in the network. The shock destroys 75% of the initial recipient's capital. The first row of estimates indicates the approximate change in GDP in percentage, while the second row expresses how much larger the change in GDP is when the shock can propagate through the network relative to the setting where the shock cannot propagate to other firms.