

Endogenous Fragility in Opaque Supply Chains

Andrea Titton*

March 26, 2024

Abstract

This paper investigates the role of supply chain unobservability in generating endogenously fragile production networks. In a simple production game, in which firms with imperfect information need to multisource to hedge against suppliers' risk, firms underdiversify vis-à-vis the social optimum. The unobservability of suppliers' relations is the driver behind this. In production networks where upstream risk is highly correlated and supplier relationships are not observable, the marginal risk reduction of adding an additional supplier is low, because this additional supplier's risk is likely to be correlated to that of existing suppliers. This channel reduces firm incentives to diversify, which gives rise to inefficiently fragile production networks.

By solving the social planner problem, I show that, if the risk reduction experienced downstream resulting from upstream diversification were to be internalised by upstream firms, endogenous production networks would be resilient to most levels of risk. Furthermore, I show that the opaqueness of the supply chain yields less fragile but more inefficient production networks. Despite its stylised form, the model identifies the trade-off firms face when diversifying risk and isolates the mechanism that aggregates these decisions into a production network. Furthermore, it maps the conditions of the trade-off, such as expected profits of the firm or the pairing costs, to the properties of the production network.

*Faculty of Economics and Business, University of Amsterdam (a.titton@uva.nl). I thank my supervisors Florian Wagener and Cees Diks for the patient guidance on this paper. I also thank the CeNDEF group and the Quantitative Economics section at the University of Amsterdam for helpful comments throughout the many seminars. Finally, I thank attendees of the EEA-ESESM conference, Barcelona, 2023, and the Dutch Network Science Society Symposium, Leiden, 2022, for the constructive comments.

In August 2020, hurricane Laura hit one of the world’s largest petrochemical districts, in the U.S. states of Louisiana and Texas. As polymer producers in the area were forced to halt production, up to 15% of the country’s polypropene (PE) and polypropene (PP) producers were unable to source polymer inputs, which in turn caused shortages across the economy (Vakil, 2021). Such a widespread disruption raised awareness on the role suppliers’ correlation has in destabilising production networks and its importance in firms’ sourcing decisions. Yet, the structure of the supply chain is often opaque: firms do not observe sourcing relations beyond their immediate suppliers (Williams et al., 2013). In face of this opacity, how do producers make sourcing decisions? And, should we expect these sourcing decisions to yield resilient production networks?

In this paper, I study the role of supply chain opacity in determining firms’ sourcing decisions and, in turn, the consequences on the resilience of the production network. A widespread approach to mitigate risk is to diversify it by multisourcing. This practice consists of procuring the same inputs from multiple suppliers, sometimes redundantly (Zhao and Freeman, 2019). Yet, when deciding how many suppliers from which to source, a firm faces decreasing marginal benefits in risk reduction, because each additional supplier’s failure to deliver is increasingly likely to be correlated with that of the firm’s current suppliers. In the presence of marginal costs of sourcing, for example contractual costs or higher prices, the uncertainty behind the correlation of a firm’s potential suppliers might induce it to diversify risk less than socially optimal. The wedge between endogenous firm decisions and social optimality arises because downstream firms would be willing to compensate their suppliers for increased diversification of inputs. This underdiversification can generate aggregate fragility in production networks. To understand the relationship between the opacity of the supply chain, firms’ diversification decisions, and production network fragility, I study the properties of a stylised production game. In the equilibrium of the game, unobserved correlation among suppliers generates fragility via two channels. First, it directly introduces endogenous correlation in downstream firms’ risk, which amplifies through the production network. This increases the probability of cascading failures, in which the entire production network is unable to produce. Second, it indirectly affects firms’ decisions by reducing the expected marginal gain from adding a source of input goods. The latter channel leads to firms diversifying increasingly less, such that small shocks in the production of basal goods can generate cascading failures downstream.

The role that production networks play in determining economic outcomes has been long recognised. As far back as Leontief (1936), economists have studied how networks in production can act as aggregators of firm level activity. Following a foundational paper by Hulten (1978), which showed that the first order impact of a productivity shock to an industry is independent of the production network structure, macroeconomics has since de-emphasised this role (Baqae and Farhi, 2019, p. 2). However, more recently, Baqae and Farhi (2019) illustrated how the structure of the production network can aggregate micro shocks via second order effects.⁽¹⁾ Furthermore, the degree of competition in an

⁽¹⁾These results build on a vast literature and recent literature (e.g.)

industry also interacts with the production network to aggregate shocks, which can lead to cascading failures (Baqaei, 2018). Once established that production networks play a central role in aggregating shocks, two natural questions arise. First, which networks can we expect to observe, given that firms endogenously and strategically choose suppliers? Second, are these endogenous network formations responsible for the growth or fragility that large economies display? These questions fuelled a number of recent papers studying endogenous production network formation. Focusing on growth, Acemoglu and Azar (2020) show that endogenous production networks can be a channel through which firms' increased productivity lowers costs throughout the supply chain and allows for sustained economic growth. In parallel, a vast literature dealt with studying the role of endogenous production networks and firm incentives in determining fragile or resilient economies. Erol and Vohra (2014) showed that in networks with strategic link formation, systemic endogenous fragility arises if the shocks experienced by firms are correlated. Later work, by Amelkin and Vohra (2020), shows that uncertainty in the time of production is crucial in determining whether production networks in equilibrium are sparse, hence fragile. Finally, Elliott et al. (2022) illustrate how complexity in the production process can also be a key driver of endogenous fragility in production networks. ⁽²⁾

A less understood link is that between the opacity of the supply chain, how firms deal with it, and which consequences this has on the economy. Kopytov et al. (2021) studied the effect of uncertainty in endogenous production network formation on firms' productivity and business cycles. They find that higher uncertainty can lead to lower economic growth. In contrast, this paper focuses on the role of uncertainty in generating endogenous fragility to cascading failures using a more stylised production network model, akin to that studied by Elliott et al. (2022). In line with the existing literature, in the model small idiosyncratic shocks can be massively amplified. The degree of amplification depends on the equilibrium behaviour of firms. This phenomenon holds true in vertical economies producing simple goods. The novel theoretical contribution of this paper is to extend the analysis of production network formation to an opaque environment in which firms aim to minimise risk while accounting for correlation between suppliers. To do so, I develop a tractable analytical framework that describes the propagation of idiosyncratic shocks through the supply chain when firms take sourcing decisions endogenously in an imperfect information environment. The model describes the evolution of risk through the supply chain as a dynamical system over the depth of the production network. The social planner solution shows that endogenous fragility can impose large welfare losses. Importantly, these losses might be discontinuous: an arbitrary small increase in the correlation of risk among basal firms, can generate large welfare losses. Finally, I study a benchmark case without opacity, in which firms have full information. In this case, despite each individual firm being able to achieve smaller disruption risk, the production network is maximally fragile and there is a high probability of large disruptions.

The remainder of the paper is structured as follows. Section 1 discusses the assumptions

⁽²⁾The literature on production networks is vast and it is unfortunately impossible to give a fair overview in this introduction. For a more comprehensive review of the literature I refer the reader to Carvalho and Tahbaz-Salehi (2019) and Amelkin and Vohra (2020)

on the supply chain disruptions, the problem of the firm, and establishes the results that allow the firm to make sourcing decisions. Section 2 derives the law of propagation of the disruption events through the production network. Section 3 establishes the firm's optimal sourcing strategy and how this endogenously determines the fragility of the production network. These results are then compared, in Section 4, to the social planner solution to determine the welfare losses induced by the firm's endogenous decisions. Finally, in Section 5, the role of opacity is isolated by solving the model under perfect information.

1. Model

1.1. Production Technology and the Firm Objective

Consider a vertical economy producing $K + 1$ goods, indexed by $k \in \{0, 1, \dots, K\}$. Each firm produces a single good and each good is produced by m_k firms. Production of the *basal good* $k = 0$ does not require any input, yet, it is at risk of random exogenous disruptions in the production process. A *disrupted* basal firm is unable to deliver its good as input to downstream producers. Each downstream good $k > 0$ requires only good $k - 1$ as input. If a firm producing good k is unable to source its input good $k - 1$, the firm is itself *disrupted* and hence unable to deliver downstream. In other words, the i -th firm producing good k , indexed by (k, i) , is able to produce if at least one of its suppliers is able to deliver, namely, not all of its suppliers are disrupted. To avoid being disrupted, the firm chooses which firms to source from, among the producers of its input good. Letting \mathcal{D}_k be the random set of disrupted firms in layer k and $\mathcal{S}_{k,i}$ the set of suppliers of firm (k, i) , we can say that $(k, i) \in \mathcal{D}_k$ if and only if all of its suppliers $(k - 1, j) \in \mathcal{S}_{k,i}$ are in \mathcal{D}_{k-1} . I refer to the set of the firm's suppliers $\mathcal{S}_{k,i}$ as its *sourcing strategy*. The disruption events are random and the probability that a firm is disrupted can be written as

$$P_{k,i} := \mathbb{P}((k, i) \in \mathcal{D}_k) = \mathbb{P}(\mathcal{S}_{k,i} \subset \mathcal{D}_{k-1}). \quad (1)$$

Figure 1 illustrates this mechanism.




Figure 1: The supply chain is depicted in the left panel. The left firm is sourcing its input good from all three suppliers, \mathcal{S}_1 , while the right firm only from the latter two, \mathcal{S}_2 . As a disruption occurs, some upstream firms are unable to supply the input good (white box). Unlike the left firm, the right firm is unable to source its inputs and is hence disrupted.

If a firm is not disrupted, it obtains a profit π . Implementing a given sourcing strategy

costs the firm $C(|\mathcal{S}_{k,i}|)$. The cost C is assumed to be increasing in the number $|\mathcal{S}_{k,i}|$ of suppliers. The problem of firm (k, i) is then to maximise the expected profit⁽³⁾

$$\Pi_{k,i}(\mathcal{S}_{k,i}) = \left(1 - \mathbb{P}(\mathcal{S}_{k,i} \subseteq \mathcal{D}_{k-1})\right) \pi - C(|\mathcal{S}_{k,i}|) \quad (2)$$

by picking a sourcing strategy $\mathcal{S}_{k,i}$. Before moving on with the solution of the model, it is useful to discuss the assumptions presented in this section. The production game is highly stylised: first, firms do not adjust prices but only quantities, such that failure to produce only arises in the case that no input is sourced; second, they are able to obtain profits by simply producing; third, contracting with new suppliers has a cost. There are both theoretical and empirical reasons behind these choices. Theoretically, a simpler model allows us to isolate the interplay between the variables of interest: correlation in the risk of suppliers, supply chain opacity, and the endogenous production network fragility. Empirically, these assumption capture well the rationale behind firms' multisourcing. There is strong evidence that firms, first, when faced with supply chain shocks, adjust quantities rather than prices in the short run (), second, that production shutdowns can be extremely costly (), and third, that fostering relationships with suppliers is costly, but important in guaranteeing operational performance (Cousins and Menguc, 2006). The model establishes a link between these issues faced by firms when choosing a sourcing strategy and the fragility of the production network.

1.2. Opacity of the Supply Chain

The supply chain is opaque: firms cannot observe the sourcing decisions of their potential suppliers before making their own. Furthermore, firms do not know how risky individual basal producers are, nor how their risk is correlated. Yet, firms know the distribution from which the probabilities of disruption in the basal layer are drawn. To motivate this definition of opacity, recall the introductory example of hurricane Laura. A downstream firm producing PP, might not be able to trace back the production steps from its input to individual polymer producers in Louisiana or Texas, and, hence, the exact exposure of its production process to hurricanes. Yet, it can estimate the aggregate risk the polymer industry faces in the region. Given this information about the basal layer and their own depth k in the production network, firms can derive the distribution of risk among their suppliers and make sourcing decisions based on it. By symmetry, the risk of two firms downstream sourcing from the same number of suppliers is ex-ante identical, albeit possibly correlated. The following two assumptions formalise this idea. Introduce

$$X_{k,j} := \begin{cases} 1 & \text{if } (k, j) \text{ is disrupted and} \\ 0 & \text{otherwise,} \end{cases} \quad (3)$$

and \mathcal{P}_0 the space of probability distributions over the basal disruption events.

⁽³⁾The expectation is taken over the random set \mathcal{D}_{k-1} .

Assumption 1. Fix an arbitrary symmetric measure ν over \mathcal{P}_0 , that is, ν is invariant under relabelling of basal firms. The probability distribution \mathbb{P}_0 of disruptions in the basal layer is sampled from ν . I assume ν is observed by all firms, while \mathbb{P}_0 is hidden.

Going back to the example of polymer producers, under this assumption, downstream PP producers understand how hurricane risk can impact the production of their input good, via ν , yet, they cannot estimate the risk that individual polymer producers face, nor how this risk is correlated, since they do not observe \mathbb{P}_0 .

Assumption 2. If there are multiple sourcing strategies that yield the same expected profit, the firm chooses one with equal probability.

Proposition 1. Under these assumptions, in each downstream layer $k \geq 1$, disruption events

$$X_{k,1}, X_{k,2}, X_{k,3} \dots X_{k,m_k},$$

are exchangeable, that is, their distribution is invariant under permutation.

Proposition 1, proven in Appendix B.1, asserts that, from the point of view of the firm, all suppliers are ex-ante equal, yet their risk might be correlated. Hence, the profit of the firm depends exclusively on *how many* suppliers it chooses, rather than *which* suppliers it chooses. Hence, a firm producing good k can first infer the distribution of the number $D_{k-1} := |\mathcal{D}_{k-1}|$ of disrupted firms among its potential suppliers and then choose the optimal number $s_{k,i} := |\mathcal{S}_{k-1,i}|$ of firms from which to source its input good. Furthermore, by symmetry, all firms in layer k choose the same number s_k of sources, that is,

$$s_{k,i} = s_k \text{ for all } i. \quad (4)$$

As a result, the sourcing strategies $\mathcal{S}_{k,i}$ and $\mathcal{S}_{k,j}$ of any two firms i and j are such that their disruption probabilities $P_{k,i}$ and $P_{k,j}$ are identically distributed (\cdot).

2. Disruptions Propagation

Building on the mechanisms behind firms' disruptions introduced above, this section studies how these disruptions propagate through the supply chain. To do so, I consider the case in which the number m_k of firms in each layer k grows large. To study the limit, it is first necessary to characterise how the sourcing relations $\mathcal{S}_{k,j}$ form as the number of firms in each layer increases.

Assumption 3. As a new firm is introduced in layer k , it starts establishing relations with its s_k suppliers. As soon as it pairs with a supplier, a new firm is introduced among the producers of its input good $k - 1$, which, in turn, selects its sources. This procedure continues recursively until all firms realised their sourcing strategy s_k .

Indexing by n the n -th step of this procedure, this section focuses on the limit as $n \rightarrow \infty$. Every new firm introduced in the basal layer has a disruption probability that is ν -distributed, hence, the new firm is ex-ante identical to existing firms. This ensures that, as $n \rightarrow \infty$, Assumption 1 is satisfied and the downstream sourcing decisions $s_1, s_2 \dots$ are unaffected. This, allows us to simply consider the problem of the representative firm in layer k .

To analytically characterise the disruption propagation through the production network, the only missing piece is the distribution ν of the disruption probabilities in the basal layer. As mentioned in the previous section, I assume that basal firms fail with a not necessarily independent probability P_0 . We can model this by assuming that P_0 follows a Beta distribution.

Assumption 4. *The probability of a disruption in the basal layer follows*

$$P_0 \sim \nu_0 \equiv \text{Beta for all } j. \quad (5)$$

The Beta distribution allows to flexibly model shocks that might happen due to spacial or technological proximity of basal producers, which cannot be diversified. Consider, for example, how oil extraction plants must be located nearby oil reserves and are hence all subject to correlated weather shocks that might force them to shut down. In this case, despite the small expected probability that an individual firm is disrupted, as a hurricane is a rare occurrence, disruptions are highly correlated, as when a hurricane occurs most of them are disrupted. To keep track of the expected disruption probability and the correlation of risk through the layers, I introduce the following alternative parametrisation of the Beta distribution.

Definition 1. *Let μ and ρ be respectively the mean and the overdispersion of a Beta distribution with shape parameters α and β , defined by*

$$\mu := \frac{\beta}{\alpha + \beta} \text{ and } \rho := \frac{1}{1 + \alpha + \beta}. \quad (6)$$

I write $Y \sim \text{Beta}(\mu, \rho)$.

Given Assumption 4, the following result links the probabilities of experiencing a disruption from upstream suppliers of k to downstream producers $k + 1$.

Definition 2. *A random variable Y follows a BetaPower distribution, with mean μ , overdispersion ρ , and power s if it can be written as $Y = X^s$ where X follows a Beta distribution with mean μ and overdispersion ρ .*

Proposition 2. *If the disruption probability P_k among suppliers of good k follows a BetaPower distribution, so does the downstream probability P_{k+1} .*

The proof is provided in Appendix B.2. Proposition 2 guarantees that the distribution of disrupted firms will remain in the same distribution family as risk amplifies through the production network. This result allows us to describe disruption propagation in the supply chain by mapping the evolution of the parameters μ_k and ρ_k through the layers. Furthermore, it allows firms to estimate μ_k and ρ_k and use this to determine the optimal sourcing strategy s_{k+1} . It is useful at this point to give an interpretation of μ_k and ρ_k in the context of our model. The parameter μ_k is the average number of disrupted firms. The parameter ρ_k tracks the degree of correlation in the disruption of firms operating in layer k . I illustrate this in Figure 2. This figure shows the distribution of the disruption probability P_{k+1} among downstream firms in the case the firm has a single supplier (dotted lines) or two suppliers (solid line). For low overdispersion, $\rho_k = 0.01$, the suppliers' disruptions are weakly correlated and the downstream disruption probability is concentrated around the average μ_k . If firms contract an additional supplier, the distribution of failures decreases and remains concentrated around the average. As ρ_k increases, the suppliers' disruption events become more correlated and the downstream disruption probabilities become fat-tailed, that is, a significant fraction of firms is likely to be disrupted and, as a consequence, diversification is ineffective. If firms contract an additional supplier, risk decreases, but a large probability of disruptions remains.




Figure 2: Distribution of disruption probabilities of downstream firms for different levels of upstream correlation ρ_k , in the cases of single sourcing (dotted) and multisourcing (solid). In both cases $\mu_k = \frac{1}{2}$.

Having established the link between the disruptions of layer k to layer $k + 1$, I now turn to the analysis of how these propagate through the whole supply chain, before studying how firms make decisions endogenously. The following result recursively connects downstream distributions with upstream sourcing decisions and initial conditions.

Proposition 3. *The average number of disrupted firms between one layer k and the next $k + 1$ depends on the sourcing strategy s_{k+1} via*

$$\mu_{k+1} = \eta(s_{k+1}, S_k) \mu_k, \quad (7)$$

where $S_k := \prod_{j=1}^k s_j$ is the diversification level up to layer k and η is the risk reduction factor, which is given by

$$\begin{aligned} \eta(s_{k+1}, S_k) &= \left(\mu_0 \frac{1 - \rho_0}{\rho_0} + S_k \right)^{\overline{S_k s_{k+1}}} \bigg/ \left(\frac{1 - \rho_0}{\rho_0} + S_k \right)^{\overline{S_k s_{k+1}}} \\ &= \left(\frac{\mu_0 \frac{\rho_0}{1 - \rho_0} + S_k}{\frac{\rho_0}{1 - \rho_0} + S_k} \right) \left(\frac{\mu_0 \frac{\rho_0}{1 - \rho_0} + S_k + 1}{\frac{\rho_0}{1 - \rho_0} + S_k + 1} \right) \cdots \left(\frac{\mu_0 \frac{\rho_0}{1 - \rho_0} + S_k s_{k+1} - 1}{\frac{\rho_0}{1 - \rho_0} + S_k s_{k+1} - 1} \right). \end{aligned} \quad (8)$$

This is proven in Appendix B.3. The risk reduction factor $\eta(s_{k+1}, S_k)$ governs how the firms' choice s_{k+1} , the choices along the firms' production chain S_k , and the basal conditions μ_0, ρ_0 affect the expected number of disruptions downstream. This interplay is illustrated in the following figures.

Figure 3 shows how the risk reduction factor varies with basal correlation ρ_0 for different sourcing strategies s_k , fixing the upstream diversification of $S_k = 2$. If correlation ρ_0 in the basal layer grows, to obtain a given level of risk reduction η , firms producing good k needs to source more suppliers. If $\rho_0 \rightarrow 1$, diversification becomes impossible, as $\eta \rightarrow 1$ and $\mu_{k+1} \rightarrow \mu_k$ for any sourcing strategy s_{k+1} .




Figure 3: Risk reduction factor μ_{k+1}/μ_k at different basal overdispersion levels ρ_0 and for different sourcing strategies s_{k+1} . $S_k = 2$.

As the above, Figure 4 shows the response of the risk reduction factors to different levels of basal correlation, but instead of varying the strategy s_k of the firm, it varies the level of upstream diversification S_k . For low levels of basal correlation ρ_0 , more upstream diversification S_k allows downstream producers to achieve lower risk with fewer suppliers. Yet, there is a level of basal correlation after which more diversification is detrimental for the downstream firm, as this high upstream diversification simply exacerbates tail-risk. This represents a crucial externality the upstream suppliers impose on downstream producers. For low level of correlation, sourcing downstream represents a positive externality downstream. This externality shrinks as correlation increases, until it becomes a negative externality. Section 4 explores the welfare consequences of this mechanism.




Figure 4: Risk reduction factor μ_{k+1}/μ_k at different basal overdispersion levels ρ_0 and for different upstream diversification levels strategies S_k . $s_{k+1} = 2$.

3. Firm Optimal Diversification and Competitive Equilibrium

The mechanics of disruption propagation, derived in the previous section, determine the firms' risk diversification incentives and, as a consequence, their optimal sourcing strategy. This section derives such optimal strategies. Importantly, due to Proposition 1, all firms in a given layer are ex-ante identical and so is their optimisation problem. We can hence focus of the representative firm in layer $k+1$, choosing how many suppliers in layer k to source from, based on the inferred distribution of their probability of experiencing a disruption event. This, in turn, is fully determined by the expected fraction of firms disrupted in the basal layer μ_0 , the correlation of such disruptions ρ_0 , and the sourcing strategies of the representative firm upstream $\{s_k, s_{k-1} \dots s_0\}$. For illustration purposes in the next section, I assume firms face quadratic costs from contracting suppliers, with a fixed marginal cost c , such that expected profits⁽⁴⁾ can be written as

$$\Pi_k(s_k) = \left(1 - \mathbb{E}[P_k | s_k]\right) \pi - \frac{c}{2} s_k^2, \quad (9)$$

⁽⁴⁾Expectation is taken over the random variable $D_{k-1} = |\mathcal{D}_{k-1}|$.

The optimisation problem of the firm is to then choose the optimal sourcing strategy

$$s_k = \arg \max_{s \in \{0,1,2,\dots\}} \Pi_k(s). \quad (10)$$

3.1. Limit Case: Uncorrelated Disruptions

Before turning towards the general framework, I first analyse a limit case in which suppliers' risk is not correlated, $\rho_0 \rightarrow 0$. This limit case allows to derive more results analytically, gives a useful interpretation of the incentives behind multisourcing, and allows to establish a benchmark against which to study the introduction of correlated shocks.

Proposition 4. *If risk among basal firms becomes uncorrelated $\rho_0 \rightarrow 0$, disruption events in layer k become independent and happen with probability*

$$\mu_k = \mu_{k-1}^{s_k} = (\mu_{k-2}^{s_{k-1}})^{s_k} = \dots = \mu_0^{S_k}. \quad (11)$$

Proof. Follows immediately from $P_k \rightarrow \mu_k$ as $\rho_0 \rightarrow 0$. \square

Using the law of motion (11), we can rewrite profits (9) as

$$\Pi_{k+1}(s) = (1 - \mu_k^s)\pi - \frac{c}{2}s^2. \quad (12)$$

Consider the marginal profits attained by adding an additional source of input goods,

$$\Delta\Pi_{k+1}(s) := \Pi_{k+1}(s+1) - \Pi_{k+1}(s) = \mu_k^s(1 - \mu_k)\pi - c\left(s + \frac{1}{2}\right). \quad (13)$$

A firm with a given number of suppliers, contracts an extra supplier if doing so yields a positive expected marginal profit. Hence, the optimal number of suppliers s_{k+1} is the smallest s for which the expected marginal profit is negative $\Delta\Pi_{k+1}(s) < 0$.

Definition 3. *Let \tilde{s}_{k+1} be the unique root of $\Delta\Pi_{k+1}(s)$ over $[-1/2, \infty)$. I refer to this quantity as the “diversification incentive”.*

Proposition 5. *The optimal sourcing strategy is given by $s_{k+1} = \lceil \tilde{s}_{k+1} \rceil$.*

The proof of the uniqueness of \tilde{s}_{k+1} and of Proposition 5 is given in Appendix B.4.1.

Corollary 5.1. *The firm does not source any input $s_{k+1} = 0$ if the suppliers' risk μ_k is greater than the threshold μ^0 ,*

$$\mu_k \geq \mu^0 := 1 - \frac{c/2}{\pi}. \quad (14)$$

I refer to all levels of risk $\mu_k > \mu^0$ as “collapse regime”.

Proof. Suppose firms do not source any inputs $s_{k+1} = 0$. This implies that the marginal benefit of adding the first supplier is negative, namely $\Delta\Pi_{k+1}(0) \leq 0$, which yields the desired inequality. \square

As expected, the diversification incentive \tilde{s}_{k+1} and the optimal sourcing strategy s_{k+1} are implicitly determined by the suppliers' risk μ_k and the relative marginal cost of contracting a new supplier $\frac{c/2}{\pi}$. Figure 5 shows the optimal diversification incentive \tilde{s}_{k+1} (solid line) and sourcing strategy s_{k+1} (dotted line), given a suppliers' level of risk μ_k , for three different marginal costs ratios $\frac{c/2}{\pi}$. First, as the marginal pairing costs increase, the firm's diversification incentives decrease. Second, higher levels of risk increase the diversification incentives of firms until a threshold is reached, after which, diversification incentives and, in turn, optimal diversification start decreasing. The concavity of the optimal sourcing strategy s_{k+1} in the suppliers' risk μ_k , highlights a vicious cycle which is introduced in production networks when firm decisions are endogenised: after a certain level of risk, firms have less incentives to diversify, which, in turn, "flattens" the marginal expected profits downstream. This mechanisms can lead to endogenous fragilities in production networks.




Figure 5: \tilde{s}_{k+1} (solid) and s_{k+1} (dotted)

To illustrate the mechanism behind the endogenous fragility, we can combine the law of risk propagation (4) and the optimal firm sourcing (5) to obtain the risk dynamics in the production network,

$$\mu_{k+1} = \mu_k^{s_{k+1}(\mu_k)}. \quad (15)$$

The recursive relation (15) maps risk from suppliers to firms throughout the production network, hence studying its properties allows to describe how risk propagates through the supply chain. Two natural questions arise: first, which levels of risk remain stable

through the supply chain, that is, they are neither amplified nor dampened as one moves from suppliers to firms? Second, how are initial levels of risk mapped to these stable risk levels? The former can be answered by determining levels of risk $\bar{\mu}$ which are fixed points of equation (15). The latter by looking at the basin of attraction of such points.

Proposition 6. *The stable levels of risk $\bar{\mu}$ are all points satisfying*

$$\frac{\bar{\mu}(1 - \bar{\mu})}{3} \leq \frac{c/2}{\pi}. \quad (16)$$

Proof. A steady state is attained if $\tilde{s}_{k+1} \leq 1$. This implies that the marginal benefit of multisourcing is not positive $\Delta\Pi_{k+1}(1) \leq 0$. This yields the desired inequality. \square

Corollary 6.1. *For all basal risk levels μ_0 larger than the critical threshold*

$$\mu^c := \frac{1}{2} + \sqrt{\frac{1}{4} - 3\frac{c/2}{\pi}}, \quad (17)$$

the endogenous supply chain is unable to diversify risk $\bar{\mu} = \mu_0$. I refer to this situation as “endogenous fragility”.

Propositions 6 and 6.1 (proven in Appendix B.4.2) link the firms’ relative costs $\frac{c/2}{\pi}$ and the aggregate outcomes in term of production network risk. As relative marginal costs increase, the capacity of the production network to endogenously diversify basal risk μ_0 decreases and firms’ underdiversification yields endogenous fragility. It is interesting to notice that, comparing the aggregate threshold μ^c with the firm shutdown threshold μ^0 (Figure 6), for some levels of basal risk μ_0 , despite no firm shutting down production $\mu_0 < \mu^0$, the production network as a whole is still unable to endogenously diversify risk $\mu_0 > \mu^c$. This is true even in this special case, where firms’ risk is uncorrelated. In the next section, I introduce correlation risk $\rho > 0$ and investigate how doing so changes the dynamics illustrated here.




Figure 6

3.2. Optimal Sourcing with Correlated Distributions

If disruption events are not independent, $\rho_0 > 0$, risk among suppliers throughout the production network is correlated, and firms’ optimisation incentives change. The firm problem (9) in layer $k + 1$ is to choose the number of suppliers $s_{k+1} \in \{0, 1, 2, \dots\}$ that

maximises the profits Π given an upstream diversification S_k . The firm's expected disruption probability is given by the expected probability of disruption among its suppliers μ_k , scaled by a factor $\eta(s_{k+1}, S_k)$ (Proposition 3) which depends on the firm sourcing strategy. As in the limit case analysed in the previous section, the firm will increase diversification as long as the expected reduction in profits obtained by adding an additional supplier outweighs the costs of contracting that additional supplier. The expected marginal profits are given by

$$\Delta\Pi_{k+1}(s_{k+1}) = \left(\eta(s_{k+1}, S_k) - \eta(s_{k+1} + 1, S_k) \right) \mu_k \pi - c \left(s_{k+1} + \frac{1}{2} \right). \quad (18)$$

As above, let \tilde{s}_{k+1} be the “diversification incentive”, that is, the level of s such that the marginal benefits and marginal costs of diversification are equal $\Delta\Pi_{k+1}(s) = 0$. Since the marginal profits are strictly decreasing in the number of suppliers (see Appendix B.5), the firm will choose its optimal sourcing strategy as $s_{k+1} = \lceil \tilde{s}_{k+1} \rceil$ (Proposition 5). The correlation of risk among producers of its input good changes the firm's incentives, as illustrated in Figure 7. In particular, as correlation increases, the firm needs to increase its sources to diversify risk. Yet, for large levels of correlation, the disruption of an additional source of the input good is likely correlated to a disruption among the firm's existing suppliers, which reduces the firm's multisourcing incentives. As disruptions among suppliers become perfectly correlated, $\rho_k \rightarrow 1$, the firm has no reason to multisource $\tilde{s}_{k+1} \rightarrow 0$.




Figure 7: Sourcing incentive (i.e. s such that $\Delta\Pi_{k+1}(s) = 0$) as suppliers' risk correlation ρ_k increases, for different values of suppliers' risk, μ_k

This result, proven in Appendix B.5, can be summarised as follows.

Proposition 7. *The sourcing incentive \tilde{s}_{k+1} is concave in the overdispersion among suppliers ρ_k .*

Corollary 7.1. *The optimal sourcing strategy is weakly concave in the overdispersion among suppliers ρ_k .*

Proof. If \tilde{s}_{k+1} is concave in ρ_0 (Proposition 7), then taking the next largest integer $\lceil \cdot \rceil$ yields weak concavity. \square

This result (7.1) introduces a channel through which the correlation of disruptions in the production networks, can generate externalities in the firms' choices and yield an endogenously fragile production network. To analyse these ramifications, consider the recursive relation of risk μ_k , when firms optimally source s_{k+1} (parallel to equation 15),

$$\mu_{k+1} = \eta(s_{k+1}, S_k) \mu_k = \dots = \prod_{j=0}^k \eta(s_{k+1-j}, S_{k-j}) \mu_0. \quad (19)$$

From the definition of η , it follows that downstream there exists a layer \bar{k} , such that, the distribution of disruptions is constant, or more formally, $s_l = 1$ for all $l \geq \bar{k}$. Let $\bar{\mu} := \mu_K$ be the stable downstream fraction of firms expected to fail downstream. Figure 8 illustrates this fraction for different levels of basal risk μ_0 and overdispersion ρ_0 in a low (left) and a high (right) relative pairing cost regime. First, in both cases if $\rho_0 \rightarrow 1$ there is no possible diversification since all firms are either disrupted or not, hence the disruption risk is constant across the production network $\bar{\mu} = \mu_0$, regardless of the sourcing strategy. Likewise, if risk among basal producers becomes independent, $\rho_0 \rightarrow 0$, the problem converges to the limit case discussed in the previous section. Second, for a given level of basal correlation ρ_0 , there is a critical threshold μ^c such that an arbitrary small increase in the initial risk level μ_0 has a discontinuous effect in the fraction of downstream firms $\bar{\mu}$. As discussed in the previous section, this discontinuity is induced on the production network by the firms' endogenous diversification incentives. As correlation increases, the maximum “diversifiable” level of basal risk, μ^c , becomes smaller. This lowering critical threshold suggests that, when allowed to form endogenously, production networks display a tendency towards endogenous fragility. Importantly, the model suggests that cascading failures can be triggered not only by increases in basal risk μ_0 but also by increases in basal correlation ρ_0 , even as basal risk remains constant. This fragility, unlike the one studied by Elliott et al. (2022), is not induced by the structure of the production function, but rather arises endogenously due to the decreasing diversification motives of firms.




Figure 8: Surface of the downstream fraction of non functioning firms $\bar{\mu}$, given the initial condition μ_0 and ρ_0 in a low (left) and a high (right) relative pairing costs regime.

4. Social Planner Problem

To establish a benchmark to which one can compare the competitive equilibrium analysed above, in this section I solve the model from the perspective of a social planner. The social planner attempts to, on the one hand, minimise the number of firms expected to fail, and, on the other, minimise the number of costly sourcing relations. To develop a useful benchmark, it is important to define a social planner problem that can be meaningfully compared to the decentralised firm’s problem. To do so, I make the following two assumptions.

Assumption 5. *The social planner knows the distribution of failure in the basal layer $P_0 \sim \text{Beta}(\mu_0, \rho_0)$ and makes decision before P_0 is realised.*

Assumption 6. *As in the firm problem, I consider the case $n \rightarrow \infty$. This allows the social planner to recursively, from the last layer K upwards, assign suppliers $\mathcal{S}_{k,i}$ such that there are sufficiently many firms so that no two firms share suppliers $\mathcal{S}_{k,i} \cap \mathcal{S}_{k,j} = \emptyset$.*

To understand the intuition behind Assumption 6, consider the possible supplier overlap illustrate in Figure 9: if a supplier has multiple downstream clients (dashed box), the social planner can always rewire a link towards a supplier without downstream clients (solid box). By doing so, the social planner can “diversify away” all the correlation that arises due to the network structure. Hence, the only source of risk in the model is represented by the shutdowns experienced by firms in the basal layer, which




Figure 9

happen with non-idiosyncratic probabilities P_0 (Assumption 5).

Combining Assumptions 5 and 6, the social planner problem is then to maximise average expected payoffs

$$W(\{\mathcal{S}_{k,i}\}) := \frac{1}{K} \sum_{k=0}^K \lim_{n \rightarrow \infty} \frac{1}{m_k(n)} \sum_{i=1}^{m_k(n)} \left(1 - \mathbb{P}(\mathcal{S}_{k,i} \subset \mathcal{D}_{k-1})\right) \pi - \frac{c}{2} |\mathcal{S}_{k,i}|^2, \quad (20)$$

by choosing a sourcing strategy $\mathcal{S}_{k,i} \subseteq \{1, 2, \dots\}$ for each firm in each layer such that $\mathcal{S}_{k,i} \cap \mathcal{S}_{k,j}$ is empty for all i, j . The social planner problem can be further simplified by noticing that, given that all firms in layer k are identical, if establishing an additional path from a firm in layer k to a basal firm has positive marginal benefits, then it has positive marginal benefits for all firms in layer k which share the same number of paths to basal firms. Hence, as in the decentralised firms' problem, the social planner can choose the optimal number of sources in each layer, let the firms source at random, and finally disentangle any overlapping paths. Using this, the social planner problem can be formulated recursively, by letting V_k be the maximal average welfare from layer k to the last layer K . This can be recursively defined as

$$V_k(P_{k-1}) = \max_{s_k} \left\{ \left(1 - \mathbb{E}[P_{k-1}^{s_k}]\right) \pi - \frac{c}{2} s_k^2 + \mathbb{E}[V_{k+1}(R_k)] \right\} \quad (21)$$

where the state $P_{k-1} \sim \text{BetaPower}(\mu_0, \rho_0, s_1 s_2 \dots s_{k-1})$ is the fraction of disrupted firms, which evolves as

$$P_k = P_{k-1}^{s_k}. \quad (22)$$

The average welfare in layer $K+1$ is given by $V_{K+1}(P_K) = 0$, since firms in the last layer are never sources to other firms, and an initial state condition $P_0 \sim \text{Beta}(\mu_0, \rho_0)$. This problem can be solved using standard backward induction techniques (see Appendix C). The optimum average social welfare (20) can then be written as

$$V_1(P_0) = \max_{s_1, s_2, \dots, s_{K-1}} W(\{s_1, s_2, \dots, s_{K-1}\}). \quad (23)$$

Letting $\{s_k^p\}_{k=1}^K$ be the socially optimal sourcing strategies sequence and $\{\mu_k^p\}_{k=1}^K$ be the expected disruption in each layer given by such sourcing strategies, we can compute the change in downstream risk compared to the decentralised case (as illustrated in Figure 8). Figure 10 shows this difference $\bar{\mu} - \bar{\mu}^p$ for the same two cost regimes. On the one hand, when relative pairing costs are low (left panel), the social planner achieves marginally lower risk levels of downstream risk for most initial conditions. If initial basal correlation ρ_0 is sufficiently large and basal risk μ_0 is sufficiently low, the firms overdiversify compared to the socially optimum $\bar{\mu} < \bar{\mu}^p$. On the other hand, if relative pairing costs are high (right panel), the differences between the social optimum and the competitive level of downstream risk are starker. First, for larger levels of basal risk μ_0 and lower levels of basal correlation ρ_0 the firms overdiversify compared to the social optimum, $\bar{\mu} < \bar{\mu}^p$. Second, the social planner is able to diversify risk around the critical threshold μ_c , $\bar{\mu} > \bar{\mu}^p$. This implies that the cascading failures that occur around

the critical threshold are fully attributable to firms' endogenous underdiversification motives.




Figure 10: Change in downstream expected failure probability between the firms' $\bar{\mu}$ and the social planner $\bar{\mu}^p$ equilibrium, given different initial conditions μ_0 and ρ_0 in a low (left) and a high (right) relative pairing costs regime.

The differences between the firms' sourcing strategies and the social optimum generate welfare losses in the production network. Letting W be the average firm profit in the decentralised case and W^p be the average profit achieved by the social planner, in Figure 11 I show the welfare loss due to the firms' diversification decisions $W - W^p$, in the high cost scenario ($\frac{c}{2} = 0.02\pi$). The welfare loss is largest around the critical value μ^c , where the production network suffer is endogenously fragile. At these levels of risk, firms' upstream firms' diversification incentives are weak, which creates large downstream resilience externalities. Crucially, both an increase in basal risk μ_0 and an increase in basal correlation ρ_0 can generate discontinuous welfare losses.




Figure 11: Welfare loss of the decentralised equilibrium $W - W^p$, given different initial conditions μ_0 and ρ_0 .

5. The Role of Opacity

So far I assumed that firms cannot observe the realisation of the supply chain when making sourcing decisions. To understand how this assumption affects optimal decisions and fragility within the supply chain, I now analyse the supply chain under perfect information. The following assumption clarifies what is meant by perfect information in the context of the model.

Assumption 7. *In a regime of perfect information each firm i in level k is able to perfectly estimate the disruption probability of each potential supplier and the full correlation structure of the disruption events.*

Under this perfect information regime, the firm can assign correct probabilities to its own disruption risk

$$\mathbb{P}(\mathcal{S}_{k,i} \subset \mathcal{D}_{k-1}) \text{ for all possible sourcing strategies } \mathcal{S}_{k,i}.$$

The firm can hence rank suppliers by the marginal reduction in risk they provide and source from the “safest” s_k desired suppliers. As all firms downstream are ex-ante

identical, the marginal benefits of diversification experienced by firm (i, j) are the same as those of all other firms in layer k , which implies that, in equilibrium, all firms in layer k will employ the same sourcing strategy, given that they are ex-ante identical. This outcome is beneficial for any single firm, but detrimental for the stability of the production network. The following to propositions formalise this.

Proposition 8. *Compared to the opaque scenario, for the same number of sources, firm (k, i) is (weakly) less likely to be disrupted.*

Proof. Given the same number of sources, the firm with perfect information minimises its disruption risk with fewer constraints than in the opaque scenario. \square

Proposition 9. *Under perfect information, the supply chain is maximally fragile: either all firms fail or none do.*

Proof. The equilibrium outcome under perfect information implies that a disruption in layer k affects all firm simultaneously as they all share the same sources, namely

$$X_{k,i} = X_{k,j} \text{ for all } i \text{ and } j.$$

This, in turn, removes any diversification incentives downstream: a firm producing $k+1$ cannot diversify its risk by multisourcing, hence it will single source. \square

Opacity has a dual role: on the one hand, it prevents firms from optimally choosing the best sourcing strategy; on the other hand, it mitigates endogenous fragility by forcing firm to diversify.

6. Conclusion

Risk diversification is a crucial determinant of firms' sourcing strategies. In this paper, I show that firms endogenously underdiversify risk when they have incomplete information about upstream sourcing relations. This endogenous underdiversification generates fragile production networks, in which, arbitrarily small increases in correlation among disruptions between basal producers can generate discontinuously large disruptions downstream. I do so by deriving an analytical solution to a simple production game in which firms' sole objective is to minimise the risk of failing to source input goods.

Despite its simple structure, the game identifies an important externality firms impose on the production network when making sourcing decisions: upstream multisourcing introduces correlation in firms' risk, which reduces incentives to multisource downstream. This externality exacerbates the risk of fragile production networks. Furthermore, I show that a social planner can design production networks that mitigate fully this externality.

A consequence of this result is that, in principle, it is possible to design a transfer mechanism that allows downstream firms to compensate upstream firms to internalise the diversification externalities. By analysing the welfare loss in competitive equilibrium, I show that there is a critical region of basal conditions where arbitrary small increases in suppliers' correlation, even if the expected number of disrupted firms remains constant, can generate catastrophic downstream disruptions by altering downstream firm diversification incentives. This result illustrates how an increase in correlation among basal producers, for example due to widespread offshoring to the same country, can endogenously generate fragile production networks which have a large tail risk of disruption. Surprisingly, opacity plays a role in mitigating this effect, suggesting that if firms were to acquire information about the production network, despite the individual firm being better off, this could generate further endogenous fragility.

References

- Leontief, W. (1936). Quantitative Input and Output Relations in the Economic Systems of the United States [MAG ID: 1977856930]. *The Review of Economics and Statistics*, 18(3), 105. <https://doi.org/10.2307/1927837>
- Hulten, C. R. (1978). Growth Accounting with Intermediate Inputs. *The Review of Economic Studies*, 45(3), 511–518. <https://doi.org/10.2307/2297252>
- Tan, J. S., & Kramer, M. A. (1997). A general framework for preventive maintenance optimization in chemical process operations. *Computers & Chemical Engineering*, 21(12), 1451–1469. [https://doi.org/10.1016/S0098-1354\(97\)88493-1](https://doi.org/10.1016/S0098-1354(97)88493-1)
- Kallenberg, O. (2005). *Probabilistic Symmetries and Invariance Principles*. Springer-Verlag. <https://doi.org/10.1007/0-387-28861-9>
- Cousins, P. D., & Menguc, B. (2006). The implications of socialization and integration in supply chain management. *Journal of Operations Management*, 24(5), 604–620. <https://doi.org/10.1016/j.jom.2005.09.001>
- Diaconis, P., & Janson, S. (2007). Graph limits and exchangeable random graphs [arXiv:0712.2749 [math]]. Retrieved October 19, 2023, from <http://arxiv.org/abs/0712.2749>
- di Giovanni, J., & Levchenko, A. A. (2010). Putting the Parts Together: Trade, Vertical Linkages, and Business Cycle Comovement [MAG ID: 1980074098]. *American Economic Journal: Macroeconomics*, 2(2), 95–124. <https://doi.org/10.1257/mac.2.2.95>
- Gabaix, X. (2011). The Granular Origins of Aggregate Fluctuations [MAG ID: 3122444308]. *Econometrica*, 79(3), 733–772. <https://doi.org/10.3386/w15286>
- Acemoglu, D., Carvalho, V. M., Ozdaglar, A., & Tahbaz-Salehi, A. (2012). The Network Origins of Aggregate Fluctuations. *Econometrica*, 80(5), 1977–2016. <https://doi.org/10.3982/ECTA9623>
- Williams, B. D., Roh, J., Tokar, T., & Swink, M. (2013). Leveraging supply chain visibility for responsiveness: The moderating role of internal integration. *Journal of Operations Management*, 31(7-8), 543–554. <https://doi.org/10.1016/j.jom.2013.09.003>
- Erol, S., & Vohra, R. (2014). Network Formation and Systemic Risk [MAG ID: 599238147 S2ID: 5ca3a071ddd810c2759c8aba0e153aad071f58ad]. *European Economic Review*. <https://doi.org/10.2139/ssrn.2500846>
- Hameed, A., & Khan, F. (2014). A framework to estimate the risk-based shutdown interval for a processing plant. *Journal of Loss Prevention in the Process Industries*, 32, 18–29. <https://doi.org/10.1016/j.jlp.2014.07.009>
- Macchiavello, R., & Morjaria, A. (2015). The Value of Relationships: Evidence from a Supply Shock to Kenyan Rose Exports. *American Economic Review*, 105(9), 2911–2945. <https://doi.org/10.1257/aer.20120141>
- Carvalho, V. M., Nirei, M., Saito, Y., & Tahbaz-Salehi, A. (2016). Supply Chain Disruptions: Evidence from the Great East Japan Earthquake [MAG ID: 1524026022]. <https://doi.org/10.2139/ssrn.2883800>
- Baqae, D. R. (2018). Cascading Failures in Production Networks [MAG ID: 2188284008]. *Econometrica*, 86(5), 1819–1838. <https://doi.org/10.3982/ecta15280>

- Baqae, D. R., & Farhi, E. (2019). The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten’s Theorem. *Econometrica*, 87(4), 1155–1203. <https://doi.org/10.3982/ECTA15202>
- Carvalho, V. M., & Tahbaz-Salehi, A. (2019). Production Networks: A Primer. *Annual Review of Economics*, 11(1), 635–663. <https://doi.org/10.1146/annurev-economics-080218-030212>
- Zhao, M., & Freeman, N. K. (2019). Robust Sourcing from Suppliers under Ambiguously Correlated Major Disruption Risks. *Production and Operations Management*, 28(2), 441–456. <https://doi.org/10.1111/poms.12933>
- Acemoglu, D., & Azar, P. D. (2020). Endogenous Production Networks. *Econometrica*, 88(1), 33–82. <https://doi.org/10.3982/ECTA15899>
- Amelkin, V., & Vohra, R. (2020). Strategic Formation and Reliability of Supply Chain Networks [arXiv: 1909.08021]. *arXiv:1909.08021 [cs, econ, eess, math, q-fin]*. Retrieved March 1, 2022, from <http://arxiv.org/abs/1909.08021>
- Kopytov, A., Bineet Mishra, Nimark, K., & Taschereau-Dumouchel, M. (2021). Endogenous Production Networks under Uncertainty [MAG ID: 3212064397]. *Social Science Research Network*. <https://doi.org/10.2139/ssrn.3936969>
- Vakil, B. (2021). The Latest Supply Chain Disruption: Plastics. *Harvard Business Review*. <https://hbr.org/2021/03/the-latest-supply-chain-disruption-plastics>
- Elliott, M., Golub, B., & Leduc, M. V. (2022). Supply Network Formation and Fragility. *American Economic Review*, 112(8), 2701–2747. <https://doi.org/10.1257/aer.20210220>
- Jiang, B., Rigobon, D., & Rigobon, R. (2022). From Just-in-Time, to Just-in-Case, to Just-in-Worst-Case: Simple Models of a Global Supply Chain under Uncertain Aggregate Shocks. *IMF Economic Review*, 70(1), 141–184. <https://doi.org/10.1057/s41308-021-00148-2>
- Lafrogne-Joussier, R., Martin, J., & Mejean, I. (2022). Supply Shocks in Supply Chains: Evidence from the Early Lockdown in China. *IMF Economic Review*. <https://doi.org/10.1057/s41308-022-00166-8>

A. Notation and Distributions

This appendix introduces standard notation and definitions that will be used throughout the following appendices.

For $x \in \mathbb{R}$ and $n \in \mathbb{Z}$, I denote the falling factorial as

$$x^{\overline{n}} := \underbrace{x(x-1)(x-2)\dots(x-(n+1))}_{n \text{ terms}}. \quad (24)$$

For non-integer exponents $s \in \mathbb{R}$, (24) can be extended as

$$x^{\overline{n}} := \frac{\Gamma(x+n)}{\Gamma(x)}, \quad (25)$$

where Γ is the gamma function.

Two properties of the falling factorial that are used below but not proven are the additive property of the exponent

$$x^{\overline{n+m}} = x^{\overline{n}}(x+m)^{\overline{m}}, \quad (26)$$

and that it is strictly increasing in its base

$$\frac{\partial x^{\overline{n}}}{\partial x} > 0. \quad (27)$$

B. Omitted Proofs

This appendix contains the proofs omitted from the paper.

B.1. Proof of Proposition 1

Proving Proposition 1, requires the following Lemma.

Lemma 1. *If the disruption events among upstream firms are exchangeable, then the probability that a downstream firm is disrupted depends only on the number of suppliers it picks.*

Proof. Consider the (possibly infinite) sequence of disruption events among upstream firms $X_{k,1}, X_{k,2}, X_{k,3} \dots$. We assume the sequence to be exchangeable, that is,

$$X_{k,1}, X_{k,2}, X_{k,3} \dots \stackrel{d}{=} X_{k,\sigma(1)}, X_{k,\sigma(2)}, X_{k,\sigma(3)} \dots, \quad (28)$$

for an arbitrary permutation of its indices σ . Fix two arbitrary finite subsets of firms $\mathcal{A} = \{X_{k,a_1}, X_{k,a_2} \dots X_{k,a_n}\}$ and $\mathcal{B} = \{X_{k,b_1}, X_{k,b_2} \dots X_{k,b_n}\}$ of size n . Here $a_i, b_i \in [m]$

are the indices of the original sequence corresponding to the i -th index of the subset. Let σ be a permutation that takes elements of \mathcal{A} to \mathcal{B} , namely,

$$\sigma(\mathcal{A}) = \mathcal{B} \text{ and } \sigma(\mathcal{A}^c) = \mathcal{B}^c. \quad (29)$$

Then the probability distribution over \mathcal{A} is

$$\begin{aligned} \mathbb{P}(\mathcal{A}) &= \mathbb{P}(\mathcal{A} \text{ and } \mathcal{A}^c \text{ taking any value}) \\ &= \mathbb{P}(\sigma(\mathcal{A}) \text{ and } \sigma(\mathcal{A}^c) \text{ taking any value}) \text{ then by exchangeability} \\ &= \mathbb{P}(\mathcal{B} \text{ and } \mathcal{B}^c \text{ taking any value}) = \mathbb{P}(\mathcal{B}). \end{aligned} \quad (30)$$

□

Now we can prove Proposition 1

Proof. It can be proven by induction.

The base case $k = 0$ follows from Assumption 1, as the disruption probabilities are ν -distributed and ν is symmetric.

Assume that for some layer $k - 1$ the disruption events $X_{k-1,1}, X_{k-1,2}, \dots, X_{k-1,m}$ are exchangeable. By Lemma 1 we know that the downstream expected profit $\Pi_{k,i}(\mathcal{S})$ depends only the number of suppliers $|\mathcal{S}|$. By symmetry and Assumption 2, all firms in layer k are then selecting a random subset of supplier from layer $k - 1$ with equal probability, which in turn determines their disruption risk $X_{k,i}$. This construction is independent of the downstream firm index i , hence

$$X_{k,1}, X_{k,2}, \dots, X_{k,m}$$

are exchangeable. □

B.2. Proof of Proposition 2

Proof. A firm producing good $k + 1$ sources from s_{k+1} suppliers, hence, its disruption event is given by

$$X_{i,k+1} = X_{j_1,k} X_{j_2,k} \dots X_{j_{s_{k+1}},k}, \quad (31)$$

where $\{j_1, j_2, \dots, j_{s_{k+1}}\}$ is an arbitrary subset of suppliers and $X_{j,k}$ are exchangeable Bernoulli trials with a P_k success probability, where

$$P_k \sim \text{BetaPower}.$$

Conditional on the underline distribution P_k of disruption probabilities, the trials $X_{j,k}$ are independent and identically distributed. Hence, we have

$$\begin{aligned}
P_{k+1} &= \mathbb{E}[X_{i,k+1}] = \mathbb{E}\left[\prod_{l=1}^{s_k} X_{j_l,k}\right] \\
&= \mathbb{E}\left[\mathbb{E}\left[\prod_{l=1}^{s_k} X_{j_l,k} \mid P_k = p_{j_l}\right]\right] \text{ by conditional independence,} \\
&= \mathbb{E}\left[\mathbb{E}\left[X_{j_l,k} \mid P_k = p_{j_l}\right]^{s_k}\right] \text{ by independence of the draws } p_{j,l} \\
&= \mathbb{E}\left[\mathbb{E}\left[X_{j_l,k} \mid P_k = p_{j_l}\right]^{s_k}\right] = P_k^{s_k}.
\end{aligned} \tag{32}$$

□

B.3. Mapping of risk across layers

This section derives the risk reduction factor η .

Lemma 2. *If $P_{k-1} \sim \text{BetaPower}(m, \alpha, \beta, S)$ for some integer S , then*

$$P_k \sim \text{BetaPower}(m, \alpha, \beta, S s_k) \tag{33}$$

where s_k is the choice of suppliers in layer k .

Proof. Follows from the definition of BetaPower. □

Proposition 10. *The expected probability of disruption faced by a firm is given by*

$$\mathbb{E}[P_k] = \frac{B\left(\mu_0 \frac{1-\rho_0}{\rho_0} + S_{k-1} s_k, (1 - \mu_0) \frac{1-\rho_0}{\rho_0}\right)}{B\left(\mu_0 \frac{1-\rho_0}{\rho_0}, (1 - \mu_0) \frac{1-\rho_0}{\rho_0}\right)}. \tag{34}$$

Proof. It follows from rewriting the moment generating function of the beta distribution as

$$M(t) = \sum_{n=0}^{\infty} \frac{t^n}{n!} \frac{B(\alpha + n, \beta)}{B(\alpha, \beta)}. \tag{35}$$

and noticing that $P_k \sim \text{BetaPower}(\alpha, \beta, S_{k-1} s_k)$. □

To simplify notation, let $r_0 = \frac{1-\rho_0}{\rho_0}$ and

$$\eta(s, S) = \frac{(\mu_0 r_0 + S)^{\overline{S(s-1)}}}{(r_0 + S)^{\overline{S(s-1)}}} \tag{36}$$

which satisfies the recursion

$$\eta(s+1, S) = \eta(s, S) \frac{(\mu_0 r_0 + Ss)^{\bar{S}}}{(r_0 + Ss)^{\bar{S}}}. \quad (37)$$

Another property of (36) which will be use later is

$$\frac{\partial \eta}{\partial s} = \eta(s, S) S \left(\psi(\mu_0 r_0 + Ss) - \psi(r_0 + Ss) \right). \quad (38)$$

Corollary 10.1. *From equation (38) and the fact that ψ is increasing over positive values, it follows that η is decreasing in s .*

Using Proposition (10), the coefficient η allows us to write the propagation of risk recursively

$$\mu_{k+1} = \eta(s_{k+1}, S_k) \mu_k. \quad (39)$$

B.4. Limit case $\rho_0 \rightarrow 0$

In this section, for convenience, I suppress the layer indices k and $k+1$.

B.4.1. Optimal sourcing strategy

This proves Proposition 5 and that Definition 3 is well defined.

Proof. Notice that $\pi > 0$ and $0 < \mu < 1$, the expected marginal benefit $\Delta\Pi : \mathbb{R} \rightarrow \mathbb{R}$ is strictly decreasing as

$$\frac{\partial}{\partial s} \Delta\Pi(s) = \log(\mu) (1 - \mu) \mu^s \pi - c < 0. \quad (40)$$

Furthermore, $\Delta\Pi$ is continuous with

$$\Delta\Pi(-1/2) = \frac{1 - \mu}{\sqrt{\mu}} \pi > 0 \text{ and} \quad (41)$$

$$\lim_{s \rightarrow \infty} \Delta\Pi(s) = -\infty < 0. \quad (42)$$

Hence, \tilde{s} exists and is unique in the interval $[-1/2, \infty)$. Since \tilde{s} is the unique root and $\Delta\Pi$ is strictly decreasing the integer value $s = \lceil \tilde{s} \rceil$ is the smallest integer such that the marginal benefits are negative $\Delta\Pi(s) < 0$. \square

B.4.2. Properties of the Law of Motion

For the following proof I only consider non-trivial values of upstream risk $\mu < \mu^0$. If $\mu \geq \mu^0$, no firm has suppliers and the supply chain is by definition stable.

Lemma 3. *A fixed point of the law of motion $g(\bar{\mu}) = \bar{\mu}$, is attained iff $\tilde{g}(\bar{\mu}) \geq \bar{\mu}$.*

Proof. By definition $\tilde{g}(\bar{\mu}) = \bar{\mu}^{\tilde{s}(\bar{\mu})}$ and $0 \leq \bar{\mu} \leq 1$. Hence $\tilde{g}(\bar{\mu}) \geq \bar{\mu} \iff \tilde{s}(\bar{\mu}) \in (0, 1]$. By definition $s(\bar{\mu}) = \lceil \tilde{s}(\bar{\mu}) \rceil = 1$, which implies that $g(\bar{\mu}) = \bar{\mu}$. \square

Now we can prove Corollary 6.1.

Proof. We seek μ , such that $\tilde{g}(\mu) \geq \mu$, which then implies that $g(\mu) = \bar{\mu}$. This will be the case if $\tilde{s}(\mu) \in (0, 1]$. This is the case if $\Delta\Pi(1) \leq 0$ and $\Delta\Pi(0) > 0$, which yields the desired inequality. \square

B.5. General Case, $\rho_0 > 0$

This is the proof of Proposition 5 in the more general case, $\rho_0 > 0$.

Proof. It is sufficient to show that $\Delta\Pi$ is strictly decreasing in s when $\rho_0 > 0$. It is convenient to rewrite η (36) as

$$\eta(s) = \frac{\Gamma(r_0 + S)}{\Gamma(\mu_0 r_0 + S)} \frac{\Gamma(\mu_0 r_0 + Ss)}{\Gamma(r_0 + Ss)}. \quad (43)$$

Then

$$\Delta\Pi(s) = (\eta(s) - \eta(s+1))\pi\mu - c \left(s + \frac{1}{2} \right), \quad (44)$$

hence

$$\begin{aligned} \Delta\Pi'(s) = & -c - \mu\pi S \left(\eta(s+1) \left(\psi(\mu_0 r_0 + S(s+1)) - \psi(r_0 + S(s+1)) \right) - \right. \\ & \left. \eta(s) \left(\psi(\mu_0 r_0 + Ss) - \psi(r_0 + Ss) \right) \right). \end{aligned} \quad (45)$$

Then $\Delta\Pi'(s) < 0$, since ψ is increasing. Finally notice that $\Delta\Pi(-1/2) = (\eta(-1/2) - \eta(1/2))\pi\mu < 0$ and $\lim_{s \rightarrow \infty} \Delta\Pi(s) = \infty$. \square

After proving that a solution exists, I will now prove that it is concave in the level of correlation ρ_0 (Proposition 7).

Proof. Notice that another way of writing $\Delta\Pi$ is letting $P \sim \text{BetaPower}(\mu, \rho, s)$ and writing

$$\Delta\Pi(s) = \mathbb{E}\left[P^s - P^{s-1}\right]\pi - c\left(s + \frac{1}{2}\right). \quad (46)$$

The optimal incentive s , is a root of $\Delta\Pi$, hence

$$s = \frac{\mathbb{E}\left[P^s - P^{s-1}\right]}{c}\pi\mu - \frac{1}{2} = \frac{\pi\mu}{c}\left(\mathbb{E}[P^s] - \mathbb{E}[P^{s-1}]\right) - \frac{1}{2}. \quad (47)$$

□

C. Solution of the Social Planner Problem

First, notice that the terminal condition V_K is linear in P_{K-1} , hence $\mathbb{E}V_K(P_{K-1}) = V_K(\mathbb{E}P_{K-1})$. In turn, this implies that V_k is linear for all k . Hence we can rewrite the value to be a function of the state space S ,

$$V_k(S_{k-1}) = \max_s \left\{ \left(1 - \mathbb{E}[\text{BetaPower}(\mu_0, \rho_0, S_{k-1} \ s)]\right)\pi - \frac{c}{2}s^2 + V_k(S_{k-1} \ s_k) \right\}. \quad (48)$$

We can find V_k numerically. Let $\Omega = [m] \times [m^K]$ for some $m \in \mathbb{N}$ and

$$l(s, S) := \left(1 - \mathbb{E}[\text{BetaPower}(\mu_0, \rho_0, S_{k-1} \ s)]\right)\pi - \frac{c}{2}s^2. \quad (49)$$

Then we can recursively compute

$$V_K(S) = \max_s l(\Omega), \quad (50)$$

$$V_{k-1}(S) = \max_s l(\Omega) + V_K(S \ s), \quad (51)$$

$$\vdots \quad (52)$$

$$V_1(S) = \max_s l(\Omega) + V_2(S \ s). \quad (53)$$