



HAL
open science

Supply Network Formtion and Fragility

Matthew Elliott, Benjamin Golub, Matthew V Leduc

► **To cite this version:**

Matthew Elliott, Benjamin Golub, Matthew V Leduc. Supply Network Formtion and Fragility. 2021.
halshs-03359607

HAL Id: halshs-03359607

<https://shs.hal.science/halshs-03359607>

Preprint submitted on 30 Sep 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



PARIS SCHOOL OF ECONOMICS
ÉCOLE D'ÉCONOMIE DE PARIS

WORKING PAPER N° 2021 – 55

Supply Network Formation and Fragility

Matthew Elliott
Benjamin Golub
Matthew V. Leduc

JEL Codes:

Keywords:



Funded by a French government subsidy managed by the
ANR under the framework of the Investissements d'avenir
programme reference ANR-17-EURE-001

SUPPLY NETWORK FORMATION AND FRAGILITY

MATTHEW ELLIOTT, BENJAMIN GOLUB, AND MATTHEW V. LEDUC

ABSTRACT. We model the production of complex goods in a large supply network. Each firm sources several essential inputs through relationships with other firms. Individual supply relationships are at risk of idiosyncratic failure, which threatens to disrupt production. To protect against this, firms multisource inputs and strategically invest to make relationships stronger, trading off the cost of investment against the benefits of increased robustness. We find that equilibrium aggregate production is robust to idiosyncratic disruptions. Nevertheless, there is a regime in which arbitrarily small systemic shocks cause arbitrarily steep drops in output, so that the supply network is fragile. The endogenous configuration of supply networks provides a new channel for the powerful amplification of shocks.

Date Printed. February 5, 2021.

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement number #757229) (Elliott); the Joint Center for History and Economics; the Pershing Square Fund for Research on the Foundations of Human Behavior; and the National Science Foundation under grant SES-1629446 (Golub). We thank seminar participants at the 2016 Cambridge Workshop on Networks in Trade and Macroeconomics, Tinbergen Institute, Monash Symposium on Social and Economic Networks, Cambridge-INET theory workshop (2018), Binoma Workshop on Economics of Networks (2018), Paris School of Economics, Einaudi Institute for Economics and Finance, Seventh European Meeting on Networks, SAET 2019 and the Harvard Workshop on Networks in Macroeconomics. We thank Joey Feffer, Riako Granzier-Nakajima, Yixi Jiang, and Brit Sharoni for excellent research assistance. For helpful conversations we are grateful to Daron Acemoglu, Nageeb Ali, Pol Antrás, David Baqaee, Vasco Carvalho, Olivier Compte, Krishna Dasaratha, Selman Erol, Marcel Fafchamps, Emmanuel Farhi, Alex Frankel, Sanjeev Goyal, Philippe Jehiel, Matthew O. Jackson, Chad Jones, Annie Liang, Eric Maskin, Marc Melitz, Marzena Rostek, Emma Rothschild, Jesse Shapiro, Ludwig Straub, Alireza Tahbaz-Salehi, Larry Samuelson, Andrzej Skrzypacz, Juuso Välimäki and Rakesh Vohra.

1. INTRODUCTION

Complex supply networks are a central feature of the modern economy. Consider, for instance, a product such as an airplane. It consists of multiple parts, each of which is essential for its production, and many of which are sourced from suppliers. The parts themselves are produced using multiple inputs, and so on.¹ Due to the resulting interdependencies, an idiosyncratic shock can cause cascading failures and disrupt many firms. We develop a theory in which firms insure against supply disruptions by strategically forming supply networks, trading off private gains in the robustness of their production against the cost of maintaining strong supply relationships. Our main results examine how equilibrium supply networks respond to idiosyncratic and aggregate risk. We find that, in equilibrium, (i) the economy is robust to idiosyncratic shocks, yet (ii) small shocks that systemically affect the functioning of supply relationships are massively amplified. Moreover, (iii) the functioning of many unrelated supply chains is highly correlated, and (iv) the complexity of production is key to the nature of these effects and the level of aggregate volatility.

Underlying these results is a discontinuous phase transition in the structure of production networks that arises due to production being complex—reliant on multiple inputs at many stages. Thus, a theoretical contribution of our work is the study of novel equilibrium fragilities in the strategic formation of large networks, and new methods for analyzing them.

Our analysis is built on a model of interfirm sourcing relationships and their disruption by shocks. We now motivate our study of these phenomena with some examples. Firms often rely on particular suppliers to deliver customized inputs. For instance, Rolls-Royce designed and developed its Trent 900 engine for the Airbus A380; Airbus could not just buy the engine it requires off-the-shelf. Such inputs are tailored to meet the customer's specifications, and there are often only a few potential suppliers that a given firm contracts with. Thus, a particular airplane producer is exposed not just to shocks in the overall availability of each needed input, but also to idiosyncratic shocks in the operation of the few particular supply relationships it has formed. Examples of such idiosyncratic shocks include a delay in shipment, a fire at a factory, a misunderstanding by a supplier that delivers an unsuitable component, or a strike by workers. As we have noted, many of a given firm's suppliers themselves will be in a similar position, relying on customized inputs, and so idiosyncratic disruptions to individual relationships and production processes somewhere in a network can have far-reaching effects, causing damage that cascades through the supply chain and affects many downstream firms.² On the other hand, many firms multisource key inputs to reduce their dependency on any one supplier and so idiosyncratic disruptions may not cause much reduction in output at all.

Even the basic properties of such complex supply networks are not well-understood. To articulate this, we describe a simple model capturing key aspects of the above examples. There are many products (e.g., airplanes, engines, etc.). Each has many differentiated varieties, produced by small, specialized firms. A given product has a set of customized inputs that must be sourced via supply

¹For example, an Airbus A380 has millions of parts produced by more than a thousand companies (Slutsken, 2018). In addition to the physical components involved, many steps of production require specific contracts and relationships with logistics firms, business services, etc. to function properly.

²Kremer (1993) is a seminal study of some theoretical aspects of such propagation. Carvalho et al. (2020) empirically study how shocks caused by the Great East Japan Earthquake of 2011 propagated through supply networks to locations far from the initial disruption. In Section 6, we discuss further evidence that such disruptions are practically important, and can be very damaging to particular firms.

relationships and are essential to its production—e.g., an airplane requires engines, navigation systems, etc. To source these specific, compatible, varieties firms may have several substitutable sourcing options (i.e., they multisource). Each such potential supply relationship that a firm has may operate successfully or not: e.g., one engine manufacturer’s delivery may be delayed by a strike, while another is able to deliver normally. In order for a firm to be functional, it must have at least one operating supply relationship to a firm producing each of its essential inputs. To be functional, these producers must in turn satisfy the same condition, and so on—until a point in the supply chain where no customized inputs are required. Our modelling of which supply relationships work is simple: independently, each relationship operates successfully with a probability called the *relationship strength*.³ This probability represents the chance of avoiding logistical disruptions and failures of contracts in any specific relationship. Given a realization of operating relationships, the firms that are able to produce purchase their required inputs and then sell their product to other firms as well as to consumers. Social welfare is increasing in the number of firms able to produce. Firms make profits from production.

Let us, to begin with, take relationship strength to be exogenous and symmetric across the supply network, and examine aggregate output as we vary this relationship strength. The key parameters for describing the mechanics of a supply network are (i) the number of distinct inputs required in each production process (a measure of complexity); (ii) the number of potential suppliers of each input (a measure of the availability of multisourcing); and (iii) the strength of relationships. In our model there is a continuum of firms and the fraction of firms functioning is deterministic.⁴ This fraction is our main outcome of interest, and we call it the *reliability* of the supply network. Our first main result concerns a distinctive form of sensitivity that can arise in such networks—what we call a precipice. Suppose production is complex—that is, that most firms have multiple essential inputs they need to source. We find that there is a discontinuity in reliability as we vary relationship strength, holding all else fixed. When relationship strength falls below a certain threshold (defining the precipice), production drops discontinuously. Thus, if relationship strengths happen to be close to the precipice, a small, systemic, negative shock to relationship strengths is amplified arbitrarily strongly, leading to severe economic damage. We also show that a social planner choosing a level of investments in relationship strengths to maximize welfare would never choose a level such that the supply network is on a precipice.

We give a few practical examples of systemic shocks to relationship strengths. Suppose, first, that the institutions that help uphold contracts and facilitate business transactions suddenly decline in quality, for example due to a political shock. Each supply relationship then becomes more prone to the idiosyncratic disruptions discussed above.⁵ Even if the damage to any single relationship is small (e.g., because usually contracts function without enforcement being relevant), our results show that such a shock can cause widespread disruptions throughout the supply network. For a second concrete interpretation of the kinds of aggregate shocks we have in mind, consider a small

³Though the modeling of basic shocks is independent, the interdependence between firms and their suppliers makes failures correlated between firms that (directly or indirectly) transact with each other.

⁴This is by a standard diversification argument. There are enough firms and supply chains operating that none of them is systemically important. On the complementary issue of when individual firms can be systemically important, see Gabaix (2011) and the large ensuing literature.

⁵Blanchard and Kremer (1997) present evidence that the former Soviet Union suffered a large shock of this kind when it transitioned to a market-based economy.

shock to the availability of credit for businesses in the supply network. The shock matters for firms that are on the margin between getting and not getting credit that is essential for them to deliver on a commitment. The effect of such a credit shock can be modeled as any given supply relationship being slightly less likely to function (depending on ex ante uncertain realizations of whether a firm is on the relevant margin). Third, during the Covid-19 crisis, there has been much uncertainty about how different supply relationships might be affected, and this can be modeled as a systemic decrease in the probability that suppliers are able to deliver the inputs required from them.

The first result just discussed on exogenous relationship strengths shows that, starting at certain strength levels, supply network functioning *can* be very sensitive to slight shocks to strengths. However, relationship strengths are strategic choices, and the endogenous determination of these is in fact our main focus. Our question is whether a supply network will be near a precipice when relationship strengths are determined by equilibrium choices rather than exogenously or by a planner. Since production is risky, an optimizing firm will strategically choose its relationship strength to manage the risk of its production being disrupted. By investing more, a firm can increase the probability that one of its potential suppliers for each of its essential inputs is able to supply it, hence allowing the firm to produce its output and make profits. Maintaining multiple relationships that facilitate production and provide backup in case of disruption is costly. Firms trade off these costs against the benefits of increased robustness.⁶ In our model, firms choose a level of costly investment toward making their supply relationships stronger—i.e., likelier to operate.⁷

Our main findings give conditions for equilibrium relationship strength levels to put the supply network on the precipice, and show that the precipice is *not* a knife-edge outcome. Indeed, we characterize a positive-measure set of parameters (governing the profits of production and the costs of forming relationships) for which the equilibrium supply network is on the precipice. The fragility that a supply network experiences in this regime is highly inefficient: a social planner would never put a supply network on the precipice for the same parameters. As supply networks become large and decentralized one might think that the impact of uncertainty on the probability of successful production would be smoothed by firms' endogenous investments to protect against shocks, and by averaging outcomes across a continuum population. We find the opposite: in equilibrium, there is a very sharp sensitivity of aggregate productivity to relationship strength. This is in contrast to many standard production network models, where the aggregate production function is not too sensitive to small shocks at any outcome. The novelty of our framework comes from the *combination* of two features of production functions that are essential for our results. The first is complexity: In the supply networks we study, firms must source multiple essential inputs that cannot be purchased off-the-shelf. The second is the presence of idiosyncratic disruptions to supply. Jointly, these phenomena create the possibility of precipices, which underlie our analysis.

⁶Strategic responses to risk in networks is a topic that has attracted considerable attention recently. See, for instance, Bimpikis, Candogan, and Ehsani (2019a), Blume, Easley, Kleinberg, Kleinberg, and Tardos (2011), Talamàs and Vohra (2020), and Erol and Vohra (2018), and Amelkin and Vohra (2020). On the practical importance of the strength of contracts in supply relationships, see, among others, Antràs (2005) and Acemoglu, Antràs, and Helpman (2007).

⁷This can be interpreted in two ways: (1) investment on the intensive margin, e.g. to anticipate and counteract risks or improve contracts; (2) on the extensive margin, to find more partners out of a set of potential ones.

More specifically, our main results concern whether the supply network is on a precipice as we vary an aggregate productivity parameter.⁸ Depending on the value of this parameter, the supply network in equilibrium can end up in one of three configurations: (i) a *noncritical equilibrium* where the equilibrium investment is enough to keep relationship strength away from the precipice; (ii) a *critical equilibrium* where equilibrium relationship strength is on the precipice; and (iii) an *unproductive equilibrium* where positive investment cannot be sustained. These regimes are ordered. As the productivity of the supply network decreases from a high to a low level, the regimes occur in the order just given. Each regime occurs for a positive interval of values of the parameter. Equivalently, for an economy consisting of many disjoint supply networks distributed with full support over the parameter space, a positive measure of them will be in the fragile regime, and these will collapse if relationship quality is shocked throughout the economy.

Our analysis makes a conceptual, modeling, and technical contribution to the theory of economic networks. First, we introduce percolation analysis (i.e., disabling some links at random) to an otherwise standard network model of complex production—with *complex* meaning that each firm must source multiple inputs through customized relationships.⁹ That combination leads to the finding that under natural assumptions, random failures in a production network result in a *discontinuous* phase transition, where aggregate functionality abruptly disappears when relationship strengths cross a critical threshold. Second, as a modeling contribution, we demonstrate the tractability of studying equilibrium investments in links (more precisely, investments in the probability that links are operational) in such a setting. By defining a suitable model with a continuous investment choice and a continuum of nodes, investment problems are characterized by relatively tractable first-order conditions, because firms are able to average over the randomness in network realizations. We expect the modeling devices we develop to have other applications. Finally, using our equilibrium conditions to deduce the ordering of regimes discussed above requires developing some new techniques for the analysis of large network formation games.¹⁰ For example, a crucial step in our main results depends on showing that firms' investments in network formation are locally strategic substitutes at the equilibrium investment levels, which depends on subtle properties of equilibrium network structure and incentives that we characterize.

We explore some extensions and implications of our modeling. First, we examine robustness on a number of dimensions. We discuss why the basic insights about fragility extend to alternative specifications of shocks and investment in relationships. Perhaps our most important robustness analysis concerns relaxing certain symmetries of the supply network that we use to simplify our main analysis. We study how fragility manifests in a supply network with rich heterogeneity across multiple dimensions (number of inputs required, amount of multisourcing possibilities, directed multisourcing efforts, profitability, etc.). One important additional implication of our heterogeneity analysis is that a supply network is only as strong as its weakest links: as one product enters the fragile regime, all products that depend on it directly or indirectly are simultaneously pushed into the fragile regime. Second, we interpret our main results in terms of the short- and medium-run resilience of a supply network to shocks, and consider whether the fragility we identify can

⁸This parameter can be interpreted as a measure of productivity relative to the costs of maintaining relationships.

⁹The closest paper is probably Acemoglu and Azar (2020), where each node in a production network relies on one failure-prone custom supplier. In Section 7, we discuss related models on information sharing, financial contagion, and other settings.

¹⁰This differs from and complements the use of graphons in Erol, Parise, and Teytelboym (2020).

be ameliorated by natural policy interventions. Third, we show how the supply networks we have studied can be embedded in a larger economy with intersectoral linkages that do not rely on specific sourcing. Our model yields a new channel for the propagation of shocks across sectors, and their stark amplification. Fourth, while the focus of our analysis is on linking complex supply networks to aggregate volatility, we also discuss how the model can provide a perspective on some stylized facts concerning industrial development (see Section 6.1.1). After presenting our results, in Sections 6 and 7 we discuss in detail how they fit into the most closely related literatures.

2. MODEL: THE SUPPLY NETWORK

Our main object of study is a network whose nodes are a continuum of small firms producing differentiated products. These are connected to their suppliers by a network of potential supply relationships, a random subset of which are realized as operational for sourcing.

An outcome of central interest is the set of firms that is *functional*—i.e., capable of producing. The measure of functional firms will be the key quantity of interest in our model, determining welfare and incentives to invest to form the network (the latter of which we introduce in Section 4). The purpose of this section is to set up the basic structure of the production network and define the set of functional firms.

2.1. Nodes: Varieties of products. There is a finite set \mathcal{I} of *products*. For each product $i \in \mathcal{I}$, there is a continuum \mathcal{V}_i of *varieties* of i , with a typical variety v being an ordered pair $v = (i, f)$, where $f \in \mathcal{F}_i \subseteq \mathbb{R}$ is a variety index; we take $\mathcal{F}_i = [0, 1/|\mathcal{I}|]$ for all i , so that the total mass of firms (and of varieties) in the supply network is 1. Let $\mathcal{V} = \bigcup_{i \in \mathcal{I}} \mathcal{V}_i$ be the union of all the varieties. These are the nodes in our supply network. Each is associated with a small firm producing the corresponding variety.

2.2. Links: Potential and realized supply relationships. First, for each product $i \in \mathcal{I}$ there is a set of required inputs $I(i) \subseteq \mathcal{I}$. Second, each variety $v \in \mathcal{V}$ is associated with a supply chain *depth* $d(v) \in \mathbb{Z}_+$ that specifies how many steps of customized, specifically sourced production are required to produce v . Different varieties of the same product can have different depths. The measure of varieties with any depth $d \geq 0$ is denoted by $\mu(d)$.

Consider any variety $v \in \mathcal{V}_i$. For each $j \in I(i)$ (i.e., each required input) the variety v has a set of *potential suppliers*, $\text{PS}_j(v) \subseteq \mathcal{V}_j$ and a random subset of *realized suppliers* $\text{S}_j(v) \subseteq \text{PS}_j(v)$.

First consider the varieties $v \in \mathcal{V}$ such that $d(v) = 0$. Though these require inputs, specialized sourcing is not required to obtain them and any variety can be used. Thus, in this case, we take $\text{S}_j(v) = \mathcal{V}_j$ for each $j \in I(i)$.

Next, consider any variety $v \in \mathcal{V}_i$ that has depth $d(v) > 0$. For each $j \in I(i)$, the set $\text{PS}_j(v)$ is a finite set of distinct varieties $v' \in \mathcal{V}_j$ with each such v' having depth $d(v') = d(v) - 1$. The identities of these suppliers are independent draws from the set of varieties v' such that $d(v') = d(v) - 1$ (i.e., the set of varieties of compatible depth).¹¹ Specialized sourcing requirements represent the need for a customized input, the procurement of which is facilitated by relational contracts.

Each sourcing relationship between v and a variety $v' \in \text{PS}_j(v)$ is *operational* or not—a binary random outcome. There is a parameter x , called *relationship strength* (for now exogenous and

¹¹For a formal construction of the potential and realized supply networks, see Appendix A.




FIGURE 1. Here we consider a potential supply network for variety $(a, 1)$ with underlying products $\mathcal{I} = \{a, b, c, d, e\}$; the relevant input requirements are apparent from the illustration. Each variety requires two distinct inputs. When these inputs must be specifically sourced, there is an edge from the sourcing variety to its potential supplier. We abbreviate $(a, 1)$ as $a1$ (and similarly for other varieties). Here variety $a1$ has depth $d(a1) = 2$. Varieties higher up are upstream of $a1$. (Thus, orders or sourcing attempts go in the direction of the arrows, and products are delivered in the opposite direction, downstream.)


homogeneous across the supply network) which is the probability that any relationship is operational. All these realizations are independent. The set of actual suppliers $S_j(v)$ is then obtained by including each potential supplier in $PS_j(v)$ independently with probability x_v . For now, the (x_v) are taken to be arbitrary parameters of the model. Whereas the potential supply relationships define compatibilities, the realized supply network identifies which links are actually available for sourcing. The stochastic nature of availability arises, e.g., from uncertainty in delivery of orders, miscommunications about specifications, etc.¹²

We define two random networks on the set \mathcal{V} of nodes. In the *potential supply network* \mathcal{G} , each v has links directed to all its potential suppliers $v' \in \bigcup_{j \in I(i)} PS_j(v)$. (See Figure 1 for an illustration.) In the *realized supply network* \mathcal{G}' , each v has only the operational subset of links, to the realized suppliers: $v' \in \bigcup_{j \in I(i)} S_j(v)$. See the links in Figure 2 for an illustration of the subset of supply relationships that are operational.


2.3. The set of functional varieties. For a given realization of the supply network, we will inductively define which varieties are *functional*, which means that sourcing disruptions do not prevent them from producing.

Depth-0 varieties are always defined to be functional: sourcing failures can never prevent the production of such varieties simply because their sourcing is unconstrained by definition. Given functionalities of varieties of depth $d - 1$, a variety v of depth d is functional if and only if its set of realized suppliers $S_j(v)$ contains at least one functional supplier for each input j that v requires. Figure 2 provides an illustration for the potential supply network shown in Figure 1 and a given realization of operational supply links.

We let \mathcal{V}' denote the (random) set of functional varieties, and \mathcal{V}'_i to be the set of functional varieties in product i .



(A) Step 1: Assigning functionalities for depth-1 varieties.



(B) Step 2: Assigning functionalities of depth-2 varieties.

FIGURE 2. An illustration for determining the set of functional varieties given a realized supply network. Functional varieties are represented by green circles, while non-functional ones are white squares. Varieties that have not yet been assigned to be functional or not are white circles. Varieties of depth 0 are always functional. Panel (A) assigns functionalities to varieties of depth 1. Panel (B) assigns functionalities to varieties of depth 2, making $a1$ nonfunctional.

3. RELIABILITY WITH EXOGENOUS RELATIONSHIP STRENGTHS

We first present our main findings in *regular* supply networks. These are defined by two main symmetries. First, the number of essential inputs is $|I(i)| = m$ for each product i . Second, each variety of depth $d > 0$ has n potential suppliers for each input—i.e. $|\text{PS}_j(v)| = n$ whenever j is one of the required inputs for variety v . Symmetric networks of this type are depicted in Figures 1 and 2. In this section, we posit that relationship strengths are given by $x_v = x$, the same number for each v , and this x is exogenous.¹³

It will be important to characterize how the measure of functional firms \mathcal{V} varies with relationship strength x . Denote by $\tilde{\rho}(x, d)$ the probability that an arbitrary variety of depth d is functional when relationship strength is x . Zero-depth varieties are functional by our assumption that they do not need any specialized inputs: $\tilde{\rho}(x, 0) = 1$. By symmetry of the supply tree, the probability $\rho(x, d)$ does indeed only depend on x and d (and we calculate it explicitly below in Section 3.1.1). The probability that a variety selected uniformly at random is functional is called the *reliability* of the supply network. Since it depends on the distribution of depths μ , we denote this by $\rho(x, \mu)$ and define it as

$$\rho(x, \mu) = \sum_{d=0}^{\infty} \mu(d) \tilde{\rho}(x, d), \quad (1)$$

¹²In a bit more detail, x_v can capture uncertainty regarding compatibility, whether delivery can happen on time, possible misunderstanding about the required input, access to credit that may be needed to deal with unexpected costs, etc. It will depend on the context or environment in which production occurs, and also (as we explicitly model below) on the investments the firm producing v makes. See Section 6 for more details.

¹³We endogenize relationship strengths in the next section, and we relax the symmetry assumption in Section 6.2.6.

To connect this graph-theoretic quantity with economic outcomes, we introduce a relationship between reliability and output. In particular, we posit that there is a function Y such that $Y(\mathcal{V}')$ is the (expected) aggregate gross output associated with a given set \mathcal{V}' of functional firms. We assume that it satisfies the following assumption:

Property A. In a regular supply network, $Y(\mathcal{V}') = h(\rho(x, \mu))$, where $h : [0, 1] \rightarrow \mathbb{R}$ is a strictly increasing, concave function with bounded and continuous derivative. Moreover, $h(0) = 0$.

In Appendix C, we provide a microfoundation for this assumption: we define a standard production network model on top of the supply network, being explicit about production functions and markets at the micro level. There, we show that aggregate output satisfies Property A.

Our main network-theoretic results—e.g., about efficient and equilibrium networks—do not rely on any details of the microfoundations. Instead, they rely only on assumptions about benefits and costs of production that we bring out into named properties. Under Property A, we may already ask: when can output be nonnegligible (i.e., bounded away from zero) even as production networks grow deep? Equivalently, this amounts to asking when reliability is bounded away from zero, which is the question that we turn to next.

Deep supply networks: Taking limits. A focus throughout will be the case where a typical variety has large depth.¹⁴ To this end, we fix a sequence $(\mu_\tau)_{\tau=1}^\infty$ of distributions, where μ_τ places mass at least $1 - \frac{1}{\tau}$ on $[\tau, \infty)$.¹⁵ For instance, we can take μ_τ to be the geometric distribution with mean τ (in which case τ has an exact interpretation as average depth). For large τ , if inputs are single-sourced ($n = 1$) and links fail with positive probability, there will only be a very remote probability of successful production. We therefore restrict attention to the case of multisourcing ($n \geq 2$).

3.1. A discontinuity in reliability. A key implication of the model is the shape of the aggregate production function in equilibrium as we vary x . This shape underlies many of our results. Our first result is:

Proposition 1. Fix any $n \geq 2$ and $m \geq 2$. Then there exist positive numbers $x_{\text{crit}}, \bar{r}_{\text{crit}} > 0$ such that, for all large enough τ ,

- (i) if $x < x_{\text{crit}}$, we have that $\rho(x, \mu_\tau) \rightarrow 0$. That is, reliability converges to 0.
- (ii) if $x > x_{\text{crit}}$, then, for all large enough τ , we have $\rho(x, \mu_\tau) > \bar{r}_{\text{crit}}$. That is, reliability remains bounded away from 0.

In Figure 3(a) we plot the reliability function $\rho(x, \mu_\tau)$ for a fixed finite value of τ against the probability x of each relationship being operational. One can see a sharp transition in relationship strength x at a value that we call x_{crit} . This can be seen more sharply in Figure 3(b), where we plot the limit of the graph shown in (a) as $\tau \rightarrow \infty$. We use the $m = n = 2$ case here as in our illustrations above. The probability of successful production is 0 when $x < x_{\text{crit}}$, but then increases sharply to more than 70% for all $x > x_{\text{crit}}$. Moreover, the derivative of the limit reliability

¹⁴Supplementary Appendix SA4 investigates how reliability varies with investment in production trees with bounded depth.

¹⁵Note that this means varieties of low depth have many incoming edges in the potential supply network, since there are relatively few of them, but they make up a relatively large number of the nodes in a typical production tree.




FIGURE 3. Panel (A) shows how reliability varies with relationship strength x for a particular τ . Panel (B) depicts a correspondence that is the limit of the graphs $\rho(x, \mu_\tau)$ as τ tends to infinity.

graph as we approach x_{crit} from above grows arbitrarily large (i.e., $\lim_{x \downarrow x_{\text{crit}}} \lim_{\tau \rightarrow \infty} \rho'(x, \mu_\tau) = \infty$). This has important ramifications, as we will see. An immediate one is that small improvements in relationship strength x , for example through the improvement of institutions, can have large payoffs for an economy, and the net marginal returns on investment in x can change sharply from being negative to being positive and very large.

3.1.1. *The reasons for the shape of the reliability function.* To explain the logic behind the proposition, let us now calculate the probability that a given variety v with depth d is functional. Recall that we denote by $\tilde{\rho}(x, d)$ the probability that a variety of depth d is functional when relationship strength is x . We will argue that this can be expressed recursively as follows. First, $\tilde{\rho}(x, 0) = 1$, since varieties of depth 0 are sure to be functional. Then, for a suitably defined function $\mathcal{R}_x : [0, 1] \rightarrow [0, 1]$, we can write the depth- d reliability in terms of the depth- $(d - 1)$ reliability:

$$\tilde{\rho}(x, d) = \mathcal{R}_x(\tilde{\rho}(x; d - 1)). \quad (2)$$

Indeed, more explicitly, the function that makes this true is¹⁶

$$\mathcal{R}_x(r) = (1 - (1 - xr)^n)^m.$$

As we look at networks with large typical depths, it is the $\tilde{\rho}(x, d)$ for large d that will matter. Proposition 5 in Appendix B.1 shows that there is a unique correspondence $\rho(x)$ that is the limit of the graphs of $\rho(\cdot, \mu_\tau)$ in a suitable sense as $\tau \rightarrow \infty$. By analyzing this correspondence we can show that a sharp transition like that shown in Figure 3(B) occurs for any complexity $m \geq 2$ and any multisourcing level $n \geq 2$, once depths become large.

The intuition for the sharp transition can be understood by looking at Figure 4(A). Here we plot the probability that a given firm is functional against the probability that its suppliers are

¹⁶Consider the first input of a focal variety. For a given supplier of that input, by definition its reliability is the argument r , and the probability that the link to the supplier is operational is x . The probability that both events happen is xr . The probability that this combination of events happens for at least one of the n potential suppliers of the first input is therefore $1 - (1 - xr)^n$. Finally, the probability that for all m inputs, such a “good event” happens is $(1 - (1 - xr)^n)^m$.




FIGURE 4. Panel (A) shows the probability, $\mathcal{R}_x(r)$, that a focal firm is functional as a function of r , the probability that a random supplier is functional. Here we use the parameters $n = 4$, $m = 2$ and $x \in \{0.5, 0.55, 0.6, 0.65\}$. The intersections with the 45 degree line marked by the red circles represent reliability values (for deep supply trees) associated with the given exogenous parameters. Panel (B) shows, for a given value of x (namely, $x = 0.55$), the first four values of $\tilde{\rho}(x, d)$ given by equation (2). Note that the intersection is approached quite closely after a few steps.

functional. As supply networks become deep, consistency requires these numbers to be equal (since a firm and its suppliers occupy essentially equivalent positions) and the reliability levels that are fixed points of the function $\mathcal{R}_x(r)$ in equation (2) are given by intersections with the 45-degree line. For high enough x , the graph of \mathcal{R}_x intersects the 45-degree line above 0.

Repeatedly applying formula (2) starting with $\tilde{\rho}(x, 0) = 1$ yields a sequence of $\tilde{\rho}(x, d)$ converging to the largest fixed point of $\mathcal{R}_x(r)$, as shown in Figure 4(B). The reliability associated with that intersection is bounded away from 0 precisely when there is a positive fixed point of $\mathcal{R}_x(r)$. When x is below a certain critical value, the graph of \mathcal{R}_x has no nontrivial intersection with the 45-degree line and so $\tilde{\rho}(x, d)$ converges to 0. As we explain in Section 6.1.1 where we contrast complex production ($m > 1$) with simple production ($m = 1$), the convex-then-concave shape of the \mathcal{R}_x curve is essential for creating a precipice.

3.1.2. Comparative statics of the reliability function. Some straightforward comparative statics can be deduced from what we have said. If n (multisourcing) increases while all other parameters are held fixed, then one can check that \mathcal{R}_x (as illustrated in Figure 4(A)) increases pointwise on $(0, 1)$, and this implies that all the $\tilde{\rho}(x, d)$ increase. It follows that the ρ curve moves upward, and the discontinuity occurs at a lower value of x .

Similarly, when m (complexity) increases, the \mathcal{R}_x curve decreases pointwise, implying that all the $\tilde{\rho}(x, d)$ decrease. It follows that the ρ curve moves downward, and the discontinuity occurs at a higher value of x .

4. SUPPLY NETWORKS WITH ENDOGENOUS RELATIONSHIP STRENGTH

In this section, we study the endogenous determination of relationship strength. We first study a planner's problem—choosing an optimal value of strength for all relationships. We then turn to a decentralized problem in which firms invest in their own relationship strengths. Throughout, we focus on symmetric outcomes.

Our main results show that while investments that put the supply network on the precipice are very inefficient, they need not be knife-edge or even unlikely outcomes in equilibrium.

4.1. A planner's problem. We study a planner who chooses a global x that determines the values of all relationship strengths, $x_v = x$. This can be interpreted as investing in the quality of institutions, at a cost that we will introduce below. As stipulated in Property A, the gross output of the supply network is an increasing function of reliability, $Y(\mathcal{V}') = h(\rho(x, \mu))$, where reliability ρ depends on the symmetric level of relationship strengths x and depth distribution μ . The planner's cost of a given choice of x enters through subtracting a quantity $\frac{1}{\kappa}c_P(x)$ of output, where κ is a strictly positive parameter and $c_P : [0, 1] \rightarrow \mathbb{R}$ is a fixed function. Higher values of κ have the interpretation that they shift down the costs of obtaining a given level of relationship strengths, and hence a given level of productivity. The planner seeks to maximize expected net aggregate output by choosing relationship strengths x , and hence solves the **planner's problem**

$$\max_{x \in [0, 1]} h(\rho(x, \mu)) - \frac{1}{\kappa}c_P(x). \quad (3)$$

We make the following assumption concerning $c_P(x)$:

Property B. The function $c_P : [0, 1] \rightarrow \mathbb{R}$ is continuously differentiable and weakly convex, with $c_P(0) = 0$, $c_P(x_{\text{crit}}) > 0$, $c'_P(0) = 0$, and $\lim_{x \rightarrow 1} c'_P(x) = \infty$.

Substantively, these entail that there are weakly increasing marginal social costs of relationship strength and achieving the critical level of reliability requires a positive investment. The Inada conditions on derivatives ensure good behavior of optima.

Define the correspondence

$$x^{SP}(\kappa, \mu) := \operatorname{argmax}_{x \in [0, 1]} h(\rho(x, \mu)) - \frac{1}{\kappa}c_P(x).$$

This gives the values of x that solve the social planner's problem for a given κ and distribution of supply chain lengths μ . As elsewhere, we consider a sequence $(\mu_\tau)_{\tau=1}^\infty$ of depth distributions, where μ_τ places mass at least $1 - \frac{1}{\tau}$ on $[\tau, \infty)$.

Proposition 2. Fix any $n \geq 2$ and $m \geq 2$. Then there exists a number $\kappa_{\text{crit}} > 0$ such that, for any small enough $\epsilon > 0$, the following statements hold for all large enough τ :

- (i) for all $\kappa < \kappa_{\text{crit}}$, all values of $x^{SP}(\kappa, \mu_\tau)$ are at most $x_{\text{crit}} - \epsilon$, have cost less than ϵ , and yield reliability less than ϵ ;
- (ii) for all $\kappa > \kappa_{\text{crit}}$, all values of $x^{SP}(\kappa, \mu_\tau)$ are at least $x_{\text{crit}} + \epsilon$ and yield reliability at least $r_{\text{crit}} + \epsilon$;
- (iii) for $\kappa = \kappa_{\text{crit}}$, all values of $x^{SP}(\kappa, \mu_\tau)$ are outside the interval $[x_{\text{crit}} - \epsilon, x_{\text{crit}} + \epsilon]$ and reliability is outside the interval $[\epsilon, r_{\text{crit}} + \epsilon]$.

The first part of Proposition 2 says that when κ is sufficiently low, it is too costly for the social planner to invest anything in the quality of institutions, and hence reliability is very low. As κ increases, a threshold κ_{crit} is reached and at this value of κ it first becomes optimal to invest in institutional quality. At this threshold, the social planner's investment increases discontinuously. Moreover, it immediately increases to a level strictly *above* x_{crit} , and for all larger κ all solutions stay above x_{crit} .

It is worth emphasizing that the planner never chooses to invest near the critical level x_{crit} . The reason is as follows: sufficiently close to $x = x_{\text{crit}}$, the marginal social benefits of investing grow

arbitrarily large in the limit as τ gets large while marginal costs at x_{crit} are bounded, and so the social planner can always do better by increasing investment at least a little. In contrast, in Section 5 we will see that individual investment choices can put the supply network on the precipice in equilibrium, and this is not a knife-edge scenario.

4.2. Decentralized investment in relationship strengths.

4.2.1. *Setup and timing.* Now we formulate a simple, symmetric, model of decentralized choices of relationship strengths. The decision-makers in this richer model are firms. In each product i , there is a continuum of separate firms (i, f) , where $f \in [0, \frac{1}{|I|}]$. The firm (i, f) owns the corresponding variety, $v = (i, f)$.¹⁷ We often abbreviate a firm by if .

Firms simultaneously choose *investment levels* $y_{if} \geq 0$. Choosing a level y_{if} has a private cost $c(y_{if})$. The random realization of the supply network occurs after the firm chooses its investment level.¹⁸ If a firm chooses an investment level y_{if} , then all sourcing links from its variety (i, f) have relationship strength

$$x_{if} = \underline{x} + y_{if}.$$

The intercept $\underline{x} \geq 0$ is a baseline probability of relationship operation that occurs absent any costly investment. This can be interpreted as a measure of the quality of institutions—e.g., how likely a “basic contract” is to deliver.¹⁹ The main purpose of this baseline level is as a simple channel to systemically shock relationship strengths throughout the supply network.

The timing is as follows.

1. Firms simultaneously choose their investment levels.
2. The realized supply network is drawn and payoffs are enjoyed.²⁰

An *outcome* is given by relationship strengths x_{if} for all firms if . An outcome is symmetric if all firms have the same relationship strength: $x_{if} = x$ for all if .

4.2.2. *Payoffs and equilibrium.* A firm’s payoff at an outcome can be written as

$$u_{if} = G_{if} - \frac{1}{\kappa} c(x_{if} - \underline{x}).$$

This is the firm’s expected gross profit, G_{if} , minus the cost of its investment, $y_{if} = x_{if} - \underline{x}$, in relationship strength. We will now discuss the parts of this payoff function in turn.

We begin with the firm’s costs. We assume that they satisfy the following property, where $\kappa > 0$ is a parameter:

Property B’. A firm’s cost is given by $\frac{1}{\kappa} c(x_{if} - \underline{x})$, where the following conditions hold:

- (i) $\underline{x} < x_{\text{crit}}$;
- (ii) c' is increasing, continuously differentiable, and strictly convex, with $c(0) = 0$;

¹⁷Thus, our notation identifies a firm with its variety.

¹⁸The supply network realization is defined as an assignment of depths to all varieties, and the graphs \mathcal{G} and \mathcal{G}' from Section 2. The assumption that investments are made before this realization is technically convenient, as it keeps the solution of the model symmetric. For example, a firm knows that after some number of stages of production, disruption-prone contracts will not be needed by its indirect suppliers (e.g., because these suppliers are able to use generic inputs or rely on inventories). However, the firm does not know how many steps this will take.

¹⁹A natural interpretation is that this is a feature of the contracting environment—concretely, for instance, it could reflect the quality of the commercial courts.

²⁰Our results in Section 3 imply that payoffs at this stage are uniquely determined.

(iii) the Inada conditions hold: $\lim_{y \downarrow 0} c'(y) = 0$ and $\lim_{y \uparrow 1-\underline{x}} c'(y) = \infty$.

The first part of this assumption ensures that baseline relationship strength is not so high that the supply network is guaranteed to be productive even without any investment. The second part imposes assumptions on investment costs that ensure agents' optimization problems are well-behaved. The Inada conditions, as usual, ensure that investments are interior. Here κ plays the same role as it did in our social planner optimization exercise: scaling down the costs of investing in relationship strength—and hence achieving a given level of productivity.

We now turn to specifying gross profits at a given outcome. Because we will characterize symmetric equilibria, we need to specify the gross profits of a given firm only for symmetric behavior by other firms. Firms make no gross profits conditional on not producing, and their profits conditional on producing satisfy the following assumption.

Property C. At a symmetric outcome with reliability r , conditional on being functional, a firm makes gross profits $g(r)$, where $g : [0, 1] \rightarrow \mathbb{R}_+$ is a decreasing, continuously differentiable function.

Property C requires that profits are higher when fewer firms are functioning and there is less competition. Appendix C microfounds this property in the same production network model that we used to microfound Property A.

Let $P(x_{if}; x, \mu)$ be the probability with which a firm if is functional if firm if 's relationship strength is x_{if} and all other firms choose symmetric relationship strengths x inducing reliability r . Then, under Property C, we have that $G_{if} = P(x_{if}; x, \mu)g(r)$. Thus, recalling the payoff formula given at the start of this section, the net expected profit of firm if when it has relationship strength x_{if} and all other firms have relationships strengths x , resulting in reliability r , is

$$\Pi_{if} = P(x_{if}; x, \mu)g(r) - \frac{1}{\kappa}c(x_{if} - \underline{x}). \quad (4)$$

Finally, because welfare properties of equilibria will play a role in our analysis, we define social welfare (for symmetric outcomes). The gross output of production given reliability r has a value of $h(r)$, as in the previous section on the planner's problem. (Some of this goes to gross profits and some to consumer welfare, but the sum is given by $h(r)$.) The planner's cost function is simply the total of firms' costs,

$$\sum_{i \in \mathcal{I}} \int_{f \in \mathcal{F}_i} \frac{1}{\kappa} c(x - \underline{x}) df = \frac{1}{\kappa} c(x - \underline{x}),$$

where we have used our assumption that the mass of firms in each industry is $1/|\mathcal{I}|$, so that the total mass of firms is 1. Thus, if x is the relationship strength of all firms, the social welfare function is²¹

$$\text{Welfare} = h(\rho(x, \mu)) - \frac{1}{\kappa} c(x - \underline{x}). \quad (5)$$

5. EQUILIBRIUM SUPPLY NETWORKS AND THEIR FRAGILITY

We now study the equilibrium of our model: its productivity and its robustness. This section builds up to a main result: Theorem 1. We show that in the limit as production networks become deep, there are three regimes. First, for low values of the parameter κ , there is an unproductive

²¹This coincides with the welfare function of the previous section if we set $c_P(x) = c(x - \underline{x})$ for $x \geq \underline{x}$, and $c_P(x) = 0$ otherwise.

regime in which equilibrium reliability is arbitrarily low. Next, for intermediate values of κ , there is a critical regime in which equilibrium relationship strengths are very close to x_{crit} and arbitrarily small shocks to relationship strength lead to discontinuous drops in production. Finally, there is a noncritical regime in which equilibrium relationship strengths are above x_{crit} and the supply network is robust to small shocks.

It is worth writing (4) more explicitly. To do this, we calculate $P(x_{if}; x, \mu)$, the probability with which if is able to produce, as a function of firm if 's relationship strength x_{if} given that all other firms choose a symmetric relationship strength x :²²

$$P(x_{if}; x, \mu) := \mu(0) + \mathbb{E} [1 - (1 - x_{if}\tilde{\rho}(x, d - 1))^n]^m,$$

where d inside the expectation is drawn from the depth distribution μ conditional on depth being at least 1. Recall from equation (2) in Section 3.1.1 the formula for $\tilde{\rho}$.

Definition 1. We say $x \geq \underline{x}$ is a *symmetric equilibrium* if $x_{if} = x$ maximizes $\Pi_{if}(x_{if}; x, \mu)$ defined in (4) for any firm if . It is a *symmetric undominated equilibrium* if it is a symmetric equilibrium and, among symmetric equilibria, maximizes the social welfare defined in (5).

When we refer to an equilibrium in the sequel, we mean a symmetric undominated equilibrium unless otherwise noted. Note that a symmetric equilibrium is defined by the level of relationship strength $x = \underline{x} + y$ realized in it, rather than the level of investment y . This turns out to be more convenient.

Our equilibrium definition requires that all firms' investment choices are equal and are mutual best responses to each other. Lemma 5 in the appendix shows that the efficiency condition selects the symmetric equilibrium associated with the highest investment level, and hence highest reliability. Note that it is always a best response for a firm to choose zero investment when all others choose zero investment. Our equilibrium definition abstracts from potential miscoordination on the zero investment level, or other inefficient ones, by selecting the symmetric equilibrium that maximizes welfare.

In the limit, as the expected depth of the supply networks becomes large, if firms symmetrically choose investments $y_{if} = 0$ then the reliability is $\rho(\underline{x}) = 0$ as $\underline{x} < x_{\text{crit}}$ (by Property B'). Hence, for large enough τ , $x_{if} = \underline{x}$ maximizes $\Pi_{if}(x_{if}; \underline{x}, \mu_\tau)$ and so there always exists an equilibrium.

In analyzing the symmetric equilibria it is helpful to make an assumption on the environment that ensures that the first-order conditions of firms' problems are sufficient for optimality among interior solutions. We first state the assumption and then formulate a condition on primitives that is sufficient for it to hold.

Assumption 1. For any τ and any $x > x_{\text{crit}}$ the function $x_{if} \mapsto \Pi_{if}(x_{if}; x, \mu_\tau)$ has a unique interior local maximum for all large enough τ .²³

Assumption 1 will be maintained in the sequel, along with Properties A, B', and C. The following lemma shows that we may always set \underline{x} so that Assumption 1 is satisfied.

²²Note that because there is a continuum of firms the probability that a firm appears in its potential supply network upstream of itself is 0. Thus the reliability of if 's suppliers does not depend on x_{if} .

²³The assumption permits another local maximum at a corner. We rule this out separately.

Lemma 1. For any $m \geq 2$, and any $n \geq 2$, there is a number²⁴ \hat{x} , depending only on m and n , such that, for large enough τ we have: (i) $\hat{x} < x_{\text{crit}}$; and (ii) if $x \geq \hat{x}$, then Assumption 1 is satisfied.

To see why this lemma implies that Assumption 1 is satisfied for a suitable choice of \underline{x} , consider any environment where $x \in [\hat{x}, x_{\text{crit}})$. Part (i) of the lemma guarantees that the interval $[\hat{x}, x_{\text{crit}})$ is nonempty, and part (ii) guarantees that Assumption 1 is satisfied for values of x in this range.²⁵ In other words, the lemma guarantees that there is an interval of possible baseline levels of robustness which are short of the critical level (so that fragility is not ruled out a priori) but high enough to ensure that the firms' maximization problem is amenable to a first-order approach.

We now characterize the equilibrium behavior.

Theorem 1. Fix any $n \geq 2$ and $m \geq 3$.²⁶ There are thresholds $\underline{\kappa} < \bar{\kappa}$, which depend only on n and m , such that the following holds. For τ exceeding some threshold $\underline{\tau}$, there is a unique symmetric undominated equilibrium with relationship strength $x_{\tau}^*(\kappa)$. Moreover, for any $\epsilon > 0$, the threshold $\underline{\tau}$ can be chosen so that the equilibrium satisfies the following properties.

- (i) If $\kappa < \underline{\kappa}$, there is no investment: $x_{\tau}^*(\kappa) = \underline{x}$.
- (ii) For $\kappa \in [\underline{\kappa}, \bar{\kappa}]$, the equilibrium relationship strength $x_{\tau}^*(\kappa) \in [x_{\text{crit}} - \epsilon, x_{\text{crit}} + \epsilon]$. We call such equilibria *critical*.
- (iii) For $\kappa > \bar{\kappa}$, the equilibrium relationship strength satisfies $x_{\tau}^*(\kappa) > x_{\text{crit}}$. We call such equilibria *non-critical*.

Moreover, for $\tau \geq \underline{\tau}$, the function $x_{\tau}^*(\kappa)$ is increasing on the domain $\kappa > \underline{\kappa}$.

If we think of different supply networks as being parameterized by different values of κ in a compact set, Theorem 1 implies that in the limit as τ gets large, there will be a positive fraction of supply networks in which firms will choose relationship strengths converging to x_{crit} in equilibrium. This contrasts with the social planner's solution, which never selects relationship strengths near x_{crit} . It also means that a positive fraction of supply networks end up perched on the precipice, vulnerable to shocks.

Figure 5 helps give some intuition for Theorem 1. In any symmetric equilibrium, the reliability of each firm must be consistent with the symmetric investment level chosen by all other firms—we must be somewhere on the reliability curve we derived in Section 2. The thick, green graphs in panels (A)–(D) of Figure 5 illustrate the shape of the reliability curve for large τ . Further, in any symmetric equilibrium each firm's (symmetric) investment choice of x must be a best response to the reliability of its suppliers. The thinner, red curves in panels (A)–(D) depict the best-response function; these curves should be thought of as having their independent variable (others' reliability) on the vertical axis, and the best-response investment on the horizontal axis. The panels show the best-response curves for increasing values of κ . Intersections of these two curves are potential symmetric equilibria. Our equilibrium definition implies that when there are multiple intersections, we select the one associated with the highest reliability. Thus equilibrium reliability is 0 for κ sufficiently small, jumps up discontinuously to $\underline{r}_{\text{crit}}$ at $\underline{\kappa}$, and is increasing in κ thereafter.

²⁴In the proof, we give an explicit description of \hat{x} in terms of the shape of the function $P(x_{if}; x, \mu_{\tau})$.

²⁵While conditions we identify in Lemma 1 are sufficient for satisfying Assumption 1, they are not necessary.

²⁶In this result, we restrict attention to the case of $m \geq 3$. It is essential for our results that supply networks are complex ($m \geq 2$), but the case of $m = 2$ generates some technical difficulties for our proof technique so we consider $m \geq 3$. In numerical exercises, our conclusions seem to also hold for $m = 2$.

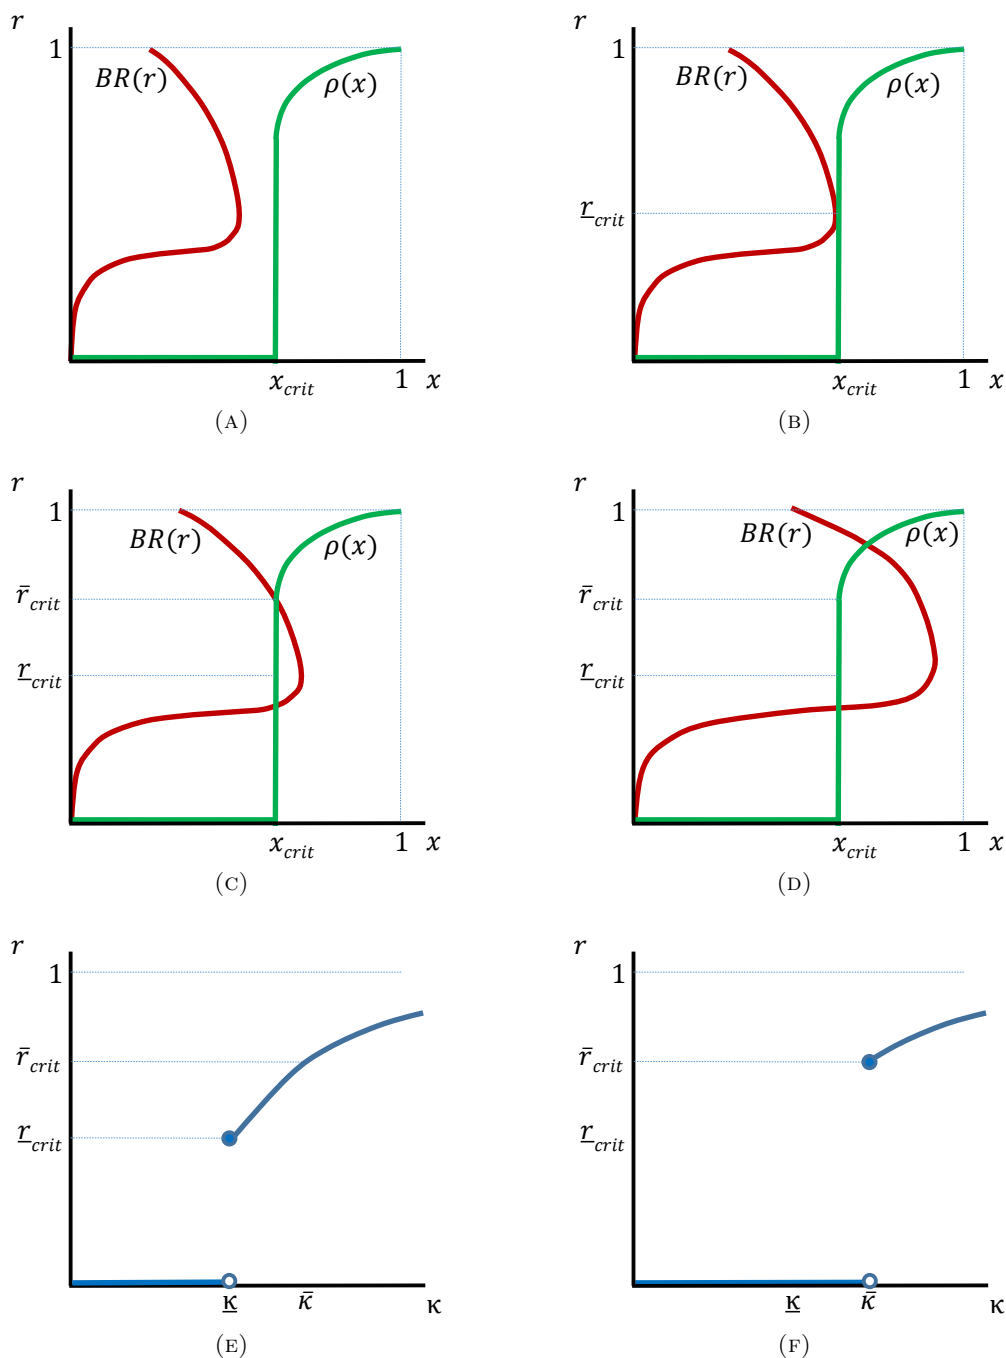


FIGURE 5. Panel (A) shows an equilibrium for $\kappa < \underline{\kappa}$. Panel (B) shows an equilibrium with $\kappa = \underline{\kappa}$. Panel (C) shows an equilibrium with $\kappa = \bar{\kappa}$. Panel (D) shows an equilibrium with $\kappa > \bar{\kappa}$. Panel (E) plots how equilibrium reliability varies with κ . Panel (F) shows reliability following an arbitrarily small negative shock to institutional quality \underline{x} as κ varies.

This is shown in panel (E). Note that although reliability increases as κ ranges over the interval $[\underline{\kappa}, \bar{\kappa}]$, because the limit reliability curve is a correspondence with $\rho(x_{crit}) = [0, \bar{r}_{crit}]$, equilibrium investment for large τ remains arbitrarily close to x_{crit} in this critical regime. In other words, equilibrium investment choices bunch around x_{crit} for an open interval of values of κ . For all $\kappa \in [\underline{\kappa}, \bar{\kappa}]$ a slight shock causing relationship strengths to diminish from \underline{x} to $\underline{x} - \epsilon$ causes relationship

strengths to fall below x_{crit} , and makes equilibrium production collapse. Panel (F) shows reliability after such a shock for different values of κ . As can be seen by comparing panels (E) and (F), there is virtually no difference in reliability for either $\kappa < \underline{\kappa}$ or $\kappa > \bar{\kappa}$. However, for $\kappa \in [\underline{\kappa}, \bar{\kappa}]$ reliability drops discontinuously to 0.

We now pause to comment on the techniques used to establish these results. While the intuition looks straightforward in our sketches, proving Theorem 1 is technically challenging. We have to get a handle on the shape of the best-response correspondence, and show that the highest intersection of it with the reliability curve moves as depicted in Figure 5. Investment incentives (which determine the shape of the best-response curve, and the implications of changing κ) are complex, and so equilibrium investment is difficult to characterize directly. We do not have an explicit expression for the best-response curve. It also does not have any global monotonicity properties that might permit standard comparative statics approaches. Investments in relationship strength are strategic complements in some regions of the parameter space, and strategic substitutes in others.²⁷ However, some control on the slopes of the intersecting curves is necessary to analyze the dependence of the intersection on parameters. A key step in our analysis is showing that if the highest intersection of best response and reliability curves is at $r \geq r_{\text{crit}}$, then the best response curve must be negatively sloped in r there, and strictly so for an intersection at $r > r_{\text{crit}}$. We do so by establishing certain properties of the system of polynomial equations that (asymptotically) characterizes the intersection. Another crucial step is to show that there is at most one point of intersection above \bar{r}_{crit} in our diagrams. This permits us to focus on a unique outcome of interest, and to sign the effects of moving κ unambiguously, which is important for the results on equilibrium fragility.

At a technical level, there are two parts to many of our proofs. One part focuses on an idealized “limit supply network,” which, in a suitable sense, has infinite depth and therefore a lot of symmetry (a firm’s suppliers look exactly like the firm itself). This symmetric network gives rise to expressions that we can manipulate to establish the key facts mentioned above. There is then a separate task of transforming these limit results into statements about the supply network with large but finite depths.

Our next result, Corollary 1, implies that the comparative statics of equilibrium as the baseline quality of institutions \underline{x} changes are analogous to those documented with respect to κ in Theorem 1. Here we explicitly include \underline{x} as an argument in x^* .

Corollary 1. Suppose $\kappa' > \kappa$. Then, for large enough τ , if $x_\tau^*(\kappa', \underline{x}) > x_\tau^*(\kappa, \underline{x})$, there is an $\underline{x}' > \underline{x}$ such that $x_\tau^*(\kappa', \underline{x}) = x_\tau^*(\kappa, \underline{x}')$.

5.1. Fragility. Critical equilibria are important because, as the example in Figure 5 shows, they create the possibility of fragility: small shocks to relationship strengths via a reduction in \underline{x} can result in a collapse of production. We formalize this idea by explicitly examining how the supply network responds to a shock to the baseline quality of institutions \underline{x} , which for simplicity is taken

²⁷When a firm’s suppliers are very unreliable, there is little incentive to invest in stronger relationships with them—there is no point in having a working supply relationship when the suppliers cannot produce their goods. On the other hand, when a firm’s suppliers are extremely reliable, a firm can free-ride on this reliability and invest relatively little in its relationships, knowing that as long as it has one working relationship for each input it requires, it is very likely to be able to source that input.

to have zero probability. (The analysis is robust to anticipated shocks that happen with positive probability—see Section 6.2.3.)

Definition 2 (Equilibrium fragility).

- There is *equilibrium fragility* at κ if for any $\eta, \epsilon > 0$, for large enough τ , we have

$$\rho(x_\tau^*(\kappa) - \epsilon, \mu_\tau) < \eta.$$

That is, a shock that reduces relationship strengths arbitrarily little (ϵ) from their equilibrium levels leads to a reliability very close to 0 (within η).

- We say there is *equilibrium robustness* at κ if there is not equilibrium fragility.

In the definition of fragility, while shocking \underline{x} , we hold investment decisions and entry choices fixed. Implicitly, we are assuming that investments in supply relationships and entry decisions are made over a sufficiently long time frame that firms cannot change the quality of their supply relationships or their entry decisions in response to a shock.

Proposition 3. Under the conditions of Theorem 1, there is equilibrium fragility at any $\kappa \in [\underline{\kappa}, \bar{\kappa}]$ and equilibrium robustness at any $\kappa > \bar{\kappa}$.

Proposition 3 follows immediately from the definition of equilibrium fragility and our previous characterization.

There are many ways in which the small systemic shock discussed in this section might arise in practice. Recall the example in the introduction. A small shock to credit markets makes many supply relationships slightly more likely to fail—in the states of the world where they happen to require financing to deal with a disruption. Another, similar, example is uncertainty about which relationships will be affected by possible compliance issues related to new trade regulations, thereby making supply relationships more prone to disruption. For a third example, consider an increasing backlog in commercial courts—a circumstance that makes contracts more costly to enforce. This again decreases the probability that contracts function in some states of the world—an uncertainty that can affect many players in the supply chain. For a final example, suppose that an epidemic occurs, increasing the costs of in-person interactions. This will likely affect some supply relationships, but not others; uncertainty around this can be captured by an increase in the probability that a given supply relationship is disrupted.

6. DISCUSSION OF THE MODEL

This section discusses our modeling choices. First, in Section 6.1 we discuss which features of our environment are essential, and motivate them. Beyond the key assumptions, we make a variety of assumptions for tractability or simplicity. In Section 6.2, we discuss the robustness of our model to various relaxations of these non-essential assumptions.

6.1. Essential features. The three essential features that we highlight are that supply relationships are specific (firm-to-firm) and subject to disruption; that firms endogenously invest to strengthen their relationships; and that production is complex in that it relies on multiple essential inputs.

Specific sourcing relationships that are subject to idiosyncratic disruptions are a crucial feature of our model. Supplier relationships have been found to play important roles in many settings—for

relationship lending between banks and firms see Petersen and Rajan (1994, 1995); for traders in Madagascar see Fafchamps and Minten (1999); for the New York apparel market see Uzzi (1997); for food supply chains see Murdoch, Marsden, and Banks (2000); for the diamond industry see Bernstein (1992); for Japanese electronics manufacturers see Nishiguchi (1994)—and so on. Indeed, even in fish markets, a setting where we might expect relationships to play a minor role, they seem to be important (Kirman and Vriend, 2000; Graddy, 2006). The importance of specific sourcing relationships in supply networks is also a major concern of the management literature on supply chains (Datta, 2017). Barrot and Sauvagnat (2016) find that firms have difficulty switching suppliers even when they need to do so. These literatures examine many reasons behind firms' reliance on a small number of suppliers to source a given type of input. For example, technological compatibility and geography can limit the pool of potential suppliers; hold-up problems can make trust important; and frequent repeated interactions can help firms to avoid misunderstandings.

The specific relationships that firms maintain for sourcing facilitate the contagion of disruption. Indeed, cascading disruptions (see, for example, Carvalho et al. (2020)) are evidence of firms' reliance on failure-prone sources of supply that cannot be quickly replaced. Interesting qualitative descriptions of cascades of disruption due to idiosyncratic shocks can be found in the business literature. A fire at a Philips Semiconductor plant in March 2000 halted production, preventing Ericsson from sourcing critical inputs, causing its production to also stop (*The Economist*, 2006). Ericsson is estimated to have lost hundreds of millions of dollars in sales as a result, and it subsequently exited the mobile phone business (Norrman and Jansson, 2004). In another example, two strikes at General Motors parts plants in 1998 led 100 other parts plants, and then 26 assembly plants, to shut down, reducing GM's earnings by \$2.83 billion (Snyder et al., 2016). Though these particular examples are particularly well-documented, disruptions are a more frequent occurrence than might be expected. In a survey of studies on this subject in operations and management, Snyder et al. (2016) write, "It is tempting to think of supply chain disruptions as rare events. However, although a given type of disruption (earthquake, fire, strike) may occur very infrequently, the large number of possible disruption causes, coupled with the vast scale of modern supply chains, makes the likelihood that some disruption will strike a given supply chain in a given year quite high." An industry study recorded 1,069 supply chain disruption events globally during a six-month period in 2018 (Supply Chain Quarterly, 2018).

Given the frequency of disruptions and the impact these can have on firms' profitability,²⁸ it is natural that firms take actions to mitigate them. In practice, these investments are often "soft" in nature. An ethnographic literature in sociology helps document them—a prominent one being Uzzi (1997), who offers a detailed account of the systematic efforts and investments made by New York garment manufacturers and their suppliers to maintain good relations. These investments include practices such as building a better understanding of a supplier's or customer's capabilities by visiting their facilities, querying odd instructions to help catch mistakes, building social relationships that span the organizations, and reciprocal gift-giving. Such investments are hard to observe and even

²⁸Hendricks and Singhal (2003, 2005a,b) examine hundreds of supply chain problems reported in the business press. Even minor disruptions are associated with significant and long-lasting declines in sales growth and stock returns.

harder to verify. This makes them hard to contract on and fits the way we model investments—as a decentralized game of private investment in one’s local relationships.²⁹

It is crucial for our theoretical results that production is complex insofar as firms need multiple inputs via specific sourcing relationships. Barrot and Sauvagnat (2016) provide evidence that supports this assumption. They show that if a supplier is hit by a natural disaster it severely disrupts the production of their customers *and* also negatively impacts their customers’ other suppliers. If production were not complex, then these other suppliers would be providing substitute inputs and hence would benefit from the disruption to a competitor rather than being adversely affected.

6.1.1. *Contrasts with benchmark models without the essential features.* In this subsection, we briefly consider two benchmarks to highlight the necessity of key features discussed in the previous section.

Contrast I: Market-based sourcing. In this benchmark, each firm sources all its inputs through markets, rather than requiring pre-established relationships. The market is populated by those potential suppliers that are able to successfully produce the required input. However, upon approaching a supplier there is still a chance that sourcing fails. (A shipment might be lost or defective, or a misunderstanding could lead the wrong part to be supplied.) In other words, we now assume each firm extends relationships *only* to functional suppliers (as opposed to suppliers whose functionality is random, in our main model); but we still keep the randomness in whether the sourcing relationships work.

We now work out which firms can produce, focusing on the example where each firm requires two inputs ($m = 2$). Each firm *if* multisources by contracting with two potential suppliers of each input ($n = 2$)—selected from the functional ones. Let the probability a given attempt at sourcing an input succeeds be x , independently. The probability that both potential suppliers of a given input type fail to provide the required input is $(1 - x)^2$, and the probability that at least one succeeds is $1 - (1 - x)^2$. As the firm needs access to all its required inputs to be able to produce, and it requires two different input types, the probability the firm is able to produce is $(1 - (1 - x)^2)^2$. In Figure 6 we plot how the probability that a given firm is able to produce varies with the probability their individual sourcing attempts are successful. This probability increases smoothly as x increases.

This benchmark shows that perfect spot markets remove the discontinuities in our main analysis. On the other hand, our findings are robust to the existence of imperfect substitutes for specifically sourced inputs—see Section 6.2.4.

Contrast II: Sourcing for simple production. To emphasize that it is essential that multiple inputs are sourced through relationships, we consider a benchmark model where each firm requires only a single relationship-sourced input ($m = 1, n = 2$). We call such production simple because

²⁹The key feature here is the modeling of investment as a non-cooperative game. Even if firms were allowed to invest in others’ relationships they would typically not want to in equilibrium. Following the logic of Bergstrom, Blume, and Varian (1986) each firm will be willing to invest in each relationship (between any two firms) up to the point that the marginal return of the investment equals the marginal benefits. However, the marginal benefits are heterogeneous and it will typically only be possible for this condition to be satisfied by the firm that has the highest marginal benefit of investing.

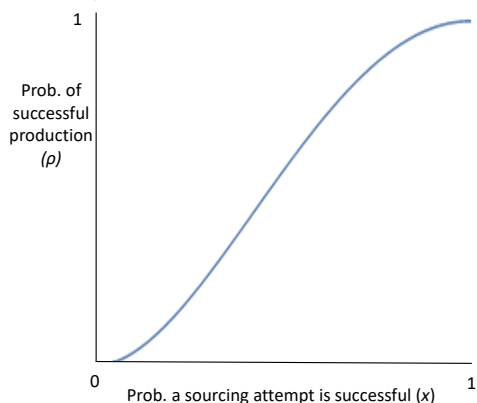


FIGURE 6. The probability of successful production for a firm as market-based sourcing attempts become more likely to succeed.




FIGURE 7. (A) The probability of successful production of a simple good ($m = 1$ distinct inputs needed) as relationship strength varies. (B) The probability, $\mathcal{R}_x(r)$, that a focal firm is functional as a function of r , the probability that a random supplier is functional. This plot is for $n = 3$ with $x \in \{0.3, 0.35, 0.4, 0.5, 0.65\}$. The plot parallels Figure 4(A), which depicts the same function for the complex production ($m > 1$) case. The marked intersections with the 45 degree line reflect limit reliability outcomes.

each firm requires only *one* type of risky input relationship to work.³⁰ We plot how the probability of successful production varies with relationship strength in Figure 7(a). In comparison to the case of complex production illustrated in Figure 3(b), there is a stark difference. For values of $x < 0.5$ the probability of successful production is 0 and for values of $x > 0.5$ the probability of successful production is strictly positive. The change at $x = 0.5$ involves the derivative changing discontinuously, but productivity itself is continuous.

The intuition is familiar from the networks literature and in particular from studies of contagion (see, for example, Elliott, Golub, and Jackson (2014) in the context of financial contagion, or Sadler (2020) for information diffusion). In the large-depth limit, production will be successful if

³⁰As a matter of interpretation, there may be more than one physical input at each stage. The key assumption is that all but one are sourced as commodities rather than through relationships, and so are not subject to disruption via shocks to these relationships.

the supply tree of functional producers upstream of a given firm continues indefinitely, rather than being extinguished due to an excess of failures. This depends on whether the rate at which new branches in the network are created is higher or lower than the rate at which existing branches die out due to failure. It turns out that when $x > 0.5$, a supply tree grows without bound in expectation, while when $x < 0.5$ it dies out.³¹ The kink in the probability of successful production around the threshold of 0.5 is related to the emergence of a giant component in an Erdős–Rényi random graph. That continuous phase transition is different from the discontinuous one driving our main results. Indeed, to see the importance of the difference in these phase transitions for equilibrium fragility, consider a productive equilibrium when production is simple. It can immediately be seen from Figure 7(a) that a small shock to reliability x will only cause a small change in the probability of successful production. In contrast to the complex production case, the reliability curve does not become arbitrarily steep as we increase the depth of the supply network.

6.2. Robustness: Extensions relaxing simplifying assumptions. In this section we discuss directions in which the analysis can be extended to show that certain simplifying assumptions are not essential to the main findings.

6.2.1. Investments on the extensive margin. Our modeling of investments in supply relationships is compatible with multiple interpretations. The first interpretation is that the set of possible suppliers is fixed, and the investment works on the intensive margin to improve the quality of these relationships (e.g., by reducing misunderstandings and so on). The second interpretation is that the investment works on the extensive margin—i.e., firms work to find a supplier willing and able to supply a given required input type, but their success is stochastic. In this interpretation there is a fixed set of n potential suppliers capable of supplying the required input to be found, and each one of them is found independently with probability x_{if} .³² Conditional on a supplier being found, the relationship is operational. In Appendix SA6 we discuss a richer extensive-margin interpretation, and also one that permits separate efforts to be directed to the extensive and intensive margins simultaneously.

6.2.2. Endogenous entry. Our model takes the varieties in the market as given, and does not include entry decisions by firms. This is done to keep the model simple. However, the key insights carry over in a model with entry.

Consider an extended model where there is an entry stage preceding all others. At this stage, a firm (i, f) pays a sunk entry cost $\Phi(f)$, where Φ is a strictly increasing function. Then the measure of varieties in each product is set equal to the measure of entering firms (i.e., each entering firm produces one variety), and the “investment game” where firms choose their relationship strengths proceeds as we described earlier. A (nontrivial) symmetric equilibrium is now one where $\bar{f} > 0$

³¹Given that each producer has two potential suppliers for the input, and each of these branches is operational with probability x , the expected number of successful relationships a given firm in this supply network has is $2x$. When $x < 1/2$, each firm links to on average less than 1 supplier, and so the rate at which branches in the supply tree fail is faster than the rate at which new branches are created. The probability that a path in the supply network reaches beyond a given tier l then goes to 0 as l gets large and production fails with probability 1. On the other hand, when $x > 1/2$, the average number of suppliers each firm has an operating relationship with is greater than 1 and so new branches appear in the supply tree at a faster rate than they die out, leading production to be successful with strictly positive probability.

³²This search technology is similar to ball-and-urn models of search, and is compatible with a matching function exhibiting constant returns to scale (see, for example, Hall (1979)).

firms enter in each product and the expected post-entry profits of the marginal firm are equal to its entry cost.

The main observation is that this model also features an open set of productivity parameters κ where equilibrium production is on the precipice. The basic logic is as follows. When the supply network is reliable and gross profits after entry are high, firms want to enter. As they enter, competition drives down gross profits and makes it less appealing to pay costs to make relationships strong. Recalling Figure 5, the key to the argument is that increasing entry moves the best-response curve leftward for a given κ . So, in equilibrium, relationships get weaker as more entry occurs. The question is where this dynamic stops. The precipice is a natural stopping point. Once the investment level is x_{crit} , reliability can adjust down until further entry is deterred, while investment remains the same. By reasoning parallel to that of our main results, this state is reached for an open set of model parameters.

6.2.3. *Anticipated shocks.* We assume that the small shock to x is unanticipated. It is straightforward to use our analysis to see that this is not essential. Suppose now all firms anticipate that a shock to x will happen with some probability p_{shock} , and suppose κ is in an open subset of the critical productivities, so that a collapse would have occurred in the baseline model. When we write the profit function for firms in the extended model, there will now be an expectation over the shock's arrival, and this will change profits and the best-response correspondence. However, for a small enough probability of the shock, this change will be small, and the best-response curve in a plot such as Figure 5(c) will move only slightly. Thus the equilibrium will still be on the precipice if it was before. Indeed, for κ not too close to the boundary of the critical productivities, the probability of the shock can be quite substantial without changing the main prediction. The key intuition is that because of externalities, while all firms may prefer a commitment to invest more to avoid being on the precipice (and avoid disruptions due to the shock), the free-riding problem might be too severe to achieve this. Indeed, note that in Figure 5, even fairly large shifts in the best-response curve can sometimes leave the equilibrium on the precipice, implying that these free-riding effects are considerable.

6.2.4. *Partial disruptions.* We assume zero-one failures: a firm can produce as long as it has an operational link to at least one functioning producer of each input. More realistically, it may be that to produce, a sufficient quantity or quality of each input is needed, and the shock is to whether the requisite level is reached. The shock need not destroy all the output, but may destroy or reduce the value of the output by some amount (say, a random fraction of the output). We consider the starkest case for simplicity; the key force is robust to these sorts of extensions.

6.2.5. *Shocks to firms.* The fundamental source of shocks in our model is at the level of links in the supply network. One could also consider shocks directly hitting firms—i.e., shocks to nodes in our supply network rather than links. The key conclusions regarding precipices are robust to adding this source of shocks, or even making node-level shocks the main source of disruptions.³³ We focus on shocks to relationships for simplicity, but the mathematical forces we identify also operate in alternative models of disruption.

³³The reliability calculations, available on request, end up being very similar, and the shape of the reliability curve has the same features as in our main model.

6.2.6. *Heterogeneities.* One might suspect that the regularity of the network structure, or some other kind of homogeneity, plays an important role in generating the sharp transition in the probability of successful production. In contrast with the key roles played by complexity and specific sourcing, homogeneities across the products are not important to our main points.

To establish this, we consider a heterogeneous analogue of our basic model with firms investing in the strength of their relationships with their suppliers. In this model, (1) the set of inputs required to produce a product i , $I(i)$, is an arbitrary set with some cardinality m_i . Thus, the number of inputs required can be different for different products (although each variety of a given product uses the same production technology). (2) For each product i and input $j \in I(i)$, there is a number n_{ij} of potential suppliers of product j that each firm has; thus n_{ij} replaces the single multisourcing parameter n . (3) Investment and link strength are input-specific: For a firm if and an input $j \in I(i)$, there is a relationship strength $x_{if,j}$ which replaces the single number x_{if} . The cost of effort is $c_{ij}(x_{if,j} - \underline{x}_{ij})$. (4) The gross profit conditional on producing product i is $G_i(r_i)$. Here G_i is a product-specific function (which can capture many different features of different product markets that affect their profitabilities) and r_i is the reliability of producers of product i . Here we work with a reduced-form specification of gross profits, though it can be microfounded paralleling our main model.

Analogously to our main model, we can introduce depths of each variety and specify the matching process so that each variety sources from varieties of smaller depth. And then, again analogously, we are interested in the limit as depths tend to infinity. As we saw in B.1 in the homogeneous case, there is a limit reliability function which corresponds to an economy with infinite depths. This reliability function governs the structure of precipices. For brevity, here we will focus directly on this limit, with the understanding that results for the high-depth supply network converge to this limit supply network as in our other analysis.

For this section, we use $x = (x_1, x_2, \dots, x_{|\mathcal{I}|})$ to describe investment profiles for the different products. Note that $x_i = (x_{i,j})_{j \in I(i)}$ is a vector describing the heterogeneous strengths chosen by a producer of product i for its relationships with firms supplying different inputs $j \in I(i)$, whereas elsewhere in this article it is a scalar. As before we identify the profile with the relationship strength achieved. We also let $x_{-i} = (x_1, x_2, \dots, x_{i-1}, x_{i+1}, \dots, x_{|\mathcal{I}|})$ denote the investment profiles of producers of products other than product i . The product i is *critical* at an investment profile x if, fixing the relationship strengths of the producers of other products x_{-i} , the probability of successful production of product i would be zero at any profile $\hat{x}_i < x_i$ (where $\hat{x}_i < x_i$ means that each entry of \hat{x}_i is weakly lower than the corresponding entry of x_i and at least one such entry is strictly lower).

The sharp transition in the mechanics of our basic model, discussed in Section 3.1, persists in the presence of heterogeneous networks where different products require different numbers of inputs, there are different numbers of suppliers for each input and relationship strengths vary by input. Indeed, paralleling the analysis of the homogeneous case, we show in Section SA5.1 of the Supplementary Appendix, that the reliability of the supply network as a function of link strength exhibits a sharp transition, even in the presence of such heterogeneities (Proposition SA1).

We next show that supply networks in the presence of heterogeneities feature a *weakest link* property (Proposition 4): When one product is critical, all products that rely directly or indirectly

on it are critical as well.³⁴ This result uses the notion of a product interdependencies graph, whose nodes are set \mathcal{I} of products, and products have directed links to the input products they require (see Figure 8 for an illustration).

Proposition 4 (Weakest link property).

- (i) Suppose product i is critical. Then any other product j on a directed path from j to i in the product interdependencies graph will fail following a shock that changes \underline{x}_{ik} to $\underline{x}_{ik} - \epsilon$ for any $k \in I(i)$.
- (ii) Let $\mathcal{I}^{SC} \subset \mathcal{I}$ be a set of products that are part of a strongly connected component of the product interdependencies graph. Then any equilibrium with positive effort is such that either all producers of products $i \in \mathcal{I}^{SC}$ are critical or no producers of products $i \in \mathcal{I}^{SC}$ are critical.




FIGURE 8. An example of a product interdependency graph. An arrow from a to b means that b is an input required in the production of a , i.e. $b \in I(a)$.

Proposition 4 asserts that supply networks suffer from a weakest link phenomenon. First, if a product is critical, then a shock to it causes the production of other products that use it as an input, directly or indirectly, to also fail. Second, if we take a strongly connected component of products where none of them are critical and, say, decrease gross profits for the producers of one product until it becomes critical, then all products in that component will also become critical at the same time. The component is only as strong as its weakest links.

In Sections SA5.3 and SA5.4 of the Supplementary Appendix, we illustrate both the discontinuities and the weakest link property with examples. We also show there that the configurations we have described are consistent with endogenous investment in relationship strength.

7. RELATED LITERATURE

We have already discussed many papers that are relevant for motivating our assumptions or interpreting our results in Section 6. In this section we review several high-level connections to related literatures not covered by our previous discussions.

³⁴Thus, as in Bimpikis, Ehsani, and Ilkılıç (2019b) the social planner has different gains from intervening in different parts of the network (see their Proposition 8).

There has been considerable recent interest in markets with non-anonymous trade mediated by relationships.³⁵ The work most closely related to ours in this area also studies network formation in the presence of shocks. This includes work in the context of production (e.g., Fafchamps (2002), Levine (2012), Brummitt, Huremović, Pin, Bonds, and Vega-Redondo (2017), Bimpikis, Candogan, and Ehsani (2019a), Yang, Scoglio, and Gruenbacher (2019), Amelkin and Vohra (2019)), work on financial networks (e.g., Cabrales, Gottardi, and Vega-Redondo (2017), Elliott, Hazell, and Georg (2018), Erol (2018), Erol and Vohra (2018), Jackson and Pernoud (2019)), and other contexts (e.g., Blume, Easley, Kleinberg, Kleinberg, and Tardos (2011) Jackson, Rodriguez-Barraquer, and Tan (2012), Talamàs and Vohra (2020)). More broadly, the aggregate implications of relying on relationships to transact have been studied across a variety of settings. For work on thin financial markets see, for example, Rostek and Weretka (2015), for buyer-seller networks see, e.g., Kranton and Minehart (2001), and for intermediation see, e.g., Gale and Kariv (2009). Our work focuses on network formation for *production* and emphasizes the distinctive network formation concerns that arise due to strong complementarities. At a methodological level, we offer an approach that may be useful more broadly. Agents make a continuous choice that determines the probability of their relationships operating successfully. The links that form may, however, fail in a “discrete” (i.e., non-marginal) way. The first feature makes the model tractable, while the second one yields discontinuities in the aggregate production function and distinguishes the predictions from models where the aggregate production function is differentiable. It might be thought that aggregating over many supply chains, these discontinuities would be smoothed out at the level of the macroeconomy; we show they are not.

There is a vibrant literature in macroeconomics on production networks. This literature dates back to investigations of the input-output structure of economies and the implications of this (Leontief, 1936). Carvalho and Tahbaz-Salehi (2019) provides a comprehensive survey.

Two recent developments in the literature are particularly relevant to our work: (i) the modeling of the endogenous determination of the input-output structure; and (ii) a firm-level approach as opposed to considering inter-industry linkages at a more aggregated level. Some of the most relevant work on these issues includes Atalay, Hortacsu, Roberts, and Syverson (2011), Oberfield (2018), Carvalho and Voigtländer (2014), Acemoglu and Azar (2020), Boehm and Oberfield (2020), Tintelnot, Kikkawa, Mogstad, and Dhyne (2018), and König, Levchenko, Rogers, and Zilibotti (2019). Baqaee and Farhi (2019) and Baqaee and Farhi (2020) focus specifically on the implications of nonlinearities, and discuss how nonlinearities in firms’ production functions propagate and aggregate up. Whereas they focus on smooth nonlinearities, we show that especially extreme nonlinearities—discontinuities—naturally come from complex supply networks. We have discussed throughout how our explicit modeling of sourcing failures at the micro level gives rise to new effects (see Section 6.1.1).

The strong complementarities in production in our model are crucial for creating fragility, and this aspect of our work builds on a large literature following the seminal work of Kremer (1993). This literature argues that complementarities can help provide a unified account of many economic phenomena. These include very large cross-country differences in production technology and aggregate productivity; rapid output increases during periods of industrialization; and the structure

³⁵A literature in sociology emphasizes the importance of business relationships, see for example Granovetter (1973) and Granovetter (1985). For a survey of related work in economics see Goyal (2017).

of production networks and international trade flows; see, among many others, Ciccone (2002), Acemoglu, Antràs, and Helpman (2007), Levchenko (2007), Jones (2011) and Levine (2012).³⁶ A possible concern with this literature is that if firms were allowed to take actions that mitigate supply risks (for example, by multisourcing), that this would endogenously dampen the crucial complementarities. Our work helps address this concern. We show that even when firms can take actions to mitigate their risk by multisourcing, a very severe form of equilibrium fragility—and hence production complementarities—arises.

At a technical level, our work is related to a recent applied mathematics literature on so-called *multilayer networks* and their phase transitions (Buldyrev, Parshani, Paul, Stanley, and Havlin, 2010). Discontinuities such as the one we study are termed first-order phase transitions in this literature.³⁷ Buldyrev, Parshani, Paul, Stanley, and Havlin (2010), and subsequent papers in this area such as Tang, Jing, He, and Stanley (2016) and Yang, Scoglio, and Gruenbacher (2019), study quite different network processes, typically with exogenous networks. We show that discontinuities of this kind arise in canonical models of equilibria in production networks, once specific sourcing relationships are taken into account. In addition, we endogenize investments in the probabilities of disruption (which are taken to be exogenous in this literature) and elucidate a new economic force endogenously putting equilibria on a precipice. Predating the recent literature on multilayer networks, Scheinkman and Woodford (1994) used insights from physics models on self-organized criticality to provide a “sandpile” model of the macroeconomy in which idiosyncratic shocks have large aggregate effects.³⁸ The setup and behavior of the model are rather different from ours: the main point of commonality is in the concern with endogenous fragility. In our model, investments in supply relationships leave the supply network robust to idiosyncratic shocks, but very sensitive to arbitrarily small systemic shocks to relationship strength.

8. CONCLUDING DISCUSSIONS

We conclude by sketching some implications of the modeling and analysis. Though we have focused on descriptive analysis of fragility, we first note some basic welfare implications of the model. Next, we discuss some of the short-run and medium-run implications of our model and suggest how our forces can be embedded in larger models of a macroeconomy. Finally, we briefly mention some potential applications of the precipice phenomena in analyzing industrial development as institutions improve.

8.1. Welfare implications. In Section 4.1 we showed that a planner will never choose investments that result in the supply network being fragile. On the other hand, in Section 5 we found that decentralized investment choices often result in fragile supply networks. In this section we first discuss the externalities that create this wedge between decentralized outcomes and the planner’s solution. We then show that there is systemic underinvestment in reliability at fragile equilibria. Finally, we briefly comment on possible policy remedies.

³⁶Prior to his literature, Jovanovic (1987) examines how strategic interdependencies or complementarities can produce aggregate volatility in endogenous variables despite only seemingly “diversifiable” idiosyncratic volatility in exogenous variables.

³⁷These can be contrasted with second-order phase transitions such as the emergence of a giant component in a communication network, which have been more familiar in economics—see Jackson (2008).

³⁸Endogenously, inventories reach a state analogous to a sandpile with a critical slope, where any additional shock (grain dropped on the sandpile) has a positive probability of leading to an avalanche.

There are three externalities present when a firm chooses its reliability. First, there is the non-appropriability of consumer surplus. Output is higher when more firms function, and this is not fully internalized in a firm's reliability choice even with respect to adding *its own* variety to the mix, since the firm appropriates only part of the surplus. Second, there are reliability spillovers: firms don't internalize their contributions to helping *other* varieties function. The reliability of a firm's intermediate good production increases the reliability, and hence profitability, of those firms sourcing from it, the firms sourcing from these firms, and so on. Both the non-appropriability of consumer surplus and the reliability spillover lead to underinvestment in reliability (all else equal). On the other hand, there is also a business-stealing effect. Producers of a good make fewer sales when more producers of the same good are functional. This is a force for overinvestment, all else equal, as firms jockey to capture a larger share of the market.³⁹

How do these externalities compare, and what are the implications for policies to increase welfare? First, consider a critical equilibrium. We know from Section 4.1 that reliability investments are never efficient in a fragile equilibrium. The marginal improvement in reliability from a uniform increase in investments is arbitrarily large—indeed, this is how we concluded that the planner never chooses reliability that leads to a fragile supply network. So, in a critical equilibrium, the non-appropriability of consumer surplus and reliability spillovers together dominate the business-stealing effect, and there is underinvestment in relationships relative to an efficient benchmark. Indeed, Corollary 2 in the appendix shows that there is underinvestment in all symmetric equilibria with positive investment levels.

Underinvestment leading to a fragile equilibrium can be counteracted in various ways, although fragility remains a robust phenomenon. Suppose, for example, effort is subsidized to increase robustness so that the cost of a given investment is $1 - \theta$ times its original cost. More precisely, we make the gross profit equal to

$$\Pi_{if} = \underbrace{P(x_{if}; x)}_{\text{prob. functional}} \underbrace{g(r)}_{\substack{\text{gross profit} \\ \text{if functional}}} - \underbrace{\left(\frac{1 - \theta}{\kappa}\right) c(y_{if})}_{\text{cost of effort}}.$$

Suppose the status quo were at a critical equilibrium with reliability $r = \underline{r}$ (see Figure 5(B) for an illustration). A subsidy of this form will shift the red best response curve in Figure 5(B) to the right. At the margin this will increase reliability, but have a very limited effect on equilibrium investments for high values of τ . Thus the equilibrium will remain fragile. The same argument can be used for increasing θ . Indeed, a similar analysis applies for many other interventions that shift the best response curve to the right—including, for example, improvements in institutions.

8.2. The short run and medium run. We define the short run to be period of time over which firms are unable to adjust their relationship strengths. Evidence from Barrot and Sauvagnat (2016) suggests that this is on the order of magnitude of several quarters.⁴⁰ We define the medium run as a period of time over which adjustments in relationship strength are possible.⁴¹ Thus, following

³⁹Higher investments in reliability increase the probability a firm is able to function and compete, and so can be interpreted as an investment in (stochastic) entry. In light of this it is unsurprising that the usual non-appropriability of consumer surplus and business-stealing effects are present.

⁴⁰See their Figure IV.

⁴¹In this discussion we hold fixed the production technology, particularly its complexity. These are dimensions that might change in the long run. For a brief discussion of changes on this time horizon see Section 8.4 below.

a shock to relationship strengths, in the short run a collapse in production will occur when the supply network is in the fragile regime—that is, holding relationship strengths at their old values. In the medium run (i.e., the time scale on which investments are adjustable), it is possible that firms can choose investment levels that allow the supply network to resume being productive.⁴²

In the medium run, it is also interesting to consider how the supply network (via equilibrium investment) reacts to improvements in institutional quality or productivity. By Theorem 1 and Corollary 1, such changes will increase the strength of supply relationships and equilibrium reliability. In the fragile regime, the impact of any improvement in κ is almost entirely incident on equilibrium reliability, as opposed to equilibrium relationship strength. (That is, we move *up* along the precipice, and not at all to the right.) Relatedly, in the fragile regime, reductions in institutional quality (lower \underline{x}) induce an almost one-for-one increase in investment, so that relationship strength is close to constant, although reliability may still fall substantially.

There are interesting equilibrium differences in both reliability and fragility between supply networks of different complexity (number of inputs m) and different levels of multisourcing (n), which can be seen as a measure of density. Given the same relationship strength and the same multisourcing possibilities, more complex supply networks will be less reliable. Thus while one supply network might be in the non-fragile regime, a more complex supply network can be in the fragile regime for the same parameter values.⁴³ Also, when there are fewer multisourcing opportunities, so that production networks are sparser, then, all else equal, production will be less reliable (see Section 3.1.2).⁴⁴

8.3. Aggregate volatility. So far we have focused on a single complex supply network with particular parameters, interconnected by specific-sourcing relationships. The larger economy can be thought of as consisting of many such supply networks, each one small relative to the economy. Our question in this section is what the fragility of some of these particular supply networks implies for the reaction of the entire economy to shocks. Our main finding is that while the severe amplification of shocks remains, the conclusion is nuanced by embedding individual supply networks in a richer setting.

Suppose there are many supply networks operating independently of each other, with heterogeneity across supply networks but, for simplicity, homogeneity within each network. The parameters of these different supply networks, including their complexities m and multisourcing numbers n are drawn from a distribution. We know from the above that a small shock to relationship strength can discontinuously reduce the production of some of these supply networks. We now point out that a small shock can have a large macroeconomic effect, and that the structure of the fragile regime is essential for this.

⁴²Our equilibrium selection here is optimistic: it focuses on the outcome where, in recovering from the shock, firms coordinate on the most productive investment equilibrium, and thus limit the losses. We could, however, work with less optimistic assumptions. It is consistent once production has become disrupted for all firms to stop investing altogether. If such coordination problems occur, disruptions can be longer-lasting and consistent with long-lasting productivity damage following a shock.

⁴³Similarly, in more complex production networks firms need to source more inputs and so are more susceptible to weak institutions undermining the strength of their supply relationships. This is congruent with findings in Boehm (2020).

⁴⁴Herskovic (2018) argues, in an asset pricing framework, that this should imply a premium on the equity values of firms embedded in sparser production networks because more of the risks they face will be idiosyncratic, and finds evidence of this.

For simplicity, fix the function $c(\cdot)$ to be the same throughout the economy.⁴⁵ A given supply network is then described by a tuple $\mathfrak{s} = (m, n, \kappa)$. We consider the space of these networks induced by letting the parameters m, n and κ vary. In particular, we let \mathcal{M} be the set of possible values of m , the set of integers between 1 and M ; we let \mathcal{N} be the set of possible values of n , integers between 1 and N , and we allow $\kappa \in \mathcal{K} = [0, K]$. The space of possible supply networks is now $\mathcal{S} = \mathcal{M} \times \mathcal{N} \times \mathcal{K}$. We let Ψ be a distribution over this space, and assume that it has full support.

In some supply networks there will not exist an equilibrium with positive production, which we henceforth call a productive equilibrium (for example, when κ is sufficiently low fixing the other parameters). Consider now those supply networks for which there is a productive equilibrium. There are two possibilities. It may be that the only supply networks for which there is a positive equilibrium have $m = 1$.⁴⁶ That is, the only supply networks with positive reliability are simple. In this case, there is no aggregate fragility.

But if, in contrast, \mathcal{S} contains supply networks where production is *not* simple, then we will have macroeconomic fragility. Indeed, an immediate consequence of Theorem 1 is that if there are some complex ($m \geq 2$) supply networks with positive equilibria, then some of the lower- κ networks with the same (n, m) —which are included in \mathcal{S} —are in the fragile regime. The measure that Ψ assigns to supply networks in the fragile regime is positive. Thus, a shock to relationship strengths will cause a discontinuous drop in expected aggregate output. On the other hand, in contrast to the situation with a single supply network, the reduction in output need not be all the way down to zero: it will be a discontinuous loss of some fraction of the output.

So far we have looked at the case in which the different supply networks operate independently and all business-to-business transactions occur through supply relationships confined to their respective supply networks. As we have emphasized, such specific relationships correspond to inputs tailored to the specifications of the business purchasing them. They are not products that can be purchased off-the-shelf. However, many other inputs are sourced in different ways. For example, most business use computers, and buy them off-the-shelf rather than through the specific-sourcing relationships we have focused on. So far we have abstracted from any interdependencies between businesses created by such arm's-length purchases. However, these interdependencies might matter. If a small aggregate shock causes the collapse of some supply networks, the inputs available to other supply networks that managed to remain functional become scarcer and more costly. This effectively damages the productivity of these other supply networks, and when they reoptimize, some of them that were not previously on a precipice will now find themselves there. Thus, they will be sensitive to further aggregate shocks. The key is that being on the precipice is not a fixed attribute of a supply network's structure, but in fact dependent on its productivity. Thus, even with market-mediated spillovers, the productivity damage of collapses leads to domino effects where, iteratively, previously robust parts of the economy become fragile.⁴⁷

8.4. Another application: Some simple implications for industrial development. The comparison between the production of complex and simple products developed in Section 6.1.1 has interesting implications for the complexity of technologies used across countries, and for industrial

⁴⁵This could also be drawn from a distribution, but the notation would be more cumbersome.

⁴⁶For example, it might be that $M = 1$ so that only simple production is feasible, or K might be sufficiently low that only simple production has a positive probability of being successful.

⁴⁷In Appendix SA3 we provide a simulated illustration of this phenomenon.




FIGURE 9. A contrast of the $m = 1$ and $m = 2$ cases with the degree of multisourcing being held at $n = 2$. Expected revenues are on the vertical axis. This is a product of the probability of successful production and price of goods. The case in which the complex good retails for a price of 1 while the simple good retails for a price of $1/4$ is illustrated.

development. While a full analysis of this is beyond the scope of this paper, Figure 9 illustrates the main point, bringing together the the reliability curve for complex products (Figure 3) with that of simple products (Figure 6.1.1(a)). The curves are scaled to reflect that complex products are more valuable, so that the value of a certain level of reliability is scaled up. Small increases in quality of commercial institutions can make more complex production technologies viable and yield discontinuous benefits to an industry if it enables the industry to transition from simple to complex production.

Even this very rudimentary theory of development and industrialization fits a number of stylized facts: (i) Industrialization, when it occurs, is rapid and economic output increases dramatically. (ii) At the same time, the share of the value of total production that can be attributed to intermediate inputs increases quickly (Chenery, Robinson, and Syrquin, 1986). (iii) The quality of institutions, and particularly those related to contracting, can help explain what kinds of production different economies can support (Nunn, 2007) and hence cross-country differences in development (Acemoglu and Johnson, 2005), wages and productivities (Jones, 2011). (iv) Finally, more complex supply chains are associated with higher rates of disruption (Craighead, Blackhurst, Rungtusanatham, and Handfield, 2007).⁴⁸

A full study of these issues obviously requires more detailed modeling. Nevertheless the discontinuities we identify may have a useful role to play in theories relating the quality of institutions and an economy's production possibilities.

REFERENCES

- ACEMOGLU, D., P. ANTRÀS, AND E. HELPMAN (2007): "Contracts and Technology Adoption," *American Economic Review*, 97, 916–943.
- ACEMOGLU, D. AND P. D. AZAR (2020): "Endogenous Production Networks," *Econometrica*, 88, 33–82.
- ACEMOGLU, D. AND S. JOHNSON (2005): "Unbundling Institutions," *Journal of Political Economy*, 113, 949–995.

⁴⁸The unscaled reliability curve for complex production lies below the reliability curve of simple production. So complex production is more valuable but less reliable.

- AMELKIN, V. AND R. VOHRA (2019): “Strategic Formation and Reliability of Supply Chain Networks,” *arXiv preprint arXiv:1909.08021*.
- (2020): “Strategic Formation and Reliability of Supply Chain Networks,” in *Proceedings of the 21st ACM Conference on Economics and Computation*, 77–78.
- ANTRÀS, P. (2005): “Incomplete Contracts and the Product Cycle,” *American Economic Review*, 95, 1054–1073.
- ATALAY, E., A. HORTACSU, J. ROBERTS, AND C. SYVERSON (2011): “Network Structure of Production,” *Proceedings of the National Academy of Sciences*, 108, 5199–5202.
- BAQAEE, D. R. AND E. FARHI (2019): “The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten’s Theorem,” *Econometrica*, 87, 1155–1203.
- (2020): “Productivity and misallocation in general equilibrium,” *Quarterly Journal of Economics*, 135, 105–163.
- BARROT, J.-N. AND J. SAUVAGNAT (2016): “Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks,” *The Quarterly Journal of Economics*, 131, 1543–1592.
- BERGSTROM, T., L. BLUME, AND H. VARIAN (1986): “On the Private Provision of Public Goods,” *Journal of Public Economics*, 29, 25–49.
- BERNSTEIN, L. (1992): “Opting Out of the Legal System: Extralegal Contractual Relations in the Diamond Industry,” *The Journal of Legal Studies*, 21, 115–157.
- BIMPIKIS, K., O. CANDOGAN, AND S. EHSANI (2019a): “Supply Disruptions and Optimal Network Structures,” *Management Science*, 65, 5504–5517.
- BIMPIKIS, K., S. EHSANI, AND R. ILKILIÇ (2019b): “Cournot Competition in Networked Markets,” *Management Science*, 65, 2467–2481.
- BLANCHARD, O. AND M. KREMER (1997): “Disorganization,” *The Quarterly Journal of Economics*, 112, 1091–1126.
- BLUME, L., D. EASLEY, J. KLEINBERG, R. KLEINBERG, AND É. TARDOS (2011): “Network Formation in the Presence of Contagious Risk,” in *Proceedings of the 12th ACM Conference on Electronic Commerce*, ACM, 1–10.
- BOEHM, C. E., A. FLAAEN, AND N. PANDALAI-NAYAR (2019): “Input Linkages and the Transmission of Shocks: Firm-Level Evidence from the 2011 Tōhoku Earthquake,” *Review of Economics and Statistics*, 101, 60–75.
- BOEHM, J. (2020): “The Impact of Contract Enforcement Costs on Outsourcing and Aggregate Productivity,” *Review of Economics and Statistics*, forthcoming.
- BOEHM, J. AND E. OBERFIELD (2020): “Misallocation in the Market for Inputs: Enforcement and the Organization of Production,” *Quarterly Journal of Economics*, forthcoming.
- BRUMMITT, C. D., K. HUREMOVIĆ, P. PIN, M. H. BONDS, AND F. VEGA-REDONDO (2017): “Contagious Disruptions and Complexity Traps in Economic Development,” *Nature Human Behaviour*, 1, 665.
- BULDYREV, S. V., R. PARSHANI, G. PAUL, H. E. STANLEY, AND S. HAVLIN (2010): “Catastrophic Cascade of Failures in Interdependent Networks,” *Nature*, 464, 1025.
- CABRALES, A., P. GOTTARDI, AND F. VEGA-REDONDO (2017): “Risk Sharing and Contagion in Networks,” *The Review of Financial Studies*, 30, 3086–3127.
- CARVALHO, V. M., M. NIREI, Y. U. SAITO, AND A. TAHBAZ-SALEHI (2020): “Supply Chain Disruptions: Evidence from the Great East Japan Earthquake,” University of Cambridge.
- CARVALHO, V. M. AND A. TAHBAZ-SALEHI (2019): “Production networks: A primer,” *Annual Review of Economics*, 11, 635–663.
- CARVALHO, V. M. AND N. VOIGTLÄNDER (2014): “Input Diffusion and the Evolution of Production Networks,” Tech. rep., National Bureau of Economic Research.
- CHENERY, H. B., S. ROBINSON, AND M. SYRQUIN (1986): *Industrialization and Growth*, world Bank Washington.
- CICCONE, A. (2002): “Input Chains and Industrialization,” *The Review of Economic Studies*, 69, 565–587.
- CRAIGHEAD, C. W., J. BLACKHURST, M. J. RUNGTUSANATHAM, AND R. B. HANDFIELD (2007): “The Severity of Supply Chain Disruptions: Design Characteristics and Mitigation Capabilities,” *Decision Sciences*, 38, 131–156.
- DATTA, P. (2017): “Supply Network Resilience: a Systematic Literature Review and Future Research,” *The International Journal of Logistics Management*.
- ELLIOTT, M., B. GOLUB, AND M. O. JACKSON (2014): “Financial Networks and Contagion,” *American Economic Review*, 104, 3115–53.
- ELLIOTT, M., J. HAZELL, AND C.-P. GEORG (2018): “Systemic Risk-Shifting in Financial Networks,” *Available at SSRN 2658249*.
- EROL, S. (2018): “Network Hazard and Bailouts,” *Available at SSRN 3034406*.
- EROL, S., F. PARISE, AND A. TEYTELBOYM (2020): “Contagion in Graphons,” *Available at SSRN*.
- EROL, S. AND R. VOHRA (2018): “Network Formation and Systemic Risk,” *Available at SSRN 2546310*.
- FAFCHAMPS, M. (2002): “Spontaneous Market Emergence,” *Topics in Theoretical Economics*, 2.
- FAFCHAMPS, M. AND B. MINTEN (1999): “Relationships and Traders in Madagascar,” *The Journal of Development Studies*, 35, 1–35.
- GABAIX, X. (2011): “The Granular Origins of Aggregate Fluctuations,” *Econometrica*, 79, 733–772.

- GALE, D. AND S. KARIV (2009): "Trading in Networks: A Normal Form Game Experiment," *American Economic Journal: Microeconomics*, 1, 114–132.
- GOYAL, S. (2017): "Networks and markets," in *Advances in Economics and Econometrics: Volume 1: Eleventh World Congress*, Cambridge University Press, vol. 58, 215.
- GRADY, K. (2006): "Markets: the Fulton Fish Market," *Journal of Economic Perspectives*, 20, 207–220.
- GRANOVETTER, M. (1973): "The Strength of Weak Ties," *American Journal of Sociology*, 78, 1360–1380.
- (1985): "Economic Action and Social Structure: the Problem of Embeddedness," *American Journal of Sociology*, 91, 481–510.
- HALL, R. E. (1979): "A Theory of the Natural Unemployment Rate and the Duration of Employment," *Journal of Monetary Economics*, 5, 153–169.
- HENDRICKS, K. B. AND V. R. SINGHAL (2003): "The Effect of Supply Chain Glitches on Shareholder Wealth," *Journal of operations Management*, 21, 501–522.
- (2005a): "Association between Supply Chain Glitches and Operating Performance," *Management Science*, 51, 695–711.
- (2005b): "An Empirical Analysis of the Effect of Supply Chain Disruptions on Long-Run Stock Price Performance and Equity Risk of the Firm," *Production and Operations Management*, 14, 35–52.
- HERSKOVIC, B. (2018): "Networks in Production: Asset Pricing Implications," *The Journal of Finance*, 73, 1785–1818.
- JACKSON, M. O. (2008): *Social and Economic Networks*, Princeton University Press.
- JACKSON, M. O. AND A. PERNOD (2019): "What Makes Financial Networks Special? Distorted Investment Incentives, Regulation, and Systemic Risk Measurement," *Distorted Investment Incentives, Regulation, and Systemic Risk Measurement (March 1, 2019)*.
- JACKSON, M. O., T. RODRIGUEZ-BARRAQUER, AND X. TAN (2012): "Social Capital and Social Quilts: Network Patterns of Favor Exchange," *American Economic Review*, 102, 1857–97.
- JONES, C. I. (2011): "Intermediate Goods and Weak Links in the Theory of Economic Development," *American Economic Journal: Microeconomics*, 3, 1–28.
- JOVANOVIC, B. (1987): "Micro Shocks and Aggregate Risk," *The Quarterly Journal of Economics*, 102, 395–409.
- KIRMAN, A. P. AND N. J. VRIEND (2000): "Learning to be Loyal. A Study of the Marseille Fish Market," in *Interaction and Market structure*, Springer, 33–56.
- KÖNIG, M. D., A. LEVCHENKO, T. ROGERS, AND F. ZILIBOTTI (2019): "Aggregate Fluctuations in Adaptive Production Networks," Tech. rep., Mimeo.
- KRANTON, R. E. AND D. F. MINEHART (2001): "A theory of Buyer-Seller Networks," *American Economic Review*, 91, 485–508.
- KREMER, M. (1993): "The O-ring Theory of Economic Development," *The Quarterly Journal of Economics*, 108, 551–575.
- LEONTIEF, W. W. (1936): "Quantitative Input and Output Relations in the Economic Systems of the United States," *The Review of Economic Statistics*, 105–125.
- LEVCHENKO, A. A. (2007): "Institutional Quality and International Trade," *The Review of Economic Studies*, 74, 791–819.
- LEVINE, D. (2012): "Production Chains," *Review of Economic Dynamics*, 15, 271–282.
- MURDOCH, J., T. MARSDEN, AND J. BANKS (2000): "Quality, Nature, and Embeddedness: Some Theoretical Considerations in the Context of the Food Sector," *Economic Geography*, 76, 107–125.
- NISHIGUCHI, T. (1994): *Strategic Industrial Sourcing: The Japanese advantage*, Oxford University Press on Demand.
- NORRMAN, A. AND U. JANSSON (2004): "Ericsson's Proactive Supply Chain Risk Management Approach After a Serious sub-Supplier Accident," *International Journal of Physical Distribution & Logistics Management*, 34, 434–456.
- NUNN, N. (2007): "Relationship-Specificity, Incomplete Contracts, and the Pattern of Trade," *The Quarterly Journal of Economics*, 122, 569–600.
- OBERFIELD, E. (2018): "A Theory of Input-Output Architecture," *Econometrica*, 86.
- PETERSEN, M. A. AND R. G. RAJAN (1994): "The Benefits of Lending Relationships: Evidence from Small Business Data," *The Journal of Finance*, 49, 3–37.
- (1995): "The Effect of Credit Market Competition on Lending Relationships," *The Quarterly Journal of Economics*, 110, 407–443.
- ROSTEK, M. AND M. WERETKA (2015): "Dynamic Thin Markets," *The Review of Financial Studies*, 28, 2946–2992.
- SADLER, E. (2020): "Diffusion Games," *American Economic Review*, 110, 225–270.
- SCHEINKMAN, J. A. AND M. WOODFORD (1994): "Self-Organized Criticality and Economic Fluctuations," *The American Economic Review Paper and Proceedings*, 84, 417–421.
- SLUTSKEN, H. (2018): "Four million parts, 30 countries: How an Airbus A380 comes together," *CNN*.
- SNYDER, L. V., Z. ATAN, P. PENG, Y. RONG, A. J. SCHMITT, AND B. SINSOYAL (2016): "OR/MS Models for Supply Chain Disruptions: A Review," *IIE Transactions*, 48, 89–109.

- SUPPLY CHAIN QUARTERLY (2018): “Supply Chain Disruptions Hit Record High,” *Supply Chain Quarterly*, <https://www.supplychainquarterly.com/news/20181210-supply-chain-disruptions-hit-record-high/>.
- TALAMÁS, E. AND R. VOHRA (2020): “Go Big or Go Home: A Free and Perfectly Safe but Only Partially Effective Vaccine Can Make Everyone Worse Off,” *Games and Economic Behavior*, 122, 277–289.
- TANG, L., K. JING, J. HE, AND H. E. STANLEY (2016): “Complex Interdependent Supply Chain Networks: Cascading Failure and Robustness,” *Physica A: Statistical Mechanics and its Applications*, 443, 58–69.
- The Economist* (2006): “When the Chain Breaks,” .
- TINTELNOT, F., A. K. KIKKAWA, M. MOGSTAD, AND E. DHYNE (2018): “Trade and Domestic Production Networks,” Tech. rep., National Bureau of Economic Research.
- UZZI, B. (1997): “Social Structure and Competition in Interfirm Networks: The Paradox of Embeddedness,” *Administrative Science Quarterly*, 35–67.
- YANG, Q., C. SCOGLIO, AND D. GRUENBACHER (2019): “Discovery of a Phase Transition Behavior for Supply Chains against Disruptive Events,” *arXiv preprint arXiv:1908.02616*.

APPENDIX A. FORMAL CONSTRUCTION OF THE SUPPLY NETWORK

We now formally construct the random supply network which was introduced in Sections 2.1 and 2.2; this construction provides the foundation for calculating the set of functional firms, etc.

Endow each \mathcal{V}_i with the Borel σ -algebra and a scaling of Lebesgue measure, denoted λ , and let \mathcal{V} be the disjoint union of these spaces, with total measure 1. Fix positive integers m and n , as well as a distribution μ over the nonnegative integers. A (symmetric) *potential supply network with parameters m , n , μ* is a random graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ satisfying the following properties.

- Its nodes are the set \mathcal{V} .
- Edges are ordered pairs (v, v') where $v' = (j, f')$ for $j \in I(i)$ —the meaning is that v can potentially source from v' . We depict such an edge as an arrow from v to v' .
- The measure of nodes v with $d(v) = d'$ is $\mu(d')$.
- Consider any $v \in \mathcal{V}_i$ with $d(v) > 0$.
 - For each $j \in I(i)$ there are n edges (v, v') to n distinct varieties $v' \in \mathcal{V}_j$. For any variety, define its neighborhood $N_v = \{v' : (v, v') \in \mathcal{E}\}$.
 - The elements of $N_v \cap \mathcal{V}_j$ are independently drawn from an atomless distribution over \mathcal{V}_j conditioned on $d(v') = d(v) - 1$.
 - For any countable set of varieties $\widehat{\mathcal{V}}$, its neighborhoods $(N_v)_{v \in \widehat{\mathcal{V}}}$ are independent.

Now fix *relationship strengths* $(x_v)_{v \in \mathcal{V}}$.⁴⁹ Define \mathcal{G}' to be a random subgraph of \mathcal{G} in which each edge from a node v of positive depth is kept independently, with probability x_v . More formally, define for every edge vv' a random variable $O_{vv'} \in \{0, 1\}$ (whether the edge is operational) such that

- $\mathbf{P}[O_{vv'} = 1 \mid \mathcal{G}] = x_v$ for every $vv' \in \mathcal{G}$ and,
- for any countable subset E of edges in \mathcal{G} , the random variables $(O_e)_{e \in E}$ are independent conditional on \mathcal{G} .

The out-neighborhood of the depth-0 varieties in \mathcal{G}' is $\bigcup_{j \in I(i)} \mathcal{V}_j$, since these can source from anyone.

A subset $\widehat{\mathcal{V}} \subseteq \mathcal{V}$ is defined to be *consistent* if, for each $v \in \widehat{\mathcal{V}}$ the following holds: for each product j that $v = (i, f)$ requires as inputs ($\forall j \in I(i)$), there is an operational edge $(v, v') \in \mathcal{G}'$ with $v' \in \widehat{\mathcal{V}}$. There may be many consistent sets, but by Tarki’s theorem, there will be a maximal one, \mathcal{V}' , which is a superset of any other consistent set. For any given variety, this can be found by the simple iterative procedure in Section 2.3.

Since any countable set of edges is independent, we can make computations about the relevant marginal probabilities in our model (e.g., reliability of any variety) as we would if there only one tree.

APPENDIX B. PROOFS OF MAIN RESULTS

B.1. The reliability curve: Basic precipice results. The first building block of our analysis is characterizing the shape of the reliability curve $x \mapsto \rho(x, \mu_\tau)$ when τ is large (so that depths are large). In Figure 10, we reproduce Figure 3 (adding one additional panel), which will help us describe what we are doing. First, we define a correspondence that will be central to our analysis. It is a limit of the functions

⁴⁹Here $v \mapsto x_v$ is a given measurable function.




FIGURE 10. Panel (A) shows how reliability varies with the investment level x when supply chains have a finite depth. Panel (B) shows the limit correspondence of the relationships between reliability r and the investment level x as the depth gets large. Panel (C) shows how reliability varies with the investment level x for infinite depth.

$x \mapsto \rho(x, \mu_\tau)$, in the sense that their graphs converge to the graph of the correspondence $\rho(x)$ as $\tau \rightarrow \infty$ (see panels (A) and (B)). More formally:

Definition 3. Let $\rho : [0, 1] \rightrightarrows [0, 1]$ be a correspondence such that⁵⁰

- (i) for any $r \in \rho(x)$, there is a sequence $\{x_\tau\}_{\tau=1}^\infty \rightarrow x$ such that $\lim_{\tau \rightarrow \infty} \rho(x_\tau, \mu_\tau) = r$;
- (ii) for any sequence $\{x_\tau\}_{\tau=1}^\infty \rightarrow x$ we have $\lim_{\tau \rightarrow \infty} \rho(x_\tau, \mu_\tau) \in \rho(x)$.

We will show that this limit correspondence ρ exists and is uniquely defined, and that it has the shape sketched in panel (B): zero until some point, with a vertical rise up to a certain critical level of reliability, followed by a concave ascent to reliability 1 at $x = 1$. With these results in hand, this subsection will culminate in the proofs of Proposition 1 and Proposition 2.

Our main result on the reliability correspondence is:

Proposition 5 (Shape of the limit reliability correspondence ρ). Let the complexity of production be $m \geq 2$ and the number of potential suppliers for each input be $n \geq 2$. Then there is a unique ρ satisfying Definition 3, and it has the following properties. There exists an $x_{\text{crit}} \in (0, 1)$ such that

- (i) $\rho(x)$ is single-valued for all $x \neq x_{\text{crit}}$;
- (ii) $\rho(x) = 0$ for all $x < x_{\text{crit}}$;
- (iii) there is a value $0 < \bar{r}_{\text{crit}} < 1$ such that $\rho(x_{\text{crit}}) = [0, \bar{r}_{\text{crit}}]$;
- (iv) $\rho(x)$ is strictly increasing in x for all $x > x_{\text{crit}}$;
- (v) $\lim_{x \downarrow x_{\text{crit}}} \rho'(x) = \infty$.

Sections B.1.1 and B.1.2 are devoted to proving this result. A closely related (discontinuous) function, sketched in panel (c), will play a role in this analysis, as we describe next. Once this is done, we will deduce Propositions 1 and 2.

B.1.1. The reliability function with infinite depth: A tool to study ρ . In order to prove Proposition 5 it is helpful to define and characterize some properties of a certain function $\hat{\rho} : [0, 1] \rightarrow [0, 1]$, sketched in Figure 10(C). It will turn out that for all $x \neq x_{\text{crit}}$, $\rho(x) = \hat{\rho}(x)$. We define $\hat{\rho}(x)$ as the largest r solving the following equation

$$\hat{r} = (1 - (1 - x\hat{r})^n)^m. \quad (6)$$

(For intuition, note that this is equation (2), which relates reliability at depths d and $d - 1$, but here with the same reliability r on both sides. This is an intuitive condition “at depth infinity” as reliability is not decreased by adding one layer of depth.)

Lemma 2, which is proved in Section SA1.1 of the Supplementary Appendix, identifies several key properties of $\hat{\rho}(x)$.

⁵⁰Equivalently, we can define ρ by saying that it is a correspondence such that the graphs of the functions $\rho(\cdot, \mu_\tau) : [0, 1] \rightarrow [0, 1]$ converge to the graph of ρ (in the Hausdorff set-distance metric).

Lemma 2. Suppose the complexity of the supply network is $m \geq 2$ and there are $n \geq 1$ potential input suppliers of each firm. For $r \in (0, 1]$ define

$$\chi(r) := \frac{1 - \left(1 - r^{\frac{1}{m}}\right)^{\frac{1}{n}}}{r}. \quad (7)$$

Then there are values $x_{\text{crit}}, \bar{r}_{\text{crit}} \in (0, 1]$ such that:

- (i) $\hat{\rho}(x) = 0$ for all $x < x_{\text{crit}}$;
- (ii) $\hat{\rho}$ has a (unique) point of discontinuity at x_{crit} ;
- (iii) $\hat{\rho}$ is strictly increasing for $x \geq x_{\text{crit}}$;
- (iv) the inverse of $\hat{\rho}$ on the domain $x \in [x_{\text{crit}}, 1]$, is given by χ on the domain $[\bar{r}_{\text{crit}}, 1]$, where $\bar{r}_{\text{crit}} = \hat{\rho}(x_{\text{crit}})$;
- (v) χ is positive and quasiconvex on the domain $(0, 1]$;
- (vi) $\chi'(\bar{r}_{\text{crit}}) = 0$.

The proof of Lemma 2 is in Section SA1.1 of the Supplementary Appendix. While the manipulations to establish these properties are a bit involved, they amount to studying the function $\hat{\rho}$ and the pseudo-inverse of it that we have defined, χ , using calculus. Figure 11(a) depicts χ as a function of r .⁵¹




FIGURE 11. Panel (A) plots the function χ as r varies, and then in Panel (B) we show how switching the axes and taking the largest r value on the graph (corresponding to the largest solution of equation (6)) generates $\hat{\rho}(x)$.

Recall the functions $\tilde{\rho}(x, d)$: consider a depth- d tree where each firm in each tier requires m kinds of inputs and has n potential suppliers of each input. We denote by $\tilde{\rho}(x, d)$ the probability of successful production at the most-downstream node of a depth- d tree with these properties. This is defined as

$$\tilde{\rho}(x, d) = (1 - (1 - x\tilde{\rho}(x, d-1))^n)^m$$

with $\tilde{\rho}(x, 0) = 1$, since the most-upstream tier nodes obtain their inputs without the possibility of disruption.

It will be useful throughout that $\tilde{\rho}(x, d)$ and its derivative converge uniformly to the function $\hat{\rho}(x)$ everywhere except near x_{crit} . Indeed, this will directly imply that reliability curve and the marginal returns to investment in our model also converge (at all points but x_{crit}) to their corresponding values in the “infinite-depth” model that defines $\hat{\rho}$. The following lemma formalizes these statements.

Lemma 3.

- (i) For all $d \geq 1$, the function $x \mapsto \tilde{\rho}(x, d)$ defined for $x \in (0, 1)$ is strictly increasing and infinitely differentiable.
- (ii) On any compact set excluding x_{crit} , the sequence $(\tilde{\rho}(x, d))_{d=1}^{\infty}$ converges uniformly to $\hat{\rho}(x)$.

⁵¹Since by definition the *largest* r satisfying (6) is the one that determines $\hat{\rho}(x)$, it follows that the increasing part of the function, where $r \in [\bar{r}_{\text{crit}}, 1]$ is the part relevant for determining equilibrium reliability—see Figure 11(b), where the light gray branch is not part of $\hat{\rho}$.

(iii) On any compact set excluding x_{crit} , the sequence $(\rho(x, \mu_\tau))_{\tau=1}^\infty$ converges uniformly to $\widehat{\rho}(x)$.

Proof. The sequence $(\widetilde{\rho}(\cdot, d))_{d=1}^\infty$ is a monotone sequence of increasing, infinitely differentiable functions,⁵² converging pointwise to $\widehat{\rho}$. We know that $\widehat{\rho}$ is continuous on any compact set excluding x_{crit} . Therefore, by Dini's theorem, the functions $\widetilde{\rho}(\cdot, d)$ converge uniformly to $\widehat{\rho}$.

Note the functions $\rho(\cdot, \mu_\tau)$ are strictly increasing and infinitely differentiable, since they are the averages of such functions. Moreover, they clearly converge pointwise to $\widehat{\rho}(x)$: since all the $\widehat{\rho}(\cdot, d)$ are uniformly bounded, the vanishing probability mass on low- d realizations makes only a negligible contribution to the value of $\rho(x, \mu_\tau)$ for large τ . From this the conclusions of Lemma 3 apply equally well to the sequence of functions $(\rho(\cdot, \mu_\tau))_{\tau=1}^\infty$. \square

B.1.2. Proof of Proposition 5. We will now establish a close relationship between $\widehat{\rho}$ and our correspondence of interest, ρ , which will allow us to use Lemmas 2 and 3 to prove Proposition 5.

Define $\rho(x) = \{\widehat{\rho}(x)\}$ for $x \neq x_{\text{crit}}$. By Lemma 3(iii) $\rho(x, \mu_\tau) \rightarrow \widehat{\rho}(x)$ as $\tau \rightarrow \infty$ for any $x \neq x_{\text{crit}}$. Thus Part (ii) of Definition 3 is satisfied. Since $\rho(x)$ is single-valued for $x \neq x_{\text{crit}}$, Part (i) of Definition 3 also holds. Lemma 2 then implies points (i), (ii), (iv) and (v) of Proposition 5.

Next, define $\rho(x_{\text{crit}}) = [0, \bar{r}_{\text{crit}}]$. We will show that the remaining conditions of Definition 3 (those pertaining to x_{crit}) also hold. First, we will check Part (i) of Definition 3 for $r \in \rho(x_{\text{crit}})$. For any $r \in [0, \bar{r}_{\text{crit}}]$, we can construct a sequence $\{x_\tau\}_{\tau=1}^\infty \rightarrow x_{\text{crit}}$ such that $\lim_{\tau \rightarrow \infty} \rho(x_\tau, \mu_\tau) = r$. To see this, simply note that for any $r \in (0, \bar{r}_{\text{crit}}]$ and any τ , there is an x_τ such that $\rho(x_\tau, \mu_\tau) = r$ since $\rho(x, \mu_\tau)$ is a continuous and increasing function of x whose image is $[0, 1]$. Moreover, since for any $x < x_{\text{crit}}$, $\lim_{\tau \rightarrow \infty} \rho(x, \mu_\tau) = 0$ and for any $x > x_{\text{crit}}$ there exists $\epsilon > 0$ such that $\lim_{\tau \rightarrow \infty} \rho(x, \mu_\tau) = \bar{r}_{\text{crit}} + \epsilon$, it follows that these x_τ must indeed converge to x_{crit} . In the case of $r = 0$, there exists a sequence $\{x_\tau\}_{\tau=1}^\infty \rightarrow x_{\text{crit}}$ such that $\rho(x_\tau, \mu_\tau) \downarrow 0$ since for any $x < x_{\text{crit}}$, $\rho(x, \mu_\tau) \downarrow 0$.

Finally, we show that every sequence $\{x_\tau\}_{\tau=1}^\infty \rightarrow x_{\text{crit}}$ satisfies $\lim_{\tau \rightarrow \infty} \rho(x_\tau, \mu_\tau) \in [0, \bar{r}_{\text{crit}}]$ (Part (ii) of Definition 3). This amounts to showing that the limit is at most r_{crit} . Suppose otherwise, that such a limit is $r' > r_{\text{crit}}$. Let x' be such that $\widehat{\rho}(x') = r'$, which exists by Lemma 2. Note for all $x_\tau < x'$, we have $\widehat{\rho}(x_\tau) < r' - \epsilon$ for some positive ϵ , and so $\rho(x_\tau, \mu_\tau) < r'$ for sufficiently large τ ; this is a contradiction to the hypothesis about the limit being r' .

These claims together establish the remaining content of the claim that the ρ we have defined is the limit satisfying Definition 3. That implies part (iii) of Proposition 5. \square

B.1.3. Proof of Proposition 1 (Discontinuity in reliability) using the limit reliability correspondence. We can now use the results established to prove the proposition about the precipice.

Part (i): The fact that $\rho(x, \mu_\tau) \rightarrow 0$ for $x < x_{\text{crit}}$ follows from Definition 3 and Proposition 5(ii).

Part (ii): From Proposition 5(iii) and (iv), it follows that $\rho(x) > \bar{r}_{\text{crit}}$ for any $x > x_{\text{crit}}$ and thus, for τ large enough, $\rho(x, \mu_\tau) > \bar{r}_{\text{crit}} > 0$. \square

B.2. Proof of Proposition 2 (Social planner's solution). Now we can use the properties of the reliability curve to prove our result on the social planner's solution. Recall that $x^{SP}(\kappa, \mu_\tau)$ is the set of all values of x maximizing the planner's objective.

First consider the following "limit" planner's problem, defined for $x \in [0, 1]$:

$$\max \left[h(r) - \frac{1}{\kappa} c_P(x) \right] \text{ subject to } r \in \rho(x). \quad (8)$$

We claim that, if $\delta > 0$ is chosen small enough, then any solution has $x \notin [x_{\text{crit}} - \delta, x_{\text{crit}} + \delta]$. To rule out solutions with $x \in [x_{\text{crit}} - \delta, x_{\text{crit}})$, note that by Property B and continuity of c_P that the cost $c_P(x)$ is positive in that interval (if δ is fixed at a small enough value) while $\rho(x)$ is zero by Proposition 1, so that $x = r = 0$ does better. To rule out solutions with $x \in [x_{\text{crit}}, x_{\text{crit}} + \delta]$, first observe that for $x = x_{\text{crit}}$, the only r that can be a solution is $r = r_{\text{crit}}$, so we may assume the solution lies on the graph of $\widehat{\rho}$. Next, note that if δ is chosen small enough, then for any κ , we have $\widehat{\rho}'(x) > \frac{1}{\kappa} c_P(x)$ in the interval $[x_{\text{crit}}, x_{\text{crit}} + \delta]$ since $\widehat{\rho}'(x)$ tends to $+\infty$ as $x \downarrow x_{\text{crit}}$ by Lemma 3.

Let $\kappa_{\text{crit}} > 0$ be the minimum value of κ such that that (8) has a solution with $r > 0$. This κ_{crit} exists because for large enough κ we have that the maximand is positive at $x = x_{\text{crit}}$, and so there must be a solution

⁵²From eq. (2), it is clear that increasing d decreases $\widetilde{\rho}(x, d)$, while increasing x increases it. Differentiability is straightforward from the iterative definition.

with positive r ; similarly, κ_{crit} is positive since for small κ , the maximand is negative for all $x \geq x_{\text{crit}}$.⁵³ By the previous paragraph, any solution at $\kappa = \kappa_{\text{crit}}$ satisfies $x > x_{\text{crit}}$. Since in the domain $x > x_{\text{crit}}$, we have that $\rho(x)$ is concave while c_P is convex, for $\kappa \geq \kappa_{\text{crit}}$ there is a unique, strictly positive solution $(x(\kappa), r(\kappa))$ of (8) and both x and r are increasing in κ . For $\kappa < \kappa_{\text{crit}}$, we have that $\rho(x) = 0$ at the optimum and hence cost is zero (otherwise $r = x = 0$ would do better). Thus, by what we have said above,

- (i) for all $\kappa < \kappa_{\text{crit}}$, all solutions of (8) have $x \leq x_{\text{crit}} - \delta$, have cost equal to 0, and yield reliability 0;
- (ii) for all $\kappa > \kappa_{\text{crit}}$, all solutions of (8) have $x \geq x_{\text{crit}} + \delta$ and yield reliability $r \geq r_{\text{crit}}$;
- (iii) for $\kappa = \kappa_{\text{crit}}$, all solutions of (8) are outside the interval $[x_{\text{crit}} - \delta, x_{\text{crit}} + \delta]$ and reliability is either 0 or strictly above r_{crit} .

Returning to the main model rather than the limit case above: for any κ and τ , consider the following rewriting of the planner's problem:

$$\max \left[h(r) - \frac{1}{\kappa} c_P(x) \right] \text{ subject to } r = \rho(x; \mu_\tau). \quad (9)$$

For any κ , it is straightforward to deduce from the definition of the correspondence ρ that a sequence of solutions $(x(\tau), \rho(x; \mu_\tau))$ of (9) must converge to a solution of (8) as $\tau \rightarrow \infty$.⁵⁴ Thus, choosing ϵ small enough and using the convergence of $\rho(x; \mu_\tau)$ to $\rho(x)$ per Definition 3 (along with the properties of ρ established in Proposition 1), the conclusions of the present proposition follow. \square

B.3. Best responses and equilibrium behavior: Overview and preliminaries. Our goal now is to work toward Theorem 1 on the ordered regimes and the properties of equilibrium.

For notational convenience, when analyzing equilibria, we will multiply through all firms' profit functions by κ , so that a firm's cost is $c(x_{if} - \underline{x})$ and its gross profit is $\kappa g(r)P(x_{if}; r)$. Note this does not change any best responses. We maintain this convention throughout the remainder of the proofs, and in the Supplementary Appendix.

In Figure 12, we reproduce part of Figure 5, which will guide our analysis. Before beginning, we outline the plan. At a high level, it amounts to bringing the best-response curve into the picture along with the precipice graph, and ultimately formalizing the graphical intuitions in Figure 12 concerning what happens to the best-response curve as κ moves around. Throughout, a recurring strategy is to prove a result for the game "at the $\tau = \infty$ " limit (in a sense we make precise), and then extend the conclusion to large values of τ , where most of the mass is on deep supply trees.

To study this $\tau = \infty$ limit, We establish some notation. First, define

$$\widehat{P}(x_{if}; r) = (1 - (1 - x_{if}r)^n)^m. \quad (10)$$

This is the probability with which if is able to produce, as a function of its investment choice x_{if} given that all its suppliers have reliability r . This is simpler than our large- τ problem because all firms are exactly symmetric (which is only approximately true in the large- τ model).

Define

$$MB(x_{if}; r, \kappa) = \kappa g(r) \frac{\partial \widehat{P}(x_{if}; r)}{\partial x_{if}} \quad (11)$$

$$MC(x_{if}) = c'(x_{if} - \underline{x}). \quad (12)$$

These are the marginal benefit and marginal cost, respectively, to a firm of investing in relationship strength. Again we use the mechanics of the system "at the limit"; this is useful because we will find that the actual, finite- τ equilibrium quantities converge to these asymptotic ones.

A first essential step in the analysis is making best responses at $\tau = \infty$ tractable by expressing them as the solution to a first-order condition equating marginal benefits and marginal costs, which is unique under suitable assumptions. This is done in Section B.3.3, by establishing that the marginal benefit function is quasiconcave in the range of interest.

Second, we show in Section B.3.2 that if we are interested in equilibria that are undominated in terms of social surplus, we may focus on the higher points of intersection in Figure 12.

⁵³Since $c_P(x_{\text{crit}}) > 0$ and c_P is convex, $c_P(x) > 0$ for all $x \geq x_{\text{crit}}$. Thus by making κ small enough, the costs exceed the bounded benefits. We can say "minimum" rather than "infimum" in the definition of κ_{crit} because the correspondence ρ is continuous, as is c_P .

⁵⁴This is a version of the Theorem of the Maximum on upper hemicontinuity of optimization.




FIGURE 12. Panel (A) shows an equilibrium for $\kappa < \bar{\kappa}$. Panel (B) shows an equilibrium with $\kappa = \bar{\kappa}$. Panel (C) shows an equilibrium with $\kappa = \bar{\kappa}$. Panel (D) shows an equilibrium with $\kappa > \bar{\kappa}$. Panel (E) plots how equilibrium reliability varies with κ . Panel (F) shows reliability following an arbitrarily small negative shock to institutional quality \underline{x} as κ varies.

Third, we show in Section B.3.3 that in the limit model there is a unique intersection between the two curves above reliability \bar{r}_{crit} , as in Figure 12(d).

Together, these ingredients set up the proof of Theorem 1 in Section B.4, allowing us to formalize the graphical intuition.

The remaining subsections carry out this plan.

B.3.1. Proof of Lemma 1 (Sufficient condition for unique interior solution to the firm's problem). We first study the firm's problem "at $\tau = \infty$." Afterward, we establish that the same uniqueness property holds for large, finite τ .

For the extended domain $x_{if} \in [0, 1/r]$, we define

$$Q(x_{ik}; r) := \frac{\partial}{\partial x_{if}} \hat{P}(x_{if}; r), \quad (13)$$

which can be calculated to be $Q(x_{ik}; r) = mn(1 - (1 - x_{if}r)^n)^{m-1}(1 - x_{if}r)^{n-1}r$.

We will need two steps to prove Lemma 1. The first step consists of establishing Lemma 4 on the basic shape of $Q(x_{if}; r)$. Figure 13 illustrates the shape of $Q(x_{if}; r)$ implied by Lemma 4.

Lemma 4. Fix any $m \geq 2$, $n \geq 2$, and $r \geq r_{crit}$. There are uniquely determined real numbers x_1, x_2 (depending on m, n , and r) such that $0 \leq x_1 < x_2 < 1/r$ and so that:

0. $Q(0; r) = Q(1/r; r) = 0$ and $Q(x_{if}; r) > 0$ for all $x_{if} \in (0, 1/r)$;
1. $Q(x_{if}; r)$ is increasing and convex in x_{if} on the interval $[0, x_1]$;
2. $Q(x_{if}; r)$ is increasing and concave in x_{if} on the interval $(x_1, x_2]$;




FIGURE 13. The shape of $Q(x_{if}, x)$ in green and $c'(x_{if} - x)$ in red. There can be only one intersection between the two curves, which corresponds to the maximizer of $\Pi(x_{if}; r)$. (Here we normalized $\kappa g(r) = 1$ for simplicity, but the illustration remains valid if the green curve is scaled.)

3. $Q(x_{if}; r)$ is decreasing in x_{if} on the interval $(x_2, 1]$.
4. $x_1 < x_{\text{crit}}$.

The proof of Lemma 4 is in Section SA1.2 of the Supplementary Appendix.

We now complete the proof of Lemma 1 by setting $\hat{x} = x_1$. By Lemma 4(4), the interval $(\hat{x}, x_{\text{crit}})$ is non-empty. Thus we just need to show that Assumption 1 (the sufficiency of the first-order condition for interior optima) is satisfied when $\underline{x} \in (\hat{x}, x_{\text{crit}})$. Note that by Property B', we have that $c'(0) = 0$ and c' is weakly increasing and weakly convex otherwise. Since, by Lemma 4, $P'(x_{if}; x)$ is first concave and increasing (possibly for the empty interval) and then decreasing (possibly for the empty interval) over the range $x_{if} \in [\underline{x}, 1]$, it follows that there is at most a single intersection between the curves $P'(x_{if}; x)$ and $c'(x_{if} - \underline{x})$. This intersection corresponds to the first-order condition $Q(x_{if}; r) = c'(x_{if} - \underline{x})$, yielding the unique maximizer of $\Pi(x_{if}; x)$, as illustrated in Figure 13. If such an intersection not exist, $y_{if} = 0$ is a local and global maximizer of the profit function.

Now we study the quantities introduced above, but for finite a τ and show that when τ is large enough, we still obtain a unique solution to the firm's problem.

Recall that

$$P(x_{if}; x, \mu_\tau) = \mu_\tau(0) + \sum_d \mu_\tau(d) (1 - (1 - x_{if} \tilde{\rho}(x, d-1))^n)^m.$$

Define

$$Q(x_{if}; x, \mu_\tau) = \frac{\partial}{\partial x_{if}} P'(x_{if}; x, \mu_\tau) = \sum_d \mu_\tau(d) m n (1 - (1 - x_{if} \tilde{\rho}(x, d-1))^n)^{m-1} (1 - x_{if} \tilde{\rho}(x, d-1))^{n-1} \tilde{\rho}(x, d-1).$$

Recall that for all $x \neq x_{\text{crit}}$ we have

$$\tilde{\rho}(x, d) \rightarrow_d \hat{\rho}(x) \text{ and } \rho(x, \mu_\tau) \rightarrow_\tau \hat{\rho}(x).$$

It follows from this and the expression for $Q(x_{if}; x, \mu_\tau)$ above that

$$P(x_{if}; x, \mu_\tau) \rightarrow_\tau \hat{P}(x_{if}; \hat{\rho}(x)) \text{ and } Q(x_{if}; x, \mu_\tau) \rightarrow_\tau Q(x_{if}; \hat{\rho}(x)).$$

It follows that marginal costs and marginal benefits in the $\tau \rightarrow \infty$ limit are arbitrarily close to the functions studied above. We then conclude that for τ large enough, there will also be at most a single intersection between the curves $Q(x_{if}; x, \mu_\tau)$ and $c'(x_{if} - \underline{x})$ and thus the results stated in the limit also hold for τ large enough. \square

B.3.2. A lemma on symmetric undominated equilibria. This subsection establishes that the symmetric equilibrium maximizing social surplus is the one with greatest relationship strengths x and highest reliability. In other words, it is the higher intersection in Figure 12 when there are several. Normalizing $\kappa = 1$ without

loss of generality, recall our notation that at a symmetric equilibrium with reliability r , gross output is $h(r)$ and gross profits are $g(r)$.

Lemma 5. Consider two symmetric equilibria with relationship strengths $x_1 < x_2$ and reliabilities $r_1 < r_2$. Then $h(r_2) - c_P(x_2) > h(r_1) - c_P(x_1)$. That is, the higher-investment equilibrium has the greater net social surplus.

Proof. Let $V(x) = h(\rho(x, \mu)) - c_P(x)$. We will first show that V is increasing at $x = x_2$, and then use this to deduce the conclusion.

Consider an outcome with firm investments x_V and a reliability of r . Note that the gross profit per functional firm is $g(r)$ so that the gross profit integrated across all firms (functional or nonfunctional) is

$$\text{GP} = rg(r).$$

We will write net social surplus as follows, and then split it up into consumer and producer surplus:

$$\begin{aligned} V &= h(r) - \int_{\mathcal{V}} c(x_v) dv \\ &= [h(r) - \text{GP}] + \left[\text{GP} - \int_{\mathcal{V}} c(x_v) dv \right] \\ &= \underbrace{[h(r) - rg(r)]}_{\text{CS}} + \text{PrS}. \end{aligned}$$

Here we define total producer surplus, PrS as gross profits net of investment costs.

Now, it will be useful to have two different expressions for producer surplus, which are identically equal:

$$\text{PrS}_1 = rg(r) - \int_{\mathcal{V}} c(x_v) dv \tag{14}$$

$$\text{PrS}_2 = \int_{\mathcal{V}} [\Pi_v(x_v; r) - c(x_v)] dv. \tag{15}$$

Now make all the variables— x_v 's and r 's—functions of a parameter t , according to $x_v = x + t$ for small t , and r is a function of the x_v (and thus of t). We want to totally differentiate V in t . We claim the derivative is positive. First consider

$$\frac{d}{dt} \text{CS} = \frac{d}{dt} [h(r) - rg(r)].$$

Now we turn to producer surplus. Here, we need to totally differentiate in t . The endogenous variables are the continuum of x_v 's and r :

$$\frac{d}{dt} \text{PrS} = \frac{dr}{dt} \cdot \frac{\partial}{\partial r} \text{PrS} + \frac{dx_v}{dt} \int_{\mathcal{V}} \frac{\partial}{\partial x_v} \text{PrS}.$$

Now, the key idea of the argument is that we use the PrS_1 expression for PrS in the first summand, and the PrS_2 expression for PrS in the second summand, which is legitimate since the expressions are identically equal. The second expression evaluates to 0 because for a fixed v , we have

$$\frac{\partial}{\partial x_v} \text{PrS}_2 = \frac{\partial}{\partial x_v} [\Pi_v(x_v; r) - c(x_v)],$$

which is 0 by the firm's first-order condition.

Using $V = \text{CS} + \text{PrS}$ and our calculations above,

$$\frac{dV}{dt} = \frac{dr}{dt} h'(r) = \frac{d}{dx} \rho(x, \mu) h'(\rho(x, \mu)).$$

Both factors are clearly positive, so V is increasing in t .

Note that for $x \geq x_{\text{crit}}$, we have that h is concave by Property A and increasing in x while $c_P(x)$ is convex, so

$$h(\rho(x, \mu)) - c_P(x)$$

is single-peaked. The above argument implies both equilibria are socially inefficient, and thus social surplus is increasing in x at both of them. This combined with single-peakedness shows that the one with higher relationship strength is more efficient. \square

An immediate corollary that follows from the proof of Lemma 5 is that there is underinvestment in relationship strengths in equilibrium.

Corollary 2. Suppose we have a symmetric equilibrium with relationship strengths $x \geq x_{\text{crit}}$. Then⁵⁵ $x < x^{SP}(\kappa, \mu)$.

B.3.3. A limit uniqueness result for equilibria above \bar{r}_{crit} . In this section we prove an important lemma that will be used to ensure that the relationship strengths and reliability achieved at symmetric undominated equilibria are uniquely determined. As before, we study the “ $\tau = \infty$ ” model as a stepping stone to our large- τ results. Visually, the lemma establishes that there is at most one intersection in Figure 12 that lies above \bar{r}_{crit} .

To motivate the statement, consider the question: when can there be an equilibrium outcome with $r \geq \bar{r}_{\text{crit}}$ for the limit payoffs at $\tau = \infty$? Recall the definitions of marginal benefits and marginal costs from (11) and (12) above. Under Assumption 1, an interior firm optimum is characterized by the first order condition (recall equations (11) and (12) for the definitions)

$$MB(x; r, \kappa) = MC(x), \quad (\text{OI})$$

which we have labeled OI, for optimal investment. In addition, an outcome of the $\tau = \infty$ model also satisfies $r = \hat{\rho}(x)$ for $r \geq \bar{r}_{\text{crit}}$: the reliability is the one induced by firms’ choices of relationship strengths. Recalling the definition of χ from Lemma 2, this entails $x = \chi(r)$. Then we have the equation

$$MB(\chi(r); r, \kappa) = MC(\chi(r)).$$

We will show that (for some $\epsilon > 0$) there is at most one solution (x, r) satisfying $r \in [\bar{r}_{\text{crit}} - \epsilon, 1]$ and $x = \chi(r)$ simultaneously. This implies that there is only one solution to the first-order conditions above \bar{r}_{crit} , and the fact that the necessary condition cannot hold even for slightly lower values of \bar{r}_{crit} will be technically useful later.

Lemma 6. Fix any $n \geq 2$ and $m \geq 3$. There exists $\epsilon > 0$ such that the equation $MB(\chi(r); r, \kappa) = MC(\chi(r))$ has at most one solution r^* in the range $[\bar{r}_{\text{crit}} - \epsilon, 1]$.

Proof. In the following argument, we defer technical steps to lemmas, which are proved in the Supplementary Appendix.




FIGURE 14. Panel (A) shows the relationship between r and x implied by physical consistency. Panel (B) plots the function $\mathfrak{b}(r)$ discussed in the proof.

⁵⁵Note that $x^{SP}(\kappa, \mu)$ is single-valued for such values of x , so we may treat x^{SP} as a function here, abusing notation.

Define $\mathbf{b}(r) = MB(\chi(r); r, \kappa)$ and $\mathbf{c}(r) = MC(\chi(r))$, so that we can simply study the equation

$$\mathbf{b}(r) = \mathbf{c}(r). \quad (16)$$

We can calculate

$$MB(x_{if}; r, \kappa) = \kappa g(r) r n (1 - x_{if} r)^{n-1} m (1 - (1 - x_{if} r)^n)^{m-1}. \quad (17)$$

By plugging in $x = \chi(r)$ into (17), we may rewrite the equation of interest as

$$\underbrace{\kappa g(r) m n r^{2-\frac{1}{m}} \left(1 - r^{1/m}\right)^{1-\frac{1}{n}}}_{\mathbf{b}(r)} = c' \left(\underbrace{\frac{1 - \left(1 - r^{\frac{1}{m}}\right)^{\frac{1}{n}}}{r}}_x - \underline{x} \right). \quad (18)$$

We wish to show that there is a range $[\bar{r}_{\text{crit}} - \epsilon, 1]$ for which there is at most one solution to equation (16).⁵⁶ The right-hand side of the equation is increasing in r for $r \in [\bar{r}_{\text{crit}}, 1]$.⁵⁷ If we could establish that the left-hand side, which we call $\mathbf{b}(r)$, is decreasing in r , it would follow that there is at most a unique r solving (18) for $r \in [\bar{r}_{\text{crit}}, 1]$. Moreover, by the continuity of $\mathbf{b}(r)$ and $\mathbf{c}(r)$, it would follow immediately that for some small $\epsilon > 0$, there is also at most a single solution over the range $[\bar{r}_{\text{crit}} - \epsilon, 1]$. Unfortunately, $\mathbf{b}(r)$ is not decreasing in r . However, we can show that it *is* decreasing in r for $r \geq \bar{r}_{\text{crit}}$, which is sufficient.⁵⁸ Panel (B) of Figure 14 gives a representative depiction of $\mathbf{b}(r)$, reflecting that it is decreasing to the right of \tilde{r} (which is itself to the left of \bar{r}_{crit}).

The following lemma will, once established, complete the proof.

Lemma 7. \mathbf{b} is strictly decreasing on the domain $[\bar{r}_{\text{crit}}, 1]$.

To prove Lemma 7 we write $\mathbf{b}(r)$ as a product of two pieces, $\alpha(r) := \kappa g(r)$ and

$$\beta(r) := m n r^{2-\frac{1}{m}} \left(1 - r^{\frac{1}{m}}\right)^{1-\frac{1}{n}}.$$

Note that the function $\beta(r)$ is positive for $r \in (0, 1)$. We will show that it is also strictly decreasing on $[\bar{r}_{\text{crit}}, 1]$. By assumption, $g(r)$ is positive and strictly decreasing in its argument, so $\alpha(r)$ is also positive and decreasing in r . Thus, because \mathbf{b} is the product of two positive, strictly decreasing functions on $[\bar{r}_{\text{crit}}, 1]$, it is also strictly decreasing on $[\bar{r}_{\text{crit}}, 1]$. It remains only to establish that $\beta(r)$ is strictly decreasing on the relevant domain. Two additional lemmas are helpful.

Lemma 8. The function $\beta(r)$ is quasiconcave and has a maximum at $\hat{r} := \left(\frac{(2m-1)n}{2mn-1}\right)^m$.

Lemma 9. For all $n \geq 2$ and $m \geq 3$, we have that $\hat{r} < r_{\text{crit}}$.

Lemmas 8 and 9 are proved in Sections SA1.3 and SA1.4 of the Supplementary Appendix. Together these show that $\beta(r)$ is strictly increasing and then strictly decreasing in r for $r \in (0, 1)$, with a turning point in the interval $(0, \bar{r}_{\text{crit}})$. Thus $\beta(r)$ is strictly decreasing on the domain $[\bar{r}_{\text{crit}}, 1]$, the final piece required to prove Lemma 7.

Thus, by the continuity of $\mathbf{b}(r)$ and $\mathbf{c}(r)$, it follows immediately that there exists $\epsilon > 0$ such that there is at most a single solution for $r \in [\bar{r}_{\text{crit}} - \epsilon, 1]$. This completes the proof of the lemma. \square

B.4. Proof of Theorem 1 (Classification of regimes as κ varies). Having developed the key machinery that we will use to analyze the comparative statics of equilibria, we can proceed to the main proof. We will sometimes take $\underline{x} = 0$ when this is immaterial to the arguments to keep notation uncluttered.

The proof relies mainly on analyzing the shape of the best-response correspondence: how it depends on κ and the reliability level r .

Denote by $BR(r, \kappa)$ the set of values x_{if} maximizing

$$\kappa g(r) \hat{P}(x_{if}; r) - c(x_{if} - \underline{x}). \quad (19)$$

⁵⁶We note that the functions on both sides of the equation are merely constructs for the proof. In particular, when we sign their derivatives, these derivatives do not have an obvious economic meaning.

⁵⁷This follows directly: χ is increasing on that domain and c' is increasing by assumption. (see Panel (A) of Figure 14).

⁵⁸We do this by showing that the global maximum of $\mathbf{b}(r)$ is achieved at a number \tilde{r} that we can prove is smaller than \bar{r}_{crit} .

Recalling the definition (10), this has the interpretation of best-response relationship strengths in the infinite-depth model. Let $\overline{BR}(r, \kappa) = \max\{x : x \in BR(r, \kappa)\}$ be the maximal element of the set $BR(r, \kappa)$. Analogously, let $\overline{BR}(r, \kappa, \mu_\tau) = \max\{x : x \in BR(r, \kappa, \mu_\tau)\}$. This is the finite- τ analogue of the limit object we have just defined.

Recall the optimal investment first-order condition (OI) in the limit model:

$$MB(x; r, \kappa) = MC(x), \tag{OI}$$

where the definitions are in (11) and (12). We will be using it to study interior optima.

A first lemma in the proof of Theorem 1 is monotonicity of the best response curve in κ .

Lemma 10. $\overline{BR}(r, \kappa)$ is increasing in κ , and strictly so whenever $\overline{BR}(r, \kappa) > 0$.

Proof. Let $x_{if} = \overline{BR}(r, \kappa)$ and suppose OI holds at the optimum. Suppose κ increases slightly. Then, evaluated at x_{if} , the left-hand side of OI increases and the right-hand side does not change as it does not depend on κ . So marginal benefits exceed marginal costs at x_{if} . Now consider increasing x_{if} . By the assumption that $c'(x') \rightarrow \infty$ as $x' \rightarrow 1$ and the boundedness of marginal benefits, there is an $x' > x$ so that marginal benefits are once again equal to marginal costs. Since the original best-response gave a nonnegative payoff, so must x' (since marginal benefits exceed marginal costs in moving from x to x'). By Assumption 1, the first-order condition is sufficient for a best-response.

On the other hand, if $\overline{BR}(r, \kappa) = 0$ is a strict best response, then by continuity of the value function, it remains so for a slightly higher κ . \square

We continue with two lemmas on the best response correspondence. Thereafter, we will prove the three claims of the theorem.

The first lemma concerns the local comparative statics of the intersection between the graph of $r \mapsto \overline{BR}(r, \kappa)$ and the graph of ρ .

Lemma 11. Take any κ_0 . Consider any (x_0, r_0) satisfying

$$x = \overline{BR}(r, \kappa) \text{ and } r \in \rho(x) \tag{20}$$

with $x_0 > x_{\text{crit}}$ at $\kappa = \kappa_0$. Then there is a neighborhood around κ_0 where $(X(\kappa), R(\kappa))$ is the unique pair satisfying (20), and both X and R are strictly decreasing, continuous functions of κ .

Proof. Restrict attention to a neighborhood of (x_0, r_0) where $x > x_{\text{crit}}$ and $r > r_{\text{crit}}$. Here ρ is the same as $\hat{\rho}$. Note that in the notation of Lemma 6, the conditions (20) are equivalent to

$$\mathfrak{b}(r) = \mathfrak{c}(r) \text{ and } x = \chi(r).$$

The existence of continuous functions $(X(\kappa), R(\kappa))$ extending the solution locally follows by the implicit function theorem applied to this equation. The solution $R(\kappa)$ can be visualized as the horizontal coordinate where the decreasing function $\mathfrak{b}(r)$ in Figure 14(b) intersects the increasing function $\mathfrak{c}(\chi(r) - \underline{x})$. As we perturb κ upward, the curve $\mathfrak{b}(r)$ moves up pointwise. Then $R(\kappa)$, the coordinate of the intersection, is increasing in κ . The statement about $X(\kappa)$ follows because X is an increasing function of $R(\kappa)$ in our neighborhood of interest. \square

The following lemma concerns the convergence of the maximal best-response function $\overline{BR}(r, \kappa, \mu_\tau)$ in the finite- τ model to the limit model.

Lemma 12. $\overline{BR}(r, \kappa, \mu_\tau)$ converges pointwise to $\overline{BR}(r, \kappa)$ for any r .

Proof of Lemma 12. When it is nonzero, $\overline{BR}(r, \kappa, \mu_\tau) = \max\{x : MB(x; r, \kappa, \mu_\tau) = MC(x)\}$ where

$$MB(x; r, \kappa, \mu_\tau) = \kappa g(r) m n (1 - xr)^{n-1} (1 - (1 - xr)^n)^{m-1} r \sum_{d=1}^{\infty} \mu_\tau(d).$$

Suppose first (passing to a subsequence if necessary) we have a sequence where these conditions hold.

Note that, since μ_τ puts mass at least $1 - \frac{1}{\tau}$ on $[\tau, \infty)$, it follows that $\sum_{d=1}^{\infty} \mu_\tau(d) \rightarrow 1$ as $\tau \rightarrow \infty$.

It follows that, as $\tau \rightarrow \infty$, we have the following convergence uniformly:

$$MB(x; r, \kappa, \mu_\tau) \rightarrow \kappa g(r) m n (1 - xr)^{n-1} (1 - (1 - xr)^n)^{m-1} r = MB(x; r, \kappa)$$

and thus that

$$\overline{BR}(r, \kappa, \mu_\tau) \rightarrow \overline{BR}(r, \kappa)$$

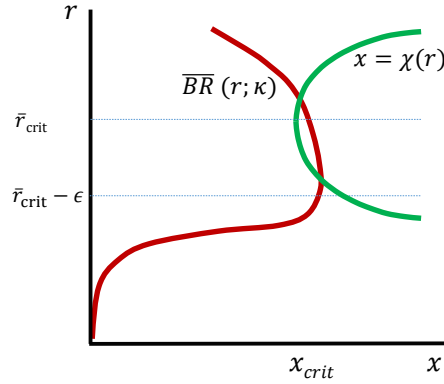


FIGURE 15. The curves $\overline{BR}(r, \kappa)$ and $x = \chi(r)$. If the two were tangent at $r = \bar{r}_{crit}$ at $\kappa = \underline{\kappa}$, then by perturbing to a slightly higher κ , we would obtain two intersections above $r = \bar{r}_{crit} - \epsilon$, contradicting Lemma 6.

for any r .

When $\overline{BR}(r, \kappa, \mu_\tau)$ is identically zero for each τ (again by passing to a subsequence if necessary), it follows from uniform convergence of the firms' value functions to their limit (19) that investing zero strictly dominates any positive level in the limit problem as well. \square

We now define some key quantities that will play a role in our proof of the theorem. It is straightforward to see by inspection of the limit maximand (19) and condition (OI) that for any $r \in (0, 1]$ and any $x_0 < 1$ if κ is large enough, all elements of $BR(r, \kappa)$ satisfy $x > x_0$. In particular, we can guarantee that they all lie above x_{crit} . What we have said implies that there is an κ_1 such that for $\kappa > \kappa_1$ the graphs of $\overline{BR}(r, \kappa)$ and the graph of ρ (both viewed as sets of points (x, r)) intersect above x_{crit} . On the other hand, at sufficiently small κ , all elements of $BR(r, \kappa)$ are bounded by a small number. Let $\underline{\kappa} > 0$ be the smallest κ such that the graph of $\overline{BR}(r, \kappa)$ intersects the graph of ρ at a nonzero value of x . What we have said implies that $\underline{\kappa}$ is a finite, nonzero number.

Consider the two graphs just mentioned at $\kappa = \underline{\kappa}$. Lemma 11 implies that any intersection of the two graphs must be such that $x \leq x_{crit}$ (and thus $r \leq \bar{r}_{crit}$), since otherwise we could find an intersection at a smaller value of κ . Thus, $\underline{\kappa}$ is such that $\overline{BR}(r, \underline{\kappa})$ just touches the correspondence $\rho(x)$ at $x = x_{crit}$. Define \underline{r}_{crit} to be the point of intersection with the highest r , as depicted in Figure 12(c). By Assumption 1, we may assume that locally, $\overline{BR}(r, \underline{\kappa})$ is exactly the set of solutions to (OI) as a function of r .⁵⁹ If it were the case that $\underline{r}_{crit} = \bar{r}_{crit}$, then $r \mapsto \overline{BR}(r, \underline{\kappa})$ would be tangent to $r \mapsto \chi(r)$ at $r = \bar{r}_{crit}$, and by increasing κ slightly we would obtain two intersections between the graphs above $\bar{r}_{crit} - \epsilon$ a contradiction to Lemma 6. Thus $\underline{r}_{crit} < \bar{r}_{crit}$.

We can now prove each part of Theorem 1.

Part (i): For any $\kappa < \underline{\kappa}$, Since $\overline{BR}(0, \kappa) = 0$, the only intersection between $\overline{BR}(r, \kappa)$ and $\rho(x)$ is at $(x, r) = (0, 0)$.

Note that for any τ , $\frac{\partial \rho(x, \mu_\tau)}{\partial x_{if}}|_{x_{if}=x=0} = 0$ and $\overline{BR}(0, \kappa) = 0$. Now, since we know $\rho(x, \mu_\tau)$ converges to $\rho(x)$ (Proposition 5) and $\overline{BR}(r, \kappa, \mu_\tau)$ converges to $\overline{BR}(r, \kappa)$ (Lemma 12), it follows that when $\kappa < \underline{\kappa}$, then there exists $\underline{\tau}$ such that for $\tau > \underline{\tau}$, the curves $\overline{BR}(r, \kappa, \mu_\tau)$ and $\rho(x, \mu_\tau)$ intersect only at a point $(x, r) = (0, 0)$.

Part (ii): Now take any $\kappa \geq \underline{\kappa}$. Since $\overline{BR}(r, \kappa) \geq \overline{BR}(r, \underline{\kappa})$ for any r (with equality holding only when $\kappa = \underline{\kappa}$ or possibly when $\overline{BR}(r, \kappa) = 0$), it follows that there is at least one point of intersection between $\overline{BR}(r, \kappa)$ and $\rho(x)$. We select the one with the highest reliability r . If κ is close enough to $\underline{\kappa}$, then this point will be (x_{crit}, r) with $r \in (\underline{r}_{crit}, \bar{r}_{crit})$.

Since we know $\rho(x, \mu_\tau)$ converges to $\rho(x)$ and $\overline{BR}(r, \kappa, \mu_\tau)$ converges to $\overline{BR}(r, \kappa)$, it then follows that for any $\epsilon > 0$, there exists $\underline{\tau}$ such that for $\tau > \underline{\tau}$, $\overline{BR}(r, \kappa, \mu_\tau)$ and $\rho(x, \mu_\tau)$ intersect at some point $x \in [x_{crit} - \epsilon, x_{crit} + \epsilon]$ and $r \in [\underline{r}_{crit} - \epsilon, \bar{r}_{crit} + \epsilon]$.

⁵⁹If the best response is indifferent to zero investment, then this indifference can be broken by changing the cost function slightly at inframarginal costs without affecting the equilibrium.

As κ keeps increasing, $\overline{BR}(r, \kappa)$ also increases and thus the point of intersection with the highest reliability will reach $(x_{\text{crit}}, \bar{r}_{\text{crit}})$ when κ reaches some value $\bar{\kappa}$. Since we know $\rho(x, \mu_\tau)$ converges to $\rho(x)$ and $\overline{BR}(r, \kappa, \mu_\tau)$ converges to $\overline{BR}(r, \kappa)$, by the same argument as before, it then follows that for any $\epsilon > 0$, there exists $\underline{\tau}$ such that for $\tau > \underline{\tau}$, $\overline{BR}(r, \kappa, \mu_\tau)$ and $\rho(x, \mu_\tau)$ intersect at some point $x \in [x_{\text{crit}} - \epsilon, x_{\text{crit}} + \epsilon]$ and $r \in [\bar{r}_{\text{crit}} - \epsilon, \bar{r}_{\text{crit}} + \epsilon]$.

Part (iii): Finally, as κ increases beyond $\bar{\kappa}$, $\overline{BR}(r, \kappa)$ keeps increasing and the point of intersection between $\overline{BR}(r, \kappa)$ and $\rho(x)$ with the highest reliability r will increase along the part of the $\rho(x)$ curve for which $x > x_{\text{crit}}$, by Lemma 11. At any such point, we also have $r > \bar{r}_{\text{crit}}$.

Since $\rho(x, \mu_\tau)$ converges to $\rho(x)$ and $\overline{BR}(r, \kappa, \mu_\tau)$ converges to $\overline{BR}(r, \kappa)$, this also holds for any τ large enough. This completes the proof. \square

B.4.1. Proof of Corollary 1 (Comparative statics in baseline institutional quality). We first consider the limit model. Equilibria are given by the highest intersection of the reliability curve with the best response curve. The reliability curve is constant as both κ and \underline{x} change. The conditions for optimal investment that must be satisfied in the equilibrium $x^* := x^*(\kappa', \underline{x})$ are that

$$MB(x^*) = \frac{1}{\kappa'} c'(x^* - \underline{x}).$$

It needs to be shown that there exists an $\underline{x}' \in (\underline{x}, x^*)$ such that

$$MB(x^*) = \frac{1}{\kappa} c'(x^* - \underline{x}').$$

Note that $c'(x^* - \underline{x}')$ is continuous in \underline{x}' holding x^* fixed, with

$$\frac{1}{\kappa} c'(0) = 0 < \frac{1}{\kappa'} c'(x^* - \underline{x}) = MB(x^*) \quad \text{and} \quad \frac{1}{\kappa} c'(x^* - \underline{x}) > \frac{1}{\kappa'} c'(x^* - \underline{x}) = MB(x^*).$$

The result then follows from the intermediate value theorem.

The extension to the limit as $\tau \rightarrow \infty$ is straightforward and analogous to other arguments. \square

B.4.2. Proof of Proposition 3 (Equilibrium fragility). When $\kappa \in [\underline{\kappa}, \bar{\kappa}]$:

From Theorem 1, for any $\epsilon > 0$, there exists $\underline{\tau}$ such that for any $\tau > \underline{\tau}$, $x^*(\mu_\tau) \in [x_{\text{crit}} - \epsilon, x_{\text{crit}} + \epsilon]$. Thus for a shock of size 2ϵ , the relationship strength after the shock is $\underline{x} - \epsilon + \xi^*(\mu_\tau) \leq x_{\text{crit}} - \epsilon$.

From Proposition 5, for any $\eta > 0$, there exists $\underline{\tau}'$ such that for all $\tau > \max\{\underline{\tau}, \underline{\tau}'\}$, $\rho(\underline{x} - \epsilon + y^*(\mu_\tau), \mu_\tau) < \eta$.

Thus, when $\kappa \in [\underline{\kappa}, \bar{\kappa}]$ the equilibrium is fragile.

When $\kappa > \bar{\kappa}$:

From Theorem 1, $x^*(\mu_\tau)$ converges to $x^* > x_{\text{crit}}$ as $\tau \rightarrow \infty$. Thus, for any $\epsilon > 0$, there may *not* exist a $\underline{\tau}$ such that $x^*(\mu_\tau) - \epsilon = \underline{x} - \epsilon + y^*(\mu_\tau) < x_{\text{crit}}$ for all $\tau > \underline{\tau}$ and from Proposition 5, for any η , there may *not* exist a $\underline{\tau}'$ such that $\rho(\underline{x} - \epsilon + y^*(\mu_\tau), \mu_\tau) < \eta$ for all $\tau > \max\{\underline{\tau}, \underline{\tau}'\}$.

Thus, when $\kappa > \bar{\kappa}$ the equilibrium is robust. \square

APPENDIX C. MICROFOUNDATIONS

In this appendix we use a canonical production network model with monopolistic competition to microfound the functionality of firms, firms' profit functions and the planner's objective function—i.e., Properties A, B, B', and C. Our goal is to keep this part of the model standard, in order to put the focus on the structure of the underlying supply network. The proofs in Appendix B above depend on the structure of production only through the properties just mentioned.

We first study production after investment has occurred and a given realized supply network is in effect; we then discuss the modeling of investment in a prior stage.

C.1. Production. In this subsection we provide a model of production that microfound the functionality of firms.

C.1.1. Intermediate and final versions of each variety. We let production of any variety $v \in \mathcal{V}_i$ be used in one of two ways. First it can be transformed into an intermediate good version, which is usable only by those varieties v' such that $v \in PS_i(v')$ —i.e., the varieties v' for which v is a potential supplier. This can be interpreted as a costless transformation made possible by the supply relationship with v' that makes v suitable for use by v' . Alternatively, v can be converted costlessly into a different, consumption good version,

denoted \underline{v} . (As we will discuss later, this transformation technology is owned by a particular firm, which earns rents from selling this differentiated consumption good.)

C.1.2. Quantities and production functions. Suppose v procures for its production $z_{v,v'}$ units of the variety $v' \in S_j(v)$. For a given required input $j \in I(i)$, let $z_{v,j}$ be the total amount of j sourced by v , summing across all of v 's suppliers for this input, and write \mathbf{z}_v for the vector of all these quantities associated with variety v . Let ℓ_v be the amount of labor used by variety v . Labor is the only factor (i.e., unproduced good) and it is inelastically supplied, with $\bar{\ell} = 1$ unit of it. The output of v is

$$\phi_v = f(\ell_v, \mathbf{z}_v) := (\ell_v)^{\varepsilon_\ell} \prod_{j \in I(i)} (z_{v,j})^{\varepsilon_z}, \quad (21)$$

where $\varepsilon_\ell + |I(i)|\varepsilon_z = 1$ so that there are constant returns to scale. Thus, all varieties in any set $S_j(v)$ are perfect substitutes. Concerning substitutability *across* varieties, note that production is not possible if one of the inputs cannot be sourced, but on the intensive margin different inputs are substitutable.⁶⁰

Let $q_{\underline{v}}$ be the quantity of \underline{v} consumed. The household consumes aggregate gross production less investments in reliability which are paid for in the aggregated final goods. Aggregate gross production of a *consumption good* is given by

$$Y = \left(\int_{\mathcal{V}} (q_{\underline{v}})^{\eta_C} d\underline{v} \right)^{1/\eta_C}, \quad (22)$$

a demand aggregator with $\frac{1}{2} < \eta_C < 1$.⁶¹ Y_C of the consumption good is allocated to consumption, while Y_I is allocated to investment. Household utility is Y_C . Given investments, efficiency corresponds to maximizing Y .

C.1.3. Functionality. Equation (21) implies that variety $v \in \mathcal{V}_i$ is able to produce if and only if it can procure quantities $(z_{v,v'})_{v,v' \in \mathcal{V}}$ from its suppliers such that $z_{v,j} > 0$ for all $j \in I(i)$. This will be possible if and only if it has at least one link in the realized supply network to a functional supplier for each input it requires—in other words, if and only if it is functional (as defined in Section 2.3). Note that depth-zero nodes have the same production functions as any other nodes; the only difference is that they are not reliant on the operation of specific supply relationships, and always have a functional firm to source their inputs from.

C.2. Equilibrium of the supply network given relationship strengths. We now define a competitive equilibrium and study the structure of it on a realized supply network. Given a realized supply network, functional firms choose how much of each variety of each input to source (taking prices as given) and what price to set for its consumer good, while the representative consumer supplies a unit of labor inelastically and chooses how much of each consumer good variety to buy (taking prices as given). We start by considering the consumer's problem, and then consider firms' problems before using these to define an equilibrium.

C.2.1. Equilibrium: Definitions. Let w denote the wage; p_v the price of variety v when used as an intermediate; and \underline{p}_v the price paid by the consumer for the final good corresponding to variety v . The numeraire is the price paid by the consumer for the final output.

As Y_I , the amount of final good devoted to investment, is sunk at the production stage, the household's problem is equivalent to choosing final good consumptions $(q_{\underline{v}})_{\underline{v}}$ to maximize equation (22) subject to the budget constraint

$$\int_{\mathcal{V}} \underline{p}_v q_{\underline{v}} d\underline{v} \leq w + \int_{\mathcal{V}} \Pi_v d\underline{v}, \quad (23)$$

where the right hand side is labor income from the consumer's inelastically supplied unit of labor at wage w and income from the profits Π_v of all firms (including those firms that are not functional). Let $q^*(\underline{p})$ denote the unique input bundle maximizing final good production, where \underline{p} is a vector of the final good prices set by

⁶⁰We view the assumption that all inputs must be sourced for successful production as mild. Indeed, Boehm, Flaaen, and Pandalai-Nayar (2019) estimate, in the context of a supply network disruption event, a production function with low elasticity of substitution, closer to Leontief than Cobb–Douglas.

⁶¹The same results would hold if we defined a suitable nested-CES aggregator, with categories of products of equal measures being the nests. The upper bound on η_C ensures that the elasticity of substitution between varieties is finite (so that variety is valuable), while the lower bound ensures that the equilibrium returns to increasing variety are concave, which is technically convenient; see the discussion after Definition 1.

different varieties, and let $q_v^*(\underline{p})$ denote the corresponding amount of variety v demanded in the production of the final good.

Each firm takes the prices of all intermediate goods as given, as well as the final good prices of other varieties. Prices for intermediate goods are competitive, i.e., equal to marginal costs of production.⁶² Our model is one of monopolistic competition in final goods: a functional firm's problem is to choose input quantities $(z_{v,v'})_{v,v' \in \mathcal{V}}$ from its functional suppliers in the realized supply network, labor demand ℓ_v and a price for its consumer good v to maximize its profits.⁶³

It is convenient to break this problem into two steps. First, for any quantity a firm produces it must choose its inputs to minimize its cost of production. As there are constant returns to scale, this requires choosing inputs to minimize the cost of producing one unit and then scaling these choices up or down to meet the production target. Thus the relative amounts of different inputs and labor used must solve

$$\min_{\mathbf{z}_v, \ell_v} \sum_j \sum_{v' \in S_j(v)} p_{v'} z_{v,v'} + w \ell_v \quad \text{subject to } f(\ell_v, \mathbf{z}_v) = 1 \quad (24)$$

Let $z_{v,v}^*$ denote the quantity of variety v' that v chooses to source per unit of its output, and ℓ_v^* the amount of labor that v sources per unit of its output.

Given the unit cost of production generated by solving (24), each variety v sets a price for its consumer good that maximizes its profits (taking reliability investment costs as sunk) and hence solves

$$\max_{p_v} q_v^*(\underline{p}) [p_v - \sum_{v'} p_{v'} z_{v,v'}^* - w \ell_v^*] \quad (25)$$

Given a realized supply network $\mathcal{G}' = (\mathcal{V}, \mathcal{E}')$, with functional firms \mathcal{V}' , a competitive equilibrium is given by a specification of intermediate good production $z_{v,v'}$ for all $(v, v') \in \mathcal{E}'$, final good production q_v for all $v \in \mathcal{V}'$, wage w , intermediate good price p_v for all $v \in \mathcal{V}'$ and final good price p_v for all $v \in \mathcal{V}'$ such that the following conditions hold:

- the household is choosing how much to consume of each consumer good variety given the prices of these varieties to maximize its utility (22) subject to its budget constraint (23);
- all functional varieties $v \in \mathcal{V}'$ are choosing inputs in ratios that minimize their unit cost given input prices (solve (24));
- all functional varieties $v \in \mathcal{V}'$ are choosing consumer good prices p_v that maximize their profits given other consumer good prices (solve (25));
- markets clear ($\phi_v = \sum_{v' \in \mathcal{V}'} z_{v,v'} + q_v^*$ for all $v \in \mathcal{V}'$).

C.2.2. Equilibrium: Analysis. Equilibrium is characterized by all firms producing the same quantities of their respective final goods and selling these at the same price. While they affect the propagation of production failures, conditional on being functional local network features do not alter equilibrium prices and quantities. This can be seen by analyzing the prices of different intermediate goods. First, because of constant returns to scale, firms sell their intermediate goods at a price equal to marginal cost. It turns out that within an industry each functional variety is subject to the same intermediate costs, and hence sells at the same intermediate price regardless of their depth. Second, because of the regularity of the supply network, a symmetry exists across industries which results in them selling at the same intermediate prices. These results are shown formally in Lemmas 13 and 14 respectively. In Lemma 15, we show these results imply that all firms sell the same amounts to consumers at the same final prices.

Lemma 13. The selling prices p_v of the intermediate goods of all varieties $v \in \mathcal{V}_i$ within any industry i are equal to some common price p_i .

⁶²Thus, the intermediate of variety v has a single price irrespective of who buys it. This pricing assumption amounts to requiring efficient production, with no distortions even when a supplier has (in a given realization) market power over a buyer. As documented by (Uzzi, 1997; Kirman and Vriend, 2000), avoiding hold-up is an important function of relational contracts in practice.

⁶³As in standard monopolistic competition models, a firm commits to a price p_v for its consumption good v before producing. These goods are sold at a markup above marginal cost; the quantity is determined by consumer demand at that price. As there are constant returns to scale, the amount of intermediate goods produced has no bearing on production costs for the final goods, and vice versa.

Proof. As there are constant returns to scale, in equilibrium each intermediate good must be sold at a price equal to their respective marginal costs of production. Thus, by constant returns to scale, the price of each variety's intermediate good, p_v , does not depend on the quantity firm v produces. We establish the conclusion of the lemma by induction on the depth of the variety v . If $d(v) = 0$, firm $v = (i, f)$ can source from any firm in each industry $j \in I(i)$ and can thus minimize its marginal costs. It is easy to deduce from this that $p_v = \min_{f \in [0,1]} p_{(i,f)}$. Now take a variety v of depth $d(v) = d$; by definition v sources its intermediate goods from firms with depth $d - 1$ in each industry j . If these produce at minimal cost so it follows that $p_v = \min_{f \in [0,1]} p_{(i,f)}$. This gives us that all firms in the same industry have the same marginal costs and thus the same intermediate prices, regardless of their depth. \square

Lemma 14. For any two industries i and j their intermediate prices are equal: $p_i = p_j$.

Proof. From Lemma 13, there is a price p_i such that all firms in industry i sell intermediate goods at price p_i . By constant returns to scale, p_i is not a function of the quantity produced and is only a function of the prices of the intermediates p_j where $j \in I(i)$. Moreover, each industry requires the same number of inputs and combines these with labor (at common wage w). Thus marginal cost of production for a firm and hence its price can be expressed as a function of just its intermediate good input prices p_j for $j \in I(i)$ and the wage w and by symmetry this is the same function for all industries. We denote this function by \mathfrak{P} ; it takes as its input m intermediate good prices and the wage. The function is invariant to permutations of intermediate good prices. We can construct a vector of input prices for industry i as $\mathbf{M}_i \mathbf{p}$, where \mathbf{M}_i is a $m \times N$ matrix such that for each row r , the entry $M_{r,k} = 0$ for all k but one which we label $j_i(r)$. This index satisfies $j_i(r) \in I(i)$; $\mathbf{M}_{r,j_i(r)} = 1$; and $\mathbf{M}_{r',j_i(r)} = 0$ for $r' \neq r$. (The matrix simply selects, in each row, one of the inputs of industry i , and the corresponding entry of $\mathbf{M}_i \mathbf{p}$ corresponds to its price.) Then, for all industries i ,

$$p_i = \mathfrak{P}(\mathbf{M}_i \mathbf{p}; w),$$

The function \mathfrak{P} is weakly increasing in each input price and strictly increasing in each intermediate input price p_j for $j \in I(i)$.

Now assume, for the sake of contradiction, that not all intermediate input prices are the same. Without loss of generality, let $p_a = \max\{p_1, \dots, p_N\}$ and $p_b = \min\{p_1, \dots, p_N\}$ with $\frac{p_a}{p_b} = r > 1$. By definition, the price of each intermediate for industry a must be at most p_a , and similarly the price of each input for industry b must be at least p_b . Thus each element of $\mathbf{M}_a \mathbf{p}$ must be less than or equal to each element of $r \mathbf{M}_b \mathbf{p}$. Now, we claim

$$p_a = \mathfrak{P}(\mathbf{M}_a \mathbf{p}; w) \leq \mathfrak{P}(r \mathbf{M}_b \mathbf{p}; w) < r \mathfrak{P}(\mathbf{M}_b \mathbf{p}; w) = r p_b = p_a.$$

To show the strict inequality, note that strictly increasing all intermediate prices relative to the wage leads to a substitution toward labor, making the increase in unit costs less than proportional. This is a contradiction. Thus, all industries must sell intermediate goods at the same price. \square

Lemma 15. For any two functional firms, $v, v' \in \mathcal{V}'$, the prices of these firms' final goods are equal and the same amount of each good is sold: $p_v = p_{v'}$ and $q_v^*(p) = q_{v'}^*(p)$.

Proof. Recall that $q_v^*(p)$ denotes the demand of the consumer for final good v at prices p . We first calculate the price elasticity of this demand for final good v .

The consumers maximize (22) subject to the constraints of (23). Using the method of Lagrange multipliers to solve for the quantity of each variety demanded in terms of prices gives that for any two varieties v, v' ,

$$\frac{q_v^*(p)}{p_v^{1/(\eta_C - 1)}} = \frac{q_{v'}^*(p)}{p_{v'}^{1/(\eta_C - 1)}} \quad (26)$$

Because there is a continuum of firms, changing the price level and quantity of firm v will not affect the right-hand side of equation (26). Thus, the left-hand side of equation (26) is a constant function of p_v . Upon differentiating this function with respect to p_v , we get

$$\frac{\partial q_v^*}{\partial p_v} \cdot p_v - \frac{q_v^*}{\eta_C - 1} = 0$$

Rearranging gives us that the price elasticity of demand for variety v at any price level is $\frac{1}{\eta_C - 1}$.

Now consider a firm's pricing problem. Given an intermediate good price p (which by Lemmas 13 and 14 is constant for all intermediate goods), there is a unique solution to the cost minimization problem of each firm (Equation (24)).

This marginal cost of production is independent of the quantity produced and is faced by all firms in each industry. Denote this marginal cost c . At a solution to this problem the firm will be charging a price p_v that satisfies the Lerner condition:

$$\frac{p_v - c}{p_v} = -\frac{1}{\epsilon_d(p_v)}, \quad (27)$$

where ϵ_d is the price elasticity of demand as a function of final good prices. Rearranging equation (27) and substituting the previously calculated elasticity, we have

$$\eta_C p_v = c \quad (28)$$

It is clear that the left-hand side is injective in p_v and there is a unique profit-maximizing price that all firms will choose. Additionally, as the price elasticity of demand is constant, equation (27) implies that all firms will charge the same constant markup regardless of the quantity of goods they are producing. \square

We have shown that in equilibrium: (i) each functional variety v has the same consumer good output q_v , which we will call q ; (ii) these goods are all priced at the same price p_v ; (iii) all intermediate goods of variety $v \in \mathcal{V}$ have the same price $p_v = p$. There are two other features of the equilibrium that are crucial:

Lemma 16.

- (1) Gross output is equal to $h(\rho(x, \mu))$, where $h : [0, 1] \rightarrow \mathbb{R}_+$ is an increasing and continuous function with bounded derivative and $h(0) = 0$.
- (2) The expected gross profit of producing each variety, conditional on that variety being functional, is equal to $g(\rho(x, \mu))$; where $g : [0, 1] \rightarrow \mathbb{R}_+$ is a decreasing function.

Proof. Note that because labor is the only unproduced input, all labor can be assigned to the production of final goods by allocating labor used to produce intermediate goods to the final good that this intermediate good is ultimately used to produce. Let L_v be the total labor assigned in this way to the production of variety v 's final good. As by Lemmas 14 and 15 all firms have equivalent production functions, face the same input prices and choose the same final good prices, they will use the same ratio of different inputs in production. This symmetry throughout the supply network implies that $L_v = L$ for all varieties v . (This does not depend on depth because the zero depth producers have production functions identical to those of other firms and buy inputs—they can just do so from any variety without needing a specific relationship.) As there is a total supply of labor equal to one and it is supplied inelastically, the amount of labor that can be assigned to the production of any given final good is $1/\lambda(\mathcal{V}')$, where $\lambda(\mathcal{V}')$ is the measure of functional firms (and hence equal to $\rho(x, \mu)$ as the measure of firms is 1). As L units of labor are required to produce one unit of any final good, this implies that

$$q_v^* = q^* = \frac{1}{L\lambda(\mathcal{V}')} \quad Y = (\lambda(\mathcal{V}')(q^*)^{\eta_C})^{1/\eta_C} = \frac{\lambda(\mathcal{V}')^{1/\eta_C-1}}{L} = \frac{\rho(x, \mu)^{1/\eta_C-1}}{L} \quad (29)$$

Setting

$$h(\rho(x, \mu)) = \frac{\rho(x, \mu)^{1/\eta_C-1}}{L},$$

completes the proof of part (1).

It can be computed that L depends only on the constants in the production functions, and not on x , which establishes (3).⁶⁴

To compute the wage w in terms of the numeraire, we express the aggregate expenditure on final goods in two different ways. On the one hand as the price of the aggregate consumer good has been normalized to 1, it is simply the quantity of aggregate output as given by equation (29). On the other hand, it is the expenditure of the consumer on all firms' final goods, which we simplify using equation (28), and noting that as all inputs are sold at cost all firms have a marginal cost of production equal to wL :

$$p_v \int_{\mathcal{V}} (q_v) dv = \rho(x, \mu) \frac{wL}{\eta_C} \frac{1}{\rho(x, \mu)L} = \frac{w}{\eta_C}$$

⁶⁴Let ℓ be the amount of labor firm v directly hires to produce a unit of its final good. Let z be the amount of each intermediate firm v uses from one of its suppliers to produce a unit of its final good. By symmetry and constant returns to scale, this firm minimizes its unit cost in terms of labor, $L = \frac{\ell}{1-mz}$ subject to the constraint that $1 = z^{m\epsilon_z} \ell^{\epsilon_\ell}$. Solving the Lagrangian yields that for some constant γ , we have $\frac{\epsilon_z}{z} = \frac{\gamma\ell}{(1-mz)^2}$ and $\frac{\epsilon_\ell}{\ell} = \frac{\gamma}{1-mz}$. Solving this yields that at the optimum $z = \epsilon_z$, that $\ell = \epsilon_z^{-m\epsilon_z/\epsilon_\ell}$ and $L = \frac{\epsilon_z^{-m\epsilon_z/\epsilon_\ell}}{\epsilon_\ell}$.

Combining these two expressions for expenditure implies that

$$w = \left(\frac{\eta_C}{L}\right) \rho(x, \mu)^{\frac{1-\eta_C}{\eta_C}}.$$

From equation (28) each firm sets a final good price of $p_v = wL/\eta_C$. Thus each firm earns gross profits

$$(p_v - c) q^* = w \left(\frac{1-\eta_C}{\eta_C}\right) \frac{1}{\lambda(\mathcal{V}')} = \left(\frac{\eta_C}{L}\right) \rho(x, \mu)^{\frac{1-\eta_C}{\eta_C}} \left(\frac{1-\eta_C}{\eta_C}\right) \frac{1}{\rho(x, \mu)} = \left(\frac{1-\eta_C}{L}\right) \rho(x, \mu)^{\frac{1}{\eta_C}-2}. \quad (30)$$

This directly implies the second part of the lemma, using the fact that $\eta_C \in (\frac{1}{2}, 1)$. \square

Lemma 16 provides microfoundations for Properties A and C. Setting $Y(\mathcal{V}') = h(\rho(x, \mu))$, part (1) shows that maximized aggregate output $Y(\mathcal{V}')$ satisfies Property A. Part (2) shows that the expected gross profit function takes the form $g(r)$ assumed in the model, with $g(r)$ satisfying Property C—i.e., that $g(r)$ is a decreasing, continuously differentiable function. For these statements it is important that L depends only on constants, not on x , which we have shown.

Part (2) of Lemma 16 follows from the household's love of variety. Because of the love of variety that households have, the profit maximizing price that firms set for their consumer goods is at a markup over marginal cost. On the other hand, constant returns to scale in the production function means that intermediate goods that can be produced (i.e., which are able to source all of their required inputs) must be priced at marginal cost in equilibrium. So a firms' profits depend just on the markup they charge to consumers, and the quantities they sell to consumers. When more other varieties are functional the consumer's love of variety reduces demand for a given variety and hence the profits obtained by that variety.

C.3. Endogenizing relationship strengths: Efficient and equilibrium outcomes. Now we turn to the earlier stage, before production, where firms invest in relationship strengths. Here our goal is to show how the reduced-form modeling of investment costs that we presented in the main text (Section 4.2.2) fits into the general equilibrium model.

C.3.1. The technology of investment. There is a technology available to each firm to invest the consumption good and increase relationship strengths. Before production, each firm if pledges some amount of the consumption good to invest. We posit that it costs a firm $\frac{1}{\kappa}c(x_{if} - \underline{x})$ units of the numeraire (the consumption good) to achieve relationship strength x with its potential suppliers, where c is a given function—an exogenous technology of investment.

Thus, a firm's expected profit when it is making its investment decision is

$$\Pi_{if} = P(x_{if}; x, \mu)g(r) - \frac{1}{\kappa}c(x_{if} - \underline{x});$$

here both expected revenues and costs are in units of the numeraire. This microfound the form of the firm profits we have assumed.

C.3.2. Planner's problem. We study the problem of efficiently choosing the investments in the first stage. The planner may choose any symmetric investments x for the firms. For a fixed choice of relationship strengths x , let $h(\rho(x, \mu))$ denote the gross production of the aggregated consumption good. As we have said, household consumption is $Y_C = Y - Y_I$, where Y_I is the value of the investment in relationship strength in terms of the quantity of the consumption good devoted to it. These are costs that are sunk prior to production. Obtaining reliability x for all such relationships costs the planner

$$Y_I = \sum_{i \in \mathcal{I}} \int_{v \in \mathcal{V}_i} \frac{1}{\kappa}c(x - \underline{x})dv = \frac{1}{\kappa}c(x - \underline{x}).$$

(Recall that the measure of the set of all varieties is $\lambda(\mathcal{V}) = 1$.) Setting $c_P(x) = c(x - \underline{x})$ the planner's cost function inherits the key properties of the individual cost function (it is continuous, increasing, and weakly convex with $c(0) = 0$, $c'(0) = 0$ and $\lim_{x \rightarrow 1} c(x) = \infty$ by Property B') and hence satisfies the properties assumed in Property B.

Maximizing $Y_C = Y - Y_I$ amounts to our planner's problem from the main text:

$$\max_{x \in [0,1]} h(\rho(x, \mu)) - \frac{1}{\kappa}c_P(x)$$