

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. Due to the risk of such supply relationships being idiosyncratically disrupted, firms multisource inputs and strategically invest to make relationships with suppliers stronger. Aggregate production is robust to idiosyncratic disruptions. However, there is a regime in which equilibrium supply networks are fragile, with small aggregate shocks to relationships causing arbitrarily steep drops in output. The endogenous configuration of supply networks provides a new channel for the powerful amplification of shocks.

Date Printed. October 8, 2020.

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 Insititute, 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, and Yixi Jiang 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, Juuso Välimäki and Rakesh Vohra.

We dedicate this paper to the memory of Emmanuel Farhi.

1. INTRODUCTION

Complex supply networks among firms 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 investing in their supply networks, trading off the gains in the robustness of their production against the cost of maintaining good supply relationships. Our main purpose is to examine the robustness of equilibrium supply networks. We find that (i) the economy is robust to idiosyncratic risk; yet (ii) small shocks that systematically affect the functioning of supply relationships are massively amplified; (iii) the functioning of many unrelated supply chains is highly correlated; and (iv) the complexity of production is key to the level of aggregate volatility.

Our analysis is built on a model of interfirm sourcing relationships, their disruption by shocks, and the downstream consequences for production. We now motivate our study of these phenomena with some examples. Firms 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. The suppliers must similarly procure tailored inputs, and so on. Thus, a particular airplane producer is exposed not just to shocks in the overall availability of each needed input product, 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. Such idiosyncratic disruptions to individual relationships can have far-reaching effects, causing damage that cascades through the supply chain and affects many downstream firms.² In Section 6, we discuss evidence that such disruptions are practically important, and can be very damaging to particular firms.

Our ultimate goal is to understand the robustness of such supply networks. In particular, we want to study how idiosyncratic and aggregate shocks to the functioning of relationships affect aggregate output and welfare. We first describe how the probability of idiosyncratic disruptions discussed in the previous paragraph shapes aggregate outcomes. Idiosyncratic realizations—whether a *particular* supply chain functions or not—average out.³ The aggregate statistic of interest is the fraction of supply chains functioning, which we call the *reliability* of the supply network. Reliability depends on aggregate parameters in a way we explicitly characterize. Some of these parameters are features of the supply network determined by technology and taken to be exogenous—e.g., how many produced

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.

³This is by a standard diversification argument. In our model, 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).

inputs a typical product requires. In contrast, the probability with which any given relationship functions—which we call *relationship strength*—is the main endogenous firm decision in our model. Since production is risky, an optimizing firm will strategically choose its relationship strength to manage this risk. By investing more, a firm can increase the probability that its potential suppliers are able to supply it with the inputs it needs, hence allowing the firm to produce its output and make profits. Airbus, for example, has two engine suppliers for its A380 airplanes. 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.⁴

Once we have a model of a supply network with equilibrium investments in relationship strength, we examine how robust it is to aggregate shocks that harm many relationships at once. We give several examples of the kinds of shocks we have in mind. First, suppose that the institutions that help uphold contracts and facilitate business transactions suddenly decline in quality. Each supply relationship then becomes more prone to the idiosyncratic disruptions discussed in the previous paragraph.⁵ For a second concrete interpretation of the kinds of aggregate shocks we have in mind, consider a small 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, for the 2019-20 outbreak of the novel coronavirus, Covid-19, there is much uncertainty about how different supply relationships might be affected, and this can be modeled as a systematic decrease in the probability that suppliers will be able to deliver the inputs required from them.

Our results show, first, that at certain configurations of the aggregate parameters, a small shock to the strength of relationships can decrease the reliability of the supply network from a large, positive level to 0. There is a *precipice* where small aggregate shocks to relationship strength unravel production in many unrelated supply chains simultaneously. The key question is then whether equilibrium supply networks will be on such a precipice. Our main findings give conditions for this to occur, and show that the precipice is *not* a knife-edge outcome. Indeed, we characterize the open 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 somehow by averaging and endogenous investment. We find the opposite: in aggregating up the uncertainty through the interdependencies of the supply network, we get a very sharp sensitivity of aggregate productivity to relationship strength. This is in contrast to many standard production network models, where the

⁴Strategic 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).

⁵Indeed, Blanchard and Kremer (1997) present evidence that the former Soviet Union suffered this kind of shock when it transitioned to a market-based economy; they argue that this contributed to the subsequent drop in output.

aggregate production function is differentiable at any point. 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 the relationships that mediate this sourcing. Jointly, these phenomena create the possibility of precipices, which underlie our analysis.

We now outline the details of our model and results more precisely. There are many products. Each has many differentiated varieties, produced by separate firms that are small relative to the overall supply network. (Each firm is associated with a single variety.) A given product has a set of input products that are essential to its production—e.g., an airplane requires engines, navigation systems, etc. Most varieties of a product have production functions requiring customized inputs—i.e., they can use only some specific, compatible, varieties of their input products. Thus, in the above example, the production function of an Airbus A380 can only use a certain number of engine varieties. These compatibilities determine the *potential supply relationships* of each firm. This aspect of the model captures *multisourcing*—the possibility of relying on any of several, substitutable, sourcing options.

Each of the potential supply relationships of a firm may operate successfully or not (we will specify the distribution shortly). The *realized supply network* consists of the subset of the potential supply relationships that operate. The realizations of which relationships operate determine which varieties are *functional*. In order for a firm’s variety to be functional, the firm must have at least one operating supply relationship to some firm producing each input product. These must in turn satisfy the same condition, and so on—until a point in the supply chain where no customized inputs are required. The firms that are able to produce purchase their required inputs from their suppliers, hire labor, and then sell their output. Each firm’s output is allocated between use as an intermediate good and sale at a profit-maximizing markup to consumers. Here our solution is a standard one—general-equilibrium monopolistic competition. Social welfare is increasing in the number of varieties produced, because of the household’s love of variety. Firms make profits from production.

Our model of how the realized supply relationships are selected from the potential ones is simple: independently, each relationship operates successfully with a probability called the *relationship strength*, whose determination we will discuss shortly.⁶ The first set of results examines aggregate welfare and firm profits when the strength of the relationships is exogenous. 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 each relationship. When production is complex (i.e., most firms have multiple essential inputs they need to source) there is a discontinuity in aggregate productivity: when relationship strength in a certain supply network falls below a certain threshold (defining the precipice), production in that network drops discontinuously to zero. This raises the prospect of fragility: a small, correlated, negative shock to relationship strengths can lead to considerable economic damage. We also show that a social

⁶The interdependence between firms and their suppliers makes failures correlated between firms that (directly or indirectly) transact with each other.

planner will always choose costly relationship strengths so that the supply network is away from a precipice.

A natural question is then whether a supply network will be near a precipice when relationship strengths are determined by equilibrium choices rather than by a planner. To analyze this, we model the incentives of firms that are attempting to produce. Firms invest in their relationships. That is, they choose a level of costly investment toward making relationships stronger—i.e., likelier to operate.⁷ Then shocks are realized, and these shocks determine which relationships can actually operate.

Our main results are on how the sensitivity of production to aggregate shocks depends on an aggregate productivity parameter. An increase in this parameter can be interpreted as an increase in the productivity of the economy relative to the costs of maintaining relationships. 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.

Fragility comes from an interaction between the mechanics of the aggregate production technology and firms' investment incentives. The mechanics of production specify how a given level of relationship strength throughout the supply network translates into equilibrium reliability. The main feature of the mechanical relationship is the precipice: the steep increase of reliability from nearly zero to a high level over a very narrow range of relationship strengths. On the incentive side, for a given level of others' reliability, firms' optimization problems determine what relationship strengths they want to choose. To provide sufficient incentives to support a given relationship strength, we show that reliability must be at a particular intermediate level.⁸ The precipice provides a natural place to find an equilibrium: because reliability goes from a low level to a high level very steeply there, the required intermediate level of reliability often corresponds to a relationship strength on the precipice.

We then explore some implications of our modeling, especially as it relates to heterogeneity within and across supply networks. First, 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.). We find that the basic fragility in production that we identify is not dependent on homogeneity, but it depends critically

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.

⁸This is intuitive: When others' reliability is low there is no point in a firm investing in its own supply relationships as these suppliers will be unable to produce. When the reliability of others is very high, there is again little point in a firm investing a great deal in its supply relationships. As all other firms are very reliable, multisourcing means that a given supply relationship failing is unlikely to matter. The incentives to invest in reliability are highest for intermediate values of others' reliability.

on complexity—the fact that production relies on multiple relationships working. Moreover, 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 show how the supply networks we have studied can be embedded in a larger economy with intersectoral linkages that come from buying inputs without specific sourcing requirements. It turns out that even at a macroeconomic scale, the fragilities we have identified are not smoothed away. Indeed, our model yields a new channel for the propagation of shocks across sectors, and their stark amplification. When we aggregate up the effects and think of the fragile supply networks as one part of an interconnected macroeconomy, the forces we identify provide a complementary perspective on phenomena recently identified in the literature on the amplification of shocks in a macroeconomics setting (Di Giovanni and Levchenko, 2012; Di Giovanni et al., 2018). Third, 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 6.1.4). After presenting our results, in Sections 6 and 7 we discuss in detail how they fit into the most closely related literatures.

2. THE MODEL OF PRODUCTION

2.1. Products and varieties. 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 $v = (i, \omega)$, where $\omega \in \Omega_i \subseteq \mathbb{R}$ is a variety index; unless otherwise noted we take $\Omega_i = [0, 1]$ for all i . Let $\mathcal{V} = \bigcup_{i \in \mathcal{I}} \mathcal{V}_i$ be the union of all the varieties. The quantity of any variety that is produced is allocated between consumption and use as an intermediate good.

2.2. The supply network. To specify the production function of a variety, we first define a supply network that describes production relationships between varieties. First, each product $i \in \mathcal{I}$ is associated with 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 . The measure of varieties with any depth $d \geq 0$ is denoted by $\mu(d)$.

First, consider any variety $v \in \mathcal{V}_i$ that requires specific sourcing—i.e., has depth $d(v) > 0$. For each $j \in I(i)$, the variety v has a finite set of *potential suppliers*, $\text{PS}_j(v) \subseteq \mathcal{V}_j$, corresponding to the fact that the variety v requires specialized sourcing. The set $\text{PS}_j(v)$ is a set of distinct varieties $v' \in \mathcal{V}_j$ with each such v' having depth $d(v') = d(v) - 1$. These are drawn uniformly and independently from the set of varieties v' such that $d(v') = d(v) - 1$ (i.e., the set of varieties of compatible depth). Specialized sourcing requirements represent the need for a customized input, the procurement of which is facilitated by relational contracts. There are thus only finitely many varieties that are compatible with the production process of v .

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 homogeneous across the supply network) which is the probability that any relationship is operational. All these realizations are independent.⁹ The set of actual suppliers $\text{S}_j(v)$ is then obtained by including each potential supplier in $\text{PS}_j(v)$ independently with probability x . Whereas the potential

⁹For 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 input products. 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.)

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.¹⁰

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 $v' \in \bigcup_{j \in I(i)} \text{PS}_j(v)$. (See Figure 1 for an illustration.) In the *realized supply network* \mathcal{G} , each v has links directed to all its (operational) suppliers $v' \in \bigcup_{j \in I(i)} \text{S}_j(v)$. See the links in Figure 2 for an illustration of the subset of supply relationships that remain operational.

It remains to consider varieties of depth 0—the $v \in \mathcal{V}$ such that $d(v) = 0$. These do not require specialized sourcing and can use any variety. Thus we take $\text{S}_j(v) = \mathcal{V}_j$ for such a variety for any $j \in I(i)$.

2.3. Production. We use a canonical monopolistic competition model to determine production and surplus on a given realized supply network.¹¹

2.3.1. Intermediate and final versions of each variety. The output of any variety $v \in \mathcal{V}_i$ can be used in one of two ways. First it can be transformed into an intermediate good version, which is usable only those varieties v' such that $v \in \text{PS}_i(v')$ —i.e., the ones compatible with it. 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 differentiated consumption goods.)

2.3.2. Quantities and production functions. Suppose v procures for its production $z_{v,v'}$ units of the variety $v' \in \text{S}_j(v)$. 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,¹² and write \mathbf{z}_v for the vector of all these

¹⁰In a bit more detail, x 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 agents make. See Section 6 for more details.

¹¹Our goal is to keep this part of the model standard, in order to put the focus on the structure of the underlying supply network.

¹²We only consider allocations where only finitely many of the $z_{v,v'}$ are positive.

quantities associated with variety v . Let ℓ_v be the amount of labor used by variety v . Then the output of v is

$$y_v = f(\ell_v, \mathbf{z}_v) := (\ell_v)^{\varepsilon_\ell} \prod_{j \in I(i)} (z_{v,j})^{\varepsilon_z}, \quad (1)$$

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. But some of each variety, as well as labor, is needed to produce any output of variety v .¹³ The only factor (i.e., unproduced good) in the economy is labor, of which there is an inelastic supply of $\bar{\ell} = 1$ unit.

Let q_v be the quantity of v consumed, and let the household's consumption be

$$C = \kappa \left(\int_{\mathcal{V}} (q_v)^{\eta_C} dv \right)^{1/\eta_C}, \quad (2)$$

a demand aggregator with $\frac{1}{2} < \eta_C < 1$.¹⁴ Household utility is equal to C . The parameter $\kappa > 0$ is a productivity multiplier.

3. EQUILIBRIUM PRODUCTION WITH EXOGENOUS RELATIONSHIP STRENGTHS

3.1. The mass of functional varieties. Given a realized supply network, it is clear that not all varieties can be produced: for example, consider a variety v with depth $d(v) > 0$ that happens to have no (operational) suppliers. Such a variety will have a production function equal to the zero function. We call a variety *functional* if there is some profile of quantities procured, $(z_{v,v'})_{v,v' \in \mathcal{V}}$, consistent with the supply network¹⁵ such that $y_v > 0$.

In this section, we describe which varieties are functional and analyze the size of this set. Because the consumer preferences feature love of variety, the economy is more productive when more varieties are functional. This motivates our analysis of which varieties *can* be produced, given the realized supply network, and what the size of this set is. We show after this that welfare (i.e., total consumption) is indeed monotonically increasing in the number of functional varieties.


For a given realization of the supply network, we will determine which varieties are functional. This can be determined inductively. First, consider the functionality of the depth-0 varieties. As they require no specific sourcing, they will always be functional. 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 $S_j(v)$ is nonempty 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 now study reliability for *regular* supply networks where the number of essential inputs is $|I(i)| = m$ for each product i and each variety of depth $d > 0$ has n potential suppliers for each


¹³The elasticity of substitution is 1 between any two inputs. Boehm, Flaaen, and Pandalai-Nayar (2019) estimate, in the context of a supply network disruption event, a production function with much lower elasticity of substitution, closer to Leontief (zero elasticity of substitution). For our study of disruptions, assuming a higher elasticity of substitution is conservative: we assume more substitution possibilities between varieties than there seem to be in times of disruption.

¹⁴The 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 less than 1 (so that variety is sufficiently 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.

¹⁵I.e., for variety $v = (i, \omega)$ the quantity $z_{v,v'}$ can be positive only if $v' \in \bigcup_{j \in I(i)} S_j(v)$, i.e., only if v' is a realized supplier of v .



(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.

input—i.e. $|\text{PS}_j(v)| = n$ whenever j is one of the required inputs for variety v . We relax this assumption in Section 6.3.5.

Denote by $\tilde{\rho}(x, d)$ the probability that a variety of depth d is functional when relationship strength is x . By symmetry of the supply tree, this probability does indeed only depend on x and d . The main outcome we will focus on is the probability that a variety selected uniformly at random is functional, which we call the *reliability* of the supply network. We denote this by $\rho(x, \mu)$ and define it as

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

recalling that μ is the distribution of depths (a probability measure on the nonnegative integers).

3.1.1. Deep supply networks: Taking limits. A focus throughout will be the case where a typical variety has large depth.¹⁶ 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 $[\tau, \infty)$.¹⁷ For instance, we can take μ_τ to be the geometric distribution with mean τ . For τ large, 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 ($n \geq 2$).

¹⁶Supplementary Appendix SA1 investigates how reliability varies with investment in production trees with bounded depth.

¹⁷Note 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.

3.2. Production equilibrium. Continuing with a fixed x , we study the structure of competitive equilibrium on a realized supply network.

A competitive equilibrium is a tuple, satisfying certain conditions, consisting of the following:

- prices: a price p_v for each intermediate variety $v \in \mathcal{V}$; a price $p_{\underline{v}}$ for the corresponding consumption good; and a wage w .
- quantities: $z_{v,v'}$ for each link in the realized supply network, as well as $q_{\underline{v}}$ for each consumption good $\underline{v} \in \mathcal{V}$.

The equilibrium conditions are that all markets clear; that production of each variety $v \in \mathcal{V}$ maximizes profits given prices; that prices $p_{\underline{v}}$ of consumer goods are set to maximize profits given all other prices; and that the consumer maximizes her utility (by spending all labor income on consumption).

3.2.1. Analysis of the production equilibrium. We briefly describe the structure of equilibrium, deferring the full reasoning behind these assertions to Appendix B.

First, in equilibrium: (i) each functional variety v has the same consumer good output $q_{\underline{v}}$, which we call \underline{q} ; (ii) these goods are all priced at some price \underline{p} ; (iii) all intermediate goods of variety $v \in \mathcal{V}$ have the same price $p_v = p$. These facts follow from the symmetry of the production technology across firms and the concavity of the household's utility.

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. All firms face the same problem for the production of their consumer goods. As a result they choose to produce the same quantity at the same marginal cost, and set the same price for it. As all profits are made on the sales of consumer goods, functional firms make symmetric profits in equilibrium which helps align their incentives to invest in reliability.

There are two other key features of the equilibrium:

Lemma 1.

- (1) The household's consumption is equal to $\kappa h(\rho(x, \mu))$, 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 $\kappa g(\rho(x, \mu))$; where $g : [0, 1] \rightarrow \mathbb{R}_+$ is a decreasing function.

The result on the household's utility is driven by love of variety: consumers are better off when the same scarce factor (labor) can be used to produce more varieties.¹⁸ The result on gross profits is more subtle.

3.2.2. Comments on the setup. As in standard monopolistic competition models, a firm commits to a single price \underline{p} for its consumption good \underline{v} before producing. These goods are sold at a markup above marginal cost; the quantity is determined by consumer demand at that price. Though a firm

¹⁸Though a larger number of functional varieties requires labor to be shared out among more varieties at once, constant returns to scale of production means that there is no loss or increase in productivity due to this.




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 $\rho(x, \tau)$ as τ tends to infinity.

would like to sell more at that price, the renegotiation in which the firm cuts the price for some consumers is unavailable: the firm posts a single price to maximize its profits.¹⁹

Concerning intermediate goods, the general equilibrium approach similarly posits that the intermediate of variety v has a single price irrespective of who buys it. Paralleling the discussion above, this forecloses some bargaining possibilities. One could imagine that some firms would be willing to pay a higher price for v than others because, in a given supply network realization, they have no alternative functional suppliers from whom they might acquire a needed input. General equilibrium pricing rules out hold-up of this sort. Our rationale is that relational contracts exist precisely to foreclose this kind of hold-up: a firm that attempts to extract more can be punished in future periods by its counterparties' forming relationships with different suppliers.²⁰

3.3. The aggregate production function: 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. Let $C^*(x, \mu_\tau)$ denote household utility in equilibrium when the distribution of depths is μ_τ and relationship strength is x . Our first result is:

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

- (i) if $x < x_{crit}$, we have that $\rho(x, \mu_\tau) \rightarrow 0$ and $C^*(x, \mu_\tau) \rightarrow 0$ as $\tau \rightarrow \infty$. That is, both reliability and consumption converge to 0.
- (ii) if $x > x_{crit}$, then, for all large enough τ , we have $\rho(x, \mu_\tau) > \bar{r}_{crit}$ and $C^*(x, \mu_\tau) \geq \underline{C} > 0$, where \underline{C} does not depend on x or τ . That is, both reliability and consumption remain bounded away from 0.

¹⁹As 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.

²⁰Indeed, one of the main advantages of close supply relationships is that they help firms commit to efficient prices and avoid holdup when a firm has temporary market power over another (Uzzi, 1997; Kirman and Vriend, 2000). This can be further microfounded in models of the determination of a relational-contract price, but we stick with a simple general-equilibrium description here.

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 graph as we approach x_{crit} from above grows 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.

3.3.1. 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:

$$\tilde{\rho}(x, d) = \mathcal{R}_x(\tilde{\rho}(x; d - 1)). \quad (3)$$

Indeed, more explicitly, the function that makes this true is²¹

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

Proposition 5 in Appendix C.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 $\tau \rightarrow \infty$, and that a sharp transition like that shown in Figure 3(b) holds for any complexity $m \geq 2$ and any multisourcing level $n \geq 2$.

The intuition for the sharp transition can be understood by looking at Figure 4. Here we plot the probability that a given firm is functional against the probability that its suppliers are 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 satisfying the fixed point condition in equation (3) 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 (3) starting with $\tilde{\rho}(x, 0) = 1$ yields a sequence of $\tilde{\rho}(x, d)$ converging to the largest such intersection. The reliability associated with that intersection is bounded away from 0. 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 6.1.4 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.

²¹Consider 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 such a “good event” happens is $(1 - (1 - xr_s)^n)^m$.

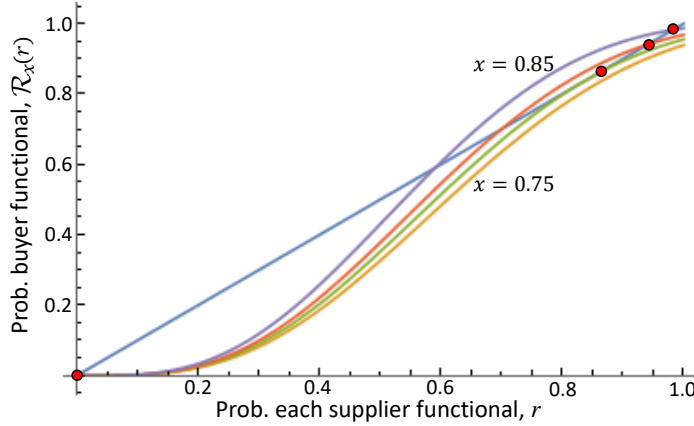


FIGURE 4. The probability, $\mathcal{R}_x(r)$, that a focal firm is functional as a function of r , the probability that a random supplier is functional. Here we use the parameters $n = 3$, $m = 4$ and $x \in \{0.75, 0.775, 0.8, 0.85\}$. The intersections with the 45 degree line marked by the red circles represent reliability values (for deep supply trees) associated with the given exogenous parameters.

3.3.2. 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) 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—setting efficient relationship strengths—and then turn to a decentralized problem. Throughout, we focus on symmetric outcomes.

The results show that while investments that put the economy 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 first consider a planner who chooses a value of x to which all relationship strengths x_{if} are set. This can be thought of as choosing the quality of institutions. The planner's problem is

$$\max_{x \in [0, 1]} C(x, \mu) - c(x). \quad (4)$$

Here $C(x, \mu)$ is equilibrium consumption²² given relationship strengths x and depth distribution μ , while c is a convex function representing the cost of maintaining institutions of quality x , paid in units of output. We assume that $c(0) = 0$, $c'(0) = 0$, and $\lim_{x \rightarrow 1} c'(x) = \infty$. Recall that κ , which we assume is positive, is a total-factor productivity multiplier in equation (1). The planner seeks to maximize expected social surplus, which is the total surplus produced by the firms that are functional minus the cost of maintaining institutions.

²²This is the same as household utility. We omit the asterisk for equilibrium, since we will work with an equilibrium of the production economy from now on.

Define the correspondence

$$x^{SP}(\kappa, \mu) := \operatorname{argmax}_{x \in [0,1]} C(x, \mu) - c(x)$$

This gives the values of x that solve the social planner's problem for a given productivity 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 for all τ sufficiently large there exists a $\kappa_{\text{crit}}(\mu_\tau) > 0$ such that

- (i) for all $\kappa < \kappa_{\text{crit}}(\mu_\tau)$, $x^{SP}(\kappa, \mu_\tau) = 0$.
- (ii) for all $\kappa > \kappa_{\text{crit}}(\mu_\tau)$, all values of $x^{SP}(\kappa, \mu_\tau)$ are strictly greater than x_{crit} .
- (iii) for $\kappa = \kappa_{\text{crit}}(\mu_\tau)$, all values of $x^{SP}(\kappa, \mu_\tau)$ are either strictly greater than x_{crit} or else equal to 0.

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. 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 at the critical level x_{crit} . The reason is as follows: at this value of x the marginal social benefits of investing are infinite 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 $f \in \Omega_i$ (the same as the index set of the varieties), which is typically $[0, 1]$. 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 firm (i, f) owns a variety $v = (i, f)$. The random realizations of the supply network occur after the firm chooses its investment level. Substantively, this means that a firm's investments occur before the realization of the randomness that governs operation of relationships in a particular time period.²³ If a firm chooses an investment level y_{if} , then all sourcing links from its variety (i, f) have relationship strength

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

²³This assumption is technically convenient, as it keeps the solution of the model symmetric. Substantively, it captures uncertainty about the depth of the supply chain of a firm's variety. The 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.

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 institution—e.g., how likely a “basic contract” is to deliver.²⁴ The main purpose of this baseline level is as a simple channel to shock relationship strength.²⁵ We will typically think of firms directly choosing the strength of their relationships, $x_{if} \in [\underline{x}, 1]$, and paying the corresponding investment cost $c(x_{if} - \underline{x})$.

The timing is as follows.

1. Firms simultaneously choose their investment levels.
2. The realized supply network is drawn, a competitive equilibrium allocation is implemented, and payoffs are enjoyed.²⁶

An *outcome* is given by relationship strengths x_{if} for all firms if .

We make the following assumptions:²⁷

Assumption 1.

- (i) $\underline{x} < x_{\text{crit}}$;
- (ii) c' is increasing and weakly convex, with $c(0) = 0$;
- (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 (free) relationship strength is not so high that the supply network is guaranteed to be productive even without any investment. The second assumption imposes assumptions on investment costs that help guarantee agents’ optimization problems are well-behaved. The Inada conditions, as usual, will be important in making investments interior.

4.2.2. Payoffs and equilibrium. We now turn to a firm’s payoffs, which—given our focus on symmetric outcomes—we must specify only for symmetric behavior of other firms. Let $P(x_{if}; x, \mu)$ be the probability with which a firm if is functional if firm if ’s investment choice is x_{if} and all other firms choose a symmetric investment level x .²⁸ By Lemma 1 in Section 3, conditional on being functional and all other firms having a reliability r , a firm earns a gross profit from sales of final goods equal to $\kappa g(r)$; where $g : [0, 1] \rightarrow \mathbb{R}_+$ is a decreasing function. Thus, the net expected profit of firm if when it makes its investment decision is

$$\Pi_{if} = \underbrace{P(x_{if}; x, \mu)}_{\text{prob. functional}} \underbrace{\kappa g(r)}_{\text{gross profit}} - \underbrace{c(x_{if} - \underline{x})}_{\text{cost of investment}}. \quad (5)$$

If a firm does not enter, its net profit is 0.

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 productivity multiplier κ there is an unproductive regime in which positive investment cannot be sustained. Next, for intermediate

²⁴This might depend, for instance, on the quality of the commercial courts.

²⁵Other specifications, such as a multiplicative one, could be used.

²⁶Our results in Section 3 imply that payoffs at this stage are uniquely determined.

²⁷Recall Proposition 1. See also Proposition 5 in Appendix C.1 for more details on this value.

²⁸We identify a firm with its variety, and thus speak of a firm being functional or not, etc.

values of κ there is a critical regime in which equilibrium investment is x_{crit} and arbitrarily small shocks to relationship strength lead to discontinuous drops in production. Finally, there is a noncritical regime in which equilibrium investment is above x_{crit} and the economy is robust to small shocks.

It is worth writing (5) 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 investment choice x_{if} given that all other firms choose a symmetric investment level x :²⁹

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

where d inside the expectation is drawn from the depth distribution μ conditional on depth being at least 1. Recall from equation (3) in Section 3.3.1 the formulas for $\tilde{\rho}$.

Definition 1. We say $x \geq \underline{x}$ is a *symmetric equilibrium* if $x_{if} = x$ maximizes $\Pi_{if}(x_{if}; x, \mu)$ for any firm if . It is a *symmetric undominated equilibrium* if it is the symmetric equilibrium maximizing the social surplus of production net of investment.³⁰

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.³¹ 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 symmetric mutual best responses maximizing output whenever possible.

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 Assumption 1). Hence, for large enough τ , $x_{if} = \underline{x}$ maximizes $\Pi_{if}(x_{if}; \underline{x}, \mu_\tau)$ and so there exists an equilibrium.

In analyzing the symmetric equilibria it is helpful to make an assumption on the environment that ensures that local optimality with respect to an investment choice implies global optimality.

Assumption 2. For any $x > x_{\text{crit}}$ the function $\Pi_{if}(x_{if}; x, \mu_\tau)$ has a unique interior local maximum for all large enough τ .³²

Assumptions 1–2 will be maintained in the sequel. The following lemma allows us to formulate a condition on primitives that is sufficient for Assumption 2 to hold.

²⁹Note 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} .

³⁰This is $C(x, \mu) - c(x)$. The latter term is the total cost of investments. We posit that firms in product i have measure $1/|\mathcal{I}|$ (i.e. that the measure on the set \mathcal{F}_i is the Lebesgue measure times $|\mathcal{I}|^{-1}$) in order to make the constant multiplying $c(x)$ equal to 1.

³¹This is the place where we use the assumption that $\eta_C > \frac{1}{2}$, which is technically convenient for this argument, though stronger than necessary.

³²The assumption permits another local maximum at a corner. We rule this out separately.

Lemma 2. For any $m \geq 2$, and any $n \geq 2$, there is a number³³ \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 $\underline{x} \geq \hat{x}$, then Assumption 2 is satisfied.

Consider any environment where $\underline{x} \in [\hat{x}, x_{\text{crit}})$. Part (i) of the lemma guarantees that the interval $[\hat{x}, x_{\text{crit}})$ is nonempty, and part (ii) guarantees that Assumption 2 is satisfied for values of \underline{x} in this range.³⁴

We now characterize the equilibrium behavior.³⁵

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

- (i) If $\kappa < \underline{\kappa}$, there is no investment: $x_{\tau}^*(\kappa) = 0$
- (ii) For $\kappa \in [\underline{\kappa}, \bar{\kappa}]$, the equilibrium investment level satisfies $x_{\tau}^*(\kappa) \in [x_{\text{crit}} - \varepsilon, x_{\text{crit}} + \varepsilon]$. We call such equilibria *critical*.
- (iii) For $\kappa > \bar{\kappa}$, the equilibrium investment level satisfies $x_{\tau}^*(\kappa) > x_{\text{crit}} + \varepsilon$. 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 about different supply networks 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 investments converging to x_{crit} in equilibrium. This contrasts with the social planner's solution, which never selects such investments $x = x_{\text{crit}}$. 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 the 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 given the reliability of its suppliers. The thin 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 $\underline{x}_{\text{crit}}$ at $\underline{\kappa}$, and is increasing in κ thereafter. 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 $\rho(x_{\text{crit}}) = [0, \bar{x}_{\text{crit}}]$, equilibrium investment for large τ remains arbitrarily close to x_{crit} . In other words, equilibrium investment choices bunch around x_{crit}

³³In the proof, we give an explicit description of \hat{x} in terms of the shape of the function $P(x_{if}; x, \mu_{\tau})$.

³⁴While conditions we identify in Lemma 2 are sufficient for satisfying Assumption 2 they are not necessary.

³⁵Throughout this section, we restrict attention to the case of $m \geq 3$. It is essential for our results that supply networks are complex and $m \geq 2$, 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 $m = 2$.

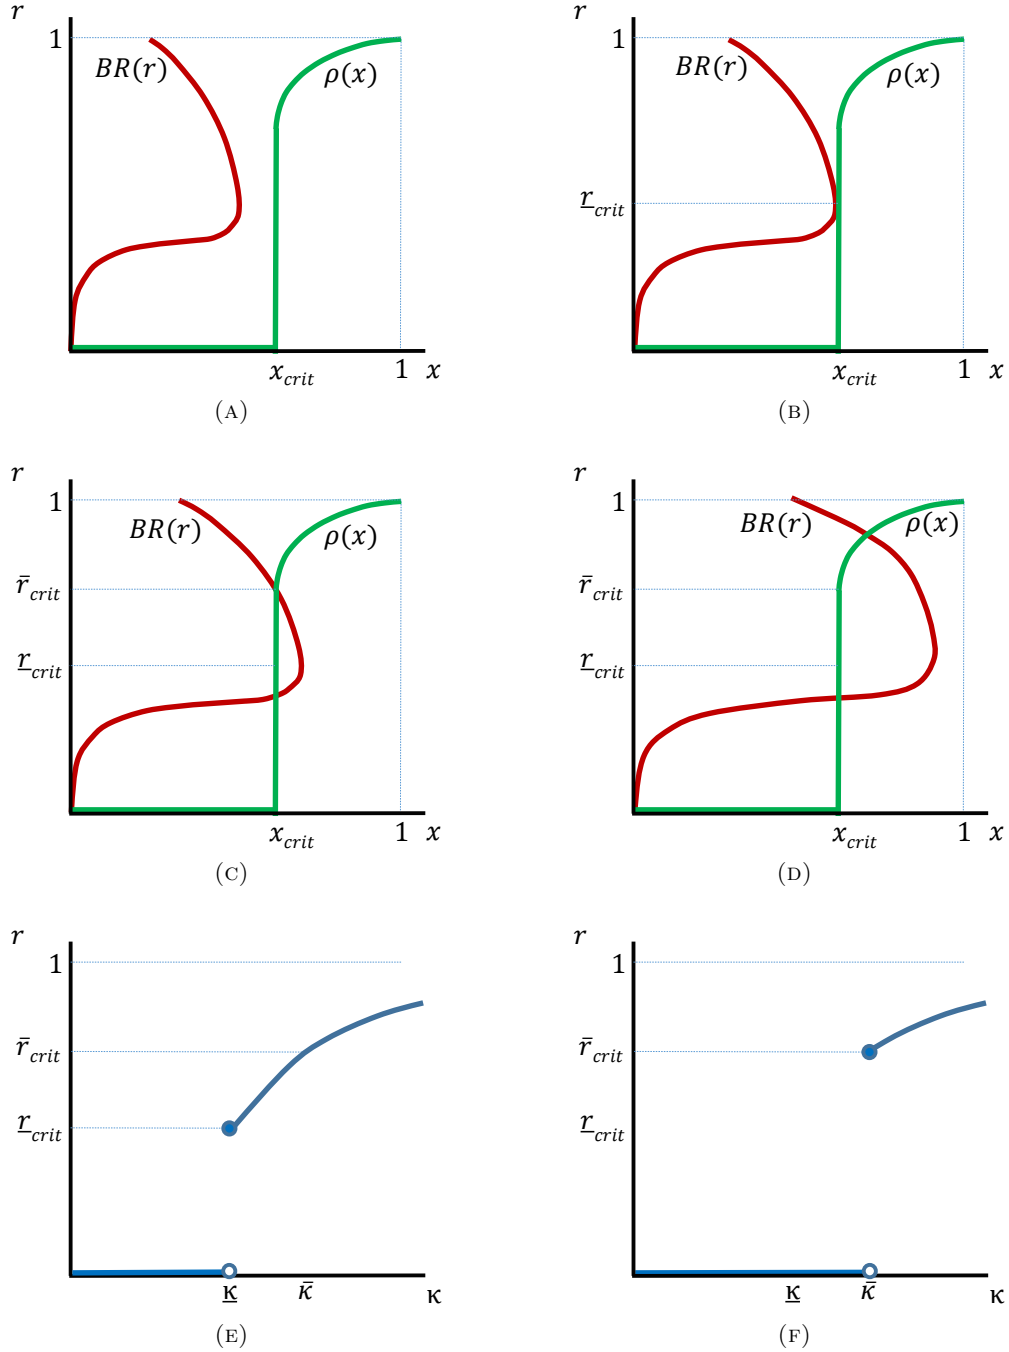


FIGURE 5. Panel (A) shows an equilibrium for $\kappa < \underline{\kappa}$. Panel (B) shows an equilibrium with $\kappa = \underline{\kappa}$. Panel (C) shows an equilibrium with $\kappa = \bar{\kappa}$. Panel (D) shows an equilibrium with $\kappa > \bar{\kappa}$. Panel (E) plots how equilibrium reliability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

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.

The key to proving Theorem 1 is showing that, as depicted in Figure 5, the highest intersection between the two curves moves downward while it is well-defined. This is not straightforward. Investment incentives (which determine the shape of the best-response curve) are complex, and so equilibrium investment is difficult to characterize directly. One reason for this is a non-monotonicity: investments in relationship strength are strategic complements in some regions of the parameter space, and strategic substitutes in others.³⁶ Thus, there are no straightforward monotone comparative statics. The problem is made more challenging by the fact that we do not have explicit expressions either for the reliability curve or the best-response curve. We establish the key property by combining a first-order condition characterizing best-response investment with the reliability condition, and showing that there can be at most one solution to both for $r \geq r_{\text{crit}}$.³⁷ Once we have uniqueness, the comparative statics follow by analyzing how the best-response curve (which we show slopes strictly downward at the point of intersection) shifts inward as we decrease κ .

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 are taken to have zero probability (though the analysis is robust to anticipated shocks that happen with sufficiently small probability).

Definition 2 (Equilibrium fragility).

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

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

That is, a shock that reduces relationship strengths arbitrarily little (ϵ) from their equilibrium levels leads to a drop to 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.

³⁶When 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.

³⁷This comes down, after a lot of massaging, to showing that there is only one root of a certain function that we *can* write explicitly. This function being zero is a necessary condition for an equilibrium.

Proposition 3. 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 common 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 short-term 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 there is a new pandemic outbreak of a disease. This will likely affect some supply relationships, but not others; uncertainty around this can be captured by a small reduction in the probability that a given supply relationship is disrupted.

6. DISCUSSION OF THE MODEL

This section discusses our modeling choices. First, we discuss which features of our environment are essential, and explain the empirical facts motivating them. Next, we present evidence from the existing literature congruent with the short-run and medium-run predictions of our model concerning the effect of shocks. Beyond the key assumptions, we make a variety of assumptions for tractability or simplicity. At the end of the section, we discuss the robustness of our model to various relaxations of these non-essential assumptions.

6.1. Essential features. The key features of our model are that (i) a typical firm uses failure-prone relationships with particular suppliers to source inputs; (ii) the investment in these relationships is voluntary and decentralized, due in part to the incompleteness of contracts; and (iii) production is complex—a typical firm’s production relies on multiple such non-commodity inputs. In this section we discuss some motivations and interpretations of these crucial assumptions. We also show that without any one of them, our main results on equilibrium fragility would not obtain.

6.1.1. Relationship-based sourcing and idiosyncratic disruptions. This subsection discusses both the limited number of suppliers that firms have and the shocks that their supply relationships are exposed to disruptions.

Specific sourcing relationships are a central feature of the modern economy. 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), while Barrot and Sauvagnat (2016) find that firms have difficulty switching suppliers even when they need to do so. There are many reasons behind firms’ reliance on a small number of suppliers to source a given type of input that are explored in this literature. Technological compatibility and geography can limit the pool of potential suppliers; hold-up problems can make trust important; repeated interactions can help firms to avoid misunderstandings.

The specific relationships that firms maintain for sourcing are imperfect; they are prone to disruptions. While the harm of some disruptions can be mitigated—e.g., through the use of inventories—this is not always possible, and so disruptions can cause cascading damage. Evidence supporting both the presence of specific sourcing relationships and their disruption comes from the presence of cascades (see, for example, Taschereau-Dumouchel (2020) and Carvalho et al. (2020)). The cascades studied in Taschereau-Dumouchel (2020) provide evidence that dependencies on specific suppliers do indeed result in contagions of disruption.³⁸

Some 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 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 found 1,069 supply chain disruption events globally during a six-month period in 2018 (Supply Chain Quarterly, 2018). In the model, a given supply chain is operating with some probability, which can be interpreted as some fraction of the time. Thus, as in reality, disruption events are frequent.

We close with a short comment on the incidence of the shocks. In our model, the ultimate source of a shock is a disruption of some particular supply relationship between two firms. As we discuss below in Section 6.3.4, the key forces in our model are robust to exactly how the shock occurs—whether shocks hit links in the supply network or nodes (i.e., firms/varieties). This is intuitive. If a node fails, then that sourcing option is lost just as it is when a link to the node fails.

6.1.2. *Non-contractible investments in relationships.* In this section we argue that the relationships just discussed are built through decentralized investment by firms in an environment of incomplete contracting.

³⁸This paper provides a theoretical account of cascades and then calibrates it to the data. While the focus is on the failure of specific firms, there is a close connection between systemic implications. In both our paper and Taschereau-Dumouchel (2020), there are correlations in the functioning of firms near each other in the production network.

Investments in supply relationships are often “soft” in nature and not easily observed. They include practices such as better understanding a supplier’s or customer’s capabilities by visiting their facilities, querying odd instructions to catch mistakes, building social relationships that span the organizations, and so on.³⁹ These types of activities have all been documented by ethnographic studies in sociology—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.

The evidence of disruptions presented in the previous section shows the potential for damaging sourcing disruptions and helps explain why investments are made. Hendricks 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.

Investments in relationships can be hard to systematically observe. However, an extreme form of such an investment is readily observable—vertical integration. As we discuss in Section 6.2, consistent with our model, Boehm and Oberfield (2020) find evidence of increased vertical integration when courts are less effective at enforcing formal contracts.⁴⁰ This behavior suggests the importance of detailed maintenance of supply relationships and the limits of arm’s-length relationships in this regard.

6.1.3. *Multiple essential inputs.* The need to source multiple inputs via specific sourcing relationships gives rise to strong complementarities in production in our model. Kremer (1993) was an influential theoretical account of the importance of complementarities, and was followed by a literature arguing 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 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).⁴¹

While this evidence is all consistent with our model, it is only suggestive of complexity in our sense (that multiple essential inputs are required at each stage); alternative models also deliver consistent complementarities. However, Barrot and Sauvagnat (2016) find direct evidence that firms need to source multiple inputs via specific supply relationships. They show that if a supplier is hit by natural disasters 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.

6.1.4. *Contrasts with benchmark models without the essential features.* In this section, we briefly consider three types of benchmark models that highlight the necessity of the key features discussed in the previous section.

³⁹Such features are also key to models such as Antràs (2005) and Acemoglu and Tahbaz-Salehi (2020).

⁴⁰See also Boehm and Sonntag (2019).

⁴¹Prior 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.

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 for one reason or another. (A shipment might be lost or defective, or a misunderstanding could lead the wrong part to be supplied.) In the metaphor we introduced earlier, we now assume each firm extends relationships *only* to functional suppliers, but we still keep the randomness in whether these relationships work.⁴²

Let the probability a given attempt at sourcing an input succeeds be x , independently. As before, in view of this idiosyncratic risk, each firm *if* multisources by contracting with two potential suppliers of each input. 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 7 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 6.3.4.

Contrast II: More complete contracts. The incompleteness of contracts and the decentralized nature of investments (modeled as a Nash equilibrium) is at the heart of our model. With fully complete multilateral contracts, efficient investments could be supported as an outcome. This would imply the absence of fragility in equilibrium, by our results on the planner’s problem (recall Proposition 2). There is a parallel between our assumptions here and those made in the financial networks literature (see, e.g., Acemoglu, Ozdaglar, and Tahbaz-Salehi (2015)): perfect contracts could facilitate the efficient management of all systemic risk, but some incompleteness seems realistic because, in both cases, it is difficult to contract on all investments relevant to robustness.

It is natural to then ask whether various other kinds of agreements—without going all the way to complete multilateral contracts—could help. It is immediate that if coalitional commitments among all firms upstream of each supplier were possible that all externalities could be internalized and hence inefficiency could be eliminated. In practice the coalitions making these commitments would have to be quite large. An interesting question is whether large coalitions are necessary, or whether sufficiently good bilateral contracts could work well. In a different but related environment for studying supply networks, Bimpikis, Fearing, and Tahbaz-Salehi (2018) find that bilateral contracts are insufficient to eliminate contracting frictions, even when investments are observable. Relatedly, the financial networks literature has made precise the sense in which bilateral contracts with realistic limitations are insufficient to restore efficiency.⁴³ In Section 6.3.1 below we also show that if a firm could voluntarily monitor those upstream of it, or invest in them, we would not expect efficiency to be restored.

⁴²Formally, in the construction of the random graph, we allow each firm’s potential supply links to be formed conditional on the realization of its supplier’s functionality; in particular, these links are extended only to functional suppliers.

⁴³See Section 4.2 of Acemoglu, Ozdaglar, and Tahbaz-Salehi (2013).

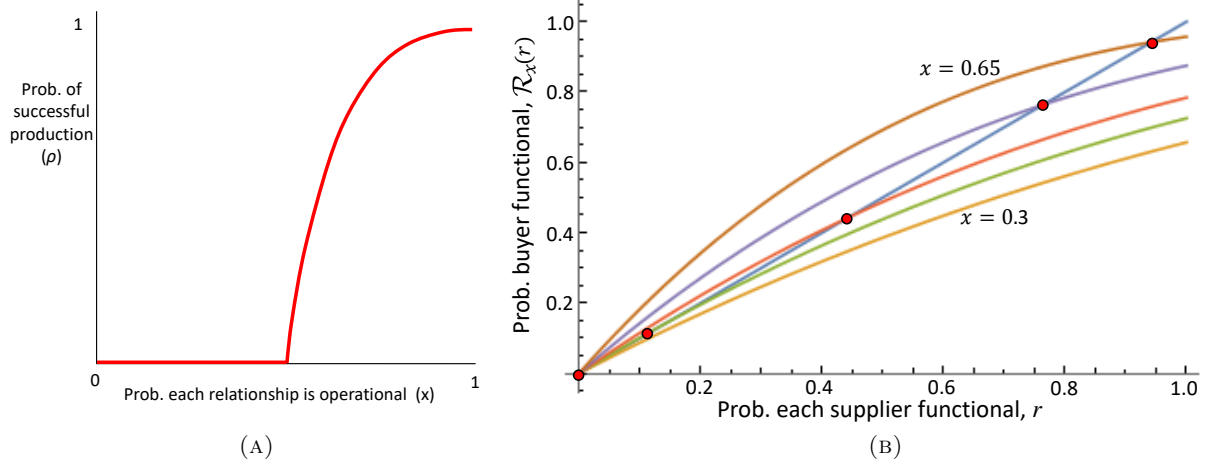


FIGURE 6. (A) The probability of successful production of a simple good ($m = 1$ distinct inputs needed) as relationship strength varies. (B) The probability, $\mathcal{R}_x(r)$, that a focal firm is functional as a function of r , the probability that a random supplier is functional. This plot is for $n = 3$ with $x \in \{0.3, 0.35, 0.4, 0.5, 0.65\}$. The plot parallels Figure 4, which depicts the same function for the complex production ($m > 1$) case.




FIGURE 7. The probability of successful production for a firm as market-based sourcing attempts become more likely to succeed.

Contrast III: 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 each firm requires only *one* type of risky input relationship to work.⁴⁴ We plot how the probability of successful production varies with relationship strength in Figure 6(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

⁴⁴As 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.

production is strictly positive. The change at $x = 0.5$ has 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). In the large-depth limit, production will be successful if the supply tree upstream of a given firm continues indefinitely. 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.⁴⁵ 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.

6.2. Short-run and medium-run predictions. In this section we discuss our model in terms of its short-run and medium-run predictions, and present congruent evidence from the existing literature. 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.⁴⁶ We define the medium run as a period of time over which adjustments in relationship strength are possible.⁴⁷ A variety of research has examined both the implications of shocks in the short run and the endogenous adjustments that firms undertake in response to shocks.

Blanchard and Kremer (1997) is an examination of the aggregate consequences of a shock to institutional quality. After the fall of the Soviet Union, the institutions in place for a planned economy were lost and market-based institutions had not been built up to adequately take their place. The short run implications of such a shock in our model is that supply networks in the fragile regime suffer severe disruption, while depending on the size of the shock, other supply networks might also suffer severe disruptions. Blanchard and Kremer (1997) argue that this type of shock can explain the large drop in aggregate output that occurred.⁴⁸

The calibration in Acemoglu and Azar (2020)⁴⁹ and the findings of Di Giovanni and Levchenko (2012) and Di Giovanni, Levchenko, and Mejean (2018) are also consistent with substantial disruptions to production arising from relatively small shocks. Such analyses are often motivated by

⁴⁵Given 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.

⁴⁶See their Figure IV.

⁴⁷In 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 8.2 below.

⁴⁸Their explanation is focused on distinctive features of the state-run economy, including single-sourcing. In contrast, we focus on economies with well developed market-based institutions. We show that such economies can endogenously arrange themselves in a way that makes them, too, extremely sensitive to institutional shocks. Instead of single-sourcing playing a key role in fragility, the interaction between endogenous multisourcing and complex production is what matters for the fragility of our decentralized market economy.

⁴⁹See their Section 4 and Appendix D.

theories of granularity (Gabaix, 2011), where idiosyncratic shocks have aggregate impacts. The empirical studies assess, for instance, shocks in one country on output in another country. Our analysis provides two complementary approaches to interpreting such effects. First, within the heterogeneous supply network version of our model discussed in Section 6.3.5, a disruption to supply relationships in one country can activate (via a “weakest link” property) a cascade of disruptions elsewhere. Even if there are no specific sourcing across countries, a shock in one first country can cause—via market prices, or other forces—a reduction in the productivity of distant supply networks, making them fragile. That can then lead to a fairly severe drop, at least temporarily, in the reliability and hence productivity of the distant supply networks. Fleshing out these points requires developing our theory to accommodate interactions across different supply networks via markets, which we discuss in Section 8.3.⁵⁰

In the medium run, investments in reliability can adjust in response to changes in productivity or institutional quality. By Theorem 1 and Corollary 1, improvements in institutional quality or productivity increase the strength of 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, reductions in institutional quality (lower \underline{x}) induce an almost one-for-one increase in investment, so that relationship strength is close to constant. These comparative statics predictions are congruent with evidence on institutional quality and investments in specific supply relationships. Comparing Indian states with better and worse court system, Boehm and Oberfield (2020) find that firms facing worse court systems respond in a way that can be interpreted as an extreme form of investing in their supply relationships—vertically integrating. Also consistent with this, Johnson, McMillan, and Woodruff (2002) provide survey evidence that reduced trust in courts makes firms that rely on relationship-specific inputs less likely to switch suppliers. This is suggestive of more investment in relationships endogenously being used to compensate for poor institutions.

In our model, given the same reliability of relationships 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. This is consistent with evidence from Boehm (2020) that weak institutions have a bigger impact when potential hold-up problems are more severe (as they are for us when production is more complex). A similar prediction of our model is that when there are fewer multisourcing opportunities, so that production networks are sparser, then, all else equal, production will be less reliable (see Section 3.3.2). Herskovic (2018) argues, in an asset pricing framework, that this should imply a premium on the equity values of firms embedded in sparser production networks, and finds evidence of this.

There is one final feature of our model that is worth remarking on. Following a shock to relationship strengths, a collapse in production occurs in the short run, i.e., holding relationship strengths at their old values. In the medium run (i.e., the time scale on which investments are adjustable), firms will frequently be able to select new investment levels that allow the supply network to resume

⁵⁰Of course, in reality, both truly granular effects stemming from shocks to large firms as well as the amplification of disruptions that we discuss are likely to be relevant. We do not take any stand on the right decomposition of these effects.

being productive. Our equilibrium selection 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 severe productivity damage following a shock. Evidence from Taschereau-Dumouchel (2020) and Acemoglu and Tahbaz-Salehi (2020) both suggest that adjustments on the extensive margin are important.

6.3. Nonessential features and extensions.

6.3.1. *Investment only in own relationships.* It is instructive to consider ways in which the actions available to firms might be expanded to improve reliability, and whether these could help overcome the contracting frictions in our setup. Suppose we allowed each firm to invest in *any* supply relationship (e.g., between their suppliers and their suppliers' suppliers). Each relationship is a public good for all those downstream in the supply network. Improvements in its reliability will strictly increase the probability that every such firm is functional and hence improve their profitability. So every such firm has an incentive to invest in it. But allowing them to invest in it does not solve the problem. In equilibrium, following the logic of Bergstrom, Blume, and Varian (1986), typically only one firm will invest in each relationship. Each firm will be willing to invest up to the point that the marginal return of the investment equals the marginal benefits. However, the marginal benefits are heterogeneous and it will only be possible for this condition to be satisfied by the firm that has the highest marginal benefit of investing. In our model, in equilibrium, the natural investor is the firm which relies directly on the supply relationship. Modeling firms as investing only in their direct relationships is a convenient simplification but, for the reasons just described, is unlikely to change the analysis markedly.

6.3.2. *Investment only on the intensive margin of relationships.* 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} .⁵¹ Conditional on a supplier being found, the relationship is operational. In Appendix F we discuss a richer extensive-margin interpretation, and also one that permits separate efforts to be directed to the extensive and intensive margins simultaneously.

6.3.3. *Fixed entry decisions.* 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

⁵¹This 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)).

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 $\bar{f} > 0$ 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 κ parameters 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. 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. To show why this is a generic (i.e., not knife-edge) situation, consider a small increase in κ . This makes entry more appealing enticing new firms to enter until the marginal firm is again indifferent.

Going beyond this, it can be shown that under suitable conditions in this richer model, an analogue of Theorem 1 holds, with an ordering of regimes. We do not pursue this extension here.

6.3.4. Details of how disruptions are modeled. 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.

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. We focus on shocks to relationships for simplicity, but the mathematical forces we identify also operate in alternative models of disruption.

Finally, 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 the precipice is driven by externalities, and while all firms may like to invest more to avoid a shock, this effect may not be large enough to overcome the underinvestment that leads to fragility.

6.3.5. *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 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 of product i , $\mathcal{I}(i)$, is an arbitrary set with some cardinality m_i . Thus, the input requirements no longer need to be homogeneous. (2) For each product i and input product $j \in \mathcal{I}(i)$, 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 $j \in \mathcal{I}(i)$, 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. We also suppose here for ease of exposition that the reliability curve is given by the limit correspondence at $\tau = \infty$, with the understanding that results for the high- τ limit can then be treated from this (fictitious) limit economy as in our other analysis.

For this section, we use $x^* = (x_1^*, x_2^*, \dots, x_{|\mathcal{I}|}^*)$ to describe equilibrium investment profiles for the different products. Note that $x_i^* = (x_{i,j}^*)_{j \in \mathcal{I}(i)}$ is a vector describing the heterogeneous equilibrium efforts invested by a producer of product i in its relationships with firms supplying different inputs $j \in \mathcal{I}(i)$, 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 . An investment profile x_i is critical 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. Product i is said to be *critical* in an equilibrium when an arbitrarily small decrease in the relationship strengths of firms producing that product result in a considerable (i.e., bounded away from zero) drop in the reliability r_i .

The sharp transition in the mechanics of our basic model, discussed in Section 3.3, 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 SA2.3 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 the production of one product is fragile, this makes the production of all products that rely directly or indirectly on it fragile as well.⁵²

⁵²Thus, 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).

Proposition 4 (Weakest link property). Denote by $c_{ij}(x_{ij} - \underline{x}_{ij}) = \gamma_{ij}\tilde{c}_{ij}(x_{ij} - \underline{x}_{ij})$ the cost for a firm producing product i of investing in a relationship with a supplier of product j . Here $\gamma_{ij} > 0$ is a scaling factor and \tilde{c}_{ij} is a cost function satisfying the same assumptions as before.

- (i) Suppose product i is critical. Then any other product j on a directed path from j to i (on the product interdependencies graph) will fail following a shock $\epsilon > 0$ to γ_{ik} (which increases γ_{ik} to $\gamma_{ik} + \epsilon$) for any $k \in \mathcal{I}(i)$.
- (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 $i \in \mathcal{I}^{SC}$ have critical relationship strengths or no producers of products $i \in \mathcal{I}^{SC}$ have critical relationship strengths.

Proposition 4 shows 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, reduce gross margin G_i 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 link.

In Sections SA2.1 and SA2.2 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 above, in Section 6. In this section we review several high-level connections to related literatures not covered by our previous discussions.

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 (2014) 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), Taschereau-Dumouchel (2017), 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 (2017) 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 discuss below how our explicit modeling of sourcing failures at the micro level gives rise to new effects.

There has been considerable recent interest in markets with non-anonymous trade mediated by relationships.⁵³ The work most closely related to ours in this area also studies network formation

⁵³A 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).

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 work in varied 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 non-anonymous trade 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 supply 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.

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.⁵⁴ 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. More importantly, 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.⁵⁵ 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 work, the supply network is robust to idiosyncratic shocks but very sensitive to arbitrarily small aggregate 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 note some basic welfare implications of the model,

⁵⁴These 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).

⁵⁵Endogenously, 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.

Next, we briefly mention some potential applications of the precipice phenomena in analyzing industrial development as institutions improve. Finally, we analyze how our modeling can play a role in the analysis of aggregate volatility.

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 systematic underinvestment in reliability at fragile equilibria. Finally, we briefly comment on possible policy remedies.

There are three externalities present when a firm chooses its reliability. First, there is the non-appropriability of consumer surplus. Consumers value variety, and this is not fully internalized in a firm’s reliability choice. Second, there are reliability spillovers. 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.

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 infinite—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.

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

$$\Pi_{if} = \underbrace{\mathbf{E}[F_{if}]}_{\text{prob. functional}} \underbrace{\kappa g(r)}_{\text{gross profit}} - \underbrace{(1 - \theta)c(y_{if})}_{\text{cost of effort}}.$$

Suppose the status quo was 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. Some simple implications for industrial development. The comparison between the production of complex and simple products developed in Section 6.1.4 has interesting implications for the complexity of technologies used across countries, and for industrial development. While a full analysis of this is beyond the scope of this paper, Figure 8 illustrates the main point—if




FIGURE 8. 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.

complex products are more valuable, small increases in quality of commercial institutions can make more complex production technologies viable and yield discontinuous benefits to an industry.

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) More complex supply chains are associated with higher rates of disruption (Craighead, Blackhurst, Rungtusanatham, and Handfield, 2007).

8.3. Aggregate volatility. So far we have focused on a single complex supply network with particular parameters. The larger economy can be thought of as consisting of many such supply networks, each one small relative to the economy. We study whether the forces making some of these small supply networks fragile can lead to fragility in the aggregate for such an economy.

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 previous results on the fragile regime, especially Proposition 6, are essential to this.

For simplicity, fix the function $c(\cdot)$.⁵⁶ A given supply network is then described by a tuple $\mathbf{s} = (m, n, \kappa)$. 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 $m = 1$.⁵⁷ 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.

In this setting of multiple supply networks with heterogeneous primitives, one might hope that aggregate shocks have a silver lining: a “housecleaning” or “forestfire” effect: Though a small aggregate shock will cause the fragile supply networks to collapse, the remaining ones networks will be robust to further shocks. We next show that this is too optimistic: collapse can beget more fragility and further collapses.

To flesh out this point, we must enrich our model a bit. 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. We think of these relationships as mediating the supply of inputs that are 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,

⁵⁶This could also be drawn from a distribution, but the notation would be more cumbersome.

⁵⁷For 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.

the productivity damage of collapses leads to domino effects where, iteratively, previously robust parts of the economy become fragile.⁵⁸

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.
- ACEMOGLU, D., A. OZDAGLAR, AND A. TAHBAZ-SALEHI (2013): “Systemic Risk and Stability in Financial Networks,” Tech. rep., National Bureau of Economic Research, working Paper No. 18727.
- (2015): “Systemic Risk and Stability in Financial Networks,” *American Economic Review*, 105, 564–608.
- ACEMOGLU, D. AND A. TAHBAZ-SALEHI (2020): “Firms, Failures, and Fluctuations: The Macroeconomics of Supply Chain Disruptions,” Tech. rep., National Bureau of Economic Research.
- 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.
- BAQAEI, D. R. (2018): “Cascading Failures in Production Networks,” *Econometrica*, 86, 1819–1838.
- BAQAEI, D. R. AND E. FARHI (2017): “Productivity and Misallocation in General Equilibrium.” Tech. rep., National Bureau of Economic Research.
- (2019): “The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten’s Theorem,” *Econometrica*, 87, 1155–1203.
- 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.
- BIMPIKIS, K., D. FEARING, AND A. TAHBAZ-SALEHI (2018): “Multisourcing and Miscoordination in Supply Chain Networks,” *Operations Research*, 66, 1023–1039.
- 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.

⁵⁸In Appendix D we provide a simulated illustration of this phenomenon.

- 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.
- BOEHM, J. AND J. SONNTAG (2019): “Vertical Integration and Foreclosure: Evidence from Production Network Data,” Centre for Economic Performance, LSE.
- 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. (2014): “From Micro to Macro via Production Networks,” *Journal of Economic Perspectives*, 28, 23–48.
- 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 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*.
- DI GIOVANNI, J. AND A. A. LEVCHENKO (2012): “Country size, international trade, and aggregate fluctuations in granular economies,” *Journal of Political Economy*, 120, 1083–1132.
- DI GIOVANNI, J., A. A. LEVCHENKO, AND I. MEJEAN (2018): “The Micro Origins of International Business-Cycle Comovement,” *American Economic Review*, 108, 82–108.
- 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. 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.

- GRADDY, 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. PERNOUD (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.
- JOHNSON, S., J. McMILLAN, AND C. WOODRUFF (2002): “Property Rights and Finance,” *American Economic Review*, 92, 1335–1356.
- JONES, C. I. (2011): “Intermediate Goods and Weak Links in the Theory of Economic Development,” *American Economic Journal: Macroeconomics*, 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.
- 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. SINSOYSAL (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/>.
- TALAMÀS, 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.
- TASCHEREAU-DUMOUCHEL, M. (2017): “Cascades and Fluctuations in an Economy with an Endogenous Production Network,” *Mimeo*.
- (2020): “Cascades and Fluctuations in an Economy with an Endogenous Production Network,” *Available at SSRN 3115854*.
- 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 formalize some details concerning the model in Sections 2.2 and 2.1.

We describe the supply Endow each \mathcal{V}_i with the Borel σ -algebra. A *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} .
- In the next several bullets we specify the edges of those $v \in \mathcal{V}$ with $d(v) > 0$. Edges are ordered pairs (v, v') where $v' = (j, \omega')$ for $j \in I(i)$ —the meaning is that v can potentially source from v' , and we depict such an edge as an arrow from v to v' .
- For each $j \in I(i)$ there are n edges (v, v') to n distinct varieties v' of product j , i.e. elements of \mathcal{V}_j . For any variety, define its neighborhood $N_v = \{v' : (v, v') \in \mathcal{E}\}$.
- For any variety v , the elements of $N_v \cap \mathcal{F}_j$ are independently drawn from an atomless distribution over \mathcal{V}_j conditioned on $d(v') = d(v) - 1$, and these realizations are independent.
- For any countable set of varieties \mathcal{V}' , neighborhoods N_{i_f} are independent.

Define \mathcal{G}'_+ to be a random subgraph⁵⁹ of \mathcal{G} in which each edge is kept independently, with probability x . More formally, define for every edge e a random variable $O_e \in \{0, 1\}$ (whether the edge is operational) such that

- $\mathbf{P}[O_e = 1 \mid \mathcal{G}] = x$ for every $e \in \mathcal{G}$ 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} .

A subset $\widehat{\mathcal{V}} \subseteq \mathcal{V}$ is defined to be *consistent* if, for each $v \in \widehat{\mathcal{V}}$ with $d(v) > 0$ the following holds: for each product j that $v = (i, \omega)$ 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, $\overline{\mathcal{V}}$, which is a superset of any other consistent set. For any given variety, this can be found by the simple iterative procedure in Section 3.1.

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. The formal construction above allows us to extend this to a continuum economy.

APPENDIX B. DETAILS AND PROOFS FOR MICROFOUNDATIONS OF PRODUCTION

We begin with some basic observations about equilibrium and then turn to proving Lemma 1. Throughout this analysis, we take $\rho = \rho(x, \mu)$ as a given parameter, let $N = |\mathcal{I}|$, and normalize the wage to 1. The first welfare theorem holds and equilibrium production is efficient—i.e., maximizes the quantity C subject to physical constraints.

We take labor to be the numeraire.

We start by showing that all firms produce the same quantity of goods for consumption.

Lemma 3. The selling prices of the intermediate goods of all varieties within an industry are equal.

Proof. Intermediates are sold at the marginal cost of the inputs required to produce them. By constant returns to scale, the price of each variety's intermediate good, p_v , is a strictly increasing function of the prices of each of firm v 's inputs and 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 $v = (i, \omega)$ 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 $p_v = \min_{\omega \in [0, 1]} p_{(i, \omega)}$. 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 . These produce at minimal cost so it follows that $p_v = \min_{\omega \in [0, 1]} p_{(i, \omega)}$. This gives us that all firms in the same industry have the same marginal costs and thus the same intermediate prices. \square

Lemma 4. All firms produce the same quantity of goods for consumption.

Proof. From the previous lemma, let 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 $j \in I(i)$. Let \mathfrak{P} be a function that expresses an industry's intermediate price in terms of its input industries' intermediate prices and let M_i be 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 $j_i(r) \in I(i)$; $M_{r,j_i(r)} = 1$; and $M_{r',j_i(r)} = 0$ for $r' \neq r$. Then, for all industries i ,

$$p_i = \mathfrak{P}(M_i \mathbf{p})$$

We have that \mathfrak{P} is increasing in each input as an increase in input prices will increase the price of the good. Additionally,

$$\mathfrak{P}(\alpha M_i \mathbf{p}) < \alpha \mathfrak{P}(M_i \mathbf{p}) \text{ if } \alpha > 1$$

⁵⁹The subscript $+$ refers to the positive-depth varieties: the entire graph \mathcal{G}' of supply relationships is obtained by setting the out-neighborhood of the depth-0 to be $\bigcup_{j \in I(i)} \mathcal{V}_j$, since they can source from anyone.

This is true because if all input prices (including the wage) are scaled, then the price of the good would be scaled by the same factor. But because the wage is fixed, the effect of this scaling is smaller.

Now assume, for the sake of contradiction, that not all 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 $\frac{p_a}{p_b} = r > 1$. Because the prices of the inputs of industry a must be at least p_a and the prices of the inputs of industry b must be at most p_b , each element of $M_a \mathbf{p}$ must be less than or equal to each element of $r M_b \mathbf{p}$. Now,

$$p_a = \mathfrak{P}(M_a \mathbf{p}) \leq \mathfrak{P}(r M_b \mathbf{p}) < r \mathfrak{P}(M_b \mathbf{p}) = r p_b = p_a$$

This is a contradiction. Thus, all industries must sell intermediate goods at the same price. \square

We now turn to the proof of Lemma 1 and normalize $\kappa = 1$. Note that because labor is the only unproduced input, the price of an intermediate good is equal to the total amount of labor used to produce a unit of it (its total factor content), which, in equilibrium, we denote L_i^* (recall labor is the numeraire). Thus, all firms sell at some universal price $p = L_i^*$. Because p does not depend on the reliability r , we can write this amount q^* and equilibrium consumption as follows:

$$1 = \sum_{i=1}^N p r q^*, \text{ which implies } q^* = \frac{1}{p N r}$$

$$C = \left(\sum_{i=1}^N r q^* \right)^{1/\eta_C} = \frac{r^{1/\eta_C - 1}}{p N} \quad (6)$$

We have that each firm produces $\frac{1}{p N r}$ and sells at a price marked up $1/\eta_C$ over marginal cost and thus earns gross profit

$$\left(\frac{1}{\eta_C} - 1 \right) \frac{1}{N r} \quad (7)$$

Because $0 < \eta_C < 1$, we have that consumption increases as r increases and profits decrease as r increases.

B.1. A lemma on symmetric undominated equilibria. This subsection uses our results above to establish that the symmetric equilibrium maximizing social surplus is the one with greatest relationship strengths x and highest reliability.

Recall our notation that consumption is $C = \kappa h(r)$ and gross profits are $\kappa g(r)$. The above calculations show we may write

$$h(r) = \frac{r^{1/\eta_C - 1}}{p N} \text{ and } g(r) = \left(\frac{1}{\eta_C} - 1 \right) \frac{1}{N r}.$$

Lemma 5. Suppose $\eta_C \in (\frac{1}{2}, 1)$ and we have two symmetric equilibria with relationship strengths $x_1 < x_2$ and reliabilities $r_1 < r_2$. Then $C(r_2) - c(x_2) > C(r_1) - c(x_1)$. That is, the higher-investment equilibrium maximizes net social surplus.

Proof. Suppose one firm slightly increases its x_{if} . The effects on social surplus can be divided into four parts:

- (1) marginal investment costs incurred by if : $c'(x_{if}) dx_{if}$ (a negative contribution);
- (2) marginal producer surplus received by if ;
- (3) marginal producer surplus received by firms other than if , which is negative;
- (4) marginal consumer surplus.

By the Nash equilibrium condition, the effects (1) and (2) cancel: the sum of these contributions is zero. By the formula for $g(r)$, (3) is zero to first order. The contribution (4) is clearly positive from the formula for $h(r)$ above. Now integrate across all firms f to obtain the effect of a small increase in all x (noting that second-order terms cancel). Our reasoning above implies a strict increase in social surplus.

Note that for $x \geq x_{\text{crit}}$, we have that $C(\rho(x, \mu))$ is concave and increasing in x while $c(x)$ is convex, so

$$C(\rho(x, \mu)) - c(x)$$

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

APPENDIX C. OMITTED PROOFS AND ANALYSIS FOR MAIN RESULTS

C.1. The reliability curve in the deep network limit. In this section we define and characterize the reliability function in the limit as $\tau \rightarrow \infty$.

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

- (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 $\{x_\tau\}_{\tau=1}^\infty \rightarrow x$ the limit $\lim_{\tau \rightarrow \infty} \rho(x_\tau, \mu_\tau) \in \rho(x)$.

Proposition 5 (Shape of the reliability correspondence). Let the complexity of production be $m \geq 2$ and the number of potential suppliers for each input be $n \geq 2$, then $\rho(x)$ is uniquely defined and 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) $\lim_{x \downarrow x_{\text{crit}}} \rho'(x) = \infty$.

C.1.1. Proof of Proposition 5. In order to prove Proposition 5 it is helpful to define and characterize a new function $\hat{\rho}(x)$. 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 fixed point of the following equation

$$\hat{r}(x) = (1 - (1 - x\hat{r}(x))^n)^m. \quad (8)$$

Lemma 6, which is proved in Section SA3.1 of the Supplementary Appendix, identifies several key properties of $\hat{\rho}(x)$.

Lemma 6. Suppose the complexity of the economy is $m \geq 2$ and there are $n \geq 1$ potential input suppliers of each firm. For $r \in (0, 1]$ define

$$\chi(r) := \frac{1 - \left(1 - r^{\frac{1}{m}}\right)^{\frac{1}{n}}}{r}. \quad (9)$$

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 $\bar{r}_{\text{crit}} = \rho(x_{\text{crit}})$;
- (v) χ is positive and quasiconvex on the domain $(0, 1]$;
- (vi) $\chi'(\bar{r}_{\text{crit}}) = 0$.

We now use Lemma 6 to prove Proposition 5.

First, consider a depth- d tree where each firm in each tier requires m kinds of inputs and has n potential suppliers of each input. The nodes at the most-upstream tier are functional for sure. 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 do not need to obtain inputs.




FIGURE 9. 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 (PC)) generates $\hat{\rho}(x)$.

Figure 9(a) depicts χ as a function of r . Since by definition the *largest* r satisfying (PC) (Eq. (8)) is the one that determines reliability $\hat{\rho}(x)$, it follows that the increasing part of the function, where $r \in [\bar{r}_{crit}, 1]$ is the part relevant for determining equilibrium reliability—see Figure 9(b), where the light gray branch is not part of $\hat{\rho}$.

A first observation is that $\tilde{\rho}(x, d)$ and its derivative converge uniformly to the function $\hat{\rho}(x)$ everywhere except near x_{crit} .

Lemma 7.

- (1) 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.
- (2) On any compact set excluding x_{crit} , the sequence $(\tilde{\rho}(x, d))_{d=1}^{\infty}$ converges uniformly to $\hat{\rho}(x)$ and $(\tilde{\rho}'(x, d))_{d=1}^{\infty}$ converges uniformly to $\hat{\rho}'(x)$.

Proof. The sequence $(\tilde{\rho}(\cdot, d))_{d=1}^{\infty}$ is a sequence of monotonically increasing⁶⁰ and differentiable functions, converging pointwise to $\hat{\rho}$. We know that $\hat{\rho}$ is continuous and has finite derivative on any compact set excluding x_{crit} . Therefore, by Dini's theorem, the functions $\tilde{\rho}(\cdot, d)$ converge uniformly to $\hat{\rho}$; the analogous statement holds with derivatives. \square

Now we will show that $\rho(x, \mu_{\tau}) \rightarrow \hat{\rho}(x)$ as $\tau \rightarrow \infty$ for any $x \neq x_{crit}$. Points (i), (ii), (iv) and (v) of Proposition 5 are then corollaries of Lemma 6.

First note that $\{\tilde{\rho}(\cdot, d)\}_{d=1}^{\infty}$ is a sequence of monotonically increasing and differentiable functions, converging pointwise to $\hat{\rho}(\cdot)$.

As an intermediate step, also note that for any $\eta > 0$, $\underline{d} > 0$ there exists $\underline{\tau}$ such that for $\tau > \underline{\tau}$,

$$\left| \mu_{\tau}(0) + \sum_{d \leq \underline{d}} \mu_{\tau}(d) ([1 - (1 - x\tilde{\rho}(x, d-1))^n]^m - \hat{\rho}(x)) \right| = \left| \mu_{\tau}(0) + \sum_{d \leq \underline{d}} \mu_{\tau}(d) (\tilde{\rho}(x, d) - \hat{\rho}(x)) \right| < \eta.$$

Indeed, as τ increases, the probability mass placed on lower d 's becomes negligible (and thus $\mu_{\tau}(0) + \sum_{d \leq \underline{d}} \mu_{\tau}(d)$ becomes negligible).

⁶⁰From Eq. (3), it is clear that increasing x increases $\tilde{\rho}(x, d)$.

Thus, for any $\tau > \underline{\tau}$,

$$\begin{aligned}
|\rho(x, \mu_\tau) - \widehat{\rho}(x)| &= \left| \mu_\tau(0) + \sum_d \mu_\tau(d) \widetilde{\rho}(x, d) - \left(\mu_\tau(0) + \sum_d \mu_\tau(d) \right) \widehat{\rho}(x) \right| \\
&= \left| \mu_\tau(0) - \mu_\tau(0) \widehat{\rho}(x) + \sum_{d \leq \underline{d}} \mu_\tau(d) [\widetilde{\rho}(x, d) - \widehat{\rho}(x)] + \sum_{d > \underline{d}} \mu_\tau(d) [\widetilde{\rho}(x, d) - \widehat{\rho}(x)] \right| \\
&\leq \left| \mu_\tau(0) + \sum_{d \leq \underline{d}} \mu_\tau(d) [\widetilde{\rho}(x, d) - \widehat{\rho}(x)] + \sum_{d > \underline{d}} \mu_\tau(d) [\widetilde{\rho}(x, d) - \widehat{\rho}(x)] \right| \\
&\leq \left| \mu_\tau(0) + \sum_{d \leq \underline{d}} \mu_\tau(d) [\widetilde{\rho}(x, d) - \widehat{\rho}(x)] \right| + \left| \sum_{d > \underline{d}} \mu_\tau(d) [\widetilde{\rho}(x, d) - \widehat{\rho}(x)] \right| \\
&\leq \eta + \left| \sum_{d > \underline{d}} \mu_\tau(d) (\widetilde{\rho}(x, d) - \widehat{\rho}(x)) \right|
\end{aligned}$$

and since $\widetilde{\rho}(x, d)$ converges pointwise to $\widehat{\rho}(x)$ for any x , then we can say that, for any $\eta' > 0$, there exists $\underline{\tau}$ such that $|\sum_{d > \underline{d}} \mu_\tau(d) (\widetilde{\rho}(x, d) - \widehat{\rho}(x))| < \eta'$ and thus that $|\rho(x, \mu_\tau) - \widehat{\rho}(x)| < \eta + \eta' = \epsilon$.

Thus, fixing $x \neq x_{\text{crit}}$, for any $\epsilon > 0$ there exists $\underline{\tau}$ such that for $\tau > \underline{\tau}$, we have $|\rho(x, \mu_\tau) - \widehat{\rho}(x)| < \epsilon$. In other words, $\rho(x, \mu_\tau)$ converges pointwise to $\widehat{\rho}(x)$ for any $x \neq x_{\text{crit}}$. Therefore, $\rho(x)$ takes the form of the continuous function $\widehat{\rho}(x)$ for $x \neq x_{\text{crit}}$ and for any fixed such x , we have $|\rho(x, \mu_\tau) - \rho(x)| < \epsilon$ for large enough τ . It follows that for any sequence $\{x_\tau\}_{\tau=1}^\infty \rightarrow x \neq x_{\text{crit}}$, $\lim_{\tau \rightarrow \infty} \rho(x_\tau, \mu_\tau) = \rho(x)$ (Part (ii) of Definition 3). Since $\rho(x)$ is single valued for $x \neq x_{\text{crit}}$, Part (i) of Definition 3 follows.

Using Lemma 6, we have thus shown points (i), (ii), (iv) and (v) of Proposition 5.

Now we analyze the case $x = x_{\text{crit}}$:

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 function⁶¹ of x whose image is $[0, 1]$. 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 $x < x_{\text{crit}}$, $\rho(x, \mu_\tau) \downarrow 0$. 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 $\lim_{\tau \rightarrow \infty} \rho(x, \mu_\tau) = \bar{r}_{\text{crit}} + \epsilon$, and since $\rho(x, \mu_\tau)$ is continuous and increasing in x , we know that (i) for every value of $r \in [0, \bar{r}_{\text{crit}}]$ there exists a sequence $\{x_\tau\}_{\tau=1}^\infty \rightarrow x_{\text{crit}}$ such that $\lim_{\tau \rightarrow \infty} \rho(x_\tau, \mu_\tau) = r$ (Part (i) of Definition 3); and (ii) every sequence $\{x_\tau\}_{\tau=1}^\infty \rightarrow x_{\text{crit}}$ will be such that $\lim_{\tau \rightarrow \infty} \rho(x_\tau, \mu_\tau) \in [0, \bar{r}_{\text{crit}}]$ (Part (ii) of Definition 3). We have thus shown that at $x = x_{\text{crit}}$ we must have $\rho(x_{\text{crit}}) = [0, \bar{r}_{\text{crit}}]$, establishing part (iii) of Proposition 5.

C.2. Proof of Lemma 2 (Sufficient condition for unique interior maximum of firm's problem). We first study the objects in the limit as $\tau \rightarrow \infty$. We will later establish convergence of the finite objects to that limit so that the properties still hold for large τ .

We establish some notation. Recall that


$$P(x_{if}; x) = (1 - (1 - x_{if}\rho(x))^n)^m.$$

For the extended domain $x_{if} \in [0, 1/\rho(x)]$, we define

$$Q(x_{ik}; x) := \frac{\partial}{\partial x_{if}} P(x_{if}; x). \quad (10)$$

and it can be expressed as $Q(x_{ik}; x) = mn(1 - (1 - x_{if}\rho(x))^n)^{m-1}(1 - x_{if}\rho(x))^{n-1}\rho(x)$.

⁶¹Indeed, it is a weighted sum of continuous and monotonically increasing functions $\widetilde{\rho}(x, d)$.

FIGURE 10. The shape of $Q(x_{if}, x)$.

We will need two steps to prove Lemma 2. The first step consists of establishing Lemma 8 on the basic shape of $Q(x_{if}, x)$.

Lemma 8. Fix any $m \geq 2$, $n \geq 2$, and $x \geq x_{\text{crit}}$ (here setting $\rho(x_{\text{crit}}) = \bar{r}_{\text{crit}}$). There are uniquely determined real numbers x_1, x_2 (depending on m, n , and x) such $0 \leq x_1 < x_2 < 1/\rho(x)$ and so that:

0. $Q(0; x) = Q(1/\rho(x); x) = 0$ and $Q(x_{if}; x) > 0$ for all $x_{if} \in (0, 1/\rho(x))$;
1. $Q(x_{if}; x)$ is increasing and convex in x_{if} on the interval $[0, x_1]$;
2. $Q(x_{if}; x)$ is increasing and concave in x_{if} on the interval $(x_1, x_2]$;
3. $Q(x_{if}; x)$ is decreasing in x_{if} on the interval $(x_2, 1]$.
4. $x_1 < x_{\text{crit}}$.

The proof of Lemma 8 is in Section SA3.2 of the Supplementary Appendix. Figure 10 illustrates the shape of $Q(x_{if}, x)$ implied by Lemma 8.

We now complete the proof of Lemma 2 by setting $\hat{x} = x_1$. By Lemma 8 property 4, the interval $(\hat{x}, x_{\text{crit}})$ is non-empty. Thus we just need show to that Assumption 2 is satisfied when $\underline{x} \in (\hat{x}, x_{\text{crit}})$. Note that by Assumption 1 $c'(0)$ is 0 and increasing and weakly convex otherwise. Since, by Lemma 8, $P'(x_{if}; x)$ 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 crossing point between the curves $P'(x_{if}; x)$ and $c'(x_{if} - \underline{x})$. This crossing point corresponds to the first-order condition⁶² $P'(x_{if}; x) - c'(x_{if} - \underline{x})$, yielding the unique maximizer of $\Pi(x_{if}; x)$, as illustrated in Fig. 11. If such a crossing does 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 single crossing.

Note that $P(x_{if}; x, \mu_\tau) = \mu_\tau(0) + \sum_d \mu_\tau(d)(1 - (1 - x_{if}\tilde{\rho}(x, d-1))^n)^m$, while $Q(x_{if}; x, \mu_\tau) = P'(x_{if}; x, \mu_\tau) = \sum_d \mu_\tau(d)mn(1 - (1 - x_{if}\tilde{\rho}(x, d-1))^n)^{m-1}(1 - x_{if}\tilde{\rho}(x, d-1))^{n-1}\tilde{\rho}(x, d-1)$ and, taking $G = 1$ for simplicity of exposure, $\Pi(x_{if}; x, \mu_\tau) = P(x_{if}; x, \mu_\tau) - c(x_{if} - \underline{x})$. These are continuous functions of $\tilde{\rho}(x, d)$. Recall that for all $x \neq x_{\text{crit}}$, $\tilde{\rho}(x, d) \rightarrow_d \rho(x)$ (and $\rho(x, \mu_\tau) \rightarrow_\tau \rho(x)$). It then follows that $P(x_{if}; x, \mu_\tau) \rightarrow_\tau P(x_{if}; x)$ since $\lim_{\tau \rightarrow \infty} \sum_{d > \underline{d}} \mu_\tau(d) = 1$ for any \underline{d} (and $\lim_{\tau \rightarrow \infty} \mu_\tau(0) = 0$). It also follows that $Q(x_{if}; x, \mu_\tau) \rightarrow_\tau Q(x_{if}; x)$ for the same reason. We then immediately have that $\Pi(x_{if}; x, \mu_\tau) \rightarrow_\tau \Pi(x_{if}; x)$. We can then conclude that for τ large enough,

⁶²We can also scale the term $P'(x_{if}; x)$ by a multiplying factor G , but this will not affect the properties discussed before and thus the uniqueness of the crossing.




FIGURE 11. $P'(x_{if}; x)$ is in blue and $c'(x_{if} - x)$ is in red. There can be only one crossing point between the two curves, which corresponds to the maximizer of $\Pi(x_{if}; x)$.

there will also be at most a single crossing between the curves $Q(x_{if}; x, \mu_\tau)$ and $c'(x_{if} - x)$ and thus the results stated in the limit also hold for τ large enough.

C.3. Proof of Proposition 1 (Discontinuity in reliability). Part (i):

The fact that $\rho(x, \mu_\tau) \rightarrow 0$ for $x < x_{\text{crit}}$ follows from Definition 3 and Proposition 5(ii). From Lemma 1, $C^*(x, \mu_\tau) = \kappa h(\rho(x, \mu_\tau))$ with h an increasing and continuous function with $h(0) = 0$. It then follows that $C^*(x, \mu_\tau) \rightarrow 0$.

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, $\rho(x, \mu_\tau) > \bar{r}_{\text{crit}} > 0$. It then follows that $C^*(x, \mu_\tau) = \kappa h(\rho(x, \mu_\tau)) > \kappa h(\bar{r}_{\text{crit}}) = \underline{C} > 0$ (from which it follows that $C^*(x, \mu_\tau) \geq \underline{C} > 0$).

C.4. Proof of Proposition 2 (Social planner's solution). Recall that $x^{SP}(\kappa, \mu_\tau)$ is the set of all values of x maximizing the planner's objective.

By Proposition 1(i), $x_{\text{crit}} > 0$, and for all $x < x_{\text{crit}}$, we have $C^*(x) = \lim_{\tau \rightarrow \infty} C^*(x, \mu_\tau) = 0$. Thus no value of $x \in (0, x_{\text{crit}})$ can be a solution to the social planner's problem in the limit. Further, we show in the proof of Proposition 5 that $\rho(x, \mu_\tau)$ converges pointwise to $\rho(x)$. Thus by Lemma 1 consumption also converges pointwise to its limit value and for all τ sufficient large there can be no solution to the social planner's problem with $0 < x < x_{\text{crit}}$.

Further, for all values of κ sufficiently small, and for all values of $x \in [0, 1]$,

$$c(x_{\text{crit}}) > \kappa \geq \kappa \rho(x, \mu_\tau).$$

Thus for sufficiently small values of κ the unique solution to the social planner's problem is $x^{SP}(\kappa) = \{0\}$.

Next note that for all values of κ sufficiently large,

$$c(x_{\text{crit}}) < \kappa r_{\text{crit}},$$

and the social planner strictly prefers choosing $x = x_{\text{crit}}$ to $x = 0$. Thus, when κ is sufficiently high all values in $x^{SP}(\kappa)$ will be weakly greater than x_{crit} .

Define $\kappa_{\text{crit}}(\mu_\tau) := \sup_{\kappa: 0 \in x^{SP}(\kappa, \mu_\tau)} \kappa$. As the social planner's unique solution is $x = 0$ for values of κ sufficiently low, and $x = 0$ is not a solution for values of κ sufficiently high, $\kappa_{\text{crit}}(\mu_\tau)$ is a finite, strictly positive number. This establishes part (i).

We now show that $1 \notin x^{SP}(\kappa, \mu_\tau)$ for all values of κ and τ . As κ is bounded, $\lim_{x \rightarrow 1} \frac{1}{\kappa} c'(x) = \infty$ by our assumption that $c'(x) = \infty$. Moreover, as $n \geq 2$, we have that $x_{\text{crit}} < 1$, so $\rho'(x)$ is bounded for values of x near 1. Thus, by the chain rule and Lemma 1, $C^{*'}(x)$ is also bounded. Hence, at $x = 1$ the social planner can always do a little better by reducing investment.

Thus for values of $\kappa > \kappa_{\text{crit}}(\mu_\tau)$ all solutions to the social planner's problem must be interior and so the following first-order condition must hold:

$$\kappa h'(\rho(x, \mu_\tau)) \rho'(x, \mu_\tau) = c'(x).$$

By Lemma 6, part (vi), $\lim_{x \downarrow x_{\text{crit}}} \lim_{\tau \rightarrow \infty} \rho'(x, \mu_\tau) = \infty$. As by assumption $c'(x)$ is bounded for interior values of x , and as h is an increasing function there cannot be an interior equilibrium in the limit as τ gets large at $x = x_{\text{crit}}$. This implies that in the limit as τ gets large all solutions to the social planner's problem for $\kappa > \kappa_{\text{crit}}$ are at values of x strictly greater than x_{crit} . As $\rho'(x, \mu_\tau)$ converges to $\rho'(x)$ pointwise, for all τ sufficiently large we then have $\lim_{x \downarrow x_{\text{crit}}} \kappa h'(\rho(x, \mu_\tau)) \rho'(x, \mu_\tau) > c'(x)$ establishing part (ii). Furthermore, as $\lim_{x \downarrow x_{\text{crit}}} \kappa h'(\rho(x, \mu_\tau)) \rho'(x, \mu_\tau) > c'(x)$, for all τ sufficiently large, there cannot be a solution to the social planner's problem at $x = x_{\text{crit}}$ for $\kappa = \kappa_{\text{crit}}$ —the social planner could always do better by increasing x a little. This establishes part (iii).

C.5. Proof of Proposition 6 (Unique symmetric undominated equilibrium).

Proposition 6. Fix any $n \geq 2$ and $m \geq 3$, and any κ and g consistent with the maintained assumptions. 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 $[\bar{r}_{\text{crit}} - \epsilon, 1]$. Here $\chi(r) = \frac{1 - (1 - r^{\frac{1}{m}})^{\frac{1}{n}}}{r}$ is defined as in Lemma 6.

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

We begin by sketching the idea of the argument. The following definitions will be helpful.

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

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

We first want to show that there is a range $[\bar{r}_{\text{crit}} - \epsilon, 1]$, for which there is at most one solution to $MB(x; r, \kappa) = MC(x)$ (OI) and $x = \chi(r)$ (PC) simultaneously. Consider any positive symmetric equilibrium x with reliability $r = \rho(x)$. Recall that physical consistency in the limit $\tau \rightarrow \infty$ entails

$$r = (1 - \underbrace{(1 - x r)^n}_m)^m. \quad (\text{PC})$$

probability a given input cannot be acquired

By Lemma 6(i), if $x < x_{\text{crit}}$, then $r = 0$ and marginal benefits from investing are 0, an impossibility by the optimal investment condition. Therefore, $x \geq x_{\text{crit}}$ and so by Lemma 6(iv) we have $x = \chi(r)$. (See see Panel (a) of Figure 12.)

Now, given Assumption 2, the optimal investment condition OI says that we have

$$MB(x; r, \kappa) = MC(x),$$

where

$$MB(x_{if}; r, \kappa) = \kappa g(r) r n (1 - x r)^{n-1} m (1 - (1 - x r)^n)^{m-1} \quad (13)$$

and $MC(x) = c'(x - \underline{x})$. Since we recently deduced $x = \chi(r)$, we can substitute out x to find that the following equation holds:

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

We can now define two auxiliary functions of r :

$$\begin{aligned} \mathbf{b}(r) &= MB(\chi(r); r, \kappa) \\ \mathbf{c}(r) &= MC(\chi(r)) \end{aligned}$$

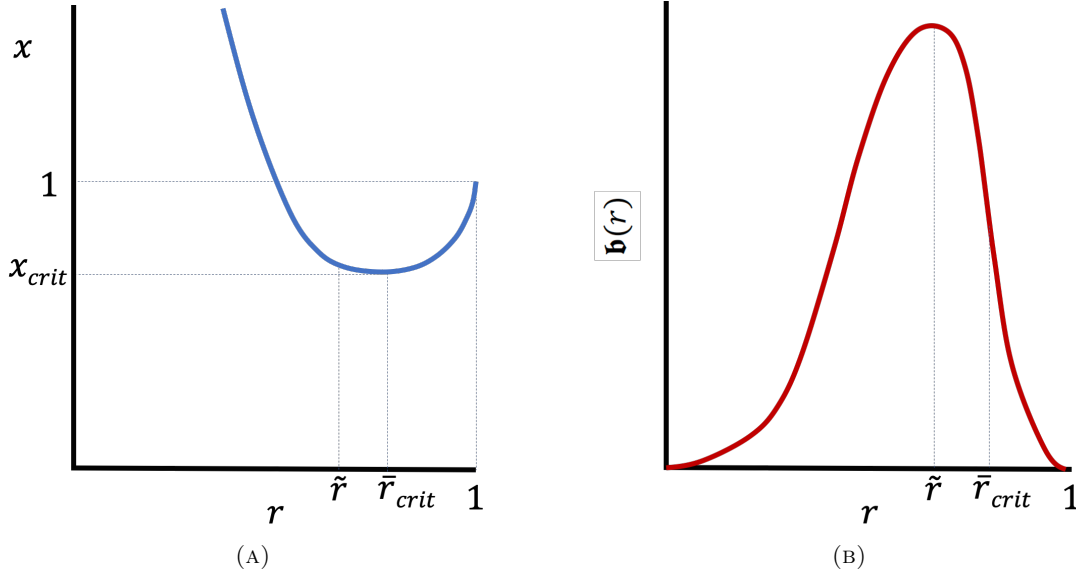


FIGURE 12. Panel (a) shows the relationship between r and x implied by physical consistency. Panel (b) plots the function $b(r)$ discussed in the proof.

so that we can simply study the equation

$$b(r) = c(r). \quad (14)$$

More explicitly,

$$\underbrace{\kappa g(r) m n r^{2-\frac{1}{m}} \left(1 - r^{1/m}\right)^{1-\frac{1}{n}}}_{b(r)} = c' \left(\underbrace{\frac{1 - \left(1 - r^{\frac{1}{m}} - x\right)^{\frac{1}{n}}}{r}}_x \right) \quad (15)$$

We wish to show that there is a range $[\bar{r}_{crit} - \epsilon, 1]$ for which there is at most one solution to equation (14).⁶³ The right-hand side of the equation is increasing in r for $r \in [\bar{r}_{crit}, 1]$.⁶⁴ If we could establish that the left-hand side, which we call $b(r)$, is decreasing in r , it would follow that there is at most a unique r solving (15) for $r \in [\bar{r}_{crit}, 1]$. Moreover, by the continuity of $b(r)$ and $c(r)$, it would follow immediately that for some small $\epsilon > 0$, there is also at most a single solution over the range $[\bar{r}_{crit} - \epsilon, 1]$. Unfortunately, $b(r)$ is not decreasing in r . However, we can show that it *is* decreasing in r for $r \geq \bar{r}_{crit}$, which is sufficient.⁶⁵ Panel (b) of Figure 12 gives a representative depiction of $b(r)$, reflecting that it is decreasing to the right of $\tilde{r} < \bar{r}_{crit}$.

By plugging in $x = \chi(r)$ into (13), we have

$$b(r) = \kappa g(r) m n r^{2-\frac{1}{m}} \left(1 - r^{\frac{1}{m}}\right)^{1-\frac{1}{n}}. \quad (16)$$

The following Lemma completes the proof.

Lemma 9. b is strictly decreasing on the domain $[\bar{r}_{crit}, 1]$.

⁶³We 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.

⁶⁴This follows directly: χ is increasing on that domain and c' is increasing by assumption. (see Panel (a) of Figure 12).

⁶⁵We do this by showing that the global maximum of $b(r)$ is achieved at a number \tilde{r} that we can prove is smaller than \bar{r}_{crit} .

To prove lemma 9 we write $\mathbf{b}(r)$ as a product of two pieces, $\alpha(r) := \kappa g(r)$ and

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

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

Lemma 11. For all $n \geq 2$ and $m \geq 3$, we have that $\hat{r} < r_{\text{crit}}$.

Lemmas 10 and 11 are proved in Sections SA3.3 and SA3.4 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 9.

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 $r \in [\bar{r}_{\text{crit}} - \epsilon, 1]$.

C.6. Proof of Theorem 1 (Classification of regimes as κ varies). Recall the optimal investment (OI) condition:

$$MB(x; r, \kappa) = MC(x). \quad (17)$$

Denote by $BR(r, \kappa) = \{x : MB(x; r, \kappa) = MC(x)\}$ the set of effort levels such that OI is satisfied at a given r . Let $\overline{BR}(r, \kappa) = \max\{x : x \in BR(r, \kappa)\}$ be the maximal element of the set $BR(r, \kappa)$. Note however that by Lemma 2 and Assumption 2, the best response is unique for $r \geq \bar{r}_{\text{crit}}$ and thus $BR(r, \kappa)$ contains a single element for $r \geq \bar{r}_{\text{crit}}$.

First note that if $\kappa' > \kappa$, then $\overline{BR}(r, \kappa') > \overline{BR}(r, \kappa)$. This is evident from condition OI, since the left-hand side of Equation (17) is increasing in κ whereas the right-hand side is constant (does not depend on κ).

Lemma 12. (i) $\overline{BR}(0, \kappa) = 0$ and (ii) for all κ such that there exists r (which must be weakly greater than \bar{r}_{crit}) solving $MB(\chi(r); r, \kappa) = MC(\chi(r))$, there exists $\epsilon > 0$ such that $\overline{BR}(r, \kappa)$ is strictly decreasing in r on the domain $[\bar{r}_{\text{crit}} - \epsilon, 1]$.

Proof of Lemma 12. We know that $MB(x; 0, \kappa) = 0$ when $r = 0$ (see Eq. (13)) and thus, from equation OI, the only best response is $x = 0$. Thus $BR(0, \kappa) = \overline{BR}(0, \kappa) = 0$.

Now consider part (ii). Fixing κ , we know from Proposition 6 that there is $\epsilon > 0$ such that there can be at most one solution r^* to $MB(\chi(r); r, \kappa) = MC(\chi(r))$ over the range $[\bar{r}_{\text{crit}} - \epsilon, 1]$. This corresponds to a point such that both $MB(x; r, \kappa) = MC(x)$ and $x = \chi(r)$ are satisfied. It follows that $\overline{BR}(r, \kappa)$ can intersect the curve $x = \chi(r)$ at only one point on that interval $[\bar{r}_{\text{crit}} - \epsilon, 1]$ (see Fig. 13).⁶⁶ We will show that $\overline{BR}(r, \kappa)$ is strictly decreasing in r on the interval $[\bar{r}_{\text{crit}} - \epsilon, 1]$, by showing that if this were not so, we would have a contradiction to Proposition 6. To see this, suppose there are \tilde{r} and \hat{r} such that $1 > \tilde{r} > \hat{r} > \bar{r}_{\text{crit}} - \epsilon$ and $\overline{BR}(r, \kappa)$ has a positive slope in r on (\tilde{r}, \hat{r}) . There would then be a cost function \hat{c} satisfying our assumptions such that there is a point of tangency between the functions $\overline{BR}(\cdot, \kappa, \hat{c})$ and $\chi(\cdot)$ for some $r \geq r_{\text{crit}}$.⁶⁷ Then, by increasing κ

⁶⁶This occurs when κ is high enough (else there cannot be a solution to $MB(x; r, \kappa) = MC(x)$ since the marginal benefits are simply lower than the marginal cost for any investment level).

⁶⁷The intuition is that as we scale down the cost function, \overline{BR} moves leftward, as can be seen by inspecting the (OI) condition. The only subtlety is that investing a positive amount may at some points become dominated by investing 0, so that \overline{BR} could jump down to 0. However, it is always possible to avoid this problem by choosing \hat{c} so that inframarginal costs are sufficiently low while matching derivatives so that the slope of \overline{BR} remains positive for some $r \in [\bar{r}_{\text{crit}} - \epsilon, \hat{r})$. (Indeed, Proposition 6 applies to a more general environment where negative cost functions are allowed, as long as the other assumptions are satisfied, so we can simply shift c by a constant and get a contradiction.)




FIGURE 13. The curves $\overline{BR}(r, \kappa)$ and $x = \chi(r)$ represented on the same graph for some large enough value of κ . The curve $\overline{BR}(r, \kappa)$ is strictly decreasing in r over a range $[\tilde{r}, 1]$ where \tilde{r} is some number satisfying $\tilde{r} < \bar{r}_{crit}$. If it were not the case, there would be a smaller value of κ for which there would be two crossings of the curves over the range $[\tilde{r}, 1]$, contradicting Proposition 6.

slightly, we would obtain two nearby points of intersection between the curves $\overline{BR}(\cdot, \kappa, \hat{c})$ and $\chi(\cdot)$ for $r \in [\bar{r}_{crit} - \epsilon', 1)$. This yields the desired contradiction to Proposition 6 (which applies to the modified environment constructed in the proof) and completes the argument. \square

Lemma 13. $\overline{BR}(r, \kappa, \mu_\tau)$ converges pointwise to $\overline{BR}(r, \kappa)$ for any r .

Proof of Lemma 13. $\overline{BR}(r, \kappa, \mu_\tau) = \max\{x : MB(x; r, \kappa, \mu_\tau) = MC(x)\}$ where

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

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 $\tau \rightarrow \infty$,

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

and thus that

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

for any r . \square

Part (i):

In the limit $\tau \rightarrow \infty$, let $\underline{\kappa}$ be the smallest κ such the graph of $\overline{BR}(r, \kappa)$ intersects the graph of ρ (both viewed as sets of points (x, r)).⁶⁸ Lemma 12 implies that all points of intersection satisfy $x = x_{crit}$ and $r \in \rho(x_{crit})$. By the properties of $\overline{BR}(r, \kappa)$ in Lemma 12, any such intersection will be such that $r < \bar{r}_{crit}$ since $\overline{BR}(r, \kappa)$ is decreasing in r over a range $[\bar{r}_{crit} - \epsilon, 1]$. Thus, $\underline{\kappa}$ is such that $\overline{BR}(r, \underline{\kappa})$ just touches $\rho(x)$ at $x = x_{crit}$.⁶⁹ If there are multiple points of intersection, we take the one with the highest r and we label this r by \underline{r}_{crit} .

⁶⁸By inspecting the (limit) optimal investment condition, we can see that for any $r \in (0, 1)$ and any $x_0 < 1$ if κ is large enough, all solutions to the OI problem satisfy $x > x_0$. This implies that there is such a $\underline{\kappa}$.

⁶⁹If the first intersection occurred for $r \geq \bar{r}_{crit}$, then by Lemma 12 the slope in r of $\overline{BR}(r, \kappa)$ at the point of intersection would be negative, and thus clearly it cannot be tangent to the ρ correspondence. Decreasing κ slightly would then still yield an intersection, a contradiction to the definition of $\underline{\kappa}$.

Since $\overline{BR}(0, \kappa) = 0$, then for any $\kappa < \underline{\kappa}$, 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 13), 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) where $x = 0$ and $r = 0$.

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 $\kappa = \underline{\kappa}$), then it follows that in the limit $\tau \rightarrow \infty$, there will be 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 $r \in (\underline{r}_{\text{crit}}, \bar{r}_{\text{crit}})$.

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 $x \in [x_{\text{crit}} - \epsilon, x_{\text{crit}} + \epsilon]$ and $r \in [\underline{r}_{\text{crit}} - \epsilon, \bar{r}_{\text{crit}} + \epsilon]$.

In the limit $\tau \rightarrow \infty$, 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 upper 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, in the limit $\tau \rightarrow \infty$, as κ increases at values above $\bar{\kappa}$, $\overline{BR}(r, \kappa)$ 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 $x > x_{\text{crit}}$. 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

C.7. Proof of Corollary 1 (Comparative statics in baseline institutional quality). 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. \square

C.8. 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

C.9. Proof of Proposition 4 (Weakest links). *Part (i):* Let \mathcal{P} be a directed path of length T from node T to node 1 and denote product i by 1. Since product 1 is critical, then following a shock $\epsilon > 0$ to γ_{ij} , $r_1^{*'} = 0$, where the ‘prime’ notation denotes the equilibrium quantity after the shock.

For any product $t+1$ that sources input t and such that $r_t^{*'} = 0$, we have that

$$\begin{aligned} r_{t+1}^{*'} &= \prod_{l \in \mathcal{I}(t+1)} (1 - (1 - x_{t+1,l}^* r_l^{*'})^{n_{t+1,l}}) \\ &= 0 \end{aligned}$$

since $t \in \mathcal{I}(t+1)$.

Since $r_1^{*'} = 0$, it then follows by induction that the production of all products $t \in \mathcal{P}$ will fail.

Part (ii): Suppose production of some product $i \in \mathcal{I}^{SC}$ is critical and consider another product $k \in \mathcal{I}^{SC}$ that is an input for the production of product i (that is, $k \in \mathcal{I}(i)$). As an investment equilibrium x^* is being played, the strength x_{kj} of a link from a producer of product k to a supplier of input $j \in \mathcal{I}(k)$ must satisfy the following condition

$$MB_{kj} = \kappa g(r_k) \prod_{l \in \mathcal{I}(k), l \neq j} (1 - (1 - x_{kl} r_l)^{n_{kl}}) n_{kj} (1 - x_{kj} r_j)^{n_{kj}-1} r_j = \gamma_{kj} \tilde{\mathcal{C}}(x_{kj} - \underline{x}_{kj}) = MC_{kj}$$

Rearranging this equation yields

$$\frac{\kappa g(r_k^*)}{\gamma_{kj}} \prod_{l \in \mathcal{I}(k), l \neq j} (1 - (1 - x_{kl}^* r_l^*)^{n_{kl}}) n_{kj} r_j^* = \frac{\tilde{\mathcal{C}}(x_{kj}^* - \underline{x}_{kj})}{(1 - x_{kj}^* r_j^*)^{n_{kj}-1}}.$$

The right hand side is strictly increasing in x_{kj}^* , while the left hand side is constant.

Consider now a shock $\epsilon > 0$ that changes the value of γ_{kj} to $\gamma'_{kj} = \gamma_{kj} + \epsilon$. This strictly reduces $\kappa g(r_k^*)/\gamma_{kj}$, and hence the new equilibrium investment level satisfies $x_{kj}^{*'} < x_{kj}^*$. This in turn implies that $r_k^{*'} < r_k^*$, and so

$$r_i^{*'} = \prod_{l \in \mathcal{I}(i)} (1 - (1 - x_{i,l}^* r_l^{*'})^{n_{i,l}}) \quad (18)$$

$$< \prod_{l \in \mathcal{I}(i)} (1 - (1 - x_{i,l}^* r_l^*)^{n_{i,l}}) \quad (19)$$

$$= r_i^* \quad (20)$$

since $k \in \mathcal{I}(i)$. Thus $r_i^{*'} = 0$ and the production of product i fails.

Now since both products k and i are part of a strongly connected component, there is also a directed path from k to i . From part (i), it follows that every product t on such a directed path (i.e. every product that uses input i either directly or indirectly through intermediate products) will also have $r_t^{*'} = 0$. This is true namely for product k and thus $r_k^{*'} = 0$. We therefore conclude that, following a small decrease in its sourcing effort from the initial x_k^* , production of product k fails. Product k was thus necessarily critical and we must have had x_k^* be critical.

Proceeding similarly for any other product $e \in \mathcal{I}(k)$ such that $e \in \mathcal{I}^{SC}$, we get that e is also critical. By induction it follows that all products in a strongly connected component \mathcal{I}^{SC} of the product interdependencies graph are critical when one of them is. We conclude that either production of all the products in \mathcal{I}^{SC} are critical or else production of all the products in \mathcal{I}^{SC} are noncritical. \square

APPENDIX D. INTERDEPENDENT SUPPLY NETWORKS AND CASCADING FAILURES

We now posit an interdependence among supply networks wherein each firm's profit depends on the *aggregate* level of output in the economy, in addition to the functionality of the suppliers with whom it has supply relationships. Formally, suppose, $\kappa_{\mathfrak{s}} = K_{\mathfrak{s}}(Y)$, where $K_{\mathfrak{s}}$ is a strictly increasing function and Y is the integral across all sectors of equilibrium output:

$$Y = \int_S \rho(x_{\mathfrak{s}}^*) d\Phi(\mathfrak{s}).$$

Here we denote by $x_{\mathfrak{s}}^*$ the unique positive equilibrium in sector \mathfrak{s} . The output in the sector is the reliability in that sector, $\rho(x_{\mathfrak{s}}^*)$.

The interpretation of this is as follows: When a firm depends on a different sector, a specific supply relationship is not required, so the idiosyncratic failure of a given producer in the different sector does not matter—a substitute product can be readily purchased via the market. Indeed, it is precisely when substitute products are not readily available that the supply relationships we model are important. However, if some sectors experience a sudden drop in output, then other sectors suffer. They will not be able to purchase inputs, via the market, from these sectors in the same quantities or at the same prices. For example, if financial markets collapse, then the productivity of many real businesses that rely on these markets for credit are likely to see their effective productivity fall. In these situations, dependencies will result in changes to other sectors' profits even if purchases are made via the market. Our specification above takes interdependencies to be highly symmetric, so that only aggregate output matters, but in general these interdependencies would correspond to the structure of an intersectoral input-output matrix, and K would be a function of sector level outputs, indexed by the identity of the sourcing sector.

This natural interdependence can have very stark consequences. Consider an economy characterized by a distribution Ψ in which the subset of sectors with $m \geq 2$ has positive measure, and some of these have positive equilibria. Suppose that there is a small shock to \underline{x} . As already argued, this will directly cause a positive measure of sectors to fail. The failure of the fragile sectors will cause a reduction in aggregate output. Thus $\kappa_{\mathfrak{s}} = K_{\mathfrak{s}}(Y)$ will decrease in other sectors *discontinuously*. This will take some other sectors out of the robust regime. Note that this occurs due to the other supply chains failing and not due to the shock itself. As these sectors are no longer robust, they topple too following an infinitesimal shock to \underline{x} . Continuing this logic, there will be a domino effect that propagates the initial shock. This domino effect could die out quickly, but need not. A full study of such domino effects is well beyond our scope, but the forces in the very simple sketch we have presented would carry over to more realistic heterogeneous interdependencies.

Fig. 14 shows⁷⁰ how an economy with 100 interdependent sectors responds to small shocks to \underline{x} . In this example, sectors differ only in their initial κ 's. The technological complexity is set to $m = 5$ and the number of potential suppliers for each firm is set to $n = 3$. The cost function⁷¹ for any firm if is $c(x_{if} - \underline{x}) = \frac{0.01}{(1 - (x_{if} - \underline{x}))^2}$ while the gross profit function is $g(\rho(x)) = 5(1 - \rho(x))$. This setup yields values $\underline{\kappa} = 0.963$ and $\bar{\kappa} = 3.585$ delimiting the region corresponding to critical (and therefore fragile) equilibria, as per Theorem 1.

In Fig. 14(a), the productivity shifter of a given sector is distributed uniformly, i.e. $\kappa_{\mathfrak{s}} \sim U(\underline{\kappa}, 25)$, so that many sectors have high enough productivity to be in a robust equilibrium while a small fraction have low enough productivity to be in a fragile equilibrium. A small shock to the \underline{x} of all sectors thus causes the failure of the fragile sectors (there are initially 12 of them). This then

⁷⁰Note that in this example, the cascade dynamics is as follows: At step 1, firms in sectors with a κ in the fragile range fail due to an infinitesimal shock to \underline{x} . The initial economy-wide output Y_1 is then decreased to Y_2 and the κ 's are updated using an updating function $K(Y)$ increasing in Y . Only then, are the firms in the surviving sectors allowed re-adjust x_{if} . At step 2, infinitesimal shocks hit again and the firms newly found in the fragile regime fail. This process goes on at each step until no further firm fails, at which point the cascade of failures stops.

⁷¹For simplicity, we set $\underline{x} = 0$. An infinitesimal shock to \underline{x} has the effect of causing the firms of sectors in the fragile regime to fail, but does not affect the value of \underline{x} , which remains at 0.

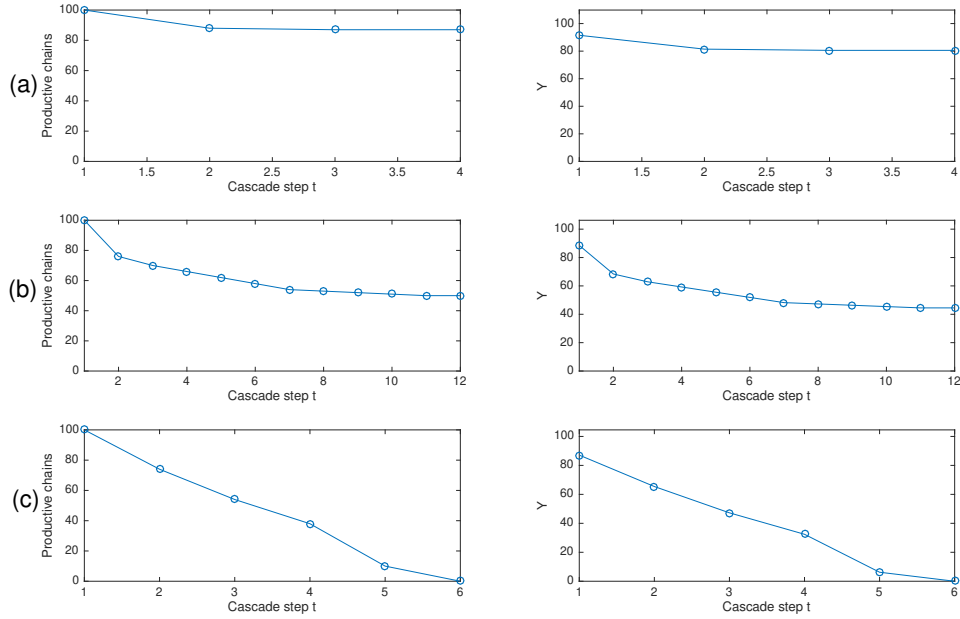


FIGURE 14. Number of sectors that remain productive (left) and economy-wide output Y (right) for each step of a cascade of failures among 100 interdependent sectors. For all sectors: $n = 3$, $m = 5$, $c(x_{if} - \underline{x}) = \frac{0.01}{(1 - (x_{if} - \underline{x}))^2}$, $g(\rho(x)) = 5(1 - \rho(x))$. This yields $\underline{\kappa} = 0.963$ and $\bar{\kappa} = 3.585$. In row (a), 100 sectors have κ_s initially distributed according to $U(\underline{\kappa}, 25)$; In row (b), 100 sectors have κ_s initially distributed according to $U(\underline{\kappa}, 13)$; In row (c), 100 sectors have κ_s initially distributed according to $U(\underline{\kappa}, 10)$.

decreases the output Y across the whole economy, but only to a small extent (as seen in the right panel). The resulting decrease in the productivities of the robust sectors is thus only enough to bring one robust sector into the fragile regime and thus to cause it to fail as well upon a small shock. In the end, a total of only 13 sectors have failed.

In contrast, Fig. 14(b) shows an economy where $\kappa_s \sim U(\underline{\kappa}, 13)$, so that more sectors have low enough productivity to be in a fragile equilibrium. A small shock to the \underline{x} of all sectors causes the failure of the fragile sectors (now initially 24). These have a larger effect on decreasing the output Y across the whole economy (as seen in the right panel). The resulting decrease in the productivities of the robust sectors is now enough to bring many of them into the fragile regime and to cause them to fail upon an infinitesimal shock to \underline{x} . This initiates a cascade of sector failures, ultimately resulting in 50 sectors ceasing production.

Fig. 14(c) shows an economy where $\kappa_s \sim U(\underline{\kappa}, 10)$, so that even more sectors have low enough productivity to be in a fragile equilibrium. A small shock to the \underline{x} of all sectors causes the failure of the fragile sectors (now initially 26) and this initiates a cascade of sector failures which ultimately brings down all 100 sectors of the economy.

The discontinuous drops in output caused by fragility, combined with the simple macroeconomic interdependence that we have outlined, come together to form an amplification channel reminiscent, e.g., of Elliott, Golub, and Jackson (2014) and Baqaee (2018). Thus, the implications of those studies apply here: both the cautions regarding the potential severity of knock-on effects, as well as the importance of preventing first failures before they can cascade.




FIGURE 15. The dynamics of unraveling (with the same parameters as in Figure 4, as discussed in Section E).

APPENDIX E. HOW PRODUCTION UNRAVELS WHEN RELATIONSHIP STRENGTH IS TOO LOW

Figure 3(b) shows that when x drops below x_{crit} , the mass of firms that can consistently function falls discontinuously to $\rho(x) = 0$. While we will typically just work with the fixed point as the outcome of interest, the transition will not be instantaneous in practice. How then might the consequences of a shock to x actually play out?

In Figure 15, we work through a toy illustration to shed some light on the dynamics of collapse. Using the same parameters as our previous example, suppose relationship strength starts out at $x = 0.8$. The higher curve in panel (a) is $\mathcal{R}(\cdot; x)$ for this value of x . The reliability of the economy here is r_0 , a fixed point of \mathcal{R} , which is the mass of functioning firms. Now suppose that a shock occurs, and all relationships become weaker, operating with the lower probability $x = 0.7$. The \mathcal{R} curve now shifts, becoming the lower curve.

To consider the dynamics of how production responds, we must specify a few more details. We sketch one dynamic, and only for the purposes of this subsection. We interpret idiosyncratic link operation realizations as whether a given relationship works in a given period. Before the shift in x , a fraction r_0 of the firms are functional. Let $\tilde{\mathcal{F}}(0)$ be the random set of functional firms at the time of the shock to x . Now x shifts to 0.7; we can view this as a certain fraction of formerly functional links failing, at random. Then firms begin reacting over a sequence of stages. Let us suppose that at stage s a firm can source its inputs if it has a functional link to a supplier who was functional in the last stage, $s - 1$.⁷² Let $\tilde{\mathcal{F}}(s)$ be the set of these functional firms. By the same reasoning as in the previous subsection, we can see that the mass of $\tilde{\mathcal{F}}(s)$, which we call r_s , is $\mathcal{R}(r_{s-1})$. Iterating the process leads to more and more firms being unable to produce as their suppliers fail to deliver essential inputs. After stage 1, the first set of firms that lost access to an essential input run out of stock and are no longer functional. This creates a new set of firms that cannot access an essential input, and these firms will be unable to produce at the end of the subsequent stage, and so on. The mass of each r_s can be described via the graphical procedure of Figure 15: take steps between the \mathcal{R} curve and the 45-degree line.

This discussion helps make three related points. First, even though the disappearance of the positive fixed point—and thus the possibility of a positive mass of consistently functional firms—is

⁷²For example, firms might hold inventories that enable them to maintain production for a certain amount of time, even when unable to source an essential input. Even if firms engage in just-in-time production and do not maintain inventories of essential inputs, there can be a lag between shipments being sent and arriving.

sudden, the implications can play out slowly under natural dynamics.⁷³ The first few steps may look like a few firms being unable to produce, rather than a sudden and total collapse of output.

The second point is more subtle. Suppose that when the dynamic of the previous paragraph reaches r_2 , the shock is reversed, x again becomes 0.8, and \mathcal{R} again becomes the higher curve. Then, with some supply links reactivated, some of the firms that were made non-functional as the supply chain unravelled will become functional again, and this will allow more firms to become functional, and so on. Such dynamics could take the system back to the r_0 fixed point if sufficiently many firms remain functional at the time the shock is reversed. Thus, our theory predicts that *sufficiently persistent* shocks to relationship strength lead to eventual collapses of production, but, depending on the dynamics, the system may also be able to recover from sufficiently transient shocks.

The third point builds on the second. Suppose that a shock is anticipated and expected to be temporary. Then firms may take actions that slow the unravelling to reduce their amount of downtime. For example, they may build up stockpiles of essential inputs. If all firms behave in this way, the dynamics can be substantially slowed down and the possibility of recovery will improve.

Having illustrated some of the basic forces and timing involved in unraveling, we do not pursue here a more complete study of the dynamics of transient shocks, endogenous responses, etc.—an interesting subject in its own right. Instead, from now on we will focus on the size, $\rho(x)$, of the consistent functional set, which is the steady-state outcome under a relationship strength x .

APPENDIX F. INTERPRETATION OF INVESTMENT

F.1. Effort on both the extensive and intensive margins. This section supports the claims made in Section 6.3.2 that our model is easily extended to allow firms to make separate multi-sourcing effort choices on the intensive margin (quality of relationships) and the extensive margin (finding potential suppliers).

Suppose a firm if chooses efforts $\hat{e}_{if} \geq 0$ on the extensive margin and effort $\tilde{e}_{if} \geq 0$ on the intensive margin, and suppose that $x_{if} = h(\hat{e}_{if}, \tilde{e}_{if})$. Let the cost of investment be a function of $\hat{e}_{if} + \tilde{e}_{if}$ instead of y_{if} . This firm problem can be broken down into choosing an overall effort level $e_{if} = \hat{e}_{if} + \tilde{e}_{if}$ and then a share of this effort level allocated to the intensive margin, with the remaining share allocated to the extensive margin. Fixing an effort level e , a firm will choose $\hat{e}_{if} \in [0, e]$, with $\tilde{e}_{if} = e - \hat{e}_{if}$, to maximize x_{if} . Let $\hat{e}_{if}^*(e)$ and $\tilde{e}_{if}^*(e) = e - \hat{e}_{if}^*(e)$ denote the allocation of effort across the intensive and extensive margins that maximizes x_{if} given overall effort e . Given these choices, define $h^*(e) := h(\hat{e}_{if}^*(e), \tilde{e}_{if}^*(e))$. As h^* is strictly increasing in e , choosing e is then equivalent to choosing x_{if} directly, with a cost of effort equal to $c(h^{*-1}(e))$. Thus, as long as the cost function $\hat{c}(e) := c(h^{*-1}(e))$ continues to satisfy our maintained assumptions on c , everything goes through unaffected.

F.2. A richer extensive margin model. In the previous subsection we gave an extensive margin search effort interpretation of x_{if} . In some ways this interpretation was restrictive. Specifically, it required there to be exactly n suppliers capable for supplying the input and that each such supplier be found independently with probability x_{if} . This alternative interpretation is a minimal departure from the intensive margin interpretation, which is why we gave it. However, it is also possible, through a change of variables, to see that our model encompasses a more general and standard search interpretation.

Fixing the environment a firms faces, specifically the probability other firms successfully produce $r > 0$ and a parameter n that will index the ease of search, suppose we let each firm if choose directly the probability that, through search, it finds an input of given type. When $r = 0$ we

⁷³Indeed, in a more realistic dynamic, link realizations might be revised asynchronously, in continuous time, and firms would stop operating at a random time when they can no longer go without the supplier (e.g., when inventory runs out). Then the dynamics would play out “smoothly,” characterized by differential equations rather than discrete iteration.

suppose that all search is futile and that firms necessarily choose $\hat{x}_{if} = 0$. Denote the probability firm if finds a supplier of a given input type by \hat{x}_{if} . Conditional on finding an input, we let it be successfully sourced with probability 1 so all frictions occur through the search process. Implicitly, obtaining a probability \hat{x}_{if} requires search effort, and we suppose that cost of achieving probability \hat{x}_{if} is $\hat{c}(\hat{x})$, where \hat{c} is a strictly increasing function with $\hat{c}(0) = 0$.

We suppose firms choose \hat{x}_{if} taking the environment as given. In particular, firms take as given the probability that suppliers of the inputs they require successfully produce. When many potential suppliers of an input produce successfully we let it be relatively easy to find one, and if none of these suppliers produce successfully then it is impossible to find one. In addition, the parameter n shifts how easy it is to find a supplier.

Given this set up we can let the probability of finding a supplier have the functional form $\hat{x} := 1 - (1 - x_{ifr})^n$, and the cost of achieving this probability be given by $\hat{c}(\hat{x}) := c\left(\frac{1 - (1 - \hat{x})^{1/n}}{r}\right)$. Although these functional form assumptions might seem restrictive, we still have freedom to use any function c satisfying our maintained assumptions. This degree of freedom is enough for the model to be quite general as all that matters is the size of the benefits of search effort relative to its cost, and not the absolute magnitudes. Further, these functional form assumptions satisfy all the desiderata we set out above. As $1 - (1 - x_{ifr})^n$ is the key probability throughout our analysis, all our results then go through with this interpretation.