NBER WORKING PAPER SERIES

ARE SUPPLY NETWORKS EFFICIENTLY RESILIENT?

Agostino Capponi
Chuan Du
Joseph E. Stiglitz

Working Paper 32221
http://www.nber.org/papers/w32221

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2024, Revised January 2025

Agostino Capponi: Research supported, in part, by the National Science Foundation under the NSF CAREER award 1752326 Joseph Stiglitz: Supported, in part, by INET. Chuan Du: Research undertaken whilst employed at the Bank of England, and the Board of Governors of the Federal Reserve System. The views in this paper are solely the authors' and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System, any other person associated with the Federal Reserve System, or the National Bureau of Economic Research. A significant part of this work was undertaken while Chuan Du was employed at the Bank of England. The views expressed in this paper do not necessarily reflect the views of the Bank of England or any of its committees. We thank our discussants and conference participants at the Allied Social Science Association Annual Meeting 2023 and the International Economic Association World Congress 2023, as well as seminar participants at the Bank of England and the Federal Reserve Board, for their constructive comments and suggestions.

Are Supply Networks Efficiently Resilient?
Agostino Capponi, Chuan Du, and Joseph E. Stiglitz
NBER Working Paper No. 32221
March 2024, Revised January 2025
JEL No. D21, D24, D25, D43, D85, E23, L13

ABSTRACT

We show that supply networks are inefficiently, and insufficiently, resilient. Upstream firms can adjust capacity investments to hedge against supply and demand shocks. However, the social benefits of such investments are not internalized, because of market incompleteness and market power. Upstream firms underinvest in resilience, passing on the costs to downstream firms, and drive trade excessively toward the spot markets. Policies designed to incentivize capacity investment, reduce reliance on spot markets, and enhance competition ameliorate the externality.

Agostino Capponi
Mudd Hall 535-G
Columbia University
500 W 120th St.
New York, NY 10027
ac3827@columbia.edu

Chuan Du
Board of Governors of the
Federal Reserve System
1801 K Street NW
Washington, D.C. 20006
chuan.du@aya.yale.edu

Joseph E. Stiglitz
5th Floor, Kravis Hall
Columbia University
665 West 130th St.
New York, NY 10027
and NBER
jes322@columbia.edu

The shortages and spikes in prices of certain intermediate goods during the COVID-19 pandemic demonstrated the fragility of supply chains. Prominent examples included a global shortage of semiconductors that led to a dramatic rise in the price of secondhand cars in the U.S. and an unprecedented demand for hand sanitizer and personal protective gear that triggered supply shortages in their respective, as well as interlinked, industries. Policymakers reacted strongly by taking industry-specific actions to repair linkages and improve resilience. For example, the Biden-Harris Administration worked in partnership with Congress to provide new legislation to alleviate specific supply chain disruptions and promote greater resilience in future situations.[1]

These experiences with supply network disruptions left open the question: Had firms invested too little in resilience ex ante? The pandemic was an extreme event, and, in general firms should not be expected to anticipate and plan for every possible contingency. Doing so would almost surely be inefficient, entailing excessive focus on resilience. We show here, however, that given market power and market incompleteness, one should expect markets to underinvest in resilience relative to a constrained efficient benchmark.

We formulate a tractable theoretical model whereby a collection of intermediate and final goods producers form supply linkages to meet uncertain consumer demand and accommodate supply shocks. Each final goods producer (the downstream firm) can source differentiated inputs from one or more suppliers of the intermediate goods (the upstream firms). Intermediate goods producers engage in price - that is, Bertrand - competition with differentiated products, taking the prices set by competitors as given. Lowering the price charged allows an intermediate goods producer to increase demand on the extensive margin (by attracting more final goods producers).[2] Intermediate goods producers face uncertainty in demand and supply conditions. They invest in non-scalable production capacity ("K") which cannot be increased in the second period (capturing the idea that some factors of production cannot be readily adjusted at short notice),[3] as well as a scalable factor ("L") that can be adjusted in response to shocks. Given the structural frictions in the economy - namely, the lags in production and the uncertainty around future market conditions - over-investment in capacity can be inefficient just as under-investment can be. A supply network

[1] While the large and small supply chain disruptions during COVID-19 had propelled the issue into popular discourse, the cracks had been evident before the pandemic. Hanjin Shipping, a world's top 10 container carrier, filed for bankruptcy in September 2016 because of sluggish freight rates caused by weak demand and soaring global capacity. The bankruptcy affected global supply chains, because half of Hanjin's container ships were denied access to ports. Major U.S. retailers, such as J.C. Penney and Walmart, began to divert and switch carriers for their containers to other suppliers. Similarly, the failure of Carillion in January 2018, once the second-largest construction company in the U.K., brought down many of its suppliers. A more academic account can be found in Baqaee and Farhi (2022), Guerrieri et al. (2022) and Di Giovanni et al. (2022).

[2] In a more general case, lowering the price may also affect the intensive margin.

[3] Semiconductors are an example of an important intermediate goods that requires significant capacity investment upfront. In the European Union, the European Chips Act (2023) aims to provide additional public and private investments of more than EUR 15 billion.

that is *efficiently resilient* strikes the optimal balance.

Using the model, we demonstrate the existence of a market failure in decentralized supply networks, whereby upstream firms do not fully internalize the social benefits of building production capacity. When upstream firms over-invest in capacity, part of the cost savings are passed on to downstream firms via lower prices; but when firms underinvest, they can defend their profit margins despite mounting costs by charging higher prices. There is a pecuniary externality from upstream firms' investment in capacity, and whenever markets are incomplete, as is obviously the case here, these pecuniary externalities matter (see Greenwald and Stiglitz (1986)).

Firms do not have access to the full set of Arrow-Debreu securities, and instead must trade either on the pre-order market, or on the spot market once the shocks have realized. The pre-order market offers partial insurance to both the upstream and downstream firms. For the upstream firms, pre-orders establish a minimum level of demand for their outputs, and help with their upfront non-scalable capacity investment decision. For the downstream final goods producers, a pre-order contract locks in an agreed price for the intermediary inputs in their production, shielding them from cost shocks in the upstream sector. If realized demand for final goods exceeds what can be fulfilled through pre-orders, the downstream firm can then source the extra inputs required from the spot market. As we observe in practice, the spot and pre-order markets are insufficient to deal with the full spectrum of possible shocks, and thus are unable to provide full insurance against supply network disruptions.[4] In the model, in the absence of complete markets, agents demonstrate an over-reliance on the spot market.

In our model, there is one further source of market failure: market power. Upstream intermediate good producers exhibit market power because: (a) there are only a finite number of such firms; and (b) the intermediate goods they produce are imperfect substitutes of each other. Studying resilience is especially relevant in this context, where the inherent decentralized market structure leads to monopolistic competition. However, this analysis departs from the Dixit-Stiglitz framework prevalent in the macroeconomic literature. Here, mark-ups are endogenous and vary across firms and economic states, providing a more dynamic examination of market behaviors.

Taken together, we show that the market-based network invests too little in production capacity (K^*) relative to a constrained optimal benchmark (K^{SP}) with a social planner facing the same informational and technological constraints as the private market. Even under the constrained benchmark, it is not optimal to build enough capacity to account for all contingencies, so there will be times when firms ex post have considerable market power, which, obviously, the social planner would not take advantage of but private firms would. In short, market-based supply networks are

[4] It is obvious that such full insurance *does not* exist. Given the range of shocks that could occur – some of which are now not even really conceivable – the incompleteness of insurance markets is inevitable. Theories of asymmetric information provide further explanations of the absence of a full set of insurance markets. See Greenwald and Stiglitz (1986) and Stiglitz (1982).

inefficiently resilient: $K^* < K^{SP}$.

Remarkably, this wedge between the decentralized and centralized solution arises even when rare large shocks are absent, and the economy operates in a "full production" equilibrium whereby supply capacity is sufficiently agile to accommodate all possible demands. Our results do not depend on an arbitrary specification of the distribution of shocks - for example, we do not require a threshold for the probability of large negative shocks. The wedge between the decentralized solution and the constrained optimal benchmark exists as long as firms cannot perfectly adjust their non-scalable production capacity in response to unanticipated shocks. We also do not need to impose a level of risk aversion on the part of private agents or the social planner. Capacity investment is suboptimally low, even when every agent – including the constrained social planner – is risk neutral.

Extending the analysis to account for rare disasters (in the online appendix), we show that the response of market-based supply networks to shocks can be highly nonlinear. Private supply networks are seemingly resilient during normal times and can comfortably withstand small to moderate shocks, but they are fragile to rare large shocks, when real rigidities prevent suppliers from fully meeting the needs of the market.[5] With a large enough shock, there is a transition from a monopolistically competitive regime to a local monopoly regime, whereby upstream firms are no longer pricing to compete and each downstream firm will receive only one credible offer for inputs. In other words, in a crisis, individual suppliers prioritize the needs of their local market but with increased margins.[6] Supply network fragility can lead to an increase in market power (in our model, reflected in suboptimal retrenchment in market coverage), especially when demand is at its greatest.

The size of the wedge between the decentralized and centralized solution depends endogenously on firms' reliance on the spot market, and exogenously on the structural parameters of the economy. An economy exhibiting greater scalability (production functions that rely less on non-scalable capacity investments), higher substitutability (intermediate goods inputs that are more interchangeable) and more competition (more upstream firms) will be more efficiently resilient.

Therefore, there are broadly three avenues for narrowing the wedge. First, a direct governmental subsidy targeting investment in production capacity could serve as the most pragmatic remedy. Second, enhancing incentives for the use of pre-order markets can offer upstream firms the assurance of recouping initial costs. We show that an overreliance on the spot market contributes

[5] By "seemingly resilient," we mean that demand can be fully met *at some price*. It is still the case that there is too little capacity.

[6] The surge in demand for COVID vaccines in 2021 and the frantic pursuit of natural gas during the European energy crisis in 2022 serve as illustrative examples. Global supply constraints often lead to redirection toward wealthier nations, leaving less-affluent developing markets economically disadvantaged during challenging times. During the post-COVID recovery, evidence suggests a marked increase in market power (markups) associated with the supply chain interruptions. See Konczal and Lusiani (2022).

to fragility in the supply network.[7] Third, the government can promote structural changes in the economy to enhance scalability, substitutability and competition. Enhancing competition is good in its own right, and doubly so when making supply networks more efficiently resilient.

The rest of the paper is organized as follows. Section 1 relates our paper to the existing literature. Section 2 sets up the model economy. Section 3 constructs the social planner benchmarks, and characterizes the constrained-optimal level of capacity investment (K^{SP}). Section 4 characterizes the decentralized equilibrium and the market solution for capacity (K^*). Section 5 presents our core result that firms' investment in capacity is insufficient - $K^* < K^{SP}$ - and discusses potential policy interventions. Section 6 concludes with suggestions for further research. We present formal derivations and proofs in the Appendix. In the separate Online Appendix, we discuss an extension of the analysis to rare large shocks pushing the economy away from full production.

1 Related literature

The literature on the resilience of supply networks to shocks can be roughly categorized into two branches. The first focuses on analyzing the mechanisms through which idiosyncratic shocks propagate and amplify within a fixed network of firms with pre-specified relationships. Acemoglu and Tahbaz-Salehi (Acemoglu and Tahbaz-Salehi) examine the effect of productivity shocks on the distribution of economic surplus, firm failures, and the amplification of shocks through disruptions. Acemoglu et al. (2012) propose a model that explains how micro shocks can be magnified into macro fluctuations through input-output linkages. Carvalho et al. (2021) use data from the 2011 Japanese earthquake to demonstrate the significant macroeconomic implications of idiosyncratic shocks. Barrot and Sauvagnat (2016) reveal evidence of fragility caused by the propagation of firm-specific shocks, using data on natural disasters. We refer to Carvalho (2014) and Carvalho and Tahbaz-Salehi (2019) for a thorough review of such mechanisms.

That markets would not be prepared for *every* shock they confront is not a surprise. The analytically interesting question is the normative one: Relative to an appropriate benchmark, do they adequately prepare for shocks? The failure of each firm in a competitive environment to take account of how capacity decisions affect the distribution of prices in the spot markets is one of the two central market failures that we identify.

The second branch of literature focuses on firms' strategic responses to mitigate the negative effects of supply chain disruptions. Birge et al. (2023) explore how firms in a supply chain network strategically react post-disruption by optimally switching demand and rerouting supply from

[7] For instance, in 2021 and 2022, more than 30 energy companies in the U.K. failed as a result of to a rapid increase in wholesale natural gas prices and inadequate hedging through futures/forward contracts by the energy companies. For details, see https://www.forbes.com/uk/advisor/energy/failed-uk-energy-suppliers-update.

defaulted firms. Amelkin and Vohra (2024) examine the competing retailers' decisionmaking process when selecting suppliers, taking into account factors such as prices and suppliers' reliability as measured by yield uncertainty and congestion.[8]

Our work is closely related to Elliot et al. (2022) and that by Grossman et al. (2023), which (also) examine supply network formation and fragility. In their models, downstream firms source customized inputs from upstream firms. To insure against possible supply disruptions, downstream firms strategically invest in relationships with multiple potential suppliers.[9] One might infer from their analyses that systemic fragility should be reduced if inputs were more (albeit still imperfectly) substitutable, and there existed a common spot market for such inputs. We show that not only would such a spot market be insufficient to eliminate supply network fragility, but that market participants' overreliance on spot market transactions would actually amplify the inherent externalities. In our model, fragility within the supply network is not a consequence of a catastrophic breakdown of upstream suppliers or a failure in supplier diversification but due to a more structural combination of market power and incomplete markets.

On the empirical side, Atalay et al. (2011) estimate a model of firms' buyer-supplier relationships using microdata on firms' customers. Crosignani et al. (2019) investigate the consequences of supply shocks resulting from NotPetya, one of the most severe cyberattacks in history. They observe that the affected downstream customers were more inclined to establish new relationships with alternative suppliers while terminating existing relationships with the directly affected firms. Lastly, Baldwin and Freeman (2022) examine the cross-border dimensions of resilience in global supply chains.

A separate branch of literature has studied the role of market power in amplifying demand or supply shocks. Noticeable contributions in this direction include Franzoni et al. (2024), who study the mechanisms by which supply chain shortages influence industry competition. Their empirical analysis demonstrates that "superstar" firms, when faced with supply chain shortages, incur comparatively smaller cost increases. This advantage enables them to expand their market share and enhance profitability. The study by Acharya et al. (2023) shows that elevated household inflation expectations enabled firms with high market power to transfer cost shocks directly to prices, with more pronounced effects observed among these firms. Unlike these studies, our findings reveal that market power held by intermediate good producers leads them to underinvest in production capacity, resulting in a supply network that is inefficiently resilient.

[8]A few other studies from the operations management literature analyze the mechanisms through which multi-sourcing strategies and supplier selection can help mitigate risk in supply chains. See Anupindi and Akella (1993), Tomlin (2006), Babich et al. (2012) and Babich et al. (2007).

[9]See also Elliott and Golub (2022) for a survey on supply chain disruptions and their macroeconomic implications.

2 Model

Consider an economy with two types of goods: final goods (the consumption numeraire) and intermediate goods used in the production of the final goods. There is a continuum of final goods producers (that is, downstream firms, indexed $i \in I = [0,1]$) and $n \geq 2$ intermediate goods producers (that is, upstream firms, indexed $j \in J = \{0,1,\ldots,n-1\}$), all located around a circle with unit circumference. The positions of the intermediate goods producers around the circle are represented by nodes, which divide the continuum of final goods producers into n "market segments." Figure 2.1 illustrates a simplified example of such an economy with $n = 3$ intermediate goods producers. Distance is quantified along the circle's circumference, ensuring that the maximum distance separating any two points is $\frac{1}{2}$.

Figure 2.1: Illustrative Economy

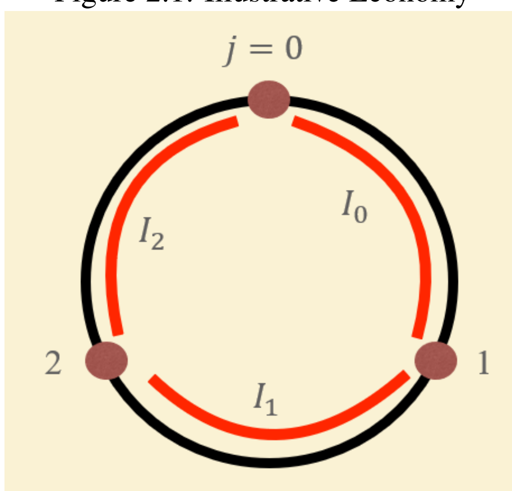

Consider an illustrative economy with three intermediate goods firms ($j \in \{0,1,2\}$). The intermediate goods firms are located equi-distant from each other, separating the circle into three equal market segments $\{I_0, I_1, I_2\}$. In a typical equilibrium, firms $j = 0$ and $j = 1$ compete over final goods firms located in the market segment I_0.

Intermediate goods producers $j \in J$ are price-setters. They set prices $\{p_j\}$ to compete over final goods producers in their two neighboring market segments.[10] The mass of final goods producers in each market segment is denoted as $\{m_k\}_{k=0,\ldots,n-1}$. To fulfill the endogenous demand for intermediate goods, each intermediate goods producer j operates a Cobb-Douglas production function with partial delay: $Y_{j,t} = L_{j,t}^{\alpha_j} K_{j,t-1}^{1-\alpha_j}$, where L_j denotes the scalable input factors in production with factor price $w_j > 0$, and K_j the non-scalable capacity investments that must be installed one period in advance at unit price $r_j > 0$. The key distinction is that non-scalable inputs K_j can-

[10] It is possible for any particular intermediate goods producer to price so aggressively as to capture demand from market segments further afield. This possibility corresponds to the "super-competitive" region of the demand curve in a circular economy (see Salop (1979)). For the purpose of the present analysis, our closed-form solutions focus on a symmetric equilibrium in which all intermediate goods producers find it optimal to set the same price, thus ruling out competition outside of the neighboring market segments.

not be adjusted in the short run.[11] The parameter $\alpha_j \in (0,1)$, the exponent of L, measures the *scalability* of each sector j. Crucially, intermediate goods producers j must decide on the level of non-scalable capacity investments K_j before the realization of shocks to the economy. As we will discuss in greater detail below, the intermediate goods producer's capacity investment (K_j), and pricing decisions (on both the spot and futures market) form the core of our model.

We model the final goods producers in a more reduced-form fashion. Specifically, final goods producers $i \in I$ are price-takers. Each atomistic final goods producer i faces an exogenous demand Q_i for its output, valued at unit price v.[12] These producers convert intermediate goods into the final goods using a linear production function $\tilde{Y}_i = \sum_j \frac{q_{ij}}{f(d(i,j))}$, where $\tilde{Y}_i$ denotes the final goods output of firm i, q_{ij} is the quantity of intermediate goods input firm i sources from firm j, and $f(d(i,j))$ is a penalty function that depends on the distance ($d(i,j) \in \left[0, \frac{1}{2}\right]$) between the two firms.

One way to think about this distance-based penalty function is that for every unit of intermediate goods j purchased by i, only a fraction $\frac{1}{f(d(i,j))}$ is usable. The remainder, $\left(1 - \frac{1}{f(d(i,j))}\right)$, "perishes in transit." A second interpretation of $f(\cdot)$ is a valuation-based penalty function. For any given valuation v, the effective valuation of the final goods i that uses inputs j is given by $\frac{v}{f(d(i,j))}$. Therefore, the function $f(\cdot)$ can also account for heterogeneous valuations of final goods. Specifically, a final goods firm i producing outputs using more "distant" intermediate goods would experience a diminished valuation for its output. A third interpretation (and the one we focus upon in the discussion below) is that the different intermediate goods are imperfect substitutes for each other. The production at any place in the circle is designed for a certain type of intermediate goods but can use other intermediate goods, though they yield less output per unit of input. (Think of an oil refinery designed to refine oil of a specific gravity and sulfur content. It can refine oil with other characteristics, but less efficiently). For ease of exposition, we will refer to $f(\cdot)$ henceforth as the *distance-based penalty function* (distance, in this interpretation, refers to distance in the product space).[13]

We assume that $f(\cdot)$ is an increasing function, normalized such that $f(0) = 1$. This penalty function, $f(\cdot)$, combined with the starting distance between firms, $d(i,j)$, captures the extent of *substitutability* among intermediate goods. The greater the distance $d(i,j)$ between two firms, and

[11] For brevity, we will henceforth drop the time subscripts, and note simply that K must be precommitted in advance of production.

[12] In our model, final goods firms form expectations over the level of demand Q_i, taking the price v as a fixed constant, whereas, more generally, shocks to final goods demand would affect both (their desired) equilibrium quantity Q_i and price v_i. We simplify the analysis by taking the integral over the distribution of Q_i only, instead of the joint distribution over both Q_i and v_i. This simplification offers greater analytical tractability and highlights the critical market failures, while preserving the essential economics of resilience. One can think of this modeling approach either as: (1) a stylized portrayal of final goods demand - a demand curve with demand equal to Q for price equal or less than v, and zero demand for price above v, or (2) a description of specific markets (like that for electricity) in which all firms have signed contracts to deliver output at price v regardless of the level of demand that materializes.

[13] For a discussion of the measurement of distance in product space, see, for example, Stiglitz (1986).

the steeper the slope $f'(x)$ of the penalty function, the more inefficient it becomes for final goods producer i to source inputs from intermediate goods producer j. For brevity, let $f_{ij} := f(d(i,j))$ and $\mathbf{f}_i := (f_{i0}, \ldots, f_{i,n-1})'$ be the corresponding $n \times 1$ column vector of penalties for final goods producer i.

Figure 2.2: Model Timeline

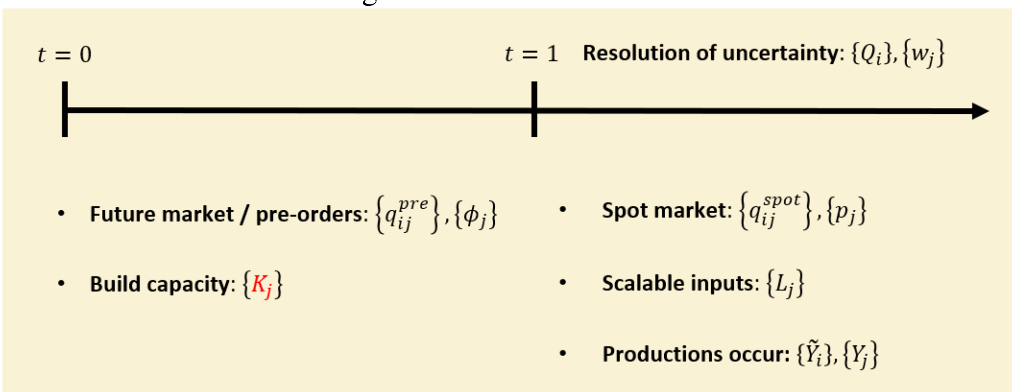

Timeline of events, decisions and actions undertaken by intermediate and final goods producers.

Figure 2.2 summarizes the timeline of the model. At period 0, there is uncertainty around the demand and supply conditions that will prevail in period 1. Specifically, the uncertainty around the demand for final goods produced by firm i is captured by the random variable Q_i. Q_i is distributed between $\left[\underline{Q}_i, \bar{Q}_i\right]$, with cumulative density function (c.d.f.) $G_i(\cdot)$ and associated probability density function (p.d.f.) $g_i(\cdot)$. There is also uncertainty around $\{w_j\}_{j \in J}$, the price of the scalable input factor, which affects the supply of the intermediate goods j. w_j is distributed between $\left[\underline{w}_j, \bar{w}_j\right]$, with c.d.f. $H_j(\cdot)$ and p.d.f. $h_j(\cdot)$.[14] Supply shocks are assumed to be independent of demand shocks. In our formulation, there is no uncertainty about the price of the final goods, v.

In period 0, to hedge against these demand and supply shocks, each final goods producer i decides whether to enter into a supplier contract with each intermediate goods producer j, placing pre-orders $\mathbf{q}_i^{pre} := \left[q_{i0}^{pre}, \ldots, q_{ij}^{pre}, \ldots, q_{in-1}^{pre}\right]'$. Each intermediate goods producer j sets pre-order price ϕ_j. Concurrently, firm j make a cost-minimizing decision on the level of non-scalable capacity K_j, incurring associated costs denoted by $r_j K_j$. The pre-order contracts between final goods and intermediate goods producers define the *endogenous* network formed in period 0.

In period 1, firms observe the realization of the demand and supply shocks. Final goods producer i submits spot-market orders $\mathbf{q}_i^{spot} := \left[q_{i0}^{spot}, \ldots, q_{ij}^{spot}, \ldots, q_{in-1}^{spot}\right]'$. The total cost of pre-orders and spot-market orders for firm i is given by $\left[\phi \cdot \mathbf{q}_i^{pre} + \mathbf{p} \cdot \mathbf{q}_i^{spot}\right]$, where $\phi := \left[\phi_0, \ldots, \phi_j, \ldots \phi_{n-1}\right]'$ denote the vector of pre-order prices, and $\mathbf{p}$ the vector of spot-market prices. At period 1, intermediate goods producer j takes pre-committed capacity K_j as given, solves for the cost-minimizing

[14] Without loss of generality, let $\infty > \bar{Q}_i > \underline{Q}_i > 0, \forall i \in [0,1]$; and $\infty > \bar{w}_j > \underline{w}_j > 0, \forall j \in J$.

scalable input L_j, and sets prices p_j to maximize profits. Production occurs, and contracts are settled. The excess of production over the contracted pre-orders is sold on the spot market.

In our model, the final goods producers can buy from any intermediate goods producer at the posted price. This flexibility stands in contrast to much of the network literature discussed in Section 1 (for example, Elliot et al. (2022)), where final goods producers can only buy from the firms with whom they have previous relations, so shocks to those firms obviously get passed on strongly through the network. Here, in effect, the ex ante and ex post networks can be different. [15] This is a crucial distinction. Even when there are established networks in "normal" situations, firms can turn to markets when there are large shocks. The real world is one in which there are both networks and markets.

For analytical tractability, we impose symmetry on the primitives of the model and derive closed-form solutions for the resulting symmetric equilibrium.

Assumption *A*1 [**Symmetry**]: $\alpha_j = \alpha$ and $r_j = r$, $\forall j \in J$; $Q_i = Q$, $\forall i \in I$; $w_j = w$, $\forall j \in J$; $m_k = \frac{1}{n}$, $\forall k \in \{0, \ldots, n-1\}$

By assumption, all intermediate goods producers share a common Cobb-Douglas production function: $\alpha_j = \alpha, \forall j \in J$; and face the same non-scalable input costs in period 0: $r_j = r$, $\forall j \in J$. We also assume that the shocks to the economy are symmetric and identical. The realization of final goods demand is the same for all final goods firms: $Q_i = Q$, $\forall i \in I$; and the realization of scalable input cost is also the same for all intermediate goods firms: $w_j = w$, $\forall j \in J$.[16] This symmetry captures an economy that is subject to systemic, correlated shocks. For instance, a symmetric demand shock might resemble the surge in demand for vaccines amid a pandemic, whereas a symmetric supply shock could be akin to a military conflict causing a spike in energy prices that affects all manufacturing sectors. Lastly, $m_k = \frac{1}{n}$, $\forall k$ implies that the sizes of each market segment are equal. The intermediate goods producers are uniformly distributed around the unit circle at equidistant intervals.

It is important to note that fully symmetric shocks to final goods demand (Q) and intermediate goods supply (w) do not immediately imply fully symmetric equilibrium outcomes. For instance, final goods producers that are further away from intermediate goods supplier nodes (that is, those with less substitutable inputs) will need to order more of a given input - compared with another final goods firm that is closer - to meet the same level of final goods demand. In practice, perfectly

[15] We assume that there are no costs to establishing a new link ex post. Our result may be generalized by assuming either that there is a fixed cost to going to the market or to buying from any specific firm with whom one does not have a previous relation. The problem would become analytically more challenging, but the main insights would stay qualitatively the same

[16] This is a slight abuse of notation. We use Q_i and w_j to represent both the random variable and its realized value. The intended meaning should be clear within the given context.

correlated shocks are the most challenging for resilience, which makes them a key "test case" to examine.

Before we dive into the formal equations that define the decentralized equilibrium, it is useful to first explore the more straightforward problems of an unconstrained and constrained social planner. The planner solutions will serve as our benchmarks for comparison.

3 The social planner benchmarks

We characterize the symmetric equilibrium outcomes for two separate benchmarks. In the first, the social planner can perfectly observe the realization of the state variables (Q,w) before committing to intermediate goods production across the network. The planner can therefore perfectly adjust both input factors (L,K) in line with market conditions. We call this unconstrained planner's solution the *first-best perfect foresight benchmark*. We re-introduce the informational and technological constraints faced by private agents in the second - *constrained optimal* - social planner's benchmark. Of the two, the constrained optimal benchmark provides a more appropriate basis for comparison. However, the perfect foresight benchmark serves a valuable role in isolating the effects of real-world frictions - such as uncertainties around states and limitations in production technology - from those associated with the distortions that arise as a result of market externalities and other imperfections.

There are two key distinctions between the social planner (under both benchmarks) and the decentralized market. First, a social planner can directly allocate order flows $\left\{q_{ij}\right\}$ without the need to use price signals $(\mathbf{p},\phi)$ as a coordinating mechanism. Second, a social planner maximizes the welfare of the economy as a whole, whereas individual private agents maximize their own profit or utilities. Thus, the social planner internalizes any externalities that may arise.

We restrict attention to a full production equilibrium, where the total demand for final goods can be met in a socially profitable way - that is where, at the margin, the value of the final goods exceeds the marginal cost of production. This setting further underscores that our core findings are not contingent on the occurrence of rare, large-scale shocks. Formally, a symmetric economy $\mathscr{E}=\{f(\cdot),\alpha,w,r,Q;v\}$ admits a *full production equilibrium* if there exists an equilibrium whereby $\tilde{Y}_i(Q,w)=Q$, $\forall i\in[0,1]$, and, for all states of the world $(Q,w)\subset\mathbb{R}_+^2$.

Assumption $A2$ **[Full production]**: We provide conditions on the model primitives that ensure the attainment of a full production equilibrium in a symmetric economy. More specifically, we assume that at every point around the circle (that is, $\forall i\in I=[0,1]$), the marginal benefits of producing final goods will at least match or exceed the marginal costs in all possible scenarios:

$$\frac{v}{f\left(\frac{1}{2n}\right)} \geq \left(\frac{\bar{w}}{\alpha}\right)^{\alpha} \left(\frac{r}{1-\alpha}\right)^{1-\alpha} \left(\frac{\bar{w}\bar{Q}^{\frac{1}{\alpha}}}{E\left[wQ^{\frac{1}{\alpha}}\right]}\right)^{1-\alpha} \qquad \text{for } n \geq 2.$$

Assumption A2 states that the marginal benefit of delivering intermediate goods to the final goods producer located farthest from the nearest node (at a distance of $\frac{1}{2n}$) is weakly greater than the marginal cost of producing the intermediate goods in equilibrium $\left(\frac{\bar{w}}{\alpha}\right)^{\alpha} \left(\frac{r}{1-\alpha}\right)^{1-\alpha} \left(\frac{\bar{w}\bar{Q}^{\frac{1}{\alpha}}}{E\left[wQ^{\frac{1}{\alpha}}\right]}\right)^{1-\alpha}$, even when the negative supply shock is at its most extreme ($w = \bar{w}$), and demand is at its upper bound ($Q = \bar{Q}$).[17] For any given v, this assumption is equivalent to a restriction on the range of the demand and supply shocks. The assumption guarantees full production under the constrained optimal benchmark, where the social planner faces the same informational and technological constraints as the decentralized market.[18] The corresponding condition for the perfect foresight benchmark is $\frac{v}{f\left(\frac{1}{2n}\right)} \geq \left(\frac{\bar{w}}{\alpha}\right)^{\alpha} \left(\frac{r}{1-\alpha}\right)^{1-\alpha}$ for $n \geq 2$, where the marginal cost of production is lower because the social planner can fully adjust both inputs of production (K as well as L) in response to shocks (that is, $\bar{w}\bar{Q}^{\frac{1}{\alpha}} > E\left[wQ^{\frac{1}{\alpha}}\right]$ by construction). Assumption A2 is therefore a sufficient condition for full production under both social planner benchmarks. See appendix C for a more detailed discussion, and for sufficient conditions for a full production equilibrium in the decentralized solution.[19]

Relaxing Assumption A2 leads to cases where some segment of the economy (farthest away from the intermediate goods producers) might become shut out from the final goods market under adverse supply conditions. In such instances, intermediate goods suppliers operate as localized monopolies rather than as direct competitors, each prioritizing the needs of their local markets (at higher margins) and leaving demand from more "distant" firms unfulfilled. The emergence of local monopolies introduces an extra layer of distortion to the decentralized market solution, which further strengthens our core argument that there is insufficient investment in non-scalable production capacity. We discuss the consequences of relaxing this assumption in greater detail in online appendix G.

[17] In general, the marginal cost of production for intermediate goods should be a function of K. The expression we use here represents the marginal cost in equilibrium, i.e., when K is chosen optimally.

[18] See appendix B.2 for details.

[19] On a technical note, the full production assumption also helps us circumvent issues related to non-differentiability in the demand function.

3.1 The perfect foresight (PF) benchmark

Consider the first-best problem for a social planner with a fully scalable production function and perfect foresight. The social planner operates a standard Cobb-Douglas production function for intermediate goods: $Y_{j,t} = L_{j,t}^{\alpha} K_{j,t}^{1-\alpha}$.[20] The planner can also dictate input choices $\left\{K_j, L_j\right\}_{j\in J}$ and order flows $\left\{q_{ij}\right\}_{i\in I, j\in J}$ for all firms after observing the realization of final goods demand Q and scalable input cost w. Although production is delayed until period 1, there is no uncertainty. At period 0, firms know the realization of the shocks that arrive at period 1. Mathematically, this is equivalent to all decisions being made in a single period optimization problem, where the objective is to maximize the value of production net of its costs.

[**Optimization Problem PF**]:

$$W(Q,w) = \max_{\left\{K_j\right\}_{j\in J}, \left\{L_j\right\}_{j\in J}, \left\{q_{ij}\right\}_{i\in I, j\in J}} \left\{ v \int_0^1 \left[\min\left\{Q, \tilde{Y}_i\right\}\right] di - \sum_{j\in J} \left[rK_j + wL_j\right] \right\}, \tag{3.1}$$

$$\text{s.t. } \tilde{Y}_i = \sum_{j\in J} \frac{1}{f_{ij}} q_{ij} \quad \text{[Production function for final good } i\text{]},$$

$$Y_j = L_j^{\alpha} K_j^{1-\alpha} \quad \forall j \in J \quad \text{[Production function for intermediate good } j\text{]},$$

$$\int_0^1 q_{ij} di \leq Y_j \quad \forall j \in J \quad \text{[Feasibility of intermediate goods order flow]},$$

$$q_{ij} \geq 0 \quad \forall i \in [0,1], \forall j \in J \quad \text{[Nonnegative inputs]}.$$

The solution is simple and intuitive. In the perfect foresight benchmark, the planner would meet final goods demand by sourcing intermediate goods inputs from the cheapest supplier and produce the required intermediate goods at minimal cost by optimizing the ratio between scalable and non-scalable inputs in every state.

Lemma 1. ***[Full production symmetric equilibrium under perfect foresight]***

1. *The social planner allocates sufficient intermediate goods j to each final goods firm i to meet consumer demand Q, accounting for any imperfect substitutability f_{ij}. The required intermediate goods inputs will be sourced from the lowest effective-cost supplier(s) for each i, whenever the value of production v exceeds the marginal cost of production:*

$$q_{ij}^{PF}(Q,w) = \begin{cases} f_{ij}Q & \text{if } j \in \underline{J}(i) \text{ and } v \geq f_{ij}\left(\frac{w}{\alpha}\right)^{\alpha}\left(\frac{r}{1-\alpha}\right)^{1-\alpha} \\ 0 & \text{otherwise} \end{cases}, \tag{3.2}$$

where $\underline{J}(i) := \left\{\tilde{j} \in J | f_{i\tilde{j}}\left(\frac{w}{\alpha}\right)^{\alpha}\left(\frac{r}{1-\alpha}\right)^{1-\alpha} \leq f_{ij}\left(\frac{w}{\alpha}\right)^{\alpha}\left(\frac{r}{1-\alpha}\right)^{1-\alpha} \quad \forall j \in J\right\}$ is the set of low-

[20] We suppress the t subscript henceforth to simplify the notation.

est effective-cost supplier(s).

2. *The planner's input choices in intermediate goods production satisfy the optimality condition:*

$$\alpha r K^{PF}(Q,w) = (1-\alpha) w L^{PF}(Q,w). \tag{3.3}$$

From the optimality condition above we can derive the explicit solutions for K^{PF} and L^{PF} :

$$K^{PF}(Q,w) = \left(\frac{w}{r}\frac{(1-\alpha)}{\alpha}\right)^{\alpha}\left(2Q\int_0^{\frac{1}{2n}} f(i)\,di\right),$$

$$L^{PF}(Q,w) = \left(\frac{r}{w}\frac{\alpha}{1-\alpha}\right)^{1-\alpha}\left(2Q\int_0^{\frac{1}{2n}} f(i)\,di\right),$$

where $f(i) = f_{i0} := f(d(i,0))$ is the shorthand for the distance penalty between final goods firm i and intermediate goods firm 0, and, by symmetry, $K_j^{PF} = K^{PF}$ and $L_j^{PF} = L^{PF}$ for all $j \in J$.

3.2 The constrained optimal social planner (SP) benchmark

Next, we consider the constrained optimal problem, whereby a social planner can dictate production choices $\{L_j, K_j\}$ and order flow $\{q_{ij}\}_{i\in I, j\in J}$ but faces the same informational and technological constraints as the private sector. We solve the constrained optimal problem through backward induction.

In period 1, the social planner takes the pre-committed non-scalable capacity $K_j = K,\ \forall j \in J$ as given and chooses the scalable input factor $\{L_j\}_{j\in J}$ and order flows $\{q_{ij}\}_{i\in I, j\in J}$ to maximize aggregate welfare for any given realization of demand and supply conditions (Q,w). For given intermediate goods output $Y_j = \int_0^1 q_{ij}di$, we can express the cost-minimizing level of the scalable factor as $L_j = \left(\int_0^1 q_{ij}di\right)^{\frac{1}{\alpha_j}} K_j^{-\frac{(1-\alpha_j)}{\alpha_j}}$. Substituting out L_j and imposing symmetry (Assumption A1), we can express the optimization problem [SP1] in terms of the order flows $\{q_{ij}\}_{i\in I, j\in J}$ only:

[**Optimization problem SP1**]:

$$W^{SP}\left(K|Q,w\right)=\max_{\left\{q_{ij}\right\}_{i\in I,j\in J}}\left\{v\int_0^1\left(\sum_{j\in J}\frac{1}{f_{ij}}q_{ij}\right)di-\sum_{j\in J}\left[rK+w\left(\left(\int_0^1 q_{ij}di\right)^{\frac{1}{\alpha}}K^{-\frac{(1-\alpha)}{\alpha}}\right)\right]\right\} \tag{3.4}$$

$$\text{s.t. } Q\geq\sum_{j\in J}\frac{1}{f_{ij}}q_{ij}\quad\forall i\in[0,1],\quad\text{[Demand cap]}$$

$$q_{ij}\geq 0\quad\forall i\in[0,1],j\in J,\quad\text{[Nonnegative inputs]}$$

where $v\int_0^1\left(\sum_{j\in J}\frac{1}{f_{ij}}q_{ij}\right)di$ is the aggregate value derived from the production of final goods and $\sum_{j\in J}\left[rK+w\left(\left(\int_0^1 q_{ij}di\right)^{\frac{1}{\alpha}}K^{-\frac{(1-\alpha)}{\alpha}}\right)\right]$ the aggregate cost of producing the necessary intermediate inputs. The demand cap reflects that any production in excess of the realized demand Q will be wasted.[21]

Back in period 0, the social planner chooses non-scalable inputs $\left\{K_j\right\}_{j\in J}$ to maximize expected welfare in period 1, accounting for the probability distribution of demand and supply shocks (Q,w).

[**Optimization problem SP0**]:

$$W^{SP}=\max_K E\left[W^{SP}\left(K|Q,w\right)\right] \tag{3.5}$$

The solution resembles that of the perfect foresight scenario, but with important distinctions, arising from the necessity of committing to a specific level of capacity investment in period 0, before the realization of states in period 1.

Lemma 2. ***[Full production symmetric equilibrium in the constrained optimal benchmark]***

1. *In period 1, the social planner allocates sufficient intermediate goods to each final goods firm i to meet consumer demand Q, accounting for imperfect substitutability. The required intermediate goods inputs will be sourced from the lowest effective-cost supplier(s) for each i, whenever the value of production v exceeds the marginal cost of production:*

$$q_{ij}^{SP}\left(Q,w\right)=\begin{cases}f_{ij}Q & \text{if } j\in\underline{J}(i) \text{ and } v\geq f_{ij}\widetilde{MC}\\ 0 & \text{otherwise}\end{cases}, \tag{3.6}$$

where $\widetilde{MC}:=\left(\frac{w}{\alpha}\right)^{\alpha}\left(\frac{r}{1-\alpha}\right)^{1-\alpha}\left(\frac{wQ^{\frac{1}{\alpha}}}{E\left[wQ^{\frac{1}{\alpha}}\right]}\right)^{1-\alpha}$ *is the marginal cost of producing the interme-*

[21] In this analysis, we deliberately exclude the effect of inventory management because of the framework's static, one-shot nature. See Ferrari (2022) for a network model with inventories.

diate goods in the symmetric equilibrium and $\underline{J}(i) := \left\{ \tilde{j} \in J | f_{i\tilde{j}} \widetilde{MC} \leq f_{ij} \widetilde{MC} \quad \forall j \in J \right\}$ *is the set of lowest effective-cost supplier(s).*

The optimal level of scalable input is given by

$$L^{SP}(Q,w) = \left(\int_0^1 q_{ij}^{SP}(Q,w)\, di \right)^{\frac{1}{\alpha}} \left(K^{SP} \right)^{-\frac{(1-\alpha)}{\alpha}}. \tag{3.7}$$

2. *In period 0, the optimal level of non-scalable production capacity* K^{SP} *satisfies the optimality condition:*

$$\alpha r K^{SP} = (1-\alpha) E\left[w L^{SP} \right], \tag{3.8}$$

which can be solved explicitly to give

$$K^{SP} = \left(\frac{1}{r} \frac{1-\alpha}{\alpha} \right)^{\alpha} \left(2 \int_0^{\frac{1}{2n}} f(i)\, di \right) \left(E\left[w Q^{\frac{1}{\alpha}} \right] \right)^{\alpha}.$$

3. *The relationship between capacity investment across the two benchmark scenarios can be summarized as follows:*

$$K^{SP} = K^{PF}(Q,w) \left(\frac{E\left[w Q^{\frac{1}{\alpha}} \right]}{w Q^{\frac{1}{\alpha}}} \right)^{\alpha}, \tag{3.9}$$

and by Jensen's inequality we have

$$K^{SP} \geq E\left[K^{PF}(Q,w) \right].$$

The first part of the proposition relating to the optimal order flow (q_{ij}^{SP}) and the level of scalable capacity (L^{SP}) is straight-forward. Here, we will concentrate our discussion on the underlying intuition of the constrained optimal solution for capacity investment K^{SP}. The choice of non-scalable capacity at period 0, K^{SP}, influences aggregate welfare in period 1 through two primary mechanisms. First, any increase in K^{SP} generates a *direct cost* given by r. This cost, however, is partly offset by the resultant decrease in the scalable input L^{SP} needed to achieve a given output Y, thus offering a *direct benefit*. Second, a rise in capacity K^{SP} may increase aggregate intermediate goods production Y, and indirectly improve welfare through this output channel $\frac{dY}{dK}$. However, in a full-production equilibrium where the demand for final goods is always met (that is, the demand cap is binding), there can be no further welfare gains from increasing aggregate intermediate goods production. Therefore, the indirect effect of K on welfare is exactly zero.[22] We are left with the familiar optimality condition that is typical for Cobb-Douglas production functions, $\alpha r K^{SP} = (1-\alpha) E\left[wL^{SP}\right]$, albeit with an expectation function to account for the ex ante uncertainty.

Finally, equation 3.9 illustrates the relationship between the level of capacity investment across the two benchmarks. Under the constrained optimal benchmark, the social planner must commit to a given level of capacity K^{SP} before observing the shocks. Hence, capacity investment is lower than that in the perfect foresight case, $K^{SP} < K^{PF}(Q,w)$, in states where marginal costs exceed expectations (when w, Q, or both are higher than expected). Conversely, $K^{SP} > K^{PF}(Q,w)$ when marginal costs fall below expectations. Importantly, this result implies that the constrained social planner recognizes that investing in a level of production capacity that accommodates all contingencies ($K^{PF}\left(\bar{Q},\bar{w}\right)$) would give rise to a supply network that is inefficiently resilient. Nevertheless, the constrained social planner invests in more capacity than its counterpart with perfect foresight does on average, $K^{SP} \geq E\left[K^{PF}(Q,w)\right]$, as a way to insure against uncertainty.[23]

4 The decentralized solution: Equilibrium in the spot and pre-order markets

In the decentralized market equilibrium, firms adjust production in response to prices in both the pre-order and spot markets. We solve the model through backward induction.

4.1 Period 1 equilibrium in the spot market

In period 1, each final goods producer can turn to the spot market to acquire additional intermediate goods beyond those that have been pre-ordered. Formally, each final goods producer i takes realized demand for final goods Q_i, prior commitments $\mathbf{q}_i^{pre}$, pre-order and spot intermediate goods prices $(\phi, \mathbf{p})$ as given[24], and purchases intermediate goods $\mathbf{q}_i^{spot}$ from intermediate goods

[22] In online appendix F, we show that the optimality condition for non-scalable production capacity K^{SP} (equation 3.8) remains unchanged when we relax the full production assumption. We can safely ignore the indirect effects of K on welfare through changes in output Y. This result bears resemblance to the Envelope Theorem, in which the total derivative of the value function with respect to the parameters of the model is equal to its partial derivative. Here K is the choice variable, but the total derivative of $W^{SP}(K|Q,w)$ with respect to K is also equal to its partial derivative.

[23] A formal exposition of this result can be found in appendix B.3. This result may be contingent on our assumption of a Cobb-Douglas production function. In a somewhat simpler context of a utility having to meet a fixed demand, Rothschild and Stiglitz show that whether uncertainty increases or decreases capacity depends on the elasticity of substitution between the scalable and nonscalable factors of production. See Rothschild and Stiglitz (1971).

[24] As is conventional in the literature on Bertrand equilibria, each firm assumes it can buy as much on the spot market as it wishes. This assumption is particularly important for the analysis of firms' decision making at time 0.

producers on the spot market in order to maximize profit:

$$\Pi_i\left(\mathbf{q}_i^{pre},\phi,\mathbf{p}\right)=\max_{\mathbf{q}_i^{spot}}\left\{v\min\left\{Q_i,\tilde{Y}_i\right\}-\tilde{C}_i\left(\mathbf{q}_i^{pre},\mathbf{q}_i,\phi,\mathbf{p}\right)\right\} \quad (4.1)$$
$$\text{s.t.}\quad \mathbf{q}_i^{spot}\geq 0\quad \text{[No-default constraint]},$$

where final goods production and total costs are given by:

$$\tilde{Y}_i=\sum_{j\in J}\frac{1}{f_{ij}}\left(q_{ij}^{spot}+q_{ij}^{pre}\right)$$
$$\tilde{C}_i\left(\mathbf{q}_i^{pre},\mathbf{q}_i,\phi,\mathbf{p}\right)=\phi\cdot\mathbf{q}_i^{pre}+\mathbf{p}\cdot\mathbf{q}_i^{spot}.$$

We interpret $\mathbf{q}_i^{spot}\geq 0$ as a "no-default constraint" because it implies that the total volume of intermediate goods orders will never fall below the pre-ordered amount: $\mathbf{q}_i:=\mathbf{q}_i^{spot}+\mathbf{q}_i^{pre}\geq\mathbf{q}_i^{pre}$. The final goods producers cannot renege on the promises made in period 0. In principle, a firm could also resell its pre-order to some other firm, so that the level of input could be less than the pre-ordered level. In a symmetric equilibrium, however, that never occurs.[25]

If the spot market were perfectly competitive, each intermediate supplier would produce up to the point where the price of the intermediate goods (on the spot market) were equal to the *marginal* cost of production, and the demand for intermediate goods would be determined in the usual way, with equilibrium in the spot market occurring at the price where demand equals supply. Instead, this is a highly differentiated market for intermediate goods, and each intermediate goods producer acts in a monopolistically competitive way, setting a spot price p_j and taking its non-scalable production capacity K_j, and the price of its competitors $\mathbf{p}_{-j}$ as given. Pre-order contracts $\left\{q_{ij}^{pre}\right\}_{i\in I}$ are honored at the agreed price ϕ_j. The profit of firm j is given by its pre-order revenue plus spot-market revenue, minus the total costs of production:

$$\Pi_j\left(K_j,\left(\phi_j,q_j^{pre}\right),\mathbf{p}_{-j}\right)=\max_{p_j}\left\{\left[\phi_jY_j^{pre}\right]+\left[p_jY_j^{spot}\right]-\left[w_jL_j^*+r_jK_j\right]\right\}, \quad (4.2)$$

where $Y_j^{pre}:=\int_0^1 q_{ij}^{pre}di$ and $Y_j^{spot}:=\int_0^1 q_{ij}^{spot}di$ are the level of intermediate goods production required to meet pre-order demand and spot market-demand, respectively.

Similar to our treatment of the social planner benchmarks, we restrict attention to a full production symmetric equilibrium for analytical tractability. In a symmetric setting, all intermediate

[25]Conceptually, we could imagine an equilibrium where, say, in some states, those in one set of locations sold excess pre-orders to those in another set of locations. Our assumption of perfectly correlated shocks rules out this scenario. Alternatively, even with imperfectly correlated shocks, reselling excess orders can be assumed away, for example, because there are some (not fully specified here) adaptations of production to each producer, which make such sales impossible. In practice resale of pre-ordered inputs do occur, though they are likely limited in scale.

goods firms $j \in J$ share the same characteristics $\alpha_j = \alpha$, $r_j = r$ and $w_j = w$; and every final goods firm $i \in [0,1]$ will face the same exogenous demand $Q_i = Q$. In equilibrium, input choices will be the same across intermediate goods firms: $K_j = K$ and $L_j = L$, $\forall j \in J$; and final goods firms will fulfill the same proportion of the realized demand for final goods through the pre-order market: $Q_i^{pre} := \sum_j \frac{1}{f_{ij}} q_{ij}^{pre} = Q^{pre}\ \forall i \in [0,1]$.

It is important to note that Q_i^{pre} denotes the level of *final goods* demand fulfilled through pre-orders and not the quantity of *intermediate goods* pre-ordered $\mathbf{q}_i^{pre}$. The link between the two is given by $Q_i^{pre} := \sum_j \frac{1}{f_{ij}} q_{ij}^{pre}$, where $\frac{1}{f_{ij}}$ accounts for imperfect substitutability. Later, in Section 4.2, we show that $Q_i^{pre} = Q^{pre}\ \forall i \in [0,1]$ is indeed an optimal equilibrium strategy in period 0, but this strategy implies pre-orders for intermediate goods q_{ij}^{pre} are not equalized across i's.

Lemma 3. ***[Full Production Symmetric Equilibrium in the spot market]*** *In period 1, taking period 0 choices* $\left(\left\{\mathbf{q}_i^{pre,*}\right\}, K^*, \phi^*\right)$ *as given:*

1. *Final goods firms order intermediate goods on the spot market from the supplier offering the lowest effective-prices* $j \in \underline{J}(i;\mathbf{p}) := \left\{\tilde{j} \in J : f_{i\tilde{j}} p_{\tilde{j}} = \min\{\mathbf{f}_i \circ \mathbf{p}\}\right\}$*:*

$$q_{ij}^{spot,*} = \begin{cases} f_{ij}\left(Q - Q_i^{pre}\right) & \textit{if } Q \geq Q_i^{pre}, j \in \underline{J}(i;\mathbf{p}), \textit{ and } v \geq f_{ij} p_j \\ 0 & \textit{otherwise} \end{cases} \quad \forall i \in [0,1]. \tag{4.3}$$

2. *Intermediate goods firms:*

   - *purchase the cost-minimizing level of scalable inputs:*

$$L_j^* = L^* = \left(Y^{pre,*} + Y^{spot,*}\right)^{\frac{1}{\alpha}} (K^*)^{-\frac{1-\alpha}{\alpha}} \quad \forall j \in J, \tag{4.4}$$

   - *set spot-market prices at a markup over marginal costs:*

$$p_j^* = p^* = \underbrace{(1+\mu)}_{\textit{mark-up} \geq 1} MC \quad \forall j \in J, \tag{4.5}$$

   where

     - $Y^{pre,*} := \int_0^1 q_{ij}^{pre,*} di$ *and* $Y^{spot,*} := \int_0^1 q_{ij}^{spot,*} di$ *are the level of intermediate goods production required to meet equilibrium pre-order demand and spot market demand, respectively;*

– *μ is the proportional mark-up over marginal costs given by:*

$$\mu := \frac{2f'\left(\frac{1}{2n}\right)\int_0^{\frac{1}{2n}} f(i)\,di}{\left(f\left(\frac{1}{2n}\right)\right)^2 - 2f'\left(\frac{1}{2n}\right)\int_0^{\frac{1}{2n}} f(i)\,di}, \tag{4.6}$$

– *MC is the marginal cost faced by intermediate goods suppliers:*

$$MC = \frac{w}{\alpha}\left(\frac{Y^{pre,*} + Y^{spot,*}}{K^*}\right)^{\frac{1-\alpha}{\alpha}}. \tag{4.7}$$

In the period 1 equilibrium, each final goods producer first evaluates whether its pre-committed orders for intermediate goods will be adequate to satisfy the existing demand for final goods - that is, whether $Q_i^{pre} := \sum_j \frac{1}{f_{ij}} q_{ij}^{pre} \geq Q_i$. Should the pre-orders prove sufficient, the final goods producer i will eschew the spot market, setting $\mathbf{q}_i^{spot} = 0$. Otherwise, additional intermediate goods will be purchased on the spot market to meet realized demand, provided that the cost of doing so is less than the value of the output v. Spot-market purchases are made from the cheapest intermediate goods producer, adjusting for the distance-based penalties (equation 4.3).

For intermediate goods producers, L^* is the cost-minimizing choice for given capacity investment K^* (equation 4.4). Equation 4.5 characterizes the optimal spot-market pricing. Intermediate goods producers engage in monopolistic competition and charge a mark-up over marginal costs. This mark-up is higher when substitutability is poor for the marginal buyer (that is, when $f'\left(\frac{1}{2n}\right)$ is high); and lower when competition is fierce (that is, when n is large). In the limit, as n approaches infinity - such that the distance between nodes shrinks to zero and intermediate goods become perfect substitutes - equation 4.5 simplifies down to price equals marginal cost (perfect competition). We explicitly assume that intermediate goods producers cannot engage in price discrimination, charging those at a greater distance less than those nearby. This assumption is natural in this context: intermediate goods producers may not fully observe the characteristics of the firms that seek to buy from them.

From equations 4.5 and 4.7, we see that non-scalable capacity K plays a key role through the marginal cost function. Higher capacity investments by any firm j in period 0 reduce its marginal cost of production in every state in period 1 (though more so in some states than in others). However, this decrease in marginal cost does not directly translate into proportionate increases in profit, especially if competing firms also expand their capacities, which would drive down the equilibrium spot price and pass on gains to final goods producers. This price response has important implications for investment in capacity, as the next section shows.

4.2 Period 0 equilibrium in the pre-order market

In period 0, the final goods producers take pre-order prices ϕ as given, form expectations over the state contingent distribution of spot prices at period 1, and submit pre-orders for intermediate goods $\mathbf{q}_i^{pre}$ to maximize their expected profit:

$$\begin{aligned}&\max_{\mathbf{q}_i^{pre}} E\left[\Pi_i\left(\mathbf{q}_i^{pre};\phi,\mathbf{p}^*\right)\right]\\ =&vE\left[Q\right]-\Pr\left(Q>Q_i^{pre}\right)E\left[\mathbf{p}^*\left(Q,w\right)\cdot\mathbf{q}_i^{spot,*}\left(Q,w\right)|Q>Q_i^{pre}\right]-\phi\cdot\mathbf{q}_i^{pre},\end{aligned} \tag{4.8}$$

where Π_i is the profit of firm i in period 1 (eqn. 4.1), $\mathbf{q}_i^{pre} := \left(q_{i0}^{pre}, q_{i1}^{pre}, \ldots, q_{i,n-1}^{pre}\right)'$ is the vector of pre-orders for intermediate goods, $\phi := \left(\phi_0, \phi_1, \ldots, \phi_{n-1}\right)'$ is the menu of pre-order prices, and $Q_i^{pre} := \sum_j \frac{1}{f_{ij}} q_{ij}^{pre}$ is the volume of final goods demand that can be met through pre-orders. The final goods producer anticipates that the realized demand for final goods Q may fall short of what could be produced from pre-orders Q_i^{pre} with probability $\left(1-\Pr\left(Q>Q_i^{pre}\right)\right)$. In such a scenario, the final goods producer will eschew the spot market in period 1, and not incur any additional costs beyond those associated with the pre-orders.[26] With complement probability $\Pr\left(Q>Q_i^{pre}\right)$, the final goods producer will need to purchase additional intermediate inputs on the spot market at expected cost $E\left[\mathbf{p}^*\left(Q,w\right)\cdot\mathbf{q}_i^{spot,*}\left(Q,w\right)|Q>Q_i^{pre}\right]$.

Simultaneously, each intermediate goods producer j sets pre-order price ϕ_j taking its competitors' prices ϕ_{-j} as given and commits to a level of non-scalable input factor K_j in order to maximize expected profit in period 1,

$$\max_{K_j,\phi_j} E\left[\Pi_j\left(K_j,\left\{\phi_j,\phi_{-j}\right\},\mathbf{q}_j^{pre,*}\right)\right]=\left[\phi_j Y_j^{pre,*}-rK_j\right]+E\left[p_j^* Y_j^{spot,*}-wL_j^*\right], \tag{4.9}$$

where $Y_j^{pre,*} := \int_0^1 q_{ij}^{pre,*} di$ and $Y_j^{spot,*} := \int_0^1 q_{ij}^{spot,*} di$ are the intermediate goods output required to meet equilibrium pre-orders and spot-market orders respectively.

We show that in a full-production symmetric equilibrium, the optimal pre-order price is equal to the unconditional expectation of spot-market prices. Without a discount over expected spot market prices, final goods firms pre-order only what is necessary to cover the lowest realization of demand. This limited demand for pre-orders affects the intermediate goods producer's incentive to invest in non-scalable production capacity.

Lemma 4. ***[Full production symmetric equilibrium in the pre-order market]*** *In period 0,*

1. *Each final goods producer i pre-orders only what is necessary to cover the lowest realization*

[26] As discussed in the previous section, we have imposed a constraint $\mathbf{q}_i^{spot} \geq 0$ ruling out the resale of pre-ordered intermediate goods.

of final goods demand from its nearest intermediat e goods supplier:

$$q_{ij}^{pre,*} = \begin{cases} f_{ij}\underline{Q} & \text{if } j \in \underline{J}(i;\phi), \text{ and } v \geq f_{ij}\phi_j \\ 0 & \text{otherwise} \end{cases} \quad \forall i \in [0,1], \tag{4.10}$$

where $\underline{J}(i;\phi) := \left\{ \tilde{j} \in J : f_{i\tilde{j}}\phi_{\tilde{j}} = \min\{f_i \circ \phi\} \right\}$ *denote the set of suppliers that provides the lowest effective pre-order price for i, which is equivalent under symmetry to the set of the nearest suppliers.*

2. *Each intermediate goods producer j:*

   (a) *sets pre-order prices to the unconditional expectation of spot-market prices*

$$\phi^* = E\left[p^*(Q,w)\right], \tag{4.11}$$

   (b) *invests in a level of non-scalable capacity K^* given by the optimality condition:*

$$\alpha \left(r + E\left[\underbrace{-\frac{dp^*}{dK}}_{<0} Y^{spot} \right] \right) K^* = (1-\alpha) E\left[wL^*\right]. \tag{4.12}$$

There is an important intermediate step to show why final goods firms find it optimal in equilibrium to pre-order only what is sufficient to meet the lowest realization of final goods demand. In appendix E (Lemma 6), we characterize final goods firms' demand for pre-orders ($Q_i^{pre,*} := \sum_j \frac{1}{f_{ij}} q_{ij}^{pre,*}$) in terms of the equation:

$$\phi = \Pr\left(Q > Q_i^{pre,*}\right) E\left[p^*(Q,w) \,|\, Q > Q_i^{pre,*}\right], \tag{4.13}$$

where $\left\{Q > Q_i^{pre,*}\right\}$ is the set of states in which the final goods firms need to purchase additional intermediate goods from the spot market in period 1. For every (effective) unit of intermediate goods pre-ordered in period 0, the final goods firm will need to order one fewer unit on the spot market, but only in states where $Q > Q_i^{pre,*}$. Thus, for a given pre-order price ϕ, final goods firms will pre-order just enough intermediate goods such that the ϕ is equal to the expected marginal savings on the spot market, accounting for the fact that larger pre-orders reduce the probability that spot-market purchases will be required.

The demand function for pre-orders (characterized by equation 4.13) has two immediate implications. First, aggregate pre-orders must be equalized across i in equilibrium ($Q_i^{pre,*} = Q^{pre,*}$ for all i). Second, the maximum sustainable pre-order price is the unconditional expectation of the spot market price $\phi = E[p^*]$. As final goods firms are risk neutral, they will not pre-order if $\phi > E[p^*]$. Likewise, intermediate goods firms do not have incentives to offer a discount on pre-orders (that is, pay a premium for insurance) by setting $\phi < E[p^*]$. Intermediate goods producers do not have incentives to reduce ϕ below $E[p^*]$ to attract more pre-orders because they expect to make more marginal profit on the spot market. Critically, any extra marginal costs incurred from lower capacity investments can also be passed on to final goods firms on the spot market along with a mark-up. With market power on the spot market, intermediate goods firms see no need to promote pre-orders to insure against correlated adverse supply shocks.

In equilibrium, therefore, we have a corner solution with $\phi = E[p^*]$ and $Q^{pre,*} = \underline{Q}$. Intermediate goods firms set pre-order prices at the level that makes final goods firms indifferent between no pre-orders at all and pre-ordering only what is necessary to cover the lowest realization of demand $\underline{Q}$. In effect, intermediate goods firms sets the highest possible pre-order price that drives the final goods firms to their participation constraint.[27]

Having characterized the equilibrium quantity and price of pre-orders, the intermediate goods suppliers determine the amount of production required to meet pre-orders ($Y^{pre,*}$) and forecast expected prices (p^*) and production on the spot market ($Y^{spot,*}$). The intermediate goods suppliers then invest in a level of non-scalable capacity K^* that minimizes expected costs for the anticipated level of production (equation 4.12). This optimality condition for K^* is similar to its analogues under the social planner benchmarks (equations 3.3 and 3.8 for the unconstrained and constrained cases, respectively), except for the addition of a final term $E\left[\frac{\partial p^*}{\partial K} Y^{spot}\right]$ that distorts the price of the non-scalable input factor. This final term captures the pecuniary externality that arises from enhanced market power and the overreliance on spot markets. It plays an important role in explaining the wedge between the decentralized market solution and the constrained optimal benchmark.

5 Decentralized solution versus constrained optimal benchmark

5.1 Under-investment in resilience

We can now prove the core proposition of the paper. The level of investment in the non-scalable capacity in a decentralized market setting (K^*) is suboptimally low when compared with the level

[27] Both the full production and the symmetry assumption play an important role here. We no longer have $\phi = E[p^*]$ as an equilibrium condition when these assumptions are relaxed. In this case, the analysis becomes more complex but the underlying economic intuition remains unchanged. See online appendix G for details. Likewise, we will also move away from this corner solution if agents are risk-averse, though the presence and qualitative properties of the market failures we identify are likely to be the same. That is, while with risk aversion there is likely to be more investment in capacity (greater resilience) in the market equilibrium, with more risk averse agents, (constrained) Pareto optimality also requires greater resilience, and a gap will remain between the two.

in the constrained optimal benchmark (K^{SP}).

Proposition 1. ***[Sub-optimal non-scalable capacity investment]*** $K^* < K^{SP}$.

Proof. We prove $K^* < K^{SP}$ by contradiction. This proof is instructive because it highlights the importance of the pecuniary externality $\frac{dp^*}{dK}$ and the overreliance on the spot market Y^{spot} as the main drivers behind the underinvestment in capacity.

First, by the full-production assumption, we know that the level of intermediate goods production is the same under both the decentralized solution and the constrained benchmark, $Y^*(Q,w) = Y^{SP}(Q,w)$, in all states of the world (Q,w).

The above equality implies that if $K^* = K^{SP}$, then $L^*(Q,w) = L^{SP}(Q,w)$ in every state, leading to a contradiction:

$$\begin{aligned}\alpha r K^{SP} &= (1-\alpha) E\left[wL^{SP}\right] = (1-\alpha) E\left[wL^*\right] \\ &= \alpha\left(r + E\left[-\underbrace{\frac{dp^*}{dK}}_{<0} Y^{spot}\right]\right) K^* > \alpha r K^*\end{aligned}$$

If instead $K^* > K^{SP}$, then $L^*(Q,w) < L^{SP}(Q,w)$ in every state, again giving rise to a contradiction:

$$\begin{aligned}\alpha r K^* &= E\left[\underbrace{\frac{dp^*}{dK}}_{<0} Y^{spot}\right] \alpha K^* + (1-\alpha) E\left[wL^*\right] \\ &< (1-\alpha) E\left[wL^*\right] < (1-\alpha) E\left[wL^{SP}\right] = \alpha r K^{SP}\end{aligned}$$

□

The proposition reveals that intermediate goods producers underinvest in capacity upfront because they are unable to fully capture the cost savings generated by increased investment. Specifically, each dollar saved through efficiency gains from capacity investment does not yield a corresponding one-dollar increase in profits, because a part of these gains is transferred to final goods producers through lower spot-market prices. The key term of interest is $E\left[\frac{dp^*}{dK} Y^{spot}\right]$, which captures the interaction between the pecuniary externality ($\frac{dp^*}{dK}$ the sensitivity of spot market prices to capacity investment) and the degree of reliance on the spot market (Y^{spot}).

Focus first on the price sensitivity term $\frac{dp^*}{dK}$ and recall that equilibrium spot prices can be expressed as a proportional mark-up over marginal costs: $p^* = (1+\mu) MC$. All else being equal, higher capacity investment K, lowers the marginal cost ($MC = \frac{w}{\alpha}\left(\frac{Y^*}{K^*}\right)^{\frac{1-\alpha}{\alpha}}$) in every possible state and thus lowers spot prices. The extent to which K matters depends on the *scalability* of the economy (α). As scalability improves and $\alpha \to 1$, the less important is K in production, and the externality shrinks.

The effect of K on marginal costs is amplified by the markup ($\mu = \frac{2f'\left(\frac{1}{2n}\right)\int_0^{\frac{1}{2n}} f(i)di}{f\left(\frac{1}{2n}\right)-2f'\left(\frac{1}{2n}\right)\int_0^{\frac{1}{2n}} f(i)di}$). The size of the mark-up depends on the *substitutability* between sectors, as measured by the distanced-based penalty function f (evaluated at the marginal buyer $i = \frac{1}{2n}$). Higher substitutability between sectors lowers mark-up and reduces the wedge between the decentralized solution and the constrained optimal benchmark in equilibrium. Lastly, another important way to reduce the wedge is through enhanced competition (that is, a larger n), which also reduces the amplification of marginal cost changes by reducing mark-ups.

Equally as important, the wedge results from an over-reliance on the spot market. Unlike an Arrow-Debreu economy, in which agents can trade contingent claims for every conceivable state of the world, in our model - much like real-world conditions - the set of contracts that can feasibly be written and traded is much smaller than the set of possible states. As a result, the pre-order, forwards, and futures markets will fall short of providing adequate risk insurance for intermediate goods producers. Downstream final goods producers fail to sufficiently compensate their suppliers for the pecuniary externality arising from the benefits of increased capital investment.[28] In the extreme case where the pre-order market for intermediate goods doesn't exist, all intermediate goods must be sourced from the spot market. The volume of transactions on the spot market, Y^{spot}, therefore takes its maximum value under full production, $Y^{spot} = 2Q\int_0^{\frac{1}{2n}} f_{ij}di$, and the gap between the decentralized solution and the constrained optimal benchmark widens.

5.2 The role of uncertainty, market power, and pre-order markets

The decentralized market will under-invest in non-scalable capacity as long as there is uncertainty around the future states of the world, and K cannot be adjusted instantaneously in response to shocks. The presence of market power, as represented by the price sensitivity term $\frac{dp^*}{dK}$, and market incompleteness, as represented by the reliance on the spot market (Y^{spot}),[29] determines the size of the wedge. We explore the role of each in greater detail below.

[28] In a sense, this pecuniary externality is a special case of the general pecuniary externality arising in economies without a complete set of AD securities analyzed by Greenwald and Stiglitz (1986) and first discussed in Stiglitz (1982).

[29] Note that since final goods firms pre-orders only what is necessary to cover the lowest realization of demand $\underline{Q}$ in equilibrium, the *expected* volume of transaction on the spot market is always strictly positive.

5.2.1 The model without uncertainty

We demonstrate that the presence of uncertainty is critical to the existence of the wedge in capacity investments between the decentralized solution and the planner's solution. If instead the future state of the world was perfectly predictable in advance, or equivalently all capacity investments can be adjusted instaneously in response to shocks, then the decentralized market would invest in the same level of capacity as the planner with perfect foresight. There would be no need for pre-order markets, and the existence of market power will only distort spot market prices (p^*) without affecting capacity investment (K).

Proposition 2. *[**Model with certainty**]: If every firm can determine their capacity investment after observing the state of the world in period* 1*, then:*

$$\begin{aligned} K_j^c(Q,w) &= K^{PF}(Q,w), \quad \forall j \in J, \forall (Q,w) \in \mathscr{S}, \\ p_j^c(n) &= \left(1 + \frac{1}{\varepsilon_j^{spot}(n) - 1}\right) MC_j =: \left(1+\mu_j\right) MC_j, \quad \forall j \in J, \\ \lim_{n\to\infty} p_j^c(n) &= MC_j, \end{aligned}$$

where $\varepsilon_j^{spot}(n)$ is the elasticity of demand for intermediate goods in the spot market when there are n intermediate goods producers.

Proof. The analysis is straightforward. First, observe that without uncertainty, the model collapses to a simple one-period problem, where all agents optimize for given state (Q,w). Second, the final goods firms' behavior is unchanged; they still source intermediate inputs from the lowest-priced supplier:

$$q_{ij}^c(\mathbf{p}) = \begin{cases} f_{ij}Q & \text{if } j \in \underline{J}(i;\mathbf{p}), \text{ and } v \geq f_{ij}p_j \\ 0 & \text{otherwise} \end{cases} \qquad \forall i \in [0,1].$$

Third, for the intermediate goods firms, their optimization problem becomes:

$$\Pi_j\left(\mathbf{p}_{-j}\right) = \max_{p_j,K_j,L_j} \left\{\left[p_jY_j(\mathbf{p})\right] - \left[w_jL_j + r_jK_j\right]\right\},$$

where $Y_j(\mathbf{p}) = \int_0^1 q_{ij}^c(\mathbf{p})\,di$ and $Y_j(\mathbf{p}) \leq L_j^{\alpha_j} K_j^{1-\alpha_j}$. From the first-order conditions to this problem, we derive the familiar optimality condition for inputs given the Cobb-Douglas production function, $\left(1-\alpha_j\right) w_jL_j^c = \alpha_j r_j K_j^c$, which gives an identical level of K to the first-best solution. Furthermore, spot market prices are set at a mark-up to marginal costs, $p_j^c = \left(1 + \frac{1}{\varepsilon_j^{spot}-1}\right) MC_j$, which depends on the elasticity of demand ε_j^{spot}. In the limit, as $n \to \infty$, we arrive at perfect competition with price equals to marginal costs. □

Thus, market power is not integral to the existence of under-investment in capacity; the inability to adjust non-scalable inputs to shocks is. In the presence of uncertainty, however, market power and the reliance on the spot market amplifies the distortion.

5.2.2 The model without the pre-order market

We now consider the model with uncertainty, but without a pre-order market in period 0. Capacity investments K still have to be put in place in period 0, before firms observe the realization of shocks. The analysis is essentially identical to that in the main text, except with $Q_i^{pre,*} = 0$. This implies that all demand for intermediate goods now need to be fulfilled on the spot market, so:

$$Y^{spot,only} = 2\int_0^{\frac{1}{2n}} f(i)\, Q > 2\int_0^{\frac{1}{2n}} f(i)\,(Q - Q^{pre}) = Y^{spot}.$$

Substituting the above into the optimality condition for capacity investments,

$$\alpha\left(r - E\left[\underbrace{\frac{dp^*}{dK}}_{<0} Y^{spot,only}\right]\right) K^* = (1-\alpha)\, E\left[wL^*\right],$$

it is straight-forward to see that capacity investments would be lower in the absence of pre-order markets.[30]

5.2.3 The model without market power

Lastly, we observe that even in the limiting case with $n \to \infty$, and $p^* = MC$, the impact of capacity investments on spot market prices is still strictly positive due to its direct effect on marginal costs (i.e., $\frac{dMC}{dK} > 0 \Rightarrow \frac{dp^*}{dK} > 0$). The wedge in capacity investments remains. Market power is not required for the existence of the distortion, it simply amplifies it (through mark-ups).

5.3 Policy implications

Using our model, it is possible to identify a number of ways to narrow the wedge between the supply network delivered by unfettered markets and the efficiently resilient network characterized under a constrained optimal benchmark.

[30] Note that since all production occurs in period 1 and the total amout of intermediate good produced is unchanged, the marginal cost of production - for given K - is unaffected by the presence/absence of the pre-order market. Similarly, the spot-market mark-up depends only on the substitutability of each intermediate goods. Thus $\frac{dp^*}{dK}$ remains unchanged when we eliminate the pre-order market.

First, the most straightforward strategy to address the externality in the model is to offer subsidies for capacity investments, thereby lowering the effective cost r incurred by intermediate goods producers for non-scalable capacity. Second, the government might extend tax benefits to downstream firms that engage in pre-orders or transact in the futures market or, alternatively, levy additional taxes on spot-market transactions to reduce dependency on spot markets. A third avenue is to reduce the sensitivity of spot prices to changes in capacity investments. This approach could entail structural economic reforms such as lowering entry barriers (including trade barriers), enacting stronger competition policies, and enhancing the substitutability of intermediate products, all of which could reduce supplier markups. Similarly, technological advancements in production scalability (through industrial policies) could shift the focus toward other input factors that can be more readily adjusted on short notice.

In practice, it may be hard to devise practical, implementable interventions. Directly subsidizing capacity investments offers a straightforward strategy, yet distinguishing such investments from other types of capital expenditure can be difficult, particularly in certain sectors. The government may want to intervene only in certain critical industries - for example computer chip production, where downstream externalities are especially significant and resilience is more important - by for instance, offering lower taxes for firms operating with excess capacity. While tax incentives for spot and pre-order markets can be effective in sectors like electricity, with its well-defined spot and futures markets, this approach becomes less straightforward in industries where market boundaries are more blurred. Industrial policies aimed at innovations increasing scalability and substitutability may be among the most practicable policies.

6 Concluding remarks

Since the pandemic and subsequent supply chain interruptions, the question of resilience has moved to the fore. Of course, we do not expect markets to be prepared for every shock, regardless of size, as doing so would be extraordinarily expensive. The overarching question is whether firms make appropriate preparations, measured against an appropriate benchmark? There are many reasons to think that they might not, critics of the market, for instance, complain about "short-termism". Moreover, intermediate firms may systematically underestimate the magnitude of upside potential demand.

We examine the normative question of resilience in a world with fully rational expectations and in which firms do not suffer from short-termism, showing that, nonetheless, there is a bias toward excessive vulnerability due to insufficient ex ante capacity investments by upstream intermediate goods producers. This shortfall arises because these producers cannot fully capture the returns on their capacity investments: A portion of the economic gains is transferred downstream to final goods producers through reduced spot-market prices.

To the best of our knowledge, our study is the first to incorporate interactions in both the spot and futures markets in such a normative analysis of supply networks, which is essential for addressing the question at hand. Performing this analysis in the context of differentiated competition necessarily entails a certain degree of complexity. For tractability and ease of exposition, we have introduced a number of simplifications, however, in online appendix G, we show how the results hold under significantly more general conditions. Most notably, we show that if there are very large shocks, such that the cost of meeting the market demand is so high that there are "unserved" customers (that is, Assumption A2 *Full Production* is not satisfied), then the analysis still holds.

The events of the past few years have made it clear that economists have paid insufficient attention to resilience. This paper is intended as a contribution to the nascent literature attempting to understand better why markets may have underinvested in resilience.

References

Acemoglu, D., V. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi (2012). The network origins of aggregate fluctuations. *Econometrica 80*(5), 1977–2016.

Acemoglu, D. and A. Tahbaz-Salehi. The macroeconomics of supply chain disruptions. *Review of Economic Studies, forthcoming*.

Acharya, V. V., M. Crosignani, T. Eisert, and C. Eufinger (2023). How do supply shocks to inflation generalize? evidence from the pandemic era in europe. *NBER Working Paper 31790*.

Amelkin, V. and R. Vohra (2024). Yield uncertainty and strategic formation of supply chain networks. *Networks 83(1)*.

Anupindi, R. and R. Akella (1993). Diversification under supply uncertainty. *Management Science 39*(8), 944–963.

Atalay, E., A. Hortacsu, J. Roberts, and C. Syverson (2011). Network structure of production. *Proceedings of the National Academy of Sciences 108*(13), 5199–5202.

Babich, V., G. Aydın, P.-Y. Brunet, J. Keppo, and R. Saigal (2012). Risk, financing and the optimal number of suppliers. In *Supply Chain Disruptions*, pp. 195–240. Springer.

Babich, V., A. N. Burnetas, and P. H. Ritchken (2007). Competition and diversification effects in supply chains with supplier default risk. *Manufacturing & Service Operations Management 9*(2), 123–146.

Baldwin, R. and R. Freeman (2022). Risks and global supply chains: What we know and what we need to know. *Annual Review of Economics 14*, 153–180.

Baqaee, D. and E. Farhi (2022). Supply and demand in disaggregated keynesian economies with an application to the covid-19 crisis. *American Economic Review 112*(5), 1397–1436.

Barrot, J.-N. and J. Sauvagnat (2016). Input specificity and the propagation of idiosyncratic shocks in production networks. *The Quarterly Journal of Economics 131*(3), 1543–1592.

Birge, J., A. Capponi, and P. Chen (2023). Disruption and rerouting in supply chain networks. *Operations Research 71*(2), 750–767.

Carvalho, V. (2014). From micro to macro via production networks. *Journal of Econometric Perspectives 28*(4), 23–48.

Carvalho, V., N. Makoto, Y. Saito, and A. Tahbaz-Salehi (2021). Production networks: A primer. *The Quarterly Journal of Economics 136*(2), 1255–1321.

Carvalho, V. and A. Tahbaz-Salehi (2019). Production networks: A primer. *Annual Review of Economics 11*(1), 635–663.

Crosignani, M., M. Macchiavelli, and A. Sliva (2019). Pirates without borders: The propagation of cyberattacks through firms' supply chains. *Journal of Financial Economics 147*(12), 5504–5517.

Di Giovanni, J., Ş. Kalemli-Özcan, A. Silva, and M. A. Yildirim (2022). Global supply chain pressures, international trade, and inflation. Technical report, National Bureau of Economic Research.

Elliot, M., B. Golub, and M. Leduc (2022). Supply network formation and fragility. *American Economic Review 112*(8), 2701–2747.

Elliott, M. and B. Golub (2022). Networks and economic fragility. *Annual Review of Economics 14*, 665–696.

Ferrari, A. (2022). Inventories, demand shocks propagation and amplification in supply chains. *arXiv preprint arXiv:2205.03862*.

Franzoni, F., M. Giannetti, and T. Roberto (2024). Supply chain shortages, large firmsâ market power, and inflation. *Swiss Finance Institute Research Paper No. 23-105*.

Greenwald, B. C. and J. E. Stiglitz (1986). Externalities in economies with imperfect information and incomplete markets. *The Quarterly Journal of Economics 101*(2), 229–264.

Grossman, G. M., E. Helpman, and A. Sabal (2023). Resilience in vertical supply chains. *NBER Working Paper 31739*.

Guerrieri, V., G. Lorenzoni, L. Straub, and I. Werning (2022). Macroeconomic implications of covid-19: Can negative supply shocks cause demand shortages? *American Economic Review 112*(5), 1437–1474.

Konczal, M. and N. Lusiani (2022). *Prices, profits, and power: an analysis of 2021 firm-level markups*. Roosevelt Institute New York.

Rothschild, M. and J. E. Stiglitz (1971). Increasing risk ii: Its economic consequences. *Journal of Economic theory 3*(1), 66–84.

Salop, S. C. (1979). Monopolistic competition with outside goods. *The Bell Journal of Economics*, 141–156.

Stiglitz, J. E. (1982). The inefficiency of the stock market equilibrium. *The Review of Economic Studies 49*(2), 241–261.

Stiglitz, J. E. (1986). *Prices, Competition, & Equilibrium*, Chapter Toward a More General Theory of Monopolistic Competition, pp. 22–69. Oxford: Philip Allan / Barnes & Noble Books.

Tomlin, B. (2006). On the value of mitigation and contingency strategies for managing supply chain disruption risks. *Management Science 52*(5), 639–657.

Appendix

A Proof for Lemma 1 - perfect foresight benchmark

We start with the optimization problem for the perfect foresight benchmark [PF] characterized in the main text. With perfect foresight, both L and K can be set as a function of the realized state (Q,w) in period 1. This is equivalent to saying that both L and K can be adjusted flexibly and simultaneously as the need arise. We thus have a standard Cobb-Douglas production for intermediate goods, with optimal input choices characterized by $(1-\alpha)wL = \alpha rK$, cost function $C\left(Y_j\right) = Y_j \left(\frac{w}{\alpha}\right)^{\alpha} \left(\frac{r}{1-\alpha}\right)^{1-\alpha}$, and constant marginal cost of production $\frac{\partial C}{\partial Y_j} = \frac{\partial C}{\partial q_{ij}} = \left(\frac{w}{\alpha}\right)^{\alpha} \left(\frac{r}{1-\alpha}\right)^{1-\alpha}$.[31]

[31] The first equality $\frac{\partial C}{\partial Y_j} = \frac{\partial C}{\partial q_{ij}}$ holds when the "feasibility of intermediate goods order flow" in the optimization problem [PF] constraint binds with equality in equilibrium.

Substituting the optimal input choices and the associated cost function into the original optimization problem [PF] reduces the dimension of the problem to one in order flows $\left\{q_{ij}\right\}$ only:

[Optimization Problem PF*]

$$W(Q,w)=\max_{\left\{q_{ij}\right\}_{i\in I,j\in J}}\left\{v\int_0^1\left[\sum_{j\in J}\frac{1}{f_{ij}}q_{ij}\right]di-\sum_{j\in J}\left[\left(\int_0^1 q_{ij}di\right)\left(\frac{w}{\alpha}\right)^{\alpha}\left(\frac{r}{1-\alpha}\right)^{1-\alpha}\right]\right\}$$

$$\text{s.t. }\sum_{j\in J}\frac{1}{f_{ij}}q_{ij}\leq Q_i\quad\forall i\in[0,1],\quad\text{[Demand cap]}$$

$$q_{ij}\geq 0\quad\forall i\in[0,1],j\in J.\quad\text{[Non-negative inputs]}$$

We can set up the Kuhn Tucker Lagrangian for Problem PF* as:

$$\mathscr{L}^{PF}=v\int_0^1\left[\sum_{j\in J}\frac{1}{f_{ij}}q_{ij}\right]di-\sum_{j\in J}\left[\left(\int_0^1 q_{ij}di\right)\left(\frac{w}{\alpha}\right)^{\alpha}\left(\frac{r}{1-\alpha}\right)^{1-\alpha}\right]-\sum_{i\in I}\lambda_i\left[\sum_{j\in J}\frac{1}{f_{ij}}q_{ij}-Q\right],$$

where by symmetry we have $Q_i=Q,\quad\forall i\in[0,1]$.

The first-order conditions (FOCs) with the corresponding complementary slackness conditions are given by:

$$q_{ij}^{PF}\frac{\partial\mathscr{L}^{PF}}{\partial q_{ij}}=q_{ij}\left(\frac{v}{f_{ij}}-\left(\frac{w}{\alpha}\right)^{\alpha}\left(\frac{r}{1-\alpha}\right)^{1-\alpha}-\frac{\lambda_i}{f_{ij}}\right)=0\quad\forall i\in I,j\in J,\tag{A.1}$$

$$\lambda_i\frac{\partial\mathscr{L}^{PF}}{\partial\lambda_i}=\lambda_i\left[\sum_{j\in J}\frac{1}{f_{ij}}q_{ij}-Q\right]=0\quad\forall i\in I.\tag{A.2}$$

We observe from the FOCs that for each final goods i, the corresponding Lagrangian multiplier λ_i, when strictly positive, is determined by the supplier $\underline{j}\in J$ that can provide the inputs most cheaply to i:

$$\begin{aligned}\lambda_i&=v-\min_{j\in J}\left\{f_{ij}\left(\frac{w}{\alpha}\right)^{\alpha}\left(\frac{r}{1-\alpha}\right)^{1-\alpha}\right\}\\&=v-f_{i\underline{j}}\left(\frac{w}{\alpha}\right)^{\alpha}\left(\frac{r}{1-\alpha}\right)^{1-\alpha},\end{aligned}\tag{A.3}$$

where $\underline{j}(i)\in\underline{J}(i):=\left\{\tilde{j}\in J|f_{i\tilde{j}}\left(\frac{w}{\alpha}\right)^{\alpha}\left(\frac{r}{1-\alpha}\right)^{1-\alpha}\leq f_{ij}\left(\frac{w}{\alpha}\right)^{\alpha}\left(\frac{r}{1-\alpha}\right)^{1-\alpha}\quad\forall j\in J\right\}$. Alternatively, if $v<\min_{j\in J}\left\{f_{ij}\left(\frac{w_j}{\alpha_j}\right)^{\alpha_j}\left(\frac{r_j}{1-\alpha_j}\right)^{1-\alpha_j}\right\}$, then $q_{ij}=0$ for all $j\in J$, $\tilde{Y}_i=0$ and $\lambda_i=0$ (i.e. it is not efficient for firm i to produce at all). This latter case is ruled out by the full production assumption

(A2).

Thus, combining equations (A.2, A.3), when $\lambda_i > 0$ final firm i will be allocated sufficient intermediate goods from its cheapest supplier to meet final demand Q:

$$q_{ij}^{PF} = \begin{cases} \frac{1}{n(\underline{J}(i))} f_{i\underline{j}} Q & \text{for } \underline{j} \in \underline{J}(i) \\ 0 & \text{for } j \neq \underline{J}(i) \end{cases},$$

where $n(\underline{J}(i))$ is the cardinality of the set $\underline{J}(i)$. In a symmetric equilibrium, the cheapest supplier(s) coincides with the closest supplier(s). With intermediate goods firms located equidistant around the circle, there are at most two closest suppliers for each i (e.g. nodes 0 and 1 for $i = \frac{1}{2n}$). In such cases when there are two closest suppliers, instead of tie-breaking by dividing order volumes in half, we assume each intermediate goods node j wins the tie-break to its right on the circle, but loses the tie-break to its left. This is loosely equivalent to imposing $n(\underline{J}(i)) = 1, \forall i \in [0,1]$; a convention we will adopt to simplify exposition without loss of generality.

Having solved for the optimal order flow $\left\{q_{ij}^{PF}\right\}$, we can now derive the aggregate output of intermediate goods. By symmetry every intermediate goods firm j will produce the same amount $Y_j = Y^{PF}$, $\forall j \in J$. So we can compute Y^{PF} from the perspective of firm $j = 0$, who is able to capture the two equal market segments to its left and right-hand side, $i \in \left[0, \frac{1}{2n}\right]$ and $\left[1 - \frac{1}{2n}, 1\right]$:

$$Y^{PF}(Q,w) = \int_0^1 q_{i0}^{PF} di = 2Q \int_0^{\frac{1}{2n}} f_{i0} di. \tag{A.4}$$

Finally, we can substitute the equilibrium intermediate goods production Y^{PF} into the Cobb-Douglas production function, combined with the optimality condition for inputs ($(1-\alpha)wL = \alpha K$) to derive explicit solutions for K^{PF} and L^{PF}:

$$K^{PF}(Q,w) = \left(\frac{w}{r}\frac{(1-\alpha)}{\alpha}\right)^{\alpha} \left[2Q \int_0^{\frac{1}{2n}} f_{i0} di\right],$$
$$L^{PF}(Q,w) = \left(\frac{r}{w}\frac{\alpha}{1-\alpha}\right)^{1-\alpha} \left[2Q \int_0^{\frac{1}{2n}} f_{i0} di\right].$$

This completes the proof for Lemma 1.

B Proof for Lemma 2 - social planner's constrained optimal problem

B.1 Period 1 optimization

Taking a similar approach to the Perfect Foresight benchmark, we form the corresponding Kuhn Tucker Lagrangian for the social planner's constrained optimal problem in period 1 [SP1]:

$$\begin{aligned}\mathscr{L} &= \max_{\left\{q_{ij}\in\mathbb{R}_+\right\}_{i,j}} v\int_0^1 \left(\sum_{j\in J}\frac{1}{f_{ij}}q_{ij}\right)di - \sum_{j\in J}\left[rK + w_j\left(\int_0^1 q_{ij}di\right)^{\frac{1}{\alpha}}K^{-\frac{(1-\alpha)}{\alpha}}\right]\dots \\ &\quad -\int_0^1 \lambda_i\left(\sum_{j\in J}\frac{1}{f_{ij}}q_{ij} - Q\right)di.\end{aligned}$$

From the Lagrangian we obtain the first-order derivatives with shortcomplementary slackness conditions:

$$\begin{aligned}q_{ij}\frac{\partial\mathscr{L}}{\partial q_{ij}} &= q_{ij}\left(\frac{v}{f_{ij}} - \frac{w}{\alpha}\left(\frac{\int_0^1 q_{ij}di}{K}\right)^{\frac{1-\alpha}{\alpha}} - \frac{\lambda_i}{f_{ij}}\right) = 0 \quad \forall i\in[0,1], \forall j\in J \\ \lambda_i\frac{\partial\mathscr{L}}{\partial\lambda_i} &= \lambda_i\left(\sum_{j\in J}\frac{1}{f_{ij}}q_{ij} - Q\right) = 0 \quad \forall i\in[0,1].\end{aligned}$$

where $\frac{v}{f_{ij}}$ is the marginal benefit from supplying i from j (i.e., q_{ij}), and $\frac{w}{\alpha}\left(\frac{\int_0^1 q_{ij}di}{K}\right)^{\frac{1-\alpha}{\alpha}}$ is the marginal cost. Later, in the final step of this proof, we will substitute out the endogenously determined K and q_{ij} to show that the marginal cost can be expressed as $\frac{w}{\alpha}\left(\frac{\int_0^1 q_{ij}di}{K}\right)^{\frac{1-\alpha}{\alpha}} = \left(\frac{w}{\alpha}\right)^{\alpha}\left(\frac{r}{1-\alpha}\right)^{1-\alpha}\left(\frac{wQ^{\frac{1}{\alpha}}}{E\left[wQ^{\frac{1}{\alpha}}\right]}\right)^{1}$

When λ_i, the Lagrangian multiplier for final goods firm i, is strictly positive, it is determined by the intermediate goods firm $\underline{j}$ that offers the lowest effective cost:

$$\begin{aligned}\lambda_i &= v - \min_{j\in J}\left\{f_{ij}\frac{w}{\alpha}\left(\frac{\int_0^1 q_{ij}di}{K}\right)^{\frac{1-\alpha}{\alpha}}\right\}, \\ &= v - f_{i\underline{j}}\frac{w}{\alpha}\left(\frac{Y}{K}\right)^{\frac{1-\alpha}{\alpha}},\end{aligned}$$

where $\underline{j}$ belongs to the set of lowest effective cost suppliers for i:

$$\underline{j} \in \underline{J}(i) := \left\{ \tilde{j} \in J | f_{i\tilde{j}} \frac{w}{\alpha} \left(\frac{\int_0^1 q_{ij} di}{K} \right)^{\frac{1-\alpha}{\alpha}} \leq f_{ij} \frac{w}{\alpha} \left(\frac{\int_0^1 q_{ij} di}{K} \right)^{\frac{1-\alpha}{\alpha}} \quad \forall j \in J \right\}.$$

Substitute the solution for λ_i into the first-order condition for q_{ij}, we arrive at the first part of the Lemma (eqn 3.6):

$$q_{ij}^{SP}(Q,w) = \begin{cases} f_{ij}Q & \text{if } j \in \underline{J}(i) \text{ and } v \geq f_{ij} \frac{w}{\alpha} \left(\frac{Y^{SP}}{K^{SP}} \right)^{\frac{1-\alpha}{\alpha}} \\ 0 & \text{otherwise} \end{cases},$$

where $Y^{SP} = \int_0^1 q_{ij}^{SP} di$, as required.

in the last part of the Lemma, the optimal choice of the scalable input factor in period 1, $L^{SP}(Q,w) = Y^{\frac{1}{\alpha}} K^{-\frac{(1-\alpha)}{\alpha}} = \left(\int_0^1 q_{ij}^{SP} di \right)^{\frac{1}{\alpha}} \left(K^{SP} \right)^{-\frac{(1-\alpha)}{\alpha}}$, is derived directly from the cost-minimization problem for the Cobb-Douglas production with partial delay.

B.2 Period 0 optimization

Recall that the period 1 value function for given K and realization of Q and w can be expressed as the difference between the value of final goods produced and the cost of the required intermediary goods:

$$W^{SP}(K|Q,w) = v\tilde{Y}^{SP} - n \left(rK + w \left(Y^{SP} \right)^{\frac{1}{\alpha}} K^{-\frac{(1-\alpha)}{\alpha}} \right),$$

In a full-production symemtric equilibrium, both the aggregate production of final goods $\left(\tilde{Y}^{SP}\right)$ and the production of intermediate goods by each firm $\left(Y^{SP}\right)$ are independent of capacity investment K.[32] Totally differentiating the expectation of W^{SP} with respect to K yields the desired

[32] By the full production assumption, we know that the marginal benefit of production exceeds the marginal cost when final goods producers source from their nearest intermediate goods producers. Therefore, aggregate final goods production in the symmetric equilibrium is equal to aggregate final goods demand:

$$\tilde{Y}^{SP} = \int_0^1 \left(\sum_{j \in J} \frac{1}{f_{ij}} q_{ij} \right) di = Q,$$

and each intermediate firm j's intermediate goods production in a symmetric equilibrium is given by:

$$Y^{SP} = Y_j = \int_0^1 q_i di = 2Q \int_0^{\frac{1}{2n}} f(i) di, \quad \forall j \in J.$$

first-order optimality condition for non-scalable capacity in period 0:

$$\frac{dE\left[W^{SP}\right]}{dK} = -n\left(r - \frac{(1-\alpha)}{\alpha}E\left[wY^{\frac{1}{\alpha}}K^{-\frac{1}{\alpha}}\right]\right) = 0$$

$$\Leftrightarrow 0 = r - \frac{1-\alpha}{\alpha}E\left[w\left(\frac{\left(L^{SP}\right)^{\alpha}\left(K^{SP}\right)^{1-\alpha}}{K^{SP}}\right)^{\frac{1}{\alpha}}\right]$$

$$\Leftrightarrow \alpha r K^{SP} = (1-\alpha)E\left[wL^{SP}\right].$$

Using the production function to substitute out $L^{SP} = \left(Y^{SP}\right)^{\frac{1}{\alpha}}\left(K^{SP}\right)^{-\frac{(1-\alpha)}{\alpha}}$ and re-arranging yields the explicit solution for K^{SP}:

$$K^{SP} = \left(\frac{1-\alpha}{\alpha}\frac{1}{r}\right)^{\alpha}\left(E\left[w\left(2Q\int_0^{\frac{1}{2n}} f(i)\,di\right)^{\frac{1}{\alpha}}\right]\right)^{\alpha},$$

where $Y^{SP} = 2Q\int_0^{\frac{1}{2n}} f(i)\,di$, as required for part 2 of the Lemma.

Finally, to complete the proof, we want to verify that this level of capacity investment $\left(K^{SP}\right)$ indeed leads to a full production equilibrium under assumption A2. We do this by substituting out the explicit expression for K^{SP} in the marginal cost function to show that in equilibrium the marginal cost of production is always below the valuation for the final goods (adjusted for the distance-based penalty):

$$\begin{aligned}
\widetilde{MC}(Q,w) &:= \frac{w}{\alpha}\left(\frac{Y^{SP}}{K^{SP}}\right)^{\frac{1-\alpha}{\alpha}} = \left(\frac{w}{\alpha}\right)^{\alpha}\left(\frac{r}{1-\alpha}\right)^{1-\alpha}\left(\frac{wQ^{\frac{1}{\alpha}}}{E\left[wQ^{\frac{1}{\alpha}}\right]}\right)^{1-\alpha} \\
&\leq \left(\frac{\bar{w}}{\alpha}\right)^{\alpha}\left(\frac{r}{1-\alpha}\right)^{1-\alpha}\left(\frac{\bar{w}\bar{Q}^{\frac{1}{\alpha}}}{E\left[wQ^{\frac{1}{\alpha}}\right]}\right)^{1-\alpha} \qquad \forall w, Q \\
&\leq \frac{v}{f\left(\frac{1}{2n}\right)} \quad \text{by assumption A2.}
\end{aligned}$$

B.3 Relationship between K^{PF} and K^{SP}

Recall that we can express $K^{PF}(Q,w)$ and K^{SP} explicitly as:

$$K^{PF}(Q,w) = \left(\frac{w}{r}\frac{(1-\alpha)}{\alpha}\right)^{\alpha}\left(2Q\int_0^{\frac{1}{2n}} f(i)\,di\right),$$
$$K^{SP} = \left(\frac{1}{r}\frac{1-\alpha}{\alpha}\right)^{\alpha}\left(2\int_0^{\frac{1}{2n}} f(i)\,di\right)\left(E\left[wQ^{\frac{1}{\alpha}}\right]\right)^{\alpha}.$$

Taking the expectation of K^{PF} over (Q,w) and re-arrange to give:

$$E\left[K^{PF}(Q,w)\right] = \left(\frac{1}{r}\frac{(1-\alpha)}{\alpha}\right)^{\alpha}\left(2\int_0^{\frac{1}{2n}} f(i)\,di\right)E\left[w^{\alpha}Q\right].$$

Therefore:

$$K^{SP} = \frac{\left(E\left[wQ^{\frac{1}{\alpha}}\right]\right)^{\alpha}}{E\left[w^{\alpha}Q\right]}E\left[K^{PF}(Q,w)\right].$$

By Jensen's inequality, given $g(x) := x^{\alpha}$ is concave for $\alpha \in (0,1)$ and $x > 0$, we have:

$$K^{SP} \geq E\left[K^{PF}(Q,w)\right],$$

as required. This completes the proof for Lemma 2.

C Sufficient condition for full production symmetric equilibrium in the decentralized solution

First, we establish the sufficient conditions for the existence of a full-production symmetric equilibrium.

Lemma 5. ***[Existence of Full Production Symmetric Equilibrium]:*** *For every configuration of the primitives of the model with the exception of v, $\mathscr{E}_{-v} = \{f(\cdot), \alpha, w, r, Q\}$, there exist a $\bar{v} \in \mathbb{R}_{++}$ such that the economies $\mathscr{E}(v) = \{f(\cdot), \alpha, w, r, Q, v \geq \bar{v}\}$ admits a full production symmetric equilibrium.*

Intuitively, the marginal benefit of production is increasing in the valuation of the final goods v, but the marginal cost is non-increasing in v. So, for every parameterization of the model, we can find a large enough $\bar{v}$ to guarantee full production in a symmetric equilibrium.

Formally, while assumption A2 establishes the sufficient conditions for full production under the social planner benchmarks, the corresponding full-production condition for the decentralized case is given by:

$$v \geq f\left(\frac{1}{2n}\right) p^* = f\left(\frac{1}{2n}\right) \mu\left(n\right) \widetilde{MC} = f\left(\frac{1}{2n}\right) \mu\left(n\right) \frac{\bar{w}}{\alpha} \left(\frac{\bar{Q} \int_0^{\frac{1}{2n}} f\left(i\right) di}{K^*}\right)^{\frac{1-\alpha}{\alpha}},$$

where p^* is the equilibrium price for intermediate goods, $\mu\left(n\right) := \frac{f\left(\frac{1}{2n}\right)}{f\left(\frac{1}{2n}\right) - 2\frac{f'\left(\frac{1}{2n}\right)}{f\left(\frac{1}{2n}\right)} \int_0^{\frac{1}{2n}} f(i) di}$ is the mark-up over marginal costs, and K^* is the equilibrium level of non-scalable capacity. We argue that for every possible parameterization of the other primitives, there exists a $\bar{v} \in \mathbb{R}_{++}$ that guarantees full production.

Consider an arbitrary economy $\mathscr{E}\left(\tilde{v}\right) = \{f\left(\cdot\right), \alpha, w, r, Q; \tilde{v}\}$ with valuation $\tilde{v}$. We want to show that by varying $\tilde{v}$ we can always construct an economy $\mathscr{E}\left(\bar{v}\right) = \{f\left(\cdot\right), \alpha, w, r, Q; \bar{v}\}$ that supports a full production symmetric equilibrium holding all other primitives the same. To do this, we compute $K^*\left(\tilde{v}\right)$, the associated equilibrium level of capacity investment *assuming* full production, and the corresponding $\overline{MC}\left(\tilde{v}\right) = f\left(\frac{1}{2n}\right) \mu\left(n\right) \frac{\bar{w}}{\alpha} \left(\frac{\bar{Q} \int_0^{\frac{1}{2n}} f(i) di}{K^*\left(\tilde{v}\right)}\right)^{\frac{1-\alpha}{\alpha}}$, the highest possible realization of marginal costs in that economy. Note that $K^*\left(v\right)$ is a non-decreasing function of v and therefore $\overline{MC}\left(v\right)$ is a non-increasing function of v (i.e. the marginal cost of production in any full production equilibrium does not increase when the valuation increases). Then if $\tilde{v} \geq \overline{MC}\left(\tilde{v}\right)$, then the economy $\mathscr{E}\left(\tilde{v}\right)$ admits a full production symmetric equilibrium characterized by $K^*\left(\tilde{v}\right)$. If instead $\tilde{v} < \overline{MC}\left(\tilde{v}\right)$, let $\bar{v} = \overline{MC}\left(\tilde{v}\right) > \tilde{v}$. Then $\bar{v} = \overline{MC}\left(\tilde{v}\right) \geq \overline{MC}\left(\bar{v}\right)$. And every economy $\mathscr{E}\left(v\right) = \{f\left(\cdot\right), \alpha, w, r, Q; v \geq \bar{v}\}$ admits a full production symmetric equilibrium as required.

Second, we remark that the full production assumption also enables us to avoid problems of non-differentiability in the demand function. In a classical treatment of the circular economy, Salop (1979) segments the demand function for intermediate goods into three sections: a "monopoly" regime (whereby the firm acts as if it is a monopoly); a "competitive" regime (where it engages in Bertrand competition with its neighbors); and a "super-competitive" regime (where it prices so aggressively as to take over its neighbor's native market). The demand function exhibits a kink at the intersection between the monopoly and competitive regime, and makes a discontinuous jump between the competitive and super-competitive regime. We can rule out equilibria falling under the super-competitive regime by setting a sufficiently steep distance-based penalty function; and for the purpose of the main analyses in section 3 and 4, the full production assumption ensures the demand function is continuously differentiable. In appendix G we relax the full production assumption to examine the interplay between the competitive and monopoly regime.

D Proof of Proposition 3: full production symmetric equilibrium in the spot market

In period 1, the equilibrium spot market orders $q_{ij}^{spot,*}$ by final goods firms, and the purchase of scalable inputs L^* by intermediate goods firms, take a similar form to their corresponding expressions under the constrained optimal benchmark. We skip their derivations to avoid repetition, and concentrate instead on the solution for the spot market price p^*, given by equation 4.5.

To solve for p^*, we will first need to derive the demand function facing the intermediate goods firm $j = 0$ on the spot market. For now, we will also need to conjecture that the aggregate volume of pre-orders must be equalized across all final goods firms in equilibrium: $Q_i^{pre} = Q^{pre}$, $\forall i \in I$, a result that we will prove formally later in appendix E.

D.1 Finding the slope of the demand curve

The demand curve facing each intermediate goods firm is piece-wise linear (when plotted against p_j, for given $\mathbf{p}_{-j}$). To see this, note that the period 1 equilibrium is governed by two indifference thresholds. First, for given price vector $(p_0, \mathbf{p}_{-0})$, the *participation threshold* for firm $j = 0$, $\bar{i}_0$, is defined as the final goods firm that is indifferent between buying inputs from intermediate goods firm $j = 0$ and not producing at all:

$$f(\bar{i}_0)\, p_0 = v, \quad \forall p_0 \in \left[\frac{v}{f\left(\frac{1}{2}\right)}, v\right] \tag{D.1}$$

Second, the *competitive threshold* $\bar{i}_{0,1}$ is the marginal final goods producer that is indifferent from buying from supplier node $j = 0$ and $j = 1$:

$$f(d(\bar{i}_{0,1}, 0))\, p_0 = f(d(\bar{i}_{0,1}, 1))\, p_{-0} \tag{D.2}$$

Hence the demand cuve facing firm $j = 0$ depends on the lower envelope of the participation and competitive threshold functions:

$$Y_0^{spot} = 2\int_0^{\bar{i}_0^*} f(i) \cdot (Q - Q^{pre})\, di \tag{D.3}$$

where

$$\bar{i}_0^* := \min\left\{\bar{i}_0(p_0), \bar{i}_{0,1}(p_0, p_{-0})\right\} \tag{D.4}$$

When the slope of the demand curve is well-defined (i.e., away from the knife-edge case when

$\bar{i}_0(p_0) = \bar{i}_{0,1}(p_0, p_{-0})$), it is given by:

$$\frac{dY_0^{spot}}{dp_0} = 2\left[\frac{d\bar{i}_0^*}{dp_0} f(\bar{i}_0^*)(Q - Q^{pre})\right] \tag{D.5}$$

Given the full production assumption, we have $\bar{i}_0^* := \bar{i}_{0,1}(p_0, p_{-0})$. So

$$\begin{aligned}
\frac{d\bar{i}_0^*(p_0 = p^*, p_{-0} = p^*)}{dp_0} = \frac{d\bar{i}_{0,1}(p^*, p^*)}{dp_0} &= -\frac{\partial\left(\frac{f(\bar{i}_{0,1})}{f\left(\frac{1}{n} - \bar{i}_{0,1}\right)} p_0\right)/\partial p_0}{\partial\left(\frac{f(\bar{i}_{0,1})}{f\left(\frac{1}{n} - \bar{i}_{0,1}\right)} p_0\right)/\partial \bar{i}_{0,1}} \\
&= -\frac{1}{\left(\frac{f'(\bar{i}_{0,1})}{f(\bar{i}_{0,1})} + \frac{f'\left(\frac{1}{n} - \bar{i}_{0,1}\right)}{f\left(\frac{1}{n} - \bar{i}_{0,1}\right)}\right) p_0}
\end{aligned}$$

Impose symmetry $\bar{i}_0^* = \bar{i}_{0,1}(p^*, p^*) = \frac{1}{2n}$ to get

$$\frac{d\bar{i}_{0,1}(p^*, p^*)}{dp_0} = -\frac{f\left(\frac{1}{2n}\right)}{2f'\left(\frac{1}{2n}\right) p_0} \tag{D.6}$$

D.2 Solving for the optimal spot market price

We can derive the following first-order condition with respect to p_0 from the intermediate goods producer $j = 0$'s optimization problem (equation 4.2):

$$\left(p^* - \frac{w}{\alpha}\left(\frac{Y}{K}\right)^{\frac{1-\alpha}{\alpha}}\right) = Y^{spot} / \left(-\frac{dY^{spot}}{dp^*}\right) \tag{D.7}$$

where $\frac{w}{\alpha}\left(\frac{Y}{K}\right)^{\frac{1-\alpha}{\alpha}}$ is the marginal cost of production for intermediate goods; $Y := Y^{spot} + Y^{pre}$ is the total amount of intermediate goods production; and $\frac{dY^{spot}}{dp^*}$ is the slope of the demand curve in the spot market. By imposing symmetry we get $p^* = p_0^* = p_j^*$ for all $j \in J$.

Substituting equations D.5 and D.6 into equation D.7 gives the optimal spot-market price as required:

$$\begin{aligned}
(p^* - MC) &= \frac{2\int_0^{\frac{1}{2n}} f(i)(Q - Q^{pre})\,di}{-2\left[-\frac{f\left(\frac{1}{2n}\right)}{2f'\left(\frac{1}{2n}\right)p^*} f\left(\frac{1}{2n}\right)(Q - Q^{pre})\right]} \\
\Rightarrow p^* &= \left(1 + \frac{2f'\left(\frac{1}{2n}\right)\int_0^{\frac{1}{2n}} f(i)\,di}{\left(f\left(\frac{1}{2n}\right)\right)^2 - 2f'\left(\frac{1}{2n}\right)\int_0^{\frac{1}{2n}} f(i)\,di}\right) MC
\end{aligned}$$

E Proof of proposition 4: full production symmetric equilibrium in the pre-order market

E.1 Final goods producers in period 0

We will start by verifying that the conjecture $Q_i^{pre} = Q^{pre}, \forall i \in I$ is indeed an equilibrium solution. Recall that $Q_i^{pre} := \sum_j \frac{1}{f_{ij}} q_{ij}^{pre}$ denotes the level of *final goods* demand fulfilled through pre-orders.

Lemma 6. ***[Optimal Pre-orders]*** *In a full-production symmetric equilibrium, each final goods producer i will:*

1. *pre-order from the intermediate goods producers that sets the lowest effective-price for i.*

$$q_{ij}^{pre,*} = \begin{cases} f_{ij} Q_i^{pre,*} & \text{if } j \in \underline{J}(i;\phi)\text{, and } f_{ij}\phi_j \le v \\ 0 & \text{otherwise} \end{cases} \tag{E.1}$$

where $\underline{J}(i;\phi) := \left\{ \tilde{j} \in J : f_{i\tilde{j}}\phi_{\tilde{j}} = \min \{\mathbf{f}_i \circ \phi\} \right\}$ denote the set of suppliers that provides the lowest effective price for i.

2. *target to fulfill a level of final goods demands $Q_i^{pre,*}$ through pre-orders such that the marginal cost of pre-orders is equal to its expected marginal benefit.*

$$f_{i\tilde{j}}\phi_{\tilde{j}} = \Pr\left(Q > Q_i^{pre,*}\right) E\left[f_{i\hat{j}} p_{\hat{j}}^* (Q,w) \,|\, Q > Q_i^{pre,*} \right], \quad \text{for } \tilde{j} \in \underline{J}(i;\phi), \hat{j} \in \underline{J}(i;\mathbf{p}) \tag{E.2}$$

where $\underline{J}(i;\mathbf{p}) := \left\{ \hat{j} \in J : f_{i\hat{j}} p_{\hat{j}} = \min \{\mathbf{f}_i \circ \mathbf{p}\} \right\}$

Furthermore, imposing symmetry implies

$$\phi = \Pr\left(Q > Q_i^{pre,*}\right) E\left[p^* (Q,w) \,|\, Q > Q_i^{pre,*} \right] \quad \forall i \in [0,1] \tag{E.3}$$

so that the volume of final goods demand fulfilled through pre-orders must be equalized across all final goods firms:

$$Q_i^{pre,*} = Q^{pre,*} \quad \forall i \in [0,1] \tag{E.4}$$

Equation E.2 is the first-order condition of final goods producer i's period 0 optimization problem. It gives an implicit expression for the equilibrium volume of final goods demand fulfilled through pre-orders $Q_i^{pre,*}$ as a function of spot and pre-order prices (p^*, ϕ). On the left hand side of the equation, $f_{i\tilde{j}}\phi_{\tilde{j}}$ is the effective marginal cost of pre-orders. On the right hand side is the expected marginal benefit of pre-orders, which is equal to the probability that the spot market order of i will be strictly positive $\Pr\left(Q > Q_i^{pre,*}\right)$, multiplied by the conditional expectation of the lowest effective spot price, given i's spot-market order is strictly positive $E\left[f_{i\hat{j}}p^*(Q,w)\,|Q > Q_i^{pre,*}\right]$. Under symmetry, $p_j^* = p^*$ and $\phi_j = \phi$ for all $j \in J$; so the nearest intermediate goods node to i will always provide the lowest effective price on both the pre-order and spot markets: $f_{i\tilde{j}} = f_{i\hat{j}}$. Equation E.2 can thus be simplified to equation E.3, which we can also interpret as the demand function for pre-orders for given pre-order price ϕ. Equation E.3 has two immediate implications: (1) $Q_i^{pre,*}$ must be equalized across i (equation E.4); and (2) the highest sustainable pre-order price is $\phi = E[p^*]$, in which case the final goods producers will only pre-order to satisfy the minimal possible realization of final goods demand $Q^{pre,*} = \underline{Q}$. For any pre-order price greater than the unconditional expectation of the spot market price, the aggregate quantity of pre-order will be zero. So we can view equation E.3 also as a participation constraint for final goods firms on the pre-order market.

E.2 Intermediate goods producers in period 0

Recall the expected profit function for intermediate goods producers:

$$
\begin{aligned}
\max_{\phi,K} E\left[\Pi_j\right] =& E\left[p^*Y^{spot} - wL^*\right] + \phi Y^{pre} - rK \\
=& \int_{Q^{pre}}^{\bar{Q}} \int_w \left(p^*Y^{spot}\right) h(w)\, g(Q)\, dw dQ \dots \\
& - \left(\int_{\underline{Q}}^{Q^{pre}} \int_w (wL^*)\, h(w)\, g(Q)\, dw dQ + \int_{Q^{pre}}^{\bar{Q}} \int_w (wL^*)\, h(w)\, g(Q)\, dw dQ\right) \dots \\
& + \phi Y^{pre} - rK
\end{aligned}
\tag{E.5}
$$

We note that in a symmetric full production equilibrium, the aggregate production of intermediate goods $Y := Y^{pre} + Y^{spot} = 2Q\int_0^{\frac{1}{2n}} f(i)\, di$ is exogenously pinned down by the realization of final goods demand Q, and the distance-based penalty function f. But the relative importance of the spot market and the pre-order market (Y^{pre} and Y^{spot}) depends on the volume of pre-orders q^{pre}, which is in turn determined by the choice of the pre-order price ϕ. On the other hand, the level of non-scalable capacity investment K affects the period 1 equilibrium spot market price $p^*(Q,w)$ and scalable input demand $L^*(Q,w)$ in each possible state. We examine the optimality conditions for ϕ and K in turn.

First we take the derivative of expected profits with respect to ϕ. With some algebra, we can show that

$$
\begin{aligned}
\frac{dE\left[\Pi_j\left(K_j,\phi_j,q_j^{pre}\right)\right]}{d\phi} =&\phi\left(-\frac{dY^{pre}}{d\phi}\right)-\Pr\left(Q\leq Q^{pre}\right)E\left[w\frac{\partial L^*}{\partial Y}|Q\leq Q^{pre}\right]\frac{dY^{pre}}{d\phi}\dots \\
&+\left(Y^{pre}+\phi\frac{dY^{pre}}{d\phi}\right) \qquad \text{(E.6)} \\
=&-\Pr\left(Q\leq Q^{pre}\right)E\left[w\frac{\partial L^*}{\partial Y}|Q\leq Q^{pre}\right]\frac{dY^{pre}}{d\phi}+Y^{pre} \\
=&\Pr\left(Q\leq Q^{pre}\right)E\left[w\frac{\partial L^*}{\partial Y}|Q\leq Q^{pre}\right]\left(-\frac{dY^{pre}}{d\phi}\right)+Y^{pre}>0 \qquad \text{(E.7)}
\end{aligned}
$$

This imply that the equilibrium must be a corner solution. Intermediate goods producers would like to set the highest possible pre-order price subject to the participation constraint of final goods producers (eqn E.3). Thus, from Lemma 6, equilibrium pre-orders will equal to the lowest possible realization of final goods demand, and the equilibrium pre-order price will equal the unconditional expectation of the spot market price:

$$Q^{pre,*}=\underline{Q} \qquad \text{(E.8)}$$

$$\phi^*=E\left[p^*\left(Q,w\right)\right] \qquad \text{(E.9)}$$

Next we take the derivative of the expected profit with respect to K:

$$E\left[\frac{\partial p^*}{\partial K}Y^{spot}\right]-E\left[w\frac{\partial L^*}{\partial K}\right]-r=0 \qquad \text{(E.10)}$$

where $L^*=\left(Y^{pre}+Y^{spot}\right)^{\frac{1}{\alpha}}(K)^{-\frac{1-\alpha}{\alpha}}$, so

$$
\begin{aligned}
\frac{\partial L^*}{\partial K}&=-\left(\frac{1-\alpha}{\alpha}\right)\left(Y^{pre}+Y^{spot}\right)^{\frac{1}{\alpha}}K^{-\frac{1}{\alpha}} \\
&=-\left(\frac{1-\alpha}{\alpha}\right)\frac{L^*}{K}
\end{aligned}
$$

Substituting $\frac{\partial L^*}{\partial K}$ back into the first-order condition to give

$$(1-\alpha)E\left[wL^*\right]=\alpha\left(r-E\left[\frac{\partial p^*}{\partial K}Y^{spot}\right]\right)K^* \qquad \text{(E.11)}$$

as required.

Online appendix

This online appendix discusses the robustness of our results to the full production assumption imposed on the main body of the paper.

F Relaxing the full production assumption for the constrained social planner

As discussed in the main body, we can show that the optimality condition for non-scalable production capacity K^{SP} (equation 3.8) holds with and without the full production assumption. To elucidate this point, note that when the full production assumption is relaxed, there may exist states of the world $(\tilde{Q},\tilde{w})$ where some final goods firms situated far from intermediate goods production firms do not find it optimal to produce at all. In other words, let $\bar{i}_0^{SP}\left(\tilde{Q},\tilde{w}\right)$ represent a "threshold" firm in the final goods sector. This firm is indifferent between sourcing inputs from intermediate goods firm $j=0$ and opting out of production altogether in state $\left(\tilde{Q},\tilde{w}\right)$: $\frac{v}{f\left(\bar{i}_0^{SP}\right)} = \left(\frac{\tilde{w}}{\alpha}\right)^{\alpha}\left(\frac{r}{1-\alpha}\right)^{1-\alpha}\left(\frac{\tilde{w}\tilde{Q}^{\frac{1}{\alpha}}}{E\left[wQ^{\frac{1}{\alpha}}\right]}\right)^{1-\alpha}$. Hence there may exist states $\left(\tilde{Q},\tilde{w}\right)$ whereby $\bar{i}_0^{SP}\left(\tilde{Q},\tilde{w}\right) < \frac{1}{2n}$, and the market segment $\left[\bar{i}_0^{SP}\left(\tilde{Q},\tilde{w}\right),\frac{1}{2n}\right]$ on the circle produces no final goods outputs and experiences "empty shelves".

To simplify notation, we will use s as a short-hand for states (Q,w), and denote the set of "empty shelves" states by $M := \left\{s\in\mathscr{S} : \bar{i}_0^{SP}(s) < \frac{1}{2n}\right\}$. In the M subset of states, production is less than "full" and intermediate goods producers operate like monopolies in their own disjoint local market segments.

At first glance, one might expect that a constrained social planner, operating without the benefit of an economy always in full production, might want to increase capacity investment K^{SP} above the level implied by the optimality condition: $\alpha r K^{SP} = (1-\alpha)E_s\left[wL^{SP}\right]$. The extra capacity would aim to: (a) reduce the range of shocks (i.e. the set of states M) where supply networks are strained, and (b) reduce marginal costs and increase production Y in such stressed states.

However, somewhat counterintuitively, we can show that in the constrained social planner's problem the indirect effects of capacity through its impact on both (a) the probability of stressed states $\left(\frac{d\Pr(M)}{dK}\right)$, and (b) output in such states $\left(\frac{dY}{dK}|_{(Q,w)\in M}\right)$ are zero in equilibrium. Specifically, even though $\frac{d\Pr(M)}{dK}$ is strictly negative and $\frac{dY}{dK}|_{(Q,w)\in M}$ strictly positive (as expected), both terms are multiplied by another term that is zero in equilibrium.[33] Therefore, irrespective of whether the

[33]This result bears resemblance to the Envelope Theorem, in which the total derivative of the value function with respect to the parameters of the model is equal to its partial derivative. Here K is the choice variable, but the total

full-production assumption holds, $\alpha r K^{SP} = (1-\alpha) E\left[wL^{SP}\right]$. Only the direct effects of K on the choice of inputs and on the cost of production matter.

Formally, recall that the period 1 value function for given K and realization of states s can be expressed as the difference between the value of final goods produced and the cost of the required intermediary goods:

$$W^{SP}(K;s) = v\tilde{Y}^{SP}(K;s) - n\left(rK + w\left(Y^{SP}(K;s)\right)^{\frac{1}{\alpha}} K^{-\frac{(1-\alpha)}{\alpha}}\right)$$
$$= v\left(2n\int_0^{\min\left\{\frac{1}{2n},\tilde{i}_0^{SP}(K;s)\right\}} Qdi\right) - n\left(rK + w\left(Y^{SP}(K;s)\right)^{\frac{1}{\alpha}} K^{-\frac{(1-\alpha)}{\alpha}}\right),$$

where $\tilde{i}_0^{SP}(K;s)$ is the *threshold buyer* for intermediate goods 0, (implicitly) defined as the final goods firm i for which the marginal benefit of sourcing inputs from $j=0$ equals the marginal cost:

$$\frac{v}{f\left(\tilde{i}_0^{SP}\right)} = \frac{w}{\alpha}\left(\frac{Y^{SP}}{K}\right)^{\frac{1-\alpha}{\alpha}}. \tag{F.1}$$

The upper limit of integration, $\min\left\{\frac{1}{2n}, \tilde{i}_0^{SP}\right\}$, reflects the possibility of "regime switching" when the full production assumption is relaxed. When the economy operates at a full-production equilibrium, the relevant threshold buyer for intermediate goods firm $j=0$ is given by $i=\frac{1}{2n}$, the final goods firm located at the half way point between $j=0$ and $j=1$. This is a *competitive regime*, where intermediate goods firms engage in monopolistic competition. But when the indifference threshold $\tilde{i}_0^{SP}$ falls below $\frac{1}{2n}$, we have instead a *local monopolies* regime, characterized by a gap in market coverage between two supplier nodes (e.g. between $j=0$ and $j=1$). We account for the possibility of "regime switching" between the competitive regime and the local monopolies regime in the analyses that follows.

Totally differentiating the expectation of W^{SP} with respect to K will yield the desired first-order optimality condition for non-scalable capacity in period 0.

$$\max_K E\left[W^{SP}(K,s)\right] = vE\left[\tilde{Y}^{SP}(K,s)\right] - nrK - nE\left[w\left(Y^{SP}\right)^{\frac{1}{\alpha}} K^{-\frac{(1-\alpha)}{\alpha}}\right],$$

derivative of $W^{SP}(K|Q,w)$ with respect to K is also equal to its partial derivative.

with first-order condition:

$$\begin{aligned} 0 &= \frac{E\left[W^{SP}\left(K^{SP},s\right)\right]}{dK} \\ &= vE\left[\frac{d\tilde{Y}^{SP}(K,s)}{dK}\right]-nr-nE\left[\frac{w}{\alpha}\left(\frac{Y^{SP}}{K}\right)^{\frac{(1-\alpha)}{\alpha}}\frac{dY^{SP}(K,s)}{dK}-w\frac{1-\alpha}{\alpha}\left(\frac{Y^{SP}}{K}\right)^{\frac{1}{\alpha}}\right]. \quad \text{(F.2)} \end{aligned}$$

Unlike in the full-production case, both the aggregate production of final goods $\left(\tilde{Y}^{SP}\right)$ and the production of intermediate goods by each firm $\left(Y^{SP}\right)$ will depend on the choice of period 0 capacity investment K. These two terms ($E\left[\frac{d\tilde{Y}^{SP}}{dK}\right]$,$E\left[\frac{dY^{SP}}{dK}\right]$) will capture the indirect effects of K on aggregate welfare. We examine each in turn.

First, for each K, we order the set of states s in ascending order of their implied threshold for regime switch $\bar{i}_0^{SP}(K,s)$, such that states in the subset M are at the bottom. This allows us to express the expected level of final goods production as:

$$\begin{aligned} E\left[\tilde{Y}^{SP}(K,s)\right] &= 2nE\left[\int_0^{\min\left\{\frac{1}{2n},\bar{i}_0(K,s)\right\}}Qdi\right] = 2n\int_s\left[\min\left\{\frac{1}{2n},\bar{i}_0(K,s)\right\}\cdot Q\right]g(s)\,ds \\ &= 2n\left[\int_{\underline{s}}^{s^M(K)}\bar{i}_0(K,s)\,Qg(s)\,ds+\int_{s^M(K)}^{\bar{s}}\frac{1}{2n}Qg(s)\,ds\right], \end{aligned}$$

where $g(s)$ is the probability density function of state s. $s^M(K) := \left\{\tilde{s}\in\mathscr{S} : \bar{i}_0^{SP}(K,\tilde{s}) = \frac{1}{2n}\right\}$ is the state at the knife-edge intersection of the competitive regime and the local monopolies regime, with $\frac{ds^M(K)}{dK} < 0$ (i.e, higher capacity investment reduces the set of states under the local monopolies regime). Totally differentiating the above expression with respect to K yields:

$$\begin{aligned} \frac{dE\left[\tilde{Y}^{SP}(K,s)\right]}{dK} &= 2n\left[\int_{\underline{s}}^{s^M}\frac{d\bar{i}_0(K,s)}{dK}Qg(s)\,ds+\frac{ds^M}{dK}g\left(s^M\right)\underbrace{\bar{i}_0\left(K,s^M\right)}_{:=\frac{1}{2n}}Q-\frac{ds^M}{dK}g\left(s^M\right)\frac{1}{2n}Q\right] \\ &= 2n\int_{\underline{s}}^{s^M}\frac{d\bar{i}_0(K,s)}{dK}Qg(s)\,ds \\ &= 2n\Pr(M)\,E\left[\frac{d\bar{i}_0(K,s)}{dK}Q|s\in M\right]. \end{aligned}$$

In the second equality above, we observe that the terms associated with the change in the probability of monopoly states $\frac{d\Pr(M)}{dK} \equiv \frac{ds^M}{dK}g\left(s^M\right)$ cancels out because at the knife-edge the limiting local monopoly regime is identical to the competitive regime ($\bar{i}_0\left(K,s^M\right) = \frac{1}{2n}$). With a bit more algebra, we can also verify that it is safe to exchange the order of differentiation and expectation such that $E\left[\frac{d\tilde{Y}^{SP}(K,s)}{dK}\right]=\frac{dE\left[\tilde{Y}^{SP}(K,s)\right]}{dK}$. Thus, we can safely ignore the effect of K on the probability of regime change in the first-order condition (equation F.2).

Second, we repeat the same process for the expected level of intermediate goods production:

$$Y^{SP}(K,s)=2\int_0^{\min\left\{\frac{1}{2n},\bar{i}_0(K,s)\right\}}(f(i)Q)\,di.$$

Note that intermediate goods production is increasing in capacity investments only in the states $s\in M$, where we don't have full production already:

$$\frac{dY^{SP}(K,s)}{dK}=\begin{cases}0 & \text{for } s\notin M\\ 2Q\frac{d\bar{i}_0(K,s)}{dK}f(\bar{i}_0(K,s)) & \text{for } s\in M\end{cases}.$$

Therefore the overall expected impact of K on intermediate goods production Y^{SP} is given by its conditional impact in states $s\in M$, multiplied by the probability of such states:

$$\begin{aligned}\frac{dE\left[Y^{SP}(K,s)\right]}{dK}=E\left[\frac{dY^{SP}(K,s)}{dK}\right]&=\Pr(M)E\left[\frac{dY^{SP}(K,s)}{dK}|s\in M\right]\\ &=2\Pr(M)E\left[\frac{d\bar{i}_0(K,s)}{dK}f(\bar{i}_0(K,s))Q|s\in M\right].\end{aligned}$$

The final equality above allows us to establish a link between $E\left[\frac{d\tilde{Y}^{SP}}{dK}\right]$ and $E\left[\frac{dY^{SP}}{dK}\right]$ as follows:

$$\begin{aligned}\frac{dE\left[\tilde{Y}^{SP}(K,s)\right]}{dK}&=2n\Pr(M)E\left[\frac{d\bar{i}_0(K,s)}{dK}Q|s\in M\right]\\ &=n\Pr(M)E\left[\frac{1}{f(\bar{i}_0(K,s))}\frac{dY^{SP}(K,s)}{dK}|s\in M\right].\end{aligned}$$

Finally, we substitute the expressions above for $\frac{dE\left[\tilde{Y}^{SP}(K,s)\right]}{dK}$ and $\frac{dE\left[Y^{SP}(K,s)\right]}{dK}$ into the first order condition for the constrained social planner to yield main result as required:

$$\begin{aligned}
0 &= v\frac{dE\left[\tilde{Y}^{SP}\right]}{dK}-nr-nE\left[\frac{w}{\alpha}\left(\frac{Y^{SP}}{K}\right)^{\frac{(1-\alpha)}{\alpha}}\frac{dY^{SP}}{dK}-w\frac{1-\alpha}{\alpha}\left(\frac{Y^{SP}}{K}\right)^{\frac{1}{\alpha}}\right] \\
&= n\Pr(M)E\left[\frac{v}{f(\bar{i}_0)}\frac{dY^{SP}}{dK}|s\in M\right]-nr+\dots \\
&\quad -n\Pr(M)E\left[\frac{w}{\alpha}\left(\frac{Y^{SP}}{K}\right)^{\frac{(1-\alpha)}{\alpha}}\frac{dY^{SP}}{dK}|s\in M\right]+nE\left[w\frac{1-\alpha}{\alpha}\left(\frac{Y^{SP}}{K}\right)^{\frac{1}{\alpha}}\right] \\
\Leftrightarrow 0 &= \Pr(M)E\left[\left(\underbrace{\frac{v}{f(\bar{i}_0)}-\frac{w}{\alpha}\left(\frac{Y^{SP}}{K}\right)^{\frac{(1-\alpha)}{\alpha}}}_{=0}\right)\frac{dY^{SP}}{dK}|s\in M\right]-r+E\left[w\frac{1-\alpha}{\alpha}\left(\frac{Y^{SP}}{K}\right)^{\frac{1}{\alpha}}\right] \\
\Leftrightarrow r &= E\left[w\frac{1-\alpha}{\alpha}\left(\frac{L^{SP}}{K}\right)^{\frac{1}{\alpha}}\right] \quad \text{(given Cobb-Douglas production fn.)} \\
\Leftrightarrow \alpha r K^{SP} &= (1-\alpha)E\left[wL^{SP}\right].
\end{aligned}$$

The equality $\frac{v}{f(\bar{i}_0)}-\frac{w}{\alpha}\left(\frac{Y^{SP}}{K}\right)^{\frac{(1-\alpha)}{\alpha}}=0$ comes from the definition of the threshold buyer $\bar{i}_0$ for $j=0$, for whom the marginal benefit from producing the final good $\frac{v}{f(\bar{i}_0)}$ is equal to the marginal cost $\frac{w}{\alpha}\left(\frac{Y^{SP}}{K}\right)^{\frac{(1-\alpha)}{\alpha}}$. This completes the proof for the result in the main body, for the case where the full production assumption is relaxed.

G Partial production and local monopolies in the decentralized solution

In this appendix, we discuss the implications of relaxing the full production assumption. Relaxing the assumption allows for shocks that are severe enough to shut out some market segments of final goods producers from the spot market. Final goods producers that are further away from intermediate goods suppliers (i.e., those with less substitutable inputs) will experience greater difficulty adjusting to the shocks.

To see this, note that the period 1 equilibrium is governed by two indifference thresholds (which may or may not be binding). First, for given price vector (p_0, p_{-0}), where $p_j = p_{-0}\ \forall j \neq 0$, the *participation threshold* $\bar{i}_0$ is defined as the final goods firm that is indifferent between buying inputs from intermediate goods firm $j=0$ and not producing at all:

$$f(\bar{i}_0)\,p_0 = v, \quad \forall p_0 \in \left[\frac{v}{f\left(\frac{1}{2}\right)}, v\right] \tag{G.1}$$

Second, the *competitive threshold* $\bar{i}_{0,1}$ is the marginal final goods producer that is indifferent between buying from supplier node $j=0$ and $j=1$:

$$f(d(\bar{i}_{0,1},0))\,p_0 = f(d(\bar{i}_{0,1},1))\,p_{-0} \tag{G.2}$$

As Figure G.1 illustrates, the participation threshold $\bar{i}_0$ (and its counterparts for $j \neq 0$) can be visualized as the arms that reaches out from each supplier node. The participation threshold therefore represents the potential *market reach* for each intermediate goods supplier. As long as the market reach from two nearby supplier nodes overlap, the two suppliers engage in competition and the competitive threshold $\bar{i}_{0,1}$ is the binding threshold for computing demand. Under this *competitive regime*, the intermediate goods suppliers' market reach covers every market segment on the circle. The aggregate demand for final goods is met and we see "full shelves". The competitive regime always prevails under the full production assumption.

We can show further that the market reach of each intermediate goods supplier is increasing in the level of non-scalable capacity installed (K), and decreasing in the cost of the scalable input (w). For given level of non-scalable capacity K, the market reach of each supplier node gets shorter as the size of the negative cost shock increases, until eventually the participation thresholds $\bar{i}_0$ and $\bar{i}_1$ no longer overlap and the two neighboring suppliers ($j=0,1$) behave like local monopolies. Under this *local monopolies regime,* there is a gap in market coverage between the two supplier nodes, and we see "empty shelves" in some segments of the market.[34]

The optimal pricing strategy of intermediate goods suppliers therefore depend on whether they are operating under the competitive or the local monopolies regime, which in turn depends on the realization of demand and supply shocks in period 1. We formally characterize the symmetric equilibrium spot-market pricing strategy under the assumption that the distance-based penalty function $f(x)$ takes the form of an exponential function, with parameter β.

Assumption *A*3 *Exponential distance-based penalty function*: $f(d) = \exp(\beta d)$, where $\beta \in (0,1]$ governs the degree of substitutability between different intermediate goods.

[34] A third possible regime arises when the market reach of one intermediate goods supplier goes past the node of another. This is the "*super-competitive*" regime, whereby one supplier prices so aggressively as to capture the home market of their neighboring competitor. Allowing for this possibility would lead to a discontinuous jump in the demand function for intermediate goods. In the interest of tractability, we can rule out the possibility of a super-competitive regime by making the distance-based penalty function $f(\cdot)$ sufficiently punishing.

Figure G.1: Regime switching: competition vs local monopolies

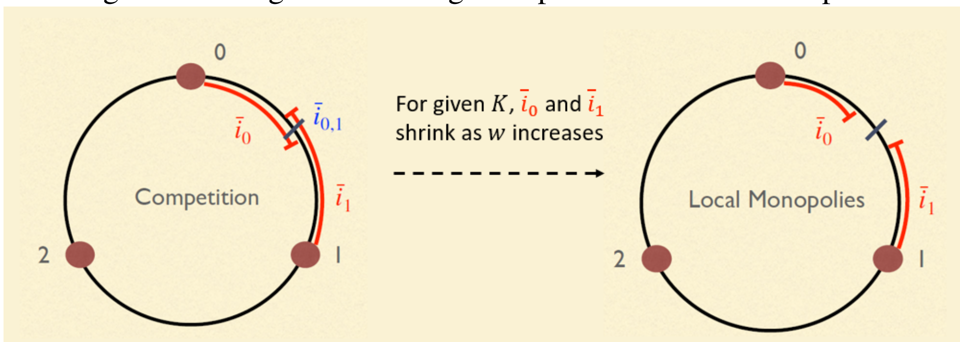

Lemma 7. *[Optimal spot-market pricing under symmetric equilibrium]*

1. *Under the competitive regime, we have* $\bar{i}_{0,1} = \frac{1}{2n} \leq \bar{i}_0$ *and the equilibrium price for the intermediate goods is given by:*

$$p_c^* = MC_c \cdot \frac{\exp\left(\frac{\beta}{2n}\right)}{2 - \exp\left(\frac{\beta}{2n}\right)} \tag{G.3}$$

where MC_c *is the marginal cost faced by intermediate goods suppliers*

$$MC_c = \frac{w}{\alpha}\left(\frac{Y^{pre} + Y_c^{spot}}{K}\right)^{\frac{1-\alpha}{\alpha}} \tag{G.4}$$

2. *Under the local monopolies regime, we have* $\bar{i}_0 < \bar{i}_{0,1} = \frac{1}{2n}$ *and the equilibrium price for the intermediate goods is given by:*

$$p_m^* = \sqrt{v \cdot MC_m} \tag{G.5}$$

where MC_m *is the marginal cost faced by intermediate goods suppliers*

$$MC_m = \frac{w}{\alpha}\left(\frac{Y^{pre} + Y_m^{spot}}{K}\right)^{\frac{1-\alpha}{\alpha}} \tag{G.6}$$

Intuitively, the first part of Proposition 7 shows that under a competitive regime, intermediate goods suppliers charge a mark-up over marginal costs.[35] The mark-up is higher when substitutability is lower (i.e., when β, the parameter governing the distance-based penalty function, is closer to 1), and lower when competition is fiercer (i.e., when n is large). In the limit, as n approaches infinity - and the distance between nodes shrinks to zero such that intermediate goods become perfectly substitutable - equation G.3 simplifies down to the familiar condition of price equals marginal cost.

[35] This part of the proposition is just a re-writing of our earlier results for this specific parameterization.

The second part of Proposition 7 shows that when intermediate goods suppliers operate as local monopolies, the price they charge is equal to the geometric average between their marginal costs (MC_m) and the highest possible price (v, the valuation of the final goods output by end consumers). Unsurprisingly, whilst intermediate goods suppliers operates as local monopolies, the number of other firms n is irrelevant to their pricing decision. Any changes in n instead influences whether the economy switches between the local monopolies regime and the competitive regime (i.e. whether $\bar{i}_0$ is less or greater than $\bar{i}_{0,1} = \frac{1}{2n}$).[36]

Other factors that influence the market pricing regime that prevails in equilibrium include the level of non-scalable production capacity in place K, and the cost of the scalable input w.

Proposition 3. *[Regime switching] $\bar{i}_0$, the participation threshold (i.e market reach) of firm $j = 0$, is increasing in K and decreasing in w:*

$$\frac{d\bar{i}_0}{dK} > 0 \tag{G.7}$$

$$\frac{d\bar{i}_0}{dw} < 0 \tag{G.8}$$

Proposition 3 formalizes our earlier discussion that, *for given non-scalable capacity K, larger negative supply shocks (larger w) increases the likelihood that the economy will end up in the local monopolies regime*. Under the local monopolies regime, the market segment ($i \in (\bar{i}_0, \bar{i}_1)$) that lies in-between the market-reach of the two nearby supplier nodes will not be able to fulfill their realized demand for final goods, and we observe "empty shelves". Intuitively, the proposition holds because a higher K, and a lower w, reduces the marginal cost of production, which increases the market reach of the intermediate goods supplier.[37]

A key implication of Proposition 3 is that the response of final goods outputs to shocks is *non-linear*. Under normal or benign market conditions, the economy might be operating under the competitive regime which ensures that demand from every market segment is met. Market reach of neighboring suppliers overlap, and continues to overlap for small perturbations in supply and demand. Under these benign conditions, the supply network appears robust. But when negative supply shocks becomes sufficiently large, the economy suddenly switches from the competitive regime to the local monopolies regime. The critical role capacity plays, therefore, is that it prevents empty shelves for a larger range of shocks. A larger K allows for a larger market-reach overlap for

[36] Clearly, this neat characterization of the monopoly price as a geometric average won't hold in general (e.g. without the exponential functional form for $f(d)$). But the other part of the proposition, that in the local monopolies regime the number of other firms is irrelevant, is more general. Even if other firms exist, they simply aren't selling in each other's "submarket".

[37] Note that this analysis is not comparative statics in the strict sense: w is an exogenous variable, but K is an endogenous variable. With regard to the latter, we are asking how firms' endogenous choice of capacity investment in period 0 affects market reach and the nature of competition on the spot market in period 1.

any given input cost w, making the entire network more robust. But since the degree of overlap is in of itself irrelevant, surplus capacity is "wasted" in the absence of large negative supply shocks.

Relaxing the full production assumption therefore reinforces our central message that $K^* < K^{SP}$. This is intuitive, because the possibility of a large shock shifting the economy to a local monopolies regime adds another distortion to the system. Ex ante capacity investment K increases network resilience by ensuring full production for a wider range of shocks, but is undervalued by market participants under business-as-usual scenarios. Robustness becomes an externality that may not be fully internalized by individual intermediate goods suppliers in their capacity decisions in period 0. Worse still, in imperfectly competitive economies, some firms may profit from the artificial scarcity that arises from a lack of resilience.