Specialization, Complexity & Resilience in Supply Chains*

Alessandro Ferrari — UPF, BSE & CEPR
Lorenzo Pesaresi — University of Zurich

March 19, 2025

Preliminary and Incomplete
Please do not cite without author permission

Abstract

Despite growing policy interest, the determinants of supply chain resilience are still not well understood. We propose a new theory of supply chain formation with compatibility frictions: only compatible inputs can be used in the final good production process. Intermediate producers choose the degree of specialization of their goods, trading off higher productivity with a lower share of compatible final producers. We model supply chains as a complex production process, in which multiple complementary inputs need to be sourced for final production to take place. Specialization choices, production complexity, and search frictions jointly determine the resilience of supply chains. Relative to the efficient allocation, the equilibrium is characterized by over-specialization due to a novel network externality arising from the interplay between frictional markets, endogenous specialization, and complex production. Intermediate producers fail to internalize how their specialization choices affect the likelihood that the final good is produced and, as a consequence, the foregone surplus from complementary inputs. In turn, over-specialization implies that supply chains are less resilient than efficient. Rising complexity exacerbates this inefficiency. We characterize how a social planner can decentralize efficient supply chain resilience through a targeted transaction subsidy.

Keywords: Supply chains, specialization, product design, resilience
JEL Codes: D21, L14, L22, L23

*We thank Pol Antràs, Vasco Carvalho, Matteo Escudé, Mathieu Parenti, and Alireza Tahbaz-Salehi for many discussions that have been instrumental to the inception of this paper. We also thank Gene Grossman, David Hémous, Fabrizio Leone, Marek Pycia, Isabelle Mejean, Alejandro Sabal, and Armin Schmutzler for their feedback. This paper also benefited from seminars at the University of Zurich and presentations at the "Workshop on International Economic Networks", "Trade, Value Chains and Financial Linkages in the Global Economy" conference at the Bank of Italy and the CEPR - Naples Trade & Development Workshop. Ferrari gratefully acknowledges financial support from the Swiss National Science Foundation (grant number 100018-215543).

1 Introduction

In modern economies, production is organized around long and complex supply chains. The substantial progress in information and transportation technologies allows firms to source specialized inputs from a large number of suppliers. The rise of supply chains has brought about significant productivity gains, as well as concerns about the frequency and economic costs of their disruptions. In recent years, policy interest has been focused on whether the organization of supply chains is efficient or if, instead, governments should intervene.[1] In particular, gains from specialization can bring about risks associated with bottlenecks or over-reliance on single providers of key inputs, which might not be socially desirable.

In this paper, we propose a new theory of frictional supply chain formation where intermediate producers choose the degree of specialization of their goods. The economy features an intuitive technological friction: inputs need to be *compatible* with the production process they are used in. Higher specialization of an input increases its value but reduces the share of compatible final producers. We model supply chains as a *complex* production process in which multiple complementary inputs need to be sourced for final production to take place. Supply chain *resilience*, that is, the time it takes for firms to return to normal operations following disruptions, is fundamentally shaped by the endogenous specialization of goods, search frictions, and production complexity. Our framework allows us to analyze the sources of inefficiency affecting supply chain resilience and address normative questions related to the optimal organization of supply chains. The intuitive trade-off that characterizes our economy is that higher specialization increases the value of production in normal times, but results in a higher cost of disruptions, as replacing specialized inputs takes longer (Miroudot, 2020).

We start by characterizing the equilibrium in a static economy where heterogeneous intermediate producers choose the degree of specialization of their goods. Ex-ante identical final producers need to source multiple complementary inputs in the market. Higher specialization increases the value of the intermediate input, but reduces the share of compatible final producers. The market for intermediate goods displays realistic sourcing frictions: each final producer contacts a finite number of intermediate producers, governed by search frictions. Within the set of contacted intermediate producers, it selects those compatible with its production process – if any – and requests a quote from each of them.[2] Intermediate producers submit a bid characterized by the surplus (profits) offered to the final producer. The final producer trades with the intermediate firm offering the highest surplus. Whenever search is unsuccessful for any of the complementary inputs, the entire production process is halted.

In this static framework, we highlight the key inefficiency arising from endogenous specialization when production is complex. Intermediate producers choose their optimal specialization

[1] In June 2021, the Biden-Harris Administration instituted the *Supply Chain Disruption Task Force* "to provide a whole-of-government response to address near-term supply chain challenges to the economic recovery" (see White House (2021)).

[2] According to a McKinsey report, "*At most organizations [. . .], it takes about three months to complete a single supplier search, with a sourcing professional logging more than 40 hours of work – and yet able to consider only a few dozen suppliers from a total population of thousands.*" (Baptista et al., 2021). We interpret this evidence as indicative of search and compatibility frictions in intermediate markets.

by trading off the private surplus they obtain upon trading against the probability of trading. Crucially, they fail to internalize how their specialization choices affect the likelihood that the final good is produced and, as a consequence, the foregone surplus from complementary inputs. For example, consider electric vehicle (EV) production. EV manufacturers produce multiple car models, each requiring a specific type of battery – such as lithium, solid-state, or nickel. Despite their critical role, these batteries are typically sourced externally. Likewise, each car model (e.g., city car, SUV) requires other externally sourced components, such as tires, cameras, and carbon fiber. Due to compatibility frictions, an EV producer may, in a given period, only locate battery suppliers offering an incompatible type. In such cases, production halts completely for that period. This mismatch stems from suppliers' product specialization choices – nickel batteries are of lower value but widely used across EV models, while lithium batteries are limited to specialized, high-value applications. Our theory predicts that battery manufacturers, when determining the level of specialization of their product, do not consider the lost value from complementary inputs caused by the inability of car manufacturers to source compatible batteries. Indeed, if EV production halts due to compatibility frictions, providers of complementary inputs, such as tires and cameras, are also unable to generate value. This example highlights the presence of a *network externality*. We show that this externality induces equilibrium *over-specialization* relative to the constrained efficient allocation. Intuitively, specialization choices fundamentally interact with the degree of complexity in the production process: if EVs were only made of batteries – that is, the production process were not complex – the *network externality* would be absent, and the equilibrium would be constrained efficient.

The *network externality* arises due to the combination of three elements in our model: i) search and compatibility frictions, which prevent final good producers from finding necessary inputs with probability one; ii) endogenous specialization decisions, which imply that the equilibrium level of compatibility is optimally chosen by agents which may or may not make socially optimal decisions; and iii) complementarity among inputs, which imply that failure to source all inputs can halt the production process. To the best of our knowledge, our paper is the first to develop a model featuring all these three elements.

Next, we study a dynamic version of our model to account for the long-term nature of procurement relationships. Once formed, procurement relationships last until the final producer – or a supplier link – faces a disruption shock, which happens stochastically. Long-term relationships bring about a further motive for specialization, that is, tailoring the product to the needs of existing customers (*customization*). On top of the network externality, endogenous specialization choices in a dynamic context introduce a *search externality* (Pissarides, 2000): intermediate producers do not internalize the effect of their specialization on the meeting probability, as more specialized goods raise the share of searching final producers. This external effect introduces a force pushing the equilibrium towards under-specialization. However, we show that, if the production process is complex, the equilibrium always displays over-specialization. Importantly, the dynamic framework allows for a model-consistent and welfare-relevant notion of supply chain resilience, that is, the average time it takes for a final producer to restore production

following a disruption. This allows us to characterize the determinants of supply chain resilience: the higher the extent of search frictions, average specialization, and complexity, the lower supply chain resilience. It follows that equilibrium over-specialization implies that supply chains are *less resilient than efficient*.

We show that, in our model, resilience and robustness of supply chains, along with the expected productivity conditional on matching, determine the value of output and welfare: since higher specialization increases the value but decreases the number of goods that are produced, a trade-off arises between *productivity* and *resilience*. More specialized economies deliver higher output in normal times, when disruptions do not occur. However, when disruptions are present, higher equilibrium specialization makes it harder to return to production in a timely fashion. A social planner, instead, by internalizing the network externality, trades part of the normal times surplus for a more resilient supply chain.

Through the lens of our model, we study three salient trends occurring in supply chains. First, supply chains have gotten more complex and longer. In our model, whenever search frictions in procurement are sufficiently strong, higher complexity induces rising inefficiencies in the form of over-specialization and under-resilience. Second, as the frequency of supply chain disruptions increases, potentially driven by geopolitical tensions, pandemics, or extreme weather events, the market equilibrium features rising over-specialization. Finally, if new technologies are introduced such that search frictions in procurement are attenuated, it is generally ambiguous whether efficiency improves. This is because firms could, in principle, respond to lowering search frictions with disproportionately higher specialization investments, thereby making the economy less resilient. We show that our conclusions remain if we allow firms to hold inventories, source the same good from multiple suppliers, produce part of the inputs in-house, and write contracts in which suppliers' remuneration is not tied to production.

We conclude by studying the normative implications of our framework. We start by showing that, if a set of firm-specific transaction subsidies is available, a planner can decentralize the constrained efficient allocation in both our static and dynamic economies. We show that the planner would like to increase the price of all transactions. This is equivalent to giving more surplus to intermediate producers for them to have *more skin in the (final production) game*, which allows the decentralization of the efficient level of specialization.[3] We also show that the optimal subsidy schedule takes a simple form comprising an aggregate component and a targeting element that depends on the average trading probability of lower-productivity intermediate competitors.

Related Literature This paper is related to several strands of literature. First, our interest in supply chains speaks to the growing literature on production network formation and its role in the economy's shock-absorption capacity. Levine (2012), Baqaee (2018), Elliott et al.

[3]While this result may sound similar to the Hosios (1990) condition, we show that the logic underlying it is different. In Hosios (1990) efficiency comes from the matching elasticity being equal to the surplus share of sellers. In our model, efficiency is obtained when the surplus share of sellers is equal to their surplus elasticity (see Mangin and Julien, 2021).

(2022), Acemoglu and Tahbaz-Salehi (2024), Carvalho et al. (2023) and Capponi et al. (2024) study the *fragility* of economies in which supply chains feature coordination problems due to complementarities, public goods problems, or market incompleteness. Carvalho et al. (2021), Meier (2020), Alessandria et al. (2023), Ferrari (2023), Carreras-Valle (2023), and Carreras-Valle and Ferrari (2025) study the dynamic response of supply chains to disruptions. Arkolakis (2010), Lenoir et al. (2022), Martin et al. (2023), Esala (2024), Huang et al. (2022) and Haasbroek (2025) study the implications of search and matching frictions in firm-to-firm trade. Closest to our work, Grossman et al. (2023b) explores market failures in vertical supply chains where firms can invest to reduce their risk of disruptions and the terms of trade in firm-to-firm transactions are bargained over. We contribute to this literature by studying the problem of endogenous specialization in frictional markets and by proposing a dynamic framework that allows us to separately identify the fragility of production (*robustness*) and its ability to recover from shocks in a timely manner (*resilience*).

Our work is also closely related to the literature on the efficient number of varieties, whose foundations date back to Spence (1976), Dixit and Stiglitz (1977) and Mankiw and Whinston (1986). Recent advances of these traditional models include Zhelobodko et al. (2012), Parenti et al. (2017), Dhingra and Morrow (2019), and Bajo et al. (2024). Grossman et al. (2023a) proposes an application to supply chains. We show that many of the insights from this literature have tight analogs in our search model with endogenous specialization choices and price posting, chief among them the offsetting effects of *appropriability* and *business-stealing* externalities.

Both methodologically and conceptually, our work is close to the strands of literature on product design (Bar-Isaac et al., 2012, 2023; Menzio, 2023; Albrecht et al., 2023), supply chain formation (Oberfield, 2018; Antràs and De Gortari, 2020; Boehm and Oberfield, 2020; Kopytov et al., 2021; Kim and Shin, 2023; Antràs, 2023) and specialization (Rauch, 1999; Nunn, 2007; Barrot and Sauvagnat, 2016; Martin et al., 2020; Ekerdt and Wu, 2023). Relative to these contributions, we study the effects of specialization choices on the resilience of supply chains. Our key novelty is the study of product design choices when goods are complements in production, which allows us to highlight the key inefficiencies associated with supply chain formation.

Finally, our model is also close in spirit to the large literature on the hold-up problem in firm-to-firm transactions (see Klein et al., 1978; Williamson, 1979; Rogerson, 1992; Hart and Moore, 1990; Tirole, 1999) and the organization of production (Antràs, 2003; Antràs and Helpman, 2006; Antràs and Chor, 2013). We study a context in which firms make sunk investments in specialization and cannot appropriate the full returns from their investment, yet our price-posting mechanism still generates bilaterally efficient outcomes in frictional markets with simultaneous search.

2 Static Model

In this section, we consider the procurement problem of a final producer that needs to source its production inputs anew on the market. Focusing on the sourcing problem of a final producer

allows us to single out the specific stage at which the specialization of intermediate inputs plays a role in driving the supply chain dynamics. As highlighted in the business literature, the more specialized intermediate inputs are, the more difficult it is for final producers to source a compatible input. We are interested in understanding how intermediate producers choose the specialization of their inputs and whether such specialization decisions are efficient. To this end, we propose a static model of sourcing from frictional markets when intermediate producers make endogenous product design decisions.

2.1 Environment

The economy is populated by a representative household, a measure 1 of perfectly competitive final producers and a measure m of intermediate producers. Final producers are ex-ante identical, whereas intermediate producers are heterogeneous in their productivity z. The market for the consumption good is perfectly competitive, whereas the markets for intermediate goods are frictional. The consumption good is the numeraire. The model is static.

Representative household. The representative household has linear preferences over an aggregate consumption good and convex disutility from labor supply. They choose whether or not to purchase the indivisible consumption goods offered by each final producer i, with associated quality $\mathcal{Q}_i$, and how many hours to work at each wage rate w, by solving the following utility maximization problem:

$$\max_{C_i \in \{0,1\}, \ell} \mathcal{U} = C + \psi \log(1 - \ell) \tag{1}$$

$$\text{s.t.} \quad \int_0^1 P_i C_i \, di = w\ell + \overline{\Pi},$$

$$C = \int_0^1 \mathcal{Q}_i C_i \, di,$$

where $\overline{\Pi}$ represents profit rebates.

Final producers. Each final producer i operates a production function combining $N > 1$ complementary and indivisible inputs into one unit of consumption good:

$$Y_i = \mathbb{1}\{\min\{y_1, \ldots, y_N\} = 1\}. \tag{2}$$

Following Elliott et al. (2022), we refer to the number of inputs sourced in the market, N, as the *complexity* of the production process.[4] We think of the N inputs as defining a recipe for final production in the spirit of Oberfield (2018). Output quality $\mathcal{Q}_i$ depends on the value of

[4]As will become apparent, the network structure of our model lends itself both to a "snake" and a "spider" interpretation (Baldwin and Venables, 2013). In the main text, we interpret the production network as spider-like. Alternatively, one can interpret the N key inputs as the number of vertical stages leading to final production.

each input j the final producer is able to source in an additively separable fashion:

$$\mathcal{Q}_i = \sum_{j=1}^{N} A_j. \tag{3}$$

The input value A_j measures the value-added created by input j in the production of the final good. In turn, the higher the value of the inputs sourced, the higher the quality of the final good. Notice that higher complexity of the production process increases the quality of the final good. This captures in reduced-form the gains from organizing production around supply chains, resulting from complementarities across different inputs.

For example, suppose a car maker needs to source a battery and a tire to assemble an electric SUV. Yet, not all batteries and tires it can find in the market fit its production process in the same way. A lithium-ion battery would be a better fit than a lead-acid battery, while a alkaline battery battery would not work at all. Similarly, a large tire would fit the production process better than a small tire. The quality function (3) assumes that the *values* of different inputs are perfect substitutes – being the inputs themselves perfect complements. Keeping with our example, a producer could make a $\$50,000$-worth SUV with a lithium-ion battery and a small tire, or with a lead-acid battery and a large tire. We relax this assumption in Appendix B, where we allow for arbitrary substitutability or complementarity in the quality function.

Each final producer i obtains profits:

$$\pi_i = \left(P_i - \sum_{j=1}^{N} p_j \right) \mathbb{1}\{Y_i = 1\}. \tag{4}$$

Each final producer i makes profits equal to the difference between the price of its consumption good net of the price it pays for the N intermediate inputs it sources. In equilibrium, the price of the consumption good reflects its quality, i.e., $P_i = \mathcal{Q}_i$. If the final producer is unable to source all the N complementary inputs, production is halted, as implied by (2).[5]

Intermediate market structure. Intermediate good markets are characterized by search and compatibility frictions. Search frictions are represented by a finite number $n \sim \text{Poisson}(\lambda)$ of intermediate producers that each final producer is able to meet, where $n \in \mathbb{N}$ and $\mathbb{E}[n] = \lambda$. Hence, λ represents the *average number of meetings* per final producer. Compatibility frictions are represented by an endogenous distribution of compatibility probability $\mathcal{H}(\phi)$, with mean $\bar{\phi}$. Hence, $\bar{\phi}$ represents the *average probability that an input provider is compatible* with the final production process. Assuming that the variance of ϕ is small relative to the mean, the number of compatible intermediate producers met by a final producer is Poisson-distributed with mean $\lambda\bar{\phi}$. Therefore, $\lambda\bar{\phi}$ represents the *average number of compatible input providers* per

[5] Our theory assumes that the N inputs cannot be produced in-house nor sourced on the spot market. We think this assumption is in line with the operation of many supply chains. According to a McKinsey report, "*At the height of the (Covid) crisis, [...] companies simply couldn't get the critical parts they needed. In 2021 and 2022, semiconductor shortages prevented the production of an estimated 13 million vehicles worldwide.*" Drentin et al. (2024).

final producer. It follows that the probability that a final producer finds a compatible input, or *input finding probability*, is given by:

$$f = 1 - \exp\{-\lambda\bar{\phi}\}. \tag{5}$$

Once the compatible intermediate producers – if there are any – are selected, final producers receive a quote. Intermediate producers submit it as a bid characterized by a surplus x offered to the final producer. This represents the amount of surplus the final good producer extracts from the transaction. Let the endogenous distribution of offered surplus be $G(x)$. The final producer chooses the intermediate producer offering the highest surplus, as in a first-price auction.

Intermediate producers. Intermediate producers differ in productivity $z \sim \Gamma(z)$, where Γ is a continuous distribution with support $[\underline{z}, \bar{z}]$. Productivity z is associated with marginal cost $c(z)$, where $c'(z) < 0$. Intermediate producers make endogenous product design decisions by choosing the specialization s of their goods, as well as a surplus x offered to final producers. On the one hand, higher specialization increases the value of the input for compatible final producers. Formally, $A = A(s; z),\ A'(s) > 0,\ A''(s) < 0$. On the other hand, higher specialization reduces the share of compatible final producers. Formally, $\phi = \phi(s)$, $\phi'(s) < 0$, $\phi''(s) < 0$. Hence, optimal specialization trades off lower matching probability against higher profits upon matching.

Each intermediate good producer can produce all N inputs through distinct product lines.[6] For each intermediate good, denote market tightness as $\theta = \frac{1}{m}$. An intermediate producer meets an expected number m_k of final producers who are in contact with $k = 0, 1, 2, \ldots$ other *compatible* intermediate producers, where:

$$m_k = \theta \sum_{n=k+1}^{\infty} n \underbrace{\frac{\lambda^n e^{-\lambda}}{n!}}_{\text{Prob. buyer meets n sellers}} \underbrace{\int_0^1 \frac{(n-1)!}{k!(n-k-1)!}(\phi)^k(1-\phi)^{n-k-1}d\mathcal{H}(\phi)}_{\text{Prob. k other sellers are compatible}} = \theta\lambda\, e^{-\lambda\bar{\phi}}\frac{(\lambda\bar{\phi})^k}{k!}.$$

This statistic summarizes the expected number of final customers *binned* by how many other compatible intermediate competitors they are in contact with.[7] For example, an intermediate producer of good j may expect to meet 3 final good producers who each are in contact with 4 other compatible intermediate producers of good j (on top of itself): this amounts to $m_4 = 3$.

The expected operating profits of an intermediate producer with productivity z choosing

[6] Since intermediate good producers make positive profits in equilibrium in each product line, they optimally choose to compete in all N products.

[7] Notice that the final equality follows from the assumption that the variance of the compatibility probability is small relative to the mean, which allows us to apply the Poisson approximation condition (Menzio, 2023).

specialization s and offering surplus x to final producers equal:

$$\begin{aligned}\Pi(s,x;z,N) &= \left(\sum_{k=0}^{\infty} m_k\ \phi(s)\ G(x)^k\right) f^{N-1}\left[a(s)-c(z)-x\right] \\ &= \underbrace{\theta\lambda}_{\substack{\text{exp \#}\\ \text{meetings}}}\ \underbrace{\mathcal{P}(s,x;N)}_{\substack{\text{trading prob.}\\ \text{meeting}}}\ \underbrace{\left[a(s)-c(z)-x\right]}_{\substack{\text{unit profit}\\ \equiv\ p(s,x)-c(z)}}.\end{aligned}$$

Expected operating profits are composed of three elements. First, every intermediate producer meets an expected number $\theta\lambda$ of final producers, as governed by the extent of search frictions. Second, an intermediate producer choosing specialization s and offering surplus x trades with the final producers it meets with probability $\mathcal{P}(s,x;N)$. Finally, an intermediate producer with productivity z and specialization s offering surplus x makes a profit of $p(s,x)-c(z)$ on each completed transaction, where $p(s,x)\equiv a(s)-x$. The input value, which determines the quality of the final good, equals the total surplus created by the match between a final producer and an intermediate producer with productivity z and specialization s, that is, $A(s;z)\equiv a(s)-c(z)$. In turn, x represents the surplus accruing to the final producer.

The endogenous compatibility frictions described in the last section are subsumed into the conditional trading probability $\mathcal{P}(s,x;N)$, which is itself a composite of three elements:

$$\mathcal{P}(s,x;N)\equiv \underbrace{\phi(s)}_{\text{prob. compatible}}\ \underbrace{\exp\left\{-\lambda\bar{\phi}\left[1-G(x)\right]\right\}}_{\substack{\text{prob. best among compatible}\\ \text{contacted suppliers}}}\ \underbrace{f^{N-1}}_{\substack{\text{prob. other complementary}\\ \text{inputs are sourced}}}.$$

Suppose an intermediate producer with specialization s, offering surplus x meets a final good producer. First, its output is compatible with the production process with probability $\phi(s)$. Second, since the number of compatible intermediate producers per final producer is Poisson-distributed, the probability that the offered surplus x is the highest among the set of compatible intermediate producers contacted by a final producer is given by $\exp\{-\lambda\bar{\phi}(1-G(x))\}$.[8] Finally, the last term states that trade only occurs if all other $N-1$ inputs are also successfully sourced. This happens with probability f^{N-1}. Hence, the trading probability of an intermediate producer of good j depends on the input finding probability of all its complementary inputs.

Let $q(s)$, $q'(s)>0$ denote the labor requirement for specialization. The intermediate producer's problem reads:

$$\max_{s,x} V(s,x;z,N)=\Pi(s,x;z,N)-wq(s). \tag{6}$$

The intermediate good producer chooses specialization and surplus to maximize the value of the firm, given by the expected operating profits minus the cost of specialization.

[8]This matching process is similar to the one proposed by Huang et al. (2022). The key difference is that, in our setting, intermediate producers can influence the share of potential buyers via endogenous product design decisions.

2.2 Equilibrium characterization

We now define our equilibrium concept and characterize the equilibrium allocation.

Definition 1 (Static equilibrium)**.** *An equilibrium consists of:*

- *Purchasing decisions, $C_i \in \{0,1\}$, for each available final good i and labor supply decisions, $\ell = \ell(w)$, which solve the utility maximization problem of the representative household (1), taking the wage rate w as given;*
- *A selection rule for the input provider of each input $j = 1, \ldots, N$ which maximizes profits (4) of the final producer, taking the surplus offered x by each of compatible intermediate producers contacted as given;*
- *Policy functions of specialization and surplus offered to final producers, $s(z)$ and $x(z)$ $\forall z$, which solve the profit maximization problem (6) of intermediate producers, taking the wage rate w and the distribution of surplus offered to final producers $G(x)$ as given;*
- *Distribution of surplus offered to final producers $G(x)$ which is consistent with the policy functions $s(z)$ and $x(z)$;*
- *Labor market clearing pinning down the wage rate such that labor supply equals labor requirement in specialization: $\ell(w) = Nm\bar{q}$, where $\bar{q} = \int q(s(z))d\Gamma(z)$;*
- *Aggregate output Y such that the aggregate resource constraint holds: $C = Y$.*

We now characterize the decision of each agent in our economy.

Representative household. The representative household optimally purchases the indivisible consumption goods offered by final producer i provided that its quality is positive, that is, $C_i = \mathbb{1}\{\mathcal{Q}_i > 0\}$. The optimal labor supply function reads $\ell(w) = 1 - \frac{\psi}{w}$.

Final producers. Final producers find it optimal to select as input provider the compatible intermediate producer contacted – if there are any – that offers the highest surplus x for each input $j = 1, \ldots, N$.

Intermediate producers. To solve the intermediate producer's problem (6), we proceed by guessing and verifying that the optimal surplus offered to final producers is increasing in productivity, i.e., $x'(z) > 0$. The policy functions $s^\star(z)$ and $x^\star(z)$ solve the following system of differential equations:

$$x^{\star\prime}(z) = \lambda\phi(s^\star(z))\gamma(z)\left[a(s^\star(z)) - c(z) - x^\star(z)\right], \tag{7}$$

$$\theta\lambda\mathcal{P}(z;N)\left[a'(s^\star(z)) + \frac{\phi'(s^\star(z))}{\phi(s^\star(z))}\left(a(s^\star(z)) - c(z) - x^\star(z)\right)\right] = wq'(s^\star(z)), \tag{8}$$

with boundary condition $x(\underline{z}) = 0$. We assume $q(s)$ is such that equation (8) is necessary and sufficient for characterizing optimal specialization.[9]

Solving the differential equation (7) allows us to characterize the optimal surplus offered to final producers.

Lemma 1 (Offered Surplus)**.** *The optimal surplus offered by an intermediate producer with productivity z to final producers equals the expected outside option of final producers when trading with the intermediate producer:*

$$x^\star(z) = \mathbb{E}_{\max\{\tilde{z}\}|\tilde{z}\leq z}[a(s^\star(\tilde{z})) - c(\tilde{z})], \tag{9}$$

where the expectation is taken with respect to the joint distribution of productivity and the number of compatible intermediate producers contacted by a final producer.

The intuition behind Lemma (1) is instructive. First, we note that an intermediate producer with productivity z loses any auction against the other compatible competitors with higher productivity $z' > z$ contacted by a final producer. Hence, an intermediate producer enjoys surplus-(or price-)setting power only in the states of the world when it has the highest productivity among the set of compatible intermediate producers contacted by a final producer. Since an intermediate producer has the power to set the surplus offered to a final producer unilaterally, the optimal choice is to make the final producer exactly indifferent between accepting and rejecting the surplus offer.[10] The final producer's outside option in case of no trade with an intermediate producer of productivity z is the expected total surplus offered by the best compatible competitor with productivity $\tilde{z} < z$. Accordingly, intermediate producers with the lowest productivity $\underline{z}$ find it optimal to offer no surplus to final producers, i.e., $x(\underline{z}) = 0$. Intuitively, those intermediate producers trade with a final producer only if the latter is not in contact with any other compatible competitor. On the other hand, intermediate producers with the highest productivity $\bar{z}$ offer the expected total surplus over the unconditional productivity distribution, i.e., $x^\star(\bar{z}) = \mathbb{E}_{\max\{\tilde{z}\}}[a(s(\tilde{z})) - c(\tilde{z})]$. Intermediate producers with productivity strictly in between offer a convex combination of these two bounds, i.e., $x^\star(z) \in (0, \mathbb{E}_{\max\{\tilde{z}\}}[a(s^\star(\tilde{z})) - c(\tilde{z})])$ $z \in (\underline{z}, \bar{z})$. This implies $x^{\star\prime}(z) > 0$, which verifies our guess.

The optimal specialization function $s^\star(z)$ is implicitly defined by Equation (8). We can single out three forces shaping optimal specialization. First, higher specialization increases the total surplus conditional on trading, by raising the input value ($a'(s) > 0$). Second, higher specialization reduces the trading probability, by lowering the compatibility probability ($\phi'(s) < 0$). Third, higher specialization increases specialization costs, by raising the labor requirement for specialization ($q'(s) > 0$). Since higher complexity reduces the trading probability $\mathcal{P}(z; N)$, optimal specialization is decreasing in complexity N. If $\lambda\bar{\phi} < 1$, higher search efficiency increases

[9] This amounts to making sure that the SOC always holds. The SOC reads $wq''(s(z)) > \theta\lambda \exp\{-\lambda\hat{\phi}(z, \bar{z})\} f^{N-1} [a''(s(z))\phi(s(z)) + \phi''(s(z))(a(s(z)) - c(z) - x(z)) \; -2(-\phi'(s(z)))a'(s(z))]$.

[10] This is the core idea behind the well-known *Diamond Paradox* (Diamond, 1971). Still, in our model, final producers are able to extract positive surplus in equilibrium due to their ability to locate multiple intermediate producers simultaneously (Burdett and Judd, 1983).

the average number of meetings per intermediate producer leading to trade, $\theta\lambda\mathcal{P}(z;N)$. Hence, optimal specialization is increasing in search efficiency λ. Finally, notice that the value function of an intermediate producer, $V(s,x;z,N)$, is supermodular in $s^\star(z)$ and $x^\star(z)$, i.e., the cross derivative at the optimal solution is positive. Hence, optimal specialization and optimal surplus offered to final producers are strategic complements. In other words, if an intermediate producer were forced to offer a higher surplus to final producers ($x(z)\uparrow$), it would find it optimal to specialize more ($s(z)\uparrow$), and viceversa.[11] Intuitively, offering a higher surplus $x(z)$ reduces the cost of foregone trade due to lower compatibility probability.

Market clearing and consistency conditions. Since $x^{\star\prime}(z)>0$, the distribution of surplus offered to final producers equals $G(x^\star(z))=\hat{\phi}(\underline{z},z)/\bar{\phi}$, where $\hat{\phi}(\underline{z},z)=\int_{\underline{z}}^{z}\phi(s^\star(\tilde{z}))d\Gamma(\tilde{z})$ and $\bar{\phi}=\hat{\phi}(\underline{z},\bar{z})$. Intuitively, the probability that a final producer meets a compatible intermediate producer with productivity lower than z equals the relative mass of intermediate producers with productivity lower than z, weighted by the respective compatibility probability.

The equilibrium wage rate equalizing labor supply and labor requirement for specialization is given by $w=\frac{\psi}{1-Nm\bar{q}}$. The aggregate resource constraint requires that aggregate output equals:

$$Y=\underbrace{f^N}_{\text{prob. active}}\quad\underbrace{N\hat{\mathbb{E}}[a(s^\star(z))-c(z)]}_{\text{expected match surplus}\,\big|\,\text{active}}, \tag{10}$$

where $\hat{\mathbb{E}}[.]$ denotes the expectation with respect to the productivity distribution of active matches. The expression for aggregate output in (10) highlights a number of important features. First, the complexity of the production process has ambiguous effects on output, as in Levine (2012). On the one hand, more complex production processes allow generating output of higher quality per active match. On the other hand, the probability that all the complementary inputs are sourced decreases. A similar argument goes through for the endogenous specialization decisions. If firms specialize more, the expected surplus per active match increases. Yet, the probability that a final producer is active decreases, thus reducing the measure of active final producers. Previewing our findings from the dynamic model in Section 3, the probability that a final producer is able to source all the complementary inputs it needs f^N is inherently linked to the notion of *resilience*. Intuitively, if a final producer were to lose contact with its suppliers, f^N would represent the probability that the final producer is able to source the inputs it needs anew and restart production, and $1/f^n$ is the expected time required to do so.

2.3 Efficiency

We study the efficiency properties of equilibrium specialization decisions by comparing the equilibrium allocation with that dictated by a benevolent social planner. The social planner is subject to the same product market frictions as market participants, so our normative criterion

[11] This result stands in contrast with standard models in which an investment is made before the parties bargain over the surplus split (*hold-up problem*). We clarify the relation with these class of models in Section 2.4.

is constrained efficiency. The social planner solves:[12]

$$\max_{s_j(z)} \mathcal{W} = C + \psi \log(1-\ell) \qquad \text{s.t.} \quad \ell = m \sum_{n=1}^{N} \int q(s_n(z)) d\Gamma(z),$$
$$C = \sum_{n=1}^{N} \prod_{v \neq n} f(\boldsymbol{s}_v(z)) \; \mathbb{E}_{\max\{z\}}[a(s_n(z)) - c(z)].$$

In words, the social planner chooses the specialization $s_j(z)$ of each intermediate producer with productivity z producing input j to maximize the utility of the representative household. The maximization problem is subject to the labor resource constraint and the output resource constraint. Unlike equilibrium specialization decisions, the social planner internalizes that higher specialization of an intermediate input reduces the expected surplus from other complementary inputs, as final producers are less likely to be active.

Since all input markets are symmetric, the first-order condition for the efficient specialization function $\mathcal{S}(z)$ is given by:

$$\theta\lambda\mathcal{P}(z;N)\Bigg[a'(\mathcal{S}(z)) + \frac{\phi'(\mathcal{S}(z))}{\phi(\mathcal{S}(z))}\Bigg([a(\mathcal{S}(z)) - c(z)] - \mathbb{E}_{\max\{\tilde{z}\}|\tilde{z}\leq z}[a(\mathcal{S}(\tilde{z})) - c(\tilde{z})] \tag{11}$$
$$+ (N-1)e^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[a(s(\tilde{z})) - c(\tilde{z})]\Bigg)\Bigg] = \frac{\psi}{1 - Nm\bar{q}}\, q'(\mathcal{S}(z)).$$

The social marginal benefit of specialization on the left-hand side of (11) consists of four elements. The first element, $\theta\lambda\mathcal{P}(z;N)a'(\mathcal{S}(z)) > 0$, represents the increase in own surplus conditional on matching by marginally increasing specialization, for given trading probability. The second element, $\theta\lambda\mathcal{P}(z;N)\phi'(\mathcal{S}(z))/\phi(\mathcal{S}(z))\ [a(\mathcal{S}(z)) - c(z)] < 0$, reflects the reduction in own expected surplus from lower trading probability by marginally increasing specialization, for given surplus conditional on matching and weighted productivity distribution, *i.e.*, $\hat{\phi}(z,\bar{z})$. These are the direct effects of specialization by an intermediate producer with productivity z on the expected surplus from trading with a final producer. As such, the same effects appear in the condition determining the privately optimal specialization (8).

The third element, $-\theta\lambda\mathcal{P}(z;N)\phi'(\mathcal{S}(z))/\phi(\mathcal{S}(z))\mathbb{E}_{\max\{\tilde{z}\}|\tilde{z}\leq z}[a(\mathcal{S}(\tilde{z})) - c(\tilde{z})] > 0$, reflects the increase in expected match surplus of other intermediate producers of the same input by marginally increasing specialization, for given surplus conditional on matching and compatibility probability of own input. Intuitively, this effect captures the increase in trading probability of intermediate producers with productivity lower than z producing the same input. The fourth element, $\theta\lambda\mathcal{P}(z;N)(N-1)\phi'(\mathcal{S}(z))/\phi(\mathcal{S}(z))\exp\{-\lambda\hat{\phi}(\underline{z},z)\}\hat{\mathbb{E}}[a(s(\tilde{z})) - c(\tilde{z})] < 0$, reflects the reduction in expected match surplus of intermediate producers of other inputs by marginally increasing specialization. This effect represents the reduction in trading probability of intermediate producers of complementary inputs. These are the external effects of specialization by an intermediate producer with productivity z on the expected surplus of *other* intermediate

[12]Notice that the problem setup implicitly assumes that the social planner makes final producers trade with the compatible input provider contacted with the highest productivity, which can be easily proven.

producers. As such, these effects are not internalized in the condition determining the privately optimal specialization (8).

The social marginal cost on the right-hand side of (11) is given by the marginal increase in hours worked induced by higher specialization, weighted by the marginal disutility of labor supply. The social marginal cost coincides with the private marginal cost.

We are interested in the efficiency properties of equilibrium specialization. Comparing (8) with (11) allows us to establish our first main result.

Proposition 1 (Efficiency)**.** *The equilibrium allocation is constrained efficient if and only if production is not complex, i.e.* $N = 1$*. If the production process is complex, i.e.* $N > 1$*, the equilibrium features over-specialization.*

We discuss the intuition behind Proposition 1 by highlighting the externalities at play in the equilibrium allocation. To do so, we focus on the marginal welfare effect of specialization. By definition, efficient specialization equalizes the marginal welfare effect of specialization to zero, i.e., $\frac{\partial \mathcal{W}}{\partial s(z)}\big|_{s(z)=\mathcal{S}(z)} = 0$. Hence, evaluating the marginal welfare effect of specialization at the equilibrium solution allows us to single out the sources and direction of the inefficiency induced by equilibrium specialization. Formally,

$$\left.\frac{\partial \mathcal{W}}{\partial s(z)}\right|_{s(z)=s^\star(z)} \propto \underbrace{\mathbb{E}_{\max\{\tilde{z}\}|\tilde{z}\leq z}\left[a(s^\star(\tilde{z})) - c(\tilde{z})\right]}_{\text{business-stealing externality}} \underbrace{- x^\star(z)}_{\substack{\text{appropriability}\\ \text{externality}}} \underbrace{- (N-1)e^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[a(s^\star(\tilde{z})) - c(\tilde{z})]}_{\text{network externality}}. \tag{12}$$

Equilibrium specialization induces three externalities. First, intermediate producers do not internalize that decreasing own specialization reduces the trading probability of all firms with lower productivity producing the same input (*business-stealing externality*). Second, intermediate producers do not internalize the total foregone surplus from lower trading probability, but only the part they able to appropriate (*appropriability externality*). Finally, intermediate producers do not internalize that raising their own specialization reduces the trading probability of all intermediate producers of complementary inputs. By decreasing the likelihood that the final producer sources the input, which in turn reduces the probability that the good is made and the complementary input providers are remunerated (*network externality*).

To evaluate the efficiency properties of equilibrium specialization, we proceed in two steps. First, we consider the case where production is not complex, *i.e.*, $N = 1$. Since production requires only one input, the network externality vanishes and the economy is efficient if and only if the surplus offered to final producers by an intermediate producer, $x^\star(z)$, is equal to the external effect it exerts on lower-productivity firms, that is if the *appropriability externality* exactly offsets the *business-stealing externality*. From Lemma 1, we know that the surplus offered to final producers in equilibrium equals to the expected outside option of final producers when trading with the intermediate producer. The latter coincides with the expected social surplus of the best firm with lower productivity producing the same input. Hence, when production is not complex, intermediate producers offer a surplus that exactly offsets their external effect on firms

with lower productivity. Therefore, under our price/surplus-posting mechanism, equilibrium specialization is efficient.[13]

When production is complex ($N > 1$), business-stealing externality and appropriability externality still cancel due to the surplus-(price-)posting mechanism, but the network externality does not vanish anymore and the marginal welfare effect of specialization, evaluated at the equilibrium solution, is negative. Hence, the planner would impose a lower level of specialization and, therefore, the equilibrium features *over-specialization*. Importantly, the network externality represents a *pure* externality, as it arises only in no-trade scenarios. In other words, higher specialization of intermediate producers of input j harms the intermediate producers of other inputs $i \neq j$ whenever the final producers they are in contact with are unable to source input j. As a consequence, there is no price that allows intermediate producers of input j to internalize this negative externality, simply because no trade occurs.

Network externality. The network externality arises from the interplay of three elements. First, search and compatibility frictions imply that final producers may not find an input they are looking for ($f < 1$). Second, endogenous product design decisions imply that the input finding probability is lower, the higher the specialization of an input ($f'(s) < 0$). Third, complementarity among inputs in final production implies that failure to source one input prevents the surplus from successful matches between final and intermediate producers to materialize ($N > 1$). We clarify the economic forces behind the network externality in the following 2-firm example.

Assume two input providers produce the only two complementary inputs needed for final production ($N = 2$). Suppose further that final producers only meet one intermediate producer per input, which implies that the latter does not have to compete with anyone else and can extract the whole surplus.[14] Denote as $f(s_i)$ the probability that the input provider i is compatible with the final producer when its specialization is s_i. Therefore, the probability that a final producer source both inputs is $f(s_1)f(s_2)$. Higher specialization raises match surplus $y(s_i)$ ($y' > 0$), but reduces the compatibility probability ($f' < 0$). Let $q(s_i)$ be the cost of specialization (in terms of final good). Social surplus reads $\mathcal{W} = f(s_1)f(s_2)[y(s_1) + y(s_2)] - [q(s_1) + q(s_2)]$. Efficient specialization of firm 1 solves:

$$\underbrace{f(\mathcal{S}_1)y'(\mathcal{S}_1)}_{\text{social MB}} = \underbrace{(-f'(\mathcal{S}_1))[y(\mathcal{S}_1) + y(\mathcal{S}_2)] + q(\mathcal{S}_1)}_{\text{social MC}}.$$

On the left-hand side is the additional surplus generated by the higher specialization, while on the right-hand side, there is the higher cost of specialization q and the expected foregone social surplus in case of no trade.

Expected private surplus of firm 1 reads $\pi_1 = f(s_1)f(s_2)y(s_1)$. Equilibrium specialization

[13] Readers familiar with CES monopolistic competition models will have noted the same coincidence of offsetting externalities. We discuss this parallel in detail at the end of this section.

[14] Since in this example each firm is the only producer of its good, there is no business-stealing externality on other firms providing the same input with lower productivity. As a consequence, the fact that the intermediate producer offers $x = 0$ surplus is not an inefficiency per se.

solves:

$$\underbrace{f(s_1^\star)y'(s_1^\star)}_{\text{private MB}} = \underbrace{(-f'(s_1^\star))y(s_1^\star) + q(s_1^\star)}_{\text{private MC}}.$$

Evaluating the efficiency condition at the equilibrium specialization, we note that we are left with the term $f'(s_1^\star)y(s_2^\star) < 0$, which captures the negative externality of firm 1 on the trading probability of firm 2. It follows that the social marginal cost is higher than the private one, and firm 1 over-specializes. By symmetry, the same holds for firm 2, as well.

2.4 Discussion

Before extending the model to dynamics, it is worth noticing how our static framework relates to existing models.

Monopolistic competition with CES production function A common framework when studying firm-to-firm relations is that of Dixit and Stiglitz (1977), as exemplified by Grossman et al. (2023a,b). In this class of models, firm entry gives rise to a *business-stealing* externality and an *appropriability* externality. Still, the equilibrium allocation with free entry (and inelastic labor supply) is constrained efficient. The reason is that the two externalities exactly offset each other, as a consequence of the elasticity of substitution – which governs the business-stealing externality – and the taste for varieties – which governs the appropriability externality – being the same parameter. In models where these do not coincide, the economy is generally inefficient.[15] In our model, endogenous product design decisions are akin to firm entry decisions. Indeed, average input specialization determines the mass of compatible varieties for each final producer. Our price-posting protocol implies that equilibrium specialization gives rise endogenously to the same offsetting externalities as in Dixit and Stiglitz (1977). However, in our model, this result arises from the optimal choice of input providers who make final good producers indifferent between themselves and the expected outside option. That is, this result arises without a parametric restriction, which imposes that the two externalities are governed by the same parameter.

Remark 1 (Monopolistic Competition & CES)**.** *The combination of endogenous specialization and price posting in frictional markets with simultaneous search provides an endogenous mechanism such that the appropriability and business-stealing externalities from designing new product varieties exactly offset each other. Models of monopolistic competition with CES final production function obtain the same result by imposing the exogenous parametric restriction that the elasticity of substitution across varieties in the final production function equals the elasticity of firm-level inverse demand.*

Bargaining over surplus shares Our baseline model of specialization posits that intermediate producers post the price of their goods (or, interchangeably, set the surplus offered to final

[15] See Dixit and Stiglitz (1975); Spence (1976); Dixit and Stiglitz (1977); Ethier (1982); Mankiw and Whinston (1986); Benassy (1996); Zhelobodko et al. (2012); Parenti et al. (2017); Dhingra and Morrow (2019); Bajo et al. (2024) for several analysis of this problem and Brakman and Heijdra (2001) for a review.

producers) before meeting. We now show how our baseline model would change if intermediate and final producers bargained over the surplus shares.

Let the structure of our model be the same up to the surplus-sharing rule. Assume that the surplus accruing to final producers is determined by generalized Nash bargaining upon meeting between a final producer and a compatible intermediate producer. Final producers can observe the total surplus v from matching with any intermediate producer they meet. Let $\xi \in [0,1]$ denotes the exogenous bargaining power weight of final producers. The surplus accruing to final producers in equilibrium solves:

$$x^\star(v) = \underset{x}{\text{argmax}} \left(x - \omega(v; \vec{v}_k)\right)^{\xi} \left(v - x\right)^{1-\xi},$$

where $\omega(v; \vec{v}_k)$ denotes the outside option of a final producer which is contact with k compatible intermediate producers whose total surplus is lower than v. While negotiating with the highest-surplus compatible intermediate producer, final producers keep the option of bargaining with the other less productive compatible intermediate producers they are in contact with. Formally, $\omega(v; \vec{v}_k) = \max\{v_1, \ldots, v_k\}$. In Appendix A.4 we show that the optimal specialization function $s_B^\star(z)$ with bargaining over surplus shares is implicitly defined by:

$$\theta\lambda\mathcal{P}(z;N)(1-\xi)\left[a'(s_B^\star(z)) + \frac{\phi'(s(z))}{\phi(s(z))}\left(a(s_B^\star(z)) - c(z) - \mathbb{E}[\omega(z)]\right)\right] = wq'(s_B^\star(z)), \tag{13}$$

where $\mathbb{E}[\omega(z)]$ is the expected outside option of a final producer trading with the intermediate producer, which equals the surplus offered to final producers in our baseline model.

The equilibrium specialization function $s_B^\star(z)$ is generally inefficient due to the network externality *and* a well known hold-up externality (Tirole, 1999). The hold-up problem is related to the appropriability externality in our baseline model. However, there are important differences between the two. Hold-up arises when intermediate and final producers split ex-post the surplus generated by the investment in specialization. Increasing investment in specialization marginally generates a' units of additional social surplus, which is split according to $(\xi,\, 1-\xi)$ shares. Hence, part of the *marginal* surplus generated by the intermediate producer's investment goes to the final producer. Because the social and private returns on investment are different, the investment level is socially inefficient. Differently, in our model with price posting, the intermediate producer commits to a *level* x of surplus offered to the final producer. Hence, any additional surplus generated by increasing investment in specialization marginally, a', only accrues to the intermediate producer itself. Since marginal social and private surpluses coincide, the investment level is efficient. As is well known, in the limit case of the bargaining model where the seller has all the bargaining power there is no hold-up problem. This is because social and private surplus coincide – not just at the margin. In this sense, our model allows for bilaterally efficient investment even in a setting in which both parties in the transaction obtain a positive surplus. This result is driven by the competitive pressure exerted by the simultaneous search, which implies that the downstream firms have a non-zero outside option.

In the bargaining setting, for specialization investment to be efficient, we would need equation

(13) to hold for all z. Clearly, this cannot be the case with a single bargaining weight ξ. If we allow bargaining power weights to depend on the productivity of the intermediate producer involved in the transaction, there exists a unique schedule of bargaining power weights that implements the efficient allocation, $\xi^\star(z)$. If intermediate producers do not hold all the bargaining power, the hold-up externality pushes equilibrium specialization to be inefficiently low. The bargaining weight schedule $\xi^\star(z)$ makes the hold-up and network externalities exactly offset each other.

Comparing Equations (13) and (8) allows us to make the following remark.

Remark 2 (General Bargaining Model)**.** *The equilibrium allocation of our baseline model is the same as that of a general bargaining model where intermediate producers hold all bargaining power, i.e.,* $\xi = 0$.

Overall, when production is not complex, our baseline model can be interpreted as the efficient benchmark of a general bargaining model.

Non-contingent contracts Our model considers a particular set of contingent contracts: transactions between intermediate and final producers that successfully match do not happen if the latter are unable to source all the complementary inputs needed for production. This kind of contract implies that final producers bear no risk of having to pay suppliers when unable to produce. We now study what would change with an alternative, non-contingent arrangement whereby suppliers are remunerated independently of the successful sourcing of all inputs. In such a setting, final producers bear all the sourcing risk. In equilibrium, all final producers that cannot successfully source all inputs incur losses that are borne by the representative household (in terms of lower profit rebates).

For intermediate producers, this means that the trading probability is independent of the complexity of final production, i.e., $\tilde{\mathcal{P}}(s, x; N) = \tilde{\mathcal{P}}(s, x) = \phi(s) \exp\{-\lambda\bar{\phi}(1 - G(x))\}$ since they get remunerated independently of successful sourcing. On the one hand, the optimal surplus offered to final producers is unaffected by the contracting scheme and is still given by equation (9). On the other hand, optimal specialization $s^\star_{NC}(z)$ differs from the baseline model with contingent contracts, and is implicitly defined by:

$$\theta\lambda\tilde{\mathcal{P}}(z)\left[a'(s^\star_{NC}(z)) + \frac{\phi'(s^\star_{NC}(z))}{\phi(s^\star_{NC}(z))}\left(a(s^\star_{NC}(z)) - x^\star_{NC}(z) - c(z)\right)\right] = wq'(s^\star_{NC}(z)). \tag{14}$$

Since the planner is indifferent as to how the total surplus is split, efficient specialization is still determined by (11). We proceed by evaluating the marginal welfare effect of specialization at

the equilibrium solution with non-contingent contracts. Formally,

$$\left.\frac{\partial \mathcal{W}}{\partial s(z)}\right|_{s(z)=s^\star_{NC}(z)} \propto \underbrace{-\left(a'(s^\star_{NC}(z)) + \frac{\phi'(s^\star_{NC}(z))}{\phi(s^\star_{NC}(z))}(a(s^\star_{NC}(z)) - c(z))\right)(1 - f^{N-1})}_{\text{contracting externality}} \tag{15}$$

$$- \frac{\phi'(s^\star_{NC}(z))}{\phi(s^\star_{NC}(z))}\Bigg[\underbrace{f^{N-1}\mathbb{E}_{\max\{\tilde{z}\}|\tilde{z}\leq z}\left[a(s^\star_{NC}(\tilde{z})) - c(\tilde{z})\right]}_{\text{business-stealing externality}} \underbrace{-x^\star_{NC}(z)}_{\substack{\text{appropriability}\\ \text{externality}}}$$

$$\underbrace{-(N-1)\exp\{-\lambda\hat{\phi}(\underline{z}, z)\}\hat{\mathbb{E}}[a(s^\star_{NC}(z)) - c(z)]}_{\text{network externality}}\Bigg].$$

First, we note that Proposition 1 also holds in an economy with non-contingent contracts: the equilibrium is efficient if and only if production is not complex, as both the *contracting* and *network externalities* vanish. However, if the production process is complex, non-contingent contracts exacerbate over-specialization. In addition to the network externality, two additional forces induce equilibrium over-specialization. First, contracting frictions imply that the private marginal benefit of specialization exceeds the social marginal benefit. Second, the appropriability externality exceeds the business-stealing externality. Intuitively, as intermediate producers are remunerated independently of whether sourcing is successful or not, they have stronger incentives to over-specialize when contracts are not contingent. We summarize our findings in the next remark.

Remark 3 (Non-contingent contracts)**.** *Economies with non-contingent contracts feature more equilibrium over-specialization than economies with contingent contracts.*

Boundaries of the firm, vertical integration, and in-house production At the heart of our model, there are frictions in sourcing inputs outside the firm. These frictions are exacerbated by the complexity of production N. We study later the problem of how firms may optimally choose complexity. Here, we are interested in understanding how the boundaries of the firm affect our results. Suppose that downstream firms can produce in-house $M < N$ inputs that are not subject to sourcing frictions. This is equivalent to vertical mergers with $M < N$ input suppliers.

Remark 4 (In-house production)**.** *In-house production reduces the inefficiencies associated with endogenous specialization by lowering the number of key inputs subject to search frictions. In the limit in which $M = N - 1$, the economy is constrained efficient.*

In the special case where only one input is left to be sourced outside the boundary of the firm, it is as if the production process is not complex. As a consequence, the result Proposition 1 for $N = 1$ applies, and the economy is constrained efficient.

Inventories and multi-sourcing Finally, we note that sensible strategies to avoid halting production following a disruption, such as inventories and multi-sourcing are ineffective in

the context of our static model. We return to these potential solutions later when we study a dynamic setting with link-specific disruption shocks.

3 Dynamic Model

In the previous section, we have studied how endogenous specialization of intermediate inputs affects the sourcing capacity of final producers. Our finding is that the equilibrium allocation features over-specialized inputs as intermediate producers do not internalize the cascading effect of halting final production on the value added produced by complementary inputs. In this section, we show that the resilience of supply chains is closely related to the dynamic sourcing capacity of final producers after shocks. To this end, we extend the static model to a dynamic framework with long-term relationships between intermediate and final producers and stochastic disruptions of final producers. This framework allows for a model-consistent and welfare-relevant notion of supply chain resilience, that is, the average time it takes for a final producer to restore production following a disruption.

3.1 Environment

Time is discrete and infinite. Agents discount the future at rate $\beta \in (0,1)$. Final producers face a disruption at the end of each period with probability δ, where $\delta \in (0,1)$. Upon facing a disruption, final producers lose contact with all their input providers. Final producers that are unable to source all the complementary inputs on the market are inactive in the respective period and change their management in the following period.[16] Final producers keep the same input providers until a disruption occurs. In other words, final producers do not search the market for replacing existing suppliers, *e.g.* due to switching costs.[17] We refer to the average time a final producer does not face any disruption, *i.e.* $1/\delta$, as *robustness*. Intermediate producers never face disruptions and search for new customers in the market in each period. The dynamic model nests the static model as final producers lose contact with all their input providers at the end of every period ($\delta \to 1$).

The timing convention is as follows: at the end of each period, a share δ of attached final producers faces a disruption, which makes them lose all their input providers. At the start of the next period, newly and already unattached final producers search for suppliers. Hence, the mass of searching final producers μ is governed by the following law of motion:

$$\mu_{t+1} = \delta \left[1 - \left(1 - \prod_{j=1}^{N} f_{j,t} \right) \mu_t \right] + \left(1 - \prod_{j=1}^{N} f_{j,t} \right) \mu_t. \tag{16}$$

[16] This implies that the decision to establish a supplier relationship with a compatible input provider embeds no option value of waiting.

[17] See Esala (2024) for the role of search for better partners in supply chains.

Market tightness equals the ratio between searching final producers and intermediate producers:[18]

$$\theta_t = \frac{\mu_t}{m}. \tag{17}$$

The representative household and final producers make only static decisions. Intermediate producers choose the specialization of their input in each period t, as well as commit to a constant surplus offered to final producers in any period, by solving the following intertemporal profit maximization problem:[19]

$$\max_{x,\{s_t\}_{t=1}^{\infty}} \sum_{t=1}^{\infty} \beta^{t-1} \left[\mathcal{D}_t(\boldsymbol{s}_{1,t}, x)\left(a(s_t) - x - c(z)\right) - w_t q(s_t)\right] \tag{18}$$

$$\begin{aligned} \text{s.t. } & \mathcal{D}_t(\boldsymbol{s}_{1,t}, x) = (1-\delta)\mathcal{D}_{t-1}(\boldsymbol{s}_{1,t-1}, x) + \theta_t \lambda \mathcal{P}(s_t, x; N), \\ & \mathcal{D}_0 = 0, \end{aligned} \tag{19}$$

where $\mathcal{D}_t$ denotes the number of final producers sourcing from an intermediate producer, or *demand*, in each period t. Since each intermediate producer serves a measure of final producers, a share δ of existing customers faces a disruption and breaks the supplier relationship in each period. At the same time, the intermediate producer acquires a measure $\theta_t \lambda \mathcal{P}(s_t, x)$ of new customers, depending on the market tightness θ_t and its own policies (s_t and x).

3.2 Stationary equilibrium

We now characterize the stationary equilibrium with balanced flows of our dynamic economy. The optimal surplus offered to final producers is the same as in the static model, given by (9) (see Appendix A.5 for the derivation). The stationary demand of a seller choosing specialization s and surplus x is:

$$\mathcal{D}(s, x; N) = \frac{\theta \lambda \mathcal{P}(s, x; N)}{\delta}. \tag{20}$$

The optimal specialization function is implicitly defined by the following condition:

$$\mathcal{D}(z; N)\left[a'(s^\star(z)) + \delta[1 + \beta(1-\delta)]\frac{\phi'(s(z))}{\phi(s(z))}\left(a(s^\star(z)) - c(z) - x^\star(z)\right)\right] = w q'(s(z)). \tag{21}$$

Increasing specialization affects old and new customers differently. For both groups $\mathcal{D}(z; N)$, more specialization implies a higher surplus. For surviving customers $(1-\delta)\mathcal{D}(s, x; N)$, higher specialization does not come at the cost of lower trading probability since the match has already occurred in the past. For potential customers $\delta\mathcal{D}(s, x; N)$ instead, we have the familiar effect

[18] Throughout, we maintain the assumption that m is fixed. In Appendix B.4, we endogenize the mass of intermediate goods producers via a free entry condition.

[19] Absent commitment to a constant surplus offered to final producers, intermediate producers would optimally choose dynamic pricing plans. Final good firms would then have to choose the supplier with the highest present discounted value of surplus streams.

of lower trading rate ϕ'. Nonetheless, the forward-looking firm understands that lower trading probability today means lower demand tomorrow. Hence, the reduction in trading probability is weighted by the intertemporal value of a marginal customer today, $1 + \beta(1 - \delta)$.

Output in this economy can be decomposed into four terms:

$$Y = \underbrace{\mu(f; N)}_{\text{market size}} \quad \underbrace{\mathcal{R}(f; N)}_{\text{resilience}} \quad \underbrace{\frac{1}{\delta}}_{\text{robustness}} \quad \underbrace{N\hat{\mathbb{E}}[a(s^\star(z)) - c(z)]}_{\text{expected match surplus| active}} .$$

First, output depends positively on market size (of investment in specialization). Market size increases if the mass of searching final producers μ increases. Second, output is increasing in the economy's resilience, *i.e.* the probability that a searching final producer finds all the complementary inputs it needs. Third, output is increasing in the economy's robustness, *i.e.* the average duration of a match between final producers and intermediate producers. Finally, output is increasing in the expected match surplus conditional on production.

Resilience. The dynamic model allows for a model-consistent and welfare-relevant notion of supply chain resilience, that is, the average time it takes for a final producer to restore production following a disruption. Given this definition, resilience equals the probability that a final producer is able to source all the complementary inputs it needs after a disruption.[20] Hence, our model allows us to single out three determinants of supply chain resilience:

$$\mathcal{R}(f; N) \equiv f^N = [1 - \exp\{-\lambda\bar{\phi}\}]^N. \tag{22}$$

First, higher search efficiency ($\lambda \uparrow$), e.g. through ICT adoption and AI-based procurement, increases resilience of supply chains.[21] Second, higher resilience can be attained by relying on less specialized inputs ($\bar{\phi} \uparrow$). This corresponds to the strategy of building resilience "through products" highlighted by Miroudot (2020), such that firms choose to source more standardized inputs in the market because they are easier to replace. Finally, higher complexity of the production process ($N \uparrow$) reduces supply chain resilience. This echoes the main insight from the network fragility literature (Elliott et al., 2022). Overall, our dynamic model provide us with an intuitive notion of supply chain resilience with realistic properties. Moreover, supply chain resilience directly affects output and, in turn, welfare. In this sense, our model justifies policymakers' interest in supply chain resilience.

Customization. Our dynamic model allows intermediate producers to re-optimize the specialization of their products in every period. This introduces a model-consistent notion of *customization*, that is, tailoring the product to the needs of the final producers *within a match*. Long-term relationships make specialization more appealing for intermediate producers: within an existing match, increasing specialization comes at no cost in terms of trading probability.

[20] In the next section, we show that a similar expression holds even if disruptions occur in individual buyer-supplier relationships.

[21] See the practicioners' reports of Baptista et al. (2021); Drentin et al. (2024).

The main prediction of our dynamic model is that intermediate producers gradually increase the specialization of their products during their life cycle as they progressively accumulate customers. The intuition is easily understood through the FOC for specialization:

$$(1-\delta)\mathcal{D}_{(-1)}a'(s)+\theta\lambda\mathcal{P}(s,x)\left[a'(s)+\frac{\phi'(s)}{\phi(s)}\left([a(s)-c(z)-x]+\beta(1-\delta)[a(s_{(+1)})-c(z)-x]\right)\right]$$
$$=wq'(s). \tag{23}$$

Optimal specialization features both a backward-and a forward-looking component. The marginal benefit of specialization on the LHS can be divided into two components. First, for the $(1-\delta)\mathcal{D}_{(-1)}$ surviving matches, higher specialization implies higher surplus, a'. For new potential matches, this benefit is traded off against the lower trading probability, governed by $\phi'<0$. As a firm accumulates customer base $\mathcal{D}$, the first term becomes more important, increasing the marginal benefit of specialization. Since intermediate producers have overcome compatibility frictions with more and more customers, intermediate producers prefer to increase match surplus with existing customers rather than attract new ones. Hence, for the share of attached customers, specialization corresponds to customization. It follows that both the value and the price of the inputs traded between firms increase over time.

3.3 Efficiency

We now study whether the efficiency properties of the static model extend to our dynamic setting. The social planner solves the following recursive problem:

$$\mathcal{W}\left(\boldsymbol{\mathcal{D}}_{(-1)}(z),\mu\right)=\max_{\boldsymbol{s}(z)} m\sum_{n=1}^{N}\int D_j(z)[a(s_j(z))-c(z)]\gamma(z)dz+\psi\log(1-\ell)+\beta\mathcal{W}\left(\boldsymbol{\mathcal{D}}(z),\mu_{(+1)}\right)$$

$$\text{s.t. } D_j(z)=(1-\delta)D_{j(-1)}(z)+\frac{\mu}{m}\lambda\phi(s_j(z))\exp\{-\lambda\hat{\phi}(z,\bar{z})\}\prod_{v\neq j}f_v,$$

$$\mu=\delta\left[1-\left(1-\prod_{n=1}^{N}f_{n(-1)}\mu_{(-1)}\right)\right]+\left(1-\prod_{n=1}^{N}f_{n(-1)}\right)\mu_{(-1)},$$

$$\ell=m\sum_{n=1}^{N}\int q(s_n(z))\gamma(z)dz.$$

In words, the social planner chooses the specialization $s_j(z)$ of each intermediate producer with productivity z producing input j to maximize the lifetime utility of the representative household. The state variables are the vector of demand for all intermediate producers one period earlier, $\boldsymbol{\mathcal{D}}_{(-1)}$, and the share of searching final producers one period earlier, $\mu_{(-1)}$. The maximization problem is subject to the laws of motion of firm-level demand and of the share of searching final producers, as well as to the labor resource constraint.

Efficient specialization of intermediate producers with productivity z is given by:

$$\mathcal{D}(z;N)\Bigg[a'(\mathcal{S}(z))+\delta[1+\beta(1-\delta)]\frac{\phi'(\mathcal{S}(z))}{\phi(\mathcal{S}(z))}\Bigg(a(\mathcal{S}(z))-c(z)-\mathbb{E}_{\max\{\tilde{z}\}|tildez\leq z}[a(\mathcal{S}(\tilde{z}))-c(\tilde{z})]$$
$$+(N-1)e^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[a(\mathcal{S}(\tilde{z}))-c(\tilde{z})]-\frac{\beta(1-\delta)}{1+\beta(1-\delta)}f^N Ne^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[a(\mathcal{S}(\tilde{z}))-c(\tilde{z})]\Bigg)\Bigg] \tag{24}$$
$$=\frac{\psi}{1-Nm\bar{q}}\,q'(\mathcal{S}(z)).$$

Comparing the dynamically efficient specialization (24) with its static counterpart (11), we observe that a further inefficiency arises in a dynamic setting: intermediate producers do not internalize the effect of their specialization on the equilibrium share of searching final producers. This force pushes the private marginal cost of specialization to exceed the social one. Rearranging (24) allows us to establish the following result.

Proposition 2 (Efficiency of the Dynamic Economy)**.** *The equilibrium allocation is constrained efficient if and only if* $N=\frac{1}{\mathcal{M}(N)}$*, where* $\mathcal{M}(N)\equiv 1-f^N\frac{\beta(1-\delta)}{1+\beta(1-\delta)}\leq 1$*. If the production process is complex, i.e.* $N>1$*, the equilibrium features over-specialization.*

Notice that $\frac{1}{\mathcal{M}(N)}=\frac{1+\beta(1-\delta)}{1+\beta(1-\delta)(1-f^N)}\in[1,2)$. Hence, our main result from the static model goes through this dynamic setting: if the production process is complex, the equilibrium allocation features over-specialization. We follow the same steps used in the static setting to derive the marginal welfare effect of equilibrium specialization. We decompose the marginal welfare effect into four terms. The first three terms represent the externalities discussed in the static model: the business-stealing, appropriability, and network externalities. A fourth effect arises in the dynamic model: intermediate producers choosing their specialization affect the market tightness by varying the probability that inputs are compatible. This effect – akin to the standard search externality in models of endogenous search effort (Pissarides, 2000) – is not internalized by intermediate producers.

$$\left.\frac{\partial\mathcal{W}}{\partial s(z)}\right|_{s(z)=s^\star(z)}\propto \underbrace{\mathbb{E}_{\max\{\tilde{z}\}|\tilde{z}\leq z}\left[a(s^\star(\tilde{z}))-c(\tilde{z})\right]}_{\text{business-stealing externality}}\underbrace{-\,x^\star(z)}_{\substack{\text{appropriability}\\ \text{externality}}}\underbrace{-\,(N-1)e^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[a(s^\star(\tilde{z}))-c(\tilde{z})]}_{\text{network externality}}$$
$$\underbrace{+\,\frac{\beta(1-\delta)}{1+\beta(1-\delta)}f^N Ne^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[a(s^\star(\tilde{z}))-c(\tilde{z})]}_{\text{search externality}}. \tag{25}$$

The decomposition of the marginal welfare effect of equilibrium specialization allows us to shed light on the economic forces behind Proposition 2. We know from Lemma 1 that the first two terms in (25) cancel out in equilibrium thanks to price posting in frictional markets with simultaneous search. As before, consider first the case of non-complex production. By setting $N=1$, we can eliminate the network externality, so we are left with the new source of inefficiency alone. This effect is positive, so the planner would like to induce further specialization. This new external effect implies that the socially optimal level of specialization is higher, all else equal.

This also implies that, for $N > 1$, the equilibrium is "less inefficient" than in the static case, as the network and search externalities partially offset each other.

Comparative statics. Over the last decades, supply chains have become more complex (Antràs, 2020; Ferrari, 2023), disruptions have become more frequent due to the occurrence of wars, trade disputes, pandemics, and extreme weather events, and the rise of new technologies for procurement may have reduced search frictions (Drentin et al., 2024). We study each of these trends through the lens of our model.

Under mild parametric restrictions, Proposition 3 shows that over-specialization of the equilibrium allocation exacerbates as the complexity of the production process increases.

Proposition 3 (Efficiency and Complexity)**.** *If search efficiency is low enough, the equilibrium allocation of the dynamic economy displays more over-specialization and under-resilience as complexity increases.*

Similarly, Proposition 4 shows that over-specialization of the equilibrium allocation exacerbates as the economy becomes more fragile, i.e., disruptions occur more frequently.

Proposition 4 (Efficiency and Robustness)**.** *As the frequency of disruption increases, the equilibrium allocation of the dynamic economy displays more over-specialization and under-resilience.*

Finally, Proposition 5 sheds light on how the endogenous response of specialization to declining search frictions shapes the effect of declining search frictions on the efficiency properties of equilibrium specialization.

Proposition 5 (Efficiency and Search Frictions)**.** *If the elasticity of the average compatibility probability to search efficiency exceeds one, the equilibrium allocation of the dynamic economy displays more over-specialization and under-resilience as search frictions decline.*

Taken together, Propositions 3 and 4 suggest that when complexity rises and disruptions become more frequent, the economy becomes ever more inefficient. Through the lens of our model, we would conclude that supply chains are less efficient today than they used to be, as they feature ever more over-specialization and under-resilience. Proposition 5, instead, states that when procurement technology improves, the efficiency effects are ambiguous. In particular, they depend on the endogenous response of firms in terms of specialization. If firms react to the technology improvement by increasing specialization more than proportionally, the economy veers towards over-specialization and under-resilience.

3.4 Dynamic Model with Link Destruction

For analytical transparency, in the dynamic model discussed so far we assumed that, once final producers face a disruption shock, they lose contact with all their input providers and have to form supplier relationships anew. This is a strong assumption in that, presumably, many supply

chain disruptions are caused by single supplier relationships breaking up. Since production inputs are complements, our model is suited for studying supply chain disruptions driven by the failure to replace single inputs, as well. To do so, we study a variant of our baseline dynamic model in which each supplier relationship breaks stochastically. It follows that it is enough for the final producer to replace the broken supplier relationships to continue production. We show that the same sources of inefficiency and key results of the baseline dynamic model go through this more realistic setting.

The main innovation with respect to the baseline dynamic model is that $\delta \in (0,1)$ now denotes the probability that one of the links between intermediate and final producers gets destroyed. Hence, a disruption happens if the final producer is unable to replace the input providers who separate each period.[22] Let $\rho(f;N)$ be the probability that an active final producer is operational next period:

$$\rho(\boldsymbol{f}) = \prod_{i=1}^{N}[1-\delta(1-f_i)]. \tag{26}$$

It follows that the probability of being operational conditional on separating with the input provider of good i equals $f_i\rho(\boldsymbol{f}_{-i})$. The market tightness in each input market i equals $\theta_{i,t} = \frac{\mu_{i,t}}{m}$. Since inputs are complements, the mass of searching final producers is the same across input markets and governed by the following law of motion:

$$\mu_{t+1} = \delta\left[1-\left(1-\prod_{j=1}^{N} f_{j,t}\right)\mu_t\right] + \left(1-\prod_{j=1}^{N} f_{j,t}\right)\mu_t + \left(1-\delta\left(1-\prod_{j=1}^{N} f_{j,t}\right)-\rho(\boldsymbol{f}_t)\right)(1-\mu_t). \tag{27}$$

The law of motion (27) differs from its baseline counterpart (16) due to the presence of the last term. Intuitively, the last term corrects the law of motion for the composition of the mass of searching final producers. Indeed, searching final producers who were active in the previous period are more likely to remain so because they have to find fewer inputs on average than those who were inactive. Hence, the law of motion (27) nests (16) if final producers that lose contact with at least one input providers need to replace all of them, in which case $\rho(\boldsymbol{f}_t) = 1-\delta\left(1-\prod_{j=1}^{N} f_{j,t}\right)$.

Demand of an intermediate producer is still given by (19). However, the trading probability incorporates the differential probability that a searching final producer is active depending on its previous state. Let $\nu_t = \nu(\boldsymbol{f}_t;\mu_t)$ denote the share of searching final producers that were active in the previous period. The probability that a searching final producer is active upon finding input i in the market equals $\tilde{f}_i(\boldsymbol{f}_t) \equiv \nu_t\rho(\boldsymbol{f}_{-i}) + (1-\nu_t)\prod_{n\neq i} f_n$. Then, the trading probability of an intermediate producer of good i is given by:

$$\mathcal{P}_i(s_t,x;\tilde{f}) = \phi(s_t)\exp\{-\lambda\bar{\phi}_t(1-G(x))\}\tilde{f}(\boldsymbol{f}_t). \tag{28}$$

[22] As before, we assume that, following a *firm-level* disruption, a final producer loses contact with all its input providers.

Up to the different definition of trading probability, the profits maximization problem of intermediate producers is the same as in the baseline dynamic model, as given by (18).

Stationary equilibrium. The stationary equilibrium is characterized by the same set of equations as the baseline dynamic model. The equilibrium specialization function differs from its dynamic model counterpart (21) only in the use of $\tilde{f}(f; N-1)$ in place of f^{N-1}.

The stationary share of searching final producers equals $\mu(f;N) = 1-(1-\delta)\frac{f^N}{1+f^N-\rho(f;N)}$ and the stationary share of searching final producers that were active in the previous period equals $\nu(f;N) = \frac{1-\delta(1-f^N)-\rho(f;N)}{f\rho(f;N-1)-f^N}\frac{f^N}{1+\delta f^N-\rho(f;N)}$. Output in this economy is given by:

$$Y = \underbrace{\mu(f;N)}_{\text{market size}} \underbrace{\tilde{\mathcal{R}}(f;N)}_{\text{resilience}} \underbrace{\frac{1}{\delta}}_{\text{robustness}} \underbrace{N\hat{\mathbb{E}}[a(s^\star(z)) - c(z)]}_{\text{expected match surplus| active}} . \tag{29}$$

Output features the same four components discussed in the baseline dynamic model. The main difference is that, if disruptions happen at the level of single supplier relationships, supply chain resilience has a different expression:

$$\tilde{\mathcal{R}}(f;N) \equiv f(N)\tilde{f}(f;N-1) = [1-e^{-\lambda\bar{\phi}}]\left(\nu[1-\delta e^{-\lambda\bar{\phi}}]^{N-1} + (1-\nu)[1-e^{-\lambda\bar{\phi}}]^{N-1}\right). \tag{30}$$

Supply chain resilience is the weighted average of two terms. First, the probability that an active final producer is able to produce next period conditional on losing one input provider, $f[1-\delta(1-f)]^{N-1}$. Second, the probability that an inactive final producer is able to produce next period, f^N. These two terms are weighted by the stationary shares of searching final producers that were active and inactive one period earlier, respectively. This expression of supply chain resilience nests the baseline (22) as the share of searching final producers that were inactive approaches one.

Efficiency. The social planner solves the following recursive problem:

$$\mathcal{W}\left(\boldsymbol{\mathcal{D}}_{(-1)}(z),\mu\right)=\max_{s_j(z)} m\sum_{n=1}^{N}\int D_j(z)[a(s_j(z))-c(z)]\gamma(z)dz+\psi\log(1-\ell)+\beta\mathcal{W}\left(\boldsymbol{\mathcal{D}}(z),\mu_{(+1)}\right)$$

$$\text{s.t. } D_j(z)=(1-\delta)D_{j(-1)}(z)+\frac{\mu}{m}\lambda\phi(s_j(z))\exp\{-\lambda\hat{\phi}(z,\bar{z})\}\tilde{f}_j(\boldsymbol{f},\nu_j),$$

$$\mu=(1-\rho(\boldsymbol{f}_{(-1)}))(1-\mu_{(-1)})+\delta\prod_{n=1}^{N}f_{n(-1)}+\left(1-\prod_{n=1}^{N}f_{n(-1)}\right)\mu_{(-1)},$$

$$\tilde{f}_j(\boldsymbol{f}_t,\nu_j)=\nu_j\rho(\boldsymbol{f}_{-j})+(1-\nu_j)\prod_{n\neq j}f_n,$$

$$\nu_j=\frac{1-\delta(1-\prod_{n=1}^{N}f_n)-\rho(\boldsymbol{f})}{(1-\delta)\left(f_j\rho(\boldsymbol{f}_{-j})-\prod_{n=1}^{N}f_n\right)}\left(\frac{1}{\mu}-1\right),$$

$$\ell=m\sum_{n=1}^{N}\int q(s_n(z))\gamma(z)dz,$$

where the second constraint follows from rearranging (27).

Efficient specialization of intermediate producers with productivity z is given by:

$$\begin{aligned}&\mathcal{D}(z;N)\Bigg[a'(\mathcal{S}(z))+\delta[1+\beta(1-\delta)]\frac{\phi'(\mathcal{S}(z))}{\phi(\mathcal{S}(z))}\Bigg(a(\mathcal{S}(z))-c(z)-\mathbb{E}_{\max\{\tilde{z}\}|\tilde{z}\leq z}[a(\mathcal{S}(\tilde{z}))-c(\tilde{z})]\\&+\chi_1(f;N)(N-1)e^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[a(\mathcal{S}(\tilde{z}))-c(\tilde{z})]-\chi_2(f;N)e^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[a(\mathcal{S}(\tilde{z}))-c(\tilde{z})]\\&=\frac{\psi}{1-Nm\bar{q}}\,q'(\mathcal{S}(z)).\end{aligned}\tag{31}$$

where the multipliers $\chi_1(f;N)\in(0,1)$ and $\chi_2(f;N)\lessgtr 0$ are defined in Appendix A.6.

Comparing (31) with (24), we observe that the model with link destruction features the same sources of externality as the baseline dynamic model. Notice, however, that the search externality has no longer a clear sign since it features an additional component capturing the marginal effect of specialization on welfare via changing composition of the pool of searching final producers. Since higher specialization increases the share of inactive firms, this additional component pushes in the opposite direction with respect to the effect going through market tightness. Moreover, the network externality is dampened in magnitude since searching final producers need to source, on average, fewer inputs in the market each period.

We characterize the efficiency properties of the economy with link destruction in the next Proposition (6).

Proposition 6 (Efficiency of the Dynamic Economy with Link destruction)**.** *The economy with link destruction is efficient if and only if* $N=1+\frac{\chi_2(f;N)}{\chi_1(f;N)}$*. If* $N>1+\frac{\chi_2(f;N)}{\chi_1(f;N)}$*, then the economy with link destruction features over-specialization.*

To sum up, the basic economic forces highlighted in the baseline dynamic model go through

modeling link destruction rather than firm-level disruptions.

Strategies to mitigate the impact of supplier breakdowns. In our model, final producers occasionally halt production when they fail to source one or more key complementary inputs. Because intermediate producers do not account for the foregone surplus from complementary inputs when their own input is missing, they tend to design overly specialized products. As a result, production disruptions become more likely. However, if final producers could continue operating despite disruptions in some supplier relationships, this network externality would not arise. We now examine two key strategies that final producers can adopt to mitigate the impact of supplier breakdowns: maintaining inventories and sourcing inputs from multiple intermediate producers (multi-sourcing). See Appendix A.7 for a formal analysis.

Inventories. Suppose that active final producers can choose to purchase more units of each input than needed for current production, setting aside inventories as a buffer.[23] If a supplier relationship breaks down, these inventories allow production to continue until a new supplier is found or until the stock is depleted. However, due to search frictions, there is always a positive probability that a final producer is unable to locate a compatible intermediate producer for any number of periods. It follows that no finite level of inventories can completely hedge final producers from disruptions.

Multi-sourcing. Suppose final producers can establish links with multiple compatible suppliers met during a supplier search. For example, a final producer might maintain an idle backup link with the second-best compatible supplier at a fixed per-period cost. If the link with the best supplier breaks down first, the final producer can immediately switch to the backup supplier and continue production (see, for example, Elliott et al., 2022; Grossman et al., 2023b). However, due to search frictions, the probability of finding multiple compatible intermediate producers is less than one. As a result, not all final producers are able to multi-source all inputs. It follows that multi-sourcing alone cannot completely hedge final producers from disruptions.

Remark 5 (Inventories & Multi-sourcing)**.** *The ability of final producers to hold inventories or to source the same input from multiple suppliers can mitigate – but not eliminate – the network externality.*

In Appendix B, we discuss several further extensions to our framework. Specifically, we consider economies in which final good producers can choose the level of *completixity* of their production process as well as invest in *redundancies* to reduce the risk of disruptions. We also extend our results to a more general quality aggregator and economies in which the number of firms is determined by a free entry condition.

[23] The role of inventories in supply chains has been the subject of a recently active literature. See, for example, Alessandria et al. (2023); Ferrari (2023); Carreras-Valle (2023); Carreras-Valle and Ferrari (2025).

4 Normative Analysis

In this section, we characterize the optimal policy that decentralizes the constrained efficient allocation in the different economies we have studied in the previous sections.

Static Economy We introduce a targeted tax instrument that allows the planner to redistribute surplus between intermediate and final producers.

From Proposition 1, it follows that intermediate producers offer too high a surplus to final producers. In particular, the efficient surplus equals:

$$\mathcal{X}(z) = x^\star(z) - (N-1)e^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[a(s(z)) - c(z)]. \tag{32}$$

We can write the wedge between the social and private surplus of intermediate producers with productivity z as $\mathcal{X}(z) = x^\star(z) - \tau^\star(z)$ and look for a transaction subsidy schedule that decentralizes the constrained efficient allocation. From (32), we can derive the following result.

Proposition 7 (Optimal Subsidy Schedule in a Static Economy)**.** *The Social Planner can decentralize the constrained efficient allocation via a targeted transaction subsidy schedule $\tau^\star(z)$ such that the price of the transactions is given by $p^\star(z) = p(z) + \tau^\star(z)$, where*

$$\tau^\star(z) = e^{-\lambda\hat{\phi}(\underline{z},z)}T^\star > 0, \tag{33}$$

and $T^\star \equiv (N-1)\, e^{-\lambda\bar{\phi}}\hat{\mathbb{E}}[a(s(z)) - c(z)] > 0$ is a function that only depends on aggregates. The subsidy schedule is decreasing in z and bounded between $[\tau^\star(\bar{z}), \tau^\star(\underline{z})] = [(1-f)^{-1}T^\star, T^\star]$.

While it is natural that $x^\star(z) > 0$, $\forall z$ in equilibrium, this needs not hold in the decentralized constrained efficient allocation. Let $\check{z} : \int_{\underline{z}}^{\check{z}}[a(s(\tilde{z})) - c(z)]\gamma^\star(\tilde{z},\check{z})d\tilde{z} = (N-1)e^{-\lambda\bar{\phi}}\hat{\mathbb{E}}[a(s(z)) - c(z)]$. A social planner would require that intermediate producers with productivity lower than $\check{z}$ offer a negative surplus to final producers for their specialization to be efficient. Notice that this can be sustained in equilibrium as long as final producers make positive total profits in expectation.

Dynamic Model By combining (21) with (24), we can follow the same steps as in the static model to obtain the surplus that decentralizes the constrained efficient allocation:

$$\mathcal{X}(z) = x^\star(z) - (\mathcal{M}(N)N - 1)\, e^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[a(s(z)) - c(z)]. \tag{34}$$

This result allows us to derive the optimal transaction subsidy schedule in the dynamic context.

Proposition 8 (Optimal Subsidy Schedule in a Dynamic Economy)**.** *The Social Planner can decentralize the efficient allocation via a targeted transaction subsidy schedule $\tau_d^\star(z)$ such that the price of the transactions is given by $p^\star(z) = p(z) + \tau_d^\star(z)$, where*

$$\tau_d^\star(z) = e^{-\lambda\hat{\phi}(\underline{z},z)}T_d^\star \lesseqgtr 0, \tag{35}$$

and $T_d^\star \equiv (\mathcal{M}(N)N-1)\, e^{-\lambda\bar{\phi}}\hat{\mathbb{E}}[a(s(z))-c(z)] \gtrless 0$ *is a function that only depends on aggregates. The subsidy schedule is decreasing in* z *and bounded between* $[\tau_d^\star(\bar{z}), \tau_d^\star(\underline{z})] = [(1-f)^{-1}T_d^\star, T_d^\star]$. *Note that* $\tau_d^\star(z) > 0$ *represents a lump-sum transaction subsidy and* $\tau_d^\star(z) < 0$ *represents a lump-sum transaction tax.*

Proposition 8 shows that in the dynamic economy, the targeting principle is the same as in the static case. However, the tax schedule is shifted since, in general, $T^\star \neq T_d^\star$.

Overall, our results show that the planner would like to increase the price of all transactions. This is equivalent to giving more surplus to intermediate producers for them to have more *skin in the (final production) game*. The higher the surplus intermediate producers appropriate, the larger the opportunity cost of marginally increasing specialization – being the marginal benefit of higher input value unchanged. Hence, a suitable transaction subsidy can decentralize the efficient level of specialization. For the government budget to be balanced, the transaction subsidy can be financed by taxing final producers through a variety of instruments (such as lump-sum taxes and profits taxes), so that surplus is genuinely reallocated along the supply chain.

Finally, note that the possibility of decentralizing the efficient allocation by redistributing surplus across trading partners relates to the well-known *Hosios condition* in matching models (Hosios, 1990). However, the economic intuition behind the two results is different. In matching models, the inefficiency of the equilibrium comes from *congestion* and *thick market* effects governed by the shape of the matching function. Accordingly, efficient entry of sellers obtains when the surplus share accruing to them equals their marginal contribution to matching (*matching elasticity*). In our framework, there is no matching function (so the matching elasticity is zero), and the key source of inefficiency, that is, the network externality, arises from the complementarity among production inputs. Hence, efficient specialization obtains when the surplus share accruing to intermediate producers compensates for the foregone surplus from complementary inputs caused by higher specialization of the respective input. Mangin and Julien (2021) provide a framework to reconcile these two different economic forces through the notion of *generalized* Hosios condition, according to which the efficient surplus share equals the sum of matching elasticity *and surplus elasticity*. Hence, efficient specialization in our model obtains when the surplus share accruing to each intermediate producer equals the elasticity of social surplus with respect to its specialization.

5 Conclusions

In this paper, we study a model of frictional supply chain formation where heterogeneous intermediate producers choose the specialization of their goods. Higher specialization increases productivity but reduces the trading probability. We model supply chains as a complex production process in which multiple complementary inputs need to be sourced for final production to take place. This framework delivers a well-defined notion of supply chain resilience, *i.e.*, the time it takes for final producers to return to normal operations following disruptions. In the model,

supply chain resilience is negatively affected by the average specialization of goods, the extent of search frictions, and the complexity of the production process.

We study the efficiency properties of equilibrium specialization. We show that the combination of endogenous specialization and price-posting in frictional markets with simultaneous search imply that appropriability and business-stealing externalities perfectly offset each other. Still, intermediate producers fail to internalize the full reduction in social surplus induced by their specialization decisions due to a novel network externality arising from the interplay between frictional markets, endogenous specialization, and complexity in production. The network externality induces over-specialization in equilibrium. In turn, over-specialization implies that supply chains are less resilient than efficient. We characterize how a social planner can decentralize efficient supply chain resilience through a targeted transaction subsidy. The optimal subsidy schedule takes a simple form comprising an aggregate component and a targeting element that depends on the average trading probability of lower-productivity intermediate competitors.

This paper sheds light on the sources of inefficiencies that are likely to affect supply chain resilience. Our analytical results call for further work to provide quantification to the qualitative mechanisms we have highlighted.

References

Acemoglu, D. and Tahbaz-Salehi, A. (2024). The macroeconomics of supply chain disruptions. *Review of Economic Studies*, page rdae038.

Albrecht, J., Menzio, G., and Vroman, S. (2023). Vertical differentiation in frictional product markets. *Journal of Political Economy Macroeconomics*, 1(3):586–632.

Alessandria, G., Khan, S. Y., Khederlarian, A., Mix, C., and Ruhl, K. J. (2023). The aggregate effects of global and local supply chain disruptions: 2020–2022. *Journal of International Economics*, page 103788.

Antràs, P. (2003). Firms, contracts, and trade structure. *The Quarterly Journal of Economics*, 118(4):1375–1418.

Antràs, P. (2020). Conceptual aspects of global value chains. *The World Bank Economic Review*, 34(3):551–574.

Antràs, P. (2023). An 'Austrian' Model of Global Value Chains. NBER Working Papers 30901, National Bureau of Economic Research, Inc.

Antràs, P. and Chor, D. (2013). Organizing the global value chain. *Econometrica*, 81(6):2127–2204.

Antràs, P. and De Gortari, A. (2020). On the geography of global value chains. *Econometrica*, 88(4):1553–1598.

Antràs, P. and Helpman, E. (2006). Contractual frictions and global sourcing.

Arkolakis, C. (2010). Market penetration costs and the new consumers margin in international trade. *Journal of Political Economy*, 118(6):1151–1199.

Bajo, I., Bennhoff, F. H., and Ferrari, A. (2024). Not-so-cleansing recessions. *Working paper*.

Baldwin, R. and Venables, A. J. (2013). Spiders and snakes: Offshoring and agglomeration in the global economy. *Journal of International Economics*, 90(2):245–254.

Baptista, M., Gordon, M., Herrmann, C., and Pratt, A. (2021). With artificial intelligence, find new suppliers in days, not months. Mckinsey's operations practice report, McKinsey & Company.

Baqaee, D. R. (2018). Cascading Failures in Production Networks. *Econometrica*, 86(5):1819–1838.

Bar-Isaac, H., Caruana, G., and Cuñat, V. (2012). Search, design, and market structure. *American Economic Review*, 102(2):1140–1160.

Bar-Isaac, H., Caruana, G., and Cuñat, V. (2023). Targeted product design. *American Economic Journal: Microeconomics*, 15(2):157–186.

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

Benassy, J.-P. (1996). Taste for variety and optimum production patterns in monopolistic competition. *Economics Letters*, 52(1):41–47.

Boehm, J. and Oberfield, E. (2020). Misallocation in the market for inputs: Enforcement and the organization of production. *The Quarterly Journal of Economics*, 135(4):2007–2058.

Brakman, S. and Heijdra, B. J. (2001). *The monopolistic competition revolution in retrospect*. Cambridge University Press.

Burdett, K. and Judd, K. L. (1983). Equilibrium price dispersion. *Econometrica: Journal of the Econometric Society*, pages 955–969.

Capponi, A., Du, C., and Stiglitz, J. E. (2024). Are supply networks efficiently resilient? Working Paper 32221, National Bureau of Economic Research.

Carreras-Valle, M.-J. (2023). Increasing inventories: The role of delivery times.

Carreras-Valle, M. J. and Ferrari, A. (2025). The cost of delivery delays. *arXiv preprint arXiv:2501.08728*.

Carvalho, V., Elliott, M., and Spray, J. (2023). Network bottlenecks and market power. *Working Paper*.

Carvalho, V. M., Nirei, M., Saito, Y. U., and Tahbaz-Salehi, A. (2021). Supply chain disruptions: Evidence from the great east japan earthquake. *The Quarterly Journal of Economics*, 136(2):1255–1321.

Dhingra, S. and Morrow, J. (2019). Monopolistic competition and optimum product diversity under firm heterogeneity. *Journal of Political Economy*, 127(1):196–232.

Diamond, P. A. (1971). A model of price adjustment. *Journal of economic theory*, 3(2):156–168.

Dixit, A. K. and Stiglitz, J. E. (1975). Monopolistic competition and optimum product diversity, university of warwick. *Economic Research Paper*, 64.

Dixit, A. K. and Stiglitz, J. E. (1977). Monopolistic competition and optimum product diversity. *The American economic review*, 67(3):297–308.

Drentin, R., Khushalani, S., Kloos, C., and Misra, P. (2024). Where procurement is going next. Mckinsey's operations practice report, McKinsey & Company.

Ekerdt, L. and Wu, K.-J. (2023). The rise of specialized firms. *Working paper*.

Elliott, M., Golub, B., and Leduc, M. V. (2022). Supply network formation and fragility. *American Economic Review*, 112(8):2701–47.

Esala, L. (2024). Supply chain dynamics with search frictions. Working paper.

Ethier, W. J. (1982). National and international returns to scale in the modern theory of international trade. *The American Economic Review*, 72(3):389–405.

Ferrari, A. (2023). Inventories, demand shocks propagation and amplification in supply chains. *Working paper*.

Grossman, G. M., Helpman, E., and Lhuillier, H. (2023a). Supply chain resilience: Should policy promote international diversification or reshoring? *Journal of Political Economy, 131(12): 3462-3496.*

Grossman, G. M., Helpman, E., and Sabal, A. (2023b). Resilience in vertical supply chains. Working Paper 31739, National Bureau of Economic Research.

Haasbroek, M. (2025). Search frictions and international trade intermediation. *Working paper*.

Hart, O. and Moore, J. (1990). Property rights and the nature of the firm. *Journal of political economy*, 98(6):1119–1158.

Hosios, A. J. (1990). On The Efficiency of Matching and Related Models of Search and Unemployment. *The Review of Economic Studies*, 57(2):279–298.

Huang, H., Manova, K. B., Pisch, F., and Perelló, O. (2022). Firm heterogeneity and imperfect competition in global production networks. *Working Paper*.

Kim, S.-J. and Shin, H. S. (2023). Theory of supply chains: a working capital approach. BIS Working Papers 1070, Bank for International Settlements.

Klein, B., Crawford, R. G., and Alchian, A. A. (1978). Vertical integration, appropriable rents, and the competitive contracting process. *The journal of Law and Economics*, 21(2):297–326.

Kopytov, A., Mishra, B., Nimark, K., and Taschereau-Dumouchel, M. (2021). Endogenous production networks under supply chain uncertainty. *Available at SSRN 3936969*.

Krishna, V. (2009). *Auction theory*. Academic press.

Lenoir, C., Martin, J., and Mejean, I. (2022). Search frictions in international goods markets. *Journal of the European Economic Association*, 21(1):326–366.

Levine, D. K. (2012). Production chains. *Review of Economic Dynamics*, 15(3):271–282.

Mangin, S. and Julien, B. (2021). Efficiency in search and matching models: A generalized Hosios condition. *Journal of Economic Theory*, 193(C).

Mankiw, N. G. and Whinston, M. D. (1986). Free entry and social inefficiency. *The RAND Journal of Economics*, pages 48–58.

Martin, J., Mejean, I., and Parenti, M. (2020). Relationship stickiness, international trade, and economic uncertainty. *CEPR Discussion Paper No. DP15609*.

Martin, J., Mejean, I., and Parenti, M. (2023). Relationship stickiness, international trade, and economic uncertainty. *The Review of Economics and Statistics*, pages 1–45.

Meier, M. (2020). Supply chain disruptions, time to build, and the business cycle. Technical report, University of Bonn and University of Mannheim, Germany.

Menzio, G. (2023). Optimal Product Design: Implications for Competition and Growth Under Declining Search Frictions. *Econometrica*, 91(2):605–639.

Miroudot, S. (2020). *COVID-19 and trade policy: Why turning inward won't work*, chapter 9, pages 117–130. CEPR Press.

Nunn, N. (2007). Relationship-specificity, incomplete contracts, and the pattern of trade. *The quarterly journal of economics*, 122(2):569–600.

Oberfield, E. (2018). A theory of input - output architecture. *Econometrica*, 86(2):559–589.

Parenti, M., Ushchev, P., and Thisse, J.-F. (2017). Toward a theory of monopolistic competition. *Journal of Economic theory*, 167:86–115.

Pissarides, C. (2000). *Equilibrium Unemployment Theory, 2nd Edition*. The MIT Press.

Rauch, J. E. (1999). Networks versus markets in international trade. *Journal of international Economics*, 48(1):7–35.

Rogerson, W. P. (1992). Contractual solutions to the hold-up problem. *The Review of Economic Studies*, 59(4):777–793.

Spence, M. (1976). Product selection, fixed costs, and monopolistic competition. *The Review of economic studies*, 43(2):217–235.

Tirole, J. (1999). Incomplete contracts: Where do we stand? *Econometrica*, 67(4):741–781.

White House (2021). Biden-harris administration announces supply chain disruptions task force to address short-term supply chain discontinuities. Technical report, White House.

Williamson, O. E. (1979). Transaction-cost economics: the governance of contractual relations. *The journal of Law and Economics*, 22(2):233–261.

Zhelobodko, E., Kokovin, S., Parenti, M., and Thisse, J.-F. (2012). Monopolistic competition: Beyond the constant elasticity of substitution. *Econometrica*, 80(6):2765–2784.

Appendix

A Additional Details and Derivations

A.1 Static Model

Throughout the text, the expectation operator $\mathbb{E}[.]$ refers to the productivity distribution $\Gamma(z)$:

$$\mathbb{E}[X] \equiv \int_{\underline{z}}^{\bar{z}} X(z)\gamma(z)dz.$$

The expectation operator $\mathbb{E}_{\max\{\tilde{z}\}|\tilde{z}\leq z}$ refers to the productivity distribution of the highest-surplus-offering compatible intermediate producer with lower productivity than z:

$$\mathbb{E}_{\max\{\tilde{z}\}|\tilde{z}\leq z}[X] \equiv \int_{\underline{z}}^{z} X(\tilde{z})e^{-\lambda\hat{\phi}(\underline{z},z)}\lambda\phi(s(\tilde{z}))\gamma(\tilde{z})d\tilde{z}.$$

Finally, the expectation operator $\hat{\mathbb{E}}[.]$ refers to the productivity distribution of active matches:

$$\hat{\mathbb{E}}[X] \equiv \frac{\mathbb{E}_{\max\{z\}}}{f} = \int_{\underline{z}}^{\bar{z}} X(\tilde{z})\frac{e^{-\lambda\hat{\phi}(\underline{z},z)}\lambda\phi(s(\tilde{z}))\gamma(\tilde{z})}{1-e^{-\lambda\bar{\phi}}}d\tilde{z}.$$

Th first-order condition of the intermediate producer problem (6) with respect to the surplus offered to final producers reads:

$$\lambda\bar{\phi}G'(x)[a(s)-x-c(z)] = 1. \tag{36}$$

Let $\mathcal{H}(z)$ denote the equilibrium distribution of surplus offered to final producers, i.e., $\mathcal{H}(z) \equiv G(x^\star(z))$. Under the guess that $x^{\star\prime}(z) > 0$, the equilibrium distribution of surplus offered to final producers equals the distribution of productivity weighted by the compatibility probability function, i.e., $\mathcal{H}(z) = \hat{\phi}(\underline{z},z)/\bar{\phi}$. It follows that $G'(x(z)) = \frac{\mathcal{H}'(z)}{x'(z)}$ and $\mathcal{H}'(z) = \frac{\phi(s(z))\gamma(z)}{\bar{\phi}}$. Substituting for these expressions into (36) yields the differential equation characterizing the policy function $x^\star(z)$ reported in the main text (7). From that expression, it is apparent that $x^{\star\prime}(z) > 0$, which allows verifying the initial guess.

A.2 Specialization Function

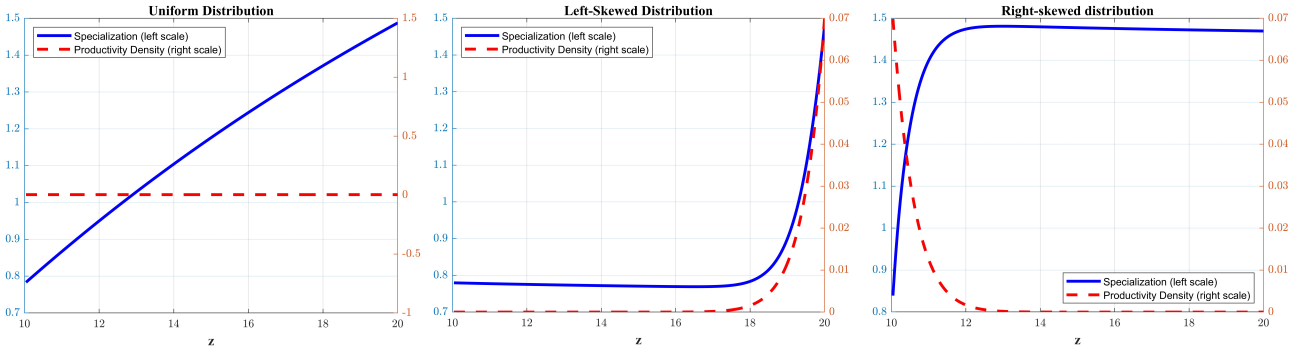


Figure A.1: Specialization function: the role of productivity distribution
Note: The three panels show the cross-sectional shape of the specialization function for different productivity distributions. The left panel considers a uniform distribution $z \sim U[10, 20]$, the middle panel a left-skewed shifted Beta distribution $z \sim \mathcal{B}(1, 18; 10, 20)$, the right panel a right-skewed shifted Beta distribution $z \sim \mathcal{B}(18, 1; 10, 20)$. All the other structural parameters are the same across panels.

Figure A.1 shows how the specialization function is influenced by the shape of the productivity distribution. With a uniform density (right panel), specialization is (approximately) linearly increasing in productivity. A left-skewed productivity distribution (middle panel) entails a nonmonotonic behavior of specialization, which is decreasing in productivity for low productivity levels and exponentially increasing thereafter. Symmetrically, a right-skewed productivity distribution (right panel) gives rise to a mirror-like, nonmonotonic behavior of specialization, which is strongly increasing in productivity for low productivity levels and then gradually decreasing. The qualitative differences across productivity distributions stress the importance of firm heterogeneity as a determinant of specialization choices.

A.3 Interpretation of Optimal Offered Surplus from Auction Theory

Our optimal offered surplus function is the solution to a first-price sealed-bid auction with an unknown number of bidders. For given specialization function, let $v(z)$ be the valuation of a successful match by an intermediate producer of productivity z. An intermediate producer with productivity z offering surplus x makes expected operating profits:

$$\begin{aligned}\Pi(x; z) &= \phi(z) \left(\sum_{k=0}^{\infty} m_k G(x)^k \right) f^{N-1}(v(z) - x) \\ &= \phi(z) f^{N-1} \theta \lambda \sum_{k=0}^{\infty} p_k G(x)^k (v(z) - x),\end{aligned}$$

where $p_k \equiv e^{-\lambda \bar{\phi}} \frac{(\lambda \bar{\phi})^k}{k!}$ denotes the probability that the number of other compatible intermediate producers contacted by a final producer is k, which follows a Poisson distribution. Setting the

first-order condition with respect to x to zero yields:

$$\sum_{k=0}^{\infty} p_k G(x)^k = (v(z) - x) \sum_{k=0}^{\infty} p_k k G(x)^{k-1} G'(x).$$

We follow the same steps as in the main text by guessing that the optimal offered surplus is monotonically increasing in productivity, i.e., $x'(z) > 0$, and defining the equilibrium offered surplus distribution as $H(z) = G(x(z))$. Substituting into the optimality condition yields:

$$x'(z) \sum_{k=0}^{\infty} p_k H(z)^k = (v(z) - x(z)) \sum_{k=0}^{\infty} p_k k H(z)^{k-1} H'(z).$$

Solving the differential equation with boundary condition $x(\underline{z}) = 0$, we obtain the standard expression for the optimal bidding strategy in first-price sealed-bid auctions with an unknown number of bidders (Krishna, 2009):

$$x^\star(z) = \frac{1}{\sum_{k=0}^{\infty} p_k(z) H^k(z)} \int_{\underline{z}}^{z} v(\tilde{z}) \sum_{k=0}^{\infty} p_k(\tilde{z}) dH^k(\tilde{z}) = \mathbb{E}[\max v(\tilde{z}) | \tilde{z} < z]. \tag{37}$$

Since $v(z) = a(s(z)) - c(z)$, it follows that the weighted productivity density $\gamma^\star(\tilde{z}, z)$ defined in the main text equals the productivity density of the highest-surplus-offering compatible intermediate producer with lower productivity than z, i.e., $\gamma^\star(\tilde{z}, z) \equiv \frac{\sum_{k=0}^{\infty} p_k dH^k(\tilde{z})}{\sum_{k=0}^{\infty} p_k H^k(z)}$.

Unlike standard auction theory, the distribution of the number of bidders is not taken as given in our model but pinned down by the extent of (exogenous) search frictions and (endogenous) compatibility frictions. Therefore, it is instructive to break down (37) further into a weighted average of states of the world where the final producer is able to locate lower-productivity competitors and where is not:

$$x^\star(z) = \Pr(m(z) = 0) \cdot 0 + \Pr(m(z) = 1) \cdot \mathbb{E}\left[\max\{a(s(\tilde{z})) - c(\tilde{z})\} \,\middle|\, \tilde{z} \leq z, m(z) = 1\right],$$

where $m(z) = \mathbb{1}\{\text{final producer contacts at least one firm w/ productivity } \leq z\}$. From this formulation, it is apparent how the optimal offered surplus equals the expected outside option of a final producer. With probability $\Pr(m(z) = 0) = e^{-\lambda \hat{\phi}(\underline{z}, z)}$, a final producer does not meet any compatible intermediate producer with productivity lower than z, in which case the final producer's outside option when meeting an intermediate producer with productivity z boils down to her reservation surplus, that is, zero. With complementary probability, the final producer meets at least one compatible intermediate producer with lower productivity than z, in which case the final producer's outside option when meeting an intermediate producer with productivity z equals the maximum surplus the best of such firms can afford offering, that is, the expected highest match value among firms with lower productivity than z.

A.4 Bargaining Model

Our baseline model of specialization posits that intermediate producers post the price of their goods (or, interchangeably, set the offered surplus) ex ante, *i.e.*, before meeting a final producer. In this section, we explore the alternative trading protocol based on ex post bargaining.

We study a model that is otherwise the same as our static baseline, but let the transaction price be determined by Nash bargaining upon meeting between a final producer and a compatible intermediate producer. We assume that final producers can observe the total surplus v from matching with any intermediate producer they meet. Let $\xi \in [0,1]$ denotes the exogenous bargaining power weight of final producers in surplus-splitting negotiations. Formally,

$$x^\star(v) = \underset{x}{\text{argmax}} \left(x - \omega(v;\vec{v}_k)\right)^{\xi} \left(v - x\right)^{1-\xi},$$

where $\omega(v;\vec{v}_k)$ denotes the outside option of a final producer which is contact with k compatible intermediate producers with less total surplus than v.

To align the final producer's outside option in the two models, we assume that, while negotiating with the highest-surplus compatible intermediate producer, final producers keep the option of bargaining with the other less productive compatible intermediate producers they are in contact with. According to the Nash bargaining protocol, the surplus accruing to the final producer therefore equals:

$$x^\star(v) = \xi v + (1-\xi)\omega(v;\vec{v}_k),$$

where the final producer's outside option in bargaining equals the total match surplus of the second best compatible intermediate producer, i.e., $\omega(v;\vec{v}_k) = \max\{v_1,\ldots,v_k\}$.[24]

According to this market structure, expected operating profits of intermediate producers with productivity z choosing total surplus v equal:

$$\begin{aligned}\Pi(v;z) &= \phi(s(v;z))\ f^{N-1}\left(\sum_{k=0}^{\infty} m_k\ \mathcal{G}\left(v\right)^k (1-\xi)\left(v - \mathbb{E}[\omega(v;\vec{v}_k)]\right)\right)\\ &= \theta\lambda\phi(s(v;z))f^{N-1}(1-\xi)\left(\sum_{k=0}^{\infty} e^{-\lambda\bar{\phi}}\frac{(\lambda\bar{\phi})^k}{k!}\mathcal{G}\left(v\right)^k \left(v - \mathbb{E}[\omega(v;\vec{v}_k)]\right)\right),\end{aligned}$$

where $\mathcal{G}(v)$ denotes the distribution of total surplus and $\mathbb{E}[\omega(v;\vec{v}_k)] \equiv \int_{\underline{v}}^{v} \tilde{v}\frac{k\mathcal{G}'(\tilde{v})\mathcal{G}(\tilde{v})^{k-1}}{\mathcal{G}(v)}d\tilde{v}$ denotes the expected final producer's outside option when bargaining with an intermediate producer of productivity z, while being in contact with k other compatible intermediate producers.

Optimal total surplus (or, interchangeably, optimal specialization) solves the following intermediate producer's problem:

$$v^\star(z) = \underset{v}{\text{argmax}}\ \ \Pi(v;z) - wq(s(v;z)).$$

[24]Since intermediate producers have zero outside option at the bargaining stage, the second best compatible intermediate producer is indifferent between not trading with the final producer and offering him the entire match surplus.

Substituting for $v(z) = a(s(z)) - c(z)$, equilibrium total surplus/specialization solves:

$$\theta\lambda\mathcal{P}(z;N)(1-\xi)\left[a'(s(z)) + \frac{\phi'(s(z))}{\phi(s(z))}\left(a(s(z)) - c(z) - \mathbb{E}[\omega(z)]\right)\right] = wq'(s(z)), \tag{38}$$

where we made use of the facts that $\sum_{k=0}^{\infty} e^{-\lambda\bar{\phi}}\frac{(\lambda\bar{\phi})^k}{k!}\frac{\partial\mathcal{G}(v)^k}{\partial v}\left(v - \mathbb{E}[\omega(v;\vec{v}_k)]\right) = \sum_{k=0}^{\infty} e^{-\lambda\bar{\phi}}\frac{(\lambda\bar{\phi})^k}{k!}k\frac{G'(v)}{G(v)}(v - \mathbb{E}[\omega(v;\vec{v}_k)]) = \sum_{k=0}^{\infty} e^{-\lambda\bar{\phi}}\frac{(\lambda\bar{\phi})^k}{k!}\mathcal{G}(v)^k\frac{\partial\mathbb{E}[\omega(v;\vec{v}_k)]}{\partial v}$ and $\sum_{k=0}^{\infty} e^{-\lambda\bar{\phi}}\frac{(\lambda\bar{\phi})^k}{k!}\mathcal{G}(v)^k\mathbb{E}[\omega(v;\vec{v}_k)] = e^{-\lambda\bar{\phi}(1-\mathcal{G}(v))}\mathbb{E}[\omega(v)]$ and $\left(\frac{\partial s(v;z)}{\partial v}\right)^{-1} = a'(s(z))$. Notice that the expected outside option $\mathbb{E}[\omega(z)]$ equals the surplus offered to final producers in our baseline model, i.e., $\mathbb{E}[\omega(z)] = \mathbb{E}_{\max\{\tilde{z}\}|\max\{\tilde{z}\}\leq z}[a(s(z)) - c(z)]$.

Comparing (38) with (8), we observe that the two coincide up to the intermediate producer's bargaining power weight multiplying the marginal benefit of specialization. Therefore, the bargaining model nests our baseline model if intermediate producers enjoy all the bargaining power, i.e., $\xi = 0$.

Since the planner is indifferent as to how the total surplus is split, efficient specialization is still determined by (11). This allows us to establish the following proposition.

Proposition A.1 (Efficiency of the Static Economy with Bargaining)**.** *The equilibrium is constrained efficient if and only if production is not complex, i.e. $N = 1$, and intermediate producers hold all bargaining power, i.e., $\xi = 0$.*

Since the network externality pushes towards over-specialization and bargaining frictions push towards under-specialization, there may be a bargaining power weight $\xi^\star$ such that the two externalities cancel out. However, notice that the bargaining power weight is independent of the trading counterpart, whereas the network externality is firm-specific. Hence, no bargaining power weight decentralizes the efficient allocation. If bargaining power weights depended on the productivity of the intermediate producer involved in the transaction, the schedule of bargaining power weights that implements the efficient allocation would be:

$$\xi^\star(z) = \frac{(N-1)\exp\{-\lambda\hat{\phi}(\underline{z},z)\}\hat{\mathbb{E}}[A(s(\tilde{z})) - c(\tilde{z})]}{A'(\mathcal{S}(z)) + \frac{\phi'(\mathcal{S}(z))}{\phi(\mathcal{S}(z))}\left([A(\mathcal{S}(z)) - c(z)] - \mathbb{E}_{\max\{\tilde{z}\}|\max\{\tilde{z}\}<z}[A(\mathcal{S}(\tilde{z})) - c(\tilde{z})]\right)}.$$

A.5 Dynamic Model

The first-order condition with respect to the surplus offered of the profit maximization problem of intermediate producers (18) is:

$$\sum_{t=1}^{\infty}\beta^t\frac{\partial\mathcal{D}_t(\boldsymbol{s}_{1,t},x)}{\partial x}[a(s_t) - x - c(z)] = \sum_{t=1}^{\infty}\beta^t\mathcal{D}_t(\boldsymbol{s}_{1,t},x)$$

where $\frac{\partial \mathcal{D}_t(\boldsymbol{s}_{1,t},x)}{\partial x} = \lambda\bar{\phi}G'(x)\sum_{v=1}^{t}(1-\delta)^{t-v}\theta_v\lambda\mathcal{P}(s_v,x)$. Evaluating the first-order condition in a stationary equilibrium yields:

$$\sum_{t=0}^{\infty}\beta^t\theta\lambda\mathcal{P}(s,x)\lambda\bar{\phi}G'(z)\sum_{v=0}^{t}(1-\delta)^{t-v}[a(s)-x-c(z)] = \sum_{t=0}^{\infty}\beta^t(1-\delta)^{t+1}\mathcal{D}_0$$
$$+\sum_{t=0}^{\infty}\beta^t\theta\lambda\mathcal{P}(s,x)\sum_{v=0}^{t}(1-\delta)^{t-v}.$$

Developing the infinite sums yields:

$$\theta\lambda\mathcal{P}(s,x)[\lambda\bar{\phi}G'(x)(a(s)-x-c(z))-1]\frac{\beta\delta}{(1-\beta)[1-\beta(1-\delta)]} = \frac{1-\delta}{1-\beta(1-\delta)}\mathcal{D}_0.$$

Letting $\mathcal{D}_0 = 0$ allows recovering the same first-order condition as in the static model (36):

$$\lambda\bar{\phi}G'(x)[a(s)-x-c(z)] = 1.$$

Hence, the equilibrium surplus offered to final producers in the dynamic model is the same as in the static model.

The equilibrium of the dynamic model in Section 3 is given by the following system of equations:

$$x^\star(z) = \mathbb{E}_{\max\{\tilde{z}\}|\tilde{z}\leq z}[a(s^\star(\tilde{z}))-c(\tilde{z})],$$
$$\mathcal{D}(z;N)\left[a'(s^\star(z))+\delta[1+\beta(1-\delta)]\frac{\phi'(s(z))}{\phi(s(z))}(a(s^\star(z))-c(z)-x^\star(z))\right] = wq'(s(z)),$$
$$1-\frac{\psi}{w} = Nm\bar{q}, \tag{39}$$
$$C = \mu(f;N)f^N\frac{1}{\delta}N\hat{\mathbb{E}}[a(s^\star(z))-c(z)]. \tag{40}$$

The equilibrium conditions represent the optimal surplus offered to final producers and optimal specialization, and the labor and goods market clearing conditions, respectively.

A.6 Dynamic Model with Link Destruction

We report here the expressions for the multipliers to the network and search externality in equation (31):

$$\chi_1(f;N) \equiv 1-\nu\frac{\rho(f;N)-\delta f\rho(f;N-2)}{\tilde{f}(f,\nu)} \in (0,1)$$

$$\begin{aligned}\chi_2(f;N) \equiv & \left(1+\frac{\delta(f\rho(f;N-1)-f^N)}{1+\delta f^N-\rho(f;N)}\right)\frac{\beta(1-\delta)}{1+\beta(1-\delta)}f^N N\frac{1}{\tilde{f}(f,\nu)}\Bigg[f^{N-1}-\frac{1-\delta(1-f^N)-\rho(f;N)}{(1-\delta)[f\rho(f;N-1)-f^N]}\\ & \left(\rho(f;N-1)-f^{N-1}\right)\Bigg]-\frac{\rho(f;N-1)-f^{N-1}}{\tilde{f}(f,\nu)}\left[\frac{\partial\nu_i}{\partial f_i}+(N-1)\frac{\partial\nu_i}{\partial f_j}\right]\\ = & \left(1+\frac{\delta(f\rho(f;N-1)-f^N)}{1+\delta f^N-\rho(f;N)}\right)\frac{\beta(1-\delta)}{1+\beta(1-\delta)}f^N N\frac{1}{\tilde{f}(f;\nu)}\left[f^{N-1}-\nu\frac{\mu}{1-\mu}(\rho(f;N-1)-f^{N-1})\right]\\ & +\frac{\rho(f;N-1)-f^{N-1}}{\tilde{f}(f,\nu)}\Bigg[1+\frac{\delta(f\rho(f;N-1)-f^N)}{1-\delta(1-f^N)-\rho(f;N)}+(N-1)\Bigg(\frac{\delta f^2\rho(f;N-2)-f^N}{f\rho(f;N-1)-f^N}\\ & +\frac{\delta(f\rho(f;N-1)-f^N)}{1-\delta(1-f^N)-\rho(f;N)}\Bigg)\Bigg]f\nu.\end{aligned}$$

The multiplier to the network externality $\chi_1(f;N) \in (0,1)$ discounts the network externality of the baseline dynamic model for the fact that a share ν of final searching producers need to source less than N inputs in the market. Indeed, as $\nu \to 0$, the network externality of the dynamic model with link destruction approaches that of the baseline dynamic model. The multiplier to the search externality features two distinct components. The first component is the exact counterpart of the search externality of the baseline dynamic model, i.e., $\frac{\beta(1-\delta)}{1+\beta(1-\delta)}f^N N$ – adjusted for (i) the different law of motion of μ – through the term $\left(1+\frac{\delta(f\rho(f;N-1)-f^N)}{1+\delta f^N-\rho(f;N)}\right)$ – and (ii) for the effect of μ on the share of active searching final producers – through the term $\frac{1}{\tilde{f}(f;\nu)}\left[f^{N-1}-\nu\frac{\mu}{1-\mu}(\rho(f;N-1)-f^{N-1})\right]$. The second component relates to the effect of the input finding probability on the probability that an attached final producer is active through its effect on the share of active searching final producers ν. Since the sign of $\delta f^2\rho(f;N-2)-f^N$ is a priori ambiguous (and likely negative), the sign of the multiplier of the search externality is ambiguous.

A.7 Strategies to mitigate the impact of supplier breakdowns

In this section we examine two key strategies that final producers can adopt to mitigate the impact of supplier breakdowns: maintaining inventories and sourcing inputs from multiple intermediate producers (multi-sourcing). Our goal is to determine whether these strategies can prevent firm-level disruptions caused by the inability of sourcing inputs. If they do, implementing the corresponding strategy would eliminate the network externality. Indeed, from the perspective of an intermediate producer, the probability that an attached final producer is active would be one regardless of the specialization of the other inputs.

A.7.1 Inventories

Suppose that active final producers can choose to purchase more units of each input than needed for current production, setting aside inventories as a buffer. If a supplier relationship breaks down, these inventories allow production to continue until a new supplier is found or until the

stock is depleted. We further assume that inputs sourced can be used either in the current period or in the next – otherwise they perish. This implies that final producers set aside either zero or one unit of input from an input provider. However, the analysis can be generalized to accommodate any finite number of periods after which inputs perish.

Active final producers choose optimal inventories by solving the following dynamic profits maximization problem:

$$\max_{\{I_j\}_{j=1}^N \in \{0,1\}^N} \underbrace{\mathcal{Q}_i - \sum_{j=1}^{N}(1+I_j)p_j}_{\pi_i\left(\{I_j^{(-1)}\}_{j=1}^N, \{I_j\}_{j=1}^N\right)} + \beta \mathbb{E}\left[\pi_i'\left(\{I_j\}_{j=1}^N\right)\right], \tag{41}$$

where $\mathcal{Q}_i = \sum_{j=1}^N \max\{A_j^{(-1)} I_j^{\star(-1)}, A_j\}$ is the final good quality and $I_j^{\star(-1)}$ denotes optimal inventories of input j chosen in the previous period.

To highlight the basic economic forces shaping the inventory decision, we now specialize on the case where $N = 2$. The optimal inventory policy function reads:

$$\begin{aligned}
&I_j^\star\left(\{A_s\}_{s=1}^N, p_j, \{I_s^\star\}_{s\neq j}^N\right) = 1 \iff \\
&p_j < \beta\delta_j \Bigg(\left[(1-\delta_{-j}) + \delta_{-j}\left(f_{-j} + (1-f_{-j})I_{-j}^\star\right)\right] f_j \Pr(A_j' \leq A_j)\mathbb{E}[A_j - A_j' | A' \leq A_j] \dots \\
&+ (1-\delta_{-j})(1-f_j)[A_{-j} + A_j - (1+I_{-j}^{\star\prime})p_{-j}] \dots \\
&+ \delta_{-j} f_{-j}(1-f_j)[\max\{A_{-j}I_{-j}^\star, A_{-j}'\} + A_j - (1+I_{-j}^{\star\prime})p_{-j}] \dots \\
&+ \delta_{-j}(1-f_{-j})(1-f_j)I_{-j}^\star[A_{-j} + A_j] \Bigg).
\end{aligned} \tag{42}$$

The marginal cost of an additional unit of inventories is the price of the input j set by the current input provider. The marginal benefit comprises four terms – all discounted by the probability that the supplier relationship with the current input provider breaks down, δ_j. First, if the final producer is able to replace the supplier relationship in the next period, she would use the unit of inventories to produce were the new input supplier of lower value than the previous one (second row). Second, if the final producer is not able to replace the supplier relationship in the next period and the other supplier relationship does not break up, she would use the unit of inventories to continue production (third row). Third, if the final producer is not able to replace the supplier relationship in the next period and the other supplier relationship gets replaced upon breaking up, she would use the unit of inventories to continue production (third row). Finally, if also the other supplier relationship breaks down and the final producer is not able to replace either of them in the next period, she would use the unit of inventories to continue production provided that she has set aside inventories of the other input, as well (fifth row).

Hence, inventories allows mitigating the impact of supplier breakdowns. However, if inputs are perishable, inventories cannot completely hedge final producers from disruptions. The reason

is that, with positive probability, the final producer will not be able to locate a compatible intermediate producer in $n < \infty$ periods.

If inputs were not perishable, a final producer may set aside a large number of units (approaching infinity) of inventories to completely hedge from disruptions. However, this strategy would not be profit-maximizing. As apparent from (42), the marginal benefit of holding inventories scales with the time-discounted probability of actually using such inventories. On the other hand, the marginal cost always equals the current price. It follows that there exists an $\tilde{n} < \infty$ such that, if $n > \tilde{n}$, the marginal benefit of the n-th unit of inventories approaches zero. Hence, holding infinite inventories cannot be a profit-maximizing strategy.

Overall, inventories do not completely hedge final producers from disruptions in equilibrium. It follows that intermediate producers will always face a probability lower than one that an attached final producer is able to source all the other complementary inputs. Since this probability is decreasing in specialization, the network externality is not eliminated by endogenous inventories.

A.7.2 Multi-sourcing

Suppose final producers can establish links with multiple compatible suppliers met during a supplier search. Specifically, assume that a final producer may choose to maintain an idle backup link with the second-best compatible supplier at a fixed per-period cost c_M. If the link with the best supplier breaks down first, the final producer can immediately switch to the backup supplier and continue production. This implies that active final producers are in contact with either one or two intermediate producers of a certain input. However, the analysis can be generalized to accommodate any finite number of suppliers per input.

Let $D_j(D_j^M)$ be an indicator of link destruction with the main (backup) provider of input j. Active final producers choose the optimal number of backup links by solving the following dynamic profits maximization problem:

$$\max_{M_j \in \begin{cases} \{0,1\} & \text{if } (D_j=1, N^c>1) \text{ or } (D_j^M=0, M_j^{\star(-1)}=1) \\ \{0\} & \text{else} \end{cases}, j=1,\ldots,N} \pi_i - c_M \sum_{j=1}^{N} M_j + \beta \mathbb{E}\left[\pi_i'\left(\{M_j\}_{j=1}^N\right)\right]. \quad (43)$$

In words, final producers make a supplier search when the supplier relationship with their main input provider breaks down ($D_j = 1$). Provided that they meet more than one compatible supplier ($N^c > 1$), they can establish a backup link with the second-best compatible supplier. If they established a backup link in a past supplier search ($M_j^{\star(-1)} = 1$), they can roll it over unless the link breaks for exogenous reasons ($D_j^M = 0$). To highlight the basic economic forces shaping the multi-sourcing decision, we now specialize on the case where $N = 2$. The optimal

multi-sourcing policy function reads:

$$
\begin{aligned}
&M_j^\star\left(\{x_s\}_{s=1}^N, \{x_s^M\}_{s=1}^N, c_M, \{M_s^\star\}_{s\neq j}^N\right) = 1 \iff \qquad (44)\\
&c_M < \beta\delta_j(1-\delta_j^M)\Bigg(\left[(1-\delta_{-j}) + \delta_{-j}\left(f_{-j} + (1-f_{-j})M_{-j}^\star\right)\right] f_j G(x_j^M)\mathbb{E}[x_j^M - x_j' | x' \leq x_j^M] \ldots \\
&+ (1-\delta_{-j})(1-f_j)\left[\pi_i(x_j^M, x_{-j}^M) - c_M \sum_j M_j^{\star\prime}\right] \ldots \\
&+ \delta_{-j} f_{-j}(1-f_j)\left[\pi_i(x_j^M, \max\{x_{-j}^M, x_{-j}'\}) - c_M \sum_j M_j^{\star\prime}\right] \ldots \\
&+ \delta_{-j}(1-f_{-j})(1-f_j)M_{-j}^\star\left[\pi_i(x_j^M, x_{-j}^M) - c_M \sum_j M_j^{\star\prime}\right]\Bigg).
\end{aligned}
$$

The optimality condition for the number of backup links shares many similarities with that of inventories (42). The marginal cost of an additional backup link is the per-period cost of maintaining an idle link alive, c_M. The marginal benefit comprises four terms – all discounted by the probability that the supplier relationship with the main input provider breaks down and that with the backup supplier does not, $\delta_j(1-\delta_j^M)$. First, if the final producer is able to replace the supplier relationship in the next period, she would upgrade the backup link to main supplier were the new input supplier of lower value, thus gaining the differential surplus offered by the two (second row). Second, if the final producer is not able to replace the supplier relationship in the next period and the other supplier relationship does not break up, she would resort to the backup link to continue production (third row). Third, if the final producer is not able to replace the supplier relationship in the next period and the other supplier relationship gets replaced upon breaking up, she would resort to the backup link to continue production (third row). Finally, if also the other supplier relationship breaks down and the final producer is not able to replace either of them in the next period, she would resort to the backup link to continue production provided that she has a backup link for the other input, as well (fifth row).

Hence, multi-sourcing allows mitigating the impact of supplier breakdowns. However, if alternative input providers can be located only when performing a supplier search, multi-sourcing cannot completely hedge final producers from disruptions. The reason is that, with positive probability, the final producer will not be able to locate more than one compatible intermediate producer in a supplier search.

If final producers could search for alternative suppliers in any period (*on-the-match* search), a final producer may build up a large number of backup links (approaching infinity) to completely hedge from disruptions. However, this strategy would not be profit-maximizing. As apparent from (44), the marginal benefit of an additional backup link scales with the time-discounted probability of actually using such links. On the other hand, the marginal cost always equals the link maintenance cost. It follows that there exists an $\tilde{n} < \infty$ such that, if $n > \tilde{n}$, the marginal benefit of the n-th backup link approaches zero. Hence, holding infinite backup links cannot be a profit-maximizing strategy.

Overall, multi-sourcing does not completely hedge final producers from disruptions in equilibrium. It follows that intermediate producers will always face a probability lower than one that an attached final producer is able to source all the other complementary inputs. Since this probability is decreasing in specialization, the network externality is not eliminated by endogenous multi-sourcing.

B Extensions

In this section, we extend our model along four dimensions. First, we study the problem of firms optimally choosing the level of complexity in their production process. Final producers choose the length of their supply chain or the number of complementary inputs in the production process. We show that, much like specialization choices, complexity is generally inefficient.

Second, we allow final producers to invest resources to reduce the likelihood of own disruptions, e.g., through prudential investment in redundancies. We show that prudential investment is generally inefficient due to a well-known hold-up problem.

Third, we consider a general quality function that allows for arbitrary substitutability or complementarity among complementary inputs. This extension brings about a further externality induced by specialization choices associated with complementarity in quality.

Finally, we note that our baseline model features arbitrage opportunities. Specifically, we have not imposed any condition that makes the relative mass of intermediate and final producers adjust to arbitrage away any difference in the expected returns from running either class of firms. In our final extension, therefore, we allow entrepreneurs to enter the market either as intermediate or final producers. This extension brings about a further externality induced by specialization choices related to entry decisions.

For simplicity, in all the extensions we assume steady-state behavior, i.e., intermediate producers use stationary demand as demand constraint (instead of its law of motion).

B.1 Endogenous Complexity

In this extension, we endogenize the problem of firms choosing the complexity of their production process. In particular, firms optimize over the number of key inputs in the production function. This approach is similar to the core idea of Oberfield (2018), Boehm and Oberfield (2020), Kopytov et al. (2021) and Kim and Shin (2023). For simplicity, in this extension, we treat N as a measure rather than an integer.

Static model. We allow final producers to choose the complexity of the production process to maximize expected profits:

$$N^\star = \underset{N \geq 1}{\operatorname{argmax}}\, f^N N \hat{\mathbb{E}}[x(z)],$$

where f is the probability of finding an input on the market, and $\hat{\mathbb{E}}[x(z)]$ is the expected surplus accruing to the final producer on each input, conditional on being active. Intuitively increasing complexity induces two opposite effects. On the one hand, adding inputs increases the value of output directly. On the other hand, it reduces the probability that the final good is produced altogether. The optimal level of complexity solves:

$$N^\star = \frac{1}{\ln(1/f)}. \tag{45}$$

Hence, complexity is increasing in the input finding probability f. Efficient complexity solves:

$$\mathcal{N} = \underset{N\geq 1}{\operatorname{argmax}} \; f^N N\hat{\mathbb{E}}[a(s(z)) - c(z)] + \psi \log\left(1 - Nm\bar{q}\right).$$

Relative to the problem of individual firms, the social planner problem features two differences. First, additional complexity is evaluated at the expected social surplus $\hat{\mathbb{E}}[a(s(z)) - c(z)]$ instead of the expected private surplus $\hat{\mathbb{E}}[x(z)]$. Second, the planner accounts for the additional specialization costs associated with additional inputs, taking the form of higher disutility from labor. Hence, efficient complexity is given by:

$$\mathcal{N} = \frac{1}{\ln(1/f)}\left(1 - \frac{w\ell}{y}\right). \tag{46}$$

where y is given by (10) and the labor share $w\ell/y$ represents the ratio between aggregate product design costs and value added. Comparing (45) with (46) allows us to establish the following proposition:

Proposition A.2 (Efficiency of the Endogenous Complexity Economy)**.** *The equilibrium is inefficient and features excess complexity, i.e., $N^\star > \mathcal{N}$, and lower sourcing capacity than efficient, i.e., $f(s^\star(\boldsymbol{z}))^{N^\star} < f(\mathcal{S}(\boldsymbol{z}))^{\mathcal{N}}$.*

Intuitively, final producers do not internalize the marginal product design cost induced by their complexity choice, thereby generating a *hold-up problem*. Choosing a higher level of complexity forces intermediate producers to invest more in product design. These costs are sunk at the time of price posting As a result, equilibrium sourcing capacity – our notion of "static resilience" – is lower than efficient.

Dynamic model. Equilibrium complexity is the same as in the static model and given by (45). Efficient complexity solves:

$$\mathcal{N} = \underset{N\geq 1}{\operatorname{argmax}} \; \mu(f,N)\left(\frac{f^{N-1}}{\delta} N\mathbb{E}[a(s(z)) - c(z)]\right) + \psi \log\left(1 - mN\bar{q}\right),$$

where $\mu(f,N) \equiv \frac{\delta}{f^N+(1-f^N)\delta}$ is the stationary share of searching final producers. Efficient complexity is given by:

$$\mathcal{N} = \frac{1}{\mu(f,N)\ln(1/f)}\left(1-\frac{w\ell}{y}\right). \tag{47}$$

Comparing (47) with (46), we observe that a further inefficiency arises in a dynamic setting. Following the same approach adopted for specialization, the marginal welfare effects of equilibrium complexity can be decomposed into two terms:

$$\left.\frac{\partial \mathcal{W}}{\partial N}\right|_{N=N^\star} \propto \underbrace{-\,w\ell/y}_{\substack{\text{hold-up}\\ \text{externality}}} \quad \underbrace{+\,(1-\mu)}_{\text{search externality}}. \tag{48}$$

First, a hold-up externality induces equilibrium complexity to be higher than efficient. The reason is that final producers fail to internalize the product design costs of each additional input, which are borne by intermediate producers before meeting. The surplus share of such product design costs is represented by the labor share. Second, final producers do not internalize the effect of their complexity choice on the equilibrium share of searching final producers. This search externality pushes the private marginal cost of complexity to exceed the social one. Rearranging (47) allows us to establish the following result:

Proposition A.3 (Efficiency of the Dynamic Endogenous Complexity Economy)**.** *The equilibrium is efficient if and only if:*

$$\begin{cases} N = \frac{1}{\mu} \\ LS = 1-\mu \end{cases}$$

where $\mu = \frac{\delta}{\delta+(1-\delta)f^N}$ *and* $LS = \frac{w\ell}{y}$*. The equilibrium features under-resilience if and only if* $LS^\star > 1-\mu^\star$*.*

In general, the equilibrium features both inefficient complexity and specialization.

B.2 Endogenous Robustness

So far, we have treated the disruption probability of final producers as an exogenous parameter. In this section, we allow final producers to make prudential investments in robustness aimed at reducing their disruption probability, *e.g.*, through redundancies. Investment in robustness comes at per-period cost $\kappa(r)$ per supply relationship, with $\kappa' > 0$, $\kappa'' < 0$. Specifically, we posit that $\delta = \delta(r)$, with $\delta' < 0$, $\delta'' > 0$. For analytical tractability, we assume that final producers choose their robustness policy before matching with their input providers.[25] We think

[25] More specifically, we assume that final producers commit to a robustness policy r at the search stage for the entire life of the firm before a disruption occurs. Implementing such a robustness policy requires paying a per-period cost $\kappa(r)$ per supplier relationship. If final producers could optimize their robustness policy once matched, intermediate producers would internalize the robustness response of downstream partners when choosing the surplus offered to final producers.

of robustness policies as the design and implementation of prudential strategies to reduce the exposure to exogenous shocks. Formally, final producers choose their robustness to maximize expected profits:

$$r = \underset{\tilde{r}}{\text{argmax}}\, f^N N \frac{\hat{\mathbb{E}}[x(z)] - \kappa(\tilde{r})}{\delta(\tilde{r})}, \tag{49}$$

where $\hat{\mathbb{E}}[x(z)]$ is the expected surplus accruing to the final producer on each input, conditional on being active. Equilibrium robustness solves:[26]

$$-\delta'(r)\frac{\hat{\mathbb{E}}[x(z)] - \kappa(r)}{\delta(r)} = \kappa'(r). \tag{50}$$

The privately optimal robustness trades off a lower disruption probability, $-\delta'$, evaluated at the expected present discounted value of each supplier relationship, $\frac{\hat{\mathbb{E}}[x(z)]-\kappa}{\delta}$, against higher private marginal costs, κ'. Importantly, note that firms discount the future with their endogenous disruption probability, δ. Under the assumption that final producers decide their robustness before matching, the expected surplus offered is constant across final producers, and so is their optimal robustness. Due to the presence of robustness costs, final producers need to break even with their share of the match surplus. As a result, the "reservation surplus", x_0, of final producers is positive and pinned down by the following expected break-even condition:

$$x_0 + (N-1)\hat{\mathbb{E}}[x(z)] = N\kappa(r). \tag{51}$$

Final producers need to source N key inputs. Whenever they meet with a potential supplier, they know that production will ultimately entail a $N\kappa(r)$ robustness cost. They are, therefore, willing to accept any surplus such that operating profits exceed that cost, where $(N-1)\hat{\mathbb{E}}[x(z)]$ represents the expected surplus they can obtain from the other $N-1$ suppliers. The reservation surplus acts as a boundary condition for the differential equation governing the equilibrium surplus offered to final producers (7). Hence, the new equilibrium surplus offered to final producers reads:

$$x^\star(z) = e^{-\lambda\hat{\phi}(\underline{z},z)}x_0 + \mathbb{E}_{\max\{\tilde{z}\}|\tilde{z}\leq z}[a(s(\tilde{z})) - c(\tilde{z})]. \tag{52}$$

Equations (51) and (52) jointly imply that the equilibrium reservation surplus offered to final producers is given by $x_0 = \kappa(r) - \frac{N-1}{N}\hat{\mathbb{E}}[\tilde{x}(z)]$, where $\tilde{x}(z) = \mathbb{E}_{\max\{\tilde{z}\}|\tilde{z}\leq z}[a(s(\tilde{z})) - c(\tilde{z})]$. Note that $x_0 < \kappa(r)$, suggesting that final producers are willing to take on negative net surplus from a single relation since they can cross-subsidize it with the other $N-1$ relationships. Up to the different formulation of the equilibrium surplus offered to final producers, equilibrium

[26] The second-order condition holds whenever $\kappa''(r) + 2\frac{(-\delta'(r))}{\delta(r)}\kappa'(r) > \frac{2\delta'^2(r)-\delta''(r)\delta(r)}{\delta(r)^2}\Big(\hat{\mathbb{E}}[x(z)] - \kappa(r)\Big)$.

specialization is determined by the counterpart of (21) with steady-state behavior:

$$\mathcal{D}(z;N)\left[a'(s^\star(z))+\frac{\phi'(s(z))}{\phi(s(z))}(a(s^\star(z))-c(z)-x^\star(z))\right]=wq'(s(z)), \tag{53}$$

Interestingly, since $\delta(r)$ discounts the marginal benefit of specialization in (53), equilibrium specialization is increasing in robustness.

Social planner problem. Next, we study the planner problem to evaluate the efficiency properties of equilibrium specialization and robustness decisions. The efficient specialization condition has the same formulation as in the model with exogenous robustness, up to the definition of total match surplus as net of robustness costs. However, the efficiency properties of equilibrium specialization differ significantly from the dynamic model with exogenous robustness due to the fact that the reservation surplus of final producers is not specific to each intermediate producer. Proposition A.4 identifies the reason why the mere presence of costs downstream makes the equilibrium surplus offered to final producers inefficient. Corollary 1 applies this general inefficiency result to the endogenous robustness economy.

Proposition A.4 (Efficiency Properties of Equilibrium Specialization with Costs Downstream)**.** *If final producers sustain other costs than input purchases, then the property of the equilibrium surplus offered to final producers of perfectly balancing appropriability and business-stealing externalities breaks.*

Corollary 1 (Special Case - Robustness Costs)**.** *Equilibrium firm-level specialization is always inefficient in the endogenous robustness economy due to the presence of robustness costs downstream.*

Intuitively, a social planner would like the reservation surplus to vary across the intermediate producers' productivity distribution for business-stealing and appropriability externalities to offset each other in the presence of robustness costs. Since the equilibrium reservation surplus is independent of productivity, misallocation of specialization always arises in the cross-section.

Efficient robustness solves the following social planner problem:

$$\mathcal{R}=\underset{r}{\text{argmax}}\ \ \mu(f,r)\frac{f^N}{\delta(r)}N\Big(\hat{\mathbb{E}}[a(s(z))-c(z)]-\kappa(r)\Big).$$

Notice that robustness costs are borne by active final producers, whose equilibrium mass is $\mu\frac{f^N}{\delta}=\frac{f^N}{\delta+(1-\delta)f^N}$. Efficient robustness is implicitly defined by:

$$-\delta'(\mathcal{R})\left(1-f^N\right)\mu(f,\mathcal{R})\frac{\hat{\mathbb{E}}[a(s(z))-c(z)-\kappa(r)]}{\delta(\mathcal{R})}=\kappa'(\mathcal{R}). \tag{54}$$

The socially optimal level of robustness trades off a lower disruption probability, $-\delta'$, evaluated at the expected present discounted social value of each supplier relationship, $\left(1-f^N\right)\mu\frac{\hat{\mathbb{E}}[a(s(z))-c(z)-\kappa]}{\delta}$, against the marginal robustness cost, κ'. Comparing (50) with (54), we observe

that equilibrium robustness is generally inefficient. Let the aggregate surplus share accruing to final producers be $\tilde{\xi} \equiv \frac{\hat{\mathbb{E}}[x(z)-\kappa]}{\hat{\mathbb{E}}[a(s(z))-c(z)-\kappa]}$. The marginal welfare effect of equilibrium robustness can be decomposed into two terms:

$$\left.\frac{\partial \mathcal{W}}{\partial r}\right|_{r=r^\star} \propto \underbrace{1-\tilde{\xi}}_{\substack{\text{hold-up}\\ \text{externality}}} \underbrace{-\left[1-(1-f^N)\mu\right]}_{\text{search externality}}. \tag{55}$$

First, a standard hold-up externality induces equilibrium investment in robustness to be lower than efficient, as the flow private marginal benefit, $\hat{\mathbb{E}}[x(z)-\kappa]$, falls short of the flow social marginal benefit, $\hat{\mathbb{E}}[a(s(z))-c(z)-\kappa]$. Second, firms do not internalize the effect on equilibrium tightness induced by their higher robustness. This search externality is evident when noting that the social planner discounts the marginal surplus at a higher rate, $\delta/[(1-f^N)\mu]$, than private firms do, δ. It follows that the search externality induces equilibrium investment in robustness to be higher than efficient. The nature of the search externality is akin to that arising in models of endogenous match destruction, where firms ignore the effect of their reservation productivity choice on equilibrium tightness. As in those models, a specific surplus sharing rule allows decentralizing the efficient allocation (Pissarides, 2000). Formally, efficient robustness obtains in equilibrium (for given aggregate specialization) if the aggregate surplus share accruing to final producers equals the equilibrium share of unattached final producers, *i.e.*, $\tilde{\xi} = (1-f^N)\mu$. Under this condition, the appropriability externality exactly offsets the search externality. Otherwise, the net effect of the hold-up and search externality is ambiguous.[27]

We now turn to determining efficient specialization. The social planner problem reads:

$$\max_{r,\ \hat{s}(z)} \frac{\delta(r)}{\delta(r)+(1-\delta(r))f^{N-1}f(\hat{s}(z))}\frac{f^{N-1}}{\delta(r)}\Bigg(\int \lambda\phi(\hat{s}(z))\exp\left\{-\lambda\hat{\phi}(z,\bar{z})\right\}\left[a(\hat{s}(z))-c(z)-\kappa(r)\right]$$
$$\gamma(z)dz+(N-1)\frac{f^{N-1}}{\delta(r)}f(\hat{s}(z))\mathbb{E}\left[a-c-\kappa(r)\right]\Bigg)+\psi\log\left(1-m\left((N-1)\bar{q}+\bar{q}(\hat{s}(z))\right)\right).$$

Efficient specialization of intermediate producers with productivity z is given by:

$$\theta\frac{f^{N-1}}{\delta(r)}\left(\frac{\partial\mathbb{E}\left[a-c-\kappa(r)\right]}{\partial s(z)}-(N-1)\frac{-f'(s(z))}{f}\mathbb{E}\left[a-c-\kappa(r)\right]\right)+\theta\frac{f^{N-1}}{\delta(r)}\frac{-f'(s(z))}{f}$$
$$\frac{(1-\delta(r))f^N}{\delta+(1-\delta(r))f^N}N\mathbb{E}[a-c-\kappa(r)]=\frac{\psi}{1-mN\bar{q}}q'(s(z)). \tag{56}$$

The only difference with the baseline dynamic model lies in the definition of total match surplus *net* of robustness costs. Evaluating the marginal welfare effect of specialization at the equilibrium

[27] Standard models of endogenous match destruction imply that either the equilibrium match destruction rate is efficient or lower than efficient. The reason why, in our model, match destruction can be excessive is the presence of the hold-up externality potentially countervailing the search externality.

specialization yields:

$$\left.\frac{\partial \mathcal{W}}{\partial s(z)}\right|_{s(z)=s^\star(z)} \propto \underbrace{\mathbb{E}\left[a(s(\tilde{z}))-c(\tilde{z})-\kappa\right]}_{\text{business stealing externality}} \underbrace{-\,x(z)}_{\substack{\text{appropriability}\\ \text{externality}}} \underbrace{-\,(N-1)e^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[a-c-\kappa]}_{\text{network externality}}$$
$$\underbrace{+\;\frac{(1-\delta(r))f^N}{\delta(r)+(1-\delta(r))f^N}e^{-\lambda\hat{\phi}(\underline{z},z)}N\hat{\mathbb{E}}[a-c-\kappa]}_{\text{search externality}}. \tag{57}$$

The only substantive difference with the dynamic model in terms of efficiency properties of equilibrium specialization is that the perfect balance between appropriability and business-stealing externalities breaks down. The reason is that the reservation surplus is not targeted to intermediate producers. Formally, the equilibrium surplus offered to final producers would keep balancing appropriability and business-stealing externalities if and only if:

$$x_0=(1-e^{-\lambda\hat{\phi}(\underline{z},z)})\kappa \iff \kappa=e^{\lambda\bar{\phi}(\underline{z},z)}\frac{N-1}{N}\hat{\mathbb{E}}[\tilde{x}(z)] \quad \forall z\in[\underline{z},\bar{z}].$$

Since this condition can never hold simultaneously for any productivity level, the net effect of appropriability and business-stealing externalities is ambiguous. This entails that firm-level equilibrium specialization is always inefficient.

We summarize our findings in the next proposition.

Proposition A.5 (Efficiency of Endogenous Robustness Equilibrium)**.** *The equilibrium can feature both under- and over-robustness. For given aggregate specialization, equilibrium robustness is efficient if and only if:*

$$\tilde{\xi}=(1-f^N)\mu\ ,$$

where $\mu\equiv\frac{\delta(r)}{\delta(r)+(1-\delta(r))f^N}$ *and* $\tilde{\xi}\equiv\frac{\mathbb{E}[x-\kappa]}{\mathbb{E}[A-c-\kappa]}$*. Since equilibrium specialization is increasing in robustness, over-robustness makes over-specialization more likely.*
The equilibrium can feature both under- and over-specialization. For given robustness, equilibrium firm-level specialization is always inefficient.

In general, the equilibrium features inefficient investment in both robustness and specialization.

B.3 Generalized Quality Function

So far, we have looked at an economy in which final producers operate a specific technology. In particular, we assumed that intermediate goods are i) perfect substitutes in quality (intensive margin) and ii) perfect complements in quantity (extensive margin). While both assumptions matter for the results derived so far, only ii) is crucial for our new network externality to arise. The intuitive reason is that, although goods are complements, an intermediate producer only cares about losing its own share of the surplus if production does not take place. As a consequence, a final producer does not have an instrument to induce bilaterally efficient specialization choices (in a way akin to the inefficiencies induced by limited liabilities). Property

ii) is at the heart of this result. If goods are not complementary at the extensive margin, then no network externality arises, and the economy would be statically (constrained) efficient. Note that the same extensive-margin complementarity is underlying some of the results in Elliott et al. (2022) and Acemoglu and Tahbaz-Salehi (2024).

We now show that our assumption of perfect substitutability in quality actually induces fewer inefficiencies than a more general quality function would. In particular, suppose that the quality function of a final producer i takes the following generic form:

$$\mathcal{Q}_i = \chi(A_1, \ldots, A_N),$$

where χ is a constant-return-to-scale aggregator, which is increasing and concave in each argument. Importantly, we impose $\frac{\partial \chi}{\partial A_j} > 0$, $\frac{\partial^2 \chi}{\partial A_j^2} < 0 \; \forall j$. Hence, we allow for arbitrary substitutability or complementarity in quality (intensive margin), while retaining complementarity in quantity (extensive margin).

Equilibrium. We analyze the equilibrium of the dynamic economy. Intermediate producers maximize expected operating profits net of product design costs:

$$V(s, x; z) = \max_{s,x} = \mathcal{D}(s, x) \left(\tilde{A}(s; \mathbf{s}_{-i}) - x \right) - wq(s).$$

where $\mathcal{D}(s, x)$ is given by (20) and $\tilde{A}(s; \mathbf{s}_{-i}) \equiv \frac{\partial \chi}{\partial A_i} A_i$. Notice that the only difference with respect to the baseline dynamic model is that match surplus potentially depends on the specialization vector of all the other input providers $\mathbf{s}_{-i}$. For a generic CES aggregator with substitution parameter $\sigma \in [0, 1]$, it holds that $\tilde{A}_i = A_i^{\sigma} \chi^{1-\sigma}$, which nests our baseline model for $\sigma = 1$. Up to the generalized definition of $\tilde{A}$ instead of A, all the equilibrium conditions are the same as in the baseline dynamic model and given by

$$\theta \frac{f^{N-1}}{\delta} \left(\frac{\partial \mathbb{E}[\tilde{A}(\boldsymbol{z})]}{\partial s(z)} - (\mu N - 1) \frac{-f'(s(z))}{f} \mathbb{E}[\tilde{A}(\boldsymbol{z})] \right) = \frac{\psi}{1 - mN\bar{q}} q'(s(z)) - (N-1)\theta \frac{f^{N-1}}{\delta} \mathbb{E}\left[\frac{\partial \tilde{A}_{-i}}{\partial s(z)}(\boldsymbol{z}) \right].$$

Social planner problem The social planner problem reads:

$$\max_{s_i(z)} C + \psi \log(1 - \ell) \quad \text{s.t.} \quad C = y \text{ and } \ell = Nm\bar{q}.$$

Efficient specialization of intermediate producers with productivity z producing input i solves:

$$\begin{aligned} s_i(z) = \underset{\hat{s}(z)}{\operatorname{argmax}} \; & \frac{\delta}{\delta + (1-\delta) f^{N-1} f(\hat{s}(z))} \left(\frac{f^{N-1}}{\delta} \mathbb{E}\left[\tilde{A}\left(\hat{s}(z), \mathbf{s}_{-i} \right) \right] \right. \\ & \left. + \sum_{j=1}^{N-1} \frac{f^{N-2}}{\delta} f(\hat{s}(z)) \mathbb{E}\left[\tilde{A}\left(s_j(z); \hat{s}(z), \mathbf{s}_{-(i,j)} \right) \right] \right) \\ & + \psi \log\left(1 - m\left((N-1)\bar{q} + \bar{q}(\hat{s}(z)) \right) \right). \end{aligned}$$

The first line represents the expected match surplus for the input of interest. The second line represents the expected match surplus of the other inputs than the one of interest. The third line represents the disutility of labor supply.

We observe that a further inefficiency arises if inputs are not perfect substitutes at the intensive margin. The reason is that intermediate producers do not internalize the effect of their specialization on the match surplus generated by other input providers:

$$\left.\frac{\partial \mathcal{W}}{\partial s(z)}\right|_{\mathcal{S}(z)=s^\star(z)} \propto \underbrace{\mathbb{E}_{\max\{\tilde{z}\}|\tilde{z}\leq z}\left[\tilde{A}(s(\tilde{z}))\right]}_{\text{business-stealing externality}} \underbrace{-\, x(z)}_{\substack{\text{appropriability}\\ \text{externality}}} \underbrace{-\,(N-1)e^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[\tilde{A}(s(z))]}_{\text{network externality}} \tag{58}$$

$$\underbrace{+\; N\frac{(1-\delta)f^N}{\delta+(1-\delta)f^N}e^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[\tilde{A}(s(z))]}_{\text{search externality}} \underbrace{+\; \frac{(N-1)e^{\lambda\bar{\phi}(z,\bar{z})}}{\lambda(-\phi'(s(z)))}\mathbb{E}\left[\frac{\partial \tilde{A}_{-i}}{\partial s(z)}\right]}_{\text{production externality}}.$$

The new production externality pushes towards equilibrium under-specialization. Hence, under the more general aggregator χ, the economy features a further inefficiency: individual specialization influences not only the trading probability of other input providers but also their match surplus. We characterize the efficiency properties of the economy with generalized quality function in the next Proposition (A.6).

Proposition A.6 (Efficiency of the Dynamic Economy: Generalized quality function)**.** *The economy with generalized quality function is efficient if and only if:*

$$-\left(\mu N-1\right)f'(s(z))\hat{\mathbb{E}}[\tilde{A}(s(z))]-(N-1)\mathbb{E}\left[\frac{\partial \tilde{A}_{-i}}{\partial s(z)}\right]=0 \quad \forall z\in[\underline{z},\bar{z}].$$

If $N<\frac{1}{\mu}$, then the economy necessarily features under-specialization.

To sum up, when inputs are complements at the extensive margin, equilibrium over-specialization is more likely the more the production process is complex (high N) and the more substitutable input quality is (the closer σ is to 1 with CES quality function).

B.4 Free Entry

Until now, we have assumed that the mass of intermediate producers m is fixed, that is, a perfectly inelastic entry margin. We now relax this assumption by positing that m is pinned down by a no-arbitrage condition between operating an intermediate or a final firm.

Equilibrium. The equilibrium is the same as in the baseline dynamic model, up to the addition of the no-arbitrage condition governing the equilibrium mass of intermediate producers.

The no-arbitrage condition reads:[28]

$$\frac{f^{N-1}}{\delta}N\mathbb{E}[x(z)] = \theta\frac{f^{N-1}}{\delta}N\mathbb{E}[a(s(z)) - x(z) - c(z)] - w\bar{q}. \tag{59}$$

The no-arbitrage condition (59), the tightness definition (17) and labor market clearing condition (39) jointly pin down the equilibrium mass of intermediate producers as the result of a second-order equation:

$$\begin{aligned} m =& \frac{(N\mathbb{E}[x(z)] + N\bar{q}\mu N\mathbb{E}[a(s(z)) - x(z) - c(z)] + \psi\bar{q}) \pm \sqrt{\Delta}}{2N\bar{q}N\mathbb{E}[x(z)]}, \\ \Delta =& (N\mathbb{E}[x(z)] + N\bar{q}\mu N\mathbb{E}[a(s(z)) - x(z) - c(z)] + \psi\bar{q})^2 - 4N\bar{q}N\mathbb{E}[x(z)]\mu N\mathbb{E}[a(s(z)) - x(z) - c(z)], \end{aligned} \tag{60}$$

Since a higher mass of intermediate producers is associated with lower leisure, the two solutions can be Pareto-ranked according to social welfare. Therefore, we assume that the mass of firms equals the lowest solution to (60). The equilibrium is therefore fully characterized by equations (7),(21),(39),(40) and (60).

Social planner problem. We consider the problem of a social planner who is constrained by the free entry condition (59). Since the free entry condition pins down m as a function of the surplus sharing rule, we formulate the social planner problem in terms of optimal choice of the surplus offered to final producers, $x(z)$, taking the equilibrium specialization function as given. The social planner problem reads:

$$\max_{x(z)\leq a(s(x(z)))-c(z)} \quad \mu\frac{f^{N-1}}{\delta}N\mathbb{E}[a(s(x)) - c(z)] + \psi\log(1 - Nm\bar{q})$$

$$\text{s.t.} \quad m(s,x) = \frac{(N\mathbb{E}[x(z)] + N\bar{q}\mu N\mathbb{E}[a(s(x)) - x(z) - c(z)] + \psi\bar{q}) - \sqrt{\Delta}}{2N\bar{q}N\mathbb{E}[x(z)]} \tag{61}$$

$$\frac{\mu\lambda}{m(s,x)}\mathcal{P}(z)\left[A'(s(x)) + \frac{\phi'(s)}{\phi(s)}(a(s(x)) - c(z) - x(z))\right] = \frac{\psi}{1 - Nm\bar{q}}q'(s), \tag{62}$$

where the constraints (61) and (62) follow from solving (59) for m and from the equilibrium specialization for given surplus, respectively. We proceed in three steps. First, we compute the full derivative of specialization with respect to the surplus offered to final producers via the implicit function theorem from (62). Second, we compute the derivative of the mass of firms with respect to s and x from (61).

The full derivative of specialization $s(x(z))$ with respect to the surplus offered to final

[28] In the formulation of eq. (59) we take a stand on the organization of intermediate production being structured along single firms producing all inputs through distinct product lines.

producers $x(z)$ reads:

$$\frac{\partial s(x)}{\partial x} = \frac{-\lambda\theta(-\phi'(s))e^{-\lambda\hat{\phi}(z,\bar{z})}\frac{f^{N-1}}{\delta} + \frac{\psi}{(1-Nm\bar{q})^2}\frac{\partial m}{\partial x}\frac{q'(s)}{m}}{\frac{\lambda f^{N-1}}{\delta}\left[\frac{\partial\theta}{\partial s}e^{-\lambda\hat{\phi}(z,\bar{z})} + \theta\lambda(-\phi'(s))\gamma(z)\right][a'(s)\phi(s) + \phi'(s)(a(s) - x(z) - c(z))]\ldots} \\ \overline{-\frac{\psi}{(1-Nm\bar{q})^2}\left(\frac{\partial m}{\partial s}\frac{1}{m} + mN\frac{\partial\bar{q}}{\partial s}\right)q'(s) + SOC(z)}. \tag{63}$$

The derivatives of the equilibrium mass of intermediate producers, $m(s(x(z)), x(z))$ with respect to $s(x(z))$ and $x(z)$ read:

$$\frac{\partial m(s,x)}{\partial s} = \frac{\left[\frac{\partial\mathbb{E}[x]}{\partial s} + \frac{\partial\bar{q}}{\partial s}\mu N\mathbb{E}[a-x-c] + N\bar{q}\left(\frac{\partial\mu}{\partial s}N\mathbb{E}[a-x-c] + \mu\frac{\partial\mathbb{E}[a-x-c]}{\partial s}\right) + \frac{\psi}{N}\frac{\partial\bar{q}}{\partial s} - \frac{\Delta^{-\frac{1}{2}}}{2}\frac{\partial\Delta}{\partial s}\right]\ldots}{2(N\bar{q}N\mathbb{E}[x])^2} \\ \frac{N\bar{q}N\mathbb{E}[x] - (N\mathbb{E}[x] + N\bar{q}\mu N\mathbb{E}[a-x-c] + \psi\bar{q})\left[\frac{\partial\bar{q}}{\partial s}N\mathbb{E}[x] + N\bar{q}\frac{\partial\mathbb{E}[x]}{\partial s}\right]}{}, \tag{64}$$

$$\frac{\partial m(s,x)}{\partial x} = \frac{\left[\frac{\partial\mathbb{E}[x]}{\partial x} + N\bar{q}\mu\frac{\partial\mathbb{E}[a-x-c]}{\partial x} - \frac{\Delta^{-\frac{1}{2}}}{2}\frac{\partial\Delta}{\partial x}\right]N\bar{q}N\mathbb{E}[x] - (N\mathbb{E}[x] + N\bar{q}\mu N\mathbb{E}[a-x-c] + \psi\bar{q})N\bar{q}\frac{\partial\mathbb{E}[x]}{\partial x}}{2(N\bar{q}N\mathbb{E}[x])^2}, \tag{65}$$

where

$$\frac{\partial\Delta(s,x)}{\partial s} = \Big(N\mathbb{E}[x] + N\bar{q}\mu N\mathbb{E}[a-x-c] + \psi\bar{q}\Big)\Big(\frac{\partial\mathbb{E}[x]}{\partial s} + \frac{\partial\bar{q}}{\partial s}\mu N\mathbb{E}[a-x-c] + N\bar{q}\Big(\frac{\partial\mu}{\partial s}N\mathbb{E}[a-x-c]\ldots \\ +\mu\frac{\partial\mathbb{E}[a-x-c]}{\partial s}\Big) + \frac{\psi}{N}\frac{\partial\bar{q}}{\partial s}\Big) - 4\Big(\Big[\frac{\partial\bar{q}}{\partial s}N\mathbb{E}[x] + N\bar{q}\frac{\partial\mathbb{E}[x]}{\partial s}\Big]\mu N\mathbb{E}[a-x-c] + N\bar{q}N\mathbb{E}[x]\Big[\frac{\partial\mu}{\partial s}N\ldots \\ \mathbb{E}[a-x-c] + \mu\frac{\partial\mathbb{E}[a-x-c]}{\partial s}\Big]\Big),$$

$$\frac{\partial\Delta(s,x)}{\partial x} = \Big(N\mathbb{E}[x] + N\bar{q}\mu N\mathbb{E}[a-x-c] + \psi\bar{q}\Big)\Big(\frac{\partial\mathbb{E}[x]}{\partial x} + N\bar{q}\mu\frac{\partial\mathbb{E}[a-x-c]}{\partial x}\Big) - 4N\bar{q}\mu N\Big(\frac{\partial\mathbb{E}[x]}{\partial x}\ldots \\ N\mathbb{E}[a-x-c] + N\mathbb{E}[x]\frac{\partial\mathbb{E}[a-x-c]}{\partial x}\Big).$$

Finally, we maximize social welfare with respect to $x(z)$, which yields:

$$\theta\frac{f^{N-1}}{\delta}\left(\frac{\partial\mathbb{E}[a-c]}{\partial s(x(z))} - (\mu N-1)\frac{-f'(s(x(z)))}{f}\mathbb{E}[a-c]\right) - \frac{\psi}{1-Nm\bar{q}}\bar{q}' \tag{66}$$
$$= \underbrace{\frac{\psi}{1-Nm\bar{q}}\left(\frac{\partial m(s(x(z)),x(z))}{\partial s(x(z))} + \frac{\partial m(s(x(z)),x(z))/\partial x(z)}{\partial s(x(z))/\partial x(z)}\right)\frac{N\bar{q}}{m}}_{\text{entry externality}}$$

The first line equals the efficient specialization condition of the baseline dynamic model (24).[29] The second line highlights the new source of inefficiency related to the external effect of individual specialization on the equilibrium mass of intermediate producers. Comparing (66) with (24), we

[29] Since $\left(\frac{\partial s(x(z))}{\partial x(z)}\right)^{-1} = \frac{\partial x(s(z))}{\partial s(z)}$ by the chain rule, maximizing with respect to $x(z)$ by taking $s(x(z))$ as given is indeed isomorphic to maximizing with respect to $s(z)$ by taking $x(z)$ as given.

observe that the former nests the latter for fixed m. The additional inefficiency term brought about by free entry crucially depends on sign $\left\{\frac{dm}{ds}\right\}$: if it is positive, the entry externality pushes towards equilibrium over-specialization; if negative, towards under-specialization. Either sign is possible in equilibrium. In any case, Proposition (A.7) shows that equilibrium specialization is inefficient whenever it affects entry decisions.

Proposition A.7 (Efficiency of the Dynamic Economy: Free entry)**.** *The economy with free entry is always inefficient.*

C Proofs

Proof of Lemma 1. To establish the result, we solve the differential equation (7) characterizing the optimal surplus offered to final producers as follows:

$$x^{\star\prime}(z) + \lambda\phi(s(z))\gamma(z)x^\star(z) = \lambda\phi(s(z))\gamma(z)[a(s^\star(z)) - c(z)]$$

$$e^{\int_{\underline{z}}^{z}\lambda\phi(s(\hat{z}))\gamma(\hat{z})d\hat{z}}\left[x^{\star\prime}(z) + \lambda\phi(s(z))\gamma(z)x^\star(z)\right] = e^{\int_{\underline{z}}^{z}\lambda\phi(s(\hat{z}))\gamma(\hat{z})d\hat{z}}\lambda\phi(s(z))\gamma(z)[a(s^\star(z)) - c(z)]$$

$$\frac{\partial}{\partial z}\left[e^{\int_{\underline{z}}^{z}\lambda\phi(s(\hat{z}))\gamma(\hat{z})d\hat{z}}x^\star(z)\right] = e^{\int_{\underline{z}}^{z}\lambda\phi(s(\hat{z}))\gamma(\hat{z})d\hat{z}}\lambda\phi(s(z))\gamma(z)[a(s^\star(z)) - c(z)]$$

$$x^\star(z) = e^{-\int_{\underline{z}}^{z}\lambda\phi(s(\hat{z}))\gamma(\hat{z})d\hat{z}}\int_{\underline{z}}^{z} e^{\int_{\underline{z}}^{\tilde{z}}\lambda\phi(s^\star(\hat{z}))\gamma(\hat{z})d\hat{z}}\lambda\phi(s(\tilde{z}))[a(s^\star(\tilde{z})) - c(\tilde{z})]\gamma(\tilde{z})d\tilde{z}$$

$$x^\star(z) = \int_{\underline{z}}^{z} e^{-\int_{\tilde{z}}^{z}\lambda\phi(s(\hat{z}))\gamma(\hat{z})d\hat{z}}\lambda\phi(s(\tilde{z}))[a(s^\star(\tilde{z})) - c(\tilde{z})]\gamma(\tilde{z})d\tilde{z}$$

$$x^\star(z) = \mathbb{E}_{\max\{\tilde{z}\}|\tilde{z}\leq z}[a(s^\star(\tilde{z})) - c(\tilde{z})],$$

where the last step follows from the definition of the expectation operator reported in Section (A.1) (see Section (A.3) for an interpretation). ■

Lemma A.1 (Comparative Statics: market share reallocation from lower search frictions)**.** *Lower search frictions ($\lambda\uparrow$) increases conditional match surplus per input for given specialization function by reallocating market shares towards more productive firms.*

Proof. Consider the direct effect of an increase in λ on conditional match surplus per input:

$$\begin{aligned}
&\frac{d}{d\lambda}\int\frac{\lambda\phi(s(z))\exp\{-\lambda\hat{\phi}(z,\bar{z})\}\gamma(z)(a(s(z)) - c(z))dz}{\int\lambda\phi(s(\tilde{z}))\exp\{-\lambda\hat{\phi}(\tilde{z},\bar{z})\}\gamma(\tilde{z})d\tilde{z}}\\
&= \int\frac{(a(s(z)) - c(z))\gamma^\star(z,\bar{z})}{f}\left[-\hat{\phi}(z,\bar{z}) + \int\hat{\phi}(\tilde{z},\bar{z})\frac{\gamma^\star(z,\bar{z})}{f}\right]dz\\
&= \frac{1}{f^2}\left(\mathbb{E}[\hat{\phi}(z,\bar{z})]\mathbb{E}[a(s(z)) - c(z)] - f\mathbb{E}[\hat{\phi}(z,\bar{z})\,(a(s(z)) - c(z))]\right)\\
&= \frac{1}{f^2}\Big((1-f)\mathbb{E}[\hat{\phi}(z,\bar{z})]\mathbb{E}[a(s(z)) - c(z)] - f\mathrm{Cov}[\hat{\phi}(z,\bar{z}),(a(s(z)) - c(z))]\Big)\\
&> 0,
\end{aligned}$$

where the last step follows from observing that $\text{Cov}[\hat{\phi}(z,\bar{z}),(a(s(z))-c(z))]<0$ as $\partial[\hat{\phi}(z,\bar{z})]/\partial z<0$ and $\partial[a(s(z))-c(z)]/\partial z>0$. ■

Proof of Proposition 1. To establish the claim, we evaluate the marginal welfare effect of specialization at the equilibrium solution, as implicitly defined by equation (8). By definition, efficient specialization, as implicitly defined by equation (11), equalizes the marginal welfare effect of specialization to zero, i.e., $\frac{\partial \mathcal{W}}{\partial s(z)}\big|_{s(z)=\mathcal{S}(z)}=0$.

$$\begin{aligned}
\frac{\partial \mathcal{W}}{\partial s(z)}\bigg|_{s^\star(z)=s(z)} &= \theta\lambda\mathcal{P}(z;N)\Bigg[a'(s(z))+\frac{\phi'(s(z))}{\phi(s(z))}\Bigg([a(s(z))-c(z)]-\mathbb{E}_{\max\{\tilde{z}\}|\tilde{z}\leq z}[a(s(\tilde{z}))-c(\tilde{z})] \\
&\quad +(N-1)e^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[a(s(\tilde{z}))-c(\tilde{z})]\Bigg)\Bigg]-\frac{\psi}{1-Nm\bar{q}}\,q'(s(z)) \\
&= \theta\lambda\mathcal{P}(z;N)\Bigg[a'(s(z))+\frac{\phi'(s(z))}{\phi(s(z))}\Bigg([a(s(z))-c(z)]-\mathbb{E}_{\max\{\tilde{z}\}|\tilde{z}\leq z}[a(s(\tilde{z}))-c(\tilde{z})] \\
&\quad +(N-1)e^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[a(s(\tilde{z}))-c(\tilde{z})]\Bigg)\Bigg]-wq'(s(z)) \\
&= \theta\lambda\mathcal{P}(z;N)\frac{\phi'(s(z))}{\phi(s(z))}\Big[x^\star(z)-\mathbb{E}_{\max\{\tilde{z}\}|\tilde{z}\leq z}[a(s(\tilde{z}))-c(\tilde{z})]+(N-1)e^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[a(s(\tilde{z}))-c(\tilde{z})]\Big] \\
&= \theta\lambda\mathcal{P}(z;N)\frac{\phi'(s(z))}{\phi(s(z))}(N-1)e^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[a(s(\tilde{z}))-c(\tilde{z})]\leq 0,
\end{aligned}$$

where the second equality follows from the labor market clearing condition, $w=\frac{\psi}{1-Nm\bar{q}}$, the third equality follows from the equilibrium specialization condition (8) and the fourth equality from the equilibrium surplus offered to final producers (9).

We establish the first statement by noting that the marginal welfare effect of equilibrium specialization is zero if and only if $N=1$. Whenever $N>1$, the marginal welfare effect of equilibrium specialization is negative since $\phi'(s)<0$. Hence, when the production process is complex, the planner would choose a lower level of specialization: the equilibrium features *over-specialization*.

■

Proof of Proposition 2. To establish the claim, we follow the same strategy as in Proposition 1. Evaluating the marginal welfare effect of specialization at the equilibrium solution yields:

$$
\begin{aligned}
\left.\frac{\partial \mathcal{W}}{\partial s(z)}\right|_{s^\star(z)=s(z)} &= \theta\lambda\mathcal{P}(z;N)\Bigg[a'(s(z)) + \frac{\phi'(s(z))}{\phi(s(z))}\Bigg([a(s(z))-c(z)] - \mathbb{E}_{\max\{\tilde{z}\}|\tilde{z}\leq z}[a(s(\tilde{z}))-c(\tilde{z})] + (N-1)\\
&\quad e^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[a(s(\tilde{z}))-c(\tilde{z})] - \frac{\beta(1-\delta)}{1+\beta(1-\delta)}f^N N e^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[a(s(\tilde{z}))-c(\tilde{z})]\Bigg)\Bigg] - \frac{\psi}{1-Nm\bar{q}}\,q'(s(z))\\
&= \theta\lambda\mathcal{P}(z;N)\Bigg[a'(s(z)) + \frac{\phi'(s(z))}{\phi(s(z))}\Bigg([a(s(z))-c(z)] - \mathbb{E}_{\max\{\tilde{z}\}|\tilde{z}\leq z}[a(s(\tilde{z}))-c(\tilde{z})] + (N-1)\\
&\quad e^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[a(s(\tilde{z}))-c(\tilde{z})] - \frac{\beta(1-\delta)}{1+\beta(1-\delta)}f^N N e^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[a(s(\tilde{z}))-c(\tilde{z})]\Bigg)\Bigg] - wq'(s(z))\\
&= \theta\lambda\mathcal{P}(z;N)\frac{\phi'(s(z))}{\phi(s(z))}\Bigg[x^\star(z) - \mathbb{E}_{\max\{\tilde{z}\}|\tilde{z}\leq z}[a(s(\tilde{z}))-c(\tilde{z})] + (N-1)e^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[a(s(\tilde{z}))-c(\tilde{z})]\\
&\quad - \frac{\beta(1-\delta)}{1+\beta(1-\delta)}f^N N e^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[a(s(\tilde{z}))-c(\tilde{z})]\Bigg]\\
&= \theta\lambda\mathcal{P}(z;N)\frac{\phi'(s(z))}{\phi(s(z))}\left[N\left(1-\frac{\beta(1-\delta)}{1+\beta(1-\delta)}f^N\right)-1\right]e^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[a(s(\tilde{z}))-c(\tilde{z})] \lesseqgtr 0.
\end{aligned}
$$

Let $\mathcal{M}(N) \equiv 1 - \frac{\beta(1-\delta)}{1+\beta(1-\delta)}f^N$. The equilibrium is efficient if and only if $N\mathcal{M}(N) = 1 \iff N = \frac{1}{\mathcal{M}(N)}$. If $N > \frac{1}{\mathcal{M}(N)}$, the equilibrium exhibits over-specialization. Notice that $\frac{1}{\mathcal{M}(N)} = \frac{1+\beta(1-\delta)}{1+\beta(1-\delta)(1-f^N)}$. Since $\beta \in (0,1), \delta \in (0,1)$, and $f \in (0,1)$, it follows that $\frac{1}{\mathcal{M}(N)} \in (1, 1+\beta(1-\delta))$, where $1+\beta(1-\delta) < 2$. Since $N \in \mathbb{N}$, if the production process is complex, the equilibrium exhibits over-specialization. ■

Proof of Proposition 3. Let $\mathcal{G}(N) \equiv N\mathcal{M}(N) - 1$. From Proposition 2, we know that the equilibrium allocation displays over-specialization if $\mathcal{G}(N) > 0$. We now analyze the behavior of $\mathcal{G}(N)$ as the complexity of the production process varies. Formally, we study the sign of

$$
\mathcal{G}(N) - \mathcal{G}(N-1) = 1 - \chi\left(Nf(N)^N - (N-1)f(N-1)^{N-1}\right),
$$

where the dependence of the input finding probability on the complexity of the production process is made explicit and $\chi \equiv \frac{\beta(1-\delta)}{1+\beta(1-\delta)}$ collects the terms independent of N. We develop our argument in three steps:

1. Find the upper bound of $\chi\left(Nf(N)^N - (N-1)f(N-1)^{N-1}\right)$

   Since $\beta \in (0,1)$ and $\delta \in (0,1)$, the upper bound of χ is $\bar{\chi} = \frac{1}{2}$. Since $f(N) = 1 - e^{-\lambda\bar{\phi}(N)}$, its the upper bound is $\bar{f} = 1 - e^{-\lambda}$. It follows that the upper bound of the expression is

   $$
   U(N;\lambda) = \frac{1}{2}N(1-e^{-\lambda})^N
   $$

2. Analyze the monotonicity condition $\mathcal{G}(N) - \mathcal{G}(N-1) > 0 \impliedby U(N;\lambda) < 1 \iff N(1-e^{-\lambda})^N < 2$

Study the behavior of $\varphi(N) = N(1-e^{-\lambda})^N$. The limiting behavior is described by:

$$\varphi(1) = \bar{f} < 1,$$
$$\lim_{N\to\infty} N\bar{f}^N = \lim_{N\to\infty} \frac{1}{N\bar{f}^{1-N}} = 0,$$

where the second limit makes use of de l'Hopital rule.
The interior behavior is described by:

$$\varphi'(N) = \bar{f}^N(1+\ln\bar{f}),$$

where $\varphi'(N) > 0 \iff N < \frac{1}{-\ln\bar{f}}$ and $\varphi'(N) < 0 \iff N > \frac{1}{-\ln\bar{f}}$. Let $\bar{N}$ denote the complexity level which maximizes $\varphi(N)$, i.e., $\bar{N} \equiv \operatorname{argmax}\varphi(N) = \frac{1}{-\ln\bar{f}}$.

3. Find a parametric restriction on λ such that $U(N;\lambda) \leq 1\ \forall N$.

   A sufficient condition for the condition to hold is $U(\bar{N};\lambda) \leq< 1$. Formally,

$$\frac{(1-e^{-\lambda})^{\frac{1}{-\ln(1-e^{-\lambda})}}}{-\ln(1-e^{-\lambda})} < 2.$$

   The left-hand side ranges from 0 (for $\lambda \to 0$) to ∞ (for $\lambda \to \infty$). Moreover, for reasonable values of λ, the left-hand side is continuous and monotonically increasing in λ. Hence, $\exists!\ \bar{\lambda}$ such that $U(\bar{N};\lambda) \leq< 1\ \forall\lambda \leq \bar{\lambda}$.

Hence, if search efficiency is low enough, the equilibrium allocation of the dynamic economy displays more over-specialization and under-resilience as complexity increases. ■

Proof of Proposition 4. Let $\mathcal{G}(\delta) \equiv N\mathcal{M}(N;\delta) - 1$. From Proposition 2, we know that the equilibrium allocation displays more over-specialization the higher $\mathcal{G}(\delta)$ is. We now analyze the behavior of $\mathcal{G}(\delta)$ as the disruption probability varies.

$$\begin{aligned}
\mathcal{G}(\delta) &= N\left(1 - \frac{\beta(1-\delta)}{1+\beta(1-\delta)}f^N\right) - 1 \\
\mathcal{G}'(\delta) &= \frac{-\beta(1-f^N)[1+\beta(1-\delta)] + \beta[1+\beta(1-\delta)(1-f^N)]}{[1+\beta(1-\delta)]^2}N \\
&\propto -\beta[1+\beta(1-\delta)] + \beta f^N[1+\beta(1-\delta)] + \beta[1+\beta(1-\delta)] - \beta f 6N\beta(1-\delta) \\
&\propto 1 > 0.
\end{aligned}$$

Hence, as the frequency of disruption increases, the equilibrium allocation of the dynamic economy displays more over-specialization and under-resilience. ■

Proof of Proposition 5. Let $\mathcal{G}(\lambda) \equiv N\mathcal{M}(N;\lambda) - 1$. From Proposition 2, we know that the equilibrium allocation displays more over-specialization the higher $\mathcal{G}(\lambda)$ is. We now analyze the

behavior of $\mathcal{G}(\lambda)$ as search efficiency varies.

$$\begin{aligned}\mathcal{G}(\lambda) &= N\left(1-\chi f(\lambda)^N\right)-1\\ \mathcal{G}'(\lambda) &= -N\chi N f(\lambda)^{N-1}\frac{\partial f(\lambda)}{\partial \lambda}\\ &\propto -\frac{\partial f(\lambda)}{\partial \lambda}\\ &\propto -(1+\epsilon_{\bar{\phi},\lambda}),\end{aligned}$$

where $\epsilon_{\bar{\phi},\lambda} \equiv \frac{\partial \bar{\phi}(\lambda)}{\partial \lambda}\frac{\lambda}{\bar{\phi}(\lambda)}$ is the elasticity of the average compatibility probability to search efficiency and the last equality follows from $\frac{\partial f(\lambda)}{\partial \lambda} = \left(\bar{\phi}(\lambda)+\lambda\bar{\phi}'(\lambda)\right)e^{-\lambda\bar{\phi}(\lambda)} = \bar{\phi}(\lambda)(1+\epsilon_{\bar{\phi},\lambda})e^{-\lambda\bar{\phi}(\lambda)}$.

Hence, if the elasticity of the average compatibility probability to search efficiency exceeds one, the equilibrium allocation of the dynamic economy displays more over-specialization and under-resilience as search frictions decline. ■

Proof of Proposition 6. To establish the claim, we follow the same strategy as in Proposition 1. Evaluating the marginal welfare effect of specialization at the equilibrium solution yields:

$$\begin{aligned}\left.\frac{\partial \mathcal{W}}{\partial s(z)}\right|_{s^\star(z)=s(z)} &= \theta\lambda\mathcal{P}(z;N)\Bigg[a'(s(z))+\frac{\phi'(s(z))}{\phi(s(z))}\Bigg([a(s(z))-c(z)]-\mathbb{E}_{\max\{\tilde{z}\}|\tilde{z}\leq z}[a(s(\tilde{z}))-c(\tilde{z})]+(N-1)\\ &\quad \chi_1(f;N)e^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[a(s(\tilde{z}))-c(\tilde{z})]-\chi_2(f;N)e^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[a(s(\tilde{z}))-c(\tilde{z})]\Bigg)\Bigg]-\frac{\psi}{1-Nm\bar{q}}\,q'(s(z))\\ &= \theta\lambda\mathcal{P}(z;N)\Bigg[a'(s(z))+\frac{\phi'(s(z))}{\phi(s(z))}\Bigg([a(s(z))-c(z)]-\mathbb{E}_{\max\{\tilde{z}\}|\tilde{z}\leq z}[a(s(\tilde{z}))-c(\tilde{z})]+(N-1)\\ &\quad \chi_1(f;N)e^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[a(s(\tilde{z}))-c(\tilde{z})]-\chi_2(f;N)e^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[a(s(\tilde{z}))-c(\tilde{z})]\Bigg)\Bigg]-wq'(s(z))\\ &= \theta\lambda\mathcal{P}(z;N)\frac{\phi'(s(z))}{\phi(s(z))}\Bigg[x^\star(z)-\mathbb{E}_{\max\{\tilde{z}\}|\tilde{z}\leq z}[a(s(\tilde{z}))-c(\tilde{z})]+(N-1)\chi_1(f;N)e^{-\lambda\hat{\phi}(\underline{z},z)}\\ &\quad \hat{\mathbb{E}}[a(s(\tilde{z}))-c(\tilde{z})]-\chi_2(f;N)e^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[a(s(\tilde{z}))-c(\tilde{z})]\Bigg]\\ &= \theta\lambda\mathcal{P}(z;N)\frac{\phi'(s(z))}{\phi(s(z))}\left[(N-1)\chi_1(f;N)-\chi_2(f;N)\right]e^{-\lambda\hat{\phi}(\underline{z},z)}\hat{\mathbb{E}}[a(s(\tilde{z}))-c(\tilde{z})]\lesseqgtr 0.\end{aligned}$$

The equilibrium is efficient if and only if $N = 1+\frac{\chi_2(f;N)}{\chi_1(f;N)}$. If $N > 1+\frac{\chi_2(f;N)}{\chi_1(f;N)}$, the equilibrium exhibits over-specialization. ■

Proof of Proposition 7. Given equation (32), and the definition $\mathcal{X}(z) = x^\star(z)-\tau^\star(z)$ we note that $\tau^\star(z) = e^{-\lambda\hat{\phi}(\underline{z},z)}\,(N-1)\,e^{-\lambda\bar{\phi}}\hat{\mathbb{E}}[a(s(z))-c(z)]$. Using the definition of $T^\star$ we obtain equation (33). ■

Proof of Proposition 8. Following the same steps of the proof of Proposition 7 we obtain the untargeted component $T_d^\star$. Similarly, we obtain the bounds by evaluating $e^{-\lambda\hat{\phi}(\underline{z},z)}$ at $\underline{z}$ and $\bar{z}$. ■

Proof of Proposition A.1. Since the efficient specialization is independent of the surplus-sharing

rule, the result follows directly from applying Proposition 1 to the equilibrium specialization function of the bargaining model (38). ■

Proof of Proposition A.2. Rearranging eq. (45) and (46) yield $f(s^\star(\boldsymbol{z}))^{N^\star} = \exp\{-1\}$ and $f(\mathcal{S}(\boldsymbol{z}))^{\mathcal{N}} = \exp\{-(1-\frac{w\ell}{y})\}$. It follows that, as long as there are some product design costs in the economy, the level of complexity in equilibrium exceeds the efficient one. Note also that if $w\ell = y$ then the efficient level of complexity is $N^\star = 0$, which clearly cannot be optimal for the social planner. ■

Proof of Proposition A.3. Taking the search externality term in (47) to the left-hand side yields $\theta\frac{f^{N-1}}{\delta}\mathbb{E}[a-c]\Big(1-\ \mu N\ln(1/f)\Big) = \frac{\psi}{1-mN\bar{q}}\bar{q}$. Upon substituting for the labor market clearing and product market clearing conditions (39)-(40), it follows that $\frac{\psi}{1-mN\bar{q}}\bar{q}\frac{\theta f^{N-1}\mathbb{E}[a-c]}{\delta} = \frac{w\ell}{y}$. This first-order condition differs from (46) only by the share of searching final good producers μ multiplying $N\ln(1/f)$. Hence, $(f^N)^\star = \exp\left\{-\frac{1}{\mu}\left(1-\frac{w\ell}{y}\right)\right\}$. Comparing this expression with (45), it follows that complexity is efficient for given specialization if and only if $LS = 1-\mu$. As in the baseline dynamic model, specialization is efficient for given complexity if and only if $N = \frac{1}{\mu}$. If the two efficiency conditions hold simultaneously, the endogenous complexity equilibrium is therefore efficient.
Since $f^N = \exp\{-1\}$ and $f^{N^\star} = \exp\{-\frac{1}{\mu^\star}(1-LS^\star)\}$, the equilibrium features under-resilience if and only if $LS^\star > 1-\mu^\star$. ■

Proof of Proposition A.4. Consider the social planner problem of choosing specialization in the presence of any cost K borne by final producers.

$$\max_{\hat{s}(z)} \frac{\delta}{\delta+(1-\delta)f^{N-1}f(\hat{s}(z))}\frac{f^{N-1}}{\delta}\Bigg(\int \lambda\phi(\hat{s}(z))\exp\left\{-\lambda\hat{\phi}(z,\bar{z})\right\}\left[a(\hat{s}(z))-c(z)-\kappa\right]$$
$$\gamma(z)dz + (N-1)\frac{f^{N-1}}{\delta}f(\hat{s}(z))\mathbb{E}\left[a-c-\kappa\right]\Bigg) + \psi\log\left(1-m\left((N-1)\bar{q}+\bar{q}(\hat{s}(z))\right)\right),$$

where $\kappa \equiv K/N$. Efficient specialization of intermediate producers with productivity z is given by:

$$\theta\frac{f^{N-1}}{\delta}\left(\frac{\partial\mathbb{E}\left[a-c-\kappa\right]}{\partial s(z)} - (N-1)\frac{-f'(s(z))}{f}\mathbb{E}\left[a-c-\kappa\right]\right) + \theta\frac{f^{N-1}}{\delta}\frac{-f'(s(z))}{f}$$
$$\frac{(1-\delta f^N}{\delta+(1-\delta f^N}N\mathbb{E}[a-c-\kappa] = \frac{\psi}{1-mN\bar{q}}q'(s(z)).$$

Evaluating the planner's first-order condition at the equilibrium specialization yields:

$$\left.\frac{\partial \mathcal{W}}{\partial s(z)}\right|_{\mathcal{S}(z)=s^\star(z)} \propto \underbrace{\int_{\underline{z}}^{z}\left[a(s(\tilde{z}))-c(\tilde{z})-\kappa\right]\gamma^\star(\tilde{z},z)d\tilde{z}}_{\text{business stealing externality}} \underbrace{-\ x(z)}_{\substack{\text{appropriability}\\ \text{externality}}} \underbrace{-\ (N-1)e^{-\lambda\hat{\phi}(\underline{z},z)}\mathbb{E}[a-c-\kappa]/f}_{\text{network externality}}$$
$$\underbrace{+\ \frac{(1-\delta f^N}{\delta+(1-\delta f^N}e^{-\lambda\hat{\phi}(\underline{z},z)}N\mathbb{E}[a-c-\kappa]/f,}_{\text{search externality}}$$

where $x(z)=x_0+\mathbb{E}[a(s(\tilde{z}))-c(\tilde{z}),\ \tilde{z}<z]$, with $x_0=\kappa-\frac{N-1}{N}\mathbb{E}[\tilde{x}(z)]/f$ and $\tilde{x}(z)=\mathbb{E}[a(s(\tilde{z}))-c(\tilde{z}),\ \tilde{z}<z]$, is the equilibrium offered surplus. For the equilibrium offered surplus to balance appropriability and business-stealing externalities as in the model without costs downstream, the following needs to hold:

$$x_0=\left(1-e^{-\lambda\hat{\phi}(\underline{z},z)}\right)\kappa \iff \kappa=e^{\lambda\bar{\phi}(\underline{z},z)}\frac{N-1}{N}\mathbb{E}[\tilde{x}(z)]/f \quad \forall z\in[\underline{z},\bar{z}].$$

Since this condition can never hold simultaneously for any productivity level, appropriability and business-stealing externalities are not perfectly balanced by the equilibrium offered surplus when final producers bear other costs than input purchases. ■

Proof of Proposition A.5. Comparing (50) with (54) we note that the benefits of higher robustness are evaluated differently by firms and by the planner. Firms obtain a marginal flow benefit $\mathbb{E}[x-\kappa]$ smaller than the social marginal flow benefit $\mathbb{E}[a-c-\kappa]$. This force, in isolation, would induce under-investment in robustness. On the other hand, the planner discounts at the social discount factor $\delta(r)/[(1-f^N)\mu]$, while private firms discount at their destruction rate $\delta(r)$. As a consequence, the planner discounts the marginal surplus more than private firms, which pushes the equilibrium towards investing excessively in robustness. For given aggregate specialization, under-robustness obtains if and only if $\tilde{\zeta}\equiv\frac{\mathbb{E}[x-\kappa]}{\mathbb{E}[a-c-\kappa]}<(1-f^N)\mu$. From this formulation, it is apparent that the lower the share of the match surplus accruing to final producers, the more likely under-robustness is. Equilibrium specialization is efficient if and only if: $N=\frac{1}{\mu}\left(1+\frac{(1-e^{-\lambda\hat{\phi}(\underline{z},z)})\kappa-x_0}{\mathbb{E}[A-c-\kappa]/f}\right)\ \forall z\in[\underline{z},\bar{z}]$. Since this condition never holds for more than one productivity level at a time, equilibrium specialization at the firm level is inefficient. Since the exit probability discounts the marginal benefit of specialization in (21), higher robustness induces higher equilibrium specialization, as well. Inefficiency of firm-level specialization directly follows from Corollary 1. ■

Proof of Proposition A.6. The proposition follows directly from eq. (A.6). If inputs are not perfect substitutes, *i.e.*, $\frac{\partial\tilde{A}_j}{\partial s_i(z)}\geq 0\ \forall j=\{1,\dots,N\}\setminus\{i\}$, the social marginal benefit of specialization increases with respect to the private one, thus making the under-specialization condition of the baseline dynamic model, $N<1+\frac{(1-\delta)f^N}{\delta}$, sufficient for equilibrium under-specialization in the model with generalized quality function. ■

Proof of Proposition A.7. By rearranging (66), it follows that the equilibrium is efficient if and

only if:

$$\theta \frac{f^{N-2}}{\delta}\left(\frac{\delta}{f^N+(1-f^N)\delta}N-1\right)\mathbb{E}[a-c]=-\frac{\psi}{1-Nm\bar{q}}\frac{N\bar{q}}{m(-f'(s(x(z))))}\left(\frac{\partial m(s(x(z)),x(z))}{\partial s(x(z))}+\frac{\partial m(s(x(z)),x(z))/\partial x(z)}{\partial s(x(z))/\partial x(z)}\right) \quad \forall z\in[\underline{z},\bar{z}].$$

Since the left-hand side is constant across firms and the right-hand side is firm-specific, the economy with free entry is always inefficient. ■