

Supply Chain Resilience: Should Policy Promote International Diversification or Reshoring?

Gene M. Grossman

Princeton University

Elhanan Helpman

Harvard University

Hugo Lhuillier

Princeton University

Little is known about optimal policy in the face of global supply chain disruptions. Should governments promote resilience by subsidizing backup sources of input supply in multiple countries? Should they encourage firms to source from safer domestic suppliers? We address these questions in a model of production with a critical input and exogenous risks of supply disturbances. With constant elasticity of substitution preferences, a subsidy for diversification achieves the constrained social optimum. When the demand elasticity rises with price, private investments in resilience may be socially excessive and the social planner may wish to discourage diversification while favoring sourcing from abroad.

We are grateful to David Baqaee and Kiminori Matsuyama for very helpful discussions, to Maxim Alekseev and Alejandro Sabal for excellent research assistance, and to the referees and editor for helpful comments and advice. This paper was edited by Andrew Atkeson.

Electronically published November 3, 2023

Journal of Political Economy, volume 131, number 12, December 2023.

© 2023 The University of Chicago. All rights reserved. Published by The University of Chicago Press.

<https://doi.org/10.1086/725173>

I. Introduction

The United States needs resilient, diverse, and secure supply chains to ensure our economic prosperity and national security. Pandemics and other biological threats, cyberattacks, climate shocks and extreme weather events, terrorist attacks, geopolitical and economic competition, and other conditions can reduce critical manufacturing capacity and the availability and integrity of critical goods, products, and services. Resilient American supply chains will revitalize and rebuild domestic manufacturing capacity, maintain America's competitive edge in research and development, and create well-paying jobs. (President Joseph R. Biden Jr.'s executive order on America's supply chains, February 24, 2021)

Supply chain disruptions have become the new normal. The Great East Japan Earthquake of 2011 and the massive tsunami that it triggered brought such events to the attention of economists. Since then, hardly a month passes without news of a fresh disturbance. The pace of disruptions has quickened with the advent of the COVID-19 pandemic, and now we hear regularly of supply chain breakdowns in industries as disparate as automobiles, dishwashers, plastics, copper wire, lumber, pork, and toilet paper.

Disruptions have a myriad of causes. They result from natural disasters, geopolitical disputes, transportation failures, cyberattacks, fires, power outages, labor shortages, human error, and pandemic lockdowns. McKinsey Global Institute (Lund et al. 2020), which recently conducted a series of interviews with supply chain experts, reports that disruptions lasting 1–2 weeks happen to a given company on average every second year, while those lasting 1–2 months occur every 3.7 years. The disruptions impose significant costs, presenting firms with expected losses per decade that averaged 42% of their annual pretax earnings (see exhibit E5 in Lund et al. 2020, 12).

Many commentators associate the increasing frequency and severity of supply chain disruptions with the perils of globalization.¹ Global supply chains leave firms exposed to risks in multiple countries, some quite distant from where consumption takes place. This perceived connection between supply shortages and international trade, in turn, has sparked soul searching among policy makers and a call to action in the broader public. If costly shocks reflect concentration of input supplies, wouldn't it be

¹ See, e.g., Iakovou and White III (2020), Shih (2020a, 2020b), and Baldwin and Freeman (2022).

sensible for governments to encourage firms to diversify their international sourcing? And if distance from suppliers intensifies the risk of disruption, wouldn't it be better to bring some parts of the supply chains closer to home? The preamble to President Biden's executive order suggests that "*diverse* and secure supply chains" are a prerequisite for economic prosperity and that "resilient American supply chains" will rest on "rebuil[t] *domestic* manufacturing capacity" (emphasis added).

Little is known about the efficacy of policies aimed at global supply chain management in an environment with recurring disturbances. Disruptions generate input shortages that can give rise to price spikes or even outright unavailability of downstream products. Consumers suffer from their hampered ability to purchase the products they covet. To the extent that households forfeit consumer surplus in the face of supply chain disruptions, governments may have reason to enact policies that curtail their occurrence. But production impediments impact not only consumers' surplus but also firms' bottom lines. The question for governments is not whether shortages adversely affect households but whether firms' private incentives to avoid such shortages fall short of (or exceed) what is socially desirable.

In this paper, we propose a bare-bones framework to evaluate policies that can alter the organization of global supply chains. We abstract from all complexity in the production process by assuming that home firms manufacture unique varieties of nontraded differentiated products using a single critical input. Firms face a choice of whether to procure their inputs at home or abroad. Sourcing from a foreign supplier is tempting, because production costs are assumed to be cheaper there. But foreign sourcing may be riskier than local sourcing for geopolitical, logistical, or other reasons. Thus, firms may face a trade-off between the lower costs of offshoring and the greater safety of onshoring. They might also invest in resilience by establishing multiple supply relationships. Redundancy is costly, but it allows a firm to be active in more states of the world.

Focusing on welfare in the country where the final goods are consumed, we identify three potential distortions in private sourcing decisions. First, firms typically do not capture all of the surplus generated by the availability of their product. This *consumer surplus externality* suggests that too many firms may be ready to accept the extra risk of foreign sourcing in exchange for lower expected costs and that too few firms may bear the extra costs of diversification. Meanwhile, when firms manage to avoid disruptions, some of their profits come at the expense of competitors that are also viable in the same state of the world. This *business-stealing externality* tends to lead firms to overweight the resilience of their networks. Finally, a *consumption distortion* arises when prices of differentiated products exceed marginal costs while other goods are priced competitively. In a setting with potential supply disruptions and positive markups, government policy

might be directed to encourage product availability in states of the world when markups would otherwise be especially high.

Inasmuch as the social cost of supply chain disruptions stems from loss of consumer surplus, the form of consumer preferences plays a crucial role in our policy analysis. It has become commonplace to use constant elasticity of substitution (CES) preferences in trade models with endogenous entry, but the very special properties of these preferences have been recognized since the seminal work by Dixit and Stiglitz (1977). In many contexts with CES preferences, the consumer surplus externality from extra product variety happens to exactly offset the business-stealing externality from extra competition. These considerations apply as much to investments in resilience as they do to investments in entry, so it is important for understanding the efficient organization of global supply chains to allow for more flexible forms of demand. To this end, we follow Matsuyama and Ushchev (2017, 2020b) in adopting a broader class of preferences that are *homothetic with a single aggregator* (HSA). With HSA preferences, the demand for any variety depends on its price relative to a (common) aggregator of all prices. The CES utility function is a member of this class, but, more generally, HSA preferences allow the demand elasticity to increase with price. This property of demand, which characterizes many consumer goods, has been termed Marshall's second law of demand (MSLD). When it applies, the consumer surplus and business-stealing externalities do not cancel.

We find that with CES preferences for the differentiated products and optimal subsidies to address consumption distortions arising from markup pricing, the planner need not influence incentives for diversification or those for sourcing at home versus abroad. But with more general forms of HSA preferences that obey MSLD, the planner requires not only consumption subsidies but also a policy to discourage diversification and another to alter the incentives to form chains at home versus abroad. The optimal diversification tax corrects for the business-stealing externality, which generally exceeds the consumer surplus externality under MSLD. The optimal policy to influence offshoring versus onshoring reflects that with a nonconstant elasticity of substitution, the benefits from greater competition differ across states of the world.

Turning to the second-best policy problem that arises when consumption subsidies are infeasible, we find that CES preferences dictate a subsidy for diversification as the only necessary supply chain policy. These subsidies are a second-best response to the distortion from markup pricing. However, with more general HSA preferences, the second-best policy must also take into account the dominance of the business-stealing externality relative to the consumer surplus externality as well as the different distortions from markup pricing that emerge in different states of the world. For example, if production costs abroad are nearly the same

as those at home but the risk of a supply disruption is higher, there will be reduced product availability and less competition in states of the world when procurement from foreign suppliers is disrupted than in states where disturbances happen locally. Price-cost markups will be larger in the former states absent government policy, and second-best policy will tilt supply chain formation in favor of offshoring. On the other hand, if the two countries are relatively similar in their risks but differ greatly in costs, a tax on offshoring or a subsidy for onshoring may be indicated.

Our paper fits into an earlier literature on trade disruptions in a neo-classical setting. Much of this previous work addressed optimal policy responses to potential trade embargoes. Mayer (1977) showed that production subsidies are an optimal response to threats of trade interruption in the presence of costly adjustment. Bhagwati and Srinivasan (1976) made the likelihood of a disruption a function of the volume of trade and elucidated an efficiency role for tariffs to give agents an incentive to internalize the externality arising from their effect on the probability of a trade restriction. Arad and Hillman (1979) extended these earlier papers to allow for learning by doing in the production of a good that might later be subject to an embargo. Bergström et al. (1985) developed an infinite-horizon model to study the potential role of inventories to mitigate the threat of embargo. Perhaps the most sophisticated of these early studies was that by Cheng (1989), who considered recurrent embargo threats as a stationary Markov process that plays out with constraints on the speed of intersectoral reallocation.

The main difference between our work and this earlier literature stems from our treatment of the endogenous availability of differentiated products. With perfect competition and homogeneous goods, aggregate quantities matter for welfare, but the availability of a particular firm's offering does not. If a disruption causes some import good to be unavailable, there is no harm to consumers beyond the higher price of the domestic (perfect) substitutes. Of course, higher sticker prices play a role in a world with differentiated products, but there is also a direct harm to consumers from a particular variety not being available for purchase. For this reason, we believe that endogenous determination of the set of available products should feature prominently when evaluating policy toward supply chain security.

Our paper also relates to a literature on distortions in the entry process, which began with the seminal paper by Dixit and Stiglitz (1977) and includes more recent contributions by Bilbiie et al. (2012), Baqaee and Farhi (2020), and Matsuyama and Ushchev (2020b), among others. This literature focuses on entry of new firms into an industry in settings with imperfect competition. Bilbiie et al. (2012) study pricing and entry over the business cycle, focusing on intertemporal variation in markups and their relation to intertemporal marginal rates of substitutions. Baqaee and Farhi (2020) decompose the welfare losses from inefficient pricing

and entry by calculating second-order approximations around an efficient equilibrium while mostly treating markups as given. Matsuyama and Ushchev (2020b) analyze the inefficiencies that arise in a one-sector model with endogenous entry and endogenous markups. They introduce HSA preferences to allow for nonconstant markups and MSLD, an original approach that proves very useful in our context as well.

Our paper differs from this literature inasmuch as we do not consider new entry into an industry but rather the organization of supply chains that determines product availability in different states of the world. Unlike Baqaee and Farhi (2020), we derive exact wedges that describe the gaps between social valuations and private valuations and use these wedges to characterize first-best and second-best policies. Our model is designed to address inefficiencies in global value chains in the face of supply disruptions, and so we emphasize the asymmetries in cost and risk that often characterize international trade in intermediate goods.²

The remainder of this paper is organized as follows. In section II, we develop our model of risky supply chains and describe the laissez-faire equilibrium. Then, in section III, we pose and solve the social planner's problem that arises when she has access to state- and product-specific consumption subsidies to offset the inefficiencies caused by monopoly pricing. We characterize the wedges between social and private incentives for forming supply relationships in each country relative to an alternative of diversified sourcing. Armed with this understanding of the sources of inefficiency in supply chain formation, we turn in section IV to the more realistic policy problem that arises when consumption subsidies are infeasible. We characterize the supply chain policies that achieve a constrained optimum for arbitrary cost and risk parameters when preferences take the CES form. For more general HSA preferences, we derive analytical results that apply when cost and risk parameters are not very different in the two countries. We then turn to numerical simulations in section V to extend our insights to asymmetric cost and risk parameters. Section VI concludes. The appendix, available online, contains further details, proofs, and additional numerical simulations.

II. A Simple Model of Risky Supply Chains

A. *Supply Relationships*

The home economy can produce a homogeneous numeraire good and potentially a unit measure of nontraded differentiated consumer products. Revenues from sales of the numeraire good amount to $\bar{Y}$, all of

² A tangentially related paper is Elliot et al. (2022), who study networks with idiosyncratic probabilities of breakup that depend on investment choices. They do not entertain aggregate shocks as here. Rather, they are interested in the propagation of idiosyncratic shocks up and down the supply chain.

which is paid to workers as labor income. Production of differentiated products requires no labor. Rather, firm ω in this industry converts a single customized critical input into the final good ω using the linear production technology,

$$x(\omega) = m(\omega),$$

where $x(\omega)$ is output of good ω and $m(\omega)$ is the quantity of the customized input.³ If the firm has established a supplier relationship in country i and if that supply chain is operative, then the firm can procure the customized inputs at a cost of q_i per unit, $i \in \{H, F\}$, where the subscripts denote home and foreign, respectively, and we assume that $q_F < q_H$ to capture the motive for the internationalization of the value chain.⁴

To form any supply relationship, a firm must bear a sunk investment cost, k , in units of the numeraire good. This cost represents the up-front outlays associated with searching for a partner, negotiating a contract, and designing a suitable input. It captures the costliness of resilience, inasmuch as firms that form multiple supply relationships bear extra expenses compared with those whose supply chain is more streamlined.

Once a supply relationship has been established, it is subject to two possible disruption shocks. With probability $1 - \rho \geq 0$, any particular supply chain breaks down for exogenous and idiosyncratic reasons, which might be a failure of the prototype input, a strike in the supplier factory, a localized weather event in the location where the input would be produced, or anything else that happens independent of all other supply relationships. In any of these circumstances, the downstream firm loses the ability to purchase its input from the particular supplier for the length of the period captured by our model.⁵ In the complementary event, with probability $0 < \rho \leq 1$, no idiosyncratic supply disruption occurs, and the firm can buy as much as it wants from the particular supplier, provided that the latter is located in a country that is open for business.⁶ However,

³ Note that the producer of the final good might actually be a retailer and the input might be a consumer product.

⁴ The constant cost of procurement fits best with a market structure in which the downstream firm is vertically integrated with its upstream subsidiaries. Alternatively, we can imagine a situation in which the downstream firm has all the bargaining power in its relationship with arms-length suppliers and the marginal cost of the input is constant.

⁵ We treat all disruptions as catastrophic; when they occur, they eliminate all input supply from the affected source. Alternatively, we could allow for less severe shocks that limit supply to some positive quantity or that raise the cost of purchases above q_i .

⁶ Our analysis could be conducted without the idiosyncratic shocks, i.e., with $\rho = 1$. However, $\rho < 1$ provides an additional incentive for diversification, and it implies that not all firms will be able to operate even when neither country suffers a broad-based disruption. Moreover, the number of active firms in the state of the world when neither source country is fully disrupted will depend on the sourcing strategies adopted by the firms, which then becomes a consideration in the choice of policy.

with probability $1 - \gamma_i$, a country-wide shock disrupts all chains with suppliers in country i . These shocks, which we assume to be independent across countries (to simplify the expressions but with no substantive importance), represent events such as epidemics, political conflicts between national governments, or failures of the national transportation system. The relative safety of the home country is captured by the assumption that $\gamma_H > \gamma_F$.

The realizations of the two country-wide shocks generate four possible aggregate states of the world that we denote by $J \in \{H, F, B, N\}$. In state H, which occurs with probability $\delta^H = \gamma_H(1 - \gamma_F)$, foreign sources of supply are unavailable, but a firm that has a supplier in the home country can still operate, conditional on its relationship there surviving any idiosyncratic shock. In state F, which happens with probability $\delta^F = \gamma_F(1 - \gamma_H)$, a common shock hits home input suppliers, and only those downstream firms with supply relationships abroad might operate. In state B, no broad-based disruptions occur, and every downstream firm that avoids idiosyncratic shocks to some of its supply relationships will produce positive output. This state arises with probability $\delta^B = \gamma_H\gamma_F$. Finally, with the residual probability $\delta^N = (1 - \gamma_H)(1 - \gamma_F)$, both countries suffer adverse shocks, and no production or consumption of differentiated products takes place.

B. Preferences and Demand

There is a unit mass of identical consumers in the home country. The representative consumer is risk neutral and holds quasi-linear preferences over consumption of the homogeneous good, Y , and consumption of an aggregate index of differentiated products, X . We represent her (cardinal) utility as

$$V(X, Y) = Y + U(X), \quad (1)$$

where $U(\cdot)$ has a constant elasticity $\varepsilon > 1$;⁷ that is,

$$U(X) = \frac{\varepsilon}{\varepsilon - 1} (X^{(\varepsilon-1)/\varepsilon} - 1) \text{ for } \varepsilon > 1.$$

Each consumer maximizes utility in any state of the world subject to a standard budget constraint, $Y + \int_{\omega \in \Omega} p(\omega)x(\omega)d\omega = I$, where $p(\omega)$ is the price and $x(\omega)$ the quantity purchased of variety ω , Ω is the set of varieties

⁷ A unitary elasticity can be treated as a limiting case as $\varepsilon \rightarrow 1$. We cannot allow $\varepsilon = 1$, because then $V \rightarrow -\infty$ in the state of the world when all firms face supply disruptions in both countries. We could introduce a backup technology such that firms can produce their own critical inputs at some high cost $\bar{q} > q_H$, and then we could entertain $\varepsilon = 1$ and even $\varepsilon < 1$. But introducing the additional parameter $\bar{q}$ complicates the expressions without providing additional insights. Accordingly, we choose to restrict the range of the demand elasticity, in keeping with the empirical evidence provided by Fajgelbaum et al. (2020).

available in the relevant state, and I is income. When consumers have sufficient income, the constant elasticity of $U(X)$ gives rise to a constant elasticity demand for differentiated products,

$$X = P^{-\varepsilon}, \quad (2)$$

where P is the real price index dual to U . The consumer spends $PX = P^{1-\varepsilon}$ on differentiated products and devotes residual spending of $I - P^{1-\varepsilon}$ to the homogeneous good.

Following Matsuyama and Ushchev (2017, 2020b), we assume that preferences for the bundle of differentiated products belong to a class they aptly term *homothetic with a single aggregator*. Homotheticity implies that the consumption index X is a linearly homogenous function of consumption of the individual varieties $\{x(\omega)\}_{\omega \in \Omega}$. A single aggregator, A , which is a linearly homogenous function of the set of prices $\{p(\omega)\}_{\omega \in \Omega}$, guides the substitution between a particular variety ω and all other varieties. More formally, HSA preferences require the existence of a price aggregator A and market share function $s[p(\omega)/A]$ that is nonnegative for all relative prices such that

$$\frac{d \log P}{d \log p(\omega)} = s[z(\omega)] \quad (3)$$

and

$$\int_{\omega \in \Omega} s[z(\omega)] d\omega = 1, \quad (4)$$

where $z(\omega) := p(\omega)/A$ is the price of variety ω relative to the price aggregator. Equation (3) expresses the demand for any variety ω in implicit form; the substantive assumption is that this demand depends only on the price of that variety relative to a common aggregator. Equation (4) stipulates that the market shares sum to 1.

We place some mild restrictions on the market share function, $s(z)$. First, we impose the following:

ASSUMPTION 1. The market share function $s(z)$ is strictly decreasing when positive, with $\lim_{z \rightarrow \bar{z}} s(z) = 0$ for $\bar{z} \equiv \inf\{z > 0 | s(z) = 0\}$.

The assumption that $s(z)$ is strictly decreasing ensures that all varieties in X are gross substitutes. It admits both the case when $\bar{z} < \infty$, so that demand “chokes” at some finite relative price, and the case $\bar{z} = \infty$, when positive quantities are demanded at any finite price. We also assume that the aggregator A is well defined for any measure of firms.⁸

⁸ Matsuyama and Ushchev (2020a, 2022) discuss various ways to ensure the existence of a well-defined aggregator, such as, e.g., limiting the size of the market relative to the fixed cost of entry or assuming that $s(z)$ is large enough for the lowest feasible value of z .

Equation (3) implies that the elasticity of substitution between any two goods with equal prices is a function of the common relative price and is given by

$$\sigma(z) = 1 - \frac{zs'(z)}{s(z)} > 1.$$

We further adopt the following:

ASSUMPTION 2. (i) Either $\sigma(z) = \sigma$ or $\sigma'(z) > 0$ for all $z \in (0, \bar{z})$; and (ii) $\sigma(z) > \varepsilon$ in the neighborhood of the equilibrium and the social optimum.

The first part of assumption 2 allows for the case of *symmetric CES* preferences, where $s(z) = \alpha z^{1-\sigma}$, $\alpha > 0$. For all other HSA preferences, we impose *Marshall's second law of demand (MSLD)*, namely, that the demand for a good becomes more elastic as its price rises.⁹ For example, the *symmetric translog* preferences, developed by Feenstra (2003), drawing on Diewert (1974), satisfy MSLD. These preferences can be represented by a market share function $s(z) = -\theta \log z$, $z \in (0, 1)$, $\theta > 0$. Then $\sigma(z) = 1 - 1/\log z$.

The second part of assumption 2 ensures that the demand for any variety ω increases when the aggregate price of competing brands rises. For some market share functions, this assumption might be satisfied at all values of $z \in (0, \bar{z})$. For others, we would need to verify that it is satisfied ex post, that is, after solving for the equilibrium.

Finally, we note for future reference the relationship between the price index P and the demand aggregator A that applies for any HSA preferences. Matsuyama and Ushchev (2020b) prove that

$$\log P = C + \log A - \int_{\omega \in \Omega} \int_{P(\omega)/A}^{\bar{z}} \frac{s(\xi)}{\xi} d\xi d\omega, \quad (5)$$

where C is a constant.

C. Profit Maximization

Once the state of the world has been realized, the surviving producers purchase inputs and set prices to maximize profits, taking into account the aggregate demand for differentiated profits (summarized by P^J) and the competition they face (summarized by A^J). The firm producing variety ω maximizes its profits in state J by procuring its inputs from its

⁹ Zhelobodko et al. (2012) and Mrázová and Neary (2017) introduce MSLD by assuming what they refer to as “increasing relative love of variety” and “subconvex” demand, respectively. In each case, their preferences are directly explicitly additive (DEA), which rules out homotheticity (see Matsuyama and Ushchev 2017). One advantage of the HSA class of utility functions, relative to DEA and many others used in the literature, is that it allows for a nonconstant elasticity of substitution without violating homotheticity. See Matsuyama (2022) for further discussion.

least cost viable supplier and by marking up its price relative to that cost. Specifically, a firm that pays q_k for its inputs in state J solves

$$p^{J,K}(\omega) = \arg\max_p (P^J)^{1-\varepsilon} s\left(\frac{p}{A^J}\right) p^{-1} [p - q_k]$$

for $J \in \{H, F, B\}$, $K \in \{H, F\}$, taking the state-contingent price index P^J and the state-contingent aggregator A^J as given. Profit maximization requires

$$p^{J,K}(\omega) = \frac{\sigma[z^{J,K}(\omega)]}{\sigma[z^{J,K}(\omega)] - 1} q_k \quad (6)$$

and yields operating profits

$$\pi^{J,K}(\omega) = \frac{s[z^{J,K}(\omega)]}{\sigma[z^{J,K}(\omega)]} (P^J)^{1-\varepsilon}, \quad (7)$$

where $z^{J,K}(\omega) := p^{J,K}(\omega)/A^J$.¹⁰ Notice that the price of any variety might vary across states of the world and with the source of its inputs. The markup reflects the elasticity of demand as usual, but the latter is not constant; rather, it reflects the cost and availability of inputs. Markups will be higher in states of nature with dampened competition due to supply disruptions than in states with more ample competition. Similarly, markups will be higher for goods produced with high cost (domestic) inputs than goods produced with cheaper (foreign) inputs.

D. Supply Chain Management

The identical households collectively own the unit measure of downstream (potential) producers. Since the quasi-linear utility represented by (1) implies that these households are risk neutral with respect to income shocks, the firms make their ex ante investment to maximize expected profits. We allow firms to choose among three modes of organization (plus exit). Strategy h entails investment in a single supply relationship in the home country in the hope of onshoring. Strategy f entails investment in a single relationship in the foreign country in the hope of offshoring. Strategy b (for both) involves diversification, that is, investment in supply relationships in both places with the intention of sourcing from the low-cost foreign supplier if that is possible and from the higher-cost domestic supplier if that is possible and the low-cost foreign option is not available.¹¹

¹⁰ Equations (6) and (7) reduce to the familiar pricing and profit expressions for the CES case, where σ is constant and $s(z) = \alpha z^{1-\sigma}$.

¹¹ In principle, a firm that diversifies may choose to invest in multiple supply relationships in the same country. To avoid a taxonomy, we do not consider this possibility here; it will not be an attractive option for ρ close to 1.

Firms calculate expected profits with rational expectations about prices, sales, and costs in each state of the world, in view of the fraction of their competitors that pursues each strategy in equilibrium. Let μ_j be the fraction of firms that opt for strategy j , $j \in \{h, f, b\}$, with $\sum_j \mu_j \leq 1$. In state H, when all foreign sources of supply are disrupted, only those firms that have chosen strategy h or strategy b might operate and, among those, only the ones that avoid an idiosyncratic shock to their home supplier. Each such firm faces competition from $\rho(\mu_h + \mu_b)$ others, all of which have a unit cost of q_H . Analogously, in state F, it is the firms that pursued strategy f or strategy b that might produce. Again, only a fraction ρ can do so, because the others suffer relationship-specific supply disturbances. It follows that in state F, an active firm competes with $\rho(\mu_f + \mu_b)$ others, each of which has a unit cost of q_F . Recognizing that all firms operating in state H have common costs q_H and those operating in state F have common costs q_F , we can use (6) and (7) to calculate the values of $\pi^H(\omega)$ and $\pi^F(\omega)$ that accrue to all firms operating in those states.¹²

A firm's profit in state B, in which supply chains in both countries are active, is slightly more complicated. In this state, firms that adopt either strategy f or strategy b anticipate a cost of q_F with probability ρ . Those that diversify by choosing strategy b anticipate that they will rely on their backup supplier at the higher cost q_H , with probability $\rho(1 - \rho)$. Meanwhile, firms that pursue strategy h also produce at q_H but with probability ρ . It follows that all firms anticipate competition in state B from $(\mu_f + \mu_b)\rho$ others producing at cost q_F and from $\mu_h\rho + \mu_b\rho(1 - \rho)$ others producing at cost q_H .

Let Π_j be the expected profit that a firm can earn by pursuing strategy j , $j \in \{h, f, b\}$. Recalling that δ^J is the probability of state J , $J \in \{H, F, B\}$, we have using (6) and (7) that

$$\Pi_h = \Pi_h(\mu) := \delta^H \frac{s[z^H(\mu)]}{\sigma[z^H(\mu)]} P^H[z^H(\mu)]^{1-\varepsilon} \rho + \delta^B \frac{s[z^{B,H}(\mu)]}{\sigma[z^{B,H}(\mu)]} P^B(\mu)^{1-\varepsilon} \rho - k, \quad (8)$$

$$\Pi_f = \Pi_f(\mu) := \delta^F \frac{s[z^F(\mu)]}{\sigma[z^F(\mu)]} P^F[z^F(\mu)]^{1-\varepsilon} \rho + \delta^B \frac{s[z^{B,F}(\mu)]}{\sigma[z^{B,F}(\mu)]} P^B(\mu)^{1-\varepsilon} \rho - k, \quad (9)$$

¹² Recognizing that all firms operating in state H source their inputs from H and that all firms operating in state F source their inputs from F, we henceforth omit the superscript K for any variable $\xi^{J,K}$ that applies when $J = H$ or F . For example, we use π^H in place of $\pi^{H,H}$, z^H in place of $z^{H,H}$, and so on. Also, with some abuse of notation, we write $\xi'(\mu)$ for the variable $\xi'(\omega)$ when the value of ξ' is common to all firms operating in state J and that value depends on the vector $\mu := (\mu_h, \mu_f, \mu_b)$.

and

$$\begin{aligned} \Pi_b = \Pi_b(\mu) := & \sum_{j=H,F} \delta^j \frac{s[z^j(\mu)]}{\sigma[z^j(\mu)]} P^j[z^j(\mu)]^{1-\varepsilon} \rho \\ & + \delta^B \left\{ \frac{s[z^{B,F}(\mu)]}{\sigma[z^{B,F}(\mu)]} + \frac{s[z^{B,H}(\mu)]}{\sigma[z^{B,H}(\mu)]} (1-\rho) \right\} P^B(\mu)^{1-\varepsilon} \rho - 2k. \end{aligned} \quad (10)$$

In equilibrium, if one strategy j dominates the other two, all active firms will make that choice, and so $\mu_\ell = 0$ for $\ell \neq j$. If two strategies yield positive and equally high expected profits and higher than the third, then these two will have positive fractions in equilibrium, while the third will find no takers. The fractions will be such as to generate indifference. Finally, if there exist $\mu_h > 0$, $\mu_f > 0$, and $\mu_b > 0$ such that $\Pi_h = \Pi_f = \Pi_b > 0$, then the equilibrium will exhibit a positive number of firms pursuing each of the available strategies.¹³

E. Welfare

We adopt expected indirect utility as our welfare metric, weighting utility in each aggregate state by the likelihood of that state. Indirect utility comprises labor income, profits, tax revenues (if any), and consumer surplus.

Expected welfare reflects the fractions of firms that choose each organizational mode, outcomes that can be influenced by government policy. When the fraction of firms that adopt strategies h, f, and b are μ_h , μ_f , and μ_b , respectively, aggregate expected profits (net of any subsidies or taxes received or paid by firms in recognition of their supply chain choices) amount to $\sum_{j=h,f,b} \mu_j \Pi_j(\mu)$. Consumer surplus in state J is given by $[1/(\varepsilon - 1)]P^J(\mu)^{1-\varepsilon}$ for $J \in \{H, F, B\}$. In the event that both countries are hit with supply disruptions, consumption of all differentiated products is zero ($X = 0$) and so consumer surplus vanishes. Therefore,

$$W(\mu) = \bar{Y} + \sum_{j=H,F,B} \delta^j T^j(\mu) + \sum_{j=h,f,b} \mu_j \Pi_j(\mu) + \frac{1}{\varepsilon - 1} \sum_{j=H,F,B} \delta^j P^j(\mu)^{1-\varepsilon}, \quad (11)$$

where $\bar{Y}$ is the (fixed) labor income from producing the numeraire good and $T^j(\mu)$ is tax revenues collected by the government and rebated to households in state J (possibly zero or negative) beyond what is paid to or collected from firms in connection with their supply chain choices.

F. The Laissez-Faire Equilibrium

In the absence of any government policy, we can use the fact that all strategies used in equilibrium maximize expected profits and the fact that

¹³ Since firms have the option to exit, it is possible that $\sum_j \mu_j < 1$, in which case all firms that do not exit make zero expected profits net of the fixed costs.

product markets must clear in every state to solve for μ as well as the state-contingent aggregators A^H , A^F , and A^B . Market shares must sum to 1 in every state; that is,

$$\rho(\mu_h + \mu_b)s[z^H(\mu)] = 1, \quad (12)$$

$$\rho(\mu_f + \mu_b)s[z^F(\mu)] = 1, \quad (13)$$

and

$$\rho(\mu_f + \mu_b)s[z^{B,F}(\mu)] + \rho[\mu_h + (1 - \rho)\mu_f]s[z^{B,H}(\mu)] = 1. \quad (14)$$

In state B, we need one more equation to solve for prices, which comes from comparing the optimal prices (6) for firms that source their inputs in H versus F. This gives

$$\frac{z^{B,F}(\mu)}{z^{B,H}(\mu)} = \frac{\sigma[z^{B,F}(\mu)]q_F/(\sigma[z^{B,F}(\mu)] - 1)}{\sigma[z^{B,H}(\mu)]q_H/(\sigma[z^{B,H}(\mu)] - 1)}.$$

Further details for computing the laissez-faire equilibrium can be found in sections 2.4 and 2.6 of the appendix.¹⁴

Figure 1 illustrates the fraction of firms that choose each organizational form for different values of k . The figure is drawn for the case in which production costs are roughly similar ($q_H \approx q_F$), but sourcing from the home country is significantly safer than sourcing from abroad ($\gamma_H \gg \gamma_F$). For $k < k_1$ in the figure, the fixed cost of a sourcing relationship is sufficiently small so that all firms find it worthwhile to invest in resilience, so $\mu_b = 1$ and $\mu_h = \mu_f = 0$. Next comes a range of fixed costs $k \in (k_1, k_2)$ for which some firms are diversified and others source only from the safer country H. For higher fixed costs such that $k \in (k_2, k_3)$ in the figure, each of the three strategies is deployed in equilibrium by some positive number of firms. When k surpasses k_3 , it is no longer profitable for any firm to diversify; $\mu_b = 0$. Then, for $k \in (k_3, k_4)$, a marginal change in the fixed cost has no effect on the relative profitability of strategy h versus strategy f.¹⁵ But once k increases beyond k_4 , the expected operating profits are not sufficient for the full unit measure of firms to cover fixed costs; some firms exit, so that $\mu_h + \mu_f < 1$ and $\Pi_h = \Pi_f = 0$. The total number of firms might remain positive but continue to fall with further increases in k , or else μ_f and then μ_h might reach zero for some finite values of $k > k_4$.¹⁶

¹⁴ Section numbers in the appendix correspond to the sections in the main text. Henceforth, when we refer to the appendix for formal arguments, we will note the section number only when it is an exception to this general rule.

¹⁵ Meanwhile, for $k \in [k_3, k_4]$, a greater fraction of firms sources their input from the safer country, $\mu_f < 0.5 < \mu_h$, since $\gamma_H \gg \gamma_F$ while $q_H \approx q_F$.

¹⁶ For some preferences, such as symmetric CES, operating profits in a state approach infinity as the number of varieties available in the state approaches zero, in which case some (shrinking number) of firms will invest in supply relationships in each state of the




FIG. 1.—Supply chain outcomes for $\gamma_H \gg \gamma_F$, $q_H \approx q_F$.

Figures analogous to figure 1 can be drawn for other configurations of cost and risk parameters. In what follows, we focus exclusively on circumstances that give rise to the use of all available strategies in equilibrium, that is, $\mu_h > 0$, $\mu_f > 0$, and $\mu_b = 1 - \mu_h - \mu_f > 0$. This outcome seems to resemble most closely what we observe in reality.¹⁷

III. The Unconstrained Social Optimum

Do private incentives for firms to invest in safe and resilient supply relationships align with social incentives in the absence of government policy? If the answer is no, can the social planner intervene to restore social efficiency? The answer to the second question may depend on the set of policy instruments that the government has at its disposal to influence resource allocation. As is well known, the markup pricing reflected in equation (6) creates a wedge between social and private incentives to consume

world. For other preferences, the potential operating profits per variety may be bounded. Then, for k large enough, all firms exit the industry.

¹⁷ The cases in which some strategies are not used by all firms could be analyzed similarly but do not generate any interesting insights beyond what we recount below.

differentiated products relative to the homogeneous good, because consumers face a price in excess of marginal cost for the former but a price equal to marginal cost for the latter. As in other contexts (see, e.g., Dhingra and Morrow 2019; Campolmi et al. 2021), this distortion can be eliminated, in principle, by an optimal set of consumption subsidies. In practice, such subsidies are difficult to implement and rarely observed; in our setting with nonconstant markups, the subsidies must vary with both the state of nature (H, F, or B) and with the source of the inputs embodied in the final good.¹⁸ Nonetheless, it is instructive to begin our analysis under the assumption that optimal subsidies are feasible to focus squarely on the wedges between private and social incentives for supply chain formation. In this section, we study the unconstrained (or first-best) planner's problem, leaving the more realistic, constrained (or second-best) problem for section IV. By assuming away the distortions caused by markup pricing, we are able to develop intuition for whether and when the incentives firms face to invest in safety and resilience are excessive, insufficient, or appropriate.

To characterize the optimal supply chain policies, we define the wedge between private and social incentives to pursue strategy j relative to strategy b when the existing mix of strategies is μ as

$$\tilde{w}_j(\mu) := [\tilde{\Pi}_j(\mu) - \tilde{\Pi}_b(\mu)] - \frac{d\tilde{W}(\mu)}{d\mu_j}, \quad j \in \{h, f\}, \quad (15)$$

where tildes indicate relationships that apply with optimal consumption subsidies in place and

$$\frac{d\tilde{W}(\mu)}{d\mu_j} := \frac{\partial \tilde{W}(\mu_h, \mu_f, \mu_b)}{\partial \mu_j} - \frac{\partial \tilde{W}(\mu_h, \mu_f, \mu_b)}{\partial \mu_b}, \quad j \in \{h, f\}$$

is the marginal change in welfare from a small change in μ_j at the expense of μ_b ; that is, $d\mu_j = -d\mu_b$.¹⁹

At an (interior) first-best allocation μ° , the first-order conditions for welfare maximization require $d\tilde{W}_j(\mu^\circ)/d\mu_j = 0$ for $j \in \{h, f\}$. Therefore, the first best can be achieved by a set of subsidies that satisfy

$$\tilde{w}_j(\mu^\circ) = \tilde{\Pi}_j(\mu^\circ) - \tilde{\Pi}_b(\mu^\circ) = \varphi_b - \varphi_j, \quad j \in \{h, f\}, \quad (16)$$

¹⁸ In state B, some final producers source from suppliers in the foreign country at unit cost q_f while others source from the home country at the higher cost q_h . The subsidies needed to ensure that consumers see relative prices equal to relative marginal costs will therefore vary with the sourcing of the inputs, despite the fact that all final goods enter demand symmetrically.

¹⁹ We henceforth use the notation of total derivatives, $dG(\mu)/d\mu_j$, for $j \in \{h, f\}$, to denote the marginal change in the function $G(\cdot)$ with respect to μ_j , taking into account that $\mu_b = 1 - \mu_h - \mu_f$.

where φ_j is a subsidy (possibly negative) paid unconditionally to any firm that pursues strategy j and φ_b is a subsidy (possibly negative) paid to a firm that diversifies its sourcing options. The optimal supply chain policies offset the wedges that remain (if any) when only consumption subsidies are applied.

We have given the social planner three policy instruments to counteract two wedges. Clearly, she has a degree of freedom in her policy choices. The first best can be achieved with a continuum of combinations of subsidies/taxes, including ones that eschew the use of one instrument entirely.

In the appendix, we compute the two wedges, (45) and (46), which yields

$$\tilde{w}_f^o := \tilde{w}_f(\mu^o) = \delta^H \Phi[z(\mu_f^o)] \tilde{P}^H(\mu^o)^{1-\varepsilon} \rho + \delta^B \Phi[z^{B,H}(\mu^o)] \tilde{P}^B(\mu^o)^{1-\varepsilon} \rho(1 - \rho) \quad (17)$$

and

$$\begin{aligned} \tilde{w}_h^o := \tilde{w}_h(\mu^o) = & \delta^F \Phi[z(\mu_h^o)] \tilde{P}^F(\mu^o)^{1-\varepsilon} \rho + \delta^B \Phi[z^{B,H}(\mu^o)] \tilde{P}^B(\mu^o)^{1-\varepsilon} \rho(1 - \rho) \\ & + \delta^B \{\Phi[z^{B,F}(\mu^o)] - \Phi[z^{B,H}(\mu^o)]\} \tilde{P}^B(\mu^o)^{1-\varepsilon} \rho, \end{aligned} \quad (18)$$

where

$$\Phi(z) := \int_z^{\bar{z}} \frac{s(\zeta)}{\zeta} d\zeta - \frac{s(z)}{\sigma(z) - 1}.$$

Assumption 2(i) specifies that $\sigma'(z) > 0$ for all $z \in (0, \bar{z})$ (i.e., MSLD) or else σ is constant for all $z \in (0, \bar{z})$. Taking the latter case first, it is straightforward to see that a CES implies $\Phi(z) = 0$ for all $z \in (0, \bar{z})$. This in turn implies $\tilde{w}_f^o = \tilde{w}_h^o = 0$; both wedges are zero in the market equilibrium when (only) optimal consumption subsidies are applied.

In contrast, when the elasticity of substitution rises with the relative price, we show in (25) in the appendix that $\Phi(z) < 0$ for all $z \in (0, \bar{z})$. Then (17) implies that $\tilde{w}_f^o < 0$, because both terms on the right-hand side are negative. As for strategy h, we have that $\tilde{z}^{B,H}(\mu^o) > \tilde{z}^{B,F}(\mu^o)$, because efficient relative prices are equal to relative marginal costs, and $q_H > q_F$. Together with $\Phi'(z) > 0$, (18) implies that $\tilde{w}_h^o < 0$ as well.

The negative wedges imply that with only consumption subsidies but no subsidies or taxes to influence supply chain formation, firms have excessive incentives for diversification; that is, converting a firm with an exclusive relationship in either country to one that is diversified will reduce aggregate welfare. To interpret this finding, note that firms' investments in supply chains determine the number of differentiated products available in each state of the world, as well as (with MSLD) their markups and prices. Greater diversification is like additional entry in states H and F, because more firms survive the country-specific supply disruptions. At the

same time, it generates greater competition in state B to the extent that idiosyncratic shocks ($\rho < 1$) impede product availability in that state.

When a firm chooses its investment strategy and thereby affects the number of varieties available in different states, it conveys two externalities. On the one hand, consumers love variety and they reap consumer surplus from greater availability at a given price. Firms do not take account of this positive effect of their product's availability on consumer surplus when forming their supply chains. On the other hand, more variety spells less demand and less profits for any particular product at given prices. Firms do not take account of this adverse effect of their product's availability on the profits earned by others. The sign of the wedge at the optimal allocation reflects the relative sizes of these two countervailing externalities.

Now consider, for example, how a change in the number of products n^H available in state H affects the gap between the price index, P^H , and the demand aggregator, A^H . The former fully captures the effect of product availability on consumer surplus whereas the effect on aggregate profits also depends on the latter. Using (5), (6) and (12), we calculate

$$\frac{1}{P^H} \frac{dP^H}{dn^H} - \frac{1}{A^H} \frac{dA^H}{dn^H} = \frac{s(z^H)}{\sigma(z^H) - 1} - \int_{z^H}^{\bar{z}} \frac{s(\zeta)}{\zeta} d\zeta = -\Phi(z^H).$$

When preferences satisfy MSLD, an extra variety in state H reduces the price index for that state by proportionately less than it does the demand aggregator. Thus, the positive consumer surplus externality from added availability in state H falls short of the negative business-stealing externality. In such circumstances, the private incentives for resilience exceed the social incentives.²⁰

Similar forces are at work with respect to a firm's choice between strategy h and strategy b. An increase in μ_b at the expense of μ_h means greater availability in state F (when the home suppliers are disrupted) and in state B (when diversified firms stand a better chance of avoiding the disruption from idiosyncratic shocks). However, an increase in μ_b at the expense of μ_h has a further effect on the wedge $\tilde{w}_h$, as represented by the third term on the right-hand side of (18). When μ_b rises and μ_h falls, the extra products that become available in state B are low-cost goods, whereas when μ_b rises and μ_f falls, the marginal products are high-cost goods. Extra low-cost goods take a greater toll on the profits of competitors than do extra high-cost goods, which adds the additional negative term to the wedge $\tilde{w}_h^0$.

²⁰ This result echoes that in Matsuyama and Uschev (2020b) that there is excessive entry under MSLD in a one-sector model of monopolistic competition and no supply shocks.

We turn now to the policies that the planner can introduce, alongside the optimal state- and product-contingent consumption subsidies, to implement the first best. Let us begin with the limiting case of symmetric, CES preferences. With $\tilde{w}_f^o = \tilde{w}_h^o = 0$, the optimum can be achieved without any intervention in supply chain formation whatsoever; that is, $\varphi_h = \varphi_f = \varphi_b = 0$. This is because with CES preferences, the price index P that (inversely) measures welfare is proportional to the demand aggregator A . Then the external effects of a product's availability on consumers and competitors are equal in magnitude and opposite in sign. Once optimal consumption subsidies are in place to counter the distortion created by markup pricing, there is no need for further government policy to influence the number of products available in any state.

Next consider the special case of symmetric translog preferences. We have seen that for all HSA preferences that obey MSLD, the government must discourage investments in resilience. But for the translog case, we can say more. Using $s(z) = -\theta \log z$, we find $\int_z^{\bar{z}} (s(\zeta)/\zeta) d\zeta = (1/2)(s(z)/(\sigma(z) - 1))$ for all $z \in (0, 1)$. Then using the planner's first-order conditions for the optimal choice of μ_h and μ_f along with (30)–(32) in the appendix, we find that $\tilde{w}_f^o = \tilde{w}_h^o = -k$; see lemma 3 in the appendix. The two wedges coincide for all values of the cost and risk parameters, and they are equal in absolute value to the fixed cost of forming a supply relationship. In the translog case, the planner can achieve the first best by combining the optimal consumption subsidies with a tax on diversification; $\varphi_b = -k$, with $\varphi_h = \varphi_f = 0$. Alternatively, she can leave diversified firms to face the private cost of their supply chains ($\varphi_b = 0$) while subsidizing firms that form exclusive supply relationships to the full extent of their investment costs ($\varphi_h = \varphi_f = k$). In either case, she has no reason to favor onshore investments relative to offshore investments.

These surprising results reflect a special property of symmetric translog preferences, namely, that the ratio of $\int_z^{\bar{z}} (s(\zeta)/\zeta) d\zeta$ to $s(z)/(\sigma(z) - 1)$ is independent of price and always equal to one-half. Since the consumer surplus loss from removing a variety in some state of the world is proportional to $\int_z^{\bar{z}} (s(\zeta)/\zeta) d\zeta$, while the loss in operating profits for the firm that does not produce its variety is proportional to $s(z)/(\sigma(z) - 1)$, translog preferences imply that the consumer surplus loss from switching a firm from having two suppliers to one is exactly half of the loss in operating profits. But the wedge $\tilde{w}_j^o$ is equal to the difference in total expected profits per (16), which in turn is equal to the fixed cost of an extra relationship minus the loss in operating profits. Finally, the first-order condition for maximizing $\tilde{W}(\mu)$ dictates that the marginal loss of consumer surplus from switching a firm from strategy b to strategy j should match the cost of an extra supplier, k .

With more general HSA preferences, the incentives created by optimal supply chain policy are not neutral with respect to the location of firms'

input suppliers. The optimal policies favor onshore relationships relative to offshore relationships if $|\tilde{w}_h^o| > |\tilde{w}_f^o|$ and offshore relationships relative to onshore relationships if the ranking of the two wedges is reversed.

To shed further light on the desired national bias in first-best supply chain policy, we examine the limiting case when production costs are nearly the same, $q_H \approx q_F$. Then we find in (50) in the appendix that

$$|\tilde{w}_h^o| - |\tilde{w}_f^o| \propto \Psi[z^H(\mu^o)] - \Psi[z^F(\mu^o)],$$

where $\Psi(z) := [\int_z^{\bar{z}} (s(\xi)/\xi) d\xi] / [s(z)/(\sigma(z) - 1)]$. In the appendix, we also show that $z^H(\mu^o) < z^F(\mu^o)$. It follows that with nearly equal costs, the government should encourage the less risky investments at home relative to the more risky investments abroad if $\Psi(z)$ is a decreasing function, and the reverse if it is an increasing function; see lemma 5 in the appendix.

When $q_H \gg q_F$, this simple reasoning does not apply, because the planner's preference for home sourcing on safety grounds is counteracted by her preference for foreign sourcing on cost grounds, so that $\mu_h^o \leq \mu_f^o$. Moreover, with unequal costs, an increase μ_b at the expense of μ_h makes a greater contribution to consumer surplus in state B while taking a greater toll on rivals' profits than does an increase in μ_b at the expense of μ_f ; although both increase product availability in state B by similar amounts, the former spells greater availability of low-cost products, whereas the latter generates greater availability of high-cost products.

We summarize our findings about first-best supply chain policy when all three strategies are used in the following:

PROPOSITION 1. Under assumptions 1 and 2, the unconstrained planner uses consumption subsidies to undo the markup distortion for each good in each state of nature. With symmetric CES preferences, a hands-off policy with respect to supply chain formation ($\varphi_h = \varphi_f = \varphi_b = 0$) achieves the first best. Under MSLD, the planner encourages single sourcing relationships relative to diversification. With symmetric translog preferences, the planner can achieve the first best with a tax on diversification of size k . More generally, if cost differences are small ($q_H \approx q_F$), the planner encourages onshore sourcing relative to offshore sourcing if $\Psi'(z) > 0$ for all $z \in (0, \bar{z})$ and encourages offshore sourcing relative to onshore sourcing if $\Psi'(z) < 0$ for all $z \in (0, \bar{z})$. For larger cost differences, the national bias in optimal sourcing policy hinges not only on the sign of $\Psi'(z)$ but also on the magnitudes of the cross-country cost and risk differences.

IV. The Constrained Social Optimum

In section III, we characterized the first-best allocation of resources when firms in a monopolistically competitive industry face potential supply chain disruptions. We noted that attainment of the first best requires

not only that government use policies to offset distortions in firms' incentives for forming supply relationships but also a policy to counter the distortion created by monopoly pricing of differentiated products alongside the competitive pricing of goods elsewhere in the economy. As we discussed, the requisite consumption subsidies are rarely implemented in practice. In our context, not only would they need to be adjusted in response to realized disruptions but also they would need to vary across otherwise symmetric products that differ only in the sourcing of their critical inputs. Nonetheless, by allowing for optimal consumption subsidies, we were able to lay bare the wedges between private and social incentives for supply diversification and for onshoring versus offshoring.

In this section, we consider the second-best problem that confronts a welfare-maximizing government that lacks the ability to implement state-contingent and sourcing-contingent consumption subsidies. We grant the policy maker only taxes or subsidies to encourage or discourage supply chain resilience and to influence whether sourcing partnerships are formed at home or abroad. As with the unconstrained optimum, the constrained social optimum can be achieved by a continuum of combinations of subsidies or taxes for the formation of home relationships, foreign relationships, and multiple relationships, and indeed, any two of these three instruments will suffice.

Let us begin with the limiting case in which the home and foreign countries are symmetric in terms of both input costs and disruption risks; that is, $q_H \approx q_F$ and $\gamma_H \approx \gamma_F$. In figure 2, we illustrate a laissez-faire equilibrium at E for a typical case in which all three strategies are employed by positive measures of firms. In such circumstances, we can use the equilibrium relationship $\mu_b = 1 - \mu_h - \mu_f$ to project the three-dimensional space (μ_f, μ_h, μ_b) onto two dimensions. Accordingly, the figure shows μ_h and μ_f on the vertical and horizontal axes, and the equilibrium falls inside the unit simplex.

The curve labeled $\Pi_h = \Pi_f$ represents combinations of μ_h and μ_f such that a strategy of forming a single supply relationship at home yields the same expected profit as that of forming a single supply relationship abroad. With similar production costs, the profits from strategies h and f are the same in state B. Hence, when choosing between these two strategies, firms compare realized profits in states H and F. Profits in state H are declining in the number of firms $n^H(\mu) = \rho(1 - \mu_f)$ that are active in that state. Similarly, profits in state F are declining in $n^F(\mu) = \rho(1 - \mu_h)$. It follows that firms will be indifferent between h and f only if the expected competition in the two states is the same, that is, $\mu_h = \mu_f$. Thus, we represent the $\Pi_h = \Pi_f$ curve by a 45° ray from the origin.

The downward sloping curve labeled $\Pi_h = \Pi_b$ shows combinations of μ_h and μ_f for which investing in a single relationship at home yields the same expected profits as a strategy of diversification. The downward




FIG. 2.—Equilibrium and constrained optimum for symmetric case.

slope of the curve can be understood as follows. Starting from a point on the curve, suppose that we raise μ_h and reduce μ_b so that μ_f remains constant. This does not affect the number of firms n^H active in state H. But it decreases the number of firms active in state F and also in state B, since $n^B(\mu) = \rho + \rho(1 - \rho)\mu_b$. Therefore, π^F and π^B both rise, leaving a (positive) gap between Π_b and Π_h .²¹ Now consider a fall in μ_f accompanied by an offsetting rise in μ_b that leaves μ_h unchanged. This change in composition has no effect on the number of firms active in state F but increases the number that are active in state H. However, the more intense competition in state H does not affect the relative attractiveness of strategy h

²¹ The effect on expected profits of a diversified firm relative to a home-only firm conditional on state B is equal and opposite for a given increase in μ_h and comparable decrease in μ_f . But the increase in μ_h (and accompanying decrease in μ_b) gives an added boost to the relative profitability of diversification, because it raises the expected profits for a b firm if state F arises.

versus b , because each of these strategies yields operating profits π^H with probability ρ in that state. Meanwhile, competition also intensifies in state B . With more firms active, π^B falls, which depresses expected profits more for diversified firms than for firms that have only a home supplier, because the diversified firms are more likely to survive. Thus, a decrease in μ_f offset by an increase in μ_b reduces Π_b relative to Π_h . It follows that a decrease in μ_f is needed to offset the effects of an increase in μ_h if strategies h and b are to remain equally profitable. We note further that the $\Pi_h = \Pi_b$ curve must have a slope less than 1 in absolute value.²² By an analogous argument, the curve $\Pi_f = \Pi_b$ also slopes downward in the figure, with a slope greater than 1 in absolute value.

In equilibrium, if all strategies are used, all must yield equal profits. So, the equilibrium is represented by the point E in figure 2. The figure also illustrates a constrained optimum at O . The constrained optimum maximizes W over the choice of μ in the presence of monopoly pricing of differentiated products. In the appendix, we show that the first-order conditions for a constrained maximum are satisfied when $\mu_h = \mu_f$.²³ The figure depicts some iso-welfare loci for successively lower levels of expected welfare as we move away from O . These curves are symmetric about the 45° line, thanks to the symmetry across countries.

It should be clear that if O falls on the 45° line, the constrained optimum can be achieved with a tax or subsidy on diversification alone, with $\varphi_f = \varphi_h = 0$. Such a policy shifts the equilibrium along the $\Pi_h = \Pi_f$ curve and thereby preserves the equality between μ_h and μ_f . What remains to be addressed is whether the government should encourage diversification with a subsidy for firms that form multiple relationships ($\varphi_b > 0$) or whether it should discourage diversification with a tax ($\varphi_b < 0$) on such firms. This amounts to the same question as to whether point O lies to the southwest of E along $\Pi_h = \Pi_f$ or whether it lies instead to the northeast of E .

We can answer these questions formally using methods similar to the ones we applied in section III. We begin with the planner's objective in (11). The wedge between social and private incentives for diversification is given by

²² The effect on expected profits of a diversified firm relative to a home-only firm conditional on state B is equal and opposite for a given increase in μ_h and comparable decrease in μ_f . But the increase in μ_h (and accompanying decrease in μ_b) gives an added boost to the relative profitability of diversification, because it raises the expected profits for a b firm if state F arises.

²³ This statement is valid for all HSA preferences. In the appendix, we also show that the welfare function $W(\mu)$ is globally concave when X takes a CES form and that the constrained optimum must have $\mu_h = \mu_f$ when preferences take the symmetric translog form. Evaluating the second-order conditions for more general HSA preferences is challenging, but it seems compelling that the planner would want equal numbers of firms with single relationships at home and abroad.

$$w_j^* := \Pi_j(\mu^*) - \Pi_b(\mu^*) - \frac{dW(\mu^*)}{d\mu_j}, \quad j \in \{h, f\},$$

where μ^* represents the allocation in the constrained optimum, and recall that $dG(\mu)/d\mu_j$ denotes the variation in any function $G(\mu)$ for $d\mu_j = -d\mu_b > 0$. With the first-order condition for the second-best allocation, this wedge is given by

$$w_j^* = - \sum_{i=h,f,b} \mu_i \frac{d\Pi_i(\mu^*)}{d\mu_j} - \frac{1}{\varepsilon - 1} \sum_{J=H,F,B} \delta^J \frac{d[P^J(\mu^*)^{1-\varepsilon}]}{d\mu_j}, \quad j \in \{h, f\}. \quad (19)$$

The first term on the right-hand side of (19) represents the business-stealing externality, that is, the change in other firms' profits that results from shifting a marginal firm from diversified sourcing to sole sourcing in country j . The second term represents the consumer surplus externality, that is, the change in consumer surplus that results from reduced product availability and higher prices in the three states. The difference from the analogous expressions in section III reflects the fact that the constrained policy maker needs to take account of not only the direct effects on profits and consumer surplus of changing the numbers of firms in each state (holding prices constant) but also the indirect effects on profits and consumer surplus that come from marginal adjustments in the markups. The unconstrained planner can ignore these latter effects when deciding μ^o , because the choice of optimal consumption subsidies ensures that the induced changes in markups have a negligible effect on aggregate utility.

In the appendix, we provide a general formula for the wedge w_j^* for an arbitrary share function $s(z)$ that satisfies assumptions 1 and 2. Then we turn to the symmetric case depicted in figure 2, where $w_h^* = w_f^* = w^*$. Point E lies above point O whenever $w^* > 0$ and below point O whenever $w^* < 0$.

The results for our two special cases of HSA preferences are instructive. First, with *symmetric CES preferences*, we find that $w^* > 0$ and thus $\varphi_b^* > 0$ for all $\sigma > 1$.²⁴ Recall that with optimal consumption subsidies in place, the optimal policy has $\varphi_b^o = 0$, because the extra consumer surplus generated by adding firms in a given state exactly matches the loss in aggregate profits. Now, with consumption subsidies unavailable to the policy maker, the monopoly pricing of differentiated products generates too little consumption of these goods relative to the numeraire good in the *laissez-faire* equilibrium. A subsidy for diversification increases the number of available products in every state, which reduces P^J for all $J \in \{H, F, B\}$, even though prices of marketed products do not change. The fall in the

²⁴ Equivalently, the planner can set $\varphi_b^* = 0$ and $\varphi_f^* = \varphi_h^* < 0$

price index in state J stimulates consumption of differentiated products in that state, thereby mitigating the consumption distortion.

Second, with *symmetric translog preferences*, we show that $w^* > 0$ if

$$\varepsilon > \frac{\theta\rho(2 - \rho)[1 + \theta\rho(2 - \rho)]}{1 + 3\theta\rho(2 - \rho)}$$

and $w^* < 0$ if

$$\varepsilon < \frac{\theta\rho(2 + \theta\rho)}{2(2 + 3\theta\rho)};$$

see lemma 7. Recall that a tax on diversification is needed to align social and private incentives for supply chain formation for any HSA preferences other than CES when a consumption subsidy is available to correct the distortion otherwise generated by markup pricing. In the absence of consumption subsidies, the tendency for firms to overinvest in resilience continues to figure in the policy maker's calculus, because the business-stealing externality is large relative to the consumer surplus externality. However, the distortion arising from monopoly pricing points in the opposite direction; there is too little consumption of differentiated products relative to the numeraire good, and a subsidy for diversification would boost consumption of these goods. When demand for differentiated goods is highly elastic, the distortion from monopoly pricing looms large, and the planner's imperative to encourage consumption of these goods outweighs her concern about firms' excessive investments in resilience, much as with CES preferences. In contrast, when demand for differentiated products is not so elastic, the welfare effects of the consumption distortion are muted, and the planner acts to dampen firms' excessive incentive to be present in the market.

Let us return now to the case in which input costs are lower abroad than at home ($q_H > q_F$) but foreign sourcing entails greater risk of disruption than home sourcing ($\gamma_H > \gamma_F$). Figure 3 depicts the laissez-faire equilibrium and the constrained optimum in such a setting for general HSA preferences. Again we consider fixed costs of sourcing relationships in the range that a positive measure of firms chooses each of the available investment strategies.

We observe first that for general HSA preferences, the constrained optimum, O , need not fall on any of the three equiprofitability curves. This means that, generically, the government cannot achieve the constrained optimum with a single policy instrument. For the case illustrated in figure 3, a subsidy for diversification improves welfare relative to E , but since such a policy preserves $\Pi_h = \Pi_f$, such a policy can achieve at best the utility associated with the iso-welfare curve through point D . A tax to discourage sourcing abroad ($\varphi_f < 0$) shifts the equilibrium to the left along the $\Pi_h = \Pi_b$ curve but at best can achieve the utility associated with




FIG. 3.—Equilibrium and constrained optimum for asymmetric case.

the iso-welfare curve through point F. Finally, a tax on onshore relationships with a single partner ($\varphi_h < 0$) can be used to achieve point H. For the scenario depicted in the figure, the constrained optimum could be achieved with a combination of a subsidy for diversification and a tax on sole-sourcing offshore or with a subsidy to diversification and a subsidy for sole-sourcing at home.

Although the constrained optimum can be characterized for particular preferences and parameters, in general the wedges w_h^* and w_f^* that dictate the second-best policy combinations can take any sign and a range of relative magnitudes. The reason for this reflects the complexity of the planner's constrained maximization problem. While the planner faces a general trade-off between alleviating the markup distortion by supporting greater resilience and mitigating the business-stealing externality, the size of this trade-off will vary across states of the world, and the planner

cannot separately address the trade-off state by state. For example, a policy that encourages greater diversification at the expense of sole-sourcing offshore generates an increase in the number of products available both in state H and in state B. As a result, the optimal policy is dictated by some weighted average of the trade-off between consumption distortion and business-stealing externality in each state of the world, and this depends on the exact form of preferences, the various preference parameters, and the sizes of the cross-country differences in costs and riskiness.

In section V, we resort to numerical methods to explore some of these trade-offs. Before that, we return briefly to the special case of symmetric CES preferences, for which a strong characterization of the second-best policies is possible even with asymmetric costs and risks. With CES preferences, the price index plays a dual role as both welfare metric and demand aggregator. We show in the appendix that this exceptional feature of the CES implies that the constrained optimum is characterized by $\Pi_h = \Pi_f$, much like the laissez-faire equilibrium. That is, the planner has no reason to tilt supplier relationships toward one location or the other. This means that the second best can be achieved with a single policy instrument, namely, a tax or subsidy for diversification. However, as we also show in the appendix, point O always lies below point E on the $\Pi_h = \Pi_f$ curve, so with CES preferences, it is always desirable for the government to promote resiliency with a subsidy to strategy b for all values of the cost and risk parameters. The explanation is the same as in the symmetric case; with CES preferences, the consumer surplus externality and the business-stealing externality exactly offset one another in every state of the world. What remains are the distortions that result from the fixed and positive markup of consumer prices over marginal costs. The constrained policy maker who cannot eliminate the consumption distortions directly can instead partially alleviate the distortion by promoting greater product availability in all states of the world.

We summarize our analytical findings about second-best supply chain policy in the following:

PROPOSITION 2. Suppose that consumption subsidies are infeasible. If consumers have symmetric CES preferences, the planner can achieve a constrained optimum with a single policy instrument, namely, a subsidy for diversification ($\varphi_b > 0$). If consumers have symmetric translog preferences and the input costs and disruption risks in the two countries are symmetric, the constrained optimum can again be achieved with a single policy, which must be a tax on diversification if

$$\varepsilon < \frac{\theta\rho(2 + \theta\rho)}{2(2 + 3\theta\rho)}$$

and a subsidy for diversification if

$$\varepsilon > \frac{\theta\rho(2-\rho)[1+\theta\rho(2-\rho)]}{1+3\theta\rho(2-\rho)}.$$

In other circumstances, two policy instruments are generally needed to alter both the incentives for diversification and the incentives for sourcing at home versus abroad.

V. Numerical Exploration of the Constrained Optimum

When monopolistically competitive firms form their supply chains with an eye to potential disruptions, the market equilibrium features several sources of inefficiency. The consumer surplus externality associated with product availability suggests underinvestment in resilience, whereas the business-stealing externality implies just the opposite. Meanwhile, monopoly pricing generates insufficient consumption of differentiated products relative to the numeraire good in realistic situations when fiscal policies cannot be used to align relative prices with relative marginal costs. We have been able to characterize the policy imperatives that these distortions create under CES preferences and, with more general HSA preferences, when the home and foreign suppliers are similar with respect to costs and risks. Armed with our understanding of the nature of the distortions, we turn now to numerical methods to explore the constrained optimal policies in situations when costs and risks differ in the two countries. To this end, we henceforth assume that preferences take the symmetric translog form.

Figure 4 depicts the constrained optimal fractions of firms (fig. 4A, 4C) and the policy wedges at the second-best allocation (fig. 4B, 4D) for two different values of ε , the elasticity of demand for differentiated products. The figure is drawn for the case when production costs are similar in the two countries ($q_H = q_F$), but we show in the appendix (see figs. 6, 7) that qualitatively similar patterns emerge when costs differ. Figure 4 illustrates the comparative statics of the equilibrium, constrained optimum, and optimal supply chain policies with respect to variation in the cross-country risk differential.

In figure 4A and 4B, we see the outcomes for a relatively low value of ε , namely, $\varepsilon = 1.2$. Figure 4A shows the fraction of firms that choose each of the investment strategies in the laissez-faire equilibrium (solid curves) and in the constrained optimum (dashed curves). When $\gamma_H = \gamma_F$ on the left side of the panel, the market equilibrium features excess investment in resilience ($\mu_b > \mu_b^*$) and insufficient investment in exclusive supply relationships ($\mu_h^* = \mu_f^* = \mu^* > \mu$). These numerical outcomes mirror the theoretical results from section IV. They reflect the fact that under MS�D, the business-stealing effect dominates the consumer surplus




FIG. 4.—Second-best policies: risk differences across locations. Baseline simulation has $\gamma_H = \gamma_F = 0.9$, $q_H = q_F = 0.1$, $\theta = 8.0$, and $\rho = 0.7$. Low and high epsilon correspond to $\varepsilon = 1.2$ and $\varepsilon = 1.7$, respectively. Fixed cost is chosen so that $\min(\mu_b^*, \mu_b^0) \approx 0$ in the baseline symmetric simulation. This yields $k = 0.13$ for $\varepsilon = 1.2$ and $k = 0.37$ for $\varepsilon = 1.7$. The cost discount is computed as $-(\gamma_F - \gamma_H)/\gamma_H$, where we keep γ_H constant at its baseline value.

effect. Moreover, with ε relatively small, the consumption distortion caused by markup pricing is not too severe. In figure 4B, we see that $w_h^* = w_f^* = w^* < 0$, so the constrained optimum can be achieved either with a tax on firms that diversify or with equal subsidies to firms that invest in exclusive supply relationships either at home or abroad. As we increase the risk of disruption in F—so that the safety premium in the home country rises—both the competitive equilibrium and the constrained optimum are characterized by greater fractions of diversified firms and greater fractions of firms that form relationships only onshore. These findings are intuitive, but what is less obvious is what happens to the wedges between social and private incentives. In figure 4B, we see that w_h^* rises while w_f^* falls. This implies, for example, that second-best subsidies for exclusive offshore relationships grow (φ_f^*) while subsidies for sourcing relationships at home (φ_h^*) shrink, if the planner eschews taxes or subsidies on diversification ($\varphi_b^* = 0$).

How do we understand this finding? In figure 4A, we see that the fraction of firms with sourcing relationships exclusively in the home country rises above the fraction with sourcing relationships exclusively offshore.

The increase in product diversity and in competition in state H relative to state F generates a decline in the price index P^H relative to P^F . But the monopoly distortion is more severe when the price index is high, so the consumption shortfall is greater in state F than in state H. The planner wishes to combat the higher price index in state F with a policy that tilts sourcing toward the foreign country.

As the foreign country becomes even riskier, the planner continues to discourage diversification; we continue to find the second-best fraction of diversified firms, μ_b^* , below the free market level. But the social cost of the market's misallocation between home sourcing and foreign sourcing also grows, and so the gap between the two wedges w_h^* and w_f^* widens. At some risk differential close to 20% in the figure, the planner's desire to shift the location of exclusive sourcing from the home country to the foreign country implies a second-best tax on onshore relationships, combined with an even larger subsidy for investing in a single relationship abroad.²⁵

The situation is similar for larger values of ε , such as depicted in figure 4C and 4D, except in one important respect. With a more elastic demand for differentiated products, the misallocation generated by markup pricing weighs more heavily in the planner's calculus compared with the net effect of the business-stealing and consumer surplus externalities. The optimal policy in the symmetric environment entails a net subsidy to diversification, which can be achieved with $\varphi_b > 0 = \varphi_h = \varphi_f$ or with $\varphi_h = \varphi_f < 0 = \varphi_b$. As the risk differential grows, the planner once again tilts policy in favor of exclusive sourcing relationships abroad to offset the increasingly deleterious effects of underconsumption of differentiated products when foreign supply is disrupted. For a sufficiently great probability of supply disruption in the foreign country, the planner subsidizes strategy f while still ensuring that the net effect of supply chain policy is to induce more diversification and greater resilience.

Notice too the scale of the optimal subsidies. Recall from section III that when able to implement the optimal state- and product-contingent consumption subsidies, the planner taxes diversification (or subsidizes the two strategies involving exclusive relationships) at 100% of the fixed cost k , regardless of the configuration of cost and risk parameters.²⁶ In the second best described here, the impetus to tax diversification in order to dampen incentives for business stealing is offset by an urge to subsidize diversification to stimulate consumption of differentiated goods. The offsetting forces result in second-best policies that are an order of magnitude smaller than in the first best.

²⁵ Alternatively, the planner can achieve the same allocation with a subsidy for diversification and an even larger subsidy for offshore relationships (and $\varphi_h = 0$), so that μ_c grows at the expense of both μ_h and μ_b .

²⁶ This statement applies to situations, as here, with symmetric translog preferences.




FIG. 5.—Second-best policies: cost differences across locations. Baseline simulation has $\gamma_H = \gamma_F = 0.9$, $q_H = q_F = 0.1$, $\theta = 8.0$, and $\rho = 0.7$. Low and high epsilon correspond to $\varepsilon = 1.2$ and $\varepsilon = 1.7$, respectively. Fixed cost chosen so that $\min(\mu_h^*, \mu_f^*) \approx 0$ in the baseline symmetric simulation. This yields $k = 0.13$ for $\varepsilon = 1.2$ and $k = 0.37$ for $\varepsilon = 1.7$. The risk premium is computed as $-(q_F - q_H)/q_H$, where we keep q_H constant at its baseline value.

Figure 5 depicts the comparative statics with respect to foreign production costs, holding riskiness in the two locations constant (and, in this figure, equal to one another).²⁷ The symmetric equilibrium again requires a second-best tax on diversification when ε is small (fig. 5A, 5B) and a subsidy when ε is larger (fig. 5C, 5D). A fall in the cost of producing inputs abroad, which expands the cost discount in the offshore location, reduces the price index in state F relative to that in state H.²⁸ Thus, the social benefit from promoting consumption in state H comes to exceed that in state F. In figure 5B, the planner alters the composition of exclusive supply relationships by offering a larger subsidy for onshore sourcing than for offshore sourcing (or, equivalently, a subsidy for onshoring combined

²⁷ The parameters used for this figure are the same as for fig. 4, so the baseline (symmetric) outcome is the same at the left-most point in both figures.

²⁸ In this case, there are two reasons: more firms choose strategy f than strategy h, and hence the market is more competitive in state F than in state H; and products containing inputs produced in F bear a lower cost than those produced in H.

with a tax on diversification). For a high enough cost discount, the wedge for offshore relationships actually turns positive. In figure 5C, with more elastic demand, the planner encourages resilience at the expense of exclusive relationships at home and abroad. As seen in figure 5D, the taxes on single relationships that are used to encourage diversification diverge; onshore relationships face a lower tax than their offshore counterparts (or else the planner can subsidize both diversification and onshoring). This policy combination reflects the fact that consumption distortion is more harmful in state H than in state F.

When both cost and risk parameters differ in the two possible locations for producing inputs, the ranking of the price index in states H and F is less clear-cut. Greater risk of supply disruption in the foreign country discourages private investment there, contributing to a relatively higher price index in state F. But lower production costs offshore raises the relative attractiveness of strategy f compared with strategy h, and it also has a direct effect on comparative prices in the two states, reflecting the relatively lower foreign unit cost. The relative size of the wedges, w_h^* versus w_f^* —and thus the net effect on the incentives for exclusive offshoring versus exclusive onshoring—hinges on the relative strength of these forces.

VI. Conclusion

Global supply chain disruptions are increasingly salient and often costly. Many commentators have been quick to conclude that governments ought to be doing something to promote greater market resilience. But the welfare theoretic calculus around government intervention is rather subtle. Private actors have a clear self-interest in taking measures to avoid disruptions to their production processes. Only when the private incentives for resilience fall short of the social benefits will government encouragement be warranted. Pointing in that direction is the observation that consumers capture part of the surplus created by the ongoing availability of firms' products. But firms also have an incentive to be in a position to reap extra profits when their rivals are suffering. The temptation for business stealing suggests that excess resilience is also a possible market outcome.

Surprisingly little research has addressed the desirability of government policy to promote resilience or to encourage sourcing from safer locations. In this paper, we have taken a first step. We have proposed a simple framework in which the supply of any product requires the availability of a critical input. Idiosyncratic and aggregate shocks can disrupt firms' relationships with their suppliers. Firms face the choice of where to develop a relationship and whether to protect their operation with backup sources of supply. We study the simplest case of two potential supply sources, one at home and one abroad, and focus on a situation where domestic sourcing is costlier than sourcing abroad but also less risky.

Since consumer gains from product availability reflect their preferences, the form of demand plays a critical role in the policy calculus. The CES demand system is popular and tractable for analysis such as ours. But it also introduces restrictions that color the findings. We allow for a CES utility function but also for a broader class of preferences that Matsuyama and Ushchev (2017, 2020b) have developed and termed HSA. The more general preferences admit nonconstant markups and, in particular, application of MSLD.

Our analysis yields several broad lessons. First, the government generally needs at least two supply chain policies to achieve efficient sourcing (first or second best). One instrument regulates the margin between sourcing from one location or two. The other guides the choice between sourcing at home and abroad. For example, the government might subsidize or tax supply chain diversification while subsidizing or taxing firms that source only at home. Or it might subsidize or tax diversification while subsidizing or taxing offshoring.

When preferences take the CES form, the first best can be achieved simply with a state-independent consumption subsidy and with no interference in supply chain organization. The second best requires a subsidy to diversification but no bias for home versus foreign sourcing. But with more general forms of HSA preferences that obey MSLD, the planner requires state- and product-specific consumption subsidies to achieve the first best, along with a tax on firms that diversify, and a policy that tilts sourcing to one country or the other depending on the relative sizes of the consumer surplus externality and the business-stealing externality. In the second best, the optimal policies are qualitatively similar to the ones that achieve the first best if the elasticity of demand for differentiated products is small, but a large demand elasticity tilts the policy toward smaller taxes or even subsidies for diversification.

Needless to say, there are many ways that our analysis could be enriched. For example, we could introduce a richer technology with potential substitution between manufactured inputs and primary factors of production. We could entertain more complex supply chains, with multiple inputs and with a sequencing of them such that some inputs enter the production process upstream from others. We could allow for dynamics, which would render inventories an additional tool for firms to invest in resilience and give governments additional policy instruments, such as stockpiling supplies or allowing accelerated depreciation of inventory costs. We could introduce political economy considerations that might drive a wedge between the parameters that capture the risk aversion of managers versus that of policy makers. We see all of these potential extensions as worthwhile and germane to the ultimate policy assessment. We believe that our simpler setting suggests a way to pose the question and that our analysis provides a proof of concept.

References

- Arad, Ruth W., and Arye L. Hillman. 1979. "Embargo Threat, Learning and Departure from Comparative Advantage." *J. Internat. Econ.* 9:265–75.
- Baldwin, Richard, and Rebecca Freeman. 2022. "Risks and Global Supply Chains: What We Know and What We Need to Know." *Ann. Rev. Econ.* 14:153–80.
- Baqae, David, and Emmanuel Farhi. 2020. "Entry vs. Rents: Aggregation with Economies of Scale." Working Paper no. 27140, NBER, Cambridge, MA.
- Bergström, Clas, Glenn C. Loury, and Mats Persson. 1985. "Embargo Threats and the Management of Emergency Reserves." *J.P.E.* 93 (1): 26–42.
- Bhagwati, Jagdish N., and T. N. Srinivasan. 1976. "Optimal Trade Policy and Compensation under Endogenous Uncertainty: The Phenomenon of Market Disruption." *J. Internat. Econ.* 6:317–36.
- Bilbiie, Florin O., Fabio Ghironi, and Mark J. Melitz. 2012. "Endogenous Entry, Product Variety, and Business Cycles." *J.P.E.* 120 (4): 304–45.
- Campolmi, Alessia, Harald Fadinger, and Chiata Forlati. 2021. "Trade and Domestic Policies in Models with Monopolistic Competition." Working paper.
- Cheng, Leonard K. 1989. "Intermittent Trade Disruptions and Optimal Production." *Internat. Econ. Rev.* 30 (4): 753–74.
- Dhingra, Swati, and John Morrow. 2019. "Monopolistic Competition and Optimum Product Diversity under Firm Heterogeneity." *J.P.E.* 127 (1): 196–232.
- Diewert, W. E. 1974. "Applications of Duality Theory." In *Frontiers of Quantitative Economics*, vol. 2, edited by M. Intriligator and D. Kendrick, 106–76. Amsterdam: North Holland.
- Dixit, Avinash K., and Joseph E. Stiglitz. 1977. "Monopolistic Competition and Optimum Product Diversity." *A.E.R.* 67 (3): 297–308.
- Elliot, Matthew, Benjamin Golub, and Matthew V. Leduc. 2022. "Supply Network Formation and Fragility." *A.E.R.* 112 (8): 2701–47.
- Fajgelbaum, Pablo, Pinelopi Goldberg, Patrick Kennedy, and Amit Khandelwal. 2020. "The Return to Protectionism." *Q.J.E.* 135 (1): 1–55.
- Feenstra, Robert C. 2003. "A Homothetic Utility Function for Monopolistic Competition Models, without Constant Price Elasticity." *Econ. Letters* 78 (1): 79–86.
- Iakovou, Eleftherios, and Chelsea C. White III. 2020. "How to Build More Secure, Resilient, Next-Gen U.S. Supply Chains." Commentary, TechStream, Brookings Inst., Washington, DC.
- Matsuyama, Kiminori. 2022. "Non-CES Aggregators: A Guided Tour." CEPR Discussion Paper no. 17610, Center Econ. Policy Res., Washington, DC.
- Matsuyama, Kiminori, and Philip Ushchev. 2017. "Beyond CES: Three Alternative Classes of Flexible Homothetic Demand Systems." CEPR Discussion Paper no. 12210, Center Econ. Policy Res., Washington, DC.
- . 2020a. "Constant Pass-Through." CEPR Discussion Paper no. 15475, Center Econ. Policy Res., Washington, DC.
- . 2020b. "When Does Procompetitive Entry Imply Excessive Entry." CEPR Discussion Paper No. 14991, Center Econ. Policy Res., Washington, DC.
- . 2022. "Destabilizing Effects of Market Size in the Dynamics of Innovation." *J. Econ. Theory* 200:105415.
- Mayer, Wolfgang. 1977. "The National Defense Tariff Argument Reconsidered." *J. Internat. Econ.* 7:363–77.
- Lund, Susan, James Manyika, Jonathan Woetzel, et al. 2020. "Risk, Resilience, and Rebalancing in Global Value Chains." McKinsey Global Inst.
- Mrázová, Monika, and J. Peter Neary. 2017. "Not So Demanding: Demand Structure and Firm Behavior." *A.E.R.* 107 (12): 3835–74.

- Shih, Willy C. 2020a. "Global Supply Chains in a Post-Pandemic World." *Harvard Business Rev.* 98 (5): 82–9.
- . 2020b. "Is It Time to Rethink Globalized Supply Chains?" *MIT Sloan Management Rev.* 61 (4): 16–18.
- Zhelobodko, Evgeny, Sergey Kokovin, Mathieu Parenti, and Jacques-François Thisse. 2012. "Monopolistic Competition: Beyond the Elasticity of Substitution." *Econometrica* 80 (6): 2765–84.