Trade and Technology Compatibility in General Equilibrium*

Weili Chen† Jingting Fan‡ Wenlan Luo§
Jinan IESR CKGSB Tsinghua University

This Draft: July 2026
First Draft: November 2023

Abstract

We develop a trade model where the horizontal proximity between a firm's technology and that of its suppliers shapes input efficiency. Trade policies, by expanding or restricting firms' access to foreign suppliers, influence these horizontal technology choices and, through input-output linkages, affect the technological interdependence between countries. We characterize these effects and test them using technological similarity measures derived from patent texts. Quantification of the model shows that an electronics and semiconductor embargo by the U.S. against China triggers technological decoupling between the two countries and realignment among other countries. These endogenous technological shifts double the global welfare loss caused by the embargo.

Key words: technology proximity; trade conflict; technology decoupling; production network; geopolitics

*We thank Jonathan Eaton, Lorenzo Caliendo, Thomas Chaney, Ana Cecilia Fieler, Agustín Gutiérrez, Andrei Levchenko, Nuno Limao, Ernest Liu, Dan Lu, Fernando Parro, Michael Zheng Song, Kjetil Storesletten, Chang Sun, Daniel Yi Xu, and Stephen Yeaple for helpful conversations. We also thank seminar and conference participants at Barcelona Summer Forum (geography, trade, and growth), CKGSB, EIIT (Purdue), Fudan Fanhai, Georgetown, Hong Kong University, HKUST, INSEAD, National University of Singapore, North Carolina State University, Peking University (Guanghua; HSBC), the University of Nottingham, Tsinghua Growth and Institutions Conference 2023, Yale, the Asian Bureau of Economic Research Conference 2024 (trade and growth), and China International Conference in Macroeconomics 2024 for helpful comments.

†The Institute for Economic and Social Research, Jinan University, weilichen.econ@gmail.com
‡The Cheung Kong Graduate School of Business, jtfancn@gmail.com
§Tsinghua University, luowl@sem.tsinghua.edu.cn

1 Introduction

Recent global events have driven major economies to adopt strategies aimed at reducing their economic engagement with geopolitical rivals. For instance, the European Union is diversifying its sources of essential goods to mitigate supply chain disruption risks, while the U.S. has implemented policies to limit China's access to advanced technology. Through initiatives like 'friend-shoring' and 'Made in China 2025,' both the U.S. and China are working to decrease reliance on critical goods from each other.

These shifts raise a growing concern: reduced access to, and dependence on, key foreign technologies may push firms toward alternative technological paths, potentially dividing the world into blocs of countries with incompatible technological systems. For example, U.S. export controls on advanced AI processors are pushing Chinese firms away from Nvidia's chips, and the CUDA software ecosystem built on them, toward Huawei's accelerators.[1] Similarly, concerns over the supply of cobalt, a critical battery input refined mainly in China, have accelerated the shift toward cobalt-free batteries. Both shifts force adaptations throughout the supply chain: AI developers must redesign their models and tools around the alternative ecosystem, while cobalt-free batteries, which charge at different voltages, require updates to charging protocols and even the design of charging stations and cars.[2] These examples illustrate how firms' technological choices are shaped by access to foreign inputs in their supply chains, and how policies and supply risks affecting this access can either promote technological convergence across global supply chains or fragment them into incompatible systems.

This paper makes two contributions. First, we construct a tractable model to formalize the link between trade and the horizontal technological choices of firms. This enables us, for the first time in the literature, to assess whether a trade shock leads to technological divergence or convergence. Second, using a new measure of technology similarity based on the texts of global patents, we test and quantify the model. We apply the model to assess the impacts of a U.S. embargo on Chinese electronics and semiconductor imports. We find that the embargo leads to technological decoupling between the two countries and a realignment of others. These endogenous changes exacerbate—rather than alleviate—the welfare losses from the embargo, accounting for nearly all of the U.S. losses and half of the global losses.

In our model, firms differ in vertical efficiency and their endowed horizontal technology, denoted by a point $\bar{\theta}$ in technology space $\mathbb{T}$. Following the seminal work of Lancaster

[1] Since 2022, U.S. export controls have restricted Chinese firms' access to the most advanced AI processors. Huawei's Ascend accelerators, the leading domestic alternative, use a software system incompatible with CUDA.

[2] Cobalt-free lithium-iron-phosphate cells, the leading alternative, charge to 3.65 volts rather than 4.2 and discharge at a nearly flat voltage. On manufacturers' shift toward cobalt-free batteries, see https://www.cnbc.com/2021/11/17/samsung-panasonic-and-tesla-embracing-cobalt-free-batteries-.html, accessed February, 2024.

(1979), we interpret each point in $\mathbb{T}$ as a specific 'technological path' composed of particular technical components or engineering features. Points that are closer in $\mathbb{T}$ correspond to more similar (and hence more compatible) technologies. We label technologies 'horizontal' because no technology is inherently superior to another; rather, they differ by their distance from one another. A firm's endowed technology $\bar{\theta}$ represents the path most suited to its intrinsic capabilities, shaped by factors outside our model such as accumulated know-how or key expertise of the firm. The distribution from which firms draw $\bar{\theta}$ varies across countries and sectors, reflecting countries' comparative advantage over different technologies.

Given its endowed $\bar{\theta}$, each firm chooses a production technology $\theta \in \mathbb{T}$. Choosing a technology closer to the ones used by key suppliers reduces effective input costs by enhancing compatibility, but deviating from $\bar{\theta}$ also entails an adaptation cost that grows with the distance $\|\theta - \bar{\theta}\|$. After deciding on θ, firms sample a set of suppliers and choose the one that offers the best compatibility-adjusted price for each intermediate input. They then produce and sell the output to final consumers and downstream firms. In equilibrium, given the country-specific *exogenous* distributions governing $\bar{\theta}$, firms anticipate the choices of all other firms and make their technology and supplier choices accordingly.

The equilibrium is characterized by a fixed point in the distributions of chosen technologies across firms in all countries and sectors—an infinite-dimensional object. We establish that there generically exists an equilibrium in which firms' technology choices vary continuously with their technology endowment, and that this equilibrium is unique provided the cost of technology adaptation is not too small and the benefit of compatibility is not too large. In the limiting case where endowment distributions are degenerate and compatibility yields no benefit, our model simplifies to the model of Caliendo and Parro (2015). Thus, the model retains the flexibility of canonical quantitative trade models.

Our model uncovers novel relationships between trade and technology, which we formalize and bring to the data. First, when compatibility drives firms' decisions, those whose endowed technologies are closer to a particular foreign country's technological profile are more likely to adopt similar technologies and import from that country. This results in a *correlation* between a firm's importing patterns and its technology similarity with exporting countries. Second, lowering tariffs with a specific trading partner incentivizes firms in the importing country to shift their technologies toward that partner's technologies, implying a *causal* link between bilateral trade costs and technological convergence.

On the normative side, the model highlights a key externality in technology choice. When a firm chooses its technology, it also affects the production costs of domestic and foreign downstream firms, who may face lower production costs if this upstream firm chooses a closer technology, an effect not internalized by the firm. Such externalities exist both within and across borders, are amplified by input-output linkages, and cannot be fully internalized by a monopolistic-competition markup. In a one-sector, two-symmetric-country special case

of the model, we show analytically that starting from the decentralized equilibrium, shifting one country's technology toward the other's increases welfare in *both* countries. Thus, firms tend to stay too close to their endowed technology. For this reason, the welfare effects of trade cost changes extend beyond the direct first-order effect captured by Hulten's theorem: because equilibrium technology choices are socially suboptimal, the induced technology response has a first-order welfare effect of its own. Between two symmetric countries, this effect is an amplification: liberalization brings technological convergence and additional gains, while a trade conflict brings divergence and additional losses. In a multi-country, multi-sector world, whether the technology response amplifies or dampens the direct effect, and by how much, depends on the entire map of trade and technology linkages.

We test and quantify the effect of the endogenous technology adaptation. Central to these exercises is a new measure of horizontal technology similarity across firms and countries. To construct this measure, we use the texts of worldwide patents and employ a state-of-the-art large language model to generate numerical vectors, or 'embeddings,' that summarize each patent's content. We then compute bilateral similarities at various levels of aggregation using the cosine similarity of these embeddings. We validate our measure by comparing it with other proxies of bilateral similarities. Similar technologies, however, need not be compatible ones: high similarity can also reflect firms competing head-to-head in the same technological space. To isolate compatibility, we further construct a technology *competition* score, using an embedding model finetuned for this purpose, and control for it throughout our empirical analyses. We also demonstrate that our empirical results remain robust when patent citations are used as a proxy for technology similarity.

We find that, consistent with the model's key prediction, lowering bilateral trade barriers increases bilateral technology similarity. Exploiting exogenous variation in HS-6-level tariffs arising from changes in most-favored-nation (MFN) rates, which affects some partners (those subject to the MFN tariffs) but not others, we estimate a negative and economically meaningful elasticity of bilateral technology similarity with respect to tariffs.[3] Beyond providing support for the model's prediction, this estimate also disciplines the strength of the compatibility mechanism in our quantitative exercises.

We parameterize the model using data on 29 major countries and 19 sectors. For tractability, we let $\mathbb{T}$ be the real line and assume that firms in each country and sector draw their endowment technology from a respective Normal distribution. We show that under a quadratic approximation to firms' technology choice problem, the chosen technologies in each country-sector also follow a Normal distribution, with its mean and variance determined by general equilibrium forces. This structure makes the model computationally tractable.

Calibration of the model requires mapping the horizontal technologies of countries and

[3]MFN tariffs can be endogenously chosen to promote trade. Following Boehm et al. (2023), we address this concern by excluding from the sample the largest exporter in each importer-product pair.

sectors to the real line. We choose the technology positions by matching the model's implied pairwise technology similarities to those obtained from patent texts.[4] In addition, two parameters are key in shaping the impacts of technology choice: the cost of incompatibility with suppliers and the cost of deviating from a firm's endowed technology. We identify these parameters by matching two reduced-form estimates: (i) the correlation between import behavior and technology similarity, which we estimate from matched data on Chinese manufacturing firms, their patents, and their customs transactions, and (ii) the elasticity of technology similarity with respect to tariffs. Using these parameters and the calibrated technological distribution of countries, the model implies a component of trade costs reflecting technology incompatibility. Residual iceberg trade costs are then chosen to match observed bilateral trade shares.

Our calibration suggests that technology compatibility is an important driver of sourcing decisions. The ad-valorem equivalent trade cost due to incompatibility is around 10%. Standard estimation suggests that to account for the observed bilateral trade (or the lack thereof), trade costs need to be well above 100%, often raising the question on the source of such high costs (Anderson and van Wincoop, 2004). Our model provides a partial explanation: technological incompatibility between countries.

We assess the impact of a trade conflict between the U.S. and China by implementing a counterfactual experiment, in which the U.S. and its allies impose an embargo on Chinese electronics and semiconductor imports. We find that the embargo leads to a 0.55% welfare loss for China, a 0.05% loss for the U.S., and a global loss of 0.23%. In response to the embargo, the technologies of firms in China diverge from those of other countries, leading to higher input costs for foreign downstream firms and triggering realignments of technologies in other countries, including the U.S. Contrary to the intuition that substitution adjustments to negative shocks alleviate welfare loss, the divergence between China and the U.S. amplifies the loss for both, due to the externality in technological choice. Almost all of the U.S. welfare loss and half of the global welfare loss are due to endogenous technological changes.

Our paper engages with several strands of the literature. The central idea—that firms' technology choices are shaped by compatibility requirements—has been explored in the trade literature. However, existing models often focus on stylized settings with two countries and two technologies (e.g., domestic vs. foreign), limiting their quantitative generalizabil-

[4] We do not seek to explain, for example, why the auto sector uses more similar technologies to the auto parts sector than to the apparel sector; rather, we are interested in, say, whether the U.S. auto sector is more similar to Germen or Canadian auto sector. Consequently, we remove sector-pair fixed effects from our calibration targets, isolating cross-country variation.

ity.[5,6] We allow for multiple technologies, countries, and input-output linkages. Our model builds on Lancaster (1979)'s work on endogenous product differentiation, but emphasizes compatibility with suppliers rather than consumers.

Our paper is related to recent studies on production networks and trade, such as Jones (2013), Oberfield (2018), Boehm and Oberfield (2020), Demir et al. (2024), Eaton et al. (2022), Dhyne et al. (2023) and Baqaee and Farhi (2024). Our focus on technology choice in a network relates to Demir et al. (2024), which studies network formation under (vertical) quality choice. Our contribution to this literature is to incorporate horizontal technology choice into a model of trade and production networks. By modeling endogenous technology choice, which shapes supplier relationships, our work also relates to Arkolakis et al. (2025), who endogenize the trade network through costly relationship formation. They show that the welfare effects of trade-cost changes remain proportional to observed trade flows, scaled by a coefficient common to all trading pairs. In our model, because adapting to one supplier's technology shapes compatibility with all other suppliers, the welfare effects of trade shocks are pair-specific, determined by each pair's position in the global map of trade and technologies.

The trade aspect of our model builds on the quantitative trade literature (e.g., Eaton and Kortum, 2002; Caliendo and Parro, 2015; see Costinot and Rodriguez-Clare, 2014 for an early review), with two main differences. First, it incorporates horizontal technology choice, which we show is important in shaping the effects of trade policies. To accommodate this mechanism, in our model, heterogeneous firms choose among a continuum of options and then interact with the entire distribution of other firms. We develop tools for establishing the existence and uniqueness of equilibrium in such environments, which differ from those in standard trade models where the equilibrium can be characterized solely by aggregate quantities and prices (e.g., Eaton and Kortum, 2002; Chaney, 2008; see Allen et al., 2024 for uniqueness results for some of these models).[7] Second, we provide a tractable extension of canonical quantitative trade models that allows firms to source inputs based on flexible interactions between buyer and seller characteristics.

[5]For example, Carluccio and Fally (2013) examine how supplying foreign firms using a 'modern' technology can reduce input availability for domestic firms using a 'traditional' technology; Costinot (2008) develops a two-country model to explore how horizontal standardization affects welfare. Given that there are multiple countries in the world economy, each with potentially different technologies, mapping these technologies to a binary choice in the model can be subjective.

[6]The concept of compatibility-driven technology choice has also been studied outside of trade. For example, Akcigit et al. (2016) present a growth model where firms sell ideas that diverge from their core competence, Lazear (1999) models language choice driven by communication incentives, and Basu and Weil (1998) model technology transfer considering whether the technology is appropriate for a country.

[7]Other trade models in which firms interact with the entire distribution of firms include those based on the assignment model, e.g., Costinot and Vogel (2010). Such models often imply efficient allocation, so welfare theorems can be invoked to establish the existence and uniqueness of equilibrium. This approach does not apply to our model with externalities.

Our paper also connects with the literature on the relationship between trade and technology (e.g., Eaton and Kortum, 1999; Buera and Oberfield, 2020; Liu and Ma, 2021; Aghion et al., 2021; Cai et al., 2022; Lind and Ramondo, 2023), which often uses patents to measure technology diffusion. While most existing studies focus on how trade affects vertical technology (e.g., aggregate or sectoral productivity), we focus on horizontal technology choices.[8] Motivated by this focus, we construct a new, theory-consistent, measure of technology similarity for empirical and quantitative analyses. We provide evidence for the compatibility mechanism and offer a structural interpretation of the evidence, complementing existing studies emphasizing learning or other forms of technology diffusion.

Finally, recent literature has estimated the impacts of the trade conflict between the U.S. and China, focusing on the short-run effects (e.g., Amiti et al., 2019; Fajgelbaum et al., 2020; Huang et al., 2023). Our finding suggests that an important part of the welfare losses will appear as technological adjustments take effect. Our focus on the non-tariff aspects of trade conflicts also relates to works on the interaction between trade and international relations, e.g., see Kleinman et al. (2024), Bonadio et al. (2024), and Clayton et al. (2023).

2 Model

2.1 Environment.

There are N countries, denoted by d or o, and S sectors, denoted by i or j. When referring to directional flows, we use (d,i) to denote the buying country-sector pair and (o,j) to denote the selling country-sector pair.

In each country d, a representative household supplies L_d units of labor, and a representative producer makes a *non-tradable* final good with the following production technology:

$$Q_d = \prod_{j=1}^{S} \Big[\sum_{o=1}^{N} \int_0^1 [q_{do}^j(\omega)]^{\frac{\eta-1}{\eta}} \mathrm{d}\omega \Big]^{\frac{\eta}{\eta-1}\cdot\rho_d^j}, \tag{1}$$

where $q_{do}^j(\omega)$ denotes the quantity of intermediate good ω produced in country o sector j that is used to produce the final good in d. The parameter $\eta > 1$ is the elasticity of substitution between different varieties, and ρ_d^j is the share of sector j in the final good. This final good is used for domestic consumption and domestic firms' technology adaptation.

In each country-sector pair (d,i), there is a continuum of intermediate-good firms denoted by $\nu, \omega \in [0,1]$. Each firm produces a differentiated variety using labor and inputs purchased from other intermediate-good firms. We assume that intermediate-good firms engage in monopolistic competition when selling to final-good producers—thereby earn-

[8] Lind and Ramondo (2023) study how trade affects the similarity of technologies among countries, modeling similarity through the correlation in productivity draws. In our model, technology similarity is an outcome of firms' choices.

ing positive profits—and charge marginal production costs when selling to downstream intermediate-good firms.[9] Before production, intermediate-good firms choose a horizontal technology, balancing the desire to use the technology they know best (their endowed technology) and the need to be compatible with key suppliers.

Let $\mathbb{T}$ be a metric space containing all technologies—such as a circle, the real line, or a finite-dimensional Euclidean space.[10] Each firm ν is endowed with a technology $\bar{\theta}(\nu) \in \mathbb{T}$. Given $\bar{\theta}(\nu)$, the firm selects a production technology $\theta(\nu) \in \mathbb{T}$. Deviating from $\bar{\theta}(\nu)$ incurs adaptation costs, yet may yield benefits by enhancing compatibility with certain suppliers. After choosing $\theta(\nu)$, the firm purchase input from suppliers based on the compatibility-adjusted input costs of these suppliers. Firms then produce and sell the output to other firms and final-good producers.

Technologies in our model have broad interpretations. They might represent distinct technical ecosystems—such as Android versus iOS for smartphones—or capture differences in marketing strategies, business networks, or political alignments among firms and countries. They may also embody product designs that rely on different scientific and engineering disciplines. For example, makers of internal combustion engine cars and their suppliers tend to rely predominantly on mechanical engineering, whereas makers of electric vehicles and their suppliers draw more on advances in chemistry and electronics. The same example illustrates the compatibility motive central to our model: as battery technology advances rapidly, a carmaker can tap these upstream gains only by shifting its own knowledge base toward electric-vehicle technologies, moving away from its accumulated mechanical-engineering expertise. Although our theoretical framework does not require a precise definition of technology, our empirical and quantitative analyses operationalize technology as the knowledge embedded in patent texts, using patent similarity measures to assess technological relationships between countries.

In the remainder of this section, we first describe firms' production and input-sourcing decisions, taking their chosen technology as given. We then describe how firms select their technology and characterize the equilibrium for technology choice.

[9] This assumption implies that the production equilibrium is efficient given technology choices, so that any inefficiency arises solely from technology choices. In Subsection 3.3, we show that our key results, including the efficiency implications, continue to hold when firms also charge monopolistic markups on sales to downstream firms.

[10] For most of the theoretical characterizations and later quantitative analysis, we restrict to the real line. We discuss the implications of focusing on the real line in Section 2.4.

2.2 Production and Sourcing Decisions: Intermediate-Good Firms

Consider a firm ν in country-sector (d, i) that has already chosen technology $\theta \equiv \theta(\nu)$. The firm's output is given by

$$y(\nu) = A_d^i [l(\nu)]^{\gamma^{iL}} \prod_j \left[m^j(\nu)\right]^{\gamma^{ij}}, \text{ with } \gamma^{iL} + \sum_j \gamma^{ij} = 1,$$

where $l(\nu)$ is the labor hired by firm ν, A_d^i is sectoral productivity, and the sector-j input $m^j(\nu)$ is a CES bundle of the varieties produced by intermediate-good firms around the world, with elasticity of substitution $\xi > 1$:

$$m^j(\nu) = \left(\sum_o \int_0^1 [m_o^j(\nu, \omega)]^{\frac{\xi-1}{\xi}} \mathrm{d}\omega \right)^{\frac{\xi}{\xi-1}}. \tag{2}$$

The parameter ξ implies a trade elasticity of $\xi - 1$ for intermediate-good firms. In this formulation, each firm bundles the varieties of all suppliers, so we abstract from firm heterogeneity in supplier access and from the extensive margin of network formation. In Appendix A.1, we show that the model is isomorphic to one with a random production network, in which firms draw production techniques with Pareto-tailed efficiencies and randomly sampled supplier sets with bilateral match-quality draws, and purchase each input only from the supplier offering the lowest compatibility-adjusted cost, as in Oberfield (2018) and Boehm and Oberfield (2020). That alternative microfoundation, in which sourcing adjusts entirely along the extensive margin, leads to the same aggregate equilibrium conditions.

Since firms sell to other firms at marginal cost, the compatibility-adjusted cost to firm ν of the variety of a supplier ω in (o, j) is

$$\tilde{c}_{do}^j(\nu, \omega) = c_o^j(\theta(\omega)) \cdot \tau_{do}^j \cdot t(\theta(\nu), \theta(\omega)), \tag{3}$$

where $c_o^j(\cdot)$ is the supplier's marginal cost, τ_{do}^j is the iceberg trade cost, and $t(\theta(\nu), \theta(\omega))$ is an increasing function of the distance between $\theta(\nu)$ and $\theta(\omega)$, capturing technology compatibility: intermediate goods perform better when the buyer and the suppliers use similar technologies.

Cost minimization over the sector bundle (2) yields the dual price index faced by firm ν,

$$P_d^{ij}(\theta) = \left(\sum_o \int \left[c_o^j(\tilde{\theta})\, \tau_{do}^j\, t(\theta, \tilde{\theta})\right]^{1-\xi} \mathrm{d}\Theta_o^j(\tilde{\theta}) \right)^{\frac{1}{1-\xi}}, \tag{4}$$

where the variety integral over the unit mass of (o, j)-firms is replaced by the probability measure Θ_o^j that describes the distribution of chosen technologies across firms in (o, j). With the Cobb–Douglas technology above, the marginal cost of a (d, i)-firm with technology θ

then solves the following fixed point problem:

$$c_d^i(\theta) = \frac{1}{A_d^i}[w_d]^{\gamma^{iL}} \prod_j \Big[\sum_o \int \left[c_o^j(\tilde{\theta})\, \tau_{do}^j\, t(\theta,\tilde{\theta})\right]^{1-\xi} \mathrm{d}\Theta_o^j(\tilde{\theta}) \Big]^{\frac{\gamma^{ij}}{1-\xi}}, \tag{5}$$

where w_d is the wage rate in d.

Equation (5) reveals the forces shaping production costs in (d,i). Standard elements in trade models with input-output linkages appear: (d,i)'s sectoral productivity A_d^i, country d's wage w_d, iceberg transport costs $\{\tau_{do}^j\}_{o=1,j=1}^{N,S}$, and the supplier costs $\{c_o^j(\tilde{\theta}) : \tilde{\theta} \in \mathbb{T}\}_{o=1,j=1}^{N,S}$. The novel compatibility effect enters via $t(\theta,\tilde{\theta})$. Because a firm with technology θ buys from suppliers in any country o with any technology $\tilde{\theta}$, $c_d^i(\theta)$ depends on the integral over the entire domain of Θ_o^j; and because technology compatibility affects input costs, firms with different θ face different production costs.

Because of the compatibility incentive, a firm's sourcing depends on its own technology θ. By Shephard's lemma, the firm allocates its input-j expenditure across suppliers according to their compatibility-adjusted costs: the share spent on firms in o with technology $\tilde{\theta}$ is

$$\chi_{do}^j(\theta,\tilde{\theta})\mathrm{d}\Theta_o^j(\tilde{\theta}) = \frac{[\tau_{do}^j c_o^j(\tilde{\theta}) t(\theta,\tilde{\theta})]^{1-\xi} \cdot \mathrm{d}\Theta_o^j(\tilde{\theta})}{\sum_{o'}[\tau_{do'}^j \Lambda_{o'}^j(\theta)]^{1-\xi}}, \tag{6}$$

where

$$\Lambda_o^j(\theta) \equiv \Big(\int [c_o^j(\tilde{\theta}) t(\theta,\tilde{\theta})]^{1-\xi} \mathrm{d}\Theta_o^j(\tilde{\theta}) \Big)^{\frac{1}{1-\xi}}; \tag{7}$$

and the share of the firm's input j expenditure spent on firms in country o (across all $\tilde{\theta}$) is

$$\chi_{do}^j(\theta) = \int \chi_{do}^j(\theta,\tilde{\theta})\mathrm{d}\Theta_o^j(\tilde{\theta}) = \frac{[\tau_{do}^j \Lambda_o^j(\theta)]^{1-\xi}}{\sum_{o'}[\tau_{do'}^j \Lambda_{o'}^j(\theta)]^{1-\xi}}. \tag{8}$$

Equation (6) is a gravity equation at the firm level, with trade elasticity $\xi - 1$: suppliers in (o,j) with technology $\tilde{\theta}$ capture a larger share of the firm's input-j spending if they are more numerous (larger $\mathrm{d}\Theta_o^j(\tilde{\theta})$) or offer a lower compatibility-adjusted cost $\tau_{do}^j c_o^j(\tilde{\theta}) t(\theta,\tilde{\theta})$. The index $\Lambda_o^j(\theta)$ in (7) summarizes the competitiveness of (o,j) as a source for a buyer with technology θ—higher when more of its firms offer low compatibility-adjusted costs—and equation (8) aggregates to the overall share of input j that a firm in (d,i) with technology θ sources from country o.

The model described above generalizes Caliendo and Parro (2015) to incorporate flexible firm-level sourcing decisions. In particular, if we impose a single technology for all firms and countries (i.e., make θ exogenous and homogeneous), then $c_d^i(\theta)$ defined in (5) collapses to a scalar $\bar{c}_d^i$, $\Lambda_o^j(\theta)$ defined in (7) also reduces to $\bar{c}_o^j$, and the trade shares $\chi_{do}^j(\theta)$ no longer depend on θ. Thus, our model preserves the flexibility and tractability of canonical quantitative trade models (e.g., Eaton and Kortum, 2002; Caliendo and Parro, 2015) while allowing for heterogeneous sourcing decisions that take into account compatibility requirements.

2.3 Production and Sourcing Decisions: Final-Good Firms

Intermediate-good firms charge final-good producers a constant monopolistic-competition markup $\mu_X \equiv \frac{\eta}{\eta-1}$ over marginal cost. Facing these prices, final-good producers choose input from each supplier to maximize their profits:

$$P_d Q_d - \sum_j \sum_o \int_0^1 [\mu_X \tau_{do}^{Uj} c_o^j(\theta(\omega))] q_{do}^j(\omega) \mathrm{d}\omega, \tag{9}$$

where τ_{do}^{Uj} is the iceberg trade cost for final users (superscript 'U'), which may differ from the τ_{do}^j faced by intermediate firms, and P_d is the price of the final good in d,

$$P_d \equiv \prod_j (P_d^j/\rho_d^j)^{\rho_d^j}, \quad \text{with } P_d^j \equiv \Big(\sum_o \int [\mu_X \tau_{do}^{Uj} c_o^j(\tilde{\theta})]^{1-\eta} \mathrm{d}\Theta_o^j(\tilde{\theta})\Big)^{\frac{1}{1-\eta}}.$$

Unlike intermediate-good firms, final-good producers are unaffected by technology compatibility. The assumption reflects where compatibility frictions originate: using an intermediate input requires integrating it into one's own production process, whereas a final product arrives with that integration already done. A household can drive an electric vehicle without any command of electrochemistry; a carmaker cannot integrate battery cells without it. Accordingly, our quantification disciplines compatibility with data on firms' input sourcing, not on final-goods purchases.

Profit maximization given (1) then yields the expenditure share allocated to sector-j goods produced by firms in country o with technology $\tilde{\theta}$,

$$\pi_{do}^j(\tilde{\theta}) \mathrm{d}\Theta_o^j(\tilde{\theta}) = \frac{[\tau_{do}^{Uj} c_o^j(\tilde{\theta})]^{1-\eta}}{\sum_{o'} [\tau_{do'}^{Uj} \bar{\Lambda}_{o'}^j]^{1-\eta}} \mathrm{d}\Theta_o^j(\tilde{\theta}), \tag{10}$$

$$\text{where} \quad \bar{\Lambda}_o^j \equiv \Big(\int [c_o^j(\tilde{\theta})]^{1-\eta} \mathrm{d}\Theta_o^j(\tilde{\theta})\Big)^{1/(1-\eta)}, \tag{11}$$

which summarizes the competitiveness of intermediate-good firms in (o, j) when serving final-good producers: reflecting the absence of compatibility concerns, no $t(\theta, \tilde{\theta})$ enters $\bar{\Lambda}_o^j$, although suppliers' technologies still matter through their costs $c_o^j(\tilde{\theta})$. The common markup μ_X cancels from the shares. Aggregating across suppliers, the expenditure share allocated to goods produced by firms in country o is

$$\pi_{do}^j = \int \pi_{do}^j(\tilde{\theta}) \mathrm{d}\Theta_o^j(\tilde{\theta}) = \frac{[\tau_{do}^{Uj} \bar{\Lambda}_o^j]^{1-\eta}}{\sum_{o'} [\tau_{do'}^{Uj} \bar{\Lambda}_{o'}^j]^{1-\eta}}. \tag{12}$$

2.4 Technology Choice

Upon entry, each firm in country-sector (d, i) draws an endowment technology $\bar{\theta}$ from an exogenous distribution characterized by probability measure $\bar{\Theta}_d^i$ on $\mathbb{T}$, which we label as

the *ex-ante* technology distribution. The firm then chooses a production technology θ and adapts its production process accordingly before making sourcing decisions. Adapting to a distant technology from the endowed one is costly but may increase profits through better compatibility with efficient suppliers.

Recall that firms earn profits solely through sales to final-good producers. Let $X_d^i(\theta)$ be the sales to final-good producers by a firm in (d,i) choosing technology θ. We specify the firm's technology adaptation cost as:

$$K_d^i(\theta;\bar{\theta}) \equiv \phi(\theta,\bar{\theta}) \cdot \frac{1}{\eta} X_d^i(\theta), \tag{13}$$

where $\frac{1}{\eta} X_d^i(\theta)$ represents the firm's expected profits under a monopolistic markup, and $\phi(\theta,\bar{\theta})$ denotes the fraction of profits spent on adapting to technology $\theta \neq \bar{\theta}$. The function $\phi(\theta,\bar{\theta})$ increases with the distance between θ and $\bar{\theta}$, capturing the idea that more distant technologies are costlier to adopt.[11]

Firms choose θ to maximize the expected net profits:

$$\Pi_d^i(\theta;\bar{\theta}) \equiv \max_{\theta} [1-\phi(\theta,\bar{\theta})] \cdot \frac{1}{\eta} X_d^i(\theta). \tag{14}$$

Under monopolistic competition in sales to final-good producers, whose demand is given by (10), this is equivalent to:

$$\max_{\theta} [1-\phi(\theta,\bar{\theta})] \times [c_d^i(\theta)]^{1-\eta}. \tag{15}$$

This formulation highlights the trade-off faced by a firm: choosing between a good technology (i.e., the one with low $c_d^i(\theta)$) and a close one. Although by themselves technologies are inherently horizontally differentiated, the ones shared by the most efficient suppliers end up vertically better. In the presence of productivity differences between countries, this means technologies adopted by leading countries tend to be vertically 'better.'

Assuming that the solution to (15) can be described by a policy function, $g_d^i(\bar{\theta})$ (for which Proposition 1 provides the existence condition), we can construct the following mapping between *ex-ante* technology distribution $\bar{\Theta}_d^i$ and the *ex-post* distribution:

$$\Theta_d^i(\mathcal{B}) = \int_{\bar{\theta}\in\mathbb{T}} \mathbb{I}[g_d^i(\bar{\theta}) \in \mathcal{B}] \mathrm{d}\bar{\Theta}_d^i(\bar{\theta}), \quad \text{for any measurable set } \mathcal{B}. \tag{16}$$

We now define the equilibrium for technology choice, taking country wages as given.

Definition 1. ***Equilibrium for Technology Choice.*** *Given the model primitives and countries' wages* $\{w_d\}_{d=1}^N$*, an equilibrium for technology choice consists of the policy functions that describe firms' technology choice,* $g_d^i : \mathbb{T} \to \mathbb{T}$*, and the cost functions* $c_d^i : \mathbb{T} \to \mathbb{R}^+$ *for all* (d,i)*, such that*

[11] We model the adaptation cost as a share of profit. This could be micro-founded in a model of bargaining, in which researchers bargain with the firm founder to split the profit from adaptation. If the founder knows less about more distant technologies, researchers can extract a larger share of the rents, or can slack and be less effective, so that more resources are wasted in the adaptation.

(1) Given $\{c_d^i\}$, $g_d^i(\bar{\theta})$ *solves problem* (15) *for all* (d,i) *and* $\forall \bar{\theta} \in \mathbb{T}$ *except on a zero-measure set.*

(2) Given $\{g_d^i\}$, $\{c_d^i\}$ *satisfies equation* (5)*, where* $\{\Theta_d^i\}$ *is given by* (16)*.*

The first condition ensures that firms' chosen technologies are optimal given the equilibrium production costs. The second condition guarantees that those choices, when aggregated, are consistent with the cost functions c_d^i.

Unlike in standard quantitative trade models, where similar conditions can be aggregated over firms and expressed solely in terms of aggregate prices and quantities,[12] in our setting, due to the bilateral nature of compatibility requirement, each firm's choice depends on the entire distribution of other firms' choices. We develop a procedure to characterize the conditions for the existence and uniqueness of equilibrium in such environments, which can be adapted to other settings where interactions between firms' decisions cannot be summarized by aggregate sufficient statistics. Our characterization proceeds in three steps:

First, we formulate equilibrium conditions involving policy functions $\{g_d^i\}$—as in Definition 1—rather than involving distributions $\{\Theta_d^i\}$. Doing so sidesteps the technical difficulties of applying fixed-point theorems to spaces of probability measures. It also allows us to focus on equilibria where technology choices vary smoothly with endowment technologies.[13]

Second, we stack the equilibrium objects to construct a joint mapping from their space (i.e., the space of $(\{g_d^i\}, \{c_d^i\})$) to itself. This differs from alternative methods viewing the equilibrium as a nested fixed point system—e.g., mapping from the space of $\{g_d^i\}$ to itself, nesting inside a fixed-point mapping for $\{c_d^i\}$ (see, e.g., Alvarez and Lucas, 2007 for an example of the nested fixed point approach in trade models). Our approach avoids the often challenging characterization of comparative statistics in the inner mapping necessary for establishing properties of the outer mapping. It shares the same spirit as Allen et al. (2024), but here, the equilibrium objects are infinite-dimensional. We illustrate how to choose appropriate norms and define spaces suitable for applying the Schauder fixed-point theorem to establish equilibrium existence.

Third, to establish uniqueness, we derive the Fréchet derivatives of the constructed joint mapping defined on functional space. We then determine the conditions under which the Jacobian matrix of these Fréchet derivatives has a row sum norm bounded by a constant below 1, which allows us to apply the contraction mapping theorem. The procedure for deriving and bounding these derivatives is broadly applicable to other models featuring similar interdependencies among firm choices.

[12] For instance, in Caliendo and Parro (2015), the first condition reduces to a system of equations for production costs; in Chaney (2008), firms' export decisions depend on the entire distribution of firm choices, which can be analytically aggregated into price indexes.

[13] Our characterization focuses on equilibria with policy functions that are bounded, continuously differentiable, and have Lipschitz-continuous first derivatives.

Our procedure works for any generic metric space $\mathbb{T}$. To simplify the algebra, however, we impose the following assumption:

Assumption 1. *1. The space of technology is the real line, i.e., $\mathbb{T} \equiv \mathbb{R}$.*

2. The costs of technological incompatibility and adaptation are given by, respectively,

$$t(\theta, \tilde{\theta}) = \exp(\frac{1}{2}\bar{t}(\theta - \tilde{\theta})^2) \text{ and } \phi(\theta, \bar{\theta}) = 1 - \exp(-\frac{1}{2}\bar{\phi}(\bar{\theta} - \theta)^2), \quad \text{with } \bar{t},\ \bar{\phi} > 0.$$

The first part of Assumption 1 takes a one-dimensional representation of the technology space. The second part states that both the compatibility and adaptation costs increase in the technological distance $(\theta - \tilde{\theta})^2$ with constant elasticities, $\bar{t}/2$ and $\bar{\phi}/2$.

There are two aspects to the assumption $\mathbb{T} = \mathbb{R}$: dimensionality and geometry. Regarding dimensionality, our theoretical results generalize to $\mathbb{T} = \mathbb{R}^n$ for a finite n.[14] In an $n > 1$ setting, any technology can be viewed as a combination of different characteristics, with each dimension representing one characteristic. We focus on $\mathbb{T} = \mathbb{R}$ because it suffices in conveying the intuition while substantially reducing computational burden in quantification. We discuss additional implications of this assumption on the calibration outcomes and quantitative results in Section 5.

Regarding geometry, the technology space does not have to be a line—one could use a circle, a ball, or other higher-dimensional manifolds. Different shape assumptions matter for firm choices.[15] Given the lack of evidence supporting one versus another, we choose a linear representation because it allows Gaussian distributions, which, as shown later, yield tractable solutions in our quantitative exercise.

Under Assumption 1, we establish the following results on existence and uniqueness of the technology choice equilibrium.

Proposition 1. *Suppose wages $\{w_d\}$ are given.*

1. Let $\{\bar{\Theta}_d^i\}$ have bounded support within $[-M, M]$ with density functions $\{\bar{\varsigma}_d^i\}$. If $(\xi - 1)\bar{t} < 1/M^2$, then there exists an equilibrium with firms' technology choice $\{g_d^i\}$ being continuously differentiable functions. In this equilibrium, the choice of firms from (d, i) with endowment technology $\bar{\theta}$ is characterized by the following first-order condition, which has a unique solution:

$$g_d^i(\bar{\theta}) = (1 - \omega^i)\bar{\theta} + \omega^i \sum_{j,o} \frac{\gamma^{ij}}{1 - \gamma^{iL}} \int \chi_{do}^j(g_d^i(\bar{\theta}), \tilde{\theta})\, \tilde{\theta}\, \mathrm{d}\Theta_o^j(\tilde{\theta}), \quad \forall \bar{\theta} \in [-M, M], \tag{17}$$

where $\omega^i \equiv \frac{(\eta-1)(1-\gamma^{iL})\bar{t}}{(\eta-1)(1-\gamma^{iL})\bar{t}+\bar{\phi}} < 1$.

[14] For $\mathbb{T} = \mathbb{R}^n$, the cost functions in part (2) of Assumption 1 take a quadratic form in the exponent, e.g., $t(\theta, \tilde{\theta}) = \exp(\frac{1}{2}(\theta - \tilde{\theta})' \cdot \bar{t} \cdot (\theta - \tilde{\theta}))$, where x' is the transpose of a vector x and $\bar{t}$ here is a positive definite *matrix*.

[15] For example, on a line, when a firm moves left, it moves closer to firms on its left while moving away from those on its right. In contrast, on a circle, firms moving in opposite directions eventually meet.

2. *In addition, if* $\bar{t} < \frac{1}{2M}$ *and* $\bar{\phi} > \underline{\phi}$ *for a constant* $\underline{\phi} > 0$ *determined by parameters* $(\xi - 1, \bar{t}, \eta, M, \gamma^{iL})$ *as detailed in the proof, then the equilibrium satisfying the above conditions is unique.*

Proof. See Appendix A.3. □

For an equilibrium with continuous policy functions to exist, firms' technology choices should not be overly sensitive to their own endowment technologies, which is ensured if $(\xi - 1)\bar{t}M^2 < 1$. Here, M^2 sets a limit on the variance of potential supplier technologies, which determines how much firms can improve compatibility through technology choice.[16] $\bar{t}$ governs how fast production costs rise with technological distance, and $\xi - 1$ is the trade elasticity shaping how price advantages translate to profits. Together, these three parameters bound the sensitivity of firms' technology choice to their endowments.

The first order condition shows that firms balance staying close to their endowment $\bar{\theta}$ (weighted by $1 - \omega^i$) against aligning with suppliers' average technology (weighted by ω^i and the importance of suppliers $\frac{\gamma^{ij}}{1-\gamma^{iL}}\chi^j_{do}(g^i_d(\bar{\theta}), \tilde{\theta})$).

The equilibrium for technology choice is not generically unique. In particular, in the limit as $\bar{\phi} \to 0$ or $\bar{t} \to \infty$, the incentive for compatibility outweighs the cost of adaptation, leading all firms to coordinate on the same location—and any such location can be an equilibrium. To establish uniqueness, we derive the conditions under which the joint mapping described earlier is a contraction. In the condition $\bar{t} < \frac{1}{2M}$, $2M$ enters as the upper bound of the distance between a firm's technology and the expenditure-weighted average technology of its suppliers. Intuitively, $\bar{t} \cdot 2M$ determines the scope for compatibility benefits. If these benefits are not too large, and if the cost of technology adaptation $\bar{\phi}$ is not too small, then firms' technology would not be overly sensitive to *other firms*' choices, thereby ensuring uniqueness.[17]

Beyond technology compatibility, one can view the model more generally: θ enters as a firm attribute, and $t(\theta, \tilde{\theta})$ as a bilateral interaction between buyer and supplier attributes. The trade block therefore accommodates general firm heterogeneity in sourcing, and the equilibrium existence and uniqueness characterization, which is cast in terms of policy and cost functions rather than relying on the CES algebra or the quadratic cost forms specific to our model, applies more generally to environments in which firms best respond to the entire distribution of other firms' choices.

[16] By Popoviciu's inequality on variances, $\text{var}(\theta) < M^2$ for any distribution of θ over $[-M, M]$. Larger variance means greater heterogeneity in supplier technologies, implying larger room for firms to improve supplier compatibility, hence leading to more sensitive technology choice with respect to endowed technology.

[17] This intuition behind the bound for $\bar{t}$ means that even if M is arbitrarily large, as long as firms' expenditure shares on distant suppliers are not too large, uniqueness can be ensured without $\bar{t} \to 0$. More generally, Proposition 1 requires bounded technological space because we allow for arbitrary endowment distributions. This assumption can be relaxed by restricting the endowment distribution to particular shapes. Indeed, under our calibration, the equilibrium under Normal distributions on ex-ante distributions exists and is unique.

2.5 General Equilibrium

We now embed firms' technology choice equilibrium into a general equilibrium. For a firm in (o, j) using technology $\tilde{\theta}$, let $M_o^j(\tilde{\theta})$ denote its expected sales *to downstream firms*. Market clearing in intermediate goods, stated in value of goods, reads:[18]

$$M_o^j(\tilde{\theta}) = \sum_d \sum_i \int \gamma^{ij} \left[M_d^i(\theta) + (1 - \frac{1}{\eta}) X_d^i(\theta) \right] \chi_{do}^j(\theta, \tilde{\theta}) \mathrm{d}\Theta_d^i(\theta), \tag{18}$$

where $M_d^i(\theta)$ is the expected sales to other firms by a firm in (d, i) with technology θ, $X_d^i(\theta)$ is that firm's expected sales to final producers. The sum $\left[M_d^i(\theta) + (1 - \frac{1}{\eta}) X_d^i(\theta) \right]$ equals the expected production cost of that firm (given its $\frac{\eta}{\eta - 1}$ markup to final producers), γ^{ij} is that firm's expenditure share on input from sector j, and $\chi_{do}^j(\theta, \tilde{\theta})$ is its likelihood of purchasing from a firm in (o, j) with technology $\tilde{\theta}$.

The expected sales of a firm with technology $\tilde{\theta}$ to *final-good producers*, $X_o^j(\tilde{\theta})$, satisfy

$$X_o^j(\tilde{\theta}) \equiv \sum_d \rho_d^j P_d Q_d \pi_{do}^j(\tilde{\theta}). \tag{19}$$

Here, $P_d Q_d$ is the final good demand from household consumption and adaptation costs:

$$P_d Q_d = I_d + \sum_i \int K_d^i(g_d^i(\bar{\theta}); \bar{\theta}) \mathrm{d}\bar{\Theta}_d^i(\bar{\theta}), \tag{20}$$

with technology adaptation costs $K_d^i(\theta; \bar{\theta})$ defined in (13), and g_d^i the policy function for technology choice as part of the technology choice equilibrium (Definition 1).

Household income consists of wages and the net profits of domestic firms:

$$I_d = w_d L_d + \sum_i \int \Pi_d^i(g_d^i(\bar{\theta}); \bar{\theta}) \mathrm{d}\bar{\Theta}_d^i(\bar{\theta}), \tag{21}$$

where net profits $\Pi_d^i(\theta; \bar{\theta})$ are defined in (14), and wages are determined by the labor market clearing condition, stated in expenditure:

$$w_d L_d = \sum_i \gamma^{iL} \int [M_d^i(\theta) + (1 - \frac{1}{\eta}) X_d^i(\theta)] \mathrm{d}\Theta_d^i(\theta). \tag{22}$$

Definition 2. ***Competitive Equilibrium.*** *Given parameters on geography $\{\tau_{do}^j, \tau_{do}^{Uj}, L_d\}$, preference $\{\rho_d^j, \eta\}$, production technology $\{\gamma^{ij}, \gamma^{iL}, A_d^i, \xi\}$, compatibility and adaptation costs $\{\bar{t}, \bar{\phi}\}$, and the ex-ante technology distribution $\{\bar{\Theta}_o^j\}$, a competitive equilibrium is defined as (i) wages, price index and income $\{w_d, P_d, I_d\}$, (ii) ex-post technology distribution $\{\Theta_o^j\}$, (iii) sales characterized by $\{X_o^j(\theta), M_o^j(\theta)\}$, and (iv) policy function for technology choice $\{g_d^i(\bar{\theta})\}$ and production cost function $\{c_d^i(\theta)\}$, such that*

[18] Or, in the integral form: for any set $\mathcal{B}$,

$$\int_{\tilde{\theta} \in \mathcal{B}} M_o^j(\tilde{\theta}) \mathrm{d}\Theta_o^j(\tilde{\theta}) = \int_{\tilde{\theta} \in \mathcal{B}} \Big[\sum_d \sum_i \int \gamma^{ij} \left[M_d^i(\theta) + (1 - \frac{1}{\eta}) X_d^i(\theta) \right] \chi_{do}^j(\theta, \tilde{\theta}) \mathrm{d}\Theta_d^i(\theta) \Big] \mathrm{d}\Theta_o^j(\tilde{\theta}).$$

i. Given wages $\{w_d\}$, $\{g_d^i(\bar{\theta}), c_d^i(\theta)\}$ *constitute a technology choice equilibrium defined in Definition 1, and* $\{\Theta_o^j\}$ *is consistent with equation (16).*

ii. Goods and labor markets clear, i.e., (20), (21), (22) hold by d, and (18) and (19) hold by $(o, j, \tilde{\theta})$.

3 Special Cases

Having defined the general equilibrium, we use special cases to analytically examine the interaction between technology choice and trade, and the resulting welfare implications. Throughout this section, we assume that each country–sector pair (d, i) has a degenerate ex-ante technology distribution with a unit mass at $\bar{\theta}_d^i$. We focus on the equilibrium for technology choice (i.e., taking wages as given) that is symmetric, meaning that all firms from a country-sector pair (d, i) with the same endowment technology choose the same technology. Under these assumptions, the technology chosen by firms in (d, i) is a weighted average between their endowment technology $\bar{\theta}_d^i$ and the technologies chosen by their suppliers $\{\theta_o^j\}$: the technology-choice first-order condition (17) specializes to

$$\theta_d^i = (1 - \omega^i)\bar{\theta}_d^i + \omega^i \sum_{j,o} \frac{\gamma^{ij}}{1 - \gamma^{iL}} \bar{\chi}_{do}^{ij} \theta_o^j, \tag{23}$$

where $\omega^i \equiv \frac{(\eta-1)(1-\gamma^{iL})\bar{t}}{(\eta-1)(1-\gamma^{iL})\bar{t}+\bar{\phi}}$, and $\bar{\chi}_{do}^{ij} \equiv \frac{\left[\tau_{do}^j \bar{c}_o^j \exp(\frac{1}{2}\bar{t}(\theta_d^i - \theta_o^j)^2)\right]^{1-\xi}}{\sum_{o'}\left[\tau_{do'}^j \bar{c}_{o'}^j \exp(\frac{1}{2}\bar{t}(\theta_d^i - \theta_{o'}^j)^2)\right]^{1-\xi}}$, with $\bar{c}_o^j \equiv c_o^j(\theta_o^j)$, is the share of sector-j input spent on country o by intermediate-good firms in (d, i). Because firms within a (d, i) are homogeneous, they have the same trade shares $\bar{\chi}_{do}^{ij}$.

3.1 Interaction between Technology Choice and Trade

The first-order condition (23) links technology choice and trade in both directions: trade costs move technology choices through the expenditure shares $\bar{\chi}_{do}^{ij}$, and technology gaps move the shares through the compatibility cost. We examine each direction in turn. To illustrate the first, we consider a 'small' country-sector (d, i), whose technology choice does not affect the rest of countries and sectors.

Proposition 2. *Consider a country-sector* (d, i) *that is small in the sense that its input and output account for a negligible share of all countries and sectors, including sectors in country d. Following an x % increase in the cost of* (d, i) *importing from* (o, j)*:*

1. *The squared distance between* θ_d^i *and* θ_o^j, $\|\theta_d^i - \theta_o^j\| \equiv (\theta_d^i - \theta_o^j)^2$, *changes by:*

$$\Delta\|\theta_d^i - \theta_o^j\| = -\underbrace{\frac{2(\xi-1)\tilde{\omega}^i \gamma^{ij} \bar{\chi}_{do}^{ij} \|\theta_o^j - \vartheta_d^{ij}\|}{1 - \bar{t}(\xi-1)\tilde{\omega}^i \sum_{j',o'} \gamma^{ij'} \bar{\chi}_{do'}^{ij'} \|\theta_{o'}^{j'} - \vartheta_d^{ij'}\|}}_{>0 \text{ by firms' second-order condition}} \times \frac{\theta_d^i - \theta_o^j}{\theta_o^j - \vartheta_d^{ij}} \times x,$$

where $\tilde{\omega}^i \equiv \omega^i/(1-\gamma^{iL})$*, with* ω^i *and* $\bar{\chi}^{ij}_{do}$ *defined in equation (23), and* $\vartheta^{ij}_d \equiv \sum_m \bar{\chi}^{ij}_{dm}\theta^j_m$ *is the weighted average technology of sector-j suppliers to* (d,i)*. We assume* $\gamma^{ij} > 0$ *and* $\theta^j_o \neq \vartheta^{ij}_d$*; if either fails,* θ^i_d *does not respond to the shock and the changes in both parts are zero.*

2. $\|\theta^i_d - \theta^j_o\|$ *increases relative to the expenditure-share weighted distance between* θ^i_d *and* $\theta^j_{o'}$ *across* $o' = 1, ..., N$*. Formally,*

$$\Delta\|\theta^i_d - \theta^j_o\| - \sum_{o'} \bar{\chi}^{ij}_{do'}\Delta\|\theta^i_d - \theta^j_{o'}\| = \frac{2(\xi-1)\tilde{\omega}^i\gamma^{ij}\bar{\chi}^{ij}_{do}\|\theta^j_o - \vartheta^{ij}_d\|}{1-\bar{t}(\xi-1)\tilde{\omega}^i\sum_{j',o'}\gamma^{ij'}\bar{\chi}^{ij'}_{do'}\|\theta^{j'}_{o'} - \vartheta^{ij'}_d\|} \times x > 0. \tag{24}$$

Proof. See Appendix A.4. □

Part 1 of the proposition characterizes the impact of an increase in τ^j_{do} on $\|\theta^i_d - \theta^j_o\|$. As τ^j_{do} increases, firms in (d,i) increase their dependence on other suppliers, which increases the weight of these suppliers' average technology in θ^i_d. Whether this change increases or decreases $\|\theta^i_d - \theta^j_o\|$ depends on the sign of the ratio $(\theta^i_d - \theta^j_o)/(\theta^j_o - \vartheta^{ij}_d)$, which in turn depends on the positions of θ^i_d and ϑ^{ij}_d relative to θ^j_o. If θ^i_d and ϑ^{ij}_d lie on the same side of θ^j_o, then (d,i) firms' putting more weights on other suppliers reduces their incentive to shift to θ^j_o, increasing $\|\theta^i_d - \theta^j_o\|$; conversely, $\|\theta^i_d - \theta^j_o\|$ may shrink.

Despite this ambiguity, Part 2 establishes that $\|\theta^i_d - \theta^j_o\|$ *always* increases *relative to* the expenditure-share-weighted average distance between θ^i_d and the technologies of all suppliers. The increase in the relative distance is larger if adaptation is less costly (smaller $\bar{\phi}$, which enters via $\tilde{\omega}^i$), if (d,i) relies more heavily on (o,j) for inputs (large $\gamma^{ij}\bar{\chi}^{ij}_{do}$), or if θ^j_o is farther from the average technology of suppliers (larger $\|\theta^j_o - \vartheta^{ij}_d\|$). We will leverage this intuition to discipline the model for quantification.

We now turn to the second direction: firms' endowment technologies shape their chosen technologies, which in turn shape trade. To illustrate this mechanism, we consider a change in the endowed technology of a zero-measure set of firms from the rest of the firms in (d,i).

Proposition 3. *Suppose all firms in* (d,i) *share the same endowment technology* $\bar{\theta}^i_d$*, except for a zero measure set* ν *endowed with* $\bar{\theta}(\nu)$*. If* $\bar{\theta}(\nu)$ *moves closer to* θ^j_o *relative to* $\theta^j_{o'}$*, then*

1. *Firm* ν *chooses a technology closer to* θ^j_o *relative to* $\theta^j_{o'}$*, namely,* $\|\theta^i_d(\nu) - \theta^j_o\| - \|\theta^i_d(\nu) - \theta^j_{o'}\|$ *decreases.*

2. *Firm* ν *is more likely to source intermediate goods from* (o,j) *relative to* (o',j)*; the change in the likelihood of sourcing from* (o,j) *relative to from* (o',j) *satisfies:* $\Delta\log\left(\bar{\chi}^{ij}_{do}(\nu)/\bar{\chi}^{ij}_{do'}(\nu)\right) = -\frac{1}{2}(\xi-1)\bar{t} \times \Delta\left(\|\theta^i_d(\nu) - \theta^j_o\| - \|\theta^i_d(\nu) - \theta^j_{o'}\|\right)$.

Proof. See Appendix A.5. □

The proposition indicates that as a firm's endowed technology moves toward θ_o^j, so does its chosen technology. This, in turn, makes the firm more likely to source from (o, j). The response of sourcing patterns to the technology distance between trade partners is governed by $\bar{t}$ as well as the conventional trade elasticity $\xi - 1$. Conditioning on $\xi - 1$, this response identifies $\bar{t}$, an intuition we will use in quantification to pin down $\bar{t}$.

3.2 The Welfare Implications of Technology Choice

In our model, firms choose technology to maximize their own profits, but their choices impact the profits of downstream firms, too. Such externalities propagate through input-output and trade linkages, opening a gap between the private and the social trade-off in technology choice. We characterize this gap first in a closed economy, where the externality is domestic, and then in a two-country economy, where it crosses the border; we then show how it shapes the welfare gains from trade liberalization.

Proposition 4 (Private and social trade-offs in technology choice)**.** *Consider a closed economy (dropping country subscripts) with S sectors and roundabout production, each sector having a degenerate endowment $\bar{\theta}^i$. The marginal impact of increasing θ^i on household real income (denoted $U \equiv I/P$) is given by*

$$\frac{\partial \ln U}{\partial \theta^i} = \underbrace{\frac{\rho^i \exp(-\frac{1}{2}\bar{\phi}(\theta^i - \bar{\theta}^i)^2)}{\eta - 1 + \sum_k \rho^k \exp(-\frac{1}{2}\bar{\phi}(\theta^k - \bar{\theta}^k)^2)} \bar{\phi}(\bar{\theta}^i - \theta^i)}_{\text{income effect}} - \underbrace{\bar{t}\, \lambda^i \sum_j \gamma^{ij}(\theta^i - \theta^j)}_{\text{own-sector price}}$$
$$\underbrace{- \bar{t} \sum_{m \neq i} \lambda^m \gamma^{mi}(\theta^i - \theta^m)}_{\text{downstream prices}}, \tag{25}$$

where $\lambda^i \equiv \sum_n \rho^n \Omega^{ni}$ is the Domar weight of sector i, with Ω^{ni} the (n, i) element of $(\mathbb{I}_{S\times S} - \Gamma)^{-1}$ and $\Gamma^{ij} \equiv \gamma^{ij}$.

Suppose the input–output structure is symmetric: all sectors share a common final-good weight $\rho^i = \rho^j$ and common diagonal and off-diagonal input shares, $\gamma^{ii} = \gamma^{jj}$ and $\gamma^{ij} = \gamma^{ij'} = \gamma^{jj'}$ for all distinct sectors i, j, j'. Then, at the decentralized equilibrium, $\partial \ln U/\partial \theta^i > 0$ if $\theta^i > \bar{\theta}^i$ and $\partial \ln U/\partial \theta^i < 0$ if $\theta^i < \bar{\theta}^i$: compared with the social planner, firms under-invest in adaptation.

Proof. See Appendix A.6. □

Equation (25) illustrates the mechanisms through which θ^i affects welfare. Suppose, without loss of generality, that firms in sector i have shifted left of their endowment, $\theta^i < \bar{\theta}^i$, to use inputs from sectors whose technologies lie further left. A marginal increase in θ^i from the equilibrium choice has three effects. First, it saves adaptation costs, freeing final output for household consumption (the income effect). Second, it changes sector i's own price, with

the Domar weight λ^i capturing how sector i's compatibility with its suppliers propagates through the network into consumer prices. Third, as θ^i shifts, the sectors that use sector-i output as an input gain or lose depending on whether their distance to θ^i falls or rises; this downstream effect is stronger the larger the compatibility benefit $\bar{t}$ or the linkage $\lambda^m \gamma^{mi}$.

Contrasting this with the firm's first-order condition that shapes its technology choice, the closed-economy counterpart of (23),

$$\rho^i \Big[\underbrace{\frac{1}{\eta - 1} \bar{\phi} (\bar{\theta}^i - \theta^i)}_{\text{adaptation cost}} - \underbrace{\bar{t} \sum_j \gamma^{ij} (\theta^i - \theta^j)}_{\text{production cost}} \Big] = 0, \tag{26}$$

we can see that, compared to the social effects of adaptation, individual firms (i) weigh the income effect differently, because a firm's profit is only a share of household income; (ii) weigh the own-sector price effect through the direct input shares γ^{ij} rather than the Domar weights λ^i, as they do not take general equilibrium into account; and (iii) do not internalize how their technology benefits downstream users (the last term of (25), absent from the private condition). Because of these externalities the decentralized equilibrium generally fails to maximize welfare, but the direction of the distortion is ambiguous. As a firm moves toward one sector it moves away from others, so the net social effect depends on the relative importance of the sectors on each side: both their direct linkages γ and their total influence summarized by the Domar weights λ. Under the symmetric input–output structure, the proposition shows that this ambiguity resolves: the downstream externality reinforces the own-sector effect, and firms under-invest, staying too close to their endowment.[19]

We next consider the international spillovers of technology choice, inspecting a one-sector economy (dropping sector superscripts) with roundabout production and two symmetric countries, 1 and 2, whose endowment technologies sit at opposite positions in the technology space, $\bar{\theta}_2 = -\bar{\theta}_1$. Without loss of generality, $\bar{\theta}_1 > 0$; in the symmetric equilibrium the chosen technologies then satisfy $\theta_2 < 0 < \theta_1$ with $|\theta_1| = |\theta_2|$.

Proposition 5 (Spillovers across borders)**.** *Starting from the symmetric decentralized equilibrium, a marginal increase in θ_2 toward θ_1, holding θ_1 fixed, raises welfare in both countries:*

$$\frac{\partial \ln U_2}{\partial \theta_2} = \underbrace{\frac{\exp(-\frac{1}{2}\bar{\phi}(\theta_2 - \bar{\theta}_2)^2)}{\eta - 1 + \exp(-\frac{1}{2}\bar{\phi}(\theta_2 - \bar{\theta}_2)^2)} \bar{\phi}(\bar{\theta}_2 - \theta_2)}_{\textit{income effect } (<0)} + \underbrace{\bar{t} \frac{1 - \gamma^L}{\gamma^L} \bar{\chi}\,(\theta_1 - \theta_2)}_{\textit{price effect } (>0)} > 0,$$

$$\frac{\partial \ln U_1}{\partial \theta_2} = \bar{t} \frac{1 - \gamma^L}{\gamma^L} \bar{\chi}\,(\theta_1 - \theta_2) > 0, \qquad \bar{\chi} \equiv \frac{\left(\tau \exp(\frac{1}{2}\bar{t}(\theta_1 - \theta_2)^2)\right)^{1-\xi}}{\left(\tau \exp(\frac{1}{2}\bar{t}(\theta_1 - \theta_2)^2)\right)^{1-\xi} + 1}, \tag{27}$$

[19] A planner can correct the under-investment with an adaptation subsidy. The gap between the private condition (26) and the social condition (25) varies with each sector's position in the production network and in the technology space, so, more generally, the required subsidy or tax is sector-specific; when endowment technologies differ within a sector, it must likewise differ across firms.

where τ is the common bilateral trade cost on intermediate inputs and $\bar{\chi}$ is each country's expenditure share on foreign intermediate inputs. Hence, locally, the decentralized equilibrium features too little adaptation toward the foreign technology, from the perspective of both home and foreign welfare.

Proof. See Appendix A.7. □

As in the closed-economy setting, firms' private incentive for technology adaptation does not fully internalize the broader benefits of compatibility, which here extend to the foreign country. Starting from the decentralized equilibrium, moving θ_2 closer to θ_1 lowers the production costs of both countries due to stronger compatibility, as countries use each other's output as input. Beyond this direct effect, lower production costs in country 1 benefit producers in country 2, and vice versa. As a result, *both* countries benefit from such a move.

This mechanism shapes the welfare effects of trade policy: because the externality in adaptation is first order, the technology convergence induced by liberalization, or the divergence induced by a trade conflict, has welfare effects beyond the direct first-order cost effect captured by Hulten's theorem. The next proposition examines this effect.

Proposition 6 (Gains from trade liberalization)**.** *Starting from a locally stable symmetric decentralized equilibrium,[20] a marginal bilateral change in trade costs, $\mathrm{d}\ln\tau^U = \mathrm{d}\ln\tau$, changes welfare in each country by*

$$\frac{\mathrm{d}\ln U_d}{\mathrm{d}\ln\tau} = \underbrace{-\left[\bar{\pi} + \frac{1-\gamma^L}{\gamma^L}\bar{\chi}\right]}_{\text{direct effect } (<0)} + \underbrace{\varrho\,\frac{\mathrm{d}\ln t}{\mathrm{d}\ln\tau}}_{\text{technology response } (<0)}, \qquad d = 1, 2, \tag{28}$$

where τ^U is the common bilateral trade cost on final use, $\bar{\pi} \equiv (\tau^U)^{1-\eta}/[1+(\tau^U)^{1-\eta}]$ is the final-use import share, and $\bar{\chi}$ is the intermediate-input import share defined in (27). The elasticity $\mathrm{d}\ln t/\mathrm{d}\ln\tau$ summarizes the response of technology to trade costs, with $t \equiv \exp(\frac{1}{2}\bar{t}(\theta_1-\theta_2)^2)$ the equilibrium bilateral compatibility cost:

$$\frac{\mathrm{d}\ln t}{\mathrm{d}\ln\tau} = \frac{(\xi-1)(1-\bar{\chi})}{\dfrac{1}{\bar{t}(\theta_1-\theta_2)^2} + \dfrac{1}{2\bar{t}(\theta_1-\theta_2)(\bar{\theta}_1-\theta_1)} - (\xi-1)(1-\bar{\chi})} > 0, \tag{29}$$

The coefficient ϱ converts the technology response into welfare:

$$\varrho = -\frac{1-\gamma^L}{\gamma^L}\bar{\chi}\left[1 - \frac{\gamma^L}{2}\frac{(\eta-1)\exp(-\frac{1}{2}\bar{\phi}(\theta_1-\bar{\theta}_1)^2)}{\eta-1+\exp(-\frac{1}{2}\bar{\phi}(\theta_1-\bar{\theta}_1)^2)}\right] < 0.$$

Hence the endogenous technology response amplifies the gains from trade cost reductions.

[20] Local stability means that when technologies adjust in the direction of firms' marginal incentives, a small perturbation of trade costs does not lead to more convergence, more trade, yet more convergence, and so on without bound: stability requires this feedback loop to be finite rather than explosive. Differentiating the technology-choice condition along the symmetric branch, this is equivalent to a positive denominator in (29). A sufficient condition, which does not involve $\bar{\phi}$, is $(\xi-1)(1-\bar{\chi})\,\bar{t}\,(\theta_1-\theta_2)^2 < 1$, paralleling the sensitivity bound of Proposition 1.

Proof. See Appendix A.8. □

The direct effect captures the first-order cost savings from the trade-cost change on existing trade flows. When Hulten's theorem applies, that is, when the initial equilibrium is efficient, this statistic is the entire first-order effect: the envelope theorem implies the endogenous response in technology choice is second order. Here, with a constant markup to consumers and zero markup to firms, the economy with exogenous technology is efficient; technology choice, however, is not. The gap between private and social trade-off in technology adaptation discussed earlier means that the endogenous response in technology could have a first-order welfare effect, captured by $\varrho < 0$. In this case, trade liberalization would lead to convergence in technology which brings additional gains.

The unambiguous amplification reflects the two-country setting: with a single partner, the compatibility externality points in one direction. More generally, the endogenous technology response to a trade-cost change can generate global gains or losses of its own. Suppose a third country lies to the right of country 1 in the technology space. Liberalization between countries 1 and 2 pulls θ_1 away from θ_3, and if country 1 trades heavily with country 3, the forgone compatibility with 3 can outweigh the improved compatibility with 2: the technology response then dampens rather than amplifies the direct effect. The same logic extends across sectors: liberalization concentrated in one sector or one country pair tilts each affected industry's adaptation toward the suppliers whose inputs become cheaper and away from its other suppliers, domestic and foreign, so the induced welfare effect can again take either sign.

This result is related to the finding of Arkolakis et al. (2025), who study a model of trade in which firms invest in forming supplier and buyer relationships before trading. They show that, relative to a fixed-network benchmark, the endogenous response of relationship-forming investment scales the effect of trade-cost reductions by a coefficient that is uniform across country pairs. While our model shares a similar spirit, there is a crucial difference. Here the endogenous margin is directional technology choice, and its effect is not uniform: each country's adjustment strengthens compatibility with some partners and weakens it with others, so the sign depends on each pair's position in the global trade and technology input–output map. Which case dominates in the actual network of trade and technologies depends on the type of shock and is a quantitative question. We take up this question in Section 5.

3.3 Markups on Sales to Other Firms

Firms under-invest in technology adaptation because they overlook the benefit their technology confers on downstream firms. A markup on sales to other firms lets a firm capture part of this benefit, though such markups are themselves a source of inefficiency through double

marginalization: the economy with exogenous technology is no longer efficient. Before moving to measurement, we discuss how monopolistically competitive markups on firms' sales to other firms alter the earlier results. Throughout this subsection, we assume intermediate firms charge the constant markup $\mu_M \equiv \xi/(\xi-1)$ on sales to other firms, in addition to the markup $\mu_X \equiv \eta/(\eta-1)$ they charge final-good producers. Aside from this, the rest of the setup stays the same. We discuss results below, deferring formal analysis to Appendix A.9.

Technology choice. With the markup, sales to other firms contribute to profit, so a firm's technology matters through two channels: its own cost, which determines its appeal to buyers as a supplier, and its compatibility with the technologies of the buyers it serves. The chosen technology now becomes a weighted average of the firm's endowment technology, the average technology of its suppliers, and, new relative to the baseline, the average technology of its buyers. The buyer-side weight is proportional to the share of operating profit the firm earns on sales to other firms; when that share is zero, the condition collapses to the baseline (17).

We now revisit the results discussed in Section 3.2.

Closed economy (Proposition 4). The trade-off in technology choice for the social planner is largely the same: the marginal effect of a sector's technology on real income has identical second and third terms as in (25), and the markup enters only by enlarging the profit share in income in the first term of (25). The firm's trade-off, by contrast, has a new margin: the firm now has a stake in its downstream users, a private counterpart of the downstream price term in the social trade-off. However, the firm weighs downstream users by its own sales shares and margin, not by the social weights that measure their importance for household real income.

In this setting, as in the case without the markup, whether firms over- or under-invest in adaptation is ambiguous. Under the symmetric setup of Proposition 4, however, firms always under-invest relative to social optimum.

Two symmetric countries with one sector (Propositions 5 and 6). With the markup, the cross-border spillover becomes more subtle. When country 2 adapts toward country 1, holding country 1's technology fixed, country 1 always benefits, but country 2 itself may gain or lose: the markup strengthens firms' private incentive for adaptation, so at the decentralized equilibrium the marginal cost of further unilateral adaptation, borne in home income, can exceed the home gain. The conclusion of Proposition 5 that a *unilateral* move benefits both countries now depends on model primitives. However, when both countries adapt toward each other symmetrically, both benefit. Therefore, firms still under-invest in adaptation relative to a global planner.

Following the discussion above, when trade costs between the two countries decrease, their technologies converge toward each other. This endogenous response again amplifies the gains from trade liberalization beyond the first-order cost savings on existing trade flows

as in Proposition 6.

Taking stock, in this section we examine the efficiency of firms' technology choice. Choosing a technology closer to other firms lowers not only the firm's own input costs but also the production costs of the firms that buy from it, at home and abroad. As a result, firms capture only part of the surplus generated through technology choice, leading to externalities.[21]

In symmetric cases, such externalities point in the direction of under-investment. In a multi-country, multi-sector setting, under-investment vis-à-vis one trade partner may mean over-investment toward another. Thus, whether firms under- or over-invest from a global perspective, and whether the technology response amplifies or dampens the welfare effects of trade shocks between a pair of countries, depends on the entire map of trade and technologies. We will use the quantitative model to measure such effects.

4 Evidence

We now turn to the data to examine the model's implication that declines in tariffs make importer and exporter technologies more similar. We test this implication by exploiting time-series changes in tariffs that affect certain MFN-bound exporting countries but not others. The resulting estimate will also help discipline the model for quantitative exercises.

4.1 Measuring Technology Similarity Using Patent Texts

Our empirical analysis requires a measure of compatibility between technology pairs. In the model, technologies are compatible if they are close to each other in the technology space. We use a proxy for this by constructing similarity measures from patent texts. Patent texts offer two advantages for our purpose. First, patents are embodiments of new technologies by design, and their texts are intended to accurately summarize the technical features of inventions. Second, patent data are available globally, facilitating comprehensive measurement across countries and sectors.

Measure Construction. We use the 2023 Fall release of *PATSTAT Global*, a database maintained by the European Patent Office (EPO) that covers close to the universe of world patents. To ensure consistency in language processing, we focus on patents with English abstracts. To avoid double counting inventions patented in multiple offices (e.g. European, U.S. etc.), we keep only one patent per patent family. This yields around 45 million patent abstracts.[22]

[21] A markup changes how the realized surplus is split, but it does not ensure efficiency. Unless firms can fully extract the rents from improved compatibility, or write ex-ante contracts governing both technology choice and sales, they will not fully internalize the externality.

[22] Because the EPO prioritizes abstracts in English when sourcing the data, approximately 90% of all abstracts in PATSTAT are in English. Our final sample keeps a patent family as long as one of the patents in the family is in English, so the coverage is even higher.

We process these English abstracts using a leading large language model developed by Alibaba ('gte-base-en-v1.5'), which transforms each paragraph into a 768-dimensional numerical vector—a process called 'text embedding.'[23] The embedding of a paragraph summarizes its meaning as a point in a 768-dimensional space. In the context of patents, we can view this space as representing all possible combination of technical ingredients and insights from different scientific disciplines to achieve a function. Closer points in the space represent patents that are more similar to each other.

To measure similarity between firms and countries, we take the average embedding of all patents produced by a given unit (firm or country) and compute the cosine similarity between units.[24] For example, let $\bar{c}_{\omega t}$ be the average embedding of all patents invented in firm ω in period t, and let $\bar{c}_{ot}$ be the average embedding of patents invented in country o. The similarity between firm ω and country o in period t is defined as:

$$\text{Similarity}_{\omega ot} = \frac{\bar{c}_{\omega t} \cdot \bar{c}_{ot}}{\|\bar{c}_{ot}\| \cdot \|\bar{c}_{\omega t}\|},$$

where the numerator is the inner product of the two vectors, and the denominator is the product of their norms. We construct analogous measures at other levels of aggregation, such as between country–sector pairs, and between firm and country-sector pairs. We assign patents to sectors and products (HS-6) using the crosswalks between the international patent classification and standard industry and product classifications by Lybbert and Zolas (2014).

What Does the Measure Capture? Patents that are close in the embedding space describe inventions drawing on similar scientific and engineering principles. Such technologies are more likely to work together, as components, tools, or inputs within the same product or process, which is the notion of compatibility in our model.

One may worry that textual similarity captures relations other than compatibility. Chief among them is competition: patents describing rival solutions to the same problem (say, two battery chemistries) also read alike, yet such technologies are substitutes rather than complements in production. To separate the two, we finetune the same embedding model on a large sample of patent pairs whose competitive relation is assessed by a large language model, so that the dot product of the resulting embeddings scores whether two technologies compete rather than whether their descriptions are similar. We aggregate this competition score to the same levels as our similarity measures and control for it in all analyses. Appendix

[23] See https://huggingface.co/Alibaba-NLP/gte-base-en-v1.5 for model descriptions. The model balances efficiency and quality. Deployed on a desktop GeForce RTX 4090 GPU, it produces the embeddings for 45 million patents in 2 days. In terms of quality, as of August 2024, it ranks 36 of 438 on the Massive Text Embedding Benchmark Leaderboard (Muennighoff et al., 2023). For comparison, the default model of SBERT ('all-mpnet-base-v2') ranks 282, and OpenAI's latest embedding model ('text-embedding-3-large-256') ranks 65.

[24] Since the inner product is bilinear, this is equivalent to first constructing pairwise similarities between individual patents and then averaging across all cross-unit pairs: the average pairwise inner product equals the inner product of the average embeddings. The two calculations differ only in whether the normalization by vector norms is applied before or after averaging.

B.3 details the construction and presents an example of how the finetuned model separates the two concepts.

A conventional measure of the linkage between technologies is patent citations, which record acknowledged knowledge flows from the cited to the citing patents (Jaffe et al., 1993; Liu and Ma, 2021). Citations carry information related to, but distinct from, compatibility. In particular, complementary components of the same product are under no obligation to cite each other. Inside an electric vehicle, for example, the battery's electrode chemistry determines the temperature window within which cells operate safely, so the vehicle's thermal-management system is designed around the chemistry; patents on the two technologies could nevertheless be unconnected in the citation network.[25] The text-based measure, being symmetric and independent of citing behavior, captures such relationships directly. Still, in Appendix B.2, we show that our measure correlates with bilateral citations, and all empirical results hold when citations are used to measure technology proximity instead.

Validation. Appendix B.2 provides case studies demonstrating that the embedding-based similarities capture technological compatibility between firms along the production chain.It also shows that patents within the same class are more similar than those in different classes; sectors with stronger input-output linkages exhibit greater similarity; and geographically proximate countries have more similar patents. These patterns suggest that the embeddings indeed capture meaningful variations in technology proximity.

Regression Sample. Our empirical analysis (Section 4.2) investigates the causal effect of trade barriers on technology proximity by exploiting changes in HS-6-level MFN tariffs. We tailor the empirical design to the variations in such tariffs, measuring similarity at the level of the importing country, exporting country, HS-6 product, and time (Similarity$_{do,HS6,t}$). Measures at other levels of aggregation will be used to discipline the model in Section 5.

The analysis uses data from 2000–2014, which we group into five three-year intervals. We aggregate countries in the World Input-Output Tables (WIOT) into 29 regions based on geographic and political proximity; see Table B.1 for the full list of regions. (In the rest of the paper, we use 'regions' and 'countries' interchangeably.) Table B.2 in the appendix reports summary statistics for the key similarity measures used in the paper.

Relationship between Technology and Bilateral Trade. Figure 1 takes a first look at the relationship between bilateral trade and technology similarity using country-sector-level data. The vertical axis shows log exports from sector j of country o to sector i of country d. The horizontal axis plots the text-based technology similarity between (d, i) and (o, j). Included in the controls are bilateral distance metrics and income differences between d and o, and fixed effects by d-i, o-j, and i-j.

[25] For instance, Tesla's patent on vehicle thermal management (US10967702B2) and Panasonic's patent on the electrode chemistry of lithium-ion cells (US10833317B2) cover parts engineered jointly within the same vehicles; the two patents do not cite each other (the example is discussed further in Chen et al., 2026).

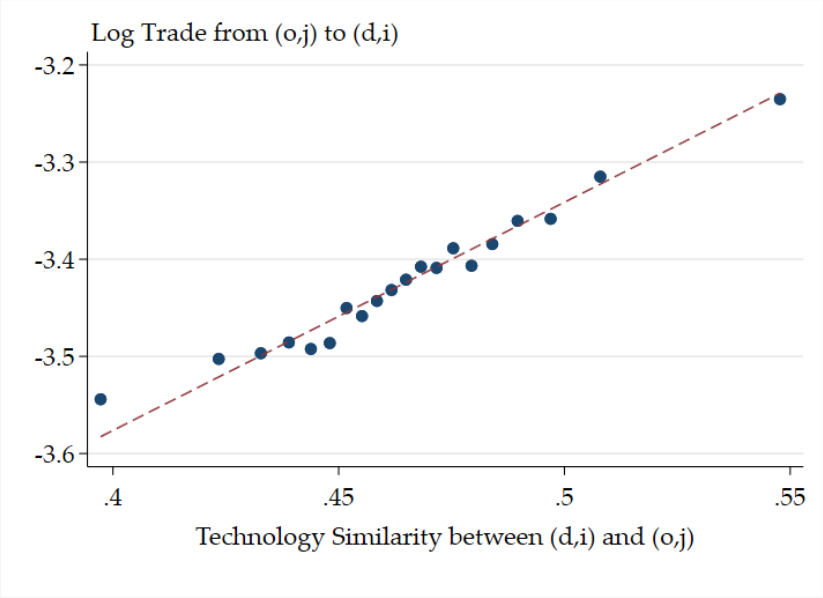


Figure 1: Technology Proximity and Trade

Notes: The figure is a binned scatter plot of log exports from country o sector j to country d sector i against the technology similarity between (d, i) and (o, j). Technology proximity is measured as the cosine similarity of the average textual embeddings of patent abstracts between (d, i) and (o, j). Trade data is from the World Input-Output Database (WIOD). The plot controls for d-i, o-j, and i-j fixed effects, as well as distance metrics (geographic, language, colonial ties) and income differences between o and d. The regression coefficient is 1.9 (s.e. 0.2 clustered by country pair).

The figure shows that buyers source more from sellers with similar technologies. This pattern contrasts with the pure Ricardian view of trade—where exogenous technology differences drive trade and greater differences lead to more trade—but is consistent with our model. In the rest of this section, we move beyond this correlation and examine how technology proximity responds to exogenous shocks in bilateral tariffs. We briefly describe the additional datasets used; further details on data preparation are provided in Appendix B.1.

4.2 Tariff Shocks and Technology Similarity

We test the following implication of our model: a decline in bilateral tariffs between country pairs o and d increases technology similarity between o and d (relative to other country pairs).

Data. Tariff data come from the UN TRAINS database, which records the ad valorem tariff faced by each $(o, HS6)$ product in destination d at time t. In addition to these applied tariffs, we collect most-favored-nation (MFN) tariffs. Pairing the tariffs with technology similarity, constructed at the same $(d, o, HS6, t)$ level as described in Section 4.1, yields a panel covering 28 regions for both d and o (excluding ROW in Table B.1) and five three-year intervals from 2000–2014.

Table 1: Tariff Shocks and Technology Similarity

	$\ln \text{Trade}_{do,HS6,t}$		$\text{Similarity}_{do,HS6,t}$	
	OLS	IV	OLS	IV
	(1)	(2)	(3)	(4)
$\ln \tau^{MFN}_{do,HS6,t}$	-1.075***		-0.007***	
	(0.105)		(0.001)	
$\ln \tau^{Applied}_{do,HS6,t}$		-1.549***		-0.009***
		(0.153)		(0.001)
$\text{Competition}_{do,HS6,t}$	-0.271**	-0.242**	0.711***	0.711***
	(0.115)	(0.115)	(0.002)	(0.002)
FE *d-o-HS6*	Yes	Yes	Yes	Yes
FE *o-HS6-t*	Yes	Yes	Yes	Yes
FE *d-HS6-t*	Yes	Yes	Yes	Yes
Observations	4,567,929	4,567,929	5,652,028	5,652,028
R^2	0.902	0.902	0.995	0.995

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the *d-o-HS6* level. Columns 1 and 3 report reduced-form regressions using OLS, while columns 2 and 4 report the second stage of 2SLS regressions using variations from MFN tariffs to instrument for applied tariffs.

Specification. We estimate the following regression:

$$\text{Similarity}_{do,HS6,t} = \beta \ln \tau_{do,HS6,t} + FE_{d,HS6,t} + FE_{o,HS6,t} + FE_{do,HS6} + \epsilon_{do,HS6,t}, \tag{30}$$

where $\tau_{do,HS6,t}$ is the applied tariff imposed by importer d on an HS-6 product from o. To address endogeneity in $\tau_{do,HS6,t}$ (e.g., countries might lower tariffs on partners from whom they seek to adopt foreign technologies), we follow Boehm et al. (2023) and instrument $\ln \tau_{do,HS6,t}$ with the interaction between (i) the change in MFN tariff rates on the product in importer d ($\ln \tau^{MFN}_{d,HS6,t}$) and (ii) an indicator of whether this MFN rate applies to exporter o in period t. This design, explained in detail in the appendix, exploits the changes in MFN tariffs that affect some of an importer's partners but not others. We exclude the largest exporter for each importer-product pair to mitigate concerns that MFN tariffs are set with the major exporters in mind.

We include three sets of fixed effects to account for time-invariant and time-varying confounding factors. $FE_{d,HS6,t}$ absorbs changes in the importer d's similarity to all producers of that HS-6 good in each period t. $FE_{o,HS6,t}$ captures the average similarity of all importers with a particular exporter-product pair $(o, HS6)$ at time t. Finally, $FE_{do,HS6}$ absorbs any pair-specific aspects of $(d, o, HS6)$ that correlate with both tariffs and technology similarity. All specifications control for the competition score $\text{Competition}_{do,HS6,t}$, so that the estimated effect of tariffs on similarity operates through the compatibility of technologies rather than competition among similar products.

Results. Table 1 reports the results. Columns 1 and 2 use the log of HS-6-level trade flow as the outcome variable and log tariffs as the explanatory variable. Column 1 estimates the

reduced-form specification, while Column 2 estimates the 2SLS specification, instrumenting $\ln \tau^{Applied}_{do,HS6,t}$ with the IV constructed above. In both cases, the estimate is between 1 and 1.6, consistent with the fixed-effect specification elasticity estimate in Boehm et al. (2023).

Columns 3 and 4 replace the outcome variable in the first two columns with technology similarity. According to the 2SLS estimate (Column 4), doubling the applied tariff reduces technology similarity by 0.009, or about 8% of the standard deviation of bilateral similarity. Since tariffs are only one of many barriers to bilateral trade (such as informational and cultural barriers), this estimate suggests that trade barriers can significantly influence technology choices.

In Appendix B.6, we show that using citations as a measure of bilateral similarity leads to similar findings. We also show that the estimates remain robust if we focus on goods primarily intended for intermediate use. This is reassuring given our model's emphasis on intermediate input and supplier compatibility.

Our findings here, and the appendix results using patent citations as a measure of similarity, relates to a literature pioneered by Jaffe et al. (1993), which uses citation flows to track technology diffusion. Most closely related, Aghion et al. (2021) show that after a French firm begins exporting to a destination country, its patents receive more citations from that country. We differ from Aghion et al. (2021) and other existing works in two aspects. First, we exploit exogenous MFN tariff variation to isolate the causal effect of trade costs on bilateral patent similarity. Second, our model interprets these findings through the lens of horizontal technology choice in production networks, complementing existing works that focus on other channels of technological diffusion.

Further evidence for the compatibility mechanism comes from deep trade agreements. In Chen et al. (2026), we extend the similarity measure to a much larger set of countries over 1980–2020 and use an event-study design to estimate the effects of preferential trade agreements, bilateral investment treaties, and WTO accession on bilateral technology similarity. Deep agreements and investment treaties lead to large and persistent increases in similarity, WTO accession to smaller ones, and shallow agreements to little: the effects are concentrated in agreements that go beyond market access to harmonize regulations and technical standards. Using data from standard-setting organizations, we show that part of the effect materializes through partner countries adopting shared standards.

5 Quantification

We use the model to quantify how trade influences firms' technology choices and how these choices affect welfare. We begin by outlining the functional-form assumptions and approximations that enable closed-form solutions for firms' technology choice. We then discuss our parametrization strategy, followed by counterfactual exercises.

5.1 Functional-Form Assumptions and Numerical Algorithm

In our model, firms directly interact with the distribution of heterogeneous firms rather than indirectly through aggregate prices and quantities. Solving for the equilibrium thus requires tracking the entire distribution of Θ_d^i across all (d,i). To maintain tractability, we assume:

Assumption 2. *For each (d,i), the endowment technology distribution is Normal with a mean of $\bar{\mu}_d^i$ and a variance of $(\bar{\sigma}^i)^2$. That is, $\bar{\Theta}_d^i \sim \mathcal{N}(\bar{\mu}_d^i, (\bar{\sigma}^i)^2)$.*

Under this assumption and a quadratic approximation to the log of $C_o^j(\theta)$ (defined in (5)) around the mean ex-post technology in each (o,j), we derive closed-form expressions for firms' technology choices and the resulting ex-post technology distributions:

Proposition 7. *Under Assumptions 2, up to a second-order approximation of $\ln C_o^j(\theta)$:*

1. *$\ln C_o^j(\theta)$ is given by the following function:*

$$\ln C_o^j(\theta) = k_{A,o}^j + m_A^j(\theta - n_{A,o}^j)^2, \tag{31}$$

where coefficients $\{m_A^j, n_{A,o}^j, k_{A,o}^j\}$ can be solved as a function of primitives and wages $\{w_o\}$.

2. *Firms' technology choice $g_o^j(\cdot)$ is characterized by*

$$g_o^j(\bar{\theta}) \equiv \alpha_o^j + \beta^j \bar{\theta}, \tag{32}$$

with $\alpha_o^j = \frac{(\eta-1)m_A^j}{\frac{1}{2}\bar{\phi}+(\eta-1)m_A^j} n_{A,o}^j$ and $\beta^j = \frac{\frac{1}{2}\bar{\phi}}{\frac{1}{2}\bar{\phi}+(\eta-1)m_A^j}$.

3. *Hence, the ex-post technology distribution Θ_o^j is also Normal:*

$$\Theta_o^j \sim \mathcal{N}(\mu_o^j, (\sigma^j)^2), \text{ with } \mu_o^j = \alpha_o^j + \beta^j \bar{\mu}_o^j \text{ and } \sigma^j = \beta^j \bar{\sigma}^j. \tag{33}$$

Proof. See Appendix C.1. □

Proposition 7 expresses in closed-form the parameters governing production costs and firms' technology choices.[26] In particular, equation (31) shows that we can recover the function $C_o^j(\theta)$ by solving for a small number of coefficients. Equation (32) is the policy function for technology choice, which circumvents the need to numerically solve the technology choice problem (15). The linearity of the policy in $\bar{\theta}$, together with Normal ex-ante technologies, also implies Normally distributed ex-post technologies, whose means and variances follow directly from the closed-form expressions in (33).

[26] Proposition 7 can be extended to allow the variance of the ex-ante technology distribution to vary not only by sector but also by country (i.e., $\bar{\sigma}_d^i$ instead of $\bar{\sigma}^i$); it can also be generalized to multi-dimensional θ with multivariate Gaussian ex-ante technology distributions. In practice, the one-dimensional setup is sufficient to capture the key features of bilateral technology similarity data, as we show below.

Building on this proposition, we devise a tractable algorithm to solve for the equilibrium. First, given wages $\{w_o\}$ and model fundamentals, solve for firms' technology choices $\{g_o^j(\bar{\theta})\}$ and cost parameters $\{C_o^j(\theta)\}$. This amounts to solving a contraction mapping problem for the coefficients in (31). Second, compute the sourcing decisions for both intermediate firms $\{\chi_{do}^j(\theta,\tilde{\theta})\}$ and final-good producers $\{\pi_{do}^j(\theta)\}$. Third, solve for industry aggregates, $\{X_o^j(\theta)\}$ and $\{M_o^j(\theta)\}$, using equations (18) to (21). Finally, check the labor market clearing condition (22). If it is satisfied, the equilibrium is found; if not, update $\{w_o\}$ and iterate from the first step. Appendix C.2 provides details for this algorithm.

5.2 Calibration Procedure

We now describe how we calibrate the model parameter, targeting the 2014 world economy.

Parameters Calibrated Externally. We set the elasticity of substitution for both final-good producers and intermediate firms to be 5. This implies a 20% markup for final-good producers' purchase from intermediate firms.

We take labor endowments $\{L_d\}$ from the Penn World Table population data. We derive production-function coefficients $\{\rho_d^i, \gamma^{ij}, \gamma^{iL}\}$ from the World Input-Output Database (WIOD), matching sectoral consumption shares in final-good production (ρ_d^i) and sectoral input shares (γ^{ij} and γ^{iL}). All these values are taken over the average of 2012-2014.

Technology Distribution. We recover the (unobserved) *ex-ante* technology distributions, $\{\bar{\mu}_d^i, \bar{\sigma}^i\}$, along with other model primitives, in two steps. In the first step, we choose the parameters governing the *ex-post* distributions, $\{\mu_d^i, \sigma^i\}$, to match the similarities between country-sector pairs and their variances as reflected in patent text embeddings. In the second step, we calibrate the production and sourcing blocks of the model by treating firms' technology as exogenously fixed at the calibrated ex-post distribution, exploiting the result that, conditional on the chosen technology, firms' sourcing decisions are independent of their endowment technology. Once this is complete, we calibrate $\bar{\phi}$ (which governs the cost of technology adaptation) and then recover the ex-ante distributions by inverting the firms' policy functions.[27]

To implement the first step, we map the technological distances in our model to a similarity metric bounded between 0 and 1, facilitating the comparison between model and empirical bilateral similarity.[28] We define the similarity between two firms with technologies θ and $\tilde{\theta}$ as:

$$sim(\theta,\tilde{\theta}) = \exp(-(\theta-\tilde{\theta})^2). \tag{34}$$

[27] Note that $\bar{\phi}$ affects only the allocation of firms' markups between profits and technology adaptation costs. For this reason, once the ex-post distributions and other general equilibrium outcomes (such as wages and prices) have been calibrated, $\bar{\phi}$ can be adjusted without altering earlier calibration steps (see Appendix C.4 for details).

[28] Although cosine similarity theoretically ranges from $[-1,1]$, the entries of the empirical similarity matrix (at the country-sectoral level) are all positive, making the logarithm transformation possible.

Analogously, the similarity between the mean technologies of two country-sector pairs is:

$$sim(\mu_d^i, \mu_o^j) = \exp(-(\mu_d^i - \mu_o^j)^2). \tag{35}$$

We choose $\{\mu_d^i\}$ by solving the following problem:

$$\min_{\{\mu_d^i\}} \sum_{d,i,o,j} \left[\ln sim(\mu_d^i, \mu_o^j) - \ln sim_{do}^{ij,Data}\right]^2, \tag{36}$$

where $\ln sim_{do}^{ij,Data}$ is the log of empirical technology similarity between country-sector pairs (d,i) and (o,j), averaged over the period 2012-2014, residualized by i-j fixed effects. With this residualization, our calibration does not attempt to explain why, for example, the auto industry is inherently closer to the auto parts industry than to agriculture. Instead, it focuses on the differences between pairs of countries within the same industry pair, e.g., whether the German auto sector is closer to the Chinese or American auto sector.[29]

Problems of the form (36) is known to be challenging.[30] We show in Appendix C.3 that by fitting the logarithms of the similarities, it can be transformed into a classical multidimensional scaling problem, which has unique solutions and can be efficiently solved for $\{\mu_d^i\}$.

To infer the standard deviation of technology, $\{\sigma^i\}$, we match the standard deviation of similarity between technologies within each (d,i). For each country-sector pair, we randomly sample 1,000 patents from year 2014, calculate the pairwise similarity of their abstract embeddings, and compute the standard deviation of these similarities by sector. We then average this statistic across countries to infer $\{\sigma^i\}$.

The second step of our calibration parameterizes the trade and production aspects of the model to match the data across a set of moments, as explained below.

Productivity and Trade Costs. We calibrate country-sectors' productivity $\{A_d^i\}$ to match the output share of (d,i) in each sector i, and trade costs $\{\tau_{do}^j, \tau_{do}^{Uj}\}$ to match the trade shares of intermediate and final goods.[31]

Technology Compatibility $\bar{t}$. The parameter $\bar{t}$ governs firms' input-sourcing decisions given their chosen technology. As Proposition 3 shows, conditional on the trade elasticity ξ, the firm-level correlation between importing and technology similarity identifies $\bar{t}$. We

[29]This approach is justified by two key reasons. First, our counterfactual analysis centers on how international trade shocks shift the technological positions of countries. Although our one-dimensional technology space may not capture all variations across sectors, the most important variation for our study is the difference between countries within a sector pair. By matching the residualized data, we match precisely these cross-country differences. Second, under the Cobb-Douglas production function of intermediate-good producers, adding exogenous technological distance between pairs of sectors amounts to adding exogenous mutilative cost shifters. Therefore, enriching the model (or its calibration) to match exogenous sectoral distance will not affect the counterfactual findings. Our calibration and counterfactuals have the interpretation of holding constant the *sectoral* technological differences while focusing on country differences.

[30]See e.g., Chapter 13 of Borg and Groenen (2007).

[31]Trade and output shares do not separately identify all $\{A_d^i\}$ and all $\{\tau_{do}^j, \tau_{do}^{Uj}\}$ jointly (Burstein and Vogel, 2017). Our counterfactual results are invariant to normalization choice in the calibration of these parameters.

Table 2: Summary of Model Parameters

Parameters	Descriptions	Value	Target/Source
A. Externally Calibrated			
$\gamma^{ij}, \gamma^{iL}, \alpha^j$	IO Structure and Consumption Share	-	WIOT; $N = 29, S = 19$
L_d	Labor Endowment	-	PWT
$\eta - 1, \xi - 1$	Trade Elasticity	4	Simonovska and Waugh (2014)
B. Just-Identified			
A_d^i	Productivity	-	Countries' output shares in each sector
$\tau_{do}^j, \tau_{do}^{Uj}$	Iceberg Trade Costs	-	Bilateral trade shares (WIOT)
$\bar{t}$	Compatibility Incentive	12.2	Firm-level import-similarity corr: 0.021
$\bar{\phi}$	Adaptation Costs	0.58	Country-sector-level similarity-tariff elas.: -0.007
C. Best Fit			
μ_d^i, σ^i	Ex-post Technology Distribution	-	Tech similarity between countries and sectors
$\bar{\mu}_d^i, \bar{\sigma}^i$	Ex-ante Technology Distribution	-	Inverted from calibrated ex-post distributions

estimate the correlation using matched data on Chinese manufacturing firms, their patents, and their customs transactions, regressing an indicator for whether a firm imports from a region on the firm's technology similarity with that region. Controlling for a rich set of fixed effects, the coefficient is precisely estimated at 0.021 (Column 3 of Table B.9); Appendix B.5 details the specification and presents evidence supporting the compatibility interpretation. We calibrate $\bar{t}$ by matching this coefficient using the simulated method of moments. Since the regression focuses on Chinese firms, our simulation centers on China too. For each sector i, we draw $10,000$ firms from the ex-post distribution Θ_{CHN}^i. We simulate whether a firm with technology θ imports from (o, j) using equation (8).[32] We then regress the import indicator on the technology similarity between θ and (o, j), defined as $sim(\theta, \mu_o^j)$, controlling for o-i and firm fixed effects. This exercise sets $\bar{t} = 12.2$, under which the model-simulated regression matches the empirical coefficient of 0.021.

Adaptation Costs $\bar{\phi}$. We choose the adaptation cost parameter $\bar{\phi}$ to match the elasticity of technology similarity with respect to tariffs, as estimated in Column 3 of Table 1. As Proposition 2 illustrates, conditional on $\bar{t}$ and other model fundamentals, this elasticity identifies $\bar{\phi}$. To calculate the model-based elasticity, we simulate changes to trade costs $\{\tau_{do}^j\}$, mimicking the size of the actual MFN tariff changes observed over time for each (d, j). This produces a counterfactual equilibrium. We then treat the calibrated equilibrium and the counterfactual as two periods and regress the model's technology similarity between d and (o, j) on the logarithm of τ_{do}^j, controlling for a range of fixed effects mirroring the empirical specification. We adjust $\bar{\phi}$ until the model's elasticity matches the empirical coefficient of -0.007, yielding a calibrated value of $\bar{\phi} = 0.58$.

[32] We interpret the firm-level regression through the random network formation microfoundation of the CES sourcing structure (Appendix A.1): each firm sources each input from the lowest-cost supplier in its randomly formed network, so equation (8) gives the probability that a firm imports from a region, an extensive-margin choice at the firm level.

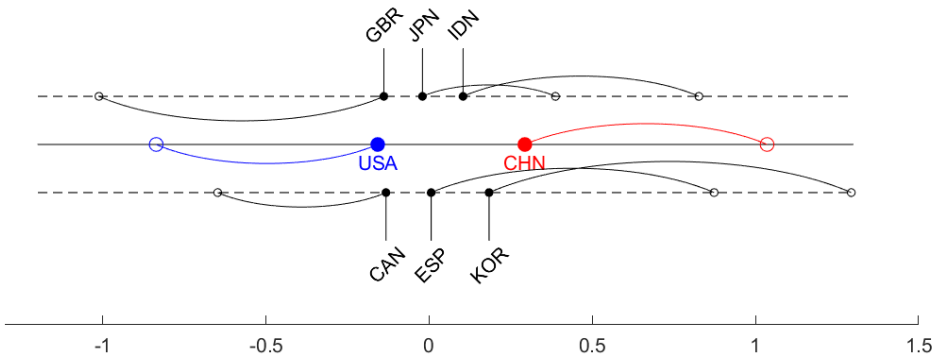


Figure 2: Estimated Technology Distributions for Selected Economies

Notes: Filled dots are the means of firms' ex-post technology distributions, $\{\mu_d^i\}$, averaged across sectors for each d for selected countries. Unfilled circles are the means of their endowment technology $\{\bar{\mu}_d^i\}$.

Numerical Implementation. We implement the calibration as follows. First, we estimate the ex-post distribution parameters $\{\mu_d^i, \sigma^i\}$ to best match the empirical technology similarity within and between country-sectors, as previously described. Given these parameters, we employ a nested algorithm for the second step. In the outer layer, we adjust $\bar{t}$ to match the firm-level extensive margin import-similarity regression coefficient. In the inner loop, we solve for the competitive equilibrium, while simultaneously calibrating $\{\tau_{do}^j, \tau_{do}^{Uj}, \Xi^i / A_d^i\}$ to match trade shares. We then simulate the model and feed the regression coefficient back into the outer loop. Once this process is complete, we calibrate $\bar{\phi}$ to match the country-sector similarity-tariff elasticity and use the results from Proposition 7 to back out the ex-ante distribution parameters $\{\bar{\mu}_d^i, \bar{\sigma}^i\}$ from $\{\mu_d^i, \sigma^i\}$. Appendix C.4 reports the detailed procedure.

Table 2 summarizes the calibration. We discuss the implications of our calibration in the next subsection.

5.3 Calibrated Technologies and Model Fit

Calibrated Technologies. Figure 2 shows the average (across sectors) of the mean locations of ex-ante and ex-post technology distributions for a subset of countries. The filled dots depict the locations of ex-post technologies. Among major economies, the technologies of the U.S. and China fall at the two ends of the spectrum, whereas the technologies of other major economies, such as Japan, Canada, and Western European countries, fall in the middle. This pattern highlights a key feature of bilateral technology similarity statistics derived from the textual analysis of patent abstracts: within sector pairs, American and Chinese technologies tend to be less similar to each other than they are to those of other countries.

The circles in Figure 2 depict the average locations of ex-ante (endowment) technologies

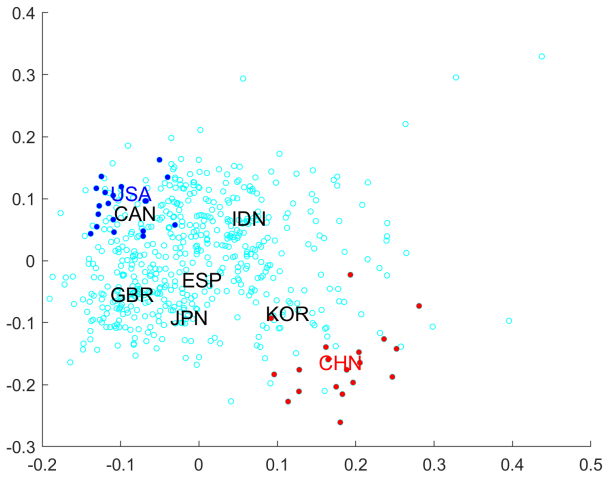


Figure 3: First Two Principal Coordinates of the Similarity Matrix

Notes: Plotted are the first two principal coordinates of the similarity matrix between country-sector pairs. The similarity measure is defined as the log of cosine similarity of the average embeddings of patent abstracts between two country-sectors (d, i) and (o, j), netting out sector-pair (i.e., i-j) fixed effects. Each dot corresponds to the principal coordinates of a country-sector pair, obtained by solving a problem analogous to (36) in the two-dimensional space for technology. Country name indicates the average coordinates of the country across sectors. Sectors in China and the U.S. are indicated in red and blue, respectively.

for these countries. Two patterns are noteworthy. First, ex-ante technologies tend to cluster spatially, with Asian economies grouped on the right and other economies clustered on the left. This clustering may reflect the influence of geographic, historical, or cultural factors on initial technological endowments. Second, with technology compatibility incentives, international trade brings countries' technologies closer. Indeed, if countries do not engage in trade, then firms will move toward the average technologies of their own countries.

Assessing Fit of Technology Similarity. Our model allows technology to be represented in a high-dimensional space. For simplicity and transparency, however, we restrict our calibration to a one-dimensional space. To demonstrate that the calibrated technology positions are not an artifact of our one-dimensional representation, Figure 3 plots the first two principal coordinates of the similarity matrix for country-sector pairs in a two-dimensional technology space, which correspond to the calibration results (up to a scaling factor) if the two-dimensional space were used. The figure shows that the divergence between U.S. and Chinese technologies, with other countries positioned in between, remains a salient feature of the data when viewed in this space.

To evaluate more systematically how well the calibration captures the cross-sectional technology similarities between country-sectors, we regress the log of empirical technology similarity between country-sectors on the log of model-predicted similarity, controlling for country-sector fixed effects. Column 1 of Table 3 reports the result. Approximately 21% of

Table 3: Bilateral Similarity Between Country-Sectors: Model v.s Data

Log Similarity in Model	Log Similarity in Data (1)	(2)
at Ex-post Tech. Dist.	0.363***	
	(0.001)	
at Ex-ante Tech. Dist.		0.001***
		(0.000)
Obs.	302,664	302,664
Within R^2	0.212	0.128
Fixed effects	*d-i*, *o-j*	*d-i*, *o-j*

Notes: This table assesses the goodness of fit in the model. Each column reports the regression of data log similarity between country-sector pairs on model log similarity. Column 1 uses the ex-post technology distributions $\{\mu_o^j, \sigma^j\}$ to construct model similarity. Column 2 uses the ex-ante distributions $\{\bar{\mu}_o^j, \bar{\sigma}^j\}$. Both columns control for *d-i* and *o-j* fixed effects. The reported within R^2 excludes the variation explained by the fixed effects.

the variation in bilateral similarities is explained by the model. Note that fully describing the bilateral similarity matrix between country-sectors would typically require $\frac{(NS)^2-NS}{2}$ parameters. Our model captures a substantial share of the variation by choosing $\{\mu_d^i\}$, totaling only $NS-1$ parameters. As another benchmark, the explanatory power of bilateral distance for bilateral trade, one of the most successful empirical model, is around 25%-50%.

Relevance in Shaping Technology. To further unpack the model's fit and highlight the role of compatibility incentives in shaping technology interdependence through trade, Column 2 of Table 3 regresses the log of data similarity on the log of model similarity implied by the calibrated endowment technologies. Here, the regression coefficient declines to almost zero, and the R^2 also falls from 21% to 13%. This comparison suggests that two-fifths of the model's explanatory power (8% out of 21%) can be attributed to endogenous technology choices driven by compatibility incentives, whereas the remaining comes from countries' endowment technology distributions.

Relevance in Shaping Trade Costs. Under our calibration, on average (weighted by trade flows), costs arising from technology incompatibility is equivalent to 9.5% ad valorem trade costs. The literature has found that, to account for the low level of international trade, iceberg trade costs must be substantially higher than the level of tariffs (Anderson and van Wincoop, 2004). Our model suggests that a modest but non-negligible part of these unobserved trade barriers could be technological incompatibility.

5.4 Counterfactual: Trade Conflict and Technology Decoupling

We use our model to evaluate the welfare costs of trade conflicts, with an emphasis on the role of endogenous technology choice. To this end, we consider a counterfactual that is intended to speak to the recent silicon blockade of the U.S. and its allies against China. In

Table 4: Welfare Costs of a Semiconductor and Electronics Embargo on China

	Import share in targeted China's sector (%)	$\Delta Sim_{US,China}$	Δ Welfare (%) China	U.S.	Others	World
U.S. only	1.18	0.000	-0.006	-0.002	-0.001	-0.002
All countries	100	-0.014	-0.546	-0.056	-0.167	-0.232

Notes: This table reports the impact of an embargo on China's imports of Computer, Electronic, and Optical Products from the U.S. and other countries. $Sim_{US,China}$ is the technology similarity between the U.S. and China as defined by equation (35), averaged across sectors. Welfare is defined as household real income.

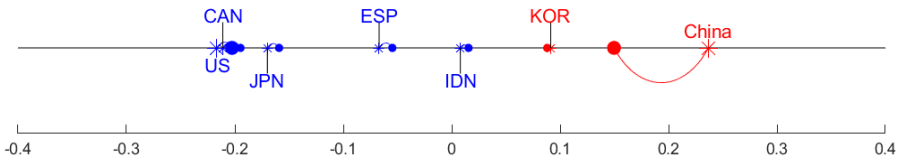
this counterfactual, we shut down China's imports of intermediate goods in the sector of Computer, Electronic, and Optical Products from the U.S. and—due to the long-arm jurisdiction of U.S. sanctions—from other countries as well. In our data, this sector accounts for approximately 19.7% of China's total imports.

To implement this counterfactual, we shut down exports in this sector to China by raising the corresponding trade costs to infinity. We investigate two cases: (i) only the U.S. imposes the embargo, and (ii) all countries impose the embargo. In both cases, countries' endowment distributions are set to be the same as their endowment distribution in the calibrated equilibrium. The comparison between these counterfactuals and the baseline equilibrium reveals what would happen when the economy starts from the same technological primitives but operates under different trade environments.[33]

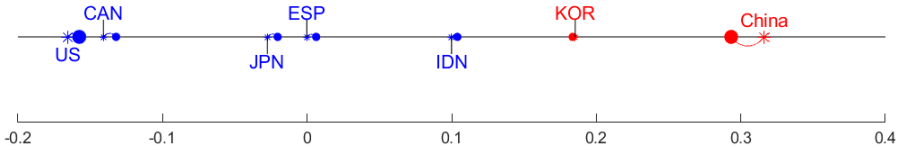
Table 4 reports the main findings. Since the U.S. accounts for only 1.2% of China's imports in this sector, when only the U.S. imposes the embargo, the impact on the technology distance between the two countries is minimal. While the embargo imposes more harm on China than the U.S., the overall economic damage is limited.

The second row of the table reports the results when all countries impose the embargo at the same time. China now sustains a welfare loss of 0.55%, about 100 times the loss when only the U.S. imposes the export ban. In addition to the larger loss for China, two other results are noteworthy. First, the distance between Chinese and U.S. technologies increases substantially more than in the first experiment, despite the fact the direct U.S. sanctions are the same in both cases. This occurs because, among major economies, China and the U.S. occupy the two ends of the technology spectrum. As Chinese technologies shift away from those of other countries due to the embargo, they also shift away from U.S. technologies. Second, the welfare cost to the U.S. is larger than before. This amplification occurs for two reasons: first, production costs in China increase due to the direct effect of the embargo, and the rest of the world must now pay higher prices for Chinese products; second, the

[33]By holding endowments in these counterfactuals fixed at the ex-ante distribution in the calibrated equilibrium, we adopt the view that countries' technology endowments are persistent. According to this view, sustained investment in technology adaptation is necessary for firms to deviate from their endowed technology, so different counterfactuals start from the same endowment.



(a) The Average Technology in the Targeted Sector



(b) The Average Technology of All Sectors

Figure 4: Technology Distributions Before and After the Embargo for Selected Economies

Notes: Dots are the ex-post means in the baseline equilibrium, and stars are the equilibrium with the embargo. Blue indicates countries moving toward the U.S. and red indicates countries moving toward China.

endogenous divergence in the technologies around the globe makes souring intermediate inputs even more costly.

The Role of Technology Decoupling. To shed light on the role of endogenous technology, Figure 4 plots the changes in countries' technologies due to the embargo. The upper panel depicts the change in the mean technology of the targeted industry. As discussed earlier, the embargo leads to a divergence of technologies between China and the U.S.[34] This divergence results in a realignment of other countries' technologies. We indicate a country in blue if its technology moves toward the U.S. and in red if it moves toward China. Among others, Japan, Mexico, and European nations move closer to the U.S., while Korea is one of the few that shift toward China. However, even as most countries' technologies move toward the U.S., the gap between their technologies and U.S. technology may still widen because the change in technology choice is smaller than that of the U.S. This widening technology gap contributes to amplified welfare costs for the U.S.

Because of the compatibility incentives, the divergence in technologies in semiconductors spreads to all other sectors via input-output linkages. The lower panel plots the mean technology for each country across all sectors, including those not directly affected by the embargo. The changes are qualitatively similar to those in the target sector.

To underscore the importance of endogenous technology decoupling on welfare, we rerun the embargo scenarios under the assumption that firms cannot change their technology

[34]The embargo directly impacts the technology choices of downstream sectors. For the Computer, Electronic, and Optical Products sector, the most affected downstream sector is itself, as it relies heavily on its own inputs.

Table 5: Mechanism Decomposition

	$\Delta Sim_{US,China}$	Δ Welfare (%)			
		China	U.S.	Others	World
Fixed Technology	0.000	-0.43	-0.01	-0.04	-0.12
+ response of targeted China's sector	-0.002	-0.33	-0.02	-0.06	-0.11
+ response of all sectors in China	-0.010	-0.51	-0.03	-0.10	-0.18
+ response of all country-sectors	-0.014	-0.55	-0.06	-0.17	-0.23

Notes: The rows show the effects of allowing different degrees of changes in firms' technology choices under the embargo. The 'Fixed Technology' row represents the scenario where technology choices remain fixed at their calibrated equilibrium values.

choice after the shock, effectively freezing technology choices at the pre-embargo equilibrium ('Fixed Technology'). The results, shown in the first row of Table 5, reveal that when technology choice is restricted, welfare losses are significantly smaller for China, the U.S., and the world. In contrast, when firms can adapt their technology, technological responses amplify welfare losses—by 28% for China (from 0.43% to 0.55%), by 500% for the U.S. (from 0.01% to 0.06%), and by 92% for the global average (from 0.12% to 0.23%). This amplification occurs because, while individual firms benefit from the ability to choose technologies most suitable for the embargo trade environment, the broader economy suffers from the negative externalities arising from technology incompatibility across sectors and countries. These externalities compound via input-output linkages, increasing the overall welfare costs of the embargo.

We decompose the role of various margins of technology choice in the remaining rows of Table 5, gradually introducing technology changes—from the most directly affected sector to all sectors in China, and to all other countries. Row 2 shows that when firms in the directly targeted sector in China adjust their technology, China's welfare losses actually decrease. These firms can re-optimize by balancing sourcing efficiency against adaptation costs in response to changes in sourcing origins, as reflected in Figure 4a. However, these unilateral adjustments create externalities for downstream sectors in other countries, contributing to additional welfare losses abroad.

Row 3 shows that as firms in non-targeted sectors in China adjust their technology choice, the negative externalities outweigh the benefits, leading to greater welfare losses for China and foreign countries compared to the 'Fixed Technology' scenario. Finally, Row 4 demonstrates that when firms in other countries adjust their technologies, technological decoupling becomes more pronounced, further increasing global welfare losses. Almost all of the U.S. loss and half of the global loss from the embargo arise because of endogenous technological decoupling.

These counterfactual findings echo the externality mechanisms discussed in Propositions 5 in the two-country setting. It is worth noting that in an environment with multiple coun-

tries, endogenous technological divergence between any two does not necessarily reduce welfare. For instance, if a trade conflict between Korea and Japan leads to a technological divergence between the two, welfare may improve—if the positive externalities of Japan aligning more closely with U.S. technology and Korea with Chinese technology outweigh the negative externalities of divergence between Korea and Japan. This possibility highlights the value of developing a quantitative framework and disciplining it with data on technology similarity.

6 Conclusion

In this paper, we construct a quantitative trade model that incorporates horizontal technology choice. In our framework, firms jointly decide their technology and where to source their inputs under compatibility incentives, with both decisions shaped by general equilibrium forces. We establish sufficient conditions for equilibrium existence and uniqueness, and show that compatibility gives rise to an externality: because firms capture only part of the surplus their technology creates, the technology response to trade shocks has a first-order welfare effect.

Using patent and trade data, we provide novel evidence at both the firm and country levels supporting the core mechanisms of our model. Our quantitative analyses reveal three main findings. First, trade-induced endogenous technology plays a significant role in shaping global technologies, accounting for two-fifths of the variation in technological proximity between countries, while differences in endowment distributions explain the remainder. Second, technology incompatibility imposes an average cost equivalent to 9.5% in ad valorem terms in international trade, underscoring the importance of this mechanism. Last, trade conflicts between the U.S. and China can lead to technological decoupling between them and trigger realignments among other countries. These endogenous responses amplify the losses from the conflict, accounting for almost all of the U.S. welfare loss and half of the global loss.

Our framework can be extended in two directions. First, in reality, firms and countries develop institutions to manage the externalities highlighted in our model. Examples include supply chain integration to internalize externalities and collaborative efforts to establish common standards. Investigating the equilibrium impacts of such mechanisms is a promising research avenue. Second, our model is static and therefore abstracts from dynamics and vertical innovation. Integrating our model into a dynamic setting can illuminate how compatibility incentive in firms' innovation decisions and supplier choices shape both industry dynamics and aggregate growth.

References

Aghion, Philippe, Antonin Bergeaud, Timothee Gigout, Matthieu Lequien, and Marc Melitz, "Exporting Ideas: Knowledge Flows from Expanding Trade in Goods," 2021. Unpublished.

Akcigit, Ufuk, Murat Alp Celik, and Jeremy Greenwood, "Buy, Keep, or Sell: Economic Growth and the Market for Ideas," *Econometrica*, 2016, *84* (3), 943–984.

Allen, Treb, Costas Arkolakis, and Xiangliang Li, "On the Equilibrium Properties of Spatial Models," *American Economic Review: Insights*, 2024, *6* (4), 472–89.

Alvarez, Fernando and Robert E Lucas, "General Equilibrium Analysis of the Eaton–Kortum Model of International Trade," *Journal of Monetary Economics*, 2007, *54* (6), 1726–1768.

Amiti, Mary, Stephen J Redding, and David E Weinstein, "The Impact of the 2018 Tariffs on Prices and Welfare," *Journal of Economic Perspectives*, 2019, *33* (4), 187–210.

Anderson, James E. and Eric van Wincoop, "Trade Costs," *Journal of Economic Literature*, 2004, *42* (3), 691–751.

Arkolakis, Costas, Federico Huneeus, and Yuhei Miyauchi, "Production Network Formation, Trade, and Welfare," CESifo Working Paper 11833, CESifo 2025.

Baqaee, David Rezza and Emmanuel Farhi, "Networks, Barriers, and Trade," *Econometrica*, 2024, *92* (2), 505–541.

Basu, Susanto and David N. Weil, "Appropriate Technology and Growth," *The Quarterly Journal of Economics*, 1998, *113* (4), 1025–1054.

Boehm, Christoph E., Andrei A. Levchenko, and Nitya Pandalai-Nayar, "The Long and Short (Run) of Trade Elasticities," *American Economic Review*, 2023, *113* (4), 861–905.

Boehm, Johannes and Ezra Oberfield, "Misallocation in the Market for Inputs: Enforcement and the Organization of Production," *The Quarterly Journal of Economics*, 2020, *135* (4), 2007–2058.

Bonadio, Barthélémy, Zhen Huo, Elliot Kang, Andrei A. Levchenko, Nitya Pandalai-Nayar, Hiroshi Toma, and Petia Topalova, "Playing with Blocs: Quantifying Decoupling," 2024. Unpublished.

Borg, Ingwer and Patrick JF Groenen, *Modern Multidimensional Scaling: Theory and Applications*, Springer Science & Business Media, 2007.

Buera, Francisco J. and Ezra Oberfield, "The Global Diffusion of Ideas," *Econometrica*, 2020, *88* (1), 83–114.

Burstein, Ariel and Jonathan Vogel, "International Trade, Technology, and the Skill Premium," *Journal of Political Economy*, 2017, *125* (5), 1356–1412.

Cai, Jie, Nan Li, and Ana Maria Santacreu, "Knowledge Diffusion, Trade, and Innovation across Countries and Sectors," *American Economic Journal: Macroeconomics*, 2022, *14* (1),

104–45.

Caliendo, Lorenzo and Fernando Parro, "Estimates of the Trade and Welfare Effects of NAFTA," *The Review of Economic Studies*, 2015, *82* (1), 1–44.

Carluccio, Juan and Thibault Fally, "Foreign Entry and Spillovers with Technological Incompatibilities in the Supply Chain," *Journal of International Economics*, 2013, *90* (1), 123–135.

Chaney, Thomas, "Distorted Gravity: the Intensive and Extensive Margins of International Trade," *The American Economic Review*, 2008, *98* (4), 1707–1721.

Chen, Weili, Jingting Fan, and Wenlan Luo, "Multilateral Institutions and Technology Fragmentation: Lessons from the Era of Global Integration," 2026. Working paper.

Clayton, Christopher, Matteo Maggiori, and Jesse Schreger, "A Framework for Geoeconomics," Working Paper 31852, National Bureau of Economic Research November 2023.

Costinot, Arnaud, "A Comparative Institutional Analysis of Agreements on Product Standards," *Journal of International Economics*, 2008, *75* (1), 197–213.

_ **and Andres Rodriguez-Clare**, "Trade Theory with Numbers: Quantifying the Consequences of Globalization," *Handbook of International Economics*, 2014, *4*, 197.

_ **and Jonathan Vogel**, "Matching and Inequality in the World Economy," *Journal of Political Economy*, 2010, *118* (4).

Demir, Banu, Ana Cecília Fieler, Daniel Yi Xu, and Kelly Kaili Yang, "O-Ring Production Networks," *Journal of Political Economy*, 2024, *132* (1), 200–247.

Dhyne, Emmanuel, Ayumu Ken Kikkawa, Xianglong Kong, Magne Mogstad, and Felix Tintelnot, "Endogenous Production Networks with Fixed Costs," *Journal of International Economics*, 2023, *145*, 103841.

Eaton, Jonathan and Samuel Kortum, "International Technology Diffusion: Theory and Measurement," *International Economic Review*, 1999, *40* (3), 537–570.

_ **and** _ , "Technology, Geography, and Trade," *Econometrica*, 2002, *70* (5), 1741–1779.

_ **, Samuel S Kortum, and Francis Kramarz**, "Firm-to-Firm Trade: Imports, Exports, and the Labor Market," Working Paper 29685, National Bureau of Economic Research 2022.

Fajgelbaum, Pablo D, Pinelopi K Goldberg, Patrick J Kennedy, and Amit K Khandelwal, "The Return to Protectionism," *The Quarterly Journal of Economics*, 2020, *135* (1), 1–55.

Huang, Yi, Chen Lin, Sibo Liu, and Heiwai Tang, "Trade Networks and Firm Value: Evidence from the US-China Trade War," *Journal of International Economics*, 2023, *145*, 103811.

Jaffe, Adam B., Manuel Trajtenberg, and Rebecca Henderson, "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations," *The Quarterly Journal of Economics*, 1993, *108* (3), 577–598.

Jones, Charles I., "Misallocation, Economic Growth, and Input–Output Economics," in Daron Acemoglu, Manuel Arellano, and Eddie Dekel, eds., *Advances in Economics and Econometrics: Tenth World Congress*, Econometric Society Monographs, Cambridge Univer-

sity Press, 2013, p. 419–456.

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy, "Measuring Technological Innovation over the Long Run," *American Economic Review: Insights*, 2021, *3* (3), 303–20.

Kleinman, Benny, Ernest Liu, and Stephen J. Redding, "International Friends and Enemies," *American Economic Journal: Macroeconomics*, 2024, *16* (4), 350–85.

Kortum, Samuel S., "Research, Patenting, and Technological Change," *Econometrica*, 1997, *65* (6), 1389–1419.

Lancaster, Kelvin, *Variety, Equity, and Efficiency: Product Variety in an Industrial Society*, New York Chichester, West Sussex: Columbia University Press, 1979.

Lazear, Edward P., "Culture and Language," *Journal of Political Economy*, 1999, *107* (S6), S95–S126.

Lind, Nelson and Natalia Ramondo, "Global Innovation and Knowledge Diffusion," *American Economic Review: Insights*, 2023, *5* (4), 494–510.

Liu, Ernest and Song Ma, "Innovation Networks and R&D Allocation," Working Paper 29607, National Bureau of Economic Research 2021.

Lybbert, Travis J and Nikolas J Zolas, "Getting Patents and Economic Data to Speak to Each Other: An 'Algorithmic Links with Probabilities' Approach for Joint Analyses of Patenting and Economic Activity," *Research Policy*, 2014, *43* (3), 530–542.

Morales, Eduardo, Gloria Sheu, and Andrés Zahler, "Extended Gravity," *The Review of Economic Studies*, 2019, *86* (6), 2668–2712.

Muennighoff, Niklas, Nouamane Tazi, Loïc Magne, and Nils Reimers, "MTEB: Massive Text Embedding Benchmark," 2023.

Oberfield, Ezra, "A Theory of Input–Output Architecture," *Econometrica*, 2018, *86* (2), 559–589.

Simonovska, Ina and Michael E Waugh, "The Elasticity of Trade: Estimates and Evidence," *Journal of International Economics*, 2014, *92* (1), 34–50.

Timmer, Marcel P, Erik Dietzenbacher, Bart Los, Robert Stehrer, and Gaaitzen J De Vries, "An Illustrated User Guide to the World Input–Output Database: the Case of Global Automotive Production," *Review of International Economics*, 2015, *23* (3), 575–605.