

A Global View of Creative Destruction

Chang-Tai Hsieh
University of Chicago and NBER

Peter J. Klenow
Stanford University and NBER

Ishan Nath *
Princeton University

May 2022

Abstract

We formulate a two-country model of trade and creative destruction by domestic and foreign firms. In the model, trade liberalization quickens the pace of creative destruction and facilitates the flow of technology across countries. The resulting dynamic gains from idea flows are at least as large as the static gains from trade. In our model, such international idea flows are essential for understanding why country technologies do not drift apart, and for matching two properties of export dynamics. First, contracting firms are more likely to lose exports than domestic sales, whereas expanding firms are more likely to gain domestic sales than to gain exports. Exports are vulnerable to foreign as well as domestic creative destruction, whereas domestic sales are comparatively insulated from foreign creative destruction by trade barriers. Second, the product composition of a country's exports exhibits ample turnover. This is consistent with our model, in which a country's comparative advantage is constantly shifting due to global creative destruction.

*We thank Jean-Felix Brouillette, Eric English, Feng Lin, Erxiao Mo, and Kazuatsu Shimizu for excellent research assistance; and both Ariel Burstein and Sam Kortum for very helpful discussions. Hsieh acknowledges support from Chicago Booth's Polsky Center and Klenow from the Stanford Institute for Economic Policy Research. Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed by the U.S. Census Bureau to ensure that no confidential information is disclosed. Contact information: chsieh@chicagobooth.edu, klenow@stanford.edu, and inath33@gmail.com.

1. Introduction

Studies by Bernard and Jensen (1999), Eaton and Kortum (2002), Melitz (2003), and others placed heterogeneous firms at the center of research on international trade. The first wave of follow-up research has mostly focused on models in which trade liberalization leads to a burst of reallocation and growth, but no long run effects on reallocation rates or growth rates.

A growing literature seeks to assess the growth effects of trade. Empirical studies include Bloom, Draca and Van Reenen (2016) and Aghion, Bergeaud, Lequien and Melitz (2020). Modeling efforts build on the foundational work of Krugman (1979), Rivera-Batiz and Romer (1991) and Grossman and Helpman (1993). Recent modeling includes Alvarez, Buera and Lucas (2017), Buera and Oberfield (2020), Perla, Tonetti and Waugh (2021), and Akcigit, Ates and Impulitti (2021). Papers with both models and empirics range from Eaton and Kortum (1999) to Arkolakis, Ramondo, Rodríguez-Clare and Yeaple (2018).

In this paper we present a two-country model on the interaction of creative destruction and trade. In our model, innovating firms improve upon existing technologies. When innovators take over the market for an existing product (creative destruction), export reallocation across countries can take place. Domestic firms can take over foreign markets for a product, and foreign firms can take over the domestic market. This is a two-economy version of the influential Klette and Kortum (2004) model of creative destruction, only with exogenous arrival rates in each country for simplicity.

We assume that innovators can build upon the technology of products sold in their market or on the blueprints of local firms. When innovators build on the technology of sellers, which includes the country's imports, trade then facilitates the flow of ideas across countries. But we also consider a model in which ideas flow across countries independent of trade, as in Ramondo, Rodríguez-Clare and Saborío-Rodríguez (2016) and others. In both versions of the model, the diffusion of ideas generates a constant reallocation of exports between the two countries

and results in the two economies growing at the same rate in the long run.

We calibrate the model to fit manufacturing moments in the U.S. vs. the rest of the OECD. We match TFP growth, relative value added per worker in the U.S. and the OECD, exports relative to all shipments (the trade share), and the sensitivity of trade to trade barriers (the trade elasticity). We also match employment in the U.S. vs the rest of the OECD. We infer higher innovation rates in the U.S. given its higher GDP per worker relative to the rest of the OECD. We pin down the dispersion of product quality of the innovation draws by fitting the dispersion in revenue per worker across manufacturing firms in the U.S.

Given the estimated dispersion in product quality we ask: how much do ideas need to flow across countries to fit a trade elasticity of 5? We estimate that spillovers must occur on most traded goods in order to match a trade elasticity of 5. Using our estimate of spillovers, we analyze the model's transition dynamics and steady state response to changes in tariffs. Because ideas flow across countries due to trade, lower tariffs not just increase trade but also increase the long run growth rate. Even taking into account the transition, the gains from trade relative to autarky from the boost in idea flows are equivalent to a permanent 31% increase in consumption in the U.S. The rest of the OECD gains even more (75%) from idea flows because the U.S. is more innovative.

In the alternative version of the model wherein idea flows are independent of trade, cutting tariffs has no effect on the growth rate. In this alternative model, increasing the flow of ideas across countries increases the long run growth rate and *reduces* trade. More idea flows lead to a narrower distribution of relative product quality across countries, thereby lowering the (standard) comparative advantage gains from trade. As in the baseline model, the rest of the OECD benefits more from idea flows than the U.S. because the U.S. is more innovative.

We also entertain the effect of trade liberalization in a model where idea flows across countries are severely limited. In this version, when a product is imported learning is almost entirely from dormant domestic producers rather than from foreign sellers into the domestic market. As a result the total gains from trade

are much closer to the static gains. In this model, however, stochastic innovation causes country technologies to drift apart. The U.S. not only grows faster than the OECD, but comparative advantage becomes very strong across products. This implies a counterfactually low trade elasticity.

We further compare our model's predictions to exporting firm dynamics in the U.S. and a number of other countries (in particular, Chile, Colombia, China and Indonesia). We find that contracting firms are much more likely to lose exports than domestic sales, consistent with creative destruction from foreign innovation. In contrast, expanding firms are much more likely to gain domestic sales than to gain exports. This occurs in our model because domestic firms carrying out innovation upon foreign firms find it easier to sell the product domestically than to overcome tariff barriers to exporting the product.

We also document ample turnover of exports across product categories, just as Hanson, Lind and Muendler (2018) do. This is consistent with our model, in which comparative advantage is constantly shifting due to global creative destruction. We find that most of the adjustment in exports within industries occurs on the extensive margin: a country gains exports in a sector primarily when new exporters enter, and declining export sectors reflect mainly firm exit from exporting. At the same time, there are many firms that exit from foreign markets in sectors where net exports increase, and many firms that enter export markets in industries where net exports fall. These facts point to creative destruction rather than demand shocks facing all firms in an industry driving the turnover of exports across industries.

Our effort is most related to four recent papers. We build on Alvarez, Buera and Lucas (2017) in having domestic firms learn from exporters into the domestic market. They analyze how this learning affects the distribution of comparative advantage and the growth rate from diffusion of an unbounded distribution of knowledge. Whereas they analyze a setting with 30 trading partners, we analyze two trading countries. They abstract from innovation, however, whereas we feature growth from the frequency of innovation, both from at home and abroad.

Perla, Tonetti and Waugh (2021) study the impact of trade on exit, entry, domestic technology diffusion, and growth in a model of symmetric countries.¹ Like us, they find large dynamic gains from trade. They derive analytical steady state solutions in a model of many countries, whereas we simulate a two-country model calibrated to evidence on export reallocation across products and firms. Our focus is on innovation, idea flows across countries, and creative destruction, whereas their focus is on how trade interacts with domestic technology diffusion.

We follow Buera and Oberfield (2020) in studying international technology diffusion in a model with Bertrand competition. They endogenously obtain Frechet distributions of productivity within countries, allowing them to characterize multilateral trade flows as in Bernard, Eaton, Jensen and Kortum (2003). They stress that the dynamic effects of trade could be small or even negative depending on whether firms learn from domestic producers or from sellers into the domestic market. Our focus is more empirical and quantitative, as we show our model matches evidence on export dynamics at the firm and industry level. We argue that these facts are consistent with knowledge flows across countries.

Like us, Akcigit, Ates and Impullitti (2021) characterize the impact of tariffs on growth in a two-country model with technology spillovers. Theirs is a step-by-step innovation model, with escape-from-competition effects through which trade can induce more innovation. They analyze transition dynamics and optimal R&D subsidies. They emphasize the convergence of patenting rates in other advanced countries toward the patent rate in the U.S. in recent decades. In our model and empirics, in contrast, we focus on how trade affects export reallocation at the firm and industry levels.

The rest of the paper is organized as follows. Section 2 lays out the details of our baseline model. Section 3 calibrates the model. Section 4 shows how the model stacks up against non-targeted evidence on firm and industry export dynamics. In Section 5 we assess the gains from trade (and idea flows more generally) in our model. Section 6 concludes.

¹Sampson (2016) is an earlier effort in the same vein as Perla et al. (2021).

2. Baseline model

This section presents a model of growth driven by creative destruction, where innovation can come from domestic or foreign firms.

2.1 Static equilibrium

The static part of our model is similar to Bernard et al. (2003), or to Dornbusch, Fischer and Samuelson (1977) only with markup heterogeneity.

Utility of the home-country representative consumer is given by consumption of a continuum of varieties C_j with measure 1:

$$U = \int_0^1 \ln C_j dj.$$

This utility function implies that consumers spend the same on each variety.²

Output of each variety is the product of labor and the quality of the blueprint for the product. We denote A_j as the “best” blueprint for j among domestic firms. A_j^* is the corresponding best blueprint for j among foreign firms. If we order products so that the index j is decreasing in A_j/A_j^* , then products $j \in [0, x]$ are traded and produced at home, $j \in [x, x^*]$ are non-traded, and $j \in [x^*, 1]$ are traded and produced abroad. The cutoff products x and x^* are defined by

$$\frac{A_x}{\tau} = \omega A_x^*$$

$$A_{x^*} = \frac{\omega A_{x^*}^*}{\tau}$$

where ω denotes the relative wage (domestic relative to foreign) and $\tau \geq 1$ is the symmetric gross tariff rate. When $\tau = 1$, $x = x^*$ and all products are traded.

The owner of the best blueprint sets their quality-adjusted price to push their closest competitor out of the market (Bertrand competition), so the gross markup

²Utility of the foreign consumer is analogously given by $U^* = \int_0^1 \ln C_j^* dj$.

is the gap between the incumbent firm's marginal cost and the cost of its closest competitor — domestic or foreign.³ The relative wage is pinned down by balanced trade:

$$I^* \cdot x = I \cdot (1 - x^*) \quad (1)$$

where I and I^* denote nominal GDP at home and abroad, respectively. The left hand side of equation (1) is the home country's exports and the right hand side is the home country's imports. Nominal GDP in each country is given by

$$I = \frac{\bar{\mu} w L}{1 - \frac{1-\tau}{\tau} \cdot (1 - x^*)}$$

$$I^* = \frac{\bar{\mu}^* w^* L^*}{1 - \frac{1-\tau}{\tau} \cdot x}$$

where $\bar{\mu}$ denotes the average gross markup, w is the nominal wage, and L the labor supply at home.⁴

We can express the real (consumption) wage as a function of the distribution of the best blueprints, markups, the cutoffs, the relative wage, and the tariff rate. The real wages at home W and in the foreign country W^* are given by

$$\ln W = \int_0^{x^*} \ln \left(\frac{A_j}{\mu_j} \right) dj + \int_{x^*}^1 \ln \left(\frac{A_j^*}{\mu_j^*} \cdot \frac{\omega}{\tau} \right) dj$$

$$\ln W^* = \int_0^x \ln \left(\frac{A_j}{\mu_j} \cdot \frac{1}{\omega \tau} \right) dj + \int_x^1 \ln \left(\frac{A_j^*}{\mu_j^*} \right) dj.$$

The home country buys $j \in [x^*, 1]$ from the foreign country, so the domestic

³See Table A1 in the Appendix for a summary of the markups implied by this model.

⁴Variables with an asterisk denote the foreign country. The average price-cost markup in the home country is $\frac{1}{\bar{\mu}} \equiv \frac{\int_0^{x^*} \frac{1}{\mu_j} dj + \frac{1}{\tau} \cdot \int_0^x \frac{1}{\mu_j^f} dj}{x^* + x/\tau}$ where μ_j^f denotes the markup of domestic firms on their exported products. The expression for the foreign firms' average markup is analogous. The expression for nominal income comes from equating nominal income to the revenue of local firms plus tariff revenue: $I = \bar{\mu} w L + (\tau - 1) \frac{1}{\tau} (1 - x^*)$ and $I^* = \bar{\mu}^* w^* L^* + (\tau - 1) \frac{I^*}{\tau} \cdot x$.

real wage is increasing in the productivity of foreign firms on these products. Likewise, the foreign country purchases $j \in [0, x]$ from the home country, so the foreign real wage increases with domestic firm productivity on these products.

2.2 Innovation

We now introduce dynamics to the model. As in Klette and Kortum (2004), a firm is a portfolio of products. An entrant has one product while incumbent firms potentially produce many varieties. Innovation takes the form of creative destruction. We posit exogenous arrival rates of innovation for simplicity.⁵ Arrivals are proportional to the number of products owned by a firm; a firm with two products is twice as likely to creatively destroy another firm's variety compared to a firm with one product. We assume that innovation builds on the existing quality of the product. Such knowledge externalities are routinely built into quality ladder models from Grossman and Helpman (1991) and Aghion and Howitt (1992) onward. See Coe, Helpman and Hoffmaister (1997, 2009) and Ayerst, Ibrahim, MacKenzie and Rachapalli (2020) for evidence consistent with learning by importing.

We depart from Klette and Kortum (2004) by allowing a product made in one country to be creatively destroyed by a firm in another country. Table 1 summarizes the arrival rates of innovation. The probability a product is improved upon by an incumbent domestic firm is λ , while η is the probability the product is improved by an entering domestic firm. Analogously, λ^* is the probability the product will be improved by a foreign incumbent firm, and η^* is the probability a foreign entrant innovates on the best blueprint. In short, a given product can be improved upon by a domestic incumbent firm, a domestic entrant, a foreign

⁵In an earlier version of the paper we endogenized arrival rates as a function of research labor. The model's steady state properties are very similar, even in response to trade liberalization. See Hsieh, Klenow and Nath (2019). Cai, Li and Santacreu (2021) also analyze the endogenous response of innovation rates to trade liberalization, although their focus is in the endogenous reallocation of research resources across sectors in a multi-sector model

Table 1: Channels of innovation

	Domestic firms	Foreign firms
Innovation by incumbents	λ	λ^*
Innovation by entrants	η	η^*

Note: The average improvement in quality for each innovation is $\frac{1}{\theta-1}$.

incumbent firm, or a foreign entrant.⁶

The improvement in product quality yielded by an innovation follows a Pareto distribution with shape parameter θ and scale parameter equal to the existing quality level. The average percent improvement in quality is thus $1/(\theta - 1) > 0$.

2.3 Trade-embodied knowledge flows

We assume innovators improve upon the products *sold* in their market with probability κ and on the blueprints of local firms (including the last domestic firm to produce the product if it is currently imported) with probability $1 - \kappa$. When $\kappa = 1$ innovators build on the blueprint of imported products they draw. The other extreme is $\kappa = 0$ when innovators only build on the blueprints owned by local firms.⁷ So $\kappa = 1$ is the case with full idea spillovers from trade and $\kappa = 0$ is the case when ideas do not move across borders.

Table 2 summarizes the odds of creative destruction in the home market (top panel) and foreign market (bottom panel) by domestic firms (first column) and foreign firms (second column). The odds of creative destruction of products sold in the home market depend on whether the product is exported (row 1), non-traded (row 2), or imported (row 3). The first row shows the arrival rate of ideas in the domestic market for an exported product. The probability such a product

⁶These are all *unconditional* probabilities of innovation. The conditional probabilities avoid duplication of arrivals on the same product in a given year.

⁷When $\kappa = 0$ growth rates diverge across countries.

is improved upon by another domestic firm is $\lambda + \eta$. A domestic innovator will always replace the incumbent firm in this market.

Table 2: Probability of creative destruction

	Domestic firm	Foreign firm
<u>Home Market</u>		
Exported by home	$\lambda + \eta$	$(\lambda^* + \eta^*) \left[\kappa \left(\frac{\omega}{\tau} \right)_m^\theta + (1 - \kappa) \left(\frac{\omega A_j^*}{\tau A_j} \right)_m^\theta \right]$
Non-traded	$\lambda + \eta$	$(\lambda^* + \eta^*) \left(\frac{\omega A_j^*}{\tau A_j} \right)_m^\theta$
Imported by home	$(\lambda + \eta) \left[\kappa \left(\frac{\tau}{\omega} \right)_m^\theta + (1 - \kappa) \left(\frac{\tau A_j}{\omega A_j^*} \right)_m^\theta \right]$	$\lambda^* + \eta^*$
<u>Foreign Market</u>		
Exported by home	$\lambda + \eta$	$(\lambda^* + \eta^*) \left[\kappa (\omega \tau)_m^\theta + (1 - \kappa) \left(\frac{\omega \tau A_j^*}{A_j} \right)_m^\theta \right]$
Non-traded	$(\lambda + \eta) \left(\frac{A_j}{\omega \tau A_j^*} \right)_m^\theta$	$\lambda^* + \eta^*$
Imported by home	$(\lambda + \eta) \left[\kappa \left(\frac{1}{\tau \omega} \right)_m^\theta + (1 - \kappa) \left(\frac{A_j}{\tau \omega A_j^*} \right)_m^\theta \right]$	$\lambda^* + \eta^*$

Note: $x_m^\theta \equiv \min [x^\theta, 1]$; κ is the probability the innovator improves upon the products sold in its market; $1 - \kappa$ is the probability the innovator improves upon the blueprint of local firms; $\lambda + \eta$ is the arrival rate of innovation from domestic firms; $\lambda^* + \eta^*$ is the arrival rate of innovation from foreign firms.

A foreign firm improves upon the domestic firm's blueprint of the same exported product with probability $\kappa (\lambda^* + \eta^*)$ and improves upon the blueprint of the current (or previous) foreign firm with probability $(1 - \kappa) (\lambda^* + \eta^*)$. Even when the quality of the foreign innovator of the exported product exceeds that of the incumbent domestic firm, however, it will not necessarily replace the domestic incumbent. Conditional on having a higher quality, the probability the foreign innovator will replace the domestic incumbent depends on the relative wage and the trade cost between the two countries. Higher domestic wages increase the probability a foreign innovator will be competitive enough to replace the domestic incumbent in the domestic market. Higher tariffs make the foreign

innovator less competitive compared to the domestic incumbent.

The expected growth in the domestic real consumption wage is the product of the rate of creative destruction from domestic and foreign firms and the increases in product quality (conditional on the product being replaced) associated with innovation on three types of products (exported, non-traded, and imported) sold in the domestic market:

$$\begin{aligned}
g = & (\lambda + \eta) x^* \phi_{[0, x^*]} \\
& + (\lambda + \eta) (1 - x^*) \left[\kappa \left(\frac{\tau}{\omega} \right)_m^\theta + (1 - \kappa) \left(\frac{\tau}{\omega} \int_{x^*}^1 (A_j / A_j^*) dj \right)_m^\theta \right] \phi_{[x^*, 1]} \\
& + (\lambda^* + \eta^*) x \left[\kappa \left(\frac{\omega}{\tau} \right)_m^\theta + (1 - \kappa) \left(\frac{\omega}{\tau} \int_0^x (A_j^* / A_j) dj \right)_m^\theta \right] \phi_{[0, x]}^* \\
& + (\lambda^* + \eta^*) \left[(x^* - x) \left(\frac{\omega}{\tau} \int_x^{x^*} (A_j^* / A_j) dj \right)_m^\theta \phi_{[x, x^*]}^* + (1 - x^*) \phi_{[x^*, 1]}^* \right]
\end{aligned} \tag{2}$$

where ϕ and ϕ^* denote the improvement in product quality from innovation by domestic and foreign firms, respectively, and the subscripts on ϕ denote the set of products over which the quality improvement applies.⁸ The expected growth rate of the foreign real consumption wage is similarly the product of the arrival rates in rows 4-6 of Table 2 and the corresponding improvements in quality.

The first line in equation (2) is the contribution of innovation by domestic firms on the varieties produced by domestic firms (exports and non-traded products); the second is the contribution of domestic innovation on imports; the third line is the contribution of foreign innovation on the home country's exports; and the fourth line is the contribution of foreign innovation on non-traded varieties and the home country's imports.

Note the growth rates in the two countries depend on the arrival rates of innovation in the two countries, the step size (θ), the spillover parameter (κ),

⁸Each ϕ depends on the step size, θ , and the frequency of knowledge spillovers, κ . Table A2 in the Appendix shows the average improvement in quality from creative destruction on each type of product.

the relative wage (ω), and the share of each products exported by each country (x and $1 - x^*$). The relative wage and the share of products made by each country are pinned down by balanced trade and the distribution of relative technologies A_j/A_j^* , where the latter is endogenous to innovation.

To illustrate how quality advances as each country builds on the innovations of the other country, it is useful to consider the case of completely free trade ($\tau = 1$) and full cross-border spillovers ($\kappa = 1$). In this case, all products are traded so the probability a domestic firm creatively destroys another firm is given by:

$$\text{Domestic creative destruction rate} = (\lambda + \eta) \cdot x^* + (\lambda + \eta) (1/\omega)_m^\theta \cdot (1 - x^*).$$

The first term is the probability a domestic firm replaces a product made by another domestic firm, and the second term is the probability a domestic firm replaces a variety produced by a foreign firm. The corresponding rate of creative destruction by a foreign firm under free trade ($\tau = 1$) is:

$$\text{Foreign creative destruction rate} = (\lambda^* + \eta^*) \cdot (1 - x^*) + (\lambda^* + \eta^*) (\omega)_m^\theta \cdot x^*.$$

Ceteris paribus, higher ω (home wage relative to foreign wage) lowers the rate of creative destruction of domestic firms and raises that of foreign firms. In steady state, the equilibrium relative wage equates the rate of creative destruction by domestic firms to that of foreign firms. So, if domestic firms are more innovative, domestic wages are higher but the creative destruction rate of domestic firms is the same as for foreign firms in equilibrium.

It is also helpful to contrast autarky and free trade when the two countries are symmetric in size and in their innovation arrival rates in the full idea spillover case ($\kappa = 1$). In this special case the relative wage $\omega = 1$ and the growth expressions become simply:

$$\text{Autarky growth rate} = (\lambda + \eta) \frac{1}{\theta - 1}$$

$$\text{Frictionless growth rate} = 2 \cdot (\lambda + \eta) \frac{1}{\theta - 1}.$$

In autarky each country benefits only from domestic arrivals. With frictionless trade, each country benefits from both domestic and foreign arrivals. The doubling of growth under free trade compared to autarky underscores the scale effect generating dynamic gains from trade in this model.

2.4 Knowledge spillovers and effect of changes in trade costs


Figure 1 illustrates the importance of the spillover parameter κ for the effect of changes in trade costs in the model. The figure shows the growth rate, the trade share, and the local trade elasticity across steady-states with different trade costs for models where $\kappa = 1$, $\kappa = 0.5$, and $\kappa = 0.01$, respectively.⁹ The plot does not include the polar case of zero spillovers ($\kappa = 0$), because in that case there is no steady state in terms of the trade shares and the local trade elasticity. The two country's TFP paths diverge in the absence of spillovers, as their long run TFP growth rates differ when $\kappa = 0$.

The left panel of Figure 1 shows the effect of trade costs on the common long run growth rate of the two economies. The higher is κ , the more sensitive is growth to tariffs. When fewer goods are traded, countries are less frequently building on each other's innovations and more frequently building on their own innovations. Changes in trade costs have no discernible effect on the growth rate when spillovers are minimal ($\kappa = 0.01$). When ideas barely flow across countries from the exchange of goods, limiting trade has little effect on the growth rate.

The middle panel in Figure 1 shows that higher tariffs lower the trade share for all values of κ . But the trade share is more sensitive to tariffs for higher values of κ . When κ is high, knowledge flows keep country technologies tethered together, thereby weakening comparative advantage and making trade more sen-

⁹The numbers in the figure are for illustrative purposes only. We discuss the precise calibration of the model in detail in a later section.

Figure 1: Effect of tariffs with varying levels of knowledge spillovers



Note: Simulations are run for symmetric countries of equal size, with $\theta = 7$ and total arrival rates ($\lambda + \eta$ for the home country and $\lambda^* + \eta^*$ for the foreign country) set to 0.12 in each country. All parameters except τ are held constant across the range of counterfactuals displayed.

sitive to tariffs.

The right panel in Figure 1 shows how the local trade elasticity responds to the trade cost. By the local trade elasticity we mean the change in the log of the import share from a local change (10 percentage point reduction) change in the log trade cost. It is local in that it is evaluated in a given year starting from the initial steady state distribution of relative quality across the two countries that exists before the tariff change.¹⁰ When spillovers are severely limited ($\kappa = 0.01$), relative qualities between the two countries drift apart because each country innovates largely on its own products. As a result, comparative advantage is strong and the trade elasticity is low for all tariff levels. When ideas flow more easily across countries with trade, however, higher tariffs hinder the flow of ideas and strengthen the degree of comparative advantage. Thus higher tariffs impede trade and idea flows and lower the trade elasticity when $\kappa = 0.5$ or $\kappa = 1$.

¹⁰Formally, we calculate the local trade elasticity as $\frac{\log\left(\frac{\text{imports}}{\text{domestic sales}}(\tau)\right) - \log\left(\frac{\text{imports}}{\text{domestic sales}}(\tau - 0.1)\right)}{\log(\tau - 0.1) - \log(\tau)}$.


Figure 2 shows why the trade elasticity varies with the tariff rate when κ is far above zero. The left panel plots the distribution across products of quality in the home relative to the foreign country (A_j/A_j^*) in a steady state with high tariffs ($\tau = 4$) versus in a steady state with low tariffs ($\tau = 1.5$) for the full spillover case ($\kappa = 1$). When ideas are embodied in trade, lower tariffs narrow the dispersion of relative quality, as ideas flow more quickly across countries with more trade. When technologies are more similar across countries, the response of trade flows to changes in tariffs (the trade elasticity) is correspondingly higher. See Alvarez et al. (2017) for a similar effect of learning by importing on the distribution of comparative advantage.

The right panel of Figure 2 shows the distribution of relative product quality in a high tariff ($\tau = 4$) steady state versus a low tariff ($\tau = 1.5$) steady state for the limited spillover case ($\kappa = 0.01$). In this case, lower tariffs have no noticeable effect on the dispersion of relative quality, as more trade does not lead to more idea flows across countries. Therefore, the strength of comparative advantage is virtually unaffected by trade costs, as is the response of trade flows to changes in trade costs (the trade elasticity).

2.5 Disembodied idea flows

In our baseline model, foreign innovators learn about domestic technologies through trade. We now consider a model in which the flow of ideas across borders is not related to trade. Consider the products $j \in [0, 1]$ sorted by the highest to lowest ratio of domestic productivity to foreign productivity, A_j/A_j^* . Suppose foreign innovators draw with probability z on a random domestic product from $j \in [0, z]$ and with probability $1 - z$ on a random foreign product from $j \in [z, 1]$. And suppose domestic innovators innovate with probability z^* on a random foreign product from $j \in [z^*, 1]$ and with probability $1 - z^*$ on a random domestic product from $j \in [0, z^*]$. Spillovers from the domestic to foreign innovators are thus increasing in z , and spillovers from foreign blueprints to domestic innova-

Figure 2: Effect of tariffs on relative quality dispersion



Note: As in Figure 1, simulations are run for symmetric countries of equal size, with $\theta = 7$ and total arrival rates ($\lambda + \eta$ for the home and $\lambda^* + \eta^*$ for the foreign country) set to 0.12 in each country. All parameters except τ are held constant across the counterfactuals displayed.


tors is increasing in z^* . We call this a “disembodied spillover” model since the knowledge spillovers are not related to trade.

In this disembodied spillover model, creative destruction from foreign innovators takes place when the foreign innovators target a domestic variety. Likewise a foreign variety is creatively destroyed when a domestic innovator targets a foreign variety. Moreover, the steady-state of the disembodied spillovers model is equivalent to that of our baseline model where idea flows are fully embodied in trade ($\kappa = 1$) if $z = x$ and $z^* = x^*$, where x and x^* are the fraction of products exported by the domestic and the foreign country, respectively.

The two models differ in that trade is essential to spillovers in our baseline model and unrelated to idea flows in this disembodied spillover model. In the former, the key parameters that determine the extent of spillovers are κ and the trade cost τ . In the latter, the key parameters that govern spillovers are the spillover thresholds z and z^* . Figure 3 illustrates the effect of the spillover

threshold by showing the growth rate, trade share, and trade elasticity in steady states with different values of the spillover threshold for foreign innovators z . Remember the spillover of domestic ideas to foreign innovators increases as z rises. The growth rate rises as ideas flow more quickly from the domestic to foreign innovators with a higher z . On the other hand, the trade share falls and the trade elasticity rises as foreign qualities hug the domestic ones more closely in response to higher z , blunting the degree of comparative advantage.

Figure 3: Effect of spillover threshold with disembodied knowledge flows



Note: Simulations are run for symmetric countries of equal size, with $\theta = 7$ and total arrival rates ($\lambda + \eta$ for the home country and $\lambda^* + \eta^*$ for the foreign country) set to 0.12 in each country. All parameters except z are held constant across the range of counterfactuals displayed.

Notice that the growth rate is negatively correlated with the trade share in the disembodied spillover model, whereas the correlation was positive in our baseline model where spillovers were embodied in trade flows. In both models the growth rate rises and technology differences between the two countries narrow as ideas flow more frequently across countries. The narrowing of technology differences in the model comes from more trade in the trade-embodied spillover model. In contrast, trade has no effect on spillovers in the disembodied model.

As a result, the dispersion of relative quality is not a function of trade costs in the disembodied spillover model.

3. Model Calibration

Our baseline model involves eight parameters: the shape θ of the Pareto distribution of innovation draws; two innovation rates (for incumbents λ and entrants η) in each country; the tariff rate τ ; the spillover parameter κ ; and relative employment in the home vs. foreign country. We infer the value of these parameters from the seven data moments listed in Table 3. We do not separately identify the arrival rate of innovations by foreign entrants vs. foreign incumbents, but rather assume this breaks down in the same way the U.S. ratio breaks down. As mentioned, the U.S. is “home” and the rest of the OECD is “foreign.”

Table 3: Data moments used for calibration

Data Moment	Source	Value	Model Fit
Standard deviation log VA/worker	U.S. Census of Manufacturing	0.108	0.104
TFP growth rate	BLS data for U.S. manufacturing	3.01%	3.01%
Value added per worker home/foreign	KLEMS for U.S. and OECD mfg.	1.29	1.29
Employment share of entrants ($\text{age} \leq 5$)	U.S. Census of Manufacturing	14.4%	14.4%
Export share of revenues (home)	U.S. Census of Manufacturing	10.2%	10.2%
Trade elasticity	Head and Mayer (2014)	5	4.82
Employment home/foreign	KLEMS for U.S. and OECD mfg.	0.389	0.389

We back out θ from the standard deviation of the log of labor productivity (revenue per worker) across firms. The higher is θ , the smaller the variance in the innovation step size and the smaller the dispersion in labor productivity across firms. In the U.S. manufacturing data the standard deviation of the log of value-added per worker across firms is 0.108.

For a given θ and relative employment L/L^* , the innovation arrival rates and the tariff rate (τ) jointly determine the growth rate, the trade share, and the relative wage. We target a growth rate of 3%, relative employment (U.S./OECD) of 0.389, a U.S. trade share of 10%, and a relative wage (U.S./OECD) of 1.29. We use the employment share of new firms in U.S. manufacturing (14.4% in the data) to pin down the ratio of innovation by entrants vs. incumbents, which we assume is the same in the two countries.

Finally, we back out the crucial spillover parameter κ by targeting a trade elasticity of 5, in line with estimates in Head and Mayer (2014). Figure 1 above showed that the trade elasticity increases with the degree of spillovers κ because the technology gaps between the two countries narrow when ideas flow more freely across countries. We ask: given the dispersion of the quality step size necessary to fit the dispersion of labor productivity in the U.S. data, how much do ideas have to move across borders such that the dispersion of technology gaps across the two countries generates a trade elasticity of 5?

To characterize the model quantitatively, we take a discrete number of products and simulate the random arrival of innovations on each variety.¹¹ Draws are randomly assigned to an existing incumbent or a new entrant. The relative wage is found that balances trade between the two countries in each year. We simulate for several hundred years until the economy settles down to a steady-state, at which point we calculate moments. We utilize a simulated annealing procedure to search for the parameter values that match the data moments.¹²

The resulting calibrated parameter values are shown in Table 4. The model estimates of the targeted data moments are shown in the last column of Table 3. The U.S. combined innovation rate for incumbents and entrants is about $\lambda + \eta = 15\%$, and the OECD combined innovation rate is roughly $\lambda^* + \eta^* = 13\%$. The U.S. innovation rate has to be higher to explain the 29% higher real wage (real

¹¹As explained further in Section 4, we run the simulation with approximately 21,000 products to match the relative volume of exports across the 264 U.S. manufacturing industries. Note that the simulated moments in Table 3 are not affected by the number of products in the simulation.

¹²Appendix B provides more details on the solution procedure.

Table 4: Model parameter estimates

Variable	Description	Value
θ	Shape parameter of innovation draws	7.11
λ	Home innovation rate from incumbents	12.1%
η	Home innovation rate from entrants	2.6%
$\lambda^* + \eta^*$	Foreign innovation rate from incumbents + entrants	13.1%
τ	Gross tariff rate	1.54
κ	Proportion trade-embodied spillovers	0.938

value added per worker) in the U.S. than in the rest of the OECD. The value of the shape parameter for innovation draws, θ , that most closely matches the dispersion of labor productivity across firms is around 7.

The value of the spillover parameter, κ , that comes closest to hitting the trade elasticity target is 0.94. This implies that our model needs to incorporate spillovers on most traded goods to match the targeted trade elasticity. It is worth noting that we have omitted any spillover on nontraded goods that could also tether productivities together and reduce the strength of comparative advantage across countries. But the key point remains: given the stochastically evolving product-level technologies in our model, idea flows are essential to keep country technologies from drifting so far apart that comparative advantage becomes too strong and the trade elasticity too low.¹³

Finally, conditional on the innovation rates, the shape parameter, the spillover parameter, and the relative size of the two economies, fitting the U.S. trade share pins down a tariff rate of about 54%.¹⁴

¹³Static models in the spirit of Eaton and Kortum (2002) and Melitz (2003) can of course simply impose a trade elasticity of 5, as the trade elasticity is governed by an exogenous Pareto or Frechet shape parameter in these models. In our model, in contrast, the distribution of productivity evolves dynamically in response to innovation draws. These stochastic innovation draws are a force for divergence that is absent from the static models.

¹⁴Eaton and Kortum (2002) and others infer high trade costs to explain bilateral trade flows.

4. Indirect evidence of idea flows


In this section, we contrast the implications of our calibrated model with strong idea flows (κ close to 1) to a model wherein idea flows are severely limited (κ close to zero). First, we show that a model with more restricted idea flows cannot come close to generating a trade elasticity of 5. Second, we use data on manufacturing plants in the U.S. and several other countries to establish that exports decline by *more* than domestic sales among contracting exporters, and grow by *less* than domestic sales among expanding exporters. Third, we document that the same manufacturing data exhibits a great deal of excess export reallocation across manufacturing industries. We show that the calibration with strong idea flows can replicate these two features of the data, whereas those calibrations with much more limited flows cannot.

As mentioned, the best fit of our model generates a trade elasticity of 4.82 with a spillover parameter κ of 0.94. The trade elasticity depends on the strength of comparative advantage, which is itself a function of the extent to which ideas flow across borders. Figure 4 shows how limiting the extent of idea spillovers lowers the trade elasticity. Here, instead of estimating κ , we *assume* a lower value of κ and re-estimate all the parameters of the model to fit the other target moments in Table 3, except for the trade elasticity and trade share. Figure 4 shows that the trade elasticity falls as we impose a lower spillover parameter κ . The dispersion of relative quality and the strength of comparative advantage increases as idea spillovers become more limited,. At the extreme where κ is almost zero, the trade elasticity is almost one.

A second fact that helps discriminate between a model with strong spillovers and one with only weak spillovers is the growth rate of exports versus domestic sales of exporting firms. Table 5 presents these growth rates over five years from the firm-level manufacturing censuses of the United States, China, Chile, Indonesia, and Colombia.¹⁵ The sample is restricted to exporting firms at the

¹⁵The data are the firm-level manufacturing censuses from 1987 to 2017 for the U.S., 1995 to

Figure 4: Effect of spillover parameter on trade elasticity



Note: Figure shows the trade elasticity in the steady-state for a given spillover parameter κ . We fix κ and calculate the model parameters to match the data moments in Table 3 except for the trade elasticity and trade share.

beginning of each five-year period, and we normalize the growth rate of total sales (domestic sales plus exports) to zero on average by subtracting the growth rate of aggregate sales of exporting firms in each five-year period. We further group firms into ones where total (normalized) sales increase over five years and ones where total sales decrease over five years. For each sample, we then calculate the growth rate of a firm's exports between year t and $t + 5$ as the ratio of the change in the firm's exports over the five years to average exports of the firm at the beginning and end of the five year period:

$$2 \cdot \frac{\text{export}_{i,t+5} - \text{export}_{i,t}}{\text{export}_{i,t+5} + \text{export}_{i,t}}$$

where $\text{export}_{i,t}$ denotes firm i 's exports at time t . The growth rate of a firm's domestic sales is calculated similarly. The growth rate of exports (domestic sales)

2007 for Chile, 1998 to 2007 for China, 1990 to 1999 for Indonesia, and 1981 to 1989 for Colombia.

of a firm that exits exporting (domestic sales) is thus -2 . Because the sample consists of firms exporting at the *beginning* of the period, entry into exporting in each five year period is not part of the calculations.

Table 5: Growth rate of exports and domestic revenues of exporting firms

	Contracting Exporters		Expanding Exporters	
	Exports	Domestic Sales	Exports	Domestic Sales
United States	-0.840	-0.408	-0.422	0.383
China	-1.469	-1.227	-0.003	0.449
Chile	-1.255	-1.083	0.041	0.291
Indonesia	-1.608	-0.684	-0.704	0.905
Colombia	-1.230	-1.001	0.007	0.290

Notes: Sample are firms with positive exports at the beginning of each five year period. Table show average growth rates of exports and domestic sales of exporters over five year periods. Growth rate measured as change in exports or domestic sales of the firm over a five-year period divided by the average of exports or domestic sales of the firm at the beginning and end of each five year period. Growth rate is -2 for firms that exit from exporting or domestic sales. Contracting firms defined as firms where total sales decrease, and expanding firms are firms where total sales increase over each five year period. The growth rate of total sales (across all firms) is normalized to zero over each five year period. Growth rates calculated over five-year periods from 1987 to 2017 for the U.S., 1995 to 2007 for Chile, 1998 to 2007 for China, 1990 to 1999 for Indonesia, and 1981 to 1989 for Colombia from the micro-data of the manufacturing censuses of these countries.

Table 5 shows the average growth rate of domestic sales and exports of contracting firms in columns 1-2 and expanding firms in columns 3-4.¹⁶ The key message we take from Table 5 is the asymmetry between contracting vs. expanding firms in the growth rate of exports vs. domestic sales. In all five countries, export sales fall among firms whose total sales decline over the five year period, and the average decline in exports is larger than the average decline in domestic sales. The opposite pattern holds among exporters whose total sales increase over the

¹⁶Table 5 shows the average growth rates over five-year periods from 1987 to 2017 for the U.S., 1995 to 2007 for Chile, 1998 to 2007 for China, 1990 to 1999 for Indonesia, and 1981 to 1989 for Colombia.

five year period. Among expanding firms, domestic sales increase as one might expect, but export growth of the same firms is either negative (in China, the U.S., and Indonesia) or essentially zero (in Chile and Colombia).

Table 6 shows what our baseline model with strong idea flows predicts for the ratio of the growth rate of exports to domestic sales for expanding vs. contracting firms.¹⁷ The ratio of the growth rate of exports to domestic sales for contracting firms is 1.92 in the model calibrated to fit the U.S. moments in Table 3 compared to 2.06 in the U.S. data. The spillover of ideas is crucial to this prediction. The spillover of ideas across borders means that foreign firms can innovate on the products of a domestic firm. When this happens, the domestic firm loses its product in the foreign market, but not necessarily in the domestic market because it is protected in the domestic market by the trade cost. As a consequence, when a firm shrinks in the model with spillovers, it is more likely that it loses the export market than the domestic market. And thus exports tend to fall more than domestic sales.

The model can also replicate the fact that, conditional on *expanding*, the expected growth rate of export sales is lower compared to domestic sales, and even negative in the U.S. The ratio of the growth rate of exports to domestic sales is -0.58 among expanding exporters in the model with ideas spillovers (Table 6, row 2); the same moment in the U.S. data is -1.1. Again, this prediction comes from creative destruction across borders due to the spillover of ideas. When a domestic firm innovates upon a foreign firm's product, it is more likely to replace the foreign firm in the domestic market than in the foreign market. The reason is because the tariff helps the domestic innovator — the foreign incumbent has to pay the tariff while the domestic innovator does not.

The last column in Table 6 shows that a model with severely limited idea

¹⁷In the model, firms are a collection of products. As in Klette and Kortum (2004), the distribution of the number of products per firm in our model is determined endogenously by the innovation rate of incumbents vs. entrants. In our baseline model, the number of products for the average firm is 2.4 and the standard deviation of the number of products per firm is 2.5. Firms expand and contract in the model as they gain and lose products.

Table 6: Growth rate of exports relative to domestic sales in the U.S.

		Model	
	U.S. Data	$\kappa = 0.94$	$\kappa = 0.01$
Contracting exporters	2.06	1.92	1.13
Expanding exporters	-1.10	-0.58	0.89

Notes: Table shows the ratio of the growth rate of exports to domestic sales among contracting exporters (row 1) and expanding exporters (row 2) for the U.S. in the data (column 1) and in the model calibrated to fit the U.S. moments (columns 2-3). Contracting exporters are ones where total sales fall over five years, and expanding exporters are ones where total sales increase over five years. The parameters of the model in column 2 are estimated to fit the data moments in Table 3. The limited spillover version of the model in column 3 assumes $\kappa = 0.01$ and estimates the model parameters to fit the same data moments, except for the trade elasticity.

spillovers cannot generate the two facts. Here we show the prediction of a model where we impose a spillover parameter of $\kappa = 0.01$ and re-estimate the other parameters to fit the same target moments in Table 3, except for the trade elasticity and trade share. In the model with limited idea spillovers, the growth rate of domestic sales is similar to the growth rate of exports. This is true for expanding as well as for contracting firms. The reason is the absence of cross-border creative destruction in the limited spillover model. When spillovers are limited, foreign firms do not improve upon a domestic firm's blueprint, and domestic firms do not improve upon a foreign firm's product. So a domestic firm shrinks primarily when another domestic firm improves upon its product, and it expands when it improves upon the product of another domestic firm. So when a firm shrinks, it loses its product both in the foreign and in the domestic market, and when it expands it gains the product in the domestic and in the foreign market. This prediction of the model with limited cross-border creative destruction is clearly at odds with the evidence in Table 5.

Our third fact is about excess export reallocation. We aggregate the firm-level data up to the industry level to calculate the aggregate rate of export reallocation

across industries in a manner akin to how Davis, Haltiwanger and Schuh (1996) calculate job reallocation rates across firms. We first net out aggregate export growth by scaling each firm's exports in year $t + 5$ by the gross growth rate of aggregate exports in the country from year t to $t + 5$. This normalization nets out aggregate changes in nominal exports, both due to real growth and changing export prices.

We calculate *excess* export reallocation across industries by summing up the increases in exports in those industries showing an increase in exports over a given five year period. We get the excess export reallocation *rate* by dividing this by aggregate exports at the beginning of the period.¹⁸

$$\frac{\sum_{j \in +} \left(\text{export}_{j,t+5} - \text{export}_{j,t} \right)}{\text{aggregate export}_t}$$

where $j \in +$ denotes the set of industries with increasing exports and $\text{export}_{j,t}$ denotes total exports of industry j at time t . We present the excess export creation rate across industries in Table 7 (column 1). It is sizable in all five countries, ranging from 16.7% for the U.S. to 41% for Indonesia. We are not the first to emphasize such dynamically evolving export patterns in the data — see Hanson, Lind and Muendler (2018).

To compare the model to the data in terms of export reallocation across industries, we need to take a stand on the number of products in each industry. We assume the smallest industry has one product and that the number of products in an industry increases at the exponential rate ϵ as one goes from smallest to the largest industries in terms of exports. We then choose the total number of products in the model and ϵ to match two numbers in the U.S. firm data: the number of industries and the ratio of the 75th to 25th percentile industry exports.

¹⁸We can also calculate the excess export destruction rate by adding the *decrease* in exports in all industries showing a decline in exports over the five year period and similarly dividing by aggregate exports at the beginning of the period. Given our normalization that the growth of total exports over the five year period is zero, however, the export destruction rate is the same as the export creation rate.

Table 7: Excess export reallocation between and within industries

	Between industries	Within <i>contracting</i> industries		Within <i>expanding</i> industries	
		Creation	Destruction	Creation	Destruction
United States	0.167	0.182	-0.349	0.340	-0.173
China	0.260	0.268	-0.529	0.510	-0.251
Chile	0.222	0.135	-0.357	0.443	-0.221
Indonesia	0.410	0.196	-0.606	0.595	-0.187
Colombia	0.295	0.083	-0.378	0.302	-0.158

Notes: Table shows export creation and destruction rates between industries (column 1) and within industries (columns 2-5) over five year periods. Between industry export creation is sum of change in exports of expanding industries over a five year period divided by total exports (across all industries) at the beginning of each period. Total growth rate of exports is normalized to zero for each five year period so the export creation rate between industries is equal to the export destruction rate between industries. Within industry export creation is sum of the change in exports of *firms* with expanding exports and exports of new exporting firms within each five year period for firms in contracting or expanding industries, all divided by total exports (across all industries) at the beginning of the five year period. Within industry export destruction rate is the sum of the change of exports of firms with decreased exports in each industry plus exports at the beginning of the period of firms that stop exporting by the end of the period for firms in contracting or expanding industries, all divided by total exports (across all industries) at the beginning of the period. Export creation and destruction rates calculated over five year periods from 1987 to 2017 for the U.S., 1995 to 2007 for Chile, 1998 to 2007 for China, 1990 to 1999 for Indonesia, and 1981 to 1989 for Colombia.

There are 264 industries consistently defined industries in the U.S. manufacturing census between 1987 and 2017. The 75/25 ratio of these industries is a factor of 27.6 for U.S. exports. We fit these two data moments with 21,000 products and $\epsilon = 2.35$.

Table 8 shows that the model calibrated to fit the U.S. moments in Table 3 predicts an export reallocation rate across industries of 10.9% in the U.S. Since our estimate of κ implies that spillovers occurs over most traded goods, industries expand their exports when firms creatively destroy the products of firms located in other countries, and industries shrink their exports when their exports are innovated upon and replaced by foreign firms. Again, creative destruction that takes place across borders is crucial to this prediction.

The excess export reallocation rate in the model that assumes limited idea flows is much lower, at 3.1%. When idea flows are limited, firms innovate upon products of other firms in the same country. If the product happens to be exported, the innovating firm gains an export but the incumbent firm loses an export. There is no net gain in exports at the industry level because higher exports of the innovating firm are offset by the export loss of the incumbent firm. There is a modest amount of excess export reallocation in the model with severely limited idea spillovers when an innovator improves upon a non-traded product and the quality improvement is large enough such that the previously non-traded product becomes exported.

Table 8: Excess export reallocation across industries in the U.S.

	U.S. Data	Model	
		$\kappa = 0.94$	$\kappa = 0.01$
Export creation/destruction rate	0.149	0.109	0.031

Notes: Table shows the export creation and destruction rate at the industry level for the U.S. over a five year period, where the growth rate of total exports (across all industries) is normalized to zero. The model parameters in column 2 are estimated to fit the U.S. data moments in Table 3. The limited spillover version of the model in column 3 assumes $\kappa = 0.01$ and estimates the model parameters to fit the same data moments, except for the trade elasticity.

To restate, the excess export churn that we observe across industries can be explained by a model with cross-border creative destruction, whereas a model with limited idea flows is not consistent with this basic fact. An alternative explanation for this excess export churn, however, is that there is volatility of export demand. Industries with increasing exports could be hit by positive demand shocks, and industries with decreasing exports could be hit by negative demand shocks. That is, export reallocation could reflect demand shocks rather than innovation and global creative destruction.

We can shed some light on the demand shock explanation by measuring ex-

port churn at the *firm* level instead of at the industry level. If export churn at the industry level is driven by industry export demand shocks, then we would expect to see exports increase among all exporting firms in industries where exports rise. Likewise, exports should fall in all firms in the declining export sectors.

Table 7 provides export creation and destruction rates at the *firm* level for contracting export sectors (columns 2-3) and expanding export sectors (columns 4-5). We calculate the firm-level export creation rate in contracting industries by adding up the increase in exports in all firms with increasing exports (including entrants) in the industries that are contracting over the five year period. We then divide this number by aggregate exports at the beginning of the period. The export destruction rate in contracting industries is the sum of the decline in exports among all firms with declining exports (including firms that exit exporting) in contracting industries over the five year period, divided by total exports at the beginning of the period. The export creation and destruction rates for industries with expanding exports are calculated similarly. By construction, the sum of the firm-level export creation and destruction rates in the contracting industries (sum of columns 2 and 3 in Table 7) is equal to the export destruction rate at the *industry* level. Likewise, the sum of the firm-level creation and destruction rates in expanding industries (sum of columns 4 and 5 in Table 7) is equal to the export creation rate at the industry level.

The basic pattern is that export creation and destruction rates at the firm level are much higher than at the industry level. In industries where net exports fall, there are still many firms where exports increase. Likewise, many firms lose their exports in industries where total exports are rising. This basic fact is consistent with our model where the within-industry export churn is driven by creative destruction across firms. It is not consistent with the view that the excess export churn at the industry level is driven by industry-wide export demand shocks, as such shocks will not generate the excess export churn within industries.

Finally, Table 9 shows the importance of entry and exit in the turnover of products across industries. Column 1 in Table 9 shows the share of gross export

Table 9: Extensive margin share of export creation and destruction

	Extensive Margin Share of	
	Export Destruction in Contracting Industries	Export Creation in Expanding Industries
United States	57.5%	57.8%
China	69.4%	90.6%
Chile	47.1%	73.1%
Indonesia	72.9%	84.0%
Colombia	35.6%	67.1%

Notes: Table shows share of gross export destruction within contracting industries due to exit from exporting (column 1) and share of gross export creation within expanding industries from entry into exporting (column 2). Gross export creation in expanding industries is sum of the change in exports of *firms* with expanding exports and exports of new exporting firms within each five year period for firms in expanding industries; gross export destruction in contracting industries is the sum of the change of exports of firms with decreased exports in each industry plus exports at the beginning of the period of firms that stop exporting by the end of the period for firms in contracting. Extensive margin shares are calculated over five year periods from 1987 to 2017 for the U.S., 1995 to 2007 for Chile, 1998 to 2007 for China, 1990 to 1999 for Indonesia, and 1981 to 1989 for Colombia.

destruction in contracting industries due to firms that exit from foreign markets, and column 2 shows the share of gross export creation in expanding industries from new exporters. The evidence shows that most of the expansion of exports within expanding industries occurs through new exporters. Likewise, most of the decline of exports within declining industries comes from firms that exit from foreign markets. This evidence also points to creative destruction rather than demand shocks that face incumbents and exiters/entrants equally as the force behind the gross creation and destruction of exports.

5. Gains from trade and idea flows

In this section, we calculate the welfare gains from trade and idea flows in the model with the baseline parameter values in Table 4 above. In this model, in which idea flows are embodied in trade, the gross tariff rate τ is the key parameter that determines the extent to which ideas flow between countries. A decrease in τ results in both static gains from trade and dynamic gains from more idea flows.

Table 10 shows the welfare gains from reducing tariffs in the model with trade-embodied idea flows. We calculate the gains as the equivalent variation in consumption (permanent percentage change) with log utility and a discount rate of 1.9% to match a real interest rate of 5% given the TFP growth rate of 3%. In this model, there are two sources of gains from trade: the standard static gains from exploiting comparative advantage, and the dynamic gains from more idea flows. The first two columns show these effects when tariffs are lowered from $\tau = 1.54$ to $\tau = 1.27$. The last two columns show the gains from moving from near autarky, where the trade share is only around 0.4%, to the baseline where the trade share is 10.2%.

The first row in Table 10 shows the static welfare gains from reducing tariffs. We calculate the static gains as the equivalent permanent gain in consumption from reducing τ while keeping fixed both idea flows and the distribution of productivity in the two countries fixed. The static gains from cutting tariffs in half are 6.2% for the U.S. and 3.6% for the rest of the OECD. The corresponding static gains from moving from autarky to our baseline tariff (implying a trade share of 10.2%) is 26.7% for the U.S. and 20.7% for the rest of the OECD. The second row in Table 10 says the dynamic gains are at least as large as the static gains in all cases. The dynamic gains are even larger for the rest of the OECD than for the U.S. Because the rest of the OECD is less innovative, it gains more ideas than it gives.

For comparison, the static gain for the U.S. implied by the ACR formula (Arkolakis et al., 2012) is 1.1% from moving from autarky and 3% from cutting tariffs

Table 10: Gains from trade with trade-embodied idea flows

	50% reduction in $(\tau - 1)$		Relative to autarky	
	U.S.	OECD	U.S.	OECD
Static gains	6.2%	3.6%	26.7%	20.7%
Dynamic gains	12.6%	20.3%	31.1%	75.0%
Static + dynamic gains	18.8%	23.9%	54.8 %	95.7%

Note: Entries show the permanent increase in consumption that yields the equivalent variation in utility as reducing tariffs from 1.50 to 1.25 (columns one and two) or reducing tariffs from 4 to 1.5 (columns three and four). The aggregate trade share at $\tau = 4$ is about 0.4%. We use a discount rate of 1.9% and log utility.

in half (starting from $\tau = 1.5$).¹⁹ Clearly, our baseline model does not fall into the ACR class in which the trade elasticity is a constant parameter. In our model, trade facilitates the flow of ideas across countries. As a result, the distribution of product quality and the comparative advantage gains from trade vary endogenously with tariffs. Recall Figure 2 above, which plotted the distribution of relative quality across products for the U.S. versus the rest of the OECD. The relative quality distribution was markedly more dispersed near autarky because relative quality drifted apart when ideas did not flow as quickly between countries. As a result, the trade elasticity was only 3 near autarky, whereas it is 5 under the baseline tariff of $\tau = 1.5$ in our model.

In addition, when going from near autarky to $\tau = 1.5$, the trade share initially leaps from 0.4% to 28% in our model. This is shown in the left panel of Figure 5 below. The trade share on impact overshoots the new steady state trade share of 10.2% precisely because of dispersed relative qualities near autarky. Applying the ACR formula to the trade share on impact, the ACR static gains are 9.8% for the U.S., compared to 3.5% with a trade share of only 10.2%. As qualities converge toward each other over time in response to higher trade flows, the trade share

¹⁹The ACR formula for welfare gains relative to autarky is $(1 - \text{trade share})^{-1/(\text{trade elasticity})}$. We use a trade elasticity of 5 based on the survey by Head and Mayer (2014).

eventually settles down to 10.2% and the trade elasticity gradually rises from about 3.8 to 4.8 (the latter is shown in the right panel in Figure 5).

Because cutting tariffs increases idea flows across countries in our baseline model, it speeds up the common steady state growth rate in the two economies. Specifically, moving away from autarky where the trade share is about 0.4% ($\tau \approx 4$) to the baseline with a trade share of 10.2% ($\tau = 1.5$) boosts the steady-state growth rate by 0.62 percentage points. This is the source of the large dynamic gains in our baseline model.

We are far from alone in finding large dynamic gains from trade, though our exact model and mechanisms differ from our predecessors. Sampson (2016) finds that growth is 15 basis points higher with trade than under autarky, lower than our 62 basis points. Alvarez, Buera and Lucas (2017) obtain at least a doubling of world output from eliminating all trade costs. Buera and Oberfield (2020) calculate that about one-fifth of global growth arises due to trade, which is remarkably close to our estimate. Perla, Tonetti and Waugh (2021) estimate that a 10% reduction in trade costs raises consumption-equivalent welfare by 11% and the long run growth rate by 24 basis points.


In the alternative model with disembodied idea flows, the parameters z and z^* govern the spillover of ideas across countries.²⁰ For fixed z and z^* , tariffs have muted effects on welfare. The first two columns of Table 11 show the effect of moving from *trade* autarky to the baseline of $\tau = 1.5$. Here there are only the static gains from exploiting fixed comparative advantage. The static gains are 11.0% for the U.S. and 8.4% for the rest of the OECD in this disembodied model.

The last two columns of Table 11 show the effect of moving from (near) *ideas* autarky, where z and z^* are about 0.01 and 0.99, to their baseline values of $z = .148$ and $z^* = .922$.²¹ Here, moving away from ideas autarky creates no static gains from trade for the simple reason that tariffs do not change, but it has a large effect

²⁰Recall these are the thresholds governing which products foreign firms learn from the U.S. ($0, z$) and which products U.S. firms learn from abroad ($z^*, 1$).

²¹Recall that the steady state of the disembodied model with $z = .148$ and $z^* = .922$ is equivalent to that in the model where idea flows are embodied in trade and $\tau = 1.5$.

Figure 5: Response of trade share and trade elasticity to reduction in τ



Note: The figure shows the response of the trade share and trade elasticity to a one-time reduction of trade cost from near autarky to $\tau = 1.5$.

on the flow of ideas across countries and thus on the growth rate. The dynamic gains from moving away from ideas autarky are 59% for the U.S. and 84.5% for the rest of the OECD. Again, the gains are larger for the rest of world because the U.S. is more innovative. This table underscores that large dynamic gains can be reaped from idea flows, whether those are facilitated by trade or occur in a disembodied fashion.

6. Conclusion

We constructed a two-country model of creative destruction, trade, and growth. In the model, foreign and domestic firms take over each other's markets more frequently when ideas flow more easily across countries. This stimulates growth in the long run under exogenous innovation rates. We find such dynamic gains

Table 11: Gains from trade and idea flows — disembodied spillover model

	Relative to trade autarky		Relative to ideas autarky	
	U.S.	OECD	U.S.	OECD
Static gains	11.0%	8.4%	0%	0%
Dynamic gains	0%	0%	59.0%	84.5%

Note: Entries give the equivalent variation change in the present discounted value of consumption in the disembodied spillovers version of the model as a result of reducing tariffs from 4 to 1.5 (columns one and two) or moving z and z^* inward from 0.01 and 0.99 to their baseline values of $z = 0.147$ and $z^* = 0.922$ (columns three and four). The aggregate trade share at $\tau = 4$ is about 0.4%. We use a discount rate of 1.9% and log utility.

from idea flow are at least as large as the usual static gains from trade.

We provided several strands of evidence consistent with idea flows across trading economies. First, idea flows are necessary to keep comparative advantage from becoming ever stronger and trade elasticities from becoming ever smaller. Second, idea flows can explain why country export patterns evolve across industries over time. Third, global creative destruction can explain why contracting firms are more likely to lose exports than domestic sales, and why expanding firms are more likely gain domestic sales than export markets. Fourth and finally, creative destruction can account for the dominant role of entering and exiting exporters in industry-level changes. We document these last three patterns among manufacturing exporters in the U.S., Chile, China, Colombia, and Indonesia in recent decades.

We see several possible directions for future research. One direction would be to explicitly incorporate frictions to reallocating workers in response to global creative destruction. These might mitigate the dynamic gains from idea flows. Another route would be to study events such as China joining the WTO and see how this affected the extent of job reallocation and innovation. A third avenue would be to obtain more direct evidence on knowledge spillovers (e.g. the frequency of imitation of rich country producers by developing country producers,

or of learning from domestic producers vs. foreign sellers in the local market). We stress again that knowledge spillovers, either embodied in trade or FDI or disembodied, may be necessary to generate realistic trade elasticities and export dynamics at the firm and industry levels. Whether trade policy or other policies have dynamic growth benefits or not, however, hinges on whether the spillovers are largely embodied or disembodied. A final direction would be to model the arrival rates of innovation. We held these fixed for simplicity. Endogenizing innovation rates would allow one to study optimal innovation policy in our setting. Because of domestic knowledge spillovers, national governments may find it optimal to subsidize domestic R&D. But they might not internalize knowledge spillovers to foreign producers who build on domestic innovations. The world might need a “Global Technical Change Accord” to internalize these positive global externalities, just as we need Global Climate Change agreements to internalize negative global pollution externalities.

References

- Aghion, Philippe and Peter Howitt, “A Model of Growth Through Creative Destruction,” *Econometrica*, 1992, 60 (2), 323–351.
- , Antonin Bergeaud, Matthieu Lequien, and Marc J Melitz, “The Heterogeneous Impact of Market Size on Innovation: Evidence from French Firm-Level Exports,” NBER Working Paper 24600, 2020.
- Akcigit, Ufuk, Sina T. Ates, and Giammario Impullitti, “Innovation and Trade Policy in a Globalized World,” 2021.
- Alvarez, Fernando E, Francisco J Buera, and Robert E Lucas, “Idea Flows, Economic Growth, and Trade,” 2017.
- Arkolakis, Costas, Arnaud Costinot, and Andres Rodriguez-Clare, “New Trade Models, Same Old Gains?,” *American Economic Review*, 2012, 102 (1), 94–130.

- , Natalia Ramondo, Andrés Rodríguez-Clare, and Stephen Yeaple, “Innovation and Production in the Global Economy,” *American Economic Review*, 2018, 108 (8), 2128–73.
- Ayerst, Stephen, Faisal Ibrahim, Gaelan MacKenzie, and Swapnika Rachapalli, “Trade and Diffusion of Embodied Technology: An Empirical Analysis,” 2020.
- Bernard, Andrew B and J Bradford Jensen, “Exceptional Exporter Performance: Cause, Effect, or Both?,” *Journal of International Economics*, 1999, 47 (1), 1–25.
- , Jonathan Eaton, J Bradford Jensen, and Samuel Kortum, “Plants and Productivity in International Trade,” *American Economic Review*, 2003, 93 (4), 1268–1290.
- Bloom, Nicholas, Mirko Draca, and John Van Reenen, “Trade Induced Technical Change? The Impact of Chinese Imports on Innovation, IT and Productivity,” *Review of Economic Studies*, 2016, 83 (1), 87–117.
- Buera, Francisco J and Ezra Oberfield, “The Global Diffusion of Ideas,” *Econometrica*, 2020, 88 (1), 83–114.
- Cai, Jie, Nan Li, and Ana Maria Santacreu, “Knowledge Diffusion, Trade, and Innovation across Countries and Sectors,” *American Economic Journal: Macroeconomics*, forthcoming, 2021.
- Coe, David T., Elhanan Helpman, and Alexander W. Hoffmaister, “North-South R&D Spillovers,” *Economic Journal*, 1997, 107 (440), 134–149.
- , —, and —, “International R&D Spillovers and Institutions,” *European Economic Review*, 2009, 53 (7), 723–741.
- Davis, Steven J, John C Haltiwanger, and Scott Schuh, *Job Creation and Destruction*, MIT Press, 1996.
- Dornbusch, Rüdiger, Stanley Fischer, and Paul Samuelson, “Comparative Advantage, Trade, and Payments in a Ricardian Model with a Continuum of Goods,” *American Economic Review*, 1977, 67 (5), 823–839.
- 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.
- Grossman, Gene M. and Elhanan Helpman, “Quality Ladders in the Theory of Growth,” *Review of Economic Studies*, 1991, 58 (1), 43–61.
- and — , *Innovation and Growth in the Global Economy*, MIT press, 1993.
- Hanson, Gordon, Nelson Lind, and Marc-Andreas Muendler, “The Dynamics of Comparative Advantage,” 2018.
- Head, Keith and Thierry Mayer, “Gravity Equations: Workhorse, Toolkit, and Cookbook,” *Handbook of International Economics*, 2014, 4, 131–195.
- Hsieh, Chang-Tai, Peter J Klenow, and Ishan B Nath, “A Global View of Creative Destruction,” NBER Working Paper 26461, 2019.
- Klette, Tor Jakob and Samuel Kortum, “Innovating Firms and Aggregate Innovation,” *Journal of Political Economy*, 2004, 112 (5), 986–1018.
- Krugman, Paul, “A Model of Innovation, Technology Transfer, and the World Distribution of Income,” *Journal of political economy*, 1979, 87 (2), 253–266.
- Melitz, Marc J, “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity,” *Econometrica*, 2003, 71 (6), 1695–1725.
- Perla, Jesse, Christopher Tonetti, and Michael E Waugh, “Equilibrium Technology Diffusion, Trade, and Growth,” *American Economic Review*, 2021, 111 (1), 73–128.
- Ramondo, Natalia, Andrés Rodríguez-Clare, and Milagro Saborío-Rodríguez, “Trade, Domestic Frictions, and Scale Effects,” *American Economic Review*, 2016, 106 (10), 3159–3184.
- Rivera-Batiz, Luis A and Paul M Romer, “Economic Integration and Endogenous Growth,” *Quarterly Journal of Economics*, 1991, 106 (2), 531–555.
- Sampson, Thomas, “Dynamic Selection: An Idea Flows Theory of Entry, Trade, and Growth,” *The Quarterly Journal of Economics*, 2016, 131 (1), 315–380.

Online Appendix

(Not for publication)

A. Markups and Quality Step Improvement

Table A1 shows the markups: μ_j for domestic firms and μ_j^* for foreign firms. A_j' and $A_j^{*'}$ denote the productivity of the second best producer in the domestic and foreign countries. These potential competitors do not produce in equilibrium but affect markups.

Table A1: Markups

	Traded Produced in home	Non-traded	Traded Produced in foreign
Home	$\frac{A_j}{\max \left[A_j', \frac{\omega A_j^*}{\tau} \right]}$	$\frac{A_j}{\max \left[A_j', \frac{\omega A_j^*}{\tau} \right]}$	$\frac{A_j^* / \tau}{\max \left[\frac{A_j^{*'}}{\tau}, \frac{A_j}{\omega} \right]}$
Foreign	$\frac{A_j / \tau}{\max \left[\frac{A_j'}{\tau}, \omega A_j^* \right]}$	$\frac{A_j^*}{\max \left[A_j^{*'}, \frac{A_j}{\omega \tau} \right]}$	$\frac{A_j^*}{\max \left[A_j^{*'}, \frac{A_j}{\omega \tau} \right]}$

Table A2 shows the average step size improvement from creative destruction in the home market for each type of product.

Table A2: Quality Improvement in Home Market

	Domestic firm	Foreign firm
Exported by home	$\phi_{[0,x^*]} \equiv \frac{1}{\theta-1}$	$\phi_{[0,x]}^* \equiv \kappa \left(\frac{\theta}{\theta-1} \cdot \max \left\{ \frac{\tau}{\omega}, 1 \right\} - 1 \right) + (1 - \kappa) \left(\frac{\theta}{\theta-1} \cdot \max \left\{ \frac{\tau}{\omega} \int_0^x (A_j / A_j^*) dj, 1 \right\} - 1 \right)$
Non-traded	$\phi_{[0,x^*]} \equiv \frac{1}{\theta-1}$	$\phi_{[x,x^*]}^* \equiv \left(\frac{\theta}{\theta-1} \cdot \max \left\{ \frac{\tau}{\omega} \int_x^{x^*} (A_j / A_j^*) dj, 1 \right\} - 1 \right)$
Imported by home	$\phi_{[x^*,1]} \equiv \kappa \left(\frac{\theta}{\theta-1} \cdot \max \left\{ \frac{\omega}{\tau}, 1 \right\} - 1 \right) + (1 - \kappa) \left(\frac{\theta}{\theta-1} \cdot \max \left\{ \frac{\omega}{\tau} \int_{x^*}^1 (A_j^* / A_j) dj, 1 \right\} - 1 \right)$	$\phi_{[x^*,1]}^* \equiv \frac{1}{\theta-1}$

Note: Table shows the average improvement in quality in the home market, conditional on creative destruction.

B. Solution steps

We calibrate the model parameters to fit the data moments as follows:

1. Guess the value of $\omega \equiv \frac{w}{w^*}$.
2. The guess for ω will pin down the set of products that are exported, non-traded, and imported (from the U.S. perspective) given the A_j, A'_j, A_j^* and $A_j^{*'} levels.$
3. Calculate markups for each variety in each market.
4. Calculate the real wage in the home country using the markups, relative wages, and realized distribution of quality.
5. Calculate the prices of each variety and the exact OECD consumer price index (the U.S. aggregate consumer price index is normalized to one).
6. Use data on U.S. and OECD export shares as initial guesses for export shares. Use them to calculate the $\bar{\mu}$ and $\bar{\mu}^*$ implied by the distribution of prices and qualities. Given data on L and L^* , the initial guess for ω , and the implied real wage at home w , we then calculate $I - T, I^* - T^*, C$ and C^* .
7. Calculate $I - T$ and $I^* - T^*$ by adding up sales of each variety.
8. Calculate $\bar{\mu}$ and $\bar{\mu}^*$ implied by the distribution of revenues and by GDP net of tariff revenues.
9. Iterate over ω until the following conditions hold:
 - (a) Trade is balanced.
 - (b) The initial guesses for $\bar{\mu}$ and $\bar{\mu}^*$ in step 4 are equal to $\bar{\mu}$ and $\bar{\mu}^*$ calculated in step 8.
 - (c) $I - T$ and $I^* - T^*$ implied by the initial guesses from step 6 are equal to $I - T$ and $I^* - T^*$ calculated in step 7.