

NBER WORKING PAPER SERIES

A GLOBAL VIEW OF CREATIVE DESTRUCTION

Chang-Tai Hsieh
Peter J. Klenow
Ishan B. Nath

Working Paper 26461
<http://www.nber.org/papers/w26461>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2019

We thank Beiling Yan of Statistics Canada for her assistance with the Canadian manufacturing data; Jean-Felix Brouillette, Eric English and Erxiao Mo for excellent research assistance; and Sam Kortum for very helpful comments. 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 or the National Bureau of Economic Research.. All results have been reviewed by the U.S. Census Bureau and Statistics Canada to ensure that no confidential information is disclosed.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Chang-Tai Hsieh, Peter J. Klenow, and Ishan B. Nath. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

A Global View of Creative Destruction
Chang-Tai Hsieh, Peter J. Klenow, and Ishan B. Nath
NBER Working Paper No. 26461
November 2019
JEL No. F11,F14,F43

ABSTRACT

In the wake of the U.S.-Canada Free Trade Agreement, both the U.S. and Canada experienced a sustained increase in job reallocation, including firms moving into exporting. The change involved big firms as much as small firms. To mimic these patterns, we formulate a model of innovation by both domestic and foreign firms. In the model, trade liberalization quickens the pace of creative destruction, thereby speeding the flow of technology across countries. The resulting dynamic gains from trade liberalization are an order of magnitude larger than the gains in a standard static model.

Chang-Tai Hsieh
Booth School of Business
University of Chicago
5807 S Woodlawn Ave
Chicago, IL 60637
and NBER
chsieh@chicagoBooth.edu

Ishan B. Nath
University of Chicago
inath33@gmail.com

Peter J. Klenow
Department of Economics
579 Serra Mall
Stanford University
Stanford, CA 94305-6072
and NBER
Klenow@Stanford.edu

A data appendix is available at <http://www.nber.org/data-appendix/w26461>

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 focused mostly on models in which trade liberalization leads to a burst of job reallocation and growth, but no medium or long run effect on either.

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 (2018). Modeling builds on the foundational work of Rivera-Batiz and Romer (1991) and Grossman and Helpman (1993). Recent efforts include Alvarez, Buera and Lucas (2013), Buera and Oberfield (2017), Akcigit, Ates and Impullitti (2018), and Perla, Tonetti and Waugh (2019).

In this paper we present facts and a model on the interaction of creative destruction and trade. We document job reallocation tied to exports in U.S. and Canadian manufacturing firms before and after the 1988 U.S.-Canada Free Trade Agreement. Exit and job destruction rates rose for big firms and small firms, a result in line with the findings of Holmes and Stevens (2014) for the U.S. in the wake of the China shock.

In our model, ideas flow across two countries through trade. Innovators draw from a Pareto distribution building on the technology of the firm selling in the local market. When innovators take over the market for an existing product (creative destruction), job reallocation takes place. Domestic firms can take over foreign markets for a product, and foreign firms can take over the domestic market. When this happens, exported products are also reallocated across countries.

The first version of the model features exogenous arrival rates of innovation as in Garcia-Macia, Hsieh and Klenow (2019). It is a two-economy version of the influential Klette and Kortum (2004) model of creative destruction, only with exogenous arrival rates. Our second version of the model endogenizes the

arrival rates. We build in diminishing returns to the stock of ideas *a la* Jones (1995) and Bloom, Jones, Van Reenen and Webb (2019), so that growth is semi-endogenous and linked to growth in the number of researchers. In both models, the two trading countries grow 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, growth in research investment, exports relative to all shipments, and the share of entrants in total employment. To pin down the Pareto shape parameter we fit the gap in revenue per worker for exporters vs. non-exporters in U.S. plants. We also target value added per worker and 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.

Once calibrated, we analyze steady states and transition dynamics in response to tariff changes. In the exogenous arrival rate version of the model, lower tariffs boost the growth rate in both the U.S. and the rest of the OECD. Because the U.S. is more innovative, the rest of the OECD benefits more and its real consumption wage rises relative to that of the U.S. Lower tariffs also lead to more job destruction. There is a spike immediately after tariffs are lowered, but job destruction remains higher in the new steady state.

In the endogenous arrival version of the model, lower tariffs boost growth only temporarily. This is because of diminishing returns in idea production. Ideas do spread faster with lower tariffs, so that each country ascends to a higher TFP path than before the liberalization. The rest of the OECD benefits more because they receive more U.S. ideas than they send to the U.S. Welfare gains from trade, in consumption-equivalent terms, are about 37% in the U.S. and 69% in the rest of the OECD, an order of magnitude higher than in a standard model with no changes in technology.

To dissect our dynamic gains from trade, we entertain alternative assumptions about idea flows across countries. When we assume countries learn partially from local *producers* rather than sellers in the local market, the gains from trade shrink toward the static gains. Thus idea flows are critical to our large

dynamic gains from trade. When we assume countries specialize in innovating on products they produce, however, the dynamic gains remain large for the U.S. Given its innovativeness, the U.S. gains a lot from specializing its draws on a subset of products. Due to limited idea flows across countries, the rest of the OECD benefits less from trade when there is research specialization. Models with limited idea flows and research specialization, however, predict much less reallocation of exported products across countries than observed in the data.

Our effort is most closely related to three recent papers. Perla, Tonetti and Waugh (2019) study the impact of trade on exit, entry, domestic technology diffusion, and growth in a model of symmetric countries. Like us, they find large dynamic gains from trade. They derive analytical solutions in a model of many countries, whereas we simulate a two-country model calibrated to evidence on trade and job flows. Our focus is innovation, idea flows across countries, and creative destruction, whereas their focus is on the interaction of trade with domestic technology diffusion.

We follow Buera and Oberfield (2017) in studying international technology diffusion in a model with Bertrand competition. They arrive at Frechet distributions of productivity within countries, allowing them to characterize multi-lateral trade flows as in Bernard, Eaton, Jensen and Kortum (2003). Our focus is more empirical, as we try to match evidence on job reallocation associated with creative destruction from trade.

Akcigit, Ates and Impullitti (2018) are similar to us in characterizing 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 that are crucial for how trade can induce more innovation. They analyze transition dynamics and optimal R&D subsidies. Their knowledge spillovers take the form of followers catching up to leaders in one big jump if they fall too far behind. They emphasize the convergence of patenting in other advanced countries toward the U.S. in recent decades. In our model and empirics, in contrast, we focus on how trade affects job reallocation.

The rest of the paper is organized as follows. Section 2 lays out nine facts from U.S. and Canadian manufacturing that we attempt to explain. In Section 3 we present a two-country model of creative destruction and growth with exogenous innovation rates. Section 4 endogenizes the innovation rates. In Section 5 we carry out additional exercises (alternative assumptions about idea flows, U.S.-Canada trade liberalization). Section 6 concludes.

2 Facts from Canadian and U.S. Manufacturing

We use the U.S. Longitudinal Business Database (LBD) and Canada's Annual Survey of Manufactures (ASM). The LBD contains administrative employment records on all nonfarm private establishments with employees in the U.S. for every year from 1977 to 2013; we look at establishments owned by firms with at least one manufacturing establishment in the given year.¹ The ASM covers all but the smallest manufacturing establishments every year from 1973 to 2012.²

From the LBD and ASM we use plant and firm identifiers, employment, and industry (four-digit SIC or six-digit NAICS). The ASM has information on exports every five years from 1974 to 1989, for 1993, and every year from 1996 to 2012. The LBD does not measure exports but this information is available in the micro-data of the U.S. manufacturing census every five years starting in 1987. We merge the establishments in the manufacturing census with the LBD to measure exports in our LBD sample.³ We aggregate establishment data in the U.S. and Canada to the firm level and highlight nine facts:

1. **Large Job Flows.** Table 1 (rows 1 and 2) presents manufacturing job creation and destruction rates over five years in Canada (from 1973 to 2012)

¹We include the non-manufacturing establishments of such firms to account for the relocation of jobs from establishments classified as manufacturing to establishments of the same firm that are classified as non-manufacturing.

²The survey threshold is currently annual sales of 30 thousand Canadian dollars.

³The LBD and the manufacturing census use the same plant identifiers.

Table 1: Job Flows in the U.S. and Canada

	U.S.	Canada
Job Creation Rate	31.4%	32.4%
Job Destruction Rate	36.6%	31.6%
Job Destruction from Large Firms	30.7%	15.3%
Job Creation from Exports	2.0%	23.3%

Note: Job creation and destruction rate are calculated over successive five year periods from 1987 to 2012 for the U.S. and from 1973 to 2012 for Canada. Jobs from exports are imputed as the product of firm employment and the ratio of exports to total shipments. “Large” refers to above-mean employment in the initial year of each five year period.

and in the U.S. (from 1987 to 2012).⁴ As in the classic work by Davis, Haltiwanger and Schuh (1996), job flows are large. The average job creation and destruction rate over five years is about 30% in Canada. The average job creation rate in U.S. manufacturing from 1973 to 2012 is also about 30%. The U.S. job destruction rate is about 5 percentage points higher.

- 2. Job destruction due to “large” firms.** Row 3 in Table 1 presents the job destruction rate among firms with above-average employment in the initial period. Such large firms account for 84% of all job destruction in the U.S. and 48% of all job destruction in Canadian manufacturing.

⁴The job creation rate between year t and $t + 5$ is defined as the ratio of (a) the sum of employment of new firms established between year t and year $t + 5$ and the change in employment among *expanding* firms between the two years; to (b) average total employment across years t and $t + 5$. The job destruction rate between years t and $t + 5$ is the sum of employment in year t of firms that exited between the two years and the change in employment between years t and $t + 5$ among *contracting* firms divided by average total employment (in the beginning and ending years). Job flows for the U.S. are calculated for every five year period from 1987 to 2012. Job creation, destruction, and job destruction from large firms for Canada are calculated every five years from 1973 to 2008. For 2008 to 2012, we multiply by 5/4 to impute the flow over five years. Job creation from exports in Canada is calculated from 1974–1979, 1979–1984, 1984–1989, 1989–1993, 1993–1998, 1998–2003, 2003–2008, and 2008–2012, where we multiply the rate from 1989–1993 and 2008–2012 by 5/4 to impute the flow over five years.

3. **Job creation due to exports.** We impute employment due to exports as the product of a firm's employment and the ratio of its exports to total shipments. Job creation from exports is the sum of imputed employment in year $t + 5$ of new exporters (firms that enter into exporting between year t and $t + 5$) and the change in imputed employment from firms where exports increased between the two years. We divide this measure of job creation from exports by the average of aggregate employment in years t and $t + 5$. The resulting number, in row 4 in Table 1, shows that the job creation rate due to exports is 23% in Canada. The job creation rate due to exports in the U.S. is much smaller at 2%.⁵

4. **Reallocation of export products across countries.** Table 2 presents two measures of the reallocation of exported products across countries.⁶ The top panel shows the probability that an exported product in a given year is no longer exported by the same country the following year. This number is about 8% for the average 4-digit manufacturing U.S. export (row 1) and 15% for the bottom half of U.S. manufacturing export products.⁷ The bottom panel replicates Hanson, Lind and Muendler (2018)'s measure of mean reversion of a country's *top* export. Row 3 shows the share of the country's top exported product in total exports in year t . Row 4 shows the ratio of the export share of the same product in year $t - 5$ relative to the share in year t . This ratio averages 66.5% for the U.S. and 86.2% for the other OECD countries.⁸

We next document how job flows changed after the Canada-U.S. Free Trade Agreement (CUSFTA). This agreement was signed on January 2, 1988, and mandated annual reductions in tariffs and other trade barriers over a 20-year period

⁵Lincoln, McCallum and Siemer (2019) estimate that 29% of U.S. exports in 2002 were by firms that had been exporting for fewer than 5 years.

⁶A product in Table 2 is one of the 540 4-digit SITC (revision 2) manufacturing industries in Feenstra, Lipsey, Deng, Ma and Mo (2005)'s World Trade Database.

⁷Rows 1 and 2 in Table 2 are the average of one-year rates from 1982–1983 to 2002–2003.

⁸These numbers are the average of five-year panels from 1982–1987 to 2002–2007.

Table 2: Reallocation of Export Products Across Countries

	U.S.	Rest of OECD
Annual Exit Rate		
All Exported Products	7.8%	8.4%
Bottom 50% in Export Sales	15.1%	15.7%
Top Export Product		
Share of total exports year t	5.7%	12.8%
Share of total exports year $t - 5$ / year t	66.5%	86.2%

Note: A product is one of 540 4-digit manufacturing industries in Feenstra et al. (2005). Rows 1 and 2 show the probability an exported product in year t is no longer exported by the country in year $t+1$, for all exported products and products in bottom half of export sales, respectively. Row 3 shows the share of the top exported product in total exports. Row 4 shows the relative share of the same product *five years before*. Entries are averages of one year panels from 1982–1983 to 2002–2003 (rows 1-2) or five year panels from 1982–1987 to 2002–2007 (rows 3-4).

starting on January 1, 1989.⁹ For Canada, we highlight three differences in 1988–2003 (Post-CUSFTA) versus 1973–1988 (Pre-CUSFTA):

5. **Job flows increased after trade liberalization.** Table 3 shows that job creation and destruction rates increased in Canada after CUSFTA.¹⁰
6. **Large firms increased job destruction after trade liberalization.** Holmes and Stevens (2014) show that large U.S. manufacturing firms were adversely affected by the surge in imports from China. Row 3 in Table 3 documents a similar fact in Canada. The job destruction rate among large (above-mean employment) firms increased by 2 percentage points after CUSFTA, out of an overall increase in job destruction of 7 percentage points.
7. **Job creation from exports increased after trade liberalization.** The last

⁹The average tariff on manufacturing imports in CUSFTA fell from over 8% to below 2% in Canada and from 4% to below 1% in the U.S.

¹⁰Our Online Appendix (<http://klenow.com/global-view-creative-appendix.pdf>) shows that job creation (destruction) rose more in Canadian 2-digit industries with bigger export (import) tariff cuts.

Table 3: Job Flows in Canada

	Pre-CUSFTA	Post-CUSFTA
Job Creation Rate	30.0%	31.3%
Job Destruction Rate	25.5%	32.7%
Job Destruction from Large Firms	22.1%	24.0%
Job Creation from Exports	9.0%	17.7%

Note: Pre-CUSFTA is 1978 to 1988. Post-CUSFTA is 1988 to 2003.

row in Table 3 shows that job creation from exports increased by almost 9 percentage points in Canada after the trade agreement.

Table 4 presents the change in job flows in U.S. manufacturing after CUSFTA. The timing of the U.S. data does not align as well with the trade agreement so here we focus on the 1972–1987 as the “pre-CUSFTA” period and 1992–2012 as the “post-CUSFTA” period. As documented by a large literature, there was also a surge of imports from China in the 1992–2012 period, so one should not interpret the changes in Table 4 as coming only from CUSFTA. Still, job destruction increased markedly after 1987, by about 6 percentage points (row 2).¹¹ The increase in job destruction was entirely driven by large firms.

We now look at differences between exporting and non-exporting firms. Figure 1 plots the distribution of employment (in the left panel) and labor productivity (revenue per worker, in the right panel) from the U.S. manufacturing census in 2012. This figure reveals two additional facts:

8. Average labor productivity and employment is higher for exporters than for non-exporters. This can easily be seen in Figure 1.

¹¹This may seem surprising given the evidence on declining dynamism in Decker, Haltiwanger, Jarmin and Miranda (2014). For U.S. manufacturing firms, this decline was concentrated in job creation and took place well after CUSFTA.

Table 4: Job Flows in the U.S.

	1977–1987	1987–1992	1992–2012
Job Creation Rate	33.7%	33.0%	31.1%
Job Destruction Rate	31.6%	31.0%	37.8%
Job Destruction from Large Firms	25.8%	24.9%	32.1%
Job Creation from Exports	–	2.0%	2.0%

Note: Calculated from U.S. manufacturing census micro-data. Job creation and destruction calculated over five year periods. “Large” firms are above average employment firms in the initial year.

9. **Overlap of labor productivity and employment between exporters and non-exporters.** Again, see Figure 1.

3 Exogenous Innovation

This section presents a model of growth driven by creative destruction, where innovation can come from domestic or foreign firms. The goal is to examine the dynamic gains from trade liberalization, and to see whether this model can mimic the nine facts described in section 2.

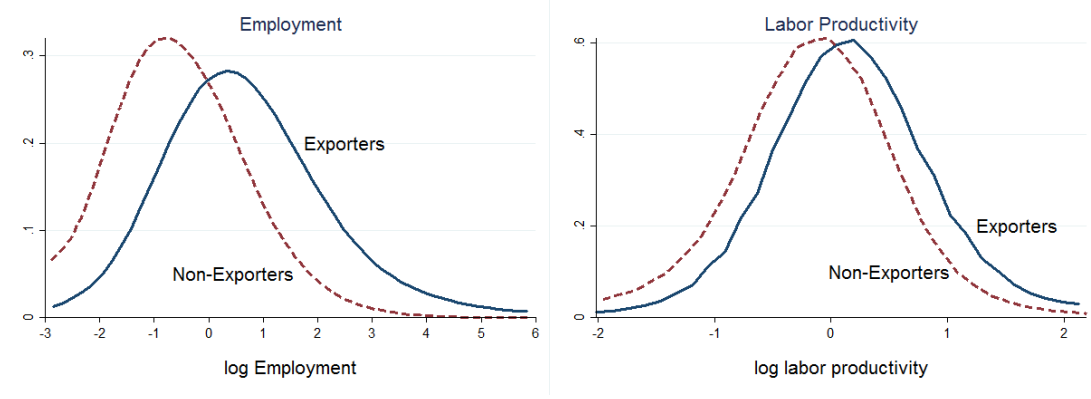
3.1 Static Equilibrium

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:

$$U = \int_0^1 \ln C_j dj. \quad (1)$$

Figure 1: Distribution of Employment and Labor Productivity



Note: The distribution of labor productivity (value-added per worker) and employment of exporting and non-exporting firms in the 2012 U.S. Census of Manufacturing.

This utility function implies that consumers spend the same on each variety.¹²

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

$$\frac{A_x}{\tau} = \omega A_x^* \quad (2)$$

$$A_{x^*} = \frac{\omega A_{x^*}^*}{\tau} \quad (3)$$

where ω denotes the relative wage (domestic relative to foreign) and $\tau \geq 1$ is the symmetric gross trade cost. 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 is the gap between the incumbent firm’s marginal cost and the cost of

¹²Utility of the foreign consumer is analogously given by $U^* = \int_0^1 \ln C_j^* dj$.

Table 5: Markups

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

its closest competitor — domestic or foreign. Table 5 summarizes 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 markets, respectively. These potential competitors do not produce in equilibrium but affect markups.

The relative wage is pinned down by balanced trade:

$$I^* \cdot x = I \cdot (1 - x^*) \quad (4)$$

where I and I^* denote nominal GDP at home and abroad, respectively. The left hand side of equation (4) 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

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

where $\bar{\mu}^*$ and $\bar{\mu}$ denote the average gross markup of foreign and domestic firms, w and w^* are the home and foreign wage, and L and L^* are labor supply at home and abroad.¹³ More exactly, the average price-cost markup in the U.S. satisfies

$$\frac{1}{\bar{\mu}} \equiv \frac{\int_0^{x^*} \frac{1}{\mu_j} dj + \frac{1}{\tau} \cdot \int_0^x \frac{1}{\mu_j^*} dj}{x^* + x/\tau}$$

¹³The expression for nominal income comes from equating nominal income to the revenue of local firms plus tariff revenue: $I = \bar{\mu} w L + (\tau - 1) \frac{I}{\tau} (1 - x^*)$ and $I^* = \bar{\mu}^* w^* L^* + (\tau - 1) \frac{I^*}{\tau} \cdot x$.

where μ_j^f denotes the markup of domestic firms on their exported products. The expression for the foreign firms' average markup is analogous.

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 trade cost. The real wage at home W and in the foreign country W^* are given by

$$\ln W = \int_0^{x^*} \ln \left(\frac{A_j}{\mu_j} \right) dj + \int_{x^*}^1 \ln \left(\frac{A_j^*}{\mu_j^*} \cdot \frac{\omega}{\tau} \right) dj \quad (5)$$

$$\ln W^* = \int_0^x \ln \left(\frac{A_j}{\mu_j} \cdot \frac{1}{\omega \tau} \right) dj + \int_x^1 \ln \left(\frac{A_j^*}{\mu_j^*} \right) dj. \quad (6)$$

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

3.2 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, and innovation only takes the form of creative destruction. Unlike Klette and Kortum, we allow trade and for creative destruction to come from a firm in another country.

For now we posit constant exogenous arrival rates for innovation. (We will endogenize arrival rates in the next section.) 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 quality of the product sold in the innovating firm's *local* market.¹⁴ Later we will entertain alternative assumptions,

¹⁴Knowledge 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) for evidence consistent with learning by importing.

Table 6: Channels of Innovation

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

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

such as learning from domestic producers only.

The quality of 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 over an existing variety (conditional on innovation) is thus $\frac{1}{\theta-1} > 0$. We add a reflecting barrier whereby the bottom ψ percent of products in each year, in terms of their quality, redraw their quality from the top $1-\psi$ percent of products. This is in the spirit of what Perla et al. (2019) obtain endogenously, and will help maintain a stationary distribution of qualities.

The notation for innovation probabilities is given in Table 6. The probability a product is improved upon by an incumbent domestic firm is λ . Conditional on not being improved by a domestic incumbent, η is the probability the product is improved by an entering domestic firm. Conditional on not being improved by *any* domestic firm, λ^* is the probability the product will be improved by a foreign incumbent firm. Finally, conditional on the product not being improved upon by either a domestic firm or by a foreign incumbent, 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 incumbent firm, or a foreign entrant.¹⁵

¹⁵The *unconditional* probabilities of each type of innovation are λ , $\tilde{\eta} \equiv \eta(1-\lambda)$, $\tilde{\lambda}^* \equiv \lambda^*(1-\lambda)(1-\tilde{\eta})$, and $\tilde{\eta}^* \equiv \eta^*(1-\lambda)(1-\tilde{\eta})(1-\tilde{\lambda}^*)$. So the unconditional probability a domestic firm (entrant or incumbent) improves a product is given by $\lambda + \tilde{\eta}$, and the unconditional probability a foreign firm innovates is $\tilde{\lambda}^* + \tilde{\eta}^*$.

Table 7 summarizes the probability of creative destruction in the domestic market (rows 1-3) and in the foreign market (rows 4-6) due to innovation by domestic firms (column 1) and by foreign firms (column 2). The first row shows the arrival rate of new ideas in the domestic market for a product that is also exported to the foreign market. The probability this product is improved upon by another domestic firm is $\lambda + \tilde{\eta}$, and a domestic innovator will always replace the incumbent firm in this market. A foreign firm also improves upon the same product with probability $\tilde{\lambda}^* + \tilde{\eta}^*$, but a successful foreign innovator does not necessarily replace the domestic incumbent. Since quality improvement follows a Pareto distribution, the probability that the quality improvement of the foreign innovator is large enough to replace the domestic incumbent is $\min \left[\left(\frac{\omega}{\tau} \right)^\theta, 1 \right]$.

For a given innovation rate by foreign firms, higher relative wages ω and lower trade costs τ increase the probability that innovation by a foreign firm benefits domestic consumers. Intuitively, higher domestic wages increase the probability a foreign innovator will be competitive enough to replace the incumbent in the domestic market. Higher trade costs make the foreign innovator less competitive compared to the domestic incumbent. Effectively, trade costs insulate domestic firms from foreign competition in the domestic market.

The expected growth rate of the real consumption wage in the domestic market is the product of the rate of creative destruction in rows 1-3 in Table 7 and the increases in product quality (conditional on the product being replaced). And the expected growth rate of the foreign real consumption wage is the product of the arrival rates in rows 4-6 in Table 7 and the corresponding improvements in quality. Real growth rates in the two countries depend on the arrival rates of innovation $\lambda + \tilde{\eta}$ and $\tilde{\lambda}^* + \tilde{\eta}^*$, the relative wage ω , and the share of each type of product (x and x^*). As discussed in the previous section, 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^* . The distribution of A_j/A_j^* is endogenous to innovation.

Table 7: Probability of Creative Destruction

Market	Product Type	Domestic Firm	Foreign Firm
Home	Exported by Home	$\lambda + \tilde{\eta}$	$(\tilde{\lambda}^* + \tilde{\eta}^*) \min \left[\left(\frac{\omega}{\tau} \right)^\theta, 1 \right]$
	Non-Traded	$\lambda + \tilde{\eta}$	$(\tilde{\lambda}^* + \tilde{\eta}^*) \min \left[\left(\frac{\omega A_j^*}{\tau A_j} \right)^\theta, 1 \right]$
	Imported by Home	$(\lambda + \tilde{\eta}) \min \left[\left(\frac{\tau}{\omega} \right)^\theta, 1 \right]$	$\tilde{\lambda}^* + \tilde{\eta}^*$
Foreign	Exported by Home	$\lambda + \tilde{\eta}$	$(\tilde{\lambda}^* + \tilde{\eta}^*) \cdot \min \left[(\omega\tau)^\theta, 1 \right]$
	Non-Traded	$(\lambda + \tilde{\eta}) \min \left[\left(\frac{A_j}{\omega\tau A_j^*} \right)^\theta, 1 \right]$	$\tilde{\lambda}^* + \tilde{\eta}^*$
	Imported by Home	$(\lambda + \tilde{\eta}) \min \left[\left(\frac{1}{\omega\tau} \right)^\theta, 1 \right]$	$\tilde{\lambda}^* + \tilde{\eta}^*$

To understand the equilibrium in the model with innovation, it is useful to consider the case of completely free trade ($\tau = 1$). In this case, all products are traded so the relevant arrival rates in Table 7 are rows 1 and 3 (for the domestic market) and 4 and 6 (for the foreign market). The probability a domestic firm creatively destroys another firm is thus given by:

$$\underline{\text{Domestic}} \text{ creative destruction rate} = (\lambda + \tilde{\eta}) \cdot x^* + (\lambda + \tilde{\eta}) \min \left[\omega^{-\theta}, 1 \right] \cdot (1 - x^*).$$

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:

$$\underline{\text{Foreign}} \text{ creative destruction rate} = (\tilde{\lambda}^* + \tilde{\eta}^*) \cdot (1 - x^*) + (\tilde{\lambda}^* + \tilde{\eta}^*) \min \left[\omega^\theta, 1 \right] \cdot x^*.$$

Ceteris paribus, higher ω 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 creative destruction rate of domestic firms is the same as that of foreign firms.

We close this subsection by contrasting autarky and free trade when the two countries are symmetric in size and in their innovation arrival rates. In this special case the relative wage $\omega = 1$ and the expressions become simply:

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

$$\underline{\text{Frictionless growth rate}} = \left(\lambda + \tilde{\eta} + \tilde{\lambda}^* + \tilde{\eta}^* \right) \frac{1}{\theta - 1}.$$

In autarky each country benefits only from domestic arrivals. With frictionless trade, each country benefits from both domestic and foreign arrivals. This underscores the scale effect generating dynamic gains from trade in this model.¹⁶

3.3 Calibration

The model is summarized by two innovation rates (for incumbents and entrants) in each country, the shape parameter of the Pareto distribution of the innovation draws, and a trade cost. In this section, we infer the value of these parameters from simple moments in the data.

For consistency with the model, we assume the world consists of the manufacturing sectors in the U.S. and the rest of the OECD (“foreign”). The shape parameter of the Pareto distribution of innovation draws (θ), relative employment (L/L^*), innovation rates in Table 6, and the trade cost (τ) jointly determine the growth rate, the trade share and the relative wage. For a given value of θ , we can back out the trade cost and innovation rates from data on total employment, the growth rate, the trade share, and the relative wage.¹⁷ We use the employment share of new firms to pin down innovation by entrants vs. incumbents. Finally,

¹⁶Section A.3 of the online Appendix shows the expected growth rate for the general case with trade frictions and asymmetric innovation rates.

¹⁷We describe in the next section how we back out θ from the gap in labor productivity (revenue per worker) between exporters and non-exporters.

Table 8: Data Moments used for Calibration

Data Moment	Source	Value
Revenue per worker exp./non-exp.	U.S. mfg	1.066
TFP growth rate	U.S. mfg	3.01%
Value added per worker home/foreign	U.S. and OECD mfg	1.29
Employment share of entrants	U.S. mfg	14.4%
Export share of revenues (home)	U.S. mfg	10.2%
Employment home/foreign	U.S. and OECD mfg	0.389
Employment growth rate	U.S. mfg	-1.1%
Trade elasticity from halving τ	Head and Mayer (2014)	-5

we use the elasticity of trade flows with respect to trade barriers to back out the reflecting barrier ψ and the implied dispersion of product qualities.

Table 8 displays the data moments we target. In steady state we fit TFP growth of 3% per year, employment shrinking at 1.3% per year, output per worker in the home country that is 29% higher than in the foreign country, and a home trade share of 10%. The innovation rates and trade cost needed to fit these facts are shown in Table 9.¹⁸

The innovation rates in Table 9 are *conditional*. The *unconditional* innovation rates are $\lambda + \tilde{\eta} = 0.160$ for domestic firms and $\tilde{\lambda}^* + \tilde{\eta}^* = 0.122$ for foreign firms. The innovation rate has to be higher for domestic firms to explain the 29% higher real wage in the home country. Conditional on the innovation rates and the relative size of the two economies, the trade share pins down the trade cost,

¹⁸We simulate the model with 5,000 varieties in each country. Each variety receives innovation draws that are randomly assigned to an existing incumbent or a new entrant. The relative wage is selected to balance trade between the two countries. 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 allow us to match the moments in the data. For more details on our calibration see Section A.2 of our Online Appendix.

Table 9: Estimates of Model Parameters

Variable	Description	Value
θ	Shape parameter of innovation draws	10.9
λ	Home innovation rate from incumbent	13.5%
η	Home innovation rate from entrants	2.95%
λ^*	Foreign innovation rate incumbents	11.5%
η^*	Foreign innovation rate entrants	3.42%
τ	Trade cost	1.49
ψ	Reflecting barrier for product quality	1.10%

a 50% tariff.¹⁹ Finally a reflecting barrier where the bottom 1.1% of products by quality redraw from the top 98.9% of products generates a trade elasticity of 5.

3.4 Firm Dynamics

Table 10 shows that the model can replicate, at least qualitatively, facts 1-4 that we presented in Section 2 (large job flows, lots of job destruction at large firms, some job creation due to exports, and high rates of export turnover). The model can also speak to facts 5-7 (job flows, job destruction at large firms, and job creation due to exports all increase after trade liberalization). More specifically:

Table 10 (rows 1 and 2, column 2) shows the job creation and destruction rates in the steady state of the model parameterized to fit the moments in Table 8. The job creation rate (over five years) is 32%, and the job destruction rate is 6% higher at about 38%. For comparison, the first column in Table 10 replicates the U.S. data. The job flows predicted by the model with the parameters in Table 9 are roughly of the same magnitude as in the data (fact 1).

¹⁹Eaton and Kortum (2002) and others infer high trade costs to explain bilateral trade flows.

The third row in Table 10 shows that, consistent with evidence from U.S. and Canadian manufacturing (fact 2), job destruction in the model also comes from large (above-mean employment) firms. The job destruction rate by large firms is 22% in the model, which is about two-thirds of the overall job destruction.

The fourth row in Table 10 shows that the job creation rate from exports in the model is 5.7%. In the U.S. data, the number is 2%. We cannot see empirically when jobs are destroyed because domestic producers are replaced by imports, but we can calculate this moment in the model. The last row shows that the job destruction rate due to imports is 7.4%, so about a fifth of the overall job destruction in the model comes from creative destruction by foreign firms.

In the data, the probability the U.S. loses an exported product to a foreign competitor in a given year is 8% for all exported SITCs and 15% for the bottom half of exported SITCs. In our baseline simulation, the annual probability that a U.S. export is reallocated to foreign producer is 13% (next to last row of Table 10). The (empirical) average relative share of the top export product five years prior is 66.5%. To mimic the top exported product, we divide our 5,000 products randomly into 125 categories, identify the top export category in each year, and calculate its average relative share 5 years prior. Our simulation yields a 77.5% average relative share of the top export product (last row in Table 10).

We next simulate the new steady state of the model with lower trade costs, holding other parameters constant. A key assumption in this exercise is that domestic and foreign innovation rates do not change when trade costs change.²⁰ We first analyze the effect of reducing trade costs for TFP growth in the two countries. Figure 2 (left panel) shows that lowering tariffs boosts the growth rate of TFP in the two countries in tandem. Intuitively lower trade costs make it easier for a country to build on ideas developed in the other country. Holding constant innovation rates, moving to a world with frictionless trade increases the steady-state TFP growth rate from 3.0% to 3.3%. Lower trade costs also raise

²⁰We relax this assumption in the next section, when we endogenize innovation rates. Section A.3 of our Online Appendix provides more comparative statics under exogenous innovation.

Table 10: Firm Dynamics, Data vs. Simulations

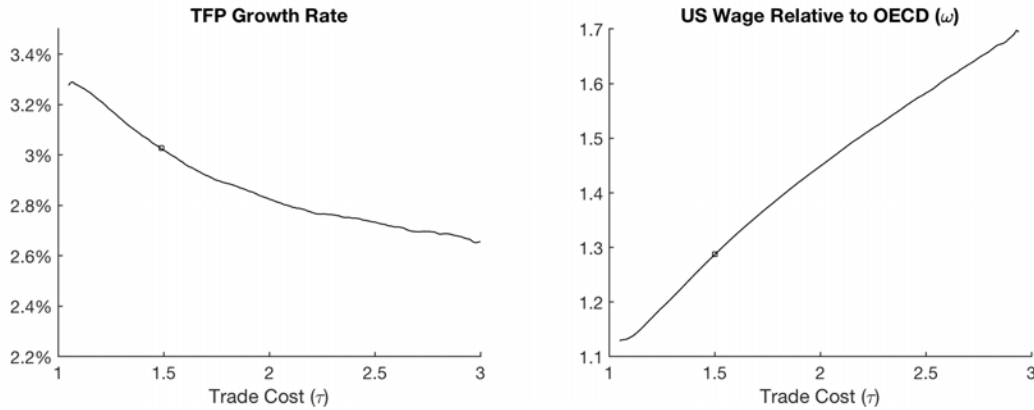
Moment	U.S. Data	Simulations
Job Creation Rate	31.4%	31.6%
Job Destruction Rate	36.6%	38.2%
Job Destruction from Large Firms	30.7%	22.3%
Job Creation from Exports	2.0%	5.7%
Job Destruction from Imports	–	7.4%
Probability of Losing an Export	15.1%	12.8%
Top Export Product Turnover	66.5%	77.5%

Note: The U.S. data is the average from 1987 to 2012. The second column shows simulated steady state moments in the model with the parameter values from Table 9.

the foreign wage relative to the domestic wage. This is shown in the right panel in Figure 2. The intuition is that a country who innovates less benefits more from trade liberalization since it is now easier for the country to “import” ideas. The home/foreign wage is 1.15 with frictionless trade, versus 1.29 in the baseline with a roughly 50% tariff rate.

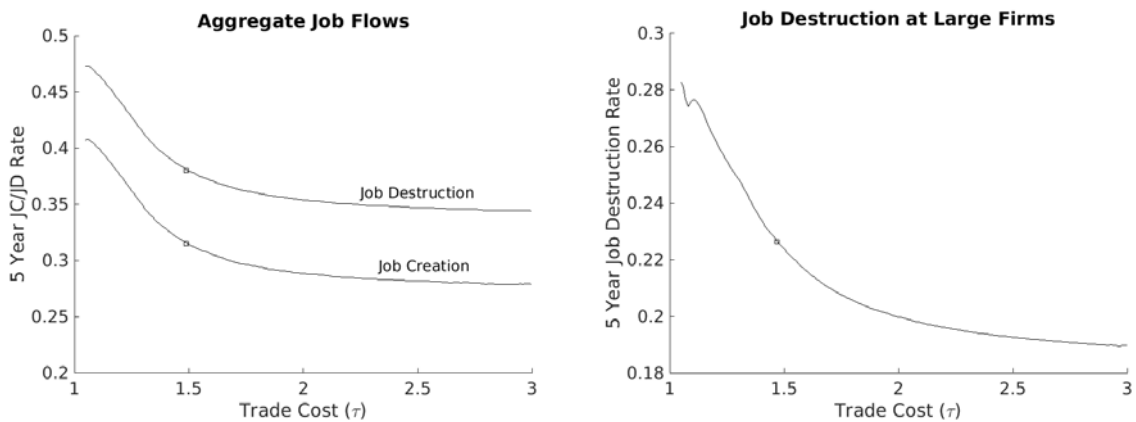
Figure 3 shows the effect of trade costs on job creation and destruction. Consistent with the evidence from Canada and the U.S. (fact 5), the model predicts that job flows rise when trade costs fall. Relative to our baseline steady state ($\tau = 1.491$), free trade would increase job creation and destruction rates by about 10 percentage points. Thus a 30 basis point increase in the growth rate is associated with a proportionally bigger impact on job reallocation. Lower trade costs facilitate creative destruction from trade for smaller step sizes (i.e., smaller quality improvements relative to foreign competitors). Figure 3 also shows the job destruction rate from large firms (those with above-average employment) for different values of τ . Consistent with the evidence from the U.S. and Canada (fact 6), in the model job destruction from large firms rises when trade costs fall.

Figure 2: Simulated TFP Growth Rate and Relative Wage vs. Trade Costs



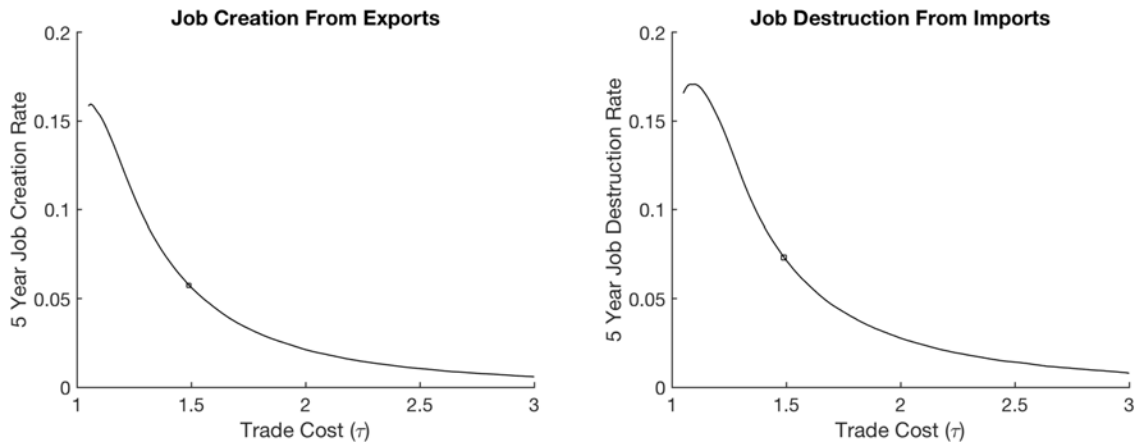
Note: The left panel shows simulated growth rate of the real wage and the right panel the relative wage (home/foreign) in the steady state for different values of τ . All other parameters of model are kept fixed at the values in Table 9.

Figure 3: Simulated Job Creation and Destruction vs. Trade Costs



Note: The figures simulate the steady state job creation and destruction rates when we vary τ but keeping constant all other parameters at the values in Table 9. The figure simulates the steady state job destruction rate from large (above mean employment) firms when we vary τ but keeping constant all other parameters at the values in Table 9.

Figure 4: Simulated Job Flows from Trade vs. Trade Costs

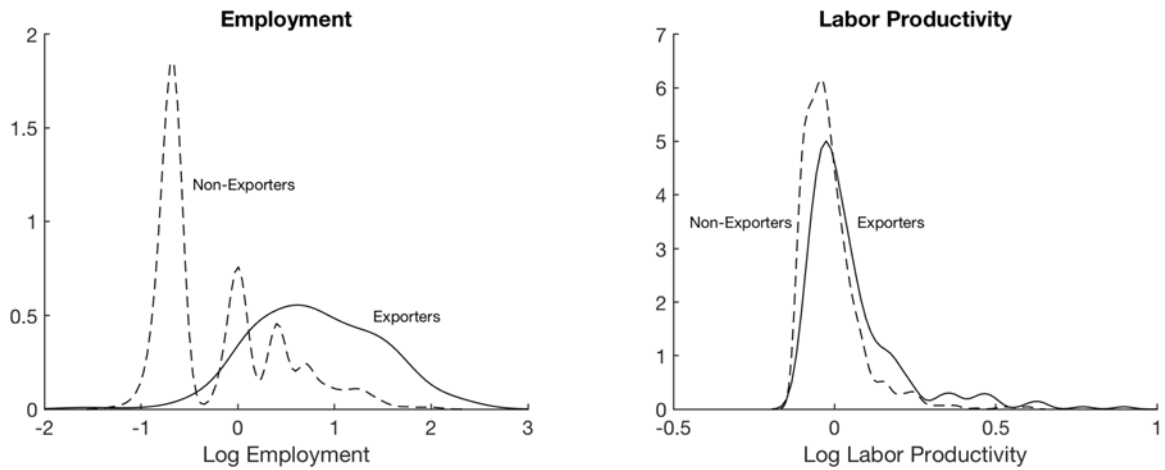


Note: The left panel shows the job creation rate due to exports and the right panel the job destruction rate due to imports in the steady state for different values of τ . All other parameters of model are kept fixed at the values in Table 9.

The model predicts an increase in job creation *from exports* (fact 7) in the aftermath of trade liberalization. In Figure 4 (left panel), the job creation rate from exports increases by 10 percentage points when trade costs fall from our benchmark value ($\tau = 1.491$) to completely free trade ($\tau = 1$). The right panel in Figure 4, meanwhile, plots the predicted job destruction rate from domestic firms who are replaced by imports. The model predicts that moving to frictionless trade would raise the overall job destruction rate by 10 percentage points. Unfortunately we do not have an empirical counterpart for this model statistic.

The left panel in Figure 5 plots the distribution of employment for exporters and non-exporters in the steady state, using our baseline parameter values from Table 9. These are the model analogues to facts 8 and 9. Firm size is determined by the number of products the firm controls, and whether the quality of the product is high enough to overcome the trade friction. A firm that exports has at least one product whose quality is high enough to overcome the trade cost. In the model, this probability is increasing in the firm's number of products. Larger firms own more products, and firms with more products are more likely to have

Figure 5: Simulated Distribution of Employment and Labor Productivity



Note: The distributions of employment (left panel) and labor productivity (right panel) in the steady state of the model with the parameters given in Table 9.

at least one product with sufficient quality to export. The gap in average size between exporters and non-exporters is not due to any fixed cost of exporting, but rather the difference in the number of products between the two groups of firms. Consistent with the empirical distribution of employment in Figure 1, the model predicts substantial overlap in the distribution of firm size of exporters and non-exporters. Although firms with fewer products are less likely to export, some of these products are high enough quality to overcome the trade cost.

The model also predicts that labor productivity is higher, on average, among exporters than among non-exporters. In the model, dispersion in labor productivity is entirely driven by markup heterogeneity. Since the markup is given by the quality gap between the best and the second best blueprint (adjusting for wages and tariffs), this gap is increasing in the quality of the best blueprint. Because a firm with high quality varieties is also more likely to export, such firms are also more likely to charge higher markups. The gap in average labor productivity between exporters and non-exporters reflects the gap in average quality between the two groups of firms. This is similar to Bernard et al. (2003).

A key parameter that governs the gap in average quality between exporters and non-exporters (and quality dispersion more generally) is the shape parameter θ of the Pareto distribution of innovation draws. We therefore calibrate this parameter to make the simulated model match the gap in average labor productivity between exporters and non-exporters in the U.S. data.²¹

Figure 5 (right panel) shows the simulated distribution of labor productivity between exporters and non-exporters in the model. While labor productivity (and thus quality when viewed through the lens of our model) of exporters is higher on average than for non-exporters, the model also generates a substantial overlap in labor productivity between the two groups of firms. The model generates such an overlap because the quality of many non-exporters can be large relative to the closest local competitor but may still not be large enough relative to the foreign firm after accounting for the trade cost and wage gap.

The empirical *dispersion* of employment and labor productivity (Figure 1) is substantially larger than in the simulated data (Figure 5). Our assumption that preferences over varieties is Cobb-Douglas implies that product quality only matters for employment when higher quality enables the firm to overcome the trade barrier. Conditional on selling in a given market at a given price, product quality has no effect on employment. We could make product quality matter more for firm employment, and thus get more employment dispersion, if we relaxed the Cobb-Douglas assumption. As for the dispersion of labor productivity, our model abstracts from differences in factor costs within countries, overhead costs, adjustment costs, and measurement error, all of which are likely present in the data and behind some of the empirical dispersion in labor productivity.²²

²¹As stated in Table 8, this gap is 6.6%.

²²Bartelsman, Haltiwanger and Scarpetta (2013) emphasize the role of overhead costs, Asker, Collard-Wexler and De Loecker (2014) the importance of adjustment costs, and Bils, Klenow and Ruane (2019) the contribution of measurement error.

4 Endogenous Innovation

A key assumption we have made so far is that innovation rates are exogenous parameters. We now consider the effect of reducing trade costs (i.e., tariffs) when innovation rates are endogenously determined.²³

Suppose the innovation rate (per variety owned) of a domestic incumbent is

$$\lambda = \left(\frac{R_i}{\gamma \chi_i \bar{A}^{(1-\phi)/\gamma}} \right)^\gamma, \quad (7)$$

where R_i denotes labor used for research (per variety owned), \bar{A} is the geometric average quality of products *sold* in the domestic market, χ_i is an efficiency parameter, $\gamma < 1$ captures the returns to research effort, and ϕ captures the external returns to the stock of ideas. As in Klette and Kortum (2004), underlying (7) is the assumption of constant returns at the firm level to research effort and the number of varieties the firm owns (i.e., elasticities of γ and $1 - \gamma$, respectively). When $\phi < 1$ we have diminishing returns to the stock of ideas so growth is semi-endogenous and linked to the population growth rate as in Jones (1995).

Similarly, suppose the unconditional innovation rate of domestic entrants is

$$\tilde{\eta} = \left(\frac{R_e}{\gamma \chi_e \bar{A}^{(1-\phi)/\gamma}} \right)^\gamma, \quad (8)$$

where R_e is labor used for research (per variety in the economy) by potential entrants and χ_e is an efficiency parameter.²⁴

The return to innovation is the product of the probability of grabbing a variety from another firm and the expected value of that variety. The new product can either be sold in both markets or only in the domestic market, and the value of this new product depends on whether it is traded or non-traded. So

²³Section A.4 of our Online Appendix provides a fuller description of this model.

²⁴The innovation rates for foreign firms are given by equations analogous to (7) and (8) with R_i and χ_i replaced by R_i^* and χ_i^* in (7), R_e and χ_e replaced by R_e^* and χ_e^* in (8), and average quality sold in the foreign market instead of in the home market.

the return to innovation is the sum of the expected value of a traded product and the expected value of a non-traded product (multiplied by the probability of getting each type of product).

It will be convenient to normalize the value of a product by $\bar{A}^{(1-\phi)/\gamma}$. We define v_x and v_n as the *expected* normalized value of a traded and non-traded product. The following arbitrage equation pins down v_x at time t :

$$\begin{aligned} v_{x,t} = & \pi_{x,t} - \gamma \chi_i \lambda_t^{\frac{1}{\gamma}} \\ & + \frac{(1+g_t)^{(1-\phi)/\gamma}}{1+r_t} [\lambda_t (\beta_{x,t} v_{x,t+1} + \beta_{n,t} v_{n,t+1})] \\ & + \frac{(1+g_t)^{(1-\phi)/\gamma}}{1+r_t} [(1-\delta_{x,t}) v_{x,t+1} - \delta'_{x,t} (v_{x,t+1} - v_{n,t+1})]. \end{aligned} \quad (9)$$

Here g denotes the growth rate of \bar{A} , r the interest rate, π_x expected profits (normalized by $\bar{A}^{(1-\phi)/\gamma}$), β_x and β_n the probability conditional on innovating of grabbing a traded and non-traded product, δ_x the probability of losing a traded variety in both markets, and δ'_x the probability of losing a traded product only in the foreign market.²⁵ The first term in (9) is profit net of research expenses, the second term is the expected value of grabbing a new product, and the last term is the expected value of an exported variety (next period) adjusted for the probability of losing the product to a competitor in the foreign market.²⁶

²⁵The conditional arrival rates are $\beta_x \equiv x + (x^* - x) \min \left[\left(\frac{A_j}{\omega \tau A_j^*} \right)^\theta, 1 \right] + (1 - x^*) \min \left[\left(\frac{1}{\omega \tau} \right)^\theta, 1 \right]$ and $\beta_n \equiv (x^* - x) \left(1 - \min \left[\left(\frac{A_j}{\omega \tau A_j^*} \right)^\theta, 1 \right] \right) + (1 - x^*) \left(\min \left[\left(\frac{\tau}{\omega} \right)^\theta, 1 \right] - \min \left[\left(\frac{1}{\omega \tau} \right)^\theta, 1 \right] \right)$. The probabilities of losing an exported product are $\delta_x \equiv (\lambda + \tilde{\eta}) + (\tilde{\lambda}^* + \tilde{\eta}^*) \min \left[\left(\frac{\omega}{\tau} \right)^\theta, 1 \right]$ and $\delta'_x \equiv (\tilde{\lambda}^* + \tilde{\eta}^*) \left(\min \left[(\omega \tau)^\theta, 1 \right] - \min \left[\left(\frac{\omega}{\tau} \right)^\theta, 1 \right] \right)$.

²⁶We assume linear utility so that the consumption Euler equation implies $r = \rho$. We set $\rho = 0.05$. Given linear utility and the same discount rate as real interest rate, consumers are indifferent about the path of consumption for a given present discounted value of consumption. We assume the consumption path in each country ensures balanced trade.

The corresponding arbitrage equation for the non-traded product v_n is

$$\begin{aligned}
v_{n,t} &= \pi_{n,t} - \gamma \chi_i \lambda_t^{\frac{1}{\gamma}} \\
&+ \frac{(1+g_t)^{(1-\phi)/\gamma}}{1+r_t} [\lambda_t (\beta_{x,t} v_{x,t+1} + \beta_{n,t} v_{n,t+1})] \\
&+ \frac{(1+g_t)^{(1-\phi)/\gamma}}{1+r_t} [(1-\delta_{n,t}) v_{n,t+1}].
\end{aligned} \tag{10}$$

where π_n denotes expected profits (normalized by $\bar{A}^{(1-\phi)/\gamma}$) of a non-traded product and δ_n is the probability a non-traded product is overtaken by another firm.²⁷

The privately optimal innovation rates are given by equating the marginal revenue from innovation to the marginal cost of innovation, which yields:

$$\lambda_t = \left(\frac{\beta_{x,t} v_{x,t} + \beta_{n,t} v_{n,t}}{\chi_i} \right)^{\frac{\gamma}{1-\gamma}} \tag{11}$$

$$\tilde{\eta}_t = \left(\frac{\beta_{x,t} v_{x,t} + \beta_{n,t} v_{x,t}}{\chi_e} \right)^{\frac{\gamma}{1-\gamma}}. \tag{12}$$

An increase in v_x and v_n raises the innovation rate with an elasticity that depends on γ . As in the model where innovation is exogenous, the equilibrium is determined by equations (2), (3), (4), and the markup formulas in Table 5, except that the innovation rates are now pinned down by (9) through (12) plus the analogues for foreign innovation.

In steady-state, β_x , β_n , v_x , and v_n are constant so the innovation rates λ and η are constant as well. Differences in innovation rates between countries now reflect differences in the innovation cost parameter χ . As before, differences in innovation rates show up as differences in the relative wage. What is new in the endogenous innovation model is that the growth rate of the real wage in steady state is ultimately given by the product of the population growth rate and $\gamma/(1-\phi)$ where $\phi < 1$.

²⁷This probability is given by $\delta_n \equiv (\lambda + \tilde{\eta}) + (\tilde{\lambda}^* + \tilde{\eta}^*) \min \left[\left(\frac{\omega A_j^*}{\tau A_j} \right)^\theta, 1 \right]$

Table 11: Estimates of Model Parameters under Endogenous Innovation

Variable	Description	Value
ϕ	Return to stock of ideas	0.120
γ	Return to research intensity	0.643
χ_e/χ_i	Home entrant/incumbent research cost	2.52
χ_i^*/χ_i	Foreign/home incumbent research cost	7.88
χ_e^*/χ_i	Foreign entrant/home incumbent research cost	16.5

We set ϕ and γ to match the growth rate of TFP, the growth rate of the “population” (growth of investments in intellectual property products), and the share of “labor” devoted to research (share of value added invested in intellectual property products).²⁸ This yields $\phi = 0.120$ and $\gamma = 0.643$. With these values in hand, we can back out the innovation cost parameters χ consistent with the arrival rates in Table 9. As shown in Table 11, the implied innovation cost parameter χ is lower in the U.S. to generate higher U.S. innovation rates and match the higher wage in the U.S. relative to the rest of the OECD.

We can now re-examine the effect of reducing trade costs, this time with endogenous arrival rates of innovation. Here, trade liberalization has no permanent effect on the long run growth rate, which is pinned down by population growth. The initial increase in the growth rate due to trade liberalization increases the level of TFP, but with $\phi < 1$ this raises the cost of innovation. Thus, in the new steady state with lower trade costs, innovation rates by each country are actually *lower* but the growth rate of TFP is the same.

Figure 6 plots the effect of a permanent, unanticipated 50% reduction of trade costs from the baseline (from $\tau = 1.491$ to $\tau = 1.245$) on innovation rates,

²⁸TFP growth averaged 3.01% per year from 1995–2008 in manufacturing according to the U.S. Bureau of Labor Statistics KLEMS data. Real intellectual property investments grew 4.12% per year from 1995–2008 according to the U.S. Bureau of Economic Analysis, and such investments averaged 9.72% of value added in U.S. manufacturing from 1997–2008.

research labor shares, real consumption, and job flows. It shows that innovation rates initially rise (fall) in the OECD (U.S.) in the aftermath of a reduction in trade costs. But as the level of TFP rises in the two countries, arrival rates fall due to the rising difficulty of innovation. This is due to diminishing returns to the stock of ideas $\phi < 1$. In the new steady state, innovation rates are *lower* compared to the initial steady state, though TFP is on a higher path (parallel to its initial path). TFP is higher despite lower arrival rates within each country because ideas flow more easily across countries with lower trade costs.

Like the arrival rates, the share of labor doing research spikes on impact. Unlike the arrival rates, however, the share of labor doing research ends up higher in the long run. The bigger market for each successful innovation makes higher research effort worthwhile despite the endogenously greater difficulty in coming up with ideas in the long run.²⁹ This result contrasts with Eaton and Kortum (2001), wherein these two forces exactly offset each other and leave research effort unchanged.

Consumption is closely tied to TFP in each country, but is also affected by tariff revenues and the share of labor diverted from production to R&D. Figure 6 plots real consumption in each country relative to the path of U.S. consumption in the absence of trade liberalization. Consumption rises on impact in the U.S. but initially falls in the OECD due to the temporary spike in the share of workers engaged in innovation rather than production. After the initial shock, the growth rate of consumption increases for a few decades, but eventually slows down as innovation becomes more difficult with higher TFP. On the new steady state path, real consumption and TFP are higher, but the growth rate is the same as in the initial equilibrium. As in the exogenous innovation case, the rest of the OECD gains more than the U.S. because the U.S. is more innovative.

Figure 6 also illustrates the effect of the reduction in trade costs on job flows. Consistent with the evidence from the U.S. and Canada after CUSFTA, job flows

²⁹Similarly, the higher share of labor devoted to research in the rest of the OECD compared to the U.S. in Figure 6 reflects the larger size of the rest of the OECD's market.

Table 12: Gains From Reducing Trade Costs

	50% Reduction in τ		Relative to Autarky	
	U.S.	OECD	U.S.	OECD
Static Gains according to the ACR formula	3.0%	1.2%	1.1%	1.0%
Static Gains in our Model	4.2%	5.0%	13.3%	15.5%
Dynamic Gains - Exogenous Innovation	9.6%	15.5%	37.0%	104.0%
Dynamic Gains - Endogenous Innovation	7.4%	8.9%	36.9%	69.1%

Entries give the percentage increase in current year consumption (static) or in the present discounted value of consumption (dynamic) as a result of reducing tariffs from 1.491 to 1.245 (columns one and two) and from 4 to 1.491 (columns three and four). The aggregate trade share at $\tau = 4$ is about 0.4%. We use a discount rate of 5%.

surge in the aftermath of a reduction in trade costs.³⁰ The pace of job flows remains elevated for decades after a tariff reduction — certainly within the 15-year window we examine after the 1988 U.S.-Canada Free Trade Agreement. This pattern drives home that there might be dynamics costs (job destruction) as well as dynamic benefits (knowledge flows) to trade liberalization.³¹

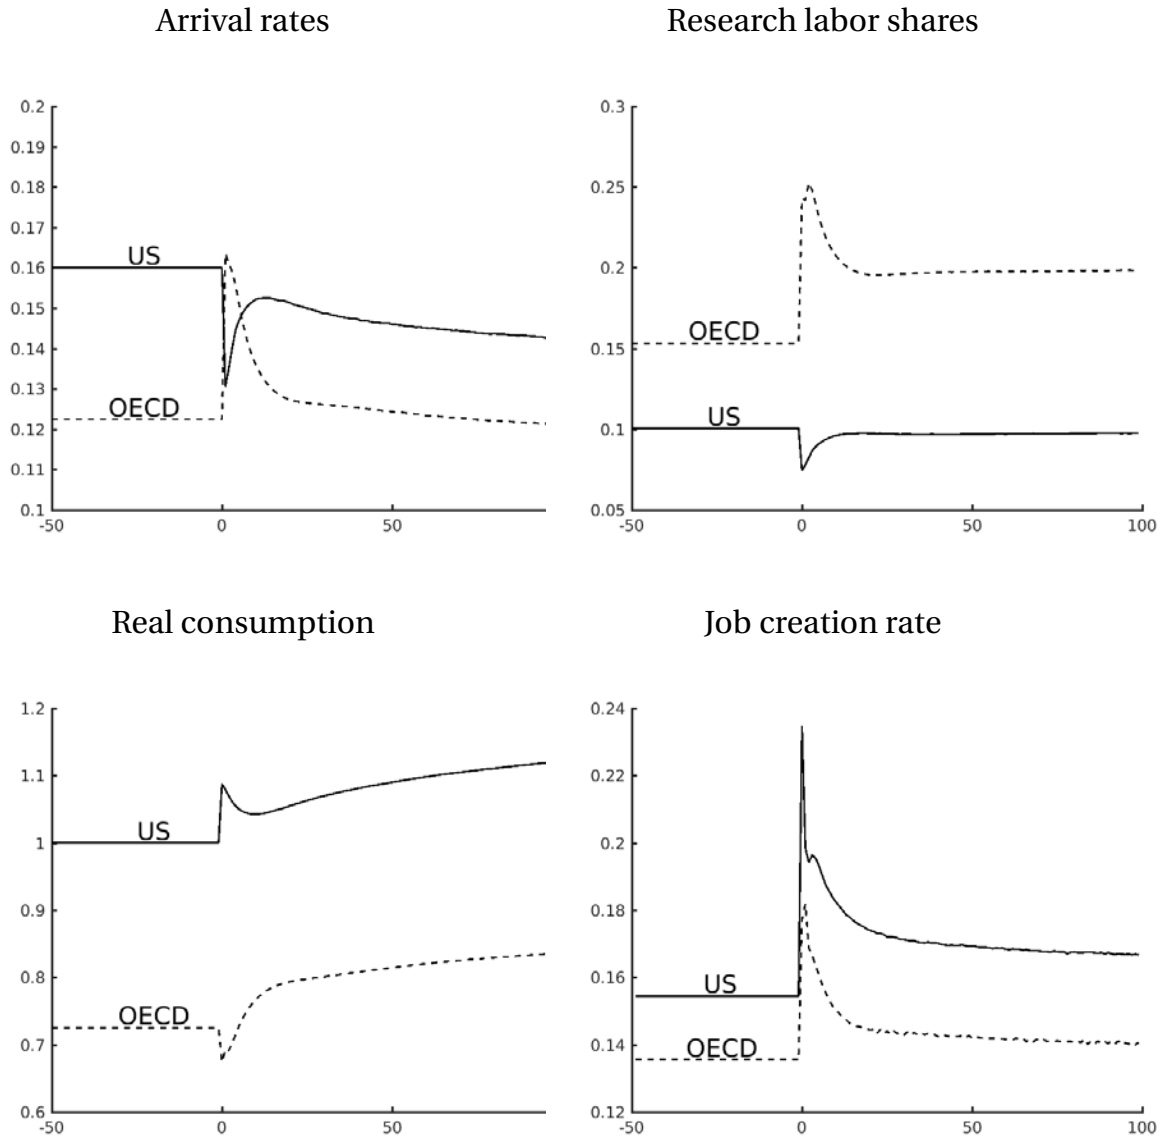
Table 12 (columns 1 and 2) displays the welfare gains from a 50% reduction in tariffs from the baseline. These are in permanent consumption-equivalent terms, which is equivalent to percentage gains in the present discounted value of consumption given our specification of linear utility. For comparison, we start with the static gains implied by the formula in Arkolakis, Costinot and Rodriguez-Clare (2012). The static gains are higher in our model than in the ACR formula, but the discrepancy is modest.

The final two rows of Table 12 present the gains from trade *including* effects on innovation. With exogenous arrival rates of innovation, the gains are more than three times as large as the ACR gains at 9.6% in the U.S. The rest

³⁰Bernard, Redding and Schott (2007) analyze a multi-sector Melitz model which also features effects of trade liberalization on steady state job flows.

³¹In our model job destruction is costless. Relaxing this assumption to weigh such dynamic costs against any dynamic gains is a useful direction for future research.

Figure 6: Simulated Effect of a 50% Reduction in Trade Costs



Note: The figure shows the simulated arrival rates ($\lambda + \tilde{\eta}$ and $\tilde{\lambda}^* + \tilde{\eta}^*$), research labor shares, real consumption (relative to the path of real consumption in the initial steady state), and job creation and destruction rates in response to an unanticipated reduction in trade costs from $\tau = 1.491$ to $\tau = 1.245$, keeping constant all other parameters at the values in Table 11.

of the OECD gains even more (15.5%) because it gets more ideas than it gives. Gains are smaller with endogenous innovation rates, but still sizable at about 7.4% in the U.S. and 8.9% in the rest of the OECD. The gains are tempered with endogenous research effort for two reasons. First, we build in diminishing returns to the stock of ideas ($\phi < 1$). This reduces the cumulative TFP gain from endogenously rising research effort and from a faster flow of ideas across countries. Second, higher research effort comes at the cost of less labor devoted to producing goods. Even in the endogenous innovation case, however, the full gains are seven times larger than the static ACR formula gains.

The gap between the static gains implied by the ACR formula and our calculations are much larger when we consider more dramatic changes in trade costs. In columns 3 and 4 in Table 12 we present the welfare gains from trade when the steady state trade share goes from 0.4% near autarky ($\tau = 4$) to 10.2% in our baseline (with $\tau = 1.491$). The ACR formula predicts a welfare gain of 1.1% for the U.S.³² As the Table shows, the static gains in our model are *much* larger than what the ACR formula implies: over 13% for the U.S.

Clearly, our model does not fall into the ACR class. In our model, trade facilitates the flow of ideas across countries. Product quality, and the comparative advantage gains from trade, varies with trade costs. Figure 7 plots the distribution of relative quality across products for the U.S. versus the rest of the OECD. The solid distribution is under our baseline ($\tau = 1.49$) and the dashed distribution is near autarky ($\tau = 4$). The relative quality distribution is markedly more dispersed near autarky. As a result, the trade elasticity is only 3 near autarky, whereas it is 5 under baseline tariff of $\tau = 1.491$.³³ The degree of U.S. versus OECD comparative advantage across products is endogenously stronger near autarky. This is because ideas are not flowing as quickly between the countries when there is so little trade, so relative qualities drift apart. If we use a trade elasticity of 3 rather than 5, the ACR gains more than triple from 1.1% to 3.5%.

³²The ACR formula for the 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).

³³We calculate local trade elasticities, varying the tariff rate by 10 percentage points.

Figure 7 also shows that the effect of trade liberalization on relative quality dispersion is smaller for small changes in trade. This can be seen by comparing the dotted distribution ($\tau = 1.245$) with the dashed distribution ($\tau = 1.491$) in the Figure. Small changes in trade costs have only a modest effect on the comparative advantage gains from trade in our model. As a consequence, the ACR gains from cutting tariffs in half are closer to the static gains in our model.


When going from near autarky to $\tau = 1.49$, the trade share initially leaps from 0.4% to 24.6%. The trade share on impact overshoots the new steady state trade share of 10.2% precisely because of divergent 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 in response to higher trade flows, the trade share eventually settles down to 10.2% and the trade elasticity rises from 3 to 5.

The remaining gap between ACR gains of 9.8% and our model static gains of 13.3% for the U.S. may be due to changes in markups. Markup dispersion within countries creates misallocation, and markup differences across countries affect the terms of trade. This is another way in which our model falls outside the ACR class: the distribution of markups is endogenous to tariffs.

Figure 8 summarizes the implication of trade costs on real consumption. The figure shows real consumption (relative to real consumption for the U.S. at the baseline of $\tau = 1.491$) and the trade elasticity in steady-states for a wide range of tariffs. The consumption path is almost always higher as a result of lower tariffs. Consumption actually falls, however, as frictionless trade is approached. At very low tariff levels, the high rate of creative destruction from imports discourages research effort so much that it outweighs the more rapid spread of ideas.

Figure 9 plots the local trade elasticity calculated at the steady-states at various trade costs. It shows that the trade elasticity increases with the *level* of trade costs. This is an important prediction of a model with idea flows, although we currently do not have evidence on whether this prediction holds empirically.

Figure 7: Relative quality distribution



Note: The figure shows the distribution across products of relative quality in the U.S. versus the rest of the OECD for steady states with a U.S. export share of 0.4%, 10%, and 22%.

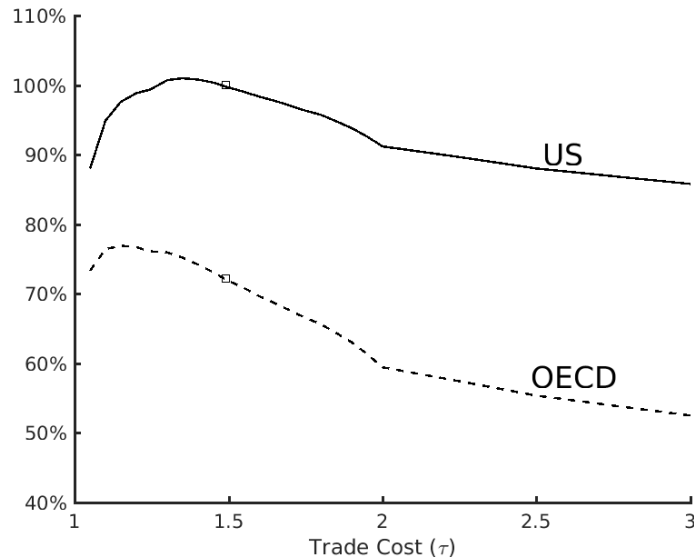
5 Models with Limited Idea Flows

Our simulations make two key assumptions about the generation of new ideas. First, we assume innovators build on the productivity level of *sellers* into the domestic market. Second, we assume innovators attempt to build on the productivity of *all* products sold in the domestic market. To gauge the importance of these assumptions for the dynamic gains from trade, we now consider two alternative assumptions about how new ideas are generated. For simplicity, these simulations involve exogenous arrival rates of innovation.

Our first alternative assumes that innovators probabilistically build on *sellers* with probability κ and on domestic *producers* with probability $1 - \kappa$.³⁴ Our second alternative assumes that innovators build on all products with probability ν and on the subset of products that are domestically produced with proba-

³⁴If a product is imported, innovators build on the quality of the last domestic producer.

Figure 8: Real consumption vs. trade costs



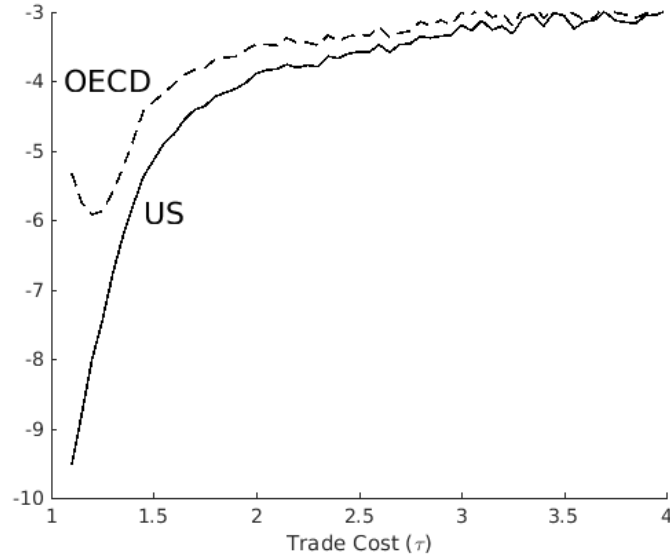
Note: The figure plots simulated equilibrium real consumption relative to the consumption path in the U.S. with $\tau = 1.491$ in steady-states with differing trade costs, keeping constant all other parameters at the values in Table 11.

bility $1 - \nu$. Such research specialization ($\nu < 1$) allows countries to experience more frequent innovations on the subset of products they produce, and more so the higher the share of products imported.

The top panel in Table 13 shows the gains from trade under these alternative assumptions about idea flows.³⁵ The first column reproduces our baseline specification ($\kappa = 1$ and $\nu = 1$). The second column reduces the frequency with which innovators learn from sellers to 30% ($\kappa = 0.3$) but does not allow research specialization ($\nu = 1$). Doing so dampens the effect of trade costs on growth rates from 47 basis points per year in the baseline to 13 basis points per year with limited idea flows. The consumption-equivalent gains from trade in the U.S. fall from 37% in the baseline to 27%, and in the OECD fall from 104% to

³⁵Specifically, Table 13 calculates the effect of changing τ to move from a steady state with a trade share of 0.4% to a trade share of 10.2% (keeping unchanged the other parameters).

Figure 9: Trade elasticity vs. trade costs



Note: The figure plots the trade elasticity in steady-states with differing trade costs, keeping constant all other parameters at the values in Table 9. The local trade elasticity in each steady state is calculated by varying trade costs by 10%.

48%.³⁶ Thus, limiting the flow of ideas across countries materially reduces the dynamic gains from trade.

The final column of Table 13 adds research specialization to limited idea flows ($\kappa = 0.3$ and $\nu = 0.3$). The effect of trade on steady-state growth (+48 basis points per year) is actually larger relative to that in the middle column (13 basis points), and is similar to that in the baseline specification (47 basis points). The U.S. gains slightly more in consumption-equivalent terms in this case, 40% versus 37% in the baseline. Thus, even if ideas flow in a more limited way, trade may still create large dynamic gains if it facilitates research specialization. The consumption-equivalent gains from trade are smaller in the OECD because it gains less from specialization (it is larger) and because it suffers more from

³⁶Recall that the rest of the OECD gains more from idea flows than the U.S. because it is less innovative in our calibration. It therefore suffers more from limiting idea flows.

Table 13: Alternative Assumptions on Idea Flows

	Data	$\kappa = 1$ $\nu = 1$	$\kappa = 0.3$ $\nu = 1$	$\kappa = 0.3$ $\nu = 0.3$
Gains from Trade (Trade share from 0.4% to 10%)				
Δ in annual Growth	–	0.47%	0.13%	0.48%
Gains from Trade U.S.	–	37.0%	27.3%	40.3%
Gains from Trade OECD	–	104.0%	47.8%	33.1%
Export Product Churn U.S.				
Prob. losing exports	15.1%	12.8%	5.9%	1.1%
Top Export Product Turnover	66.5%	77.5%	87.9%	95.4%

Note: The top panel shows the gains from trade (consumption-equivalent welfare) from changing τ to move from a steady state with a trade share of 0.4% to a trade share of 10%. The bottom panel shows the simulated export churn in a steady state with a trade share of 10%. Column two is the baseline specification. Column three reduces the frequency with which innovators learn from foreign producers to 30% ($\kappa = 0.3$); the other 70% of draws build on domestic producers. Column four allows for research specialization, so that 70% of innovation ($\nu = 0.3$) draws are in products a country produces; the other 30% draw from all products. In each case, innovation parameters are re-estimated to match the targeted moments in Table 8 while θ and ψ are held fixed at the values shown in Table 9. Arrival rates are exogenous.

reduced idea inflows given U.S. innovativeness.

How can we assess the realism of these alternative models of idea flows?³⁷ We first examine the implication of the models for export product turnover. Recall one of our facts is the high turnover rate of a country's exported products. The bottom panel in Table 13 gives our two measures of export product turnover. The model with limited idea flows and no research specialization ($\kappa = 0.3$ and $\nu = 1$) generates less export turnover compared to the data (and the baseline model): the predicted exit rate of an average product drops to 6%

³⁷We do not entertain a model in which trade fails to tether growth rates across countries because such a model would be inconsistent with the similarity of observed growth rates across OECD countries (Klenow and Rodriguez-Clare, 2005).


per year and the SITC with the highest level of exports in a given year tends to account for 87.9% of total exports five years earlier. In our baseline simulation, these two numbers are 12.8% and 77.5%, respectively. The last column shows that adding specialized research ($\nu < 1$) lowers export product turnover more dramatically. The exit rate of an average product is only 1.1% per year and the export share of the top product five years earlier is 95.4% of its current share.

In the baseline model with idea flows across countries, the trade elasticity is sensitive to the level of trade barriers. This is not the case when flows of ideas between countries are limited. Figure 10 plots the distribution of quality of the U.S. versus the rest of the OECD in the model with limited idea flows and research specialization. The figure plots the distribution of relative quality in the steady-state with a trade share of 10.2% and the near-autarky case with a trade share of 0.4%. In contrast to the model with idea flows (Figure 7), the dispersion of relative quality is about the same in the two steady-states shown in Figure 10. As a consequence, the trade elasticity will be about the same with high trade costs as with low trade costs.

The change in jobs flows in the aftermath of CUSFTA (facts 5-7) is also relevant for assessing whether trade facilitates idea flows. In Section 2, we emphasized that job flows, job destruction at large firms, and job creation associated with exports all increased noticeably in Canada after the agreement. Here we simulate the impact of CUSFTA on job flows.³⁸ We analyze the effect of lowering the gross tariff rate from 1.39 to 1.25 to match Canada's trade shares of 25% before CUSFTA (1978–1988) and 37% after CUSFTA (1988–2003). We do this for three cases: in our baseline model ($\kappa = 1$ and $\nu = 1$); under limited idea flows ($\kappa = 0.3$ and $\nu = 1$); and with both limited idea flows and research specialization ($\kappa = 0.3$ and $\nu = 0.3$). We re-estimate other parameters in each case so that the model matches the empirical wage in the U.S. relative to Canada. We set the population in the U.S. relative to Canada to match the data.

³⁸Until now we carried out simulations with the rest of the OECD because its larger size drives home the importance of scale for the dynamic gains from trade.

Figure 10: Relative quality distribution with limited idea flows



Note: The figure shows the distribution across products of relative quality in the U.S. versus the rest of the OECD for steady states with a trade share of 0.4% and 10% in the model with limited idea flows and research specialization.

Table 14 presents the predicted change in job flows in Canada after trade liberalization in the three models.³⁹ The first column shows the change in empirical moments and the second column shows the change in the simulated moments in the base case ($\kappa = 1$ and $\nu = 1$). The base case model matches the patterns in the data: rising job creation, job destruction, and job creation from exports. The model overstates job creation but well matches both job destruction and job creation from exports. Note that no parameter values were chosen to target any of the data moments in Table 14.

The last two columns in Table 14 show the simulated change in job flows with limited idea flows (column 3) and with research specialization (column 4). The increase in job flows predicted by these two models generally understates

³⁹We focus on changes in Canada rather than the U.S. because CUSFTA was a much bigger shock to the Canadian economy given its large trade share with the U.S.

Table 14: Job Flows in Canada — Post-CUSFTA versus Pre-CUSFTA

	Data	$\kappa = 1$ $\nu = 1$	$\kappa = 0.3$ $\nu = 1$	$\kappa = 0.3$ $\nu = 0.3$
Δ Job Creation Rate	+1.3%	+6.7%	+4.0%	+4.3%
Δ Job Destruction Rate	+7.2%	+6.7%	+4.0%	+4.3%
Δ Job Creation from Exports	+8.7%	+8.7%	+5.8%	+6.1%

Note: Pre-CUSFTA is 1978 to 1988. Post-CUSFTA is 1988 to 2003. Job creation and destruction are calculated over five year periods. Simulations use a version of the model estimated to match the relative wage and population of the U.S. and Canada, with a trade liberalization event that matches the 12 percentage point increase in Canadian export share from the pre-CUSFTA to post-CUSFTA period. Arrival rates are exogenous.

the increase in job flows in Canada after CUSFTA. The increase in the job destruction rate is slightly above half of the increase seen in the data. The gap would widen for lower values of κ and ν .

One could also entertain the possibility that ideas flow in ways unrelated to trade altogether. Ramondo, Rodríguez-Clare and Saborío-Rodríguez (2016) argue that ideas flowing independent of trade (or even FDI) would help explain why small countries do not tend to be poorer than large countries. It could also contribute to export product turnover. It could not, however, explain why job flows increased in a sustained way in Canada after CUSFTA.

6 Conclusion

We documented facts about trade and job reallocation in U.S. and Canadian manufacturing in recent decades. After the U.S.-Canada Free Trade Agreement in 1988, job destruction rates spiked and remained elevated through 2012 (our latest year of data). The increase in job destruction rates and exit from exporting occurred equally at big and small firms in Canada.

Motivated by these facts, we constructed a two-country model of creative destruction and trade. In the model, foreign and domestic firms take over each other's markets more frequently when trade barriers are lower. This stimulates growth in the long run under exogenous innovation rates. When we endogenize innovation and build in diminishing returns, lower tariffs boost growth only temporarily. Still, trade liberalization raises levels of productivity permanently. Compared to (near) autarky, such dynamic gains are an order of magnitude larger than the usual static gains from trade.

We see several directions for future research. One is to explicitly incorporate frictions to reallocating workers in response to trade-induced creative destruction. Another route is to study events such as China joining the WTO. 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 think idea flows will need to be tied to trade flows to explain why trade liberalization ushers in more rapid job reallocation in a sustained way.

We end with a conjecture about 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 externalities, just as we need Global Climate Change agreements to internalize negative pollution externalities.

References

- Aghion, Philippe and Peter Howitt, "A Model of Growth Through Creative Destruction," *Econometrica*, 1992, 60 (2), pp. 323–351.
- , Antonin Bergeaud, Matthieu Lequien, and Marc Melitz, "The Impact of Exports on Innovation: Theory and Evidence," 2018.
- Akcigit, Ufuk, Sina T. Ates, and Giammario Impullitti, "Innovation and Trade Policy in a Globalized World," 2018.
- Alvarez, Fernando E, Francisco J. Buera, and Robert E. Lucas, "Idea Flows, Economic Growth, and Trade," 2013.
- Arkolakis, Costas, Arnaud Costinot, and Andres Rodriguez-Clare, "New Trade Models, Same Old Gains?," *American Economic Review*, 2012, 102 (1), 94–130.
- Asker, John, Allan Collard-Wexler, and Jan De Loecker, "Dynamic Inputs and Resource (Mis)Allocation," *Journal of Political Economy*, 2014, 122 (5), 1013–1063.
- Bartelsman, Eric, John Haltiwanger, and Stefano Scarpetta, "Cross-Country Differences in Productivity: The Role of Allocation and Selection," *American Economic Review*, 2013, 103 (1), 305–334.
- 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.
- Bernard, Andrew B, Stephen J. Redding, and Peter K. Schott, "Comparative Advantage and Heterogeneous Firms," *Review of Economic Studies*, 2007, 74 (1), 31–66.
- Bils, Mark, Peter J. Klenow, and Cian Ruane, "Misallocation or Mismeasurement?," 2019.
- Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb, "Are Ideas Getting Harder to Find?," 2019.

- , 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,” 2017.
- 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 Books*, 1996.
- Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda, “The Role of Entrepreneurship in U.S. Job Creation and Economic Dynamism,” *Journal of Economic Perspectives*, 2014, 28 (3), 3–24.
- 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), 823–39.
- Eaton, Jonathan and Samuel Kortum, “Technology, Trade, and Growth: A Unified Framework,” *European Economic Review*, 2001, 45 (4-6), 742–755.
- and —, “Technology, Geography, and Trade,” *Econometrica*, 2002, 70 (5), 1741–1779.
- Feenstra, Robert C., Robert E. Lipsey, Haiyan Deng, Alyson C. Ma, and Hengyong Mo, “World Trade Flows: 1962-2000,” *NBER Working Paper 11040*, 2005.
- Garcia-Macia, Daniel, Chang-Tai Hsieh, and Peter J. Klenow, “How Destructive is Innovation?,” *Econometrica*, 2019.
- 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.
- Holmes, Thomas J. and John J. Stevens, “An Alternative Theory of the Plant Size Distribution, with Geography and Intra- and International trade,” *Journal of Political Economy*, 2014, 122 (2), 369–421.
- Jones, Charles I., “R&D-Based Models of Economic Growth,” *Journal of Political Economy*, August 1995, 103 (4), 759–784.
- Klenow, Peter J. and Andres Rodriguez-Clare, “Externalities and Growth,” *Handbook of Economic Growth*, 2005, 1, 817–861.
- Klette, Tor Jakob and Samuel Kortum, “Innovating Firms and Aggregate Innovation,” *Journal of Political Economy*, 2004, 112 (5), 986–1018.
- Lincoln, William, Andrew McCallum, and Michael Siemer, “The Great Recession and a Missing Generation of Exporters,” *IMF Economic Review*, 2019.
- 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,” 2019.
- 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.