

A Global View of Creative Destruction

Chang-Tai Hsieh

University of Chicago and NBER

Peter J. Klenow

Stanford University and NBER

Ishan Nath *

University of Chicago

June 4, 2018

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 was concentrated in industries with bigger tariff cuts, and 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 and thereby speeds 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.

*We thank Beiling Yan of Statistics Canada for her assistance with the Canadian manufacturing data and Eric English for excellent research assistance. 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 U.S. Census Bureau and Statistics Canada to ensure that no confidential information is disclosed. Contact information: chsieh@chicagobooth.edu, klenow@stanford.edu, and inath@uchicago.edu.

1 Introduction

Landmark 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 job reallocation and growth, but no medium or long run effect on either.

A growing literature seeks to assess the dynamic costs (such as time consuming job reallocation) and the dynamic benefits of trade (such as faster growth). Empirical studies on the cost side include Autor, Dorn and Hanson (2013) and Dix-Carneiro and Kovak (2017), and on the benefit side include Bloom, Draca and Van Reenen (2016) and Aghion, Bergeaud, Lequien and Melitz (2017). Efforts at modeling the growth effects of trade build on the foundational models of Rivera-Batiz and Romer (1991) and Grossman and Helpman (1993). These include Alvarez, Buera and Lucas (2013), Perla, Tonetti and Waugh (2015), Buera and Oberfield (2016), and Akcigit, Ates and Impullitti (2017).

In this paper we present facts and a model focusing on the role of creative destruction in trade. We document the magnitude of job reallocation tied to exports in U.S. and Canadian manufacturing firms before vs. after the 1988 U.S.-Canada Free Trade Agreement. Industries which saw larger import tariff cuts experienced elevated job reallocation rates for decades after the agreement. Exit and job destruction rates rose for big firms as much as for 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 – both entrants and incumbents – draw from a Pareto distribution building on the technology of the firm selling in the domestic market. When innovators take over the local market for an existing product (creative destruction), job reallocation takes place. Domestic firms can also take over foreign markets for a product, as can foreign firms the domestic market.

The first version of the model features exogenous arrival rates of innovation as in Garcia-Macia et al. (2016). 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 (2017), so that growth is semi-endogenous and linked to growth in the number of researchers. In both models, the two countries grow at the same rate in the long run.

We calibrate the model to fit moments in U.S. manufacturing vs. manufacturing in the rest of the OECD.¹ We match TFP growth, growth in research investment, exports relative to manufacturing shipments, the job reallocation rate, and job reallocation tied to newly exporting plants and plants who cease exporting. To pin down the Pareto shape parameter we fit the gap in revenue per worker for exporters vs. non-exporters in U.S. manufacturing plants. We also target manufacturing value added per worker and employment of the U.S. vs the rest of the OECD. We infer higher innovation rates in the U.S. and the rest of the OECD, either exogenously or endogenously.

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 wage relative to the U.S. rises. Just as found in the wake of the 1988 U.S.-Canada Free Trade Agreement, lower tariffs lead to more job destruction. In the model, there is a spike immediately after tariffs are lowered, but job destruction remains higher in the new steady state than before tariffs were lowered.²

In the endogenous arrival rate 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

¹Later we will try alternative calibrations such as to the U.S. vs. Canada or China.

²Our model features full employment, but the evidence cited above makes clear that job reallocation is not so seamless in the real world.

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. (We infer the rest of the OECD is less innovative given their lower GDP per worker.) Job reallocation rises more in the rest of the OECD than in the U.S. after tariff reductions.

Our effort is most closely related to three recent papers. Perla, Tonetti and Waugh (2015) 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 theirs is the interaction of trade with domestic technology diffusion.

We follow Buera and Oberfield (2016) 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 entry and exit into exporting and job reallocation associated with creative destruction from trade. They stress that the dynamic effects of trade could be negative depending on whether firms learn from domestic producers or sellers.

Akcigit, Ates and Impullitti (2017) 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. Empirically, 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 creative destruction and its implications for growth and job reallocation.

The rest of the paper is organized as follows. Section 2 lays out eight 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. Section 5 concludes.

2 Facts from Canadian and U.S. Manufacturing

We use data from the U.S. Manufacturing Census and Canada's Annual Survey of Manufactures. The U.S. data is a census of all manufacturing establishments, including the smallest ones, every five years from 1973 to 2012. The Canadian data covers all but the smallest manufacturing establishments every year from 1973 to 2012.³ The variables we use from the two datasets are the plant and firm identifiers, employment, industry (four digit SIC or six digit NAICS), and exports. This information is available in every year, except for exports which is only available in the U.S. data starting in 1987 and in the Canadian data in 1974, 1979, 1984, 1989, 1993, and every year after 1996. We aggregate the establishment data to the firm level. We highlight eight facts:

1. **Large Job Flows.** Table 1 (rows 1 and 2) presents the job creation and destruction rates over five years for Canadian (from 1973 to 2012) and U.S. manufacturing (from 1987 to 2012).⁴ As in the classic work by Davis,

³The survey threshold is currently annual sales of 30 thousand Canadian dollars.

⁴The job creation rate between year t and $t + 5$ is defined as the ratio of the sum of employment of new firms (established between the two years) in year $t + 5$ and the change in employment among *expanding* firms between the two years to average total employment (in year 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 exit in the next five 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.

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

	U.S.	Canada
Job Creation Rate	28.9%	32.4%
Job Destruction Rate	39.4%	31.6%
Job Destruction from Large Firms	31.5%	15.3%
Job Creation from Exports	3.0%	23.3%

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

Haltiwanger and Schuh (1996), job flows are enormous. 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 10 percentage points higher.

- Job destruction due to “large” firms.** Row 3 in Table 1 presents the job destruction rate due to exit or employment declines only among *large* firms. Large is defined as above-average employment in the initial period. Large firms account for about 80% of all job destruction in the U.S., and about one-half of job destruction in Canadian manufacturing.
- Exports Create Jobs.** We impute employment due to exports as the product of a firm’s employment and the ratio of the firm’s 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 among firms where imputed employment from exports increases between the two years. We

divide this measure of job creation from exports by the average of total 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. So exports account for 70% of total job creation in Canada. The job creation rate due to exports in the U.S. is much smaller at 3% (about 10% of aggregate job creation in the U.S.).⁵

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 ten-year period starting on January 1, 1989.⁶ For Canada, we focus on the difference between 1973–1988 (“Pre-CUSFTA”) and 1988–2012 (“Post-CUSFTA”). We highlight three facts in the Canadian data:

4. **Job flows increased after the trade liberalization.** Table 2 shows that job creation and destruction rates increased in Canada after the trade agreement with the U.S. Figure 1 plots the change in the job destruction rate in two digit Canadian industries from 1973–1988 to 1988–2012 against the change in the tariff rate on Canadian imports in the industry due to CUSFTA.⁷ This figure indicates that job destruction rates increased more in industries where tariffs declined the most.⁸

5. **Large firms drove increased job destruction after trade liberalization.** Holmes and Stevens (2014) show that large U.S. manufacturing firms were most adversely affected by the surge in imports from China. Row 3 in Table 2 shows a similar fact applies to Canada. The job destruction rate

⁵Lincoln, McCallum and Siemer (2017) estimate that 29% of U.S. exports in 2002 were by firms that had been exporting for fewer than 5 years.

⁶The average tariff on manufacturing imports among the CUSFTA partners fell from over 8% to below 2% in Canada and from 4% to below 1% in the U.S.

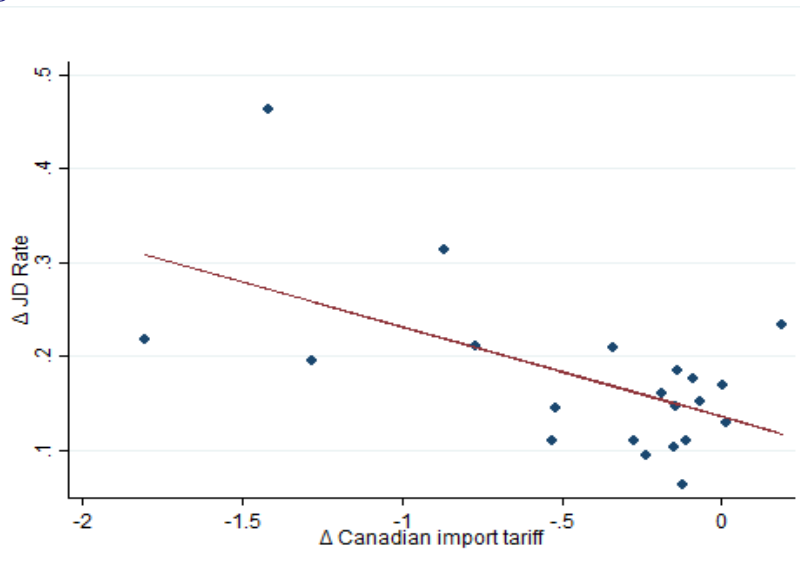
⁷We use the tariff cut measures constructed by Trefler (2004) which give changes in bilateral tariffs between Canada and the U.S. following CUSFTA net of the changes in the respective most-favored-nation tariffs.

⁸The coefficient of the OLS regression in Figure 1 is $-.096$ with a standard error of $.031$.

Table 2: Job Flows in Canada

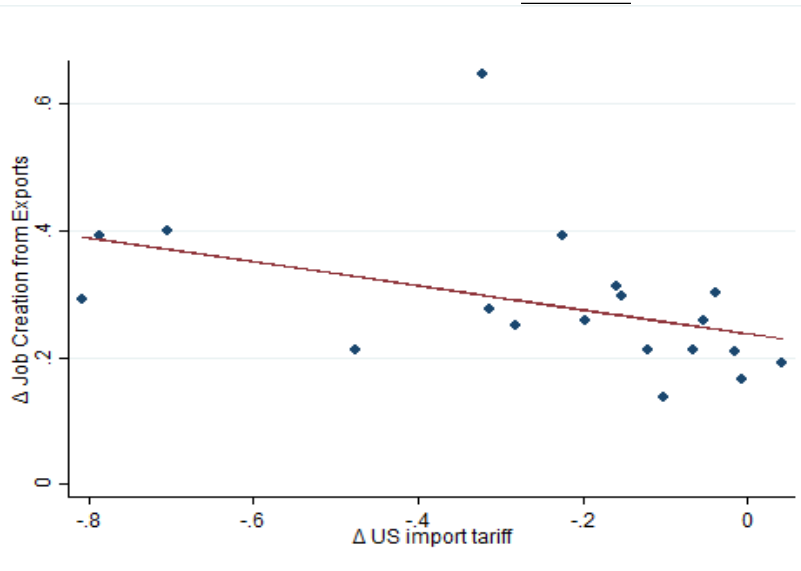
	Pre-CUSFTA	Post-CUSFTA
Job Creation Rate	28.0%	36.9%
Job Destruction Rate	26.3%	38.6%
Job Destruction from Large Firms	18.5%	29.1%
Job Creation from Exports	8.3%	32.3%

Note: Pre-CUSFTA is 1973 to 1988. Post-CUSFTA is 1988 to 2012. Job creation and destruction calculated over five year periods.

Figure 1: Δ Job Destruction in Canada vs. Δ Canadian Tariffs

Note: Each observation is a two digit Canadian industry. Δ job destruction is the difference between the average job destruction rate (calculated over five years) in 1988 to 2012 and 1973 to 1988.

Figure 2: Δ Job Creation in Canada from Exports vs. Δ U.S. Tariffs



Note: Each observation is a two digit Canadian industry. Δ job creation from exports is the difference between the average job creation rate from exports from 1989 to 2012 (calculated every five years) and 1974 to 1989.

among large (above mean employment) firms increased from 18.5% before CUSFTA to 29.1% after the trade agreement. So job destruction among large firms accounted for the bulk of the increase in the aggregate job destruction rate (from 26.3% to 38.6%).

6. **Job creation from exports increased after trade liberalization.** The last row in Table 2 shows that job creation from exports increases dramatically after the trade agreement. Figure 2 shows that a similar fact holds across two-digit Canadian industries. Job creation from exports increased by more in sectors where U.S. tariffs declined the most.⁹

Table 3 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

⁹The coefficient of the OLS regression in Figure 2 is -0.189 with a standard error of 0.096 .

Table 3: Job Flows in the U.S.

	1972–1987	1987–1992	1992–2012
Job Creation Rate	31.0%	29.9%	28.6%
Job Destruction Rate	29.3%	33.9%	40.7%
Job Destruction from Large Firms	22.3%	26.0%	32.9%
Job Creation from Exports	–	2.7%	3.1%

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.

also a surge of imports from China in the 1992–2012 period, so one should not interpret the changes in Table 3 as coming only from CUSFTA.

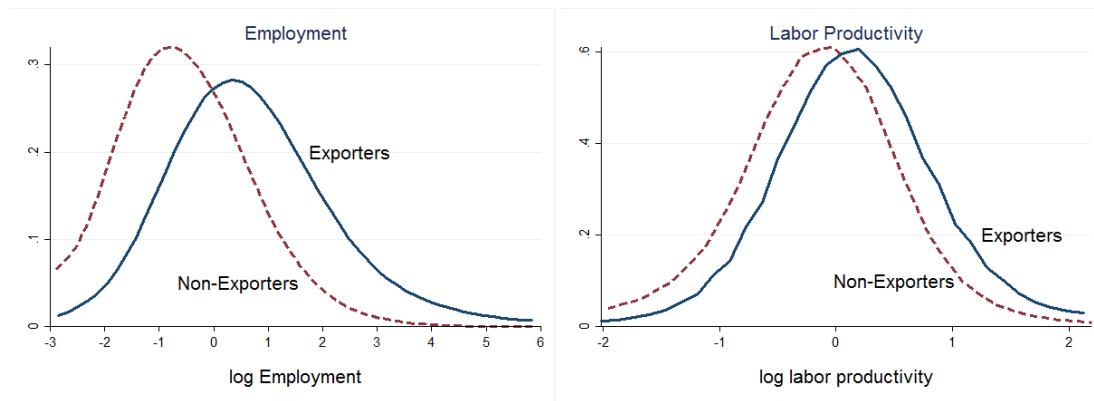
Three facts are clear from Table 3. First, job destruction increased markedly after 1987, by about 10 percentage points (row 2).¹⁰ Second, the increase in job destruction was entirely driven by large firms (job destruction by large firms increased by 10 percentage points after CUSFTA). Third, there is no export information in the U.S. manufacturing data prior to 1987, but job creation from exports increased modestly from 2.7% in 1987–1992 to 3.1% in 1992–2012.

We now look at differences between exporting and non-exporting firms. Figure 3 plots the distribution of employment (in the left panel) and labor productivity (in the right panel) from the U.S. manufacturing census in 2012. This figure reveals two additional facts:

7. **Average labor productivity and employment is higher for exporters than for non-exporters.** This can easily be seen in Figure 3.
8. **Large overlap of labor productivity and employment between exporters and non-exporters.** Figure 3 also makes clear that there is a substantial

¹⁰This may seem surprising given the evidence on declining dynamism in Decker et al. (2014). For U.S. manufacturing firms, at least, this decline was concentrated in job creation and took place well after CUSFTA.

Figure 3: Distribution of Employment and Labor Productivity



Note: The figures plot 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.

overlap in the two distributions. Many exporting firms are smaller than non-exporters, and labor productivity among many exporters is lower than for many non-exporters.

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 show that this model can mimic the eight facts described in section 2.

3.1 Static Equilibrium

Utility of the home-country representative consumer is given by the product of consumption of individual varieties C_j :

$$U = \prod_{j=1}^M C_j^{\frac{1}{M}} \quad (1)$$

where M denotes the (fixed) number of varieties.¹¹ This utility function implies that consumers spend the same share of their income on each variety.

Output of each variety is the product of labor and the quality of the blueprint for the product. We assume that there is a choice of blueprints for each product. We denote A_j as the “best” blueprint among domestic firms. A_j^* is the corresponding best blueprint among foreign firms. Furthermore, suppose the product index j is decreasing in A_j/A_j^* . Then products $j \in [0, x_1]$ are traded and produced at home, $j \in [x_1, x_2]$ are non-traded, and $j \in [x_2, M]$ are traded and produced abroad. The cutoff products x_1 and x_2 are defined by

$$\frac{A_{x_1}}{\tau} = \omega A_{x_1}^* \quad (2)$$

$$A_{x_2} = \frac{\omega A_{x_2}^*}{\tau} \quad (3)$$

where ω denotes the relative wage (domestic relative to foreign) and $\tau \geq 1$ is the symmetric gross trade cost. When $\tau = 1$, $x_1 = x_2$ and all products are traded.

The owner of the best blueprint sets the price to push its 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 its closest competitor – domestic or foreign. Table 4 summarizes the (gross) markup of domestic firms μ_i and foreign firms μ_i^* . Here A'_i and A'^*_i denote the productivity of the second best producer in the domestic and foreign markets, respectively. These poten-

¹¹We suppress the equations for the foreign country when they are the same as that of the home country. For example, utility of the foreign consumer is given by $U^* = \prod_{j=1}^M C_j^{*\frac{1}{M}}$.

Table 4: Markup

Market	Traded	
	Produced in Home	Produced in Foreign
Home	$\frac{A_i}{\max\left[A'_i, \frac{\omega A_i^*}{\tau}\right]}$	$\frac{A_i^*/\tau}{\max\left[\frac{A_i^*}{\tau}, \frac{A_i}{\omega}\right]}$
Foreign	$\frac{A_i/\tau}{\max\left[\frac{A'_i}{\tau}, \omega A_i^*\right]}$	$\frac{A_i^*}{\max\left[A_i^*, \frac{A_i}{\omega\tau}\right]}$

tial competitors do not produce in equilibrium but affect markups.

The relative wage is pinned down by the assumption that the value of exports is equal to the value of imports:

$$\bar{\mu}^* L^* \cdot \frac{x_1}{M} = \bar{\mu} \omega L \cdot \frac{M - x_2}{M} \quad (4)$$

where $\bar{\mu}^*$ and $\bar{\mu}$ denote the average gross markup of foreign and domestic firms and L^* and L the labor supply abroad and at home. The left hand side of equation (4) is the home country's exports and the right hand side is the home country's imports.

We can now express the real wage as a function of the distribution of the best blueprints, markups, the cutoff indexes, the relative wage, and the trade cost. The real wage at home W and in the foreign country W^* are given by

$$W = \prod_{j=1}^{x_2} \left(\frac{A_j}{\mu_j} \right)^{\frac{1}{M}} \prod_{j=x_2}^M \left(\frac{A_j^*}{\mu_j^*} \cdot \frac{\omega}{\tau} \right)^{\frac{1}{M}}. \quad (5)$$

$$W^* = \prod_{j=1}^{x_1} \left(\frac{A_j}{\mu_j} \cdot \frac{1}{\omega \tau} \right)^{\frac{1}{M}} \prod_{j=x_1}^M \left(\frac{A_j^*}{\mu_j^*} \right)^{\frac{1}{M}}. \quad (6)$$

Remember the home country buys products $j \in [x_2, M]$ from the foreign country so the real domestic real wage is increasing in the productivity of foreign firms of these products. Likewise, the foreign country purchases products

$j \in [0, x_1]$ from the home country so the foreign real wage depends on the productivity of domestic firms of these products.

Our static model is simply the Bernard, Eaton, Jensen and Kortum (2003) model, which is essentially the Dornbusch, Fischer and Samuelson (1977) model with heterogeneity in markups. Here equations (2), (3), (4) and Table 4 jointly determine the real wage in the two countries, the relative wage ω , the cutoff products x_1 and x_2 , and the markup for each product.

3.2 Innovation

We now introduce dynamics to the model. First, following Klette and Kortum (2004) we assume a firm is a portfolio of products, where a firm produces a product if it owns the best technology for that product. Second, we assume all growth comes from creative destruction – i.e., when a firm (another incumbent or an entrant) improves a product’s technology and steals it from the incumbent producer. Third, we assume creative destruction can also come from a firm located in another country. Our goal is to show that the trade friction τ is a key parameter that determines whether innovation in one country results in creative destruction in the other country.

We make the following assumptions about innovation. First, we assume a constant exogenous arrival rate for each type of innovation. (We will endogenize the arrival rate in the next section.) Second, we assume that arrivals are in proportion to the number of products owned by a firm. For example, a firm with two products is twice as likely to creatively destroy another firm’s variety compared to a firm with one product. Third, we assume that innovation builds on the existing quality level of the product *consumed domestically*.¹² Specifically, the quality drawn by an innovation follows a Pareto distribution

¹²If innovation was endogenous, there would be a positive externality to research unless all research was done by firms on their own products. Such knowledge externalities are routinely assumed in the quality ladder literature, such as Grossman and Helpman (1991), Aghion and Howitt (1992), Kortum (1997), and Acemoglu, Akcigit, Bloom and Kerr (2013).

Table 5: Channels of Innovation

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

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

with shape parameter θ and scale parameter equal to the existing quality level. The average percent improvement in quality of an existing variety (conditional on innovation) is thus $\frac{1}{\theta-1} > 0$.¹³ Finally, we assume that entrants have one product, while incumbent firms potentially produce many varieties.

The notation for innovation probabilities is given in Table 5. 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 a new 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. The probability a product will be improved upon by a domestic incumbent is λ . The *unconditional* probability of innovation by domestic entrant is $\tilde{\eta} \equiv \eta(1 - \lambda)$. The *unconditional* probability the product will be improved by foreign incumbent is $\tilde{\lambda}^* \equiv \lambda^*(1 - \lambda)(1 - \tilde{\eta})$. Finally, the unconditional probability of innovation by a foreign entrant is $\tilde{\eta}^* \equiv \eta^*(1 - \lambda)(1 - \tilde{\eta})(1 - \tilde{\lambda}^*)$.

¹³In the simulations described in the next section, we also implement a mean reversion process where we reduce the gap between each quality and the mean consumed quality in each country by 2% in each period. This enables us to maintain a stationary distribution of qualities.

Table 6: Probability of Creative Destruction

Home market	Exported or Non-Traded	$\lambda + \tilde{\eta} + (\tilde{\lambda}^* + \tilde{\eta}^*) \min \left[\left(\frac{\omega}{\tau} \right)^\theta, 1 \right]$
	Imported	$(\lambda + \tilde{\eta}) \min \left[\left(\frac{\tau}{\omega} \right)^\theta, 1 \right] + \tilde{\lambda}^* + \tilde{\eta}^*$
Foreign market	Exported or Non-Traded	$(\lambda + \tilde{\eta}) \min \left[\left(\frac{1}{\omega\tau} \right)^\theta, 1 \right] + \tilde{\lambda}^* + \tilde{\eta}^*$
	Imported	$\lambda + \tilde{\eta} + (\tilde{\lambda}^* + \tilde{\eta}^*) \min \left[(\omega\tau)^\theta, 1 \right]$

In closed economy models of creative destruction, such as Klette and Kortum (2004), a firm that improves a product's quality replaces the incumbent producer. In our model, this is still the case in the innovating firm's local market, but this is not necessarily true in the other market. Consider a foreign firm that improves a variety produced by a domestic firm. Since the foreign innovator has to pay a trade cost to access the domestic market, the innovating firm will replace the incumbent producer in the domestic market only if the quality improvement exceeds $\frac{\tau}{\omega}$. With a Pareto distribution of innovation draws, this occurs with probability $\min \left[\left(\frac{\omega}{\tau} \right)^\theta, 1 \right]$. So for a given relative wage, a higher trade cost lowers the probability that a foreign innovator creatively destroys the incumbent producer in the domestic market.

Similarly, relative wages help determine who produces a given product. A domestic firm with higher productivity may still lose its market to a foreign firm if tariffs and its productivity advantage are overwhelmed by lower wages in the foreign market. A domestic firm that retains its home market may not be able to penetrate the foreign market due to a combination of tariffs and wage differentials, even if it has the best blueprint in the world.

Table 6 summarizes the probability of creative destruction in the home market and in the foreign market. In an open economy, the arrival rate of new

blueprints available to domestic consumers depends on the innovation rate of domestic firms $\lambda + \tilde{\eta}$ and of foreign firms $\tilde{\lambda}^* + \tilde{\eta}^*$. The impact of a given foreign innovation on domestic consumers is increasing in the relative wage ω and decreasing in the trade cost τ . A higher domestic wage increases the probability a foreign firm that innovates will be competitive enough to replace the domestic incumbent in the domestic market. In contrast, higher trade costs make the foreign innovator less competitive compared to the domestic firm. Effectively, trade costs insulate domestic firms from foreign competition in the domestic market.

Note that trade costs also helps domestic firms that improve an imported variety replace the foreign firm in the domestic market. Intuitively, a foreign firm has to pay trade costs when selling into the domestic market. The quality gap between the domestic innovator and the foreign firm has to be greater than $\frac{\omega}{\tau}$ for the domestic firm to replace the imported product. This occurs with probability $\min \left[\left(\frac{\tau}{\omega} \right)^\theta, 1 \right]$, which is increasing in τ (until it hits probability 1).

The *expected* growth rate g of the real wage at home is given by

$$g = \underbrace{(\lambda + \tilde{\eta}) \left[\frac{x_2}{M} \cdot \frac{1}{\theta - 1} + \frac{M - x_2}{M} \left(\frac{\tau}{\omega} \right)^\theta \left(\frac{\theta}{\theta - 1} \left[\frac{\omega}{\tau} \right] - 1 \right) \right]}_{\text{domestic innovation}} + \underbrace{(\tilde{\lambda}^* + \tilde{\eta}^*) \left[\frac{M - x_2}{M} \cdot \frac{1}{\theta - 1} + \frac{x_2}{M} \left(\frac{\omega}{\tau} \right)^\theta \left(\frac{\theta}{\theta - 1} \left[\frac{\tau}{\omega} \right] - 1 \right) \right]}_{\text{foreign innovation}}. \quad (7)$$

The first term in (7) is growth from innovation by domestic firms. Remember domestic firms attempt to innovate over *all* products, including imported ones. So $\lambda + \tilde{\eta}$ is the rate at which a domestically produced variety is replaced by another domestic firm and $(\lambda + \tilde{\eta}) \left(\frac{\tau}{\omega} \right)^\theta$ is the probability an imported product is taken over by a domestic firm. Conditional on being replaced, $\frac{1}{\theta - 1}$ and $\frac{\theta}{\theta - 1} \frac{\tau}{\omega} - 1$ are the quality improvement of a domestically produced and an imported variety, respectively. The second term in (7) captures the impact of innovation by *foreign* firms on domestic consumers. The probability an imported product is

replaced by another foreign firm is $\tilde{\lambda}^* + \tilde{\eta}^*$ and the probability a domestically produced variety is replaced by a foreign firm is $(\tilde{\lambda}^* + \tilde{\eta}^*) \left(\frac{\omega}{\tau}\right)^\theta$.

The expected growth rate of the real wage in the foreign country is

$$g^* = \underbrace{\left(\tilde{\lambda}^* + \tilde{\eta}^*\right) \left[\frac{M - x_1}{M} \cdot \frac{1}{\theta - 1} + \frac{x_1}{M} (\tau\omega)^\theta \left(\frac{\theta}{\theta - 1} \left[\frac{1}{\omega\tau} \right] - 1 \right) \right]}_{\text{foreign innovation}} + (\lambda + \tilde{\eta}) \underbrace{\left[\frac{x_1}{M} \cdot \frac{1}{\theta - 1} + \frac{M - x_1}{M} \left(\frac{1}{\omega\tau} \right)^\theta \left(\frac{\theta}{\theta - 1} [\omega\tau] - 1 \right) \right]}_{\text{domestic innovation}}. \quad (8)$$

Here, foreign consumers gain from innovation by domestic firms, in the same way that domestic consumers gain from innovation by foreign firms.

In a steady state, the real wage grows at the same rate in the two countries, and differences in innovation rates show up in the relative wage.¹⁴ Figure 4 shows the domestic wage relative to the foreign wage (in the left panel) and the growth rate in *both* countries (in the panel on the right) as a function of the innovation rate of domestic incumbent firms λ .¹⁵ A faster innovation rate by domestic firms increases the domestic wage relative to the foreign wage, but increases the growth rate of the real wage equally in the two countries.

3.3 Calibration

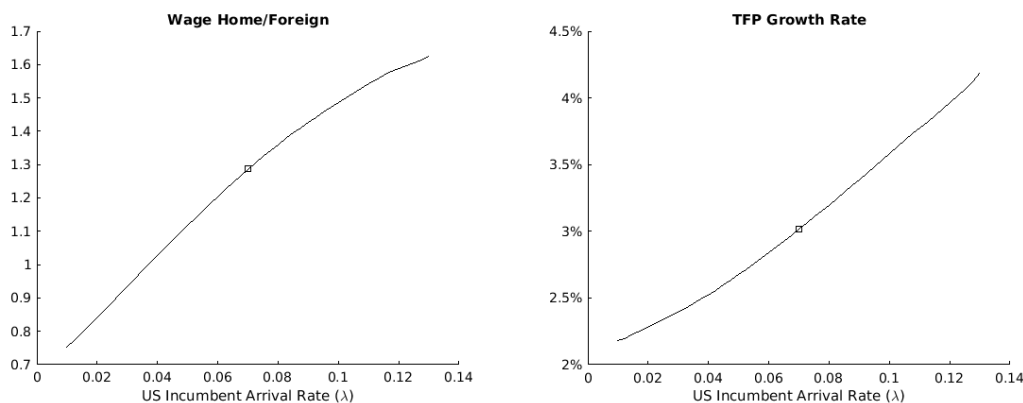
The model is summarized by two innovation rates (for incumbents and entrants) in each country, the shape parameter of the Pareto distribution of the innovation draws, and a trade cost. In this section, we infer the value of these parameters from simple moments of the data.

For consistency with the model, we assume the world consists of the manu-

¹⁴This result is reminiscent of Acemoglu and Ventura (2002).

¹⁵The values of the other parameters used in the simulation behind Figure 4 are $\eta = .021$, $\lambda^* = .0706$, $\eta^* = .0228$, $\theta = 5.455$, and $\tau = 1.515$. We explain in the next subsection where these numbers come from, although for the purposes of illustrating the effect of λ on ω and g any set of parameter values will do.

Figure 4: Effect of Home Innovation on Growth and Relative Income



Note: The left panel shows the steady-state wage at home relative to the foreign country as a function of the innovation rate of domestic incumbent firms, holding fixed the other variables. The right panel shows the effect on the steady state growth rate of aggregate TFP in *both* countries.

facturing 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^*), the innovation rates in Table 5, and the trade cost (τ) jointly determine the growth rate, the trade share and the relative wage. Therefore, 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.¹⁶ We use the employment share of new firms to pin down the share of innovation by entrants vs. incumbents.

The data moments we use are displayed in Table 7. We fit a steady state where TFP in the two countries grows at 3% per year, employment shrinks at 1.3% per year, output per worker in the home country is 29% higher than in the foreign country, and the home trade share is 10%.¹⁷ The innovation rates and

¹⁶We describe in the next section how we back out θ from the gap in labor productivity (revenue per worker) between exporters and non-exporters.

¹⁷Overall OECD employment and population was not contracting during our sample, of course. But employment was indeed shrinking within manufacturing. We target this employment contraction to help fit the observed job creation and destruction rates in the U.S. Negative employment growth poses no problems for obtaining steady state TFP growth and

Table 7: 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	10.17%
Employment home/foreign	U.S. and OECD mfg	0.389
Employment growth rate	OECD mfg	-1.3%

trade cost needed to fit these facts are shown in Table 8.¹⁸ The innovation rates in Table 8 are *conditional*. The *unconditional* innovation rate is $\lambda + \tilde{\eta} = 0.090$ for domestic firms and $\tilde{\lambda}^* + \tilde{\eta}^* = 0.083$ for foreign firms. The innovation rate of domestic firms has to be higher than that of foreign firms to explain the 29% higher real wage in the home country than in the foreign country. Conditional on the innovation rates and the relative size of the two economies, the trade share pins down the trade cost, which is roughly equivalent to a 50% tariff.¹⁹

3.4 Firm Dynamics

We now show that the model can generate the eight facts described in Section 2. The driver of growth here is the creation of better quality products and the

relative wages in the exogenous innovation case. We will revisit this issue, however, in the endogenous innovation case.

¹⁸We simulate the model for 5000 varieties in each country. Each variety receives an innovation draw and is assigned to an existing incumbent or a new entrant with random probability governed by the innovation rates. The relative wage is selected to balance trade between the two countries. We simulate this process for several hundred years until it settles down to a steady-state, at which point we calculate the firm-level and aggregate moments produced by the model. We utilize a simulated annealing procedure to search for the parameter values that match the moments in the data.

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

Table 8: Estimates of Model Parameters

Variable	Description	Value
θ	Shape parameter of innovation draws	5.45
λ	Home innovation rate incumbent	6.93%
η	Home innovation rate entrants	2.21%
λ^*	Foreign innovation rate incumbents	7.06%
η^*	Foreign innovation rate entrants	2.28%
τ	Trade cost	1.515

resulting destruction of product lines of incumbent producers. In our open economy setting, the destruction of product lines can come from firms located in other countries. This has three implications:

First, growth is associated with job creation by innovating firms and job destruction by firms whose products are innovated upon. Table 9 (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 7. The job creation rate (over five years) is 28%, and the job destruction rate is 6% higher at about 34%. For comparison, the first column in Table 9 replicates the U.S. data. The job flows predicted by the model with the parameters in Table 8 are roughly of the same magnitude as in the data (fact 1).

Second, firms shrink their employment in this model when their products are replaced by other firms. A firm exits (and all its jobs are destroyed) when it loses *all* of its products. But job destruction is not only due to exit by small firms; jobs are destroyed at large (multiproduct) firms when a subset of their products are innovated upon. The third row in Table 9 shows that, consistent with the evidence from U.S. and Canadian manufacturing (fact 2), job destruction in the model is mostly driven by large firms, where large is defined as firms with above-average employment. The job destruction rate by large firms is 22% in

Table 9: Firm Dynamics, Simulations vs. Data

Moment	U.S. Data	Simulations
Job Creation Rate	28.9%	27.7%
Job Destruction Rate	39.4%	34.2%
Job Destruction from Large Firms	31.5%	21.9%
Job Creation from Exports	3.0%	3.6%
Job Destruction from Imports	–	5.2%

The U.S. data is the average from 1987 to 2012. The second column show the simulated values in the steady state of the model with the parameters in Table 8.

the model, which is about two-thirds of the overall job destruction rate.

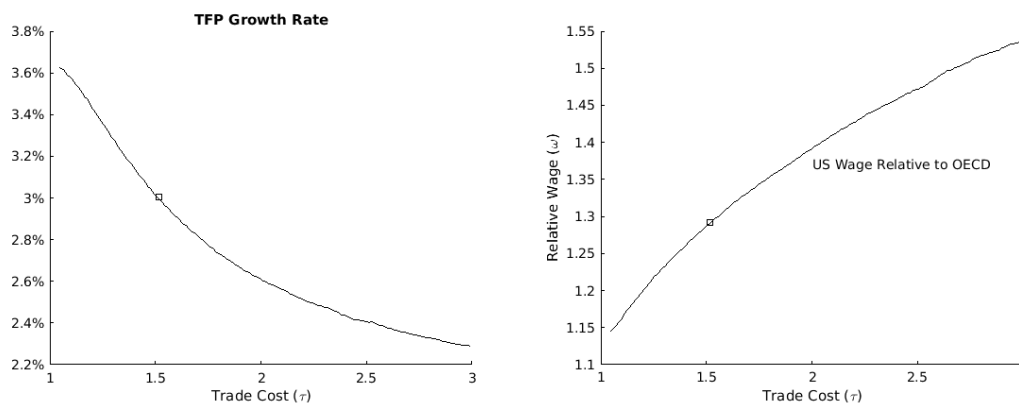
Third, some of the job destruction is the result of innovation from firms located in other countries, and some of the job creation is the result of domestic firms replacing producers in the foreign market. The fourth row in Table 9 shows that the job creation rate from exports in the model is 3.6%. In the U.S. data, the number is 3%. We cannot measure 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 5.2%, so about a quarter of the overall job destruction in the model comes from creative destruction by foreign firms.

Table 9 shows that the model can replicate, at least qualitatively, facts 1-3 in Section 2. We now show that the model can generate facts 4-6. We simulate the new steady state of the model with lower trade costs, holding constant all the other parameters. A key assumption in this exercise is that domestic and foreign innovation rates do not change when trade costs change.²⁰

We first show the consequence of reducing trade costs for TFP growth in the two countries. Lower trade costs make it more likely that innovation by a do-

²⁰We relax this assumption in the next section, when we endogenize innovation rates.

Figure 5: Simulated TFP Growth Rate and Relative Wage vs. Trade Costs



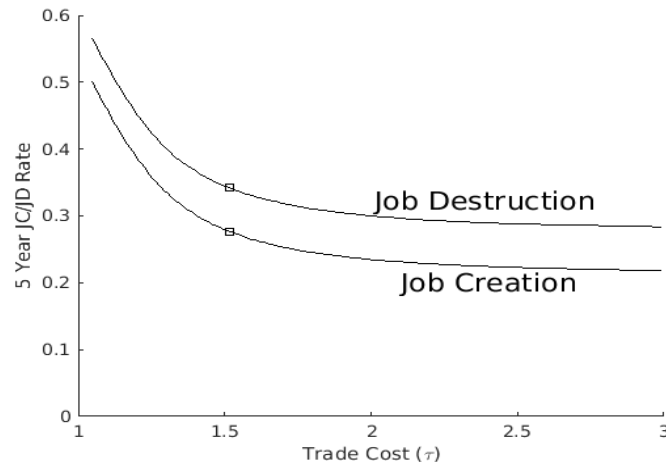
Note: The left panel shows simulated growth rate of the real wage and the right panel the relative wage (home/foreign) in the steady state for different values of τ . All other parameters of model are kept fixed at the values in Table 8.

mestic firm on a foreign product will replace the foreign producer in the foreign market, and thus raise the quality foreign consumers can consume. Similarly, lower trade costs increase the probability that a foreign firm who innovates on a domestically-produced variety will replace the domestic producer and thus improve the quality that domestic consumers enjoy. The effect of lower trade costs, shown in Figure 5 (left panel), is to boost the growth rate of TFP in the two countries in tandem. Holding constant innovation rates, moving to a world with frictionless trade increases the steady-state TFP growth rate from 3% to 3.6%.

Lower trade costs also raise the foreign wage relative to the domestic wage. This is shown in the right panel in Figure 5. The intuition is that a country that 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 6 shows the effect of the higher TFP growth rate on job creation and destruction. The model predicts that, relative to our baseline steady state ($\tau = 1.515$), free trade would increase the job creation and destruction rates by about 20 percentage points. So a 0.6 percentage point increase in the growth rate is

Figure 6: Simulated Job Creation and Destruction vs. Trade Costs



Note: The figure simulates the steady state job creation and destruction rates when we vary τ but keeping constant all other parameters at the values in Table 8.

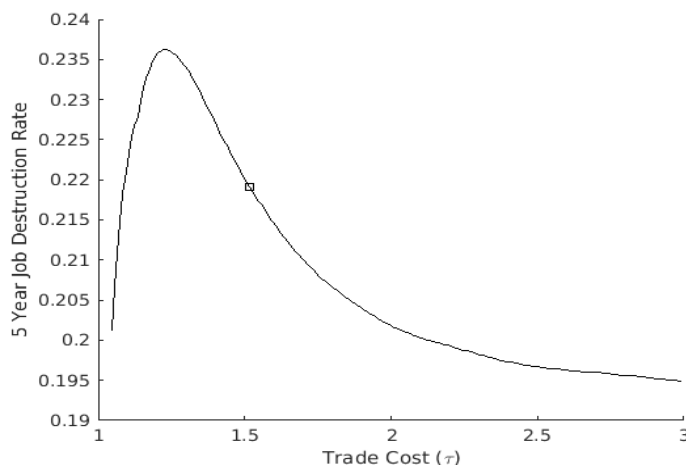
associated with a 20 percentage point increase in the job creation and destruction rate. The impact on job reallocation is more proportionally bigger than the impact on growth, because lower trade costs facilitate creative destruction from trade with smaller step sizes. Consistent with the evidence from Canada and the U.S. (fact 4), the model predicts that job flows rise when trade costs fall.

Figure 7 shows the counterfactual 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 5), in the model job destruction from large firms rises when trade costs fall – at least down to current levels of around a 50% tariff rate. Surprisingly, the model says that job destruction eventually falls as trade costs plunge toward frictionless trade.²¹

Next we show the model predicts an increase in job creation *from exports* (fact 6) in the aftermath of trade liberalization. In Figure 8 (left panel), the job

²¹We are still trying to understand this counterintuitive result, but suspect it is because the fraction of employment at large firms endogenously falls along with trade costs.

Figure 7: Simulated Job Destruction from *Large* Firms vs. Trade Costs

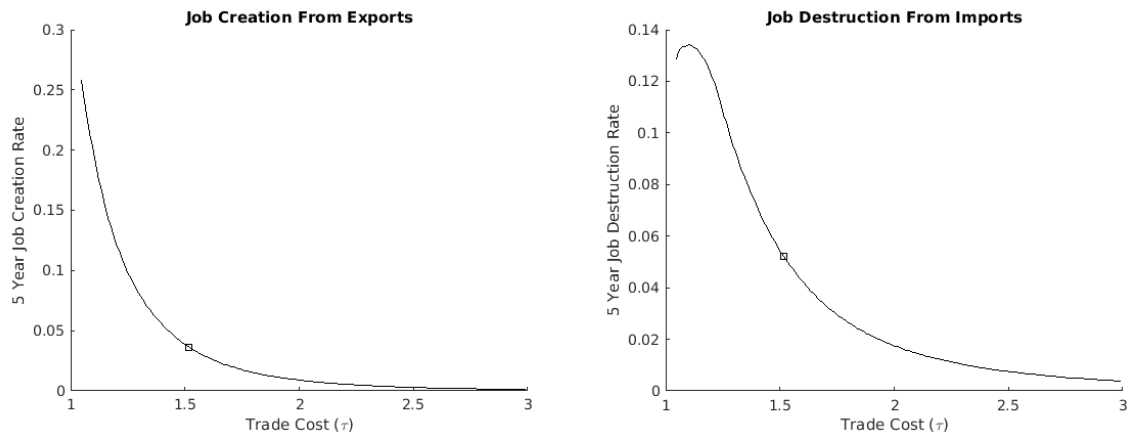


Note: The figure simulates the steady state job destruction rate from large (above mean employment) firms when we vary τ but keeping constant all other parameters at the values in Table 8.

creation rate from exports increases by 20 percentage points when trade costs fall from our benchmark value ($\tau = 1.515$) to completely free trade ($\tau = 1$). The right panel in Figure 8, meanwhile, plots the predicted job destruction rate from domestic firms who are replaced by imports. The model predicts that moving to free 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 9 plots the simulated distribution of employment for exporters and non-exporters in the steady state, using our baseline parameter values from Table 8. These are the model analogues to facts 7 and 8. 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

Figure 8: Simulated Job Flows from Trade vs. Trade Costs

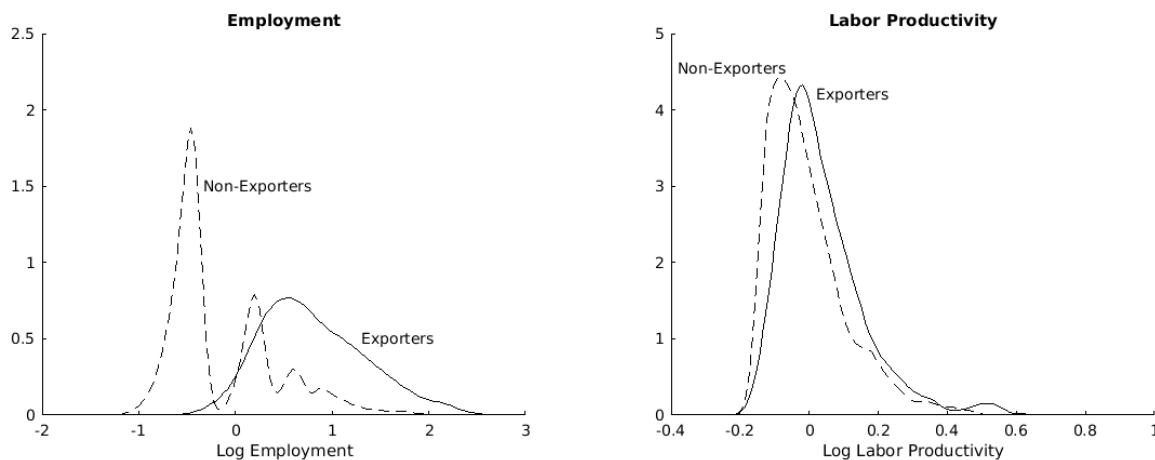


Note: The left panel shows the job creation rate due to exports and the right panel the job destruction rate due to imports in the steady state for different values of τ . All other parameters of model are kept fixed at the values in Table 8.

are more likely to have 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 and quality of products between the two groups of firms. Consistent with the empirical distribution of employment in Figure 3, 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 compared to non-exporters. This can be seen in the right panel in Figure 9, which plots the distribution of labor productivity (revenue per worker) of exporters and 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, 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

Figure 9: 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 8.

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 of the Pareto distribution of innovation draws θ . We therefore calibrate this parameter to make the simulated model match the average gap in labor productivity between exporters and non-exporters in the U.S. data.²²

Figure 3 makes clear that, while average labor productivity (and thus quality when viewed through the lens of our model) of exporters is higher than for non-exporters on average, there is substantial overlap in their distributions. As depicted in Figure 9, our model generates such an overlap because many non-exporters. This reflects overlap in markups between the two groups of firms.

To be fair, the empirical *dispersion* of employment and labor productivity (Figure 3) is substantially larger than in the simulated data (Figure 9). Our as-

²²As stated in Table 7, this gap is 6.6%.

sumption that preferences over varieties is Cobb-Douglas implies that product quality only matters for employment when higher quality helps the firm 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, overhead costs, adjustment costs and measurement error, all of which are likely present in the data and behind some of the dispersion in labor productivity in the data.²³ We leave these useful extensions for future work.

4 Endogenous Innovation

A key assumption we have made so far is that the innovation rates are exogenous parameters. We now consider the effect of a reduction in trade costs when innovation rates are endogenously determined.

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

$$\lambda = \left(\frac{R_i}{\gamma \chi_i \bar{A}^{(1-\phi)/\gamma}} \right)^\gamma, \quad (9)$$

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 (9) 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).

²³Bartelsman et al. (2013) emphasize the role of overhead costs, Asker et al. (2014) the importance of adjustment costs, and Bils et al. (2017) the contribution of measurement error.

Similarly, suppose the unconditional innovation rate of domestic entrants is

$$\tilde{\eta} = \left(\frac{R_e}{\gamma \chi_e \bar{A}^{(1-\phi)/\gamma}} \right)^\gamma, \quad (10)$$

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 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 $\bar{A}^{(1+\eta-\phi)/\eta}$. 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 :

$$\begin{aligned} v_{x,t} = & \pi_{x,t} - \gamma \chi_i \lambda_t^{\frac{1}{\gamma}} \\ & + \frac{(1+g_t)^\phi}{1+r_t} [\lambda_t (\beta_{x,t} v_{x,t+1} + \beta_{n,t} v_{n,t+1})] \\ & + \frac{(1+g_t)^\phi}{1+r_t} [(1-\delta_{x,t}) v_{x,t+1} - \delta'_{x,t} (v_{x,t+1} - v_{n,t+1})]. \end{aligned} \quad (11)$$

Here g denotes the growth rate of the real wage, r the interest rate, π_x expected profits (normalized by $\bar{A}^{(1+\eta-\phi)/\eta}$), 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 equation (11) is profit net

²⁴The innovation rates for foreign firms are given by equations analogous to (9) and (10) with R_i and χ_i replaced by R_i^* and χ_i^* in (9), R_e and χ_e replaced by R_e^* and χ_e^* in (10), and average quality sold in the foreign market instead of in the home market.

²⁵ δ_x and δ'_x are given by the creative destruction rates in Table 6. Since we assume innovation is not targeted, β_x and β_n depends on the share of each type of product and the creative

of research expenses (normalized by $W^{2-\phi}$), the second term is the expected value of grabbing a new product, and the last term is the expected value of an exported variety in the next period minus the expected value of losing the product to a competitor in the foreign market.

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

$$\begin{aligned} v_{n,t} = & \pi_{n,t} - \gamma \chi_i \lambda_t^{\frac{1}{\gamma}} \\ & + \frac{(1+g_t)^\phi}{1+r_t} [\lambda_t (\beta_{x,t} v_{x,t+1} + \beta_{n,t} v_{n,t+1})] \\ & + \frac{(1+g_t)^\phi}{1+r_t} [(1-\delta_{n,t}) v_{n,t+1}]. \end{aligned} \quad (12)$$

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

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

$$\lambda_t = \left(\frac{\beta_{x,t} v_{x,t} + \beta_{n,t} v_{n,t}}{\chi_i} \right)^{\frac{\gamma}{1-\gamma}} \quad (13)$$

$$\tilde{\eta}_t = \left(\frac{\beta_{x,t} v_{x,t} + \beta_{n,t} v_{x,t}}{\chi_e} \right)^{\frac{\gamma}{1-\gamma}} \quad (14)$$

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 4, except that the innovation rates are now pinned down by (11) through (14).²⁶

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 , but it is still the case that

destruction rates in Table 6.

²⁶Plus the corresponding optimal innovation rates for foreign firms. We assume linear utility so that the consumption Euler equation implies $r = \rho$. We set $\rho = 0.05$.

Table 10: Estimates of Model Parameters, Endogenous Innovation

Variable	Description	Value
ϕ	Return to stock of ideas	0.17
γ	Return to research intensity	0.55
χ_e/χ_i	Home entrant/incumbent research cost	2.69
χ_i^*/χ_i	Foreign/home incumbent research cost	3.50
χ_e^*/χ_i	Foreign entrant/home incumbent research cost	9.44

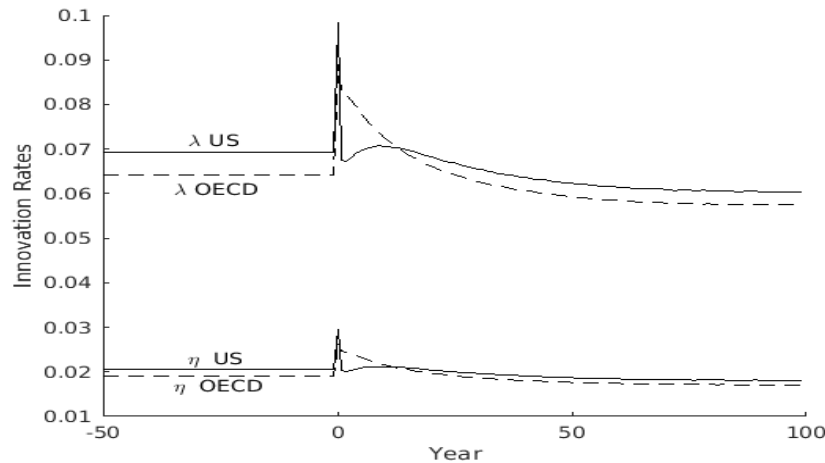
differences in innovation rates show up as differences in the relative wage. What is new in the endogenous innovation model is that the growth rate of the real wage in steady state is ultimately given by the product of the population growth rate and $\gamma/(1 - \phi)$ with $\phi < 1$.

We set ϕ and γ to match the empirical growth rate of real wages (TFP growth), 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).²⁷ This yields $\phi = 0.17$ and $\gamma = 0.55$. The calibrated values of these parameters are given in Table 10. The innovation cost parameter χ is lower in the U.S. to generate higher U.S. innovation rates and match the higher wages in the U.S. (relative to the foreign country).

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

²⁷TFP growth averaged 3.01% per year from 1995–2008 according to the U.S. Bureau of Labor Statistics KLEMS data. Real intellectual property investments grew 4.53% per year from 1995–2008 according to the U.S. Bureau of Economic Analysis, and such investments averaged 10.4% of value added in U.S. manufacturing from 1997–2008.

Figure 10: Simulated arrival rates after trade liberalization



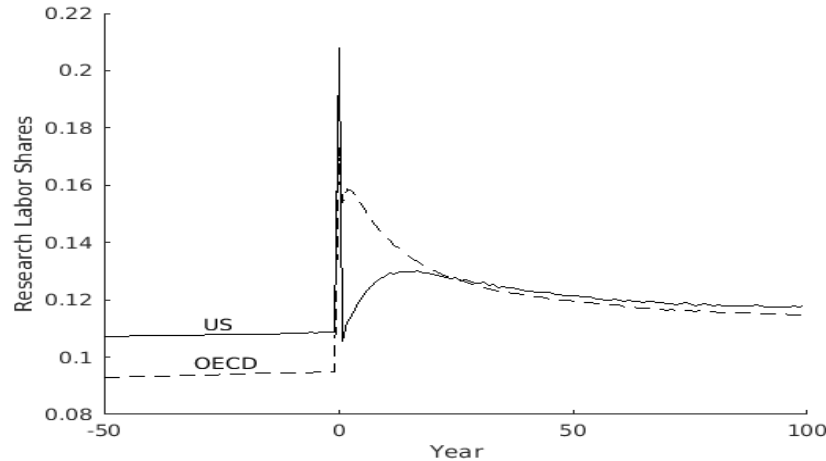
Note: The figure simulates the path of arrival rates in response to an unanticipated reduction in trade costs, keeping constant all other parameters at the values in Table 10. Here λ and η are the endogenous arrival rates of innovation by incumbents and entrants, respectively.

are actually *lower* but the growth rate of TFP is the same.

Figure 10 plots the effect of a permanent, unanticipated reduction of trade costs, from $\tau = 1.515$ to $\tau = 1.2575$, on innovation rates. It shows that innovation rates initially rise in the aftermath of a reduction in trade costs. But as the level of TFP rises, arrival rates fall due to the rising difficulty of innovating. 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 across countries more with lower trade costs.

Figure 11 plots the effect of a permanent reduction in trade costs on the share of labor devoted to research in each country. Like the arrival rates, the shares spike on impact. Unlike the arrival rates, the share of labor doing research ends up higher in the long run. The bigger market for each successful

Figure 11: Research labor shares after trade liberalization



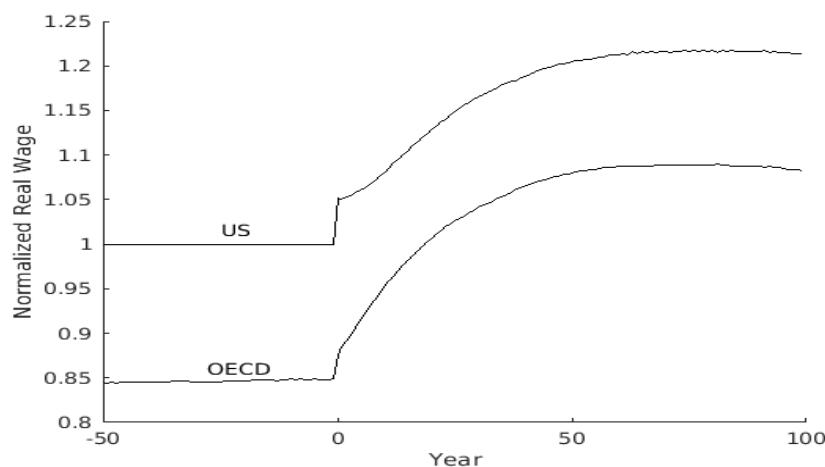
Note: The figure simulates the share of research labor in total employment in response to an unanticipated reduction in trade costs, keeping constant all other parameters at the values in Table 10.

innovation makes higher research effort worthwhile despite the endogenously greater difficulty in coming up with ideas. This result contrasts with Eaton and Kortum (2001), wherein these two forces exactly offset each other and leave research effort unchanged.

Figure 12 simulates the effect of permanently lower trade costs on the level of real wages. Real wages are closely tied to TFP in each country. The figure expresses variables relative to the path of U.S. wages in the absence of trade liberalization. The growth rate of wages increases temporarily in response to a decline in trade costs, but eventually slows down as innovation becomes more difficult with higher TFP. In the new steady state, real wages and TFP are permanently higher (compared to the initial steady state path), 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 13 maps the impact of a host of different tariff levels on the level of real wages in the long run. The real wage is almost always higher as a result of lower

Figure 12: Simulated real wages after trade liberalization



Note: The figure simulates the path of the real wage in response to an unanticipated reduction in trade costs, keeping constant all other parameters at the values in Table 10.

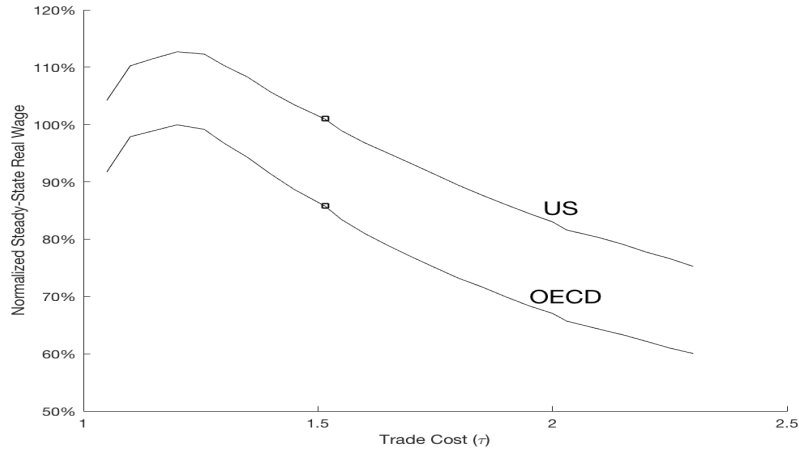
tariffs. The real wage actually falls, however, as frictionless trade is approached. At very low tariff levels, the high rate of creative destruction from imports discourages research effort so much that it outweighs the quicker spread of ideas.

Figure 14 illustrates how freer trade affects job flows. Consistent with the evidence from the U.S. and Canada after CUSFTA, job flows surge in the aftermath of a reduction in trade costs. The pace of job flows remains elevated for decades after a tariff reduction — certainly within the 24-year window we examine after the 1988 U.S.-Canada Free Trade Agreement. This pattern drives home that there can be dynamic costs (job destruction) as well as dynamic benefits (knowledge flows) to trade liberalization.²⁸

In Table 11 we present the welfare gains from lower trade barriers. Here we measure welfare as the present discounted value of consumption, consistent with linear utility and a constant real interest rate, using a discount rate of 5%.

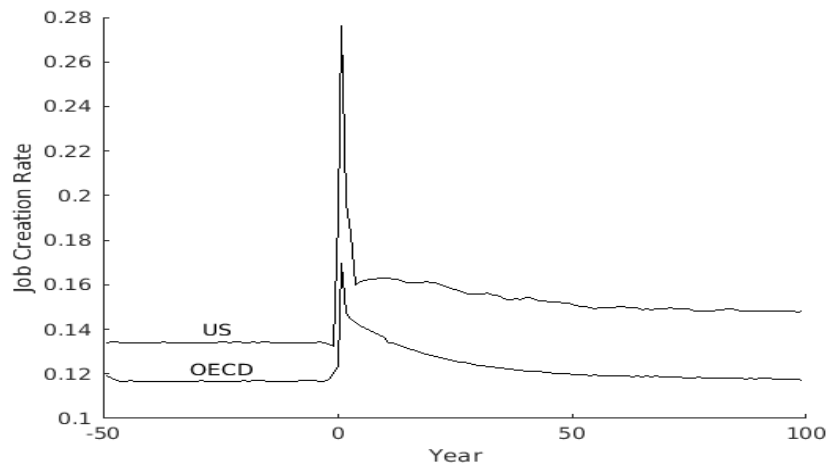
²⁸Bernard, Redding and Schott (2007) analyze a multi-sector Melitz model which also features effects of trade liberalization on steady state job flows.

Figure 13: Simulated steady-state real wages vs. various trade costs



Note: The figure plots the simulated equilibrium real wage in steady-states with differing trade costs, keeping constant all other parameters at the values in Table 10.

Figure 14: Simulated job flows from trade liberalization



Note: The figure simulates job creation and destruction rates in response to an unanticipated reduction in trade costs, keeping constant all other parameters at the values in Table 10.

In the first counterfactual, we reduce tariffs from 3.5, a high tariff at which the aggregate trade share is about 1%, to our estimated value of 1.515. For comparison, we first calculate the static gains – the jump in real wages on impact. The static gains are 1.6% in the U.S. and 1.1% in the rest of the OECD. The U.S. is the smaller economy (due to smaller population) with the larger trade share, so it gains more.²⁹

Next we present the full gains from freer trade, including effects on innovation. With exogenous arrival rates of innovation, the total gains are quite large at almost 42% in the U.S. The rest of the OECD gains even more (over 63%) because it gets more ideas than it gives. The dynamic gains are markedly smaller with endogenous innovation rates, but still sizable at about 18% in the U.S. and 26% in the rest of the OECD. The gains are much more modest with endogenous research effort for two reasons. First and foremost, we built in diminishing returns to the stock of ideas ($\phi < 1$) and a congestion externality ($\gamma < 1$). These temper the cumulative TFP gains from endogenously rising research effort and a faster flow of ideas across countries. Second, the higher research effort comes at the cost of less labor devoted to production and consumption. Even in the endogenous innovation case, however, the full gains are an order or magnitude larger than the static gains.

The second counterfactual in Table 11 reduces tariffs from our estimated value of 1.515 to 1.2575. The static gains are larger because trade flows increase sharply from 10% to over 25% in the U.S. in this counterfactual. The dynamic gains are smaller, as the ideas flowing involving endogenously smaller step sizes. The dynamic gains are still two to seven times the static gains.

²⁹Our static gains from trade of 1.6% for the U.S. are very close to the formula in Arkolakis et al. (2012). When the trade share goes from 1% to 10.2%, their formula predicts gains of 1.7% (9.2%/5.45, given our $\theta = 5.45$).

Table 11: Gains From Trade

Method	“Near” Autarky to Estimated τ		Reducing τ By Half	
	U.S.	OECD	U.S.	OECD
Static Gains	1.6%	1.1%	4.1%	2.6%
Dynamic Gains - Exogenous Innovation	41.8%	63.5%	30.8%	34.6%
Dynamic Gains - Endogenous Innovation	17.9%	25.8%	9.8%	16.4%

Entries give the percentage increase in the present discounted value of consumption as a result of reducing tariffs. In the first counterfactual, we reduce tariffs from 3.5, a high tariff at which the aggregate trade share is about 1%, to our estimated value of 1.515. In the second counterfactual, we reduce tariffs from our estimated value of 1.515 to 1.2575. We use a discount rate of 5%.

5 Conclusion

We documented eight 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.

We used these facts as motivation to construct a two-country model of creative destruction and trade. In the model, lower tariffs increase the probability that foreign firms take over domestically produced products. As in the data, this occurs in an initial burst but remains higher for decades. Foreign and domestic firms take over each other’s markets more frequently when trade barriers are lower. This dynamic disruption is a byproduct of faster growth. Growth remains high in the version of the model with exogenous innovation rates.

When we endogenize innovation rates 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 static gains from trade in our model.

Trade, like technological change more generally, brings dynamic benefits but also dynamic costs.

We see several potential 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 empirical discipline on the form and magnitude of 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 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

- Acemoglu, Daron and Jaume Ventura, “The world income distribution,” *The Quarterly Journal of Economics*, 2002, 117 (2), 659–694.
- , Ufuk Akcigit, Nicholas Bloom, and William R Kerr, “Innovation, Reallocation and Growth,” Working Paper 18993, National Bureau of Economic Research 2013.
- 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,” Technical Report, working paper 2017.
- Akcigit, Ufuk, S Ates, and Giammario Impullitti, “Innovation and Trade Policy in a Globalized World,” *University of Chicago mimeograph*, 2017.

- Alvarez, Fernando E, Francisco J Buera, and Robert E Lucas, “Idea flows, economic growth, and trade,” Technical Report, National Bureau of Economic Research 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.
- Autor, David H, David Dorn, and Gordon H Hanson, “The China syndrome: Local labor market effects of import competition in the United States,” *American Economic Review*, 2013, 103 (6), 2121–68.
- Bartelsman, Eric, John Haltiwanger, and Stefano Scarpetta, “Cross-country differences in productivity: The role of allocation and selection,” *The 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.
- , Stephen J Redding, and Peter K Schott, “Comparative advantage and heterogeneous firms,” *The Review of Economic Studies*, 2007, 74 (1), 31–66.
- Bils, Mark, Peter J. Klenow, and Cian Ruane, “Misallocation or Mismeasurement?,” 2017.
- Bloom, Nicholas, Charles I Jones, John Van Reenen, and Michael Webb, “Are Ideas Getting Harder to Find?,” 2017.
- , Mirko Draca, and John Van Reenen, “Trade induced technical change? The impact of Chinese imports on innovation, IT and productivity,” *The Review of Economic Studies*, 2016, 83 (1), 87–117.
- Buera, Francisco J and Ezra Oberfield, “The global diffusion of ideas,” Technical Report, National Bureau of Economic Research 2016.

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 US job creation and economic dynamism," *Journal of Economic Perspectives*, 2014, 28 (3), 3–24.

Dix-Carneiro, Rafael and Brian K Kovak, "Trade liberalization and regional dynamics," *American Economic Review*, 2017, 107 (10), 2908–46.

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.

Garcia-Macia, Daniel, Chang-Tai Hsieh, and Peter J Klenow, "How Destructive is Innovation?," 2016.

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.

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.

Klette, Tor Jakob and Samuel Kortum, "Innovating firms and aggregate innovation," *Journal of Political Economy*, 2004, 112 (5), 986–1018.

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

Lincoln, William, Andrew McCallum, and Michael Siemer, “The Great Recession and a Missing Generation of Exporters,” 2017.

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,” Technical Report, National Bureau of Economic Research 2015.

Rivera-Batiz, Luis A and Paul M Romer, “Economic integration and endogenous growth,” *The Quarterly Journal of Economics*, 1991, 106 (2), 531–555.

Trefler, Daniel, “The Long and Short of the Canada-U. S. Free Trade Agreement,” *American Economic Review*, September 2004, 94 (4), 870–895.