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

March 21, 2019

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 steeper 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, thereby speeding the flow of technology across countries. The resulting dynamic gains from trade liberalization are 2.5 to 5 times larger than the gains in a standard static model.

*We thank Beiling Yan of Statistics Canada for her assistance with the Canadian manufacturing data, 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. 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 (2018). 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 (2017), and Akcigit, Ates and Impullitti (2018).

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, Hsieh and Klenow (2018). 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 (2018), 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. We match TFP growth, growth in research investment, exports relative to manufacturing 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. 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. than in the rest of the OECD to match the lower GDP per worker in the latter.

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 the U.S. version. 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.[1]

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 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

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

from trade, in consumption-equivalent terms, are much higher than in a model with no changes in technology: 2.5 times bigger in the U.S., and 5 times bigger in the rest of the OECD.

We also do a calibration to fit U.S. and Canadian manufacturing. We use it to carry out a counterfactual to fit the change in trade flows after the U.S.-Canada Free Trade Agreement. We then compare the model's quantitative predictions to the changes in job reallocation that we observed in the data, and find them broadly similar. The model predicts persistently higher job reallocation rates after trade liberalization, much of it at large and newly exporting firms.

To dissect our dynamic gains from trade, we entertain alternative assumptions about idea flows across countries. When we assume countries learn partially from *domestic producers* rather than from sellers into the domestic 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 rather than import, however, the dynamic gains from trade are even larger 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. We did not try disconnecting idea flows from trade entirely, because then trade liberalization would not lead to sustained increases in job reallocation as seen in Canada after its 1988 Free Trade Agreement with the U.S.

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 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 multilateral 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. They stress that the dynamic effects of trade could be negative depending on whether firms learn from domestic producers or from sellers into the domestic market.

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 (U.S.-Canada trade liberalization, alternative assumptions about idea flows). Section 6 concludes.

2 Facts from Canadian and U.S. Manufacturing

We use data from the U.S. Longitudinal Business Database (LBD) and Canada's Annual Survey of Manufactures. The U.S. LBD is based on administrative employment records of every nonfarm private establishment in the U.S. We have this data every year from 1977 to 2013, and restrict the sample to establishments owned by firms that own at least one manufacturing establishment in the given year. We 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. The Canadian data covers all but the smallest manufacturing establishments every year from 1973 to 2012.[2] The variables we use from the U.S. LBD and Canadian manufacturing data are the plant and firm identifiers, employment and industry (four digit SIC or six digit NAICS). The Canadian data has information on exports every five years from 1974 to 1989, 1993 and every year from 1996 to 2012. The U.S. 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. We highlight nine 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, 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] The survey threshold is currently annual sales of 30 thousand Canadian dollars.

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

[4] The 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	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 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.

2. **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 84% of all job destruction in the U.S. and 48% of all job destruction in Canadian manufacturing.

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]

[5] Lincoln, 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.

Table 2: 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. Job creation and destruction calculated over five year periods.

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 1988–2003 ("Post-CUSFTA"). We highlight three facts in the Canadian data:

4. **Job flows increased after 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 1988–2003 against the change in the tariff rate on Canadian imports in those same industries over the same period due to CUSFTA.[7] Job destruction rates increased more in industries where tariffs declined the most, thought the relationship is not particularly tight.[8]

[6]The 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.

[7]We use the tariff cuts 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.

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

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

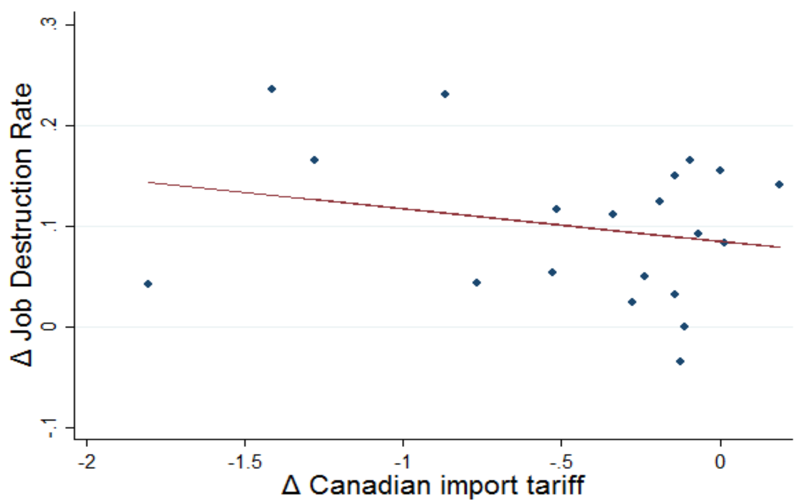

Note: Each observation is a two digit Canadian industry. Δ job destruction is the difference between the average job destruction rate (calculated over five years) in 1988 to 2003 and 1978 to 1988.

5. **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 2 documents a similar fact in Canada. The job destruction rate among large (above-mean employment) firms increased by 2 % points after CUSFTA, out of an overall increase in job destruction of 7 % points.

6. **Job creation from exports increased after trade liberalization.** The last row in Table 2 shows that job creation from exports increased by almost 9 % points in Canada after the trade agreement. Figure 2 shows that a similar fact holds across two-digit Canadian industries. Job creation from exports increased more in sectors where U.S. tariffs declined the most, though again this relationship is not a tight one.[9]

[9]The coefficient of the OLS regression in Figure 2 is $-.189$ with a standard error of $.096$.

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

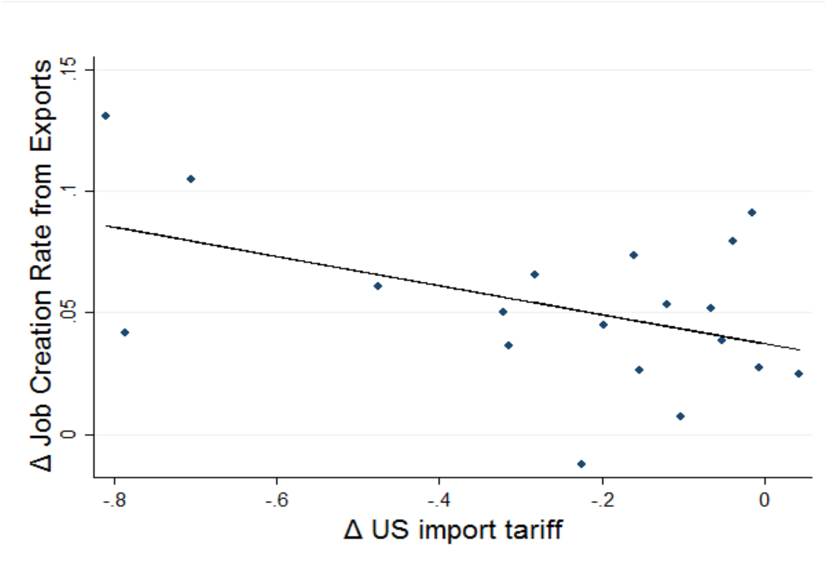

Note: Each observation is a two digit Canadian industry. Δ job creation from exports is the difference between the average job creation rate from exports from 1988 to 2003 and 1988 to 1998, both calculated every five years.

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 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 6 percentage points (row 2).[10] Second, the increase in job destruction was entirely driven by large firms. Third, there is no U.S. export data prior to 1987, but job creation from exports was stable at 2% between 1987–1992 and 1992–2012. If we only consider manufacturing establishments, job creation from exports increased modestly from 2.7% to 3.1%.

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

Table 3: 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.

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 (revenue per worker, 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. **Overlap of labor productivity and employment between exporters and non-exporters.** Figure 3 makes clear that the distributions overlap substantially. Many exporting firms are smaller than non-exporters, and many exporters have lower labor productivity than many non-exporters.

Our last fact is the rapid and continuous turnover in a country's top export products highlighted by Hanson, Lind and Muendler (2018).

9. **Churning of the top export product.** Table 4 presents two measures of churning of the top exported product for the U.S. and other OECD countries. We measure churn as the share in total exports and the rank among all exported products in year $t-5$ of the top exported product in year t. For the top U.S. export at time t, the ratio of the export share in $t-5$ to the

Figure 3: Distribution of Employment and Labor Productivity

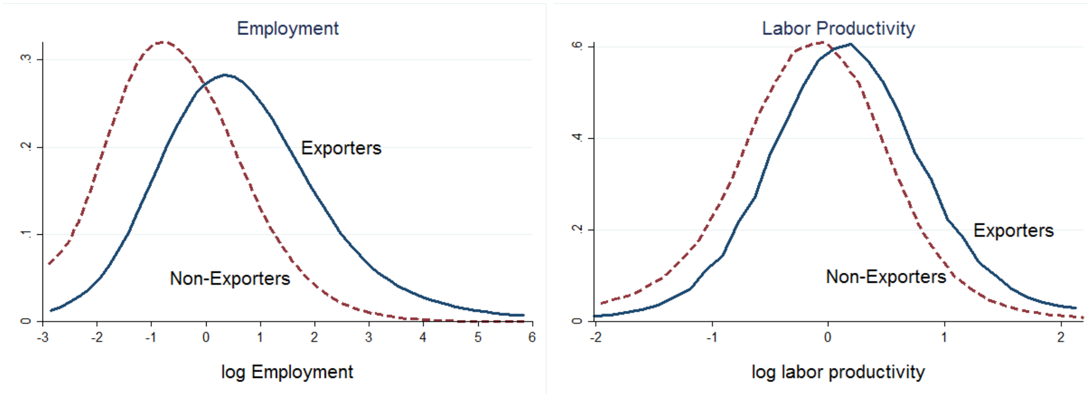

Note: The figures plot the distribution of labor productivity (value-added per worker) and employment of exporting and non-exporting firms in the 2012 U.S. Census of Manufacturing.

Table 4: Churning of Top Export Product

	U.S.	Rest of OECD
Top export share year t	5.7%	12.8%
Share top export product year $t-5$ / year t	66.5%	86.2%
Rank of top export in year $t-5$	5	6

Note: A product is one of 540 4-digit manufacturing industries in Feenstra et al. (2005)'s World Trade Flows. Row 1 shows the share in total exports of the top 4-digit exported product in a given year. Row 2 shows the ratio of the share in total exports of the top export five years before to the export share of the top product in a given year. Row 3 shows the rank of the top product in year t in year $t-5$. Entries are average of overlapping five year panels from 1982–1987 to 2002–2007.

export share of the top product in t is 66.5%. And the top exported product at time t was only the fifth largest exported product five years earlier.[11]

[11] A product in Table 4 is one of the 540 4-digit SITC (revision 2) manufacturing industries in Feenstra et al. (2005)'s World Trade Database (and the extension of this data to 2007). The table shows the average of five-year panels from 1982–1987 to 2002–2007.

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:

$$U = \prod_{j=1}^{M} C_j^{\frac{1}{M}} \tag{1}$$

where M denotes the (fixed) number of varieties.[12] 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

$$\frac{A_{x_1}}{\tau} = \omega A_{x_1}^* \tag{2}$$

$$A_{x_2} = \frac{\omega A_{x_2}^*}{\tau} \tag{3}$$

where ω denotes the relative wage (domestic relative to foreign) and $\tau \geq 1$ is the

[12] We 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 5: Markups

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

symmetric gross trade cost. When $\tau = 1$, $x_1 = x_2$ and all products are traded.

The owner of the best blueprint sets the quality-adjusted 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 5 summarizes the (gross) markup of domestic firms μ_i and foreign firms μ_i^*. Here A_i' and $A_i^{*\prime}$ 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 the assumption that the value of exports is equal to the value of imports:

$$\frac{I^*}{\tau M} \cdot x_1 = \frac{I}{\tau M} \cdot (M - x_2) \tag{4}$$

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

$$I = \frac{\overline{\mu} wL}{1 - \frac{1-\tau}{\tau} \cdot \frac{M-x_2}{M}} \quad \text{and} \quad I^* = \frac{\overline{\mu}^* w^* L^*}{1 - \frac{1-\tau}{\tau} \cdot \frac{x_1}{M}}$$

where $\overline{\mu}^*$ and $\overline{\mu}$ 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

$$\frac{1}{\bar{\mu}} \equiv \frac{\sum_{j=1}^{x_2} \frac{1}{\mu_j} + \frac{1}{\tau} \cdot \sum_{j=1}^{x_1} \frac{1}{\mu_j^f}}{x_2 + x_1/\tau}$$

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

We can now express the real (consumption) 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

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

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

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 5 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.

[13]The expression for nominal income come from equating nominal income to the revenue of local firms plus tariff revenue: $I = \overline{\mu}\, wL + (\tau - 1)\frac{I}{\tau \cdot M}(M - x_2)$ and $I^* = \overline{\mu}^* w^* L^* + (\tau - 1)\frac{I^*}{\tau \cdot M} \cdot x_1$.

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***.[14] Specifically, the quality drawn by an innovation follows a Pareto distribution with shape parameter θ and scale parameter equal to the existing quality level. The average percent improvement in quality of an existing variety (conditional on innovation) is thus $\frac{1}{\theta-1} > 0$.[15] Finally, we assume that entrants have one product, while incumbent firms potentially produce many varieties.

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

[14] If 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, Alp, Bloom and Kerr (2018). See Coe et al. (1997, 2009) for evidence consistent with learning by importing.

[15] In 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: Channels of Innovation

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

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

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 $\tilde{\eta}^* \equiv \eta^*(1-\lambda)(1-\tilde{\eta})(1-\tilde{\lambda}^*)$.

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,

Table 7: Probability of Creative Destruction

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

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.

Table 7 summarizes the probability of creative destruction in the home market and in the foreign market for each type of good. In a closed economy ($\tau=\infty$), the arrival rate of new ideas only depends on innovation by local firms, which is $\lambda+\tilde{\eta}$ in the home market and $\lambda^*+\tilde{\eta}^*$ in the foreign market. With trade, domestic consumers benefit from innovation by foreign firms and foreign consumers benefit from innovation by domestic firms. In the extreme case of completely free trade ($\tau=1$), when the innovation rate is the same in the two countries, the rate of creative destruction is simply the sum of the innovation rates in the two countries. So relative to autarky free trade doubles the rate at which new blueprints are available to consumers in both countries.

More generally, the probability that domestic consumers benefit from foreign innovation 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. 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.

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).

Table 8 summarizes the expected growth rate of the real consumption wage for the cases of complete autarky and free trade.[16] The first row in Table 8 says that under autarky, a country's growth rate only depends on its innovation rate (and the innovation step size). A country with a higher innovation rate grows faster than a country with a lower innovation rate.

The second row in Table 8 shows the growth rate with free trade. The first term is the contribution of 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 $(\lambda + \tilde{\eta}) \min\left[1, \omega^{-\theta}\right]$ is the probability an imported product is taken over by a domestic firm. The second term captures the impact of innovation by *foreign* firms on domestic consumers. The rate at which an imported variety is replaced by another foreign firm is $\lambda^{*} + \tilde{\eta}^{*}$ and the probability a variety produced by a domestic firm is replaced by a foreign firm is $(\lambda^{*} + \tilde{\eta}^{*}) \min\left[1, \omega^{\theta}\right]$.

Given innovation rates in the two countries, the relative wage ω has opposite effects on the contribution of domestic vs. foreign innovation. A higher ω lowers the contribution of innovation by domestic firms and raises the contribution of foreign firms. In a steady state, differences in innovation rates show up as dif-

[16] The expected growth rate of the consumption wage in the home country for the general case is given by $(\lambda + \tilde{\eta})\left[\frac{x_2}{M}\frac{1}{\theta-1} + \frac{M-x_2}{M}\min\left(1, \left(\frac{\tau}{\omega}\right)^{\theta}\right)\left(\frac{\theta}{\theta-1}\max\left[\frac{\omega}{\tau}, 1\right] - 1\right)\right] + \left(\tilde{\lambda}^{*} + \tilde{\eta}^{*}\right)\left[\frac{M-x_2}{M}\frac{1}{\theta-1} + \frac{x_2-x_1}{M}\min\left(1, \left(\frac{\omega}{\tau}\frac{A_j^{*}}{A_j}\right)^{\theta}\right)\left(\frac{\theta}{\theta-1}\max\left[\frac{\tau}{\omega}\frac{A_j}{A_j^{*}}, 1\right] - 1\right) + \frac{x_1}{M}\min\left(1, \left(\frac{\omega}{\tau}\right)^{\theta}\right)\left(\frac{\theta}{\theta-1}\max\left[\frac{\tau}{\omega}, 1\right] - 1\right)\right]$. In a steady state the growth rate in the foreign country is the same as in the home country.

Table 8: Growth Rate under Autarky and Free Trade

	Home	Foreign
Autarky	$(\lambda+\tilde{\eta})\left(\frac{1}{\theta-1}\right)$	$(\lambda^{*}+\tilde{\eta}^{*})\left(\frac{1}{\theta-1}\right)$
Free Trade	$\underbrace{(\lambda+\tilde{\eta})\left[\frac{x_2}{M}\cdot\frac{1}{\theta-1}+\frac{M-x_2}{M}\min\left(1,\omega^{-\theta}\right)\left(\frac{\theta}{\theta-1}\max\left[\omega,1\right]-1\right)\right]}_{\text{domestic innovation}}$ $+\underbrace{\left(\tilde{\lambda}^{*}+\tilde{\eta}^{*}\right)\left[\frac{M-x_2}{M}\cdot\frac{1}{\theta-1}+\frac{x_2}{M}\min\left(1,\omega^{\theta}\right)\left(\frac{\theta}{\theta-1}\max\left[\frac{1}{\omega},1\right]-1\right)\right]}_{\text{foreign innovation}}$	

ferences in the relative wage, but the real consumption wage grows at the same rate in the two countries.[17] Figure 4 shows the relative 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.[18] 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 manufacturing sectors in the U.S. and the rest of the OECD ("foreign"). The shape pa-

[17] This result is reminiscent of Acemoglu and Ventura (2002).

[18] The values of the other parameters used in the simulation behind Figure 4 are $\eta = .0157$, $\lambda^{*} = .083$, $\eta^{*} = .025$, $\theta = 4.84$, and $\tau = 1.474$. 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

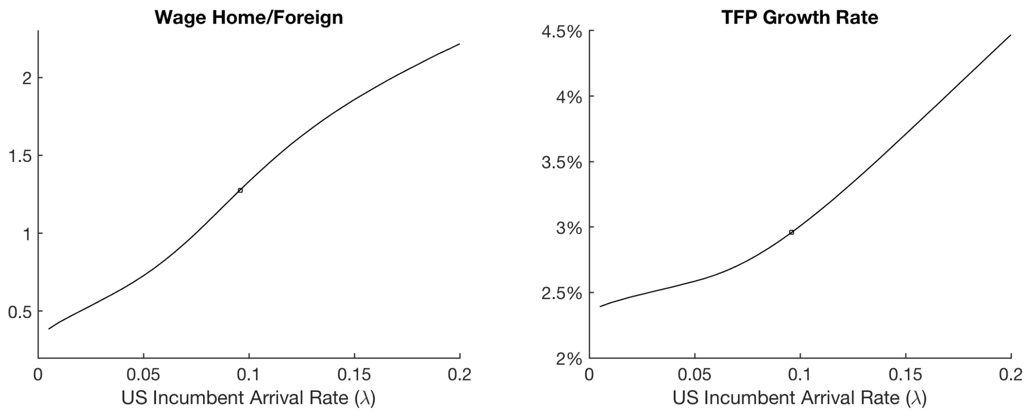

Note: The left panel shows the steady-state wage at home relative to the foreign country as a function of the innovation rate of domestic incumbent firms, holding fixed the other variables. The right panel shows the effect on the steady state growth rate of aggregate TFP in *both* countries.

rameter 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.[19] We use the employment share of new firms to pin down the innovation by entrants vs. incumbents.

The data moments we use are displayed in Table 9. 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%.[20] The innovation rates and trade cost needed to fit these facts are shown in Table 10.[21] The innovation rates

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

[20]Overall OECD employment was not contracting during our sample, but employment in manufacturing was. We target contracting employment to fit net job creation in the U.S. Negative employment growth poses no problem for obtaining steady state TFP growth and relative wages with exogenous innovation. But we will revisit this for endogenous innovation.

[21]We 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

Table 9: 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	16.9%
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

in Table 10 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.[22]

3.4 Firm Dynamics

We now show that the model can generate the nine facts described in Section 2. The driver of growth here is the creation of better quality products and the resulting destruction of product lines of incumbent producers. In our open economy setting, the destruction of product lines can come from firms located

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.

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

Table 10: Estimates of Model Parameters

Variable	Description	Value
θ	Shape parameter of innovation draws	4.84
λ	Home innovation rate incumbent	9.65%
η	Home innovation rate entrants	1.90%
μ^*	Foreign innovation rate incumbents + entrants	10.81%
τ	Trade cost	1.474
ρ	Quality mean reversion parameter	0.92

in other countries. This has four implications:

First, growth is associated with job creation by innovating firms and job destruction by firms whose products are innovated upon. Table 11 (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 9. 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 11 replicates the U.S. data. The job flows predicted by the model with the parameters in Table 10 are roughly of the same magnitude as in the data (fact 1).

Second, firms shrink 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 11 shows that, consistent with the evidence from U.S. and Canadian manufacturing (fact 2), job destruction in the model is mostly driven by large (above-mean employment) firms. The job destruction rate by large firms is 22% in the model, 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 domes-

Table 11: Firm Dynamics, Simulations vs. Data

Moment	U.S. Data	Simulations
Job Creation Rate	31.4%	30.6%
Job Destruction Rate	36.6%	37.1%
Job Destruction from Large Firms	30.7%	21.0%
Job Creation from Exports	2.0%	6.7%
Job Destruction from Imports	–	7.8%
Top Export Product Turnover	66.5%	68.4%

The U.S. data is the average from 1987 to 2012. The second column show the simulated moments in the steady state of the model with the parameter values from Table 10.

tic firms replacing producers in the foreign market. The fourth row in Table 11 shows that the job creation rate from exports in the model is 3.6%. In the U.S. data, the number is 2%. 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 7.8%, so about a fifth of the overall job destruction in the model comes from creative destruction by foreign firms.

Fourth, another implication of the fact that foreign firms replace domestic firms is some of the country's exported products are replaced by a foreign producer. If an exported product is replaced by a another domestic firm, the country will continue to export the given product. In our baseline simulation, the product with the highest level of exports tends to have lower export shares five years previously, at 68.4%. In the data, this number is 66.5%.

Table 11 shows that the model can replicate, at least qualitatively, facts 1-3 and 9 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 4-6 (job flows, job destruction at large firms, and job creation due to exports all increase after trade liberalization). 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.[23]

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 domestic 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 ($\tau = 1.474$), 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 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

[23]We relax this assumption in the next section, when we endogenize innovation rates.

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

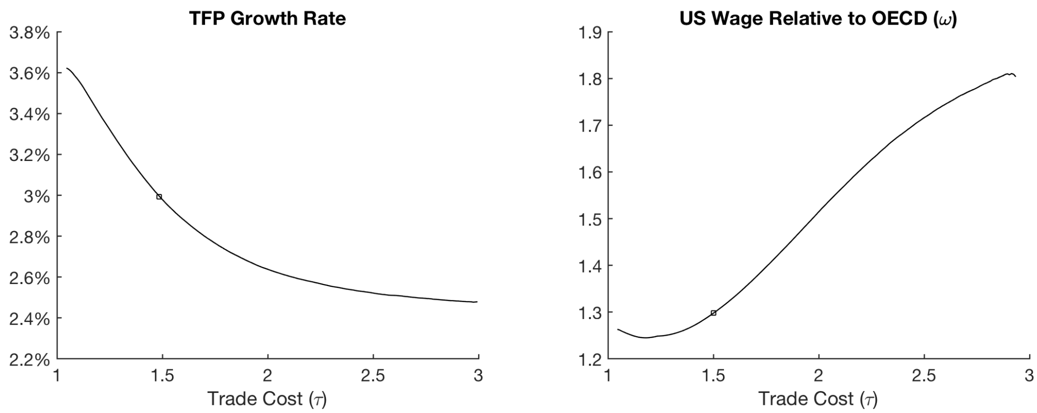

Note: The left panel shows simulated growth rate of the real wage and the right panel the relative wage (home/foreign) in the steady state for different values of τ. All other parameters of model are kept fixed at the values in Table 10.

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

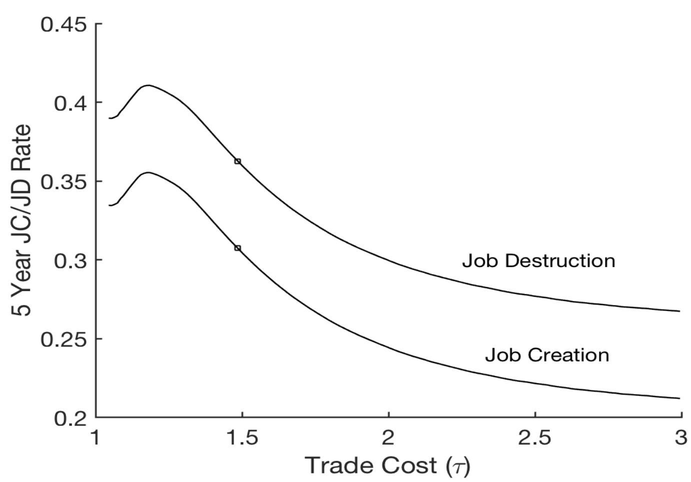

Note: The figure simulates the steady state job creation and destruction rates when we vary τ but keeping constant all other parameters at the values in Table 10.

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

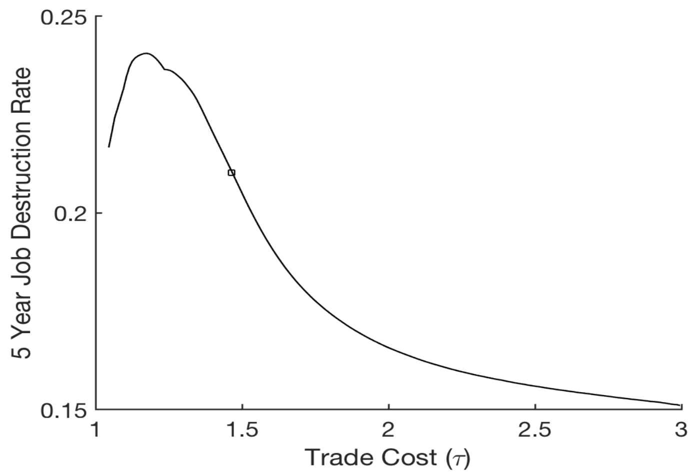

Note: The figure simulates the steady state job destruction rate from large (above mean employment) firms when we vary τ but keeping constant all other parameters at the values in Table 10.

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.[24]

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 creation rate from exports increases by 20 percentage points when trade costs fall from our benchmark value ($\tau = 1.474$) 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.

[24]We 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 8: Simulated Job Flows from Trade vs. Trade Costs

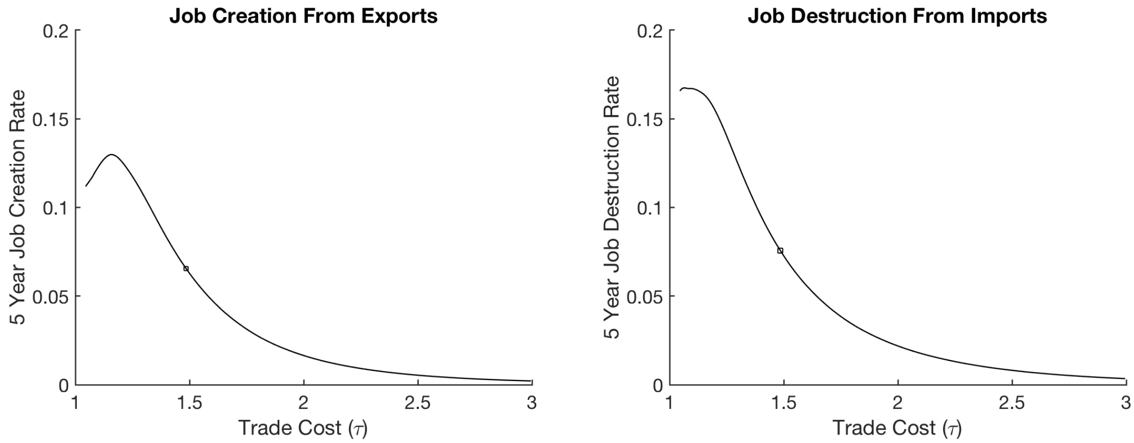

Note: The left panel shows the job creation rate due to exports and the right panel the job destruction rate due to imports in the steady state for different values of τ. All other parameters of model are kept fixed at the values in Table 10.

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 10. 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 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.

Figure 9: Simulated Distribution of Employment and Labor Productivity

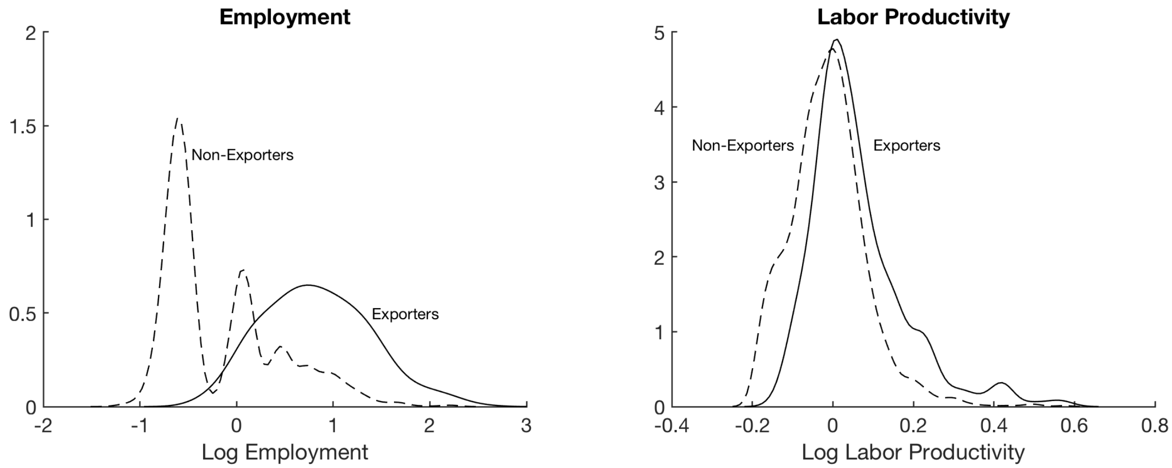

Note: The distributions of employment (left panel) and labor productivity (right panel) in the steady state of the model with the parameters given in Table 10.

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 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.[25]

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 assumption 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.[26] 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

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

[25]As stated in Table 9, this gap is 6.6%.

[26]Bartelsman, Haltiwanger and Scarpetta (2013) emphasis the role of overhead costs, Asker, Collard-Wexler and De Loecker (2014) the importance of adjustment costs, and Bils, Klenow and Ruane (2018) the contribution of measurement error.

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

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

where R_e is labor used for research (per variety in the economy) by potential entrants and χ_e is an efficiency parameter.[27]

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 $\bar{A}^{1-\phi}$. We define v_x and v_n as the *expected* normalized value of a traded and non-traded product.

[27]The 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 following arbitrage equation pins down v_x at time t:

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

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.[28] The first term in equation (9) is profit net 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

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

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

[28] δ_x and δ'_x are the creative destruction rates in Table 7. Since innovation is not targeted, β_x and β_n depend on the share of each type of product and the creative destruction rates.

revenue from innovation to the marginal cost of innovation, which yields:

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

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

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).[29]

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 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).[30] This yields $\phi = 0.17$ and $\gamma = 0.55$. The calibrated values of these parameters are given in Table 12. 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 per-

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

[30]TFP 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.

Table 12: Estimates of Model Parameters, Endogenous Innovation

Variable	Description	Value
ϕ	Return to stock of ideas	0.097
γ	Return to research intensity	0.6
χ_e/χ_i	Home entrant/incumbent research cost	3.44
χ_i^*/χ_i	Foreign/home incumbent research cost	4.44
χ_e^*/χ_i	Foreign entrant/home incumbent research cost	10.6

manent 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 10 plots the effect of a permanent, unanticipated reduction of trade costs, from $\tau = 1.474$ to $\tau = 1.237$, 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 innovation makes higher research effort worthwhile despite the endogenously

Figure 10: Simulated arrival rates after trade liberalization

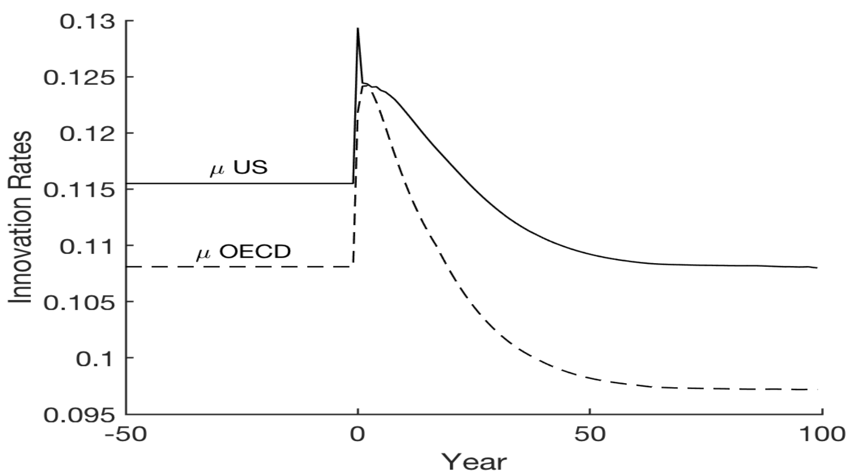

Note: The figure simulates the path of arrival rates in response to an unanticipated reduction in trade costs, keeping constant all other parameters at the values in Table 12. Here $\mu = \lambda + \eta$, where λ and η are the endogenous arrival rates of innovation by incumbents and entrants, respectively.

Figure 11: Research labor shares after trade liberalization

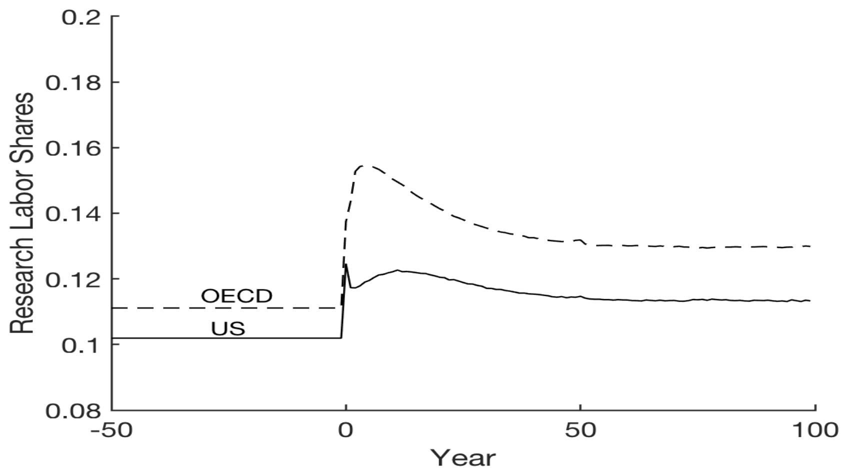

Note: The figure simulates the share of research labor in total employment in response to an unanticipated reduction in trade costs, keeping constant all other parameters at the values in Table 12.

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 tariffs. The real wage actually falls, however, as frictionless trade is approached.

Figure 12: Simulated real consumption after trade liberalization

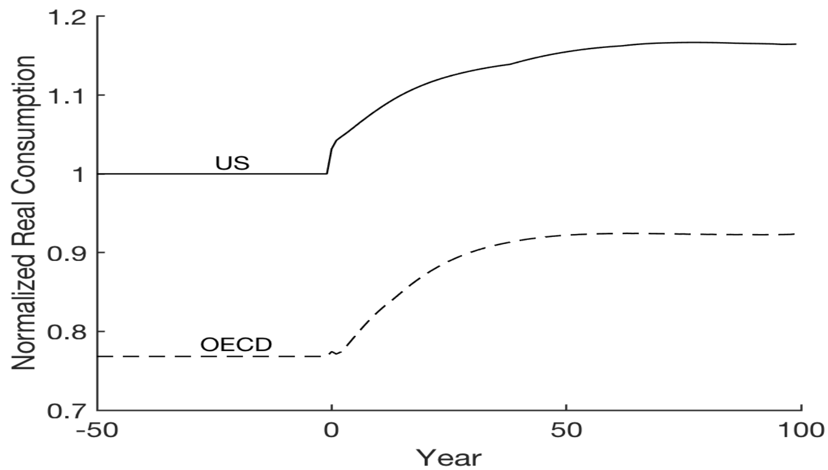

Note: The figure simulates the path of real consumption in response to an unanticipated reduction in trade costs, keeping constant all other parameters at the values in Table 12.

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 15-year window we examine after the 1988 U.S.-Canada Free Trade Agreement. This pattern drives home that there can be dynamics costs (job destruction) as well as dynamic benefits (knowledge flows) to trade liberalization.[31]

In Table 14 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%. In the first counterfactual, we reduce tariffs from 3, a high tariff at which the aggre-

[31] Bernard, 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

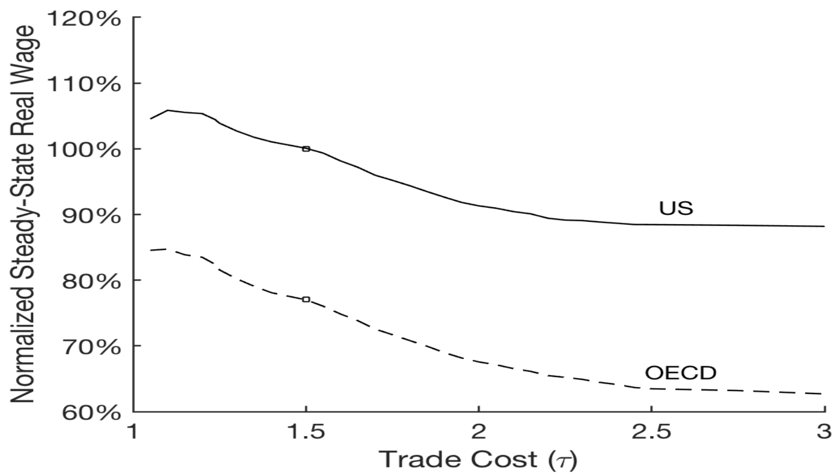

Note: The figure plots the simulated equilibrium real wage in steady-states with differing trade costs, keeping constant all other parameters at the values in Table 12.

Figure 14: Simulated job flows from trade liberalization

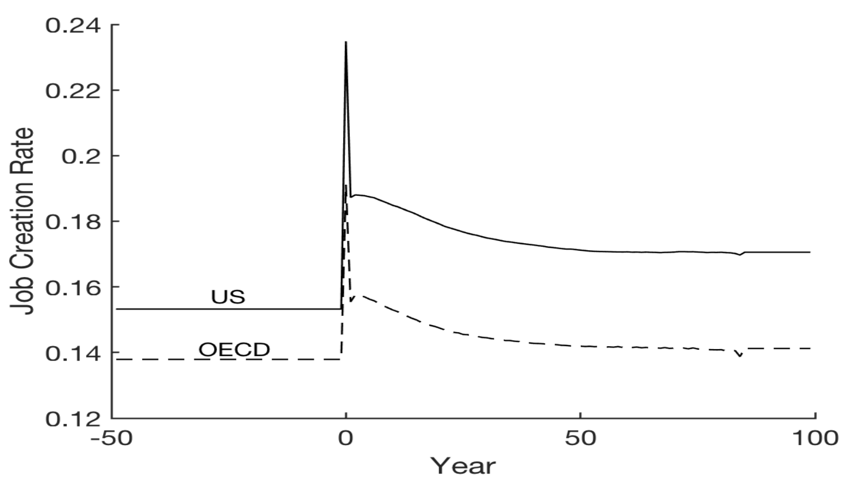

Note: The figure simulates job creation and destruction rates in response to an unanticipated reduction in trade costs, keeping constant all other parameters at the values in Table 12.

gate trade share is about 1%, to our estimated value of 1.474. For comparison, we first calculate the static gains — the jump in real consumption on impact. The static gains are 8.3% in the U.S. and 4.8% 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.[32]

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 14 reduces tariffs from our estimated value of 1.474 to 1.237. The dynamic gains are smaller here than when going from near-autarky to current trade flows. The ideas flowing after further trade liberalization involve endogenously smaller step sizes. The dynamic gains are still three to seven times bigger than the static gains.

[32]Our static gains from trade are larger than implied by the formula in Arkolakis, Costinot and Rodriguez-Clare (2012). When the trade share goes from 1% to 10.2%, their formula predicts gains of 1.9% (9.2%/5), if one uses our trade elasticity of 5. This could be because our Cobb-Douglas preferences result in less substitutability across products than in the class of models ACR examine.

Table 13: Gains From Trade

	Real Wage		Real Consumption	
Method	U.S.	OECD	U.S.	OECD
Static Gains	8.3%	5.1%	8.3%	4.8%
Dynamic Gains - Exogenous Innovation	21.8%	38.4%	25.7%	45.3%
Dynamic Gains - Endogenous Innovation	18.1%	24.1%	21.9%	26.2%

Entries give the percentage increase in the present discounted value of the real (consumption) wage and real consumption as a result of reducing tariffs from 3 to 1.474. The aggregate trade share at $\tau = 3$ is about 0.25%. In the second counterfactual, we reduce tariffs from our estimated value of 1.474 to 1.237. We use a discount rate of 5%.

Table 14: Gains From Trade Liberalization

	Real Wage		Real Consumption	
Method	U.S.	OECD	U.S.	OECD
Static Gains	6.5%	4.6%	4.9%	3.0%
Dynamic Gains - Exogenous Innovation	20.5%	23.9%	18.9%	22.6%
Dynamic Gains - Endogenous Innovation	15.5%	18.1%	12.2%	14.1%

Entries are the percentage increase in the present discounted value of the real wage and real consumption as a result of reducing tariffs from 1.474 to 1.237. We use a discount rate of 5%.

5 Additional Exercises

5.1 U.S.-Canada Simulations

The U.S.-Canada Free Trade Agreement provided three of the nine motivating facts we offered 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 a simplified version of CUSFTA in our model with endogenous arrival rates.[33] Specifically, we analyze the effect of lowering tariffs from 1.47 to 1.38 to match Canada's trade shares of 25% before CUSFTA (1978–1988) and 37% after CUSFTA (1988–2003). We re-estimate other parameter values so that the model matches the relative wage and population between the U.S. and Canada.

Table 15 presents the changes in job flow moments in Canada after trade liberalization.[34] The first column has (changes in) empirical moments, and the second column the changes in simulated moments. The model qualitatively matches the patterns in the data: rising job creation, job destruction, job destruction from large firms, and job creation from exports. The best fit is for job creation from exports, which rose 9 percentage points in both the data and simulations. The increase in job destruction in the data of 7 percentage points was somewhat larger than the 5 percentage point increase in the simulations. The simulations markedly overstate the rise in job creation and in job destruction at large firms (1-2 percentage points in the data versus 5-6 percentage points in the simulations). Note that no parameter values were chosen to target any of the data moments in Table 15.

[33] Until now, we have carried out simulations with the rest of the OECD because its larger size helped drive home the importance of scale for the dynamic gains from trade. Because the rest of the OECD is so large, trade flows with them are larger and, presumably, idea flows as well.

[34] We 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 15: Job Flows in Canada — Post-CUSFTA versus Pre-CUSFTA

	Data	Simulations
Job Creation Rate	+1.3%	+5.0%
Job Destruction Rate	+7.2%	+5.0%
Job Destruction from Large Firms	+1.9%	+6.0%
Job Creation from Exports	+8.7%	+9.3%

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

5.2 Alternative Assumptions about Idea Flows

All of our simulations so far assume that innovators build on the productivity level of *sellers* into the domestic market. To gauge the importance of this assumption for the dynamic gains from trade, we now consider two alternative assumptions about idea flows across countries. For simplicity, these simulations involve exogenous arrival rates of innovation. And they involve going from near-autarky to current levels of trade.

We first assume that innovators probabilistically build on sellers (with probability κ) or on domestic *producers* (with probability $1-\kappa$).[35] We next assume that innovators build on all products (with probability ν) or on the subset of products that are domestically produced (with probability $1-\nu$). Such research specialization ($\nu < 1$) allows countries to experience faster arrival of innovations on the subset of produces they produce, and more so the higher the share of products imported.

[35] If a product is imported, innovators build on the quality of the last domestic producer of the product. There is sufficient turnover of imported products to domestic producers so that there is always a previous (fallow) domestic producer.

Table 16 presents our findings. Column one reproduces our baseline specification of $\kappa = 1$ and $\nu = 1$ to facilitate comparison. Column two reduces the frequency with which innovators learn from sellers to 50% ($\kappa = 0.5$), but does not allow research specialization ($\nu = 1$). Doing so dampens the increase in job reallocation from trade in the U.S. and OECD, but only slightly. The U.S. growth kick from trade is reduced substantially, however, from 56 basis points per year in the baseline to 31 basis points per year with limited idea flows. The rest of the OECD likewise receives a smaller TFP growth boost of 38 basis points versus 66 basis points in the baseline.[36] The consumption-equivalent gains from trade fall from 26% to 15% in the U.S., and from 45% to 36% in the rest of the OECD.

Column three of Table 16 allows for research specialization along with limited idea flows ($\kappa = 0.5$ and $\nu = 0.5$). This has a more material effect on job reallocation: when innovators specialize more on domestically produced products, the country producing a product does not turn over as frequently. Trade boosts job reallocation in the U.S. closer to 6 % points rather than the 9 % points in the baseline. The increase in the U.S. growth rate, however, is essentially the same at 57 basis points as the baseline's 56 basis points. The U.S. actually gains more in consumption-equivalent terms in this case, 43% versus 26% in the baseline. Thus, even if ideas flow in a more limited away, trade may create large dynamic gains if they facilitate research specialization.[37]

How can we assess the realism of these assumptions? All three cases feature increases in job reallocation after trade liberalization, as seen in Canada after CUSFTA.[38] But some insight can be gleaned from product turnover. Recall that we reported in Table 4 that SITCs with the highest level of exports in a given year tend to have lower export shares five years previous (66.5%). In our baseline

[36] Recall 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.

[37] The dynamic gains are smaller for the rest of the OECD despite research specialization. Their arrival rates are lower, so the rest of the OECD benefits less from specialization and relies more on idea flows from the U.S.

[38] As we noted in the Introduction, we did not entertain a case where trade has no impact at all on growth and job reallocation. The reason is that such a model could not explain the elevated job reallocation observed after CUSFTA.

Table 16: Counterfactuals with Alternative Assumptions on Idea Flows

Near-Autarky to Estimated τ	$\kappa = 1$, $\nu = 1$	$\kappa = 0.5$, $\nu = 1$	$\kappa = 0.5$, $\nu = 0.5$
Δ Job Reallocation U.S.	9.8%	9.3%	6.4%
Δ Job Reallocation OECD	5.6%	5.2%	4.3%
Δ Growth U.S.	0.56%	0.31%	0.57%
Δ Growth OECD	0.66%	0.38%	0.54%
Gains from Trade U.S.	25.7%	14.9%	42.7%
Gains from Trade OECD	45.3%	35.7%	18.9%

Note: Column one is the baseline specification. Column two reduces the frequency with which innovators learn from foreign producers to 50% ($\kappa = 0.5$); the other 50% of learning is from domestic producers. Column three allows for research specialization in the reduced spillover case (that is, it allows 50% of innovation draws to be specialized in products a country producers ($\nu = 0.5$); the other 50% of draws are on all products. In each case, all parameters are re-estimated to match the targeted moments and counterfactuals are shown for the exogenous arrival rate near-autarky and estimated steady-states. Estimated τ is 1.47 for column one, 1.51 for column two, and 2.0 for column three. The gains from trade are in terms of consumption-equivalent welfare.

simulation, the fraction is similar at 68.4%. In simulations with $\kappa = 0.5$ and $\nu = 1$, the fraction rises to 70.3%. It will rise further if we move to even lower κ values. Specialized research ($\nu < 1$) further lowers product turnover. In the case where $\kappa = 0.5$ and $\nu = 0.5$, the export share of the top product five years earlier was 93.8% of its current share. Thus deviating from our baseline would move us further away from the data on product churn.

When we consider even lower values of κ and ν we find that growth rates diverge between the U.S. and Canada or the rest of the OECD in our simulations. Thus the similarity of observed growth rates across these countries is another piece of evidence in favor of tying them together through idea flows.[39]

[39] See Klenow and Rodriguez-Clare (2005) for such evidence.

6 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 more direct 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). As we stressed, we think idea flows will need to be tied to trade flows somehow 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 "Global Technical Change" accords 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," *Quarterly Journal of Economics*, 2002, *117* (2), 659–694.

__ , Ufuk Akcigit, Harun Alp, Nicholas Bloom, and William Kerr, "Innovation, Reallocation, and Growth," *American Economic Review*, 2018, *108* (11), 3450–91.

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.

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," *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?," 2018.

Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb, "Are Ideas Getting Harder to Find?," 2018.

_ , 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.

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

Dornbusch, Rdiger, 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?," 2018.

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.

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," 2015.

Rivera-Batiz, Luis A. and Paul M. Romer, "Economic Integration and Endogenous Growth," *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.