

Better, Faster, Stronger: Global Innovation and Trade Liberalization*

Federica Coelli[†] Andreas Moxnes[‡] Karen Helene Ulltveit-Moe[§]

May 2016

Abstract

This paper estimates the effect of trade policy during the Great Liberalization of the 1990s on innovation in over 60 countries using international firm-level patent data. The empirical strategy exploits ex-ante differences in firms' exposure to countries, allowing us to construct firm-specific measures of tariffs. This provides a novel source of variation that enables us to establish the causal impact of trade policy on innovation. Our results suggest that trade liberalization has economically significant effects on innovation and, ultimately, technical change. According to our estimates, about 7 percent of knowledge creation during the 1990s can be explained by trade policy. Furthermore, we find that the increase in patenting reflects innovation, rather than simply more protection of existing knowledge. Both improved market access and more import competition contribute to the positive innovation response.

*We thank Bjarne J. Kvam from the Norwegian Industrial Property Office for helpful clarifications and comments.

[†]University of Oslo, federica.coelli@econ.uio.no

[‡]University of Oslo, CEPR & NBER, andreas.moxnes@econ.uio.no

[§]University of Oslo & CEPR, k.h.ulltveit-moe@econ.uio.no

1 Introduction

Trade policy liberalization opens up new markets for exporters but also increases the competitive pressure in the home market. Both mechanisms are likely to affect the innovation rate in the economy, and thereby affect the rate of economic growth. But what is the net impact of trade policy on innovation and can we disentangle the impact of market access and import competition? These are the main questions of this paper.

Our approach is as follows. During the 1990s, tariffs in both developing and developed countries came down substantially, leading researchers to name the period the Great Liberalization of the 1990s (Estevadeordal and Taylor, 2013). Those reductions were predominantly a result of the GATT Uruguay Round, spanning the years 1986 to 1994 and phased in from 1995 to 2000, but also a result of regional trade agreements and unilateral liberalization. On average, developed country tariffs were cut from 10 to 5 percent, while developing country tariffs were cut from 30 to 15 percent (Estevadeordal and Taylor, 2013).

We use the Great Liberalization as a quasi-natural experiment, and estimate the causal impact of tariffs on innovation by using firm-level variation in which countries firms were exposed to before the tariff cuts. Intuitively, a firm x located in Germany and selling to the U.S. and Mexico is affected differently than a Japanese firm y selling to China and South Korea because tariff cuts vary across countries and industries. Furthermore, a German firm z selling only to Germany is again affected differently because that firm does not benefit from improved market access abroad but is potentially hurt by import competition in its home market.

The data requirements for this exercise are large: one would need a firm-level panel dataset on innovation over a long time period, along with detailed information on where firms are located and in which markets they sell in. For this, we construct a global and comprehensive firm-level dataset on patenting using PATSTAT from the European Patent Office. In our data, we observe nearly every firm worldwide that files a patent, in which country (patent office) they file, along with their industry and home country affiliation, over four decades. Unfortunately, we do not directly observe in which markets firms sell in, but we do observe where firms are patenting before the Great Liberalization. We therefore construct a firm-level measure of country exposure by using patent information up until 1985, one year before the Uruguay negotiations started. Of course, those patent weights are not perfectly correlated with sales weights, but we provide detailed evidence that the correlation is indeed high (Section 2 and Appendix Section F). Those weights are also remarkably persistent over time, even after a period of 20 years, suggesting that geographic frictions (e.g., shipping costs on the supply side, or idiosyncratic taste differences on the demand side) are severely

limiting where firms sell and file patents (Appendix Section F).

Our firm-level approach has a number of advantages. First, because ex-ante country exposure varies significantly within a country and within narrowly defined industries, we can sweep out all home country-industry trends in innovation by fixed effects. Second, because we observe aggregate patenting in all countries and industries, we can flexibly control for all other factors that are correlated with tariff cuts and also affect innovation. An example of this is market size: perhaps being exposed to a high tariff cut country fosters innovation simply because that country grows faster. Controlling for aggregate patenting in that destination country will eliminate this concern. Third, our long time period allows us to perform placebo tests; to test if treated firms (exposed to high-tariff cut countries) typically always patent more. Specifically, we check whether patenting in the 1980s responds to tariff cuts in the 1990s. The answer is no; our estimates then becomes noisy and close to zero.

Our results show that the Great Liberalization of the 1990s on net had a large positive impact on innovation. Overall, our estimates can explain roughly 7 percent of aggregate patenting over the period. This suggests that innovation was one important channel by which trade policy liberalization improved growth during this period. Furthermore, our decomposition exercise shows that both improved market access and tougher import competition have large and positive effects on innovation. Furthermore, the economic magnitude of the two mechanisms is similar.

Does increased patenting reflect increased innovation? The previous literature typically finds a strong correlation between patenting and research and development, and between patenting and other measures of innovation (Griliches, 1990). We also find a strong positive correlation between patent counts and innovation indicators in our own data (Appendix Section H). Also, our firm-level identification strategy ensures that all regulatory changes in the patent system, or differences across patent offices, are differenced out by fixed effects. But the concern remains that more trade could induce greater protection of intellectual property rights (IPR), i.e. that more patenting is simply a “lawyer effect”. To deal with this, we calculate citation counts for all firms in our dataset, and check whether average citations are falling in response to trade liberalization. This would indeed be the case if import competition induced firms to take out more patents to protect marginal inventions. The data rejects this hypothesis, if anything, average citations are rising in response to trade liberalization.

The contributions of this paper are as follows. First, we provide broad and systematic evidence of the impact of trade policy on innovation for a large set of countries over a decade with steep tariff declines. This provides external evidence compared to the current literature that has primarily focused on relatively narrow policy changes (e.g., Bustos, 2011).

Also, while there is a large literature on the impact of trade policy on firm performance (e.g., TFP or labor productivity), there is little evidence on observable output or input measures of innovation (e.g., patents and research and development, respectively). Second, we disentangle the import competition from the market access effect of trade policy, which not only informs the literature on trade policy but also the broader literature on the effects of competition on innovation (e.g., Aghion et al., 2005) and market size on innovation (e.g., Acemoglu and Linn, 2004). Third, we construct and analyze a novel, comprehensive and global firm-level patent dataset that has so far not been applied in the context of international trade.

Our analysis speaks to different strands of literature. First, our work is related to the empirical analyses of firm level data on the impact of trade on firm performance. Halpern, Koren, and Szeidl (2015) estimate a model of importers using Hungarian micro data and find that importing more varieties leads to large measured productivity effects. Recent work by Gopinath and Neiman (2013) also find large negative measured productivity effects from a collapse in imports following the Argentine crisis of 2001-2002. The empirical studies of Amiti and Konings (2007), Goldberg et al. (2010) and Khandelwal and Topalova (2011) all find that declines in input tariffs are associated with sizable measured productivity gains. Compared to our work, these papers focus on the impact of trade on firm performance but do not separately identify what are the channels that allow for the benign impact of trade on innovation.¹ Along the same line of work, but somehow closer to this paper is Boler et al. (2015) who explores the complementarities between international sourcing of intermediates and R&D investment and their joint impact on firm performance.

Second, our work relates to the literature on complementarities between exports and technology adoption. Closest in the spirit to our analysis is empirical work by Bustos (2011) and Lileeva and Trefler (2010) who show that trade integration can induce exporters to upgrade technology, Bloom, Draca, and Van Reenen (2016) who focus on the effect of imports from China on technology upgrading and productivity in OECD countries, and Teshima (2009) who examines the impact of reduced output tariffs on Mexican firms and finds that the reduction in Mexican output tariffs increased innovative activity of Mexican firms due to increased competition. What distinguishes our paper from these contributions is the fact that we focus on the global impact of multilateral trade liberalization rather than on unilateral or bilateral trade liberalization episodes. Moreover, our international firm-level data set and the high number of countries in our sample provide external validity. Finally, our paper is also related to Aghion et al. (2014) and Calel and Dechezleprêtre (2014), who

¹Note that Goldberg et al. (2010) find that lower input tariffs are associated with increased R&D expenditures.

also use PATSTAT data and a related empirical approach, although they focus on very different questions, being the impact of environmental policies on technical change.

The rest of the paper is organized as follows. Section 2 presents our empirical framework, while Section 3 describes the data and descriptive statistics. Section 4 discusses empirical results while Section 5 concludes.

2 Empirical Framework

2.1 Basic Setup

This section presents a basic economic framework that will guide the subsequent empirical work. We choose not to specify a full model here in order to save space and to focus on the key relationships that will inform the econometric specification. Consider the following profit function of a firm i ,

$$\pi_i = \sum_{n \in \Omega_i} \pi_{in} = z_i \sum_{n \in \Omega_i} \tau_n^{\beta_n} \epsilon_{in},$$

where z_i is productivity, τ_n is the iceberg-equivalent tariff in country n , β_n is an unknown parameter and ϵ_{in} is a residual capturing all other country-specific factors that also determine profits, such as demand in country n .² To ease notation, we drop industry subscripts here, but we will exploit industry variation in τ_n later in the paper. We sum over all countries n where the firm has positive sales (either through exporting or multinational production). This is the set Ω_i .

The set of countries where the firm has positive sales is varying across firms but taken as exogenous. This is motivated by the empirical fact that there is a high degree of persistence in country-specific export participation (Moxnes, 2010).³ Appendix Section F provides empirical evidence that patenting to a country is remarkably persistent over time. Note the similarity between our profit function and gross profits in a standard trade model with monopolistic competition. In that framework, ϵ_n would capture the size of the market and the aggregate price index in n , while β_n would equal $1 - \sigma$, where σ is the elasticity of substitution.

The profit function captures two main ideas. First, profits are increasing with effective market size, $\sum_{n \in \Omega_i} \tau_n^{\beta_n} \epsilon_{in}$, either by demand ϵ_{in} or by market access $1/\tau_n$. Second, the impact of tariffs on profits may be heterogeneous across countries. In particular, home country tariffs

² $\tau_n \equiv 1 + T_n$, where T_n is the ad-valorem tariff. In this part of the paper, we abstract from the possibility of bilateral tariffs due to preferential trade agreements. Our empirical strategy will, however, take this into account as well.

³Also, firms tend to enter markets based on their efficiency (Dekle et al., 2008).

are expected to have a negative impact on profits while tariffs in export markets are expected to have a positive impact, $\beta_h \neq \beta_n$, where h is the home country of firm i .

In a wide class of models, effective market size will affect a firm's innovative activity and therefore its stock of knowledge (e.g., Acemoglu and Linn, 2004). Appendix Section D also outlines a simple model where this is indeed the case. A firm's knowledge stock is therefore governed by the following equation,

$$K_i = \kappa_i \sum_{n \in \Omega_i} \tau_n^{\beta_n} \epsilon_{in},$$

where κ_i summarizes all other factors that also determine the knowledge stock. We discuss measurement of K_i below.

Now consider a change in the economic environment, κ_i , τ_n and ϵ_{in} , from $t = 0$ to $t = 1$. Using Jones' hat algebra popularized recently by Dekle et al. (2008), we get

$$\hat{K}_i = \hat{\kappa}_i \sum_{n \in \Omega_i} \omega_{in} \hat{\tau}_n^{\beta_n} \hat{\epsilon}_{in}, \quad (1)$$

where $\omega_{in} \equiv \pi_{in0}/\pi_{i0}$ and the hat notation is the change from $t = 0$ to $t = 1$, $\hat{x} \equiv x_1/x_0$. The weights ω_{in} are simply the share of profits generated in country n relative to total profits in the pre-period ($t = 0$). Appendix Section E shows that equation (1) can be approximated by

$$\Delta \ln K_i = \Delta \ln \kappa_i + \sum_{n \in \Omega_i} \beta_n \omega_{in} \Delta T_n + \varepsilon_i, \quad (2)$$

where $T_n \equiv \tau_n - 1$ is the ad-valorem tariff and $\varepsilon_i \equiv \sum_{n \in \Omega_i} \omega_{in} \Delta \ln \epsilon_{in}$. We choose this approximation because it is empirically much more convenient to work with.

2.2 Two Cases

We consider two cases; first a symmetric model, \mathbb{S} , with $\beta_n = \beta_m \equiv \beta$, so that the impact of tariffs in home and export markets is identical. Second, an asymmetric model, \mathbb{A} , where the impact is allowed to differ across home and export markets, $\beta^H \neq \beta^E$, where β^H is the home market impact and β^E is the export market impact. In the first case, we can rewrite equation (2) to

$$\Delta \ln K_i = \Delta \ln \kappa_i + \beta \Delta \bar{T}_i + \varepsilon_i, \quad (3)$$

where

$$\bar{T}_{it} \equiv \sum_{n \in \Omega_i} \omega_{in} T_{nt} \quad (4)$$

is the weighted average of tariffs across all of firm i 's markets. Hence, the knowledge stock is growing when weighted average tariffs decline (assuming $\beta < 0$) or when weighted average market size ε_i rises. As described in the Section 3 below, tariffs will be measured at the 3-digit industry-level, so that \bar{T}_i will vary both because firms are exposed to different markets and because firms belong to different industries.

The second case becomes

$$\Delta \ln K_i = \Delta \ln \kappa_i + \beta^H \omega_i^H \Delta T_i^H + \beta^E (1 - \omega_i^H) \Delta \bar{T}_i^E + \varepsilon_i, \quad (5)$$

where ω_i^H is the home market weight, T_i^H is the home market tariff and $\bar{T}_i^E \equiv [1 / (1 - \omega_i^H)] \sum_{n \in \Omega_i \setminus h} \omega_{in} T_n$ is the weighted average tariff in export markets. The asymmetric model, \mathbb{A} , allows us to separate the import competition effect from the market access effect of trade policy. Specifically, β^H will be identified by firms with a high degree of home bias, while β^E is identified by firms primarily exposed to foreign markets, and as such mostly affected by the export tariff \bar{T}_i^E .

Mechanisms. Theory and empirical evidence suggests that better market access gives rise to more innovation, i.e. that β^E is negative (e.g. Bustos, 2011). However, the sign of β^H is theoretically ambiguous; e.g. in Aghion et al (2005) competition has both a positive and negative impact on innovation. More competition may foster innovation because the pre-innovation rents decline (the escape competition effect), while competition may also hurt innovation because the expected profits from innovation decline (the Schumpeterian effect).

2.3 Measurement and Baseline Periods

Weights. The symmetric and asymmetric models shown above require data on the country weights ω_{in} . According to our model, these weights should reflect the relative importance of a country n in the firm's total profits. Profits and sales are unobserved in our data, but we do observe in which markets a firm is patenting in. We follow Aghion et al. (2014) and calculate the share of patents issued in country n relative to all the patents issued by the firm during the pre-period (to be defined below). Specifically, we define

$$\omega_{in} \equiv \frac{x_{in}}{\sum_k x_{ik}}, \quad (6)$$

where x_{in} is the number of patents issued in market n during the pre-period. Seeking intellectual property rights in a country is typically motivated by (future) profits in that market. There is strong empirical support that patent weights are highly correlated with sales weights (Aghion et al., 2014). We provide additional empirical support for this in Appendix Section F.

Baseline periods. The years 1965 to 1985 are defined as our pre-period. We use 1965 as the starting year because the number of patents in PATSTAT is limited in earlier years. 1985 is chosen as the final year because the Uruguay round negotiations started in 1986; hence the weights are not themselves affected by trade liberalization of the 1990s. The years 1992 to 2000 are defined as our baseline sample period, i.e. $\Delta \bar{T}_i = \bar{T}_{i2000} - \bar{T}_{i1992}$ and $\Delta \ln K_i = \ln K_{i2000} - \ln K_{i1992}$. This is motivated by the fact that tariff reductions agreed upon during the Uruguay Round were gradually phased in from 1995 until 2000. Starting our sample in 1992 ensures that we capture the full impact of tariff reductions. Our data (shown below) also confirms that the 1990s was unique; the overall reduction in tariffs was much greater during latter half of the 1990s compared to both earlier and later periods. Finally, we choose to work with long differences 1992-2000 in our baseline specification because we want to allow for long time lags in the innovation response to trade liberalization. Long differences also eliminate serial correlation in the errors, since the averaging over periods ignores time-series information (Bertrand et al., 2004).

Outcome variable. In our model, the outcome variable $\Delta \ln K_i$ is the change in the log knowledge stock. Our empirical counterpart is the cumulative patent count of a firm until year t ,

$$K_{it} \equiv \sum_{s=1965}^t p_{is},$$

where p_{is} is the number of patents filed by firm i in year s . The outcome variable $\Delta \ln K_{it} = \ln K_{i2000} - \ln K_{i1992}$ has a number of desirable properties. First, patenting is known to be highly correlated with R&D and innovation (Griliches, 1990). Second, focusing on the change in the stock over a long time period smooths out lumpiness and zeros in the p_{it} variable. Indeed, in a given year the median p_{it} is zero while the maximum p_{it} is very large, suggesting that linear models are not adequate to model the data generating process at the annual level.

2.4 Empirical Specification

Estimating model S and A is challenging for a number of reasons. The first econometric concern is that the weighted average tariff reduction $\Delta \bar{T}_i$ may be correlated with unobservable firm characteristics, $\Delta \ln \kappa_i$. For example, firms exposed to high-tariff reduction countries

may innovate more even in the absence of trade liberalization. We solve this by including home country-industry pair fixed effects in the regressions as well as controlling for pre-period firm characteristics.⁴ Intuitively, we compare firms within the same industry, headquartered in the same country, and with similar observed characteristics during the pre-period, but that differ in terms of their exposure to international markets, and ask whether firms exposed to high tariff-cut countries innovate more than firms exposed to low tariff-cut countries.

An alternative way of solving this problem is by differencing out idiosyncratic firm trends. Specifically, we split the sample into our main period, 1992-2000 ($t = 1$) and add a second period, 2000-2004 ($t = 2$), when the decline in tariffs was much smaller (see Figure 1 below), and estimate model \mathbb{S} by

$$\Delta \ln K_{i2} - \Delta \ln K_{i1} = \beta (\Delta \bar{T}_{i2} - \Delta \bar{T}_{i1}) + \varepsilon_i. \quad (7)$$

Idiosyncratic growth trends in innovation, $\Delta \ln \kappa_i$, that may be correlated with $\Delta \bar{T}_i$, is then differenced out. We use a similar specification for the asymmetric model \mathbb{A} . This is reminiscent of a triple differences model, as we compare the growth in the change in tariffs (two differences) across firms (third difference).

A second econometric concern is that the error term $\varepsilon_i \equiv \sum_{n \in \Omega_i} \omega_{in} \Delta \ln \varepsilon_{in}$, which is a weighted average of all other country-specific factors that determine innovation, may be correlated with trade liberalization. A case in point is the TRIPS agreement that strengthened intellectual property rights (IPR) among WTO members in the aftermath of the Uruguay round. A positive correlation between tariff reductions and IPR strengthening could therefore produce biased estimates. We solve this by using the fact that we observe aggregate patenting by industry and country, and this measure is itself determined by the unobserved shocks ε_i . Specifically, we calculate the aggregate knowledge stock by industry j and headquarters country h , $\mathcal{K}_{hjt} = \sum_{i \in \Gamma_{hj}} K_{it}$, where Γ_{hj} is the set of firms in industry j headquartered in h , and construct the weighted average

$$\tilde{\varepsilon}_i \equiv \sum_{n \in \Omega_i} \omega_{in} \Delta \ln \mathcal{K}_{nj},$$

where $\Delta \ln \mathcal{K}_{nj} = \ln \mathcal{K}_{nj2000} - \ln \mathcal{K}_{nj1992}$. While headquarters-industry pair fixed effects control for innovation trends in firm i 's *home* market, $\tilde{\varepsilon}_i$ controls for innovation trends in firm i 's *destination* markets. For example, if a US headquartered firm primarily exposed to the Indian market is innovating more because the Indian market is growing quickly (high

⁴Industries are defined at the NACE 3 digit level. Pre-sample covariates are home weights ω_i^H , the number of countries the firm is patenting in during the pre-period, $n_{i,Pre}$, and the knowledge stock of the firm in 1985, $K_{i,Pre}$.

$\Delta \ln \epsilon_{India}$), then including $\tilde{\epsilon}_i$ will control for this effect.

3 Data

3.1 Patents

Overview. We use data from the European Patent Office (EPO) Worldwide Patent Statistical Database (henceforth PATSTAT), the April 2015 version. PATSTAT offers bibliographic data, family links and citations of 90 million applications of nearly 100 countries. It contains the population of all patents globally since mid 1960s. The patent documents as provided by PATSTAT are a rich source of information. We observe the name of the applicant, date of filing and publication and whether or not the patent was granted. We know the geography of the patent in the sense that we have information on both source and destination country. Source country is the residence country of the applicant. Destination is the country of the patent authority (e.g. USPTO, EPO, JPO, etc).

Firm characteristics. PATSTAT allows us to construct an international firm-level dataset and to follow the patenting activity of a firm through time. The number of patents filed p_{it} for firm i in year t is our basic measure of the innovative activity of a firm. In our analysis, a patent corresponds to a *unique* invention, i.e. filing the same patent in multiple locations does not inflate the patent count p_{it} . Specifically, PATSTAT organizes patents into “patent families” that identifies identical inventions filed in multiple countries.⁵ We date patents by application filing year.⁶ An additional advantage of PATSTAT is that names of applicants are harmonized over the entire sample period, alleviating the concern that slight differences in the spelling of firm names generate multiple firm IDs.⁷ Unfortunately, information about firms in PATSTAT is restricted to what can be retrieved from the patent applications. Our basic firm characteristics are: Industry affiliation (NACE 3-digit), home country of the firm, as well as in which countries the firm is patenting. Appendix Section A details the precise construction of all variables.

Citations. We measure citations received by the patents filed by a firm i in year t . We do so by counting the number of citations per patent over the three years after the patent has been filed and let the count be denoted by c_p for patent p . Using a three-year window ensures that older patents do not mechanically get more citations than younger patents.

⁵We use DOCDB patent family.

⁶This is a common approach in the empirical literature because the application filing date is more closely timed with the R&D process than the patent publication and grant date. Patent applications are usually published 18 month after the first application.

⁷An applicant can be a firm or individual, but we will use the terminology firm when referring to an applicant.

3.2 Tariffs

The main source of tariff data is the UNCTAD Trade Analysis and Information System (TRAINS), which contains tariffs at the most disaggregated level of the Harmonized System (HS) for more than 150 countries. From this database we extract the average industry-level tariff (NACE 3-digit) T_{njt} for industry j , country n , for year t over the period 1992 to 2009.⁸ Appendix Section B describes the procedure to calculate industry-level tariffs. Appendix Section C provides details about the historical background for tariff reductions during the 1990s.

In addition to UNCTAD TRAINS data, we use information on regional trade agreements (RTAs) between pairs of countries. This allows us to take into account the fact that some countries are part of trade agreements, and as such cannot be treated as having the same level of protection as countries where such agreements are not in force. The information on RTAs for around 200 countries from 1948 to 2006 comes from the CEPII gravity dataset.

3.3 Final Sample

Our final sample consists of the following firms. First, firms must be observed in the pre-period (1965-1985) in order to be assigned weights ω_{in} . Second, in our baseline specification firms must issue at least one patent in 2004 or later. This is done to ensure that firms actually exist during the 1990s. For example, if $p_{it} = 0$ from 1995 and onwards, the econometrician cannot tell whether this is due to firm exit or firm survival but no patenting. Hence, our final sample is a strongly balanced dataset, with no firm entry or exit. In Section 4.3 we perform robustness tests allowing for firm exit. Third, in some cases firms have missing address and home country, or they may issue patents in countries with missing tariff data for their industry.⁹ These firms are dropped from the analysis. In sum, these restrictions reduce the sample to roughly 72,000 firms from 60 countries, filing 1.5 million patents over the 1992-2000 period. As a comparison, the raw PATSTAT data contains about 12.2 million patents over the 1992-2000 period.


3.4 Descriptives

Weighted average trade barriers. To illustrate our identification strategy, we take a closer look at the weighted average trade barriers, \bar{T}_{it} . Figure 1 shows the mean \bar{T}_{it} for firms headquartered in the U.S., Germany, Japan and the UK. There is a strong decline during

⁸Ad-valorem applied MFN rates.

⁹We drop all firms that have non-zero weights for one or more countries with missing tariff data, i.e. if T_{jnt} is missing when calculating \bar{T}_{it} from equation (4).

Figure 1: Average Weighted Import Tariffs, \bar{T}_{it} .




Note: The figure shows the annual average \bar{T}_{it} across firms according to headquarters country. The population of firms is described in Section 3.3.

the 1990s; the average firm experienced a decline of around 3 percentage points during the 1990s. Also, the decline almost stops in the year 2000, consistent with the fact that Uruguay Round concessions were phased in until that year. The averages mask a considerable amount of heterogeneity; Figure 2 shows that the whole distribution of \bar{T}_{it} 's across firms shifts markedly to the left from 1992 to 2000.

Patents and citations. Figure 3 shows average patenting p_{it} as well as average citations per patent for our final sample of firms. It is interesting to note that average patenting is increasing during the 1990s, while average citations are fairly constant over the period. Of course, these aggregate trends may not only reflect innovation, but could also reflect changes e.g. the patent systems worldwide.


Figure 4 shows the distribution of firms across home countries and industries (NACE 2-digit) in our final sample. We note the dominance of Japan and the US and by the industries machinery and equipment (28), computers, electronic and optical products (26), other manufacturing (32).

Figure 2: Density of Weighted import Tariffs, \bar{T}_{it} , in 1992 and 2000.




Note: The figure shows the histogram of \bar{T}_{it} , the weighted average import tariff in firm i 's markets, in 1992 and 2000. For expositional purposes the histogram is truncated at $\bar{T}_{it} = 20$. The population of firms is described in Section 3.3.

Figure 3: Patenting and Citations. 1980-2009.



Note: The figure shows the average number of patents per firm per year and the average number of citations per patent 3 years after the patent application date. The population of firms is described in Section 3.3.

Figure 4: Share of Firms by Country and Industry



Note: The figure shows the share of firms in the final dataset by home country and NACE 2-digit industry. Only the top 10 countries/industries are shown. The population of firms is described in Section 3.3.


4 Results

4.1 Innovation

We proceed by estimating models \mathbb{S} and \mathbb{A} from equations (3) and (5). As described in Section (2.4), all specifications include home country-industry (NACE 3-digit) pair fixed effects, which will control for aggregate (country and industry) trends in patenting. Columns (1) to (3) in Table 1 show results for model \mathbb{S} with various control variables included. Column (1) has only fixed effects, column (2) adds pre-sample firm characteristics (the home weight ω_{iH} and the number of countries the firm is patenting in during the pre-period $n_{i,Pre}$, as well as log knowledge stock in 1985, $\ln K_{i,Pre}$), while column (3) also controls for aggregate destination trends $\tilde{\varepsilon}_i$, as explained in Section 2.4. The results are highly significant and fairly constant across specifications, with an estimated coefficient in the range of -0.8 to -1.0. The final specification where we control for idiosyncratic firm trends (equation 7) in column (4) also produces a negative and significant result, although the economic magnitude is slightly smaller.

We also show the results graphically using a binned scatter plot in Figure 5. The figure plots changes in firm-specific tariffs $\Delta \bar{T}_i$ versus changes in the log stock of knowledge, $\Delta \ln K_i$, over the 1992-2000 period. The bins are created by dividing the sample into 20 equal-sized groups and taking the averages of $\Delta \bar{T}_i$ and $\Delta \ln K_i$ within each group. Both variables are first demeaned by home country-industry averages, which is equivalent to controlling for home country-industry fixed effects. It is interesting to note that the binned scatter plot, which is a non-parametric representation of the conditional expectation function, is close to

Figure 5: The Effect of Trade Policy on Innovation.



Note: The figure plots changes in firm-specific tariffs $\Delta \bar{T}_i$ versus changes in the log stock of knowledge, $\Delta \ln K_i$, over the 1992-2000 period. To construct the binned scatter plot, we first demean both the y- and x-axis variable by home country-industry fixed effects. We then group the observations into twenty equal-sized (5 percentile-point) bins based on the x-axis variable and scatter the means of the y- and x-axis variables within each bin. The solid line shows the best linear fit estimated on the underlying micro data estimated using OLS.

linear.

A semi-log elasticity of -0.9 implies that a one percentage point reduction in tariffs causes a 0.9 percent increase in the knowledge stock. Our data shows that over the period 1992 to 2000 the mean knowledge stock among firms globally went up by 45 percent (mean of $\Delta \ln K_i$), while the mean reduction in the firm specific tariff measure \bar{T}_{it} was almost three percentage points (mean of $\Delta \bar{T}_i$). Hence, our results suggest that roughly 7 (3/45) percent of the observed increase in the knowledge stock can be explained by trade policy.

The symmetric model \mathbb{S} masks the fact that liberalization in the home market can have a different effect than liberalization in export markets. Table (2) shows the results when we unpack home tariffs T_i^H and export market tariffs T_i^E . Columns (1) to (2) show the baseline model with and without the destination market control $\tilde{\varepsilon}_i$, while column (3) controls for idiosyncratic firm trends. The coefficient for $\omega_i^H \times \Delta T_i^H$ captures the differential impact of home tariffs depending on the firm's home bias. The coefficients are negative and significant, indicating that firms with high exposure to the home market innovate more, relative to firms with less exposure to the home market, when home tariffs decline. Of course, the overall

impact of home tariffs is not identified, as this effect is subsumed by the home country fixed effect. The coefficient for $(1 - \omega_i^H) \times \Delta \bar{T}_i^E$ captures the differential impact of export market tariffs depending on the firm’s export market bias. Again, we find a negative and significant number, suggesting that firms highly exposed to export markets innovate relatively more when export market tariffs decline. Moreover, the coefficient for $\Delta \bar{T}_i^E$, which captures the effect of export market tariffs for firms with close to zero export market exposure, is not significantly different than zero, consistent with the prediction in equation (5). As above, the point estimates become slightly smaller when controlling for idiosyncratic firm trends in column (3).

As discussed above, theory suggests that the market size effect of trade policy liberalization promotes innovation, while the import competition effect may encourage or discourage innovation, depending on the strength of escape-competition and Schumpeterian forces. Our results suggest that both the market size and import competition effect increase innovation, so that the net effect is unambiguously positive.

4.2 Citations

One may argue that patents are an imprecise measure of knowledge and innovation. For example, it may be that firms are taking out more patents in response to import competition. If this is the case, one should expect that firms are taking out patents on their marginal innovations, so that the average quality of their patent stock is decreasing. One should then see a drop in the average citation count for firms exposed to trade liberalization, since highly valuable inventions are more extensively cited than low value patents (Harhoff et al., 1999). We denote the number of citations three years after a patent p was filed c_p and calculate the cumulative citation count for all of firm i ’s patents,

$$C_{it} = \sum_{s=1965}^t \sum_{p \in \Omega_{is}} c_p,$$

where Ω_{is} is the set of firm i ’s patents filed in year s . The average citation count for all of i ’s patents is then $\bar{C}_{it} = C_{it}/K_{it}$. We proceed by using $\Delta \ln \bar{C}_i = \ln \bar{C}_{i2000} - \ln \bar{C}_{i1992}$ as the dependent variable. The results are reported in Table 3. The results suggest that trade liberalization did not affect the quality of patents, i.e. there is no evidence of a “lawyer effect”. If anything, the firm fixed effects results in column (3) indicate that trade policy may have increased the quality of patents.

4.3 Robustness

Falsification test. A potential concern is that firms being exposed to high tariff-cut countries always have always higher patent growth compared to other firms. To address this concern, we perform a placebo test: we regress knowledge growth during the 1980s, $\ln K_{i1988} - \ln K_{i1980}$, on trade policy changes during the 1990s, $\Delta \bar{T}_{i2000} - \Delta \bar{T}_{i1992}$.¹⁰ The results are shown in the first column of Table 4: The coefficient of interest becomes noisy and close to zero, suggesting that there are no differential pre-trends in patenting.

Country-level tariff data. Next, we take some additional steps to refine our analysis. First, industry-level tariff data may not always be the relevant tariff facing the firm, because it may also be exporting or importing products associated with other 3-digit NACE industries. We therefore test the sensitivity of our results using the simple average country tariff T_{nt} instead of T_{jnt} . The results, displayed in the second column of Table 4, confirm our main finding that a reduction of a firm’s trade barrier increases innovative activity. The estimated effect is similar in magnitude to our main specification and economically significant.

Regional trade agreements. Our main measure of tariffs is the applied MFN ad-valorem rate. This masks the fact that many firms get preferential market access through regional trade agreements (RTAs). Recognizing this, we calculate a firm-level measure of how exposed a firm is to RTA’s. Specifically, we construct $R\bar{T}A_{it}$ in a similar way as \bar{T}_{it} above,

$$R\bar{T}A_{it} \equiv \sum_{n \in \Omega_i} \omega_{in} RTA_{hn},$$

where $RTA_{hn} = 1$ if the home country h and country n have an RTA and zero otherwise.¹¹ We then add $\Delta R\bar{T}A_i = R\bar{T}A_{i2000} - R\bar{T}A_{i1992}$ to the model. The results in column (3) show that the RTA variable is insignificant, while our main variable $\Delta \bar{T}_i$ continues to be highly significant and negative.

Triadic patents. Third, we restrict our sample to triadic patents. These are patents filed at the three main patent offices, namely the European Patent Office (EPO), the Japanese Patent Office (JPO) and the United States Patents and Trademark Office (USPTO).¹² Triadic patents are commonly used in the literature to retain only highly valuable inventions and to work with a more uniform and comparable set patents. This is shown in the last column of Table 4. In this case, the sample size is significantly reduced and the standard

¹⁰The weighted average \bar{T}_{it} is now calculated using weights ω_{in} based on a firm’s patent portfolio until 1980 (not 1985 as in the baseline). This is done in order to ensure that the weights ω_{in} are not themselves determined by the dependent variable $\ln K_{i1988} - \ln K_{i1980}$.

¹¹As a matter of convention, we set $RTA_{hh} = 1$.

¹²See Dervis and Khan (2004) and Martinez (2010) for more detailed definition and explanation of how triadic patent families are constructed.

errors become very large.

Firm exit. A final consideration is that our sample is limited to firms that file at least one patent at the end of the observation period (between 2004 and 2014). Thus, by construction, all our firms survive until the end of the period of analysis. However, firms may exit the market as a result of increased competition following trade liberalization. To address this concern, we estimate the model including all firms that have at least one patent during the 1965-1992 period. The estimated effect in column (5) on this larger sample of firms is still highly significant but the magnitude is somewhat lower. This may suggest that trade policy also induced firm exit, although one would need data on actual exit to corroborate this.

Destination country trends. The variable $\tilde{\varepsilon}_i$ was included in the regressions to capture all patenting trends in destination countries. An alternative empirical strategy is to include destination country fixed effects in the regressions. Specifically, we rewrite model S in equation (3) to

$$\Delta \ln K_i = \Delta \ln \kappa_i + \beta \Delta \bar{T}_i + \sum_{n \in \Omega_i} \gamma_n + \varepsilon_i,$$

where γ_n is a fixed effect for destination n , and we sum over all countries the firm has non-zero weights during the pre-period (the set Ω_i). As an example, if all firms exposed to the Indian market (but not necessarily headquartered in India) have high $\Delta \ln K_i$, then this will be controlled for by γ . Identification of β then only comes from within-country, across-industry variation in tariffs, i.e. that among firms exposed to the Indian market, some firms experience greater tariff reductions because they belong to an industry getting large tariff cuts in India. Destination country trends will therefore control for the possibility that firms exposed to India may patent more because of unobserved factors specific to India (e.g., growth in market size or strengthening of IPR). The estimated coefficient in column (6) shows that β is still highly significant, although the economic magnitude is lower than in the baseline specification.

5 Conclusions

TBW

Table 1: Trade Policy and Knowledge Creation. Model S.

Dep. variable: $\Delta \ln K_i$	(1)	(2)	(3)	(4)
$\Delta \bar{T}_i$	-.96 ^a (.12)	-.88 ^a (.12)	-.82 ^a (.12)	-.66 ^a (.12)
Home country-industry FE	Yes	Yes	Yes	Yes
Pre-sample firm characteristics	No	Yes	Yes	Yes
$\tilde{\varepsilon}$	No	No	Yes	Yes
Firm trends	No	No	No	Yes
Number of firms	72,188	72,188	71,748	71,748

Standard errors clustered by home country-industry in parentheses.
^a p < 0.01, ^b p < 0.05, ^c p < 0.1.

Table 2: Trade Policy and Knowledge Creation. Model A.

Dep. variable: $\Delta \ln K_i$	(1)	(2)	(3)
$\omega_i^H \times \Delta T_i^H$	-.77 ^a (.13)	-.75 ^a (.13)	-.67 ^a (.22)
$(1 - \omega_i^H) \times \Delta \bar{T}_i^E$	-.83 ^a (.27)	-.75 ^a (.28)	-.56 ^b (.22)
$\Delta \bar{T}_i^E$	-.09 (.20)	-.11 (.20)	-.08 (.15)
Home country-industry FE	Yes	Yes	Yes
Pre-sample firm characteristics	Yes	Yes	Yes
$\tilde{\varepsilon}$	No	Yes	Yes
Firm trends	No	No	Yes
Number of firms	72,188	71,784	71,784

Standard errors clustered by home country-industry in parentheses.
^a p < 0.01, ^b p < 0.05, ^c p < 0.1.

Table 3: Trade Policy and Citations.

Dep. variable: $\Delta \ln \bar{C}_i$	(1)	(2)	(3)
$\Delta \bar{T}_i$	-.30 (.30)	-.30 (.30)	-.66 ^c (.35)
Home country-industry FE	Yes	Yes	Yes
Pre-sample firm characteristics	Yes	Yes	Yes
$\tilde{\varepsilon}$	No	Yes	Yes
Firm trends	No	No	Yes
Number of firms	35,773	35,556	35,556

Standard errors clustered by home country-industry in parentheses.

^a p < 0.01, ^b p < 0.05, ^c p < 0.1.

Table 4: Robustness.

Dep. variable: $\Delta \ln K_{it}$	(1) Placebo	(2) Aggregate $\Delta \bar{T}_i$	(3) RTA's	(4) Triadic	(5) Large sample	(6) Dest. trends
$\Delta \bar{T}_i$	-.18 (.24)	-.96 ^a (.15)	-.78 ^a (.13)	1.64 (3.59)	-0.42 ^a (.11)	-0.48 ^a (.14)
ΔRTA_i			.01 (.04)			
Home country-industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Pre-sample firm characteristics	Yes	Yes	Yes	Yes	Yes	Yes
$\tilde{\varepsilon}$	Yes	Yes	Yes	Yes	Yes	No
Firm trends	No	No	No	No	No	No
Destination country trends	No	No	No	No	No	Yes
Number of firms	40,574	69,837	71,121	1,522	132,034	72,188

Standard errors clustered by home country-industry in parentheses.

^a p < 0.01, ^b p < 0.05, ^c p < 0.1.

References

- Acemoglu, D. and J. Linn (2004). Market size in innovation: Theory and evidence from the pharmaceutical industry. *Quarterly Journal of Economics* 119(3), 1049–1090.
- Aghion, P., N. Bloom, R. Blundell, R. Griffith, and P. Howitt (2005). Competition and innovation: An inverted u relationship. *Quarterly Journal of Economics* 102(2), 701–728.
- Aghion, P., A. Dechezleprêtre, D. Hemous, R. Martin, and J. van Reenen (2014). Carbon taxes, path dependency and directed technical change: Evidence from the auto industry. *Journal of Political Economy* forthcoming.
- Altomonte, C. and T. Aquilante (2012). The eu-efige/bruegel-unicredit dataset. Technical report, Bruegel working paper.
- Amiti, M. and J. Konings (2007). Trade liberalization, intermediate inputs, and productivity: Evidence from Indonesia. *American Economic Review* 97(5), 1611–1638.
- Bertrand, M., E. Duflo, and S. Mullainathan (2004). How much should we trust differences-in-differences estimates? *Quarterly Journal of Economics* 119(1), 249–275.
- Bloom, N., M. Draca, and J. Van Reenen (2016). Trade induced technical change? the impact of chinese imports on innovation, it and productivity. *Review of Economic Studies* 83(1), 87–117.
- Boler, E. A., A. Moxnes, and K. H. Ulltveit-Moe (2015). R&d, international sourcing and the joint impact on firm performance. *American Economic Review* 105(12), 366–3703.
- Bustos, P. (2011). Trade liberalization, exports and technology upgrading: Evidence on the impact of mercosur on argentinean firms. *American Economic Review* 101(1), 304–340.
- Calel, R. and A. Dechezleprêtre (2014). Environmental policy and directed technological change: Evidence from the european carbon market. *Review of Economics and statistics* forthcoming.
- Dekle, R., J. Eaton, and S. Kortum (2008). Global rebalancing with gravity: Measuring the burden of adjustment. Working Paper 13846, National Bureau of Economic Research.
- Dernis, H. and M. Khan (2004). Triadic patent families methodology. OECD Science, Technology and Industry Working Papers 2004/2, OECD, Directorate for Science, Technology and Industry.

- Estevadeordal, A. and A. M. Taylor (2013). Is the washington consensus dead? growth, openness, and the great liberalization, 1970s–2000s. *Review of Economics and Statistics* 95(5), 1669–1690.
- Goldberg, P. K., A. K. Khandelwal, N. Pavcnik, and P. Topalova (2010). Imported intermediate inputs and domestic product growth: Evidence from India. *The Quarterly Journal of Economics* 125(4), 1727–1767.
- Gopinath, G. and B. Neiman (2013). Trade adjustment and productivity in large crises. *American Economic Review* (3).
- Griliches, Z. (1990). Patent statistics as economic indicators: a survey. *Journal of Economic Literature* 28(4), 1661–1707.
- Halpern, L., M. Koren, and A. Szeidl (2015). Imported inputs and productivity. *American Economic Review* 105(12), pp. 366–3703.
- Harhoff, D., F. Narin, F. M. Scherer, and K. Vopel (1999). Citation frequency and the value of patented inventions. *The Review of Economics and Statistics* 81(3), 511–515.
- Khandelwal, A. K. and P. Topalova (2011). Trade liberalization and firm productivity: The case of India. *Review of Economics and Statistics* 93(3), 995–1009.
- Lileeva, A. and D. Trefler (2010). Improved access to foreign markets raises plant-level productivity ... for some plants. *The Quarterly Journal of Economics* 125(3), 1051–1099.
- Magerman, T., B. Van Looy, and X. Song (2006). Data production methods for harmonized patent statistics: Patentee name harmonization. Eurostat Working Paper and Studies, Luxembourg.
- Martinez, C. (2010). Insight into different types of patent families. OECD Science, Technology and Industry Working Papers 2010/2, OECD, Directorate for Science, Technology and Industry.
- Moxnes, A. (2010). Are sunk costs in exporting country specific? *Canadian Journal of Economics/Revue canadienne d'Économie* 43(2), 467–493.
- OECD (2009). *Patent Statistics Manual*. OECD Publishing, Paris.
- Peeters, B., X. Song, J. Callaert, J. Grouwels, and B. Van Looy (2009). Harmonizing harmonized patentee names: an exploratory assessment of top patentees. EUROSTAT working paper and Studies, Luxembourg.

Teshima, K. (2009). Import competition and innovation at the plant level: Evidence from Mexico.

Appendix

A PATSTAT

Identify unique firms/patent holders. As described in the main text, for each patent application in PATSTAT we know the exact name of the patent applicant(s). However, looking at a long period of time and at patents filed worldwide introduces some complications. A well understood problem is that patentee names that appear in patent documents may vary both within and across patent systems. Inconsistencies might be due to spelling mistakes, typographical errors, name variants, etc. In order to identify unique patent holders, we use the ECOOM-EUROSTAT-EPO PATSTAT Person Augmented Table (EEE-PPAT). This table was developed by EUROSTAT in collaboration with ECOOM (K.U.Leuven) and Sogeti, and provides harmonized patent applicants' names obtained through an automated algorithm.¹³ These harmonized names have been included in PATSTAT TLS906_PERSON table since October 2011. We use the variable "HRM_L2_ID" from this table. This allows us to reduce the number of unique patent holders by about one third, from from 3085499 to 2130328.

Patent families. To construct the innovation variables in our analysis we use patent counts. In principle, an applicant may decide to patent an invention in one or more countries, depending on where he seeks IP protection, and he can do so contemporaneously or at subsequent times after the first application. Therefore, simply counting the number of patent filings for each patentee would result in double counting the number of unique inventions belonging to each firm. To avoid this problem, we look at patent families. A patent family identifies and groups all subsequent patent filings originating from the same initial (priority) application; hence it comprises all patents protecting the same invention.¹⁴ An example can be helpful to clarify the main idea/principle behind patent families. Suppose a German firm develops a new invention and patents it in Germany. Subsequently, it decides to seek protection for the same invention in US and in Japan and files a patent at the USPTO and at the JPO. These three applications clearly protect the same invention and thus belong to the same patent family. For the purpose of our analysis these three applications are counted as one. Notice also that a patent family is a generic term: different definitions of how to group applications can be applied, depending on the specific purpose. Throughout our analysis we use DOCBD patent families.¹⁵

¹³For more information on the method developed to arrive at harmonized patentee names see <https://www.ecoom.be/nl/eee-ppat> and Magerman et al. (2006) and Peeters et al. (2009).

¹⁴The OECD Patent Statistics Manual defines patent families as "the set of patents (or applications) filed in several countries which are related to each other by one or several common priority filings" (OECD, 2009, Ch.4, p.71).

¹⁵See also Dernis and Khan (2004) and Martinez (2010) for an overview of different types of patent families

Assigning patents to firms. We identify the list of patent applicants from PATSTAT table TLS207_PERS_APPLN. Applicants have “APPLN_SEQ_NR” greater than 0. The same table provides the correspondence between each applicant and the patents he owns. We use this built in link to assign patents to firms. Technically, patentees can be private business enterprises, universities/higher education institutions, governmental agencies, or individuals, but for simplicity we call them firms throughout the paper. At this point, one clarification is required. It is possible that several applicants co-own the same patent. In this case we proceed by assigning the patent to every co-owner of the patent application.

Identify home country of firms. In order to identify the home country of a firm we use PERSON_CTRY_CODE from TLS906_PERSON in PATSTAT. One difficulty is that the information on the applicant’s country is not always reported. For 974855 out of 2130328 unique firms in our sample we are not able to recover the home country because the information is missing. As a result, these firms have to be dropped. For the remaining sample, we consider the country associated with each applicant as its headquarter country. Notice that a firm may be associated with more than one country. We have 42574 of such cases. When this is the case, we attach to the firm the home country that has the highest frequency in the data.

Identify the industry affiliation of a firm. PATSTAT includes a table that assigns one more more industries j (NACE revision 2) to a patent application p . Industries are given weights w_{pj} that sum to one for a given application (table TLS229). We define the main industry of a firm as the industry that obtains the maximum weight across all of the firm’s applications, $\max \sum_p w_{pj}$ during the pre-period.

B Tariff Data

The main source of tariff data is the UNCTAD Trade Analysis and Information System (TRAINS), which contains tariffs at the most disaggregated level of the Harmonized System (HS) for more than 150 countries. From this database we extract the average ad-valorem industry-level tariff (NACE 3-digit) T_{njt} for industry j , country n , for year t over the period 1992 to 2009.

The details are as follows. First, we convert 6-digit HS codes to a the 6-digit HS Combined (HSC) nomenclature using a World Bank correspondence table.¹⁶ In some cases, a 6-digit tariff line is missing in year t , but non-missing in $t-1$ and $t+1$; in these cases we interpolate to get a non-missing observation in year t . We also extrapolate tariffs in those cases where tariffs

and how they are constructed.

¹⁶http://wits.worldbank.org/product_concordance.html

exist in 1995 but not in 1992-1994, or 1994 but not 1992-1993, or 1993 but not 1992. Tariff data for all EU member countries are also manually added to the database, as EU tariffs are not listed for individual EU countries in the raw data. Second, we strongly balance the raw data and drop all HSC-country combinations that are not available for all years 1992-2009. This is done to eliminate the possibility that average tariffs change simply due to sampling issues. Third, we aggregate the data to NACE revision 2 3-digit codes. To do so, we first aggregate to 4-digit ISIC revision 3.0 by using a correspondence table from the World Bank. This is then converted to 4-digit ISIC revision 3.1, then to 4-digit ISIC revision 4, which is again converted to NACE revision 2. The last three conversions use correspondences from the UN.¹⁷ In cases where e.g. several ISIC revision 3.1 codes are associated with a single NACE revision 2 code, we take the simple average across the ISIC codes. In some cases, a firm has a missing industry code or a 2-digit code instead of a 3-digit code. In those cases, we use the simple average tariff across all industries, or across 3-digit codes within a 2-digit industry, $T_{nt} = (1/N) \sum_j T_{njt}$, instead.

The final tariff dataset contains data for 96 countries, 128 3-digit industries and 12,174 country-industry combinations. Figure 6 shows average tariffs for high- and low income countries in our final tariff dataset.

C Trade Policy During the 1990s


Launched in Punta del Este, Uruguay, on 20 September 1986, the Uruguay Round of Multilateral Trade Negotiations was formally concluded in Marrakesh, Morocco, on April 15 1994, when 125 Governments and the European Communities, accounting for more than 90 percent of world trade, concluded a historical agreement to reform international trade. As stated in the Marrakesh declaration,¹⁸ the Uruguay Round achieved a global reduction by 40 percent of tariffs and wider market-opening agreements on goods, and the increased predictability and security represented by a major expansion in the scope of tariff commitments. In addition, participation in the Uruguay Round was considerably wider than in any previous multilateral trade negotiation and, in particular, developing countries played a notably active role in it. While only few developing countries took part in earlier GATT rounds, and trade barriers reduction was negligible,¹⁹ the Uruguay round achieved important tariff reductions in both developed and developing countries. Hence, after eight rounds of troubled negotiations, the most ambitious and far reaching multilateral trade negotiation

¹⁷<http://unstats.un.org/unsd/cr/registry/regot.asp?Lg=1>

¹⁸https://www.wto.org/english/docs_e/legal_e/marrakesh_decl_e.pdf

¹⁹Exceptions are represented by the East Asian NICs.

Figure 6: Average Tariffs



Note: The figure shows average tariffs for high- and low income countries according to the World Bank 1995 definition, using our final tariff dataset. Average tariffs are calculated as the simple average across countries. 3-digit NACE tariffs are aggregated to country level tariffs using simple averages.

ever started led to the biggest reform of the world’s trading system since the GATT was created. The Uruguay Round implied commitments to cut and bind tariffs on the imports of goods. The tariff reductions agreed on were explicit on both the timing and magnitude in cut. The deadlines for cut ended in 2000.

The tariff reductions

The major results of the Uruguay Round were the individual commitments of the contracting parties to cut and bind their custom duty rates on imports of goods. It is important to note that the phase-in of tariff reductions were programmed during the negotiations. This feature of the Marrakesh Agreement is important because it ensures that tariff reductions were pre-determined and therefore unlikely to be correlated with contemporaneous shocks, or to be driven by political pressure arising from the effects of trade liberalization.

For non-agricultural products the agreed tariff reductions were implemented in five equal installments.²⁰ The first cut was made on the date of entry into force of the WTO agreement,

²⁰Unless it is otherwise stated in a Member’s Schedule.

and the following four on 1 January of each subsequent year.²¹ Over the five years, this process led to a 40% tariff cut on average on industrial products in developed countries, from an average of 6.3% to an average of 3.8%.

In addition to tariff cuts, the number of “bound” tariffs²² increased significantly, from 78% to 99% in developed countries, from 21% to 73% in developing countries, and from 73% to 98% in transition economies.

D A Simple Model of Innovation

We motivate the empirical estimating equation with a simple model of innovation. Innovation is costly and the cost depends on the level of productivity z_i the firm wants to achieve. Specifically, the innovation cost is $c(z_i) = \psi z_i^k$, where ψ determines average innovation cost and k determines how quickly innovation costs rise with productivity. The firm then chooses the optimal z_i that maximizes $\pi_i - c_i$. The first order condition is

$$\frac{\partial \pi_i}{\partial z_i} - \psi k z_i^{k-1} = 0,$$

where

$$\frac{\partial \pi_i}{\partial z_i} = \alpha z_i^{\alpha-1} \sum_{n \in \Omega_i} \tau_n^{\beta_n} \epsilon_n,$$

which gives us

$$z_i^* = \kappa \left(\sum_{n \in \Omega_i} \tau_n^{\beta_n} \epsilon_n \right)^{1/(k-\alpha)},$$

where κ is a constant.²³ The second order condition is

$$\frac{\partial^2 \pi_i}{\partial z_i^2} - \psi k(k-1) z_i^{k-2} = \alpha(\alpha-1) z_i^{\alpha-2} \sum_{n \in \Omega_i} \tau_n^{\beta_n} \epsilon_n - \psi k(k-1) z_i^{k-2} < 0,$$

which holds whenever $k > \alpha$ evaluated at $z_i = z_i^*$. The optimal level of productivity, z_i^* , is therefore a function of the market access term $\sum_{n \in \Omega_i} \tau_n^{\beta_n} \epsilon_n$.

²¹see Marrakesh Protocol to the General Agreement on Tariffs and Trade 1994 for more information.

²²Bound tariffs are duty rates that are committed under WTO. Raising them above the bound rate is possible but hard: the process involves a negotiation with the most affected countries and it possibly requires a compensation for their loss of trade.

²³ $\kappa = \left(\frac{\alpha}{\psi k} \right)^{1/(k-\alpha)}$.

E Approximation of the Knowledge Production Function

The expression $\hat{K}_i = \hat{\kappa}_i \sum_{n \in \Omega_i} \omega_{in} \hat{\tau}_n^{\beta_n} \hat{\epsilon}_n$ can be approximated by equation (2) in the main text, $\Delta \ln K_i = \sum_{n \in \Omega_i} \beta_n \omega_{in} \Delta T_n + \sum_{n \in \Omega_i} \omega_{in} \Delta \ln \epsilon_n$.

Proof. The term

$$\begin{aligned} \sum_{n \in \Omega_i} \omega_{in} \hat{\tau}_n^{\beta_n} \hat{\epsilon}_n &= \sum_{n \in \Omega_i} \omega_{in} e^{\beta_n \Delta \ln \tau_n + \Delta \ln \epsilon_n} \\ &\approx \sum_{n \in \Omega_i} \omega_{in} (1 + \beta_n \Delta \ln \tau_n + \Delta \ln \epsilon_n) \\ &= 1 + \sum_{n \in \Omega_i} \omega_{in} (\beta_n \Delta \ln \tau_n + \Delta \ln \epsilon_n), \end{aligned}$$

where we used the fact that $\ln(1+x) \approx x \iff 1+x \approx e^x$ for x close to 0. Hence,

$$\begin{aligned} \Delta \ln K_i &= \ln \left[1 + \sum_{n \in \Omega_i} \omega_{in} (\beta_n \Delta \ln \tau_n + \Delta \ln \epsilon_n) \right] \\ &\approx \sum_{n \in \Omega_i} \omega_{in} (\beta_n \Delta \ln \tau_n + \Delta \ln \epsilon_n) \\ &= \sum_{n \in \Omega_i} \beta_n \omega_{in} \Delta T_n + \sum_{n \in \Omega_i} \omega_{in} \Delta \ln \epsilon_n, \end{aligned}$$

where we used $\Delta \ln \tau_n = \Delta \ln(1 + T_n) \approx \Delta T_n$ for T_n close to 0. □


F Patent and Sales Weights

This section provides empirical evidence that trade and patent flows are highly correlated. We present two approaches. The first method aggregates data at the country level to calculate the shares of patent flows into a country, and uses trade data to compare them to a country's import shares. We proceed as follows. Using the entire population of firms with non-missing headquarters country information in PATSTAT, we calculate the share of patents filed in country s that come from firms headquartered in country r , relative to all other foreign patents filed in country s ,

$$\Omega_{rst} = \frac{\text{Patents from } r \text{ to sat time } t}{\sum_{k \neq s} \text{Patents from } k \text{ to sat time } t}.$$

Similarly, by using trade data from CEPII, we calculate the import share Π_{rst} as the share of trade from r to s relative to s 's total imports. Figure 7 shows the trade and patent

Figure 7: Import Shares - Trade and Patents.




Note: The vertical axis shows the share of patents filed in U.S./Germany/Japan/Great Britain belonging to firms headquartered in source country r (log scales). The horizontal axis shows the share of total imports in U.S./Germany/Japan/Great Britain coming from source country r (log scales). The year is 2000. The population of firms is all firms in PATSTAT with non-missing headquarters country information.

import shares on the horizontal and vertical axis, respectively, on log scales, for four major economies, the U.S., Germany, Japan and Great Britain in year 2000. There is a high degree of overlap; typically the top three countries on the trade side are also the top three countries on the patent side. The correlation between Ω_{rs} and Π_{st} for all possible country-pairs is shown in Figure 8. There is a tight log linear relationship between bilateral patenting and trade, with a linear regression slope of 0.80 (s.e. 0.02). Finally, we show that patent flows adhere to a gravity model. Table 5 shows results when regressing the number of patents from r filed in s on distance and GDP in r and s (all in logs). Column (1) uses only the year 2000 cross-section, while column (2) uses all years from 1965 to 2006 and includes year and country-pair fixed effects. Just as for trade flows, bilateral patenting falls with distance and increases with the size of the home and destination country.

The second method uses survey data to calculate firm specific export shares to different country groups, and compares them to patent weights from PATSTAT. We use a sample of

Figure 8: Bilateral Trade and Patenting.



Note: The figure shows the number of patents and total trade from headquarters country r to destination country s in year 2000 (both in logs). The solid line is the local polynomial regression fit and the gray area represents the 95% confidence bands. The linear regression slope is 0.80 (s.e. 0.02). The population of firms is all firms in PATSTAT with non-missing headquarters country information.

Table 5: Patent flows.

Dep. variable: $\ln Patents_{rst}$	(1) Year 2000	(2) All years
Distance $_{rs}$	-.44 ^a (.03)	
GDP $_r$.68 ^a (.02)	.48 ^a (.05)
GDP $_s$.50 ^a (.02)	.27 ^a (.04)
Year FE	No	Yes
Source-destination FE	No	Yes
R ²	0.43	0.34
Number of observations	2,558	68,447

Robust standard errors in parentheses. ^a $p < 0.01$, ^b $p < 0.05$, ^c $p < 0.1$.

European firms that results from matching firm level data from Amadeus with data from EU-EFIGE/Bruegel-UniCredit dataset (henceforth EFIGE). The EFIGE database consists of a representative sample of about 15,000 surveyed manufacturing firms (above 10 employees) in seven countries (Germany, France, Italy, Spain, United Kingdom, Austria, Hungary). The survey was conducted in 2010, and covers the years from 2007 to 2009.²⁴ For our purpose, EFIGE contains information on firms' international activities. We use firms' self-reported export shares to different countries in 2008,²⁵ and, for each firm, we construct weights for market exposure based on the share of sales to eight groups of countries: EU 15 countries, other EU countries, other European countries not EU, China and India, other Asian countries, USA and Canada, Central and South America, and a residual category including all remaining countries.²⁶

Similarly, by using patent data from PATSTAT, we calculate weights for market exposure based on firms' patenting activity abroad. To do so, we match Amadeus-EFIGE firms with PATSTAT data. We use the patent application number of each patent owned by Amadeus firms to establish the link between Amadeus and PATSTAT.²⁷ In this way we are able to establish a correspondence between a firm in Amadeus and the same firm in PATSTAT, and to obtain relevant information on a firm's patents, in particular the application filing year and the application authority. We use patent applications for a period of ten years, from 1998 to 2008, and, using information on the application authority, we aggregate patents in the same way as export data.²⁸ If a firm does not have patents, then all its weights for all groups of countries are set to zero. Figure 9 shows a kernel-weighted local polynomial regression of patent shares on export shares for firms with at least one patent. There is a strong relationship between patent and trade weights. The corresponding linear regression slope is 0.89 (s.e. 0.008).

²⁴The EFIGE dataset is described in Altomonte and Aquilante (2012).


²⁵Specifically, we use answers to two questions. D4 asks: "Which percentage of your 2008 annual turnover did the export activities represent?" D13 asks: "If we assume that the total export activities equal to 100 which percentage goes to each of the following areas: 15 UE countries area, Other UE countries, Other European countries not UE (Switzerland, Norway, Russia, Turkey, Byelorussia, Ukraine, . . .), China and India, Other Asian countries (excluded China and India), USA and Canada, Central and South America, and Other areas.

²⁶The weight for EU 15 is computed by summing a firm's exports share to EU 15 area and the share of sales in its home market.

²⁷Specifically, from the variable `patentapplicationnumber` in Amadeus we are able to construct the `apl_n_r_epodoc` in PATSTAT, and to link each patent application in Amadeus to the same patent application in PATSTAT.

²⁸When the application authority is EPO, we assume that the patent was filed in at least one of the EU 15 countries, and include it in the EU 15 share. The motivation is that EPO filing is cost effective if the applicant wants to protect an invention in 4 or more countries, so there must be at least one application filed in one of the EU15 countries.

Figure 9: Market exposure weights - Trade and Patents.



Note: The figure shows market exposure weights based on sales and patenting activity. Sales data refer to 2008, patent weights are calculated over a ten year window, from 1998 to 2008. The solid line is the local polynomial regression fit and the gray area represents the 95% confidence bands. The linear regression slope is 0.89 (s.e. 0.008). The population of firms consists of surveyed firms in EFIGE, subsequently matched to Amadeus and, for firms holding patents, to PATSTAT. Only firms with at least one patent are included in the figure.

Table 6: Persistence in Patent Weights. Extensive Margin.


	(1) $t = 0$	(2) $t = 1$	(3) $t = 2$
$P [p_{int} p_{in0}]$ (forward)	1	0.44	0.44
	-	(.001)	(.001)
$P [p_{int} p_{in2}]$ (backwards)	0.37	0.39	1
	(.001)	(.001)	-
$P [p_{int}]$.037	.033	.035

Standard errors in parentheses. The first row ($P [p_{int} | p_{in0}]$) shows the share of firm-destinations with positive patenting in both $t = 0$ and $t = 1$ (column 2), or $t = 0$ and $t = 2$ (column 3), relative all firms-destinations with positive patenting in $t = 0$. The second row ($P [p_{int} | p_{in2}]$) shows the share of firms-destinations with positive patenting in both $t = 0$ and $t = 2$ (column 1), or $t = 1$ and $t = 2$ (column 2), relative all firms-destinations with positive patenting in $t = 2$. The final row shows the unconditional probability, i.e. the share of firm-destinations with patents in t relative to all possible firm-destinations. $t = 0$: 1965-1985, $t = 1$: 1985-1995 and $t = 2$: 1996-2005. The population of firms is described in Section 3.3.

G Persistence in Patent Weights

This section provides empirical evidence that patent weights ω_{in} are highly persistent over time. We calculate weights ω_{int} based on all patents filed during three non-overlapping time periods, $t = 0$: 1965-1985 (the weights used in the main text), $t = 1$: 1985-1995 and $t = 2$: 1996-2005. First, we calculate the likelihood of continuing to patent in a country conditional on patenting there in $t = 0$ (i.e., the extensive margin). We also calculate the likelihood of patenting in $t = 0$ and $t = 1$ conditional on patenting in the country in $t = 2$. We use the final sample of firms, which ensures that we know that all firms exist throughout the sample. Table 6 reports the results. Even after 20 years, the likelihood of remaining in a market remains high (44 percent). The same is true on the entry side; conditional on being in a market in $t = 2$, the likelihood of being in that market 20 years earlier is nearly 40 percent. These conditional probabilities are an order of magnitude higher than the unconditional probability of patenting in a market. The final row in the table shows that the unconditional probability is roughly 4 percent. Second, we calculate the correlation in weights conditional on being in that market in both t and $t + 1$ (i.e., the intensive margin). Figure 10 shows the expected weight in $t = 1$ and $t = 2$ conditional on a 1985 weight ω_{in0} . Even after 20 years there is a highly significant and positive correlation in the weights.

Figure 10: Persistence in Patent Weights. Intensive Margin.



Note: The figure shows the kernel-weighted local polynomial regression of weights ω_{int} in 1995 or 2005 (vertical axis) on weights in 1985 (horizontal axis). The two lines represent two separate regressions. Gray areas denote the 95 percent confidence bands. The sample includes all pairs $(\omega_{int}, \omega_{in,t+1})$ where both values are non-zero. The population of firms is described in Section 3.3.

H Patents as a Measure of Innovation

We test the robustness of patenting as an indicator of innovative activity by looking at the correlation between patent applications and other measures of innovation, specifically, R&D expenditures. We calculate R&D expenditures for a sample of European firms that we obtain by matching firm level data from Amadeus database with data from EU-EFIGE/Bruegel-UniCredit dataset (henceforth EFIGE).²⁹ For our purposes, EFIGE contains information of firms' technological innovation and R&D activities. We use the self-reported average share of R&D expenditures for the period 2007-2009,³⁰ and we obtain turnover data from Amadeus. In this way we are able to calculate R&D expenditures for the sample of firms.

We proceed by calculating the correlation between firm level R&D expenditures (in logs) and the number of patent applications (in logs) for each firm. We calculate firms' patents from PATSTAT raw data, after having matched Amadeus-EFIGE firms with PATSTAT.^{31,32} In order to account for the lag between the investment in R&D and the successful outcome of the R&D process and subsequent patent application, we calculate the average number of patents applied for per year by a firm by considering a window of six years. We include the survey period (2007-2009) and the three subsequent years, until 2012. Our results can be summarized as follows. On the intensive margin, higher R&D expenditures are strongly correlated with a higher number of patent applications. Figure 12 shows a kernel-weighted local polynomial regression of firms' R&D expenditures on number of patent applications (both in logs) for the survey period. The relationship between the number of patents filed by a firm and its investment in R&D is strong and positive. This relationship is not monotonic. We notice a drop for very high numbers of patent applications; but only a minor number of firms file such a high number of patent applications per year. The corresponding linear regression slope is 0.68 (s.e. 0.05).

On the extensive margin, firms with at least one patent application spend on average more on R&D than firms with no patents. We use firm level R&D expenditures and construct a binary variable, which equals one if the firm has applied for one or more patent on average in the 2007-2012 period, and zero otherwise. Figure 11 shows the histogram of average R&D expenditures for firms with positive patent applications and for firms that didn't file any patent. The shape of the distribution is very similar in the two groups, but for firms with patents the distribution is shifted to the right, suggesting a positive correlation between


²⁹See Appendix F and Altomonte and Aquilante (2012) for a more detailed description of EFIGE data.

³⁰Precisely, question C21 in EFIGE asks: "Which percentage of the total turnover has the firm invested in R&D on average in the last three years (2007-2009)?"

³¹See Appendix F for a description of how PATSTAT and Amadeus were matched.

³²Our preferred way calculates the number of patents from PATSTAT raw data, since this gives a more complete measure of patent applications per firm, but results are robust to using patent counts from Amadeus.

Figure 11: R&D expenditures and patenting




Note: The figure shows the distribution of firms' R&D expenditures (in logs) for firms with (white) and without (gray) patent applications in the period 2007-2009. R&D expenditures are calculated combining EFIGE and Amadeus data. The binary variable indicating whether a firm did, on average, apply for patents, is calculated using patent counts from PATSTAT, after matching Amadeus with PATSTAT data. The population of firms consists of surveyed firms in EFIGE, subsequently matched to Amadeus and, for firms holding patents, to PATSTAT.

R&D expenditures and patenting. For high levels of R&D investments, there is a higher share of firms with at least one patent application. Conversely, for low levels of R&D, the share of firms with no patent applications is higher. We also run a correlation between firm level R&D expenditures and the binary variable indicating whether, on average, the number of patent applications per year in the 2007-2012 period is positive.³³ We repeat the same exercise for both the level and the log of R&D expenditures. The results are reported in column one and two of Table 7 respectively. In both cases we find a positive and strong correlation between R&D expenditures and patent applications.

³³In the simplest methodology described above, firms are classified into having patents or not based on their answer to the question of whether the firm has applied for a patent during the period considered in the EFIGE survey. In this case, the binary variable for positive patent applications is calculated based on patent counts from Patstat raw data.

Figure 12: R&D expenditures and patenting: Intensive margin



Note: The figure shows the average number of patent applications per year and average R&D expenditures per year (both in logs). R&D expenditures refer to the period 2007-2009, patent counts are calculated over a six year window, from 2007 to 2012, to take the lag between R&D investment and subsequent patent applications into account. The solid line is the local polynomial regression fit and the gray area represents the 95% confidence bands. The linear regression slope is 0.68 (s.e. 0.05). The population of firms consists of surveyed firms in EFIGE, subsequently matched to Amadeus and, for firms holding patents, to PATSTAT.

Table 7: R&D expenditures and patenting

	(1)	(2)
Dependent variable	R&D expenditures (average)	R&D expenditures (logs)
has_patent2	3570.16 ^c (584.17)	1.28 ^c (0.06)
Observations	6204	6074

Standard errors in parentheses

^a $p < 0.1$, ^b $p < 0.05$, ^c $p < 0.01$

The table shows a regression of R&D expenditures on a binary variable indicating whether the firm has any patent. R&D expenditures are calculated combining EFIGE and Amadeus data. The binary variable indicating whether a firm did, on average, apply for patents, is calculated using patent counts from PATSTAT for the period 2007-2012. The population of firms consists of surveyed firms in EFIGE, subsequently matched to Amadeus and, for firms holding patents, to PATSTAT.