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

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

January 2018

Abstract

This paper investigates the effect of improved market access on worldwide innovation. Using exogenous changes in tariffs in firms' export markets during the 1990s, we find a large effect of tariff cuts on innovation as measured by patent data. These effects are not driven by the deterioration of innovation quality, and the results are robust to controlling for changes in the patent system and to industry-wide trends in innovation.

*We would like to thank Swati Dhingra and Rachel Griffith as well as numerous seminar and conference participants for helpful clarifications and valuable discussion. We also thank Bjarne J. Kvam from the Norwegian Patent Offce (Patentstyret) for guidance related to the patent data.

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

[‡]University of Oslo & CEPR, 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 abroad and therefore increases the effective size of the market. Economists have long known that the amount of invention is governed by the extent of the market.[1] However, comprehensive empirical evidence of broad and deep changes in trade policy on worldwide innovation is nevertheless scarce. 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 in large part a result of the GATT Uruguay Round, spanning the years 1986 to 1994 and phased in from 1995 to 2000. On average, developed country tariffs were cut from around 6 to 3 percent, while developing country tariffs were cut from almost 20 to 13 percent between 1990 and 2000.[2] This paper uses the Great Liberalization as a quasi-natural experiment to estimate the causal impact of improved market access on innovation among firms from more than 100 different countries, representing over 90 percent of global gross domestic product.

A major empirical concern in any study of the effect of market access on innovation is the endogeneity of tariffs, e.g. that highly innovative industries and countries may choose a low-tariff economic environment. We overcome this issue by exploiting variation in applied most favored nation (MFN) tariff cuts in firms' initial *export* markets, which are exogenous to other determinants of innovation in the home country and industry of the firm. Specifically, we link tariff data to the initial industries and foreign countries the firm is active in, in order to compute the average tariff cut faced by the firm. 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. Those tariff cuts are unlikely to be caused by the German and Japanese firm (conditional on industry-country trends). We also instrument the tariff cuts and show that the instrumental variables estimates are comparable to the baseline estimates.

The data requirements for this exercise are large; one would ideally need a firm-level panel data set on innovation, along with detailed information on where firms are located and in which markets they sell in. To achieve this, we construct a global and comprehensive firm-level data set on patenting using the data base 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. We do not directly observe in which markets firms sell in, but we observe where

[1]An early contribution is Schmookler (1966).

[2]See Section E in the Appendix for details.

firms are patenting. Hence, we follow Aghion et al. (2016) and construct firm-level measures of country exposure by using information on patent filing in the years prior to the Uruguay negotiations. The weights based on patent filing may potentially be superior measures of market exposure because they reflect the firms' *expectations* of where their future markets will be. Moreover, we provide evidence that patent weights are strongly correlated with export weights.

Our firm-level approach has a number of advantages. First, because initial country exposure varies significantly within a country and within narrowly defined industries, our global approach allows us to sweep out all home country-industry trends in innovation by fixed effects. This is crucial for three reasons. First, tariff cuts may lead to greater import competition that is known to affect innovation (e.g., Bloom et al., 2016). Second, the likelihood of patenting depends on a host of time-varying factors such as the legal framework and technological characteristics of the industry.[3] Third, changes in tariff policy often go together with other policy changes, such as product market deregulation. Using information on firms' exposure to different markets allows us to isolate the role of trade policy, sidestepping all of these issues.

A number of factors that are correlated with tariff cuts in the destination country may also affect innovation. We observe aggregate patenting in all countries and industries, and can therefore employ a control function approach to flexibly deal with such factors. An example of this is market size. Being exposed to a high-tariff cut country may be correlated with innovation simply because this country grows faster and increased market size fosters innovation. The control function approach, or alternatively including a *vector* of fixed effects for each destination market, eliminates this concern. Finally, 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.

Our results show that the Great Liberalization of the 1990s had a large positive net impact on innovation. The overall estimates mask considerable heterogeneity across countries and industries. For example, firms in developing countries experienced greater tariff cuts in their export markets compared to firms in developed countries, implying that the boost to innovation was stronger in developing countries. While our data do not include other firm-level outcomes than patenting, other researchers have found an economically significant impact of patenting on firm-level productivity (Bloom and Van Reenen, 2002). Hence, our results suggest that improved market access not only produced more innovation but also productivity growth during this period.

One may question whether increased patenting reflects more innovation. The literature

[3]Such as regulatory changes in the patent system and differences across patent offices.

typically finds a strong correlation between patenting and research and development, and between patenting and other measures of innovation. We also find a strong positive correlation between patent counts and other innovation indicators in our own data.[4] But the concern remains that more trade could induce the need for greater protection of intellectual property rights (IPR), i.e. that more patenting can simply be attributed to a "lawyer effect". To deal with this, we calculate citation counts for all firms in our data set to control for the quality of a patent, and check whether average citations are falling in response to trade liberalization. The data rejects this hypothesis, if anything, average citations are rising in response to trade liberalization. Alternative measures correlated with the economic value of patents confirm that trade liberalization has not led to a reduction in patent quality.

The contributions of this paper are as follows. First, we develop a simple theory and a novel empirical methodology consistent with the model. This allows us to produce broad and systematic evidence of the impact of market access on worldwide innovation over a decade with steep global tariff declines. This provides external validity compared to the current literature that has primarily focused on specific industries or countries. Moreover, there is a large literature on the impact of trade policy on firm performance (e.g., TFP or labor productivity), but there is little direct evidence on observable measures of innovation such as patents. Third, we construct and analyze a novel, comprehensive and global firm-level patent data set that has so far not been applied in the context of international trade.

Our analysis thus speaks to different strands of literature. Our work is related to the empirical analyses of firm level data on the impact of trade liberalization on firm performance such as Amiti and Konings (2007), Goldberg et al. (2010), Khandelwal and Topalova (2011) and Loecker et al. (2016). Our work also relates to the literature on trade, import competition and technology adoption. Bloom et al. (2016) analyze the effect of Chinese import competition on technology upgrading in Europe, while Autor et al. (2016) examine the impact of China competition on patenting in the U.S. While these studies focus on the effect of reduced import tariffs in firms' home market, we instead control for import competition while estimating the effect of improved market access.[5]

We also relate to Bustos (2011) and Lileeva and Trefler (2010), who analyze complementarities between trade liberalization and technological upgrading and innovation. What distinguishes our paper from these contributions is that (i) we develop a new identification strategy, (ii) we offer evidence at the economywide and worldwide level, providing external validity and (iii) we use patents as a direct output-based measure of innovation, rather than

[4]See Section I in the Appendix.

[5]Aghion et al. (2017), Boler et al. (2015), Gopinath and Neiman (2013) and Halpern, Koren, and Szeidl (2015) also examine the link between firm performance and trade, but do not use variation in tariffs.

input-based or survey information.[6]

Finally, our empirical approach is related to Aghion et al. (2016) and Calel and Dechezleprêtre (2016) which also use PATSTAT data, but focus on very different questions, namely the impact of environmental policies on technical change. Our choice of approach and results do not only inform the literature on trade policy but also the broader literature on the effects of the drivers of innovation (see e.g. Acemoglu and Linn, 2004, Aghion et al., 2005, Bloom et al., 2016 and Griffith et al., 2010).

The rest of the paper is organized as follows. Section 2 presents our theoretical framework. Section 3 lays out the empirical model and highlights econometric issues. Section 4 describes the data and descriptives. Section 5 presents and discusses the empirical results, Section 6 examines the underlying economic mechanism while Section 7 concludes.

2 Theoretical Framework

We aim to investigate the effect of foreign market access on firms' innovation. To do so, we start by presenting a basic economic framework to support the analysis, and proceed by developing predictions for the relationship between market access and innovation.

Consider a firm i with productivity z_i, located in country m and producing in industry j with constant returns to scale using only labor. Goods sold from m to a foreign country n in industry j are subject to an ad-valorem tariff $T_{jmn} = \tau_{jmn} - 1 \geq 0$. Preferences across varieties within an industry are CES with an elasticity of substitution σ. This gives rise to a demand function $A_{in} p_{imn}^{-\sigma}$ in country n, where p_{imn} is the price charged by firm i in n and the demand shifter A_{in} may vary across firms and countries, and is exogenous from the point of view of an individual firm.[7] Producers engage in monopolistic competition, so that the price charged by firm i in market n is $p_{imn} = [\sigma / (\sigma - 1)] \tau_{jmn} w_m / z$, where w_m is the wage of country m. For expositional clarity, we normalize the wage to one, as it will be inconsequential for the remaining analysis. The profits from serving country n is $\pi_{imn} = (z/\tau_{jmn})^{\sigma-1} B_{in}$, where $B_{in} = (1/\sigma)[(\sigma-1)/\sigma]^{\sigma-1} A_{in}$. Global profits are then

$$\Pi_i = \sum_n \left[\left(\frac{z_i}{\tau_{jmn}} \right)^{\sigma-1} B_{in} \right].$$

The firm faces the problem of how much to innovate. Consider the simplest possible case

[6] Steinwender (2015) also documents the relationship between access to export markets and productivity increases in the case of Spain.

[7] Given the CES structure, $A_{in} = E_{in}/P_{jn}^{\sigma-1}$, where E_{in} is a demand shifter and P_{jn} is the CES price index of industry j.

where productivity z is proportional to the firm's stock of knowledge K_i, $z_i = \xi K_i$. We discuss the measurement of K_i in Sections 3 and 5.2. Gaining new knowledge is costly, and we assume that the cost of obtaining a stock of knowledge K_i is $c(K_i) = \psi K_i^k$, where ψ determines average innovation cost and $k > \sigma - 1$ determines how quickly those costs rise with knowledge. The firm then chooses the optimal K_i that maximizes global net profits, $\Pi_i - c(K_i)$. Appendix A shows that the optimal knowledge stock is

$$K_i = \kappa \left(\sum_n \tau_{jmn}^{1-\sigma} B_{in} \right)^{1/[k-(\sigma-1)]}, \tag{1}$$

where κ is a positive constant.[8]

Now consider a change in τ_{jmn} from one equilibrium to the next. Using the exact hat algebra approach as popularized recently by Dekle et al. (2008), we get

$$\hat{K}_i = \left(\sum_n \omega_{in} \hat{B}_{in} \hat{\tau}_{jmn}^{1-\sigma} \right)^{1/[k-(\sigma-1)]}, \tag{2}$$

where $\omega_{in} = \pi_{in}(z) / \sum_o \pi_{io}(z)$ is the share of global profits coming from market n in the initial equilibrium, and the hat notation denotes the value in the counterfactual relative to the initial equilibrium, i.e. $\hat{x} \equiv x'/x$. Equation (2) highlights two important economic mechanisms. First, all else equal, tariff cuts (lower τ_{jmn}) lead to both higher profits and a greater knowledge stock. Intuitively, a larger effective market means that a marginal improvement in productivity or quality yields a higher return. Second, our theory shows that tariff cuts in large markets matter more for innovation compared to tariff cuts in small ones, and that the theoretically correct weight is the initial share of global profits in that market. Note that tariff cuts have general equilibrium effects that will show up in $\hat{B}_{in}$. Specifically, B_{in} is a function of the price index in the market, which is again a function of tariffs. Our empirical approach will capture both the direct impact of $\hat{\tau}$ and the indirect impact from the price index. Also note that by construction $\hat{\tau}_{jnn} = 1$, i.e. a firm is never charged a tariff when selling to its home market. Our model therefore zooms in on one particular channel through which tariff cuts affect the incentives to innovate: foreign market access.

Before concluding this section, we briefly discuss three possible extensions of the model. First, in our model, tariff cuts only matter if the firm is already exporting to a destination, i.e. if ω_{in} is strictly positive. In practice, firms may choose to both start exporting to country

[8] $\kappa \equiv \left[\xi^{\sigma-1} (\sigma - 1) / (k\psi)\right]^{1/[k-(\sigma-1)]}$. The second order condition for profit maximization is satisfied given that $k > \sigma - 1$.

n and innovate as a response to tariff cuts in n. We investigate this case theoretically in Appendix B and empirically in Section 5.3. Moreover, we show in Appendix G that there is a striking degree of persistence in ω_{in} over time, suggesting that exit or entry into new markets is limited in our dataset.

Second, the partial equilibrium approach chosen here means that we abstract from general equilibrium effects. For example, it may well be that trade policy induces entry of a particular type of firms with different innovation rates. The approach taken in this paper is guided by the empirical analysis and identification strategy. As will become clear, our unit of analysis is the firm and we require pre-sample data on ω_{in} for all the firms in the dataset (which by construction excludes entrants from the analysis).

Third, we have omitted other economic factors that may also affect innovation. For example, it is well known that a more competitive marketplace, e.g. coming from import competition, may give rise to both less or more innovation (e.g., Aghion et al., 1997 and Aghion et al., 2005). Again, our approach is guided by the empirical analysis and main question of the paper, namely how changes in a firm's foreign market access affect innovation (as opposed to changes in the economic environment in the home market). Our empirical approach will, however, flexibly control for the impact of import competition on innovation.

3 The Empirical Model

This section presents the main empirical model and discusses our identification strategy. Appendix C shows that equation (2) can be approximated by

$$\Delta \ln K_i = \beta \Delta \bar{T}_i + \varepsilon_i, \tag{3}$$

where

$$\Delta \bar{T}_i \equiv \sum_n \omega_{in} \Delta T_{jmn} \tag{4}$$

is the weighted average of tariff changes across all of firm i's export markets, T_{jmn} is the ad-valorem tariff from country m to n in industry j, $\beta \equiv (1-\sigma) / [k-(\sigma-1)]$ and $\varepsilon_i \equiv [k-(\sigma-1)]^{-1} \sum_n \omega_{in} \Delta \ln B_{in}$. We proceed with this approximation because it is empirically more convenient to work with. According to our framework, we expect that the knowledge stock is changing when weighted average tariffs in export markets decline or when weighted average demand (ε_i) rises. As demand shocks are unobserved in our data, ε_i will enter into the regression residual.

Sample period. The years 1992 to 2000 are defined as our baseline sample period. Hence, the change in average tariffs facing firm i is $\Delta\bar{T}_i = \bar{T}_{i2000} - \bar{T}_{i1992}$ and the change in the knowledge stock of firm i is $\Delta \ln K_i = \ln K_{i2000} - \ln K_{i1992}$. The choice of sample period 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 also confirms that the 1990s was unique: the overall reduction in tariffs was much greater during the 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 (see Bertrand et al., 2004).

Outcome variable. In the model presented above, 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,

$$K_{it} \equiv \sum_{s=1965}^{t} p_{is}, \tag{5}$$

where p_{is} is the number of unique granted patents filed by firm i in year s. The outcome variable $\Delta \ln K_{it}$ gives the change in the log cumulative patent count between 1992 and 2000 and provides a measure of the innovation that takes place during this time period. 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.

Econometric concerns. Estimating equation (3) is challenging for a number of reasons. The first potential issue is that tariff cuts might be endogenous. We overcome this in two alternative ways. First, $\Delta\bar{T}_i$ is constructed based on the tariff a firm is facing in its *export markets*, and in our baseline analysis we restrict the analysis to applied MFN tariff rates (see Section 4), which means that we do not utilize variation coming from bilateral or regional trade agreements. Endogeneity would then only be a concern if e.g. a Norwegian exporter might influence the MFN tariff of the U.S., which seems relatively unlikely. Formally, then, $\Delta T_{jmn} = \Delta T_{jkn}$ for $m, j \neq n$, and $\Delta T_{jnn} = 0$ (i.e., firms do not incur tariffs when selling domestically). Second, the alternative approach is to find an instrument for ΔT_{jmn}. We explore this in the robustness section below. Of course, the MFN tariff might not always be the relevant tariff facing a given firm because of regional trade agreements. We discuss this further in Section 4 and 5.3.3.

The second econometric concern is that the weighted average tariff reduction $\Delta\bar{T}_i$ may be correlated with unobservable firm characteristics. 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.[9] Intuitively, we compare firms within the same narrowly defined 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. This approach also ensures that changes in the patent system or industry-specific trends in patenting are all differenced out.

An alternative way of solving this problem is by differencing out idiosyncratic firm trends. Specifically, we split the sample into our main sample period, ($t = 1$) and add a second period ($t = 2$), and estimate the equation

$$\Delta \ln K_{i2} - \Delta \ln K_{i1} = \beta \left(\Delta\bar{T}_{i2} - \Delta\bar{T}_{i1}\right) + \varepsilon_i. \tag{6}$$

Idiosyncratic growth trends in innovation that may be correlated with $\Delta\bar{T}_i$ are then differenced out. This is reminiscent of a triple differences model, as we compare the growth in the change in tariffs (two differences) across firms (third difference).[10]

A third econometric concern is that the error term ε_i, which is a weighted average of country-specific demand shocks, 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.[11] We solve this by using a control function approach and the fact that we observe aggregate patenting by industry and country, and this measure is itself determined by the unobserved shocks B_{in}. 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,

[9] Industries 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 log knowledge stock of the firm in 1985, $\ln K_{i,Pre}$.

[10] $t = 1$ is the baseline period 1992 to 2000 and $t = 2$ is the years 2000 to 2004. While tariffs fell strongly during the first period, the decline was much smaller in $t = 2$ (see Figure 1 below).

[11] TRIPS established minimum and common standards of IP protection to be adopted by all WTO members. While the institutions in the developed countries were little affected due to already strong IP protection, developing countries had to reform and strengthen their IP protection system to comply with new WTO rules.

and construct the weighted average

$$\tilde{\varepsilon}_i \equiv \sum_{n \in \Omega_i} \omega_{in} \Delta \ln \mathcal{K}_{nj}, \tag{7}$$

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 U.S. headquartered firm primarily exposed to the Indian market is innovating more because the Indian market is growing quickly (high $\Delta \ln B_{iIndia}$), then including $\tilde{\varepsilon}_i$ will control for this effect. An alternative approach is use a *vector* of fixed effects for each of firm i's destination markets. We explore this approach, along with other robustness checks, in Section 5.3 below.

4 Data

4.1 Patents

Our main data set is based on the European Patent Office's (EPO) Worldwide Patent Statistical Database (henceforth PATSTAT).[12] PATSTAT offers bibliographic data, family links and citations of 90 million applications of more than 100 countries. It contains the population of all patents globally since the mid 1960s. The patent documents as provided by PATSTAT are a rich source of information. We observe the name of the applicant and date of filing, publication and if and when the patent was granted. We have information on citations, technology areas (IPC codes) and the research teams behind the inventions. 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 and JPO). Appendix D provides more details on the construction of our data set, while Abramovsky et al. (2008) provides a thorough review of the PATSTAT data and the patenting process.

Firm-specific knowledge stocks. PATSTAT allows us to construct an international firm-level data set and to follow the patenting activity of a firm through time. To measure the innovation activity of a firm i in year t we shall use a count of granted patents dated by application/filing year (p_{it}). Not all filed patents are granted. Hence, to account for differences in quality we limit the analysis to patents that are granted.[13] Dating the patents by application filing date is conventional in the empirical innovation literature as it is much

[12]The April 2015 version.

[13]To be granted a patent, an innovation must satisfy three key criteria: it must be novel or new, it must involve an inventive step, and it must be industrially applicable.

more closely timed with when the R&D process took place than the patent publication and grant date.[14]

Patenting is known to be highly correlated with innovation and R&D, see e.g. Griliches (1990). In the Appendix I we document a close relationship between R&D expenditure and patenting for a subsample of our data set. The advantages and limitations of patenting as a measure of innovation have been extensively discussed.[15] For our purpose there is one major advantage of using patents. It is the only source of information that allows for a comprehensive firm-level analysis of innovation at a global scale. In Section 5.2 we use different measures to control for the quality of patents as innovation indicators.

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 identify identical inventions filed in multiple countries.[16] 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.[17] Information about firms in PATSTAT is restricted to what can be retrieved from the patent applications. Our basic firm characteristics are industry affiliation (NACE Rev. 2 3-digit), home country of the firm, as well as in which countries the firm is patenting.[18]

Firm-specific weights. Based on our theoretical model, see Section 2, the impact of trade liberalization will depend on firm-specific weights. These weights reflect the relative importance of a country n in the firm's total profits. Profits and sales are unobserved in our patent data, but we do observe in which markets a firm is patenting. As pointed out by Aghion et al. (2016), a patent based weighting scheme may potentially be a superior measure because it reflects the firms' *expectations* of where their future market will be. We calculate these weights based on patent filings over the pre-period years 1965 to 1985. 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. Specifically, we define

$$\omega_{in} \equiv \frac{x_{in}}{\sum_k x_{ik}}, \tag{8}$$

[14]Patent applications are usually published 18 months after the first application.

[15]See e.g. OECD (2009), Griliches (1990) and Nagaoka et al. (2010) for reviews and discussion of patent data as innovation indicators.

[16]We use DOCDB patent family.

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

[18]Home country and industry affiliation are missing for a subset of firms. These observations are dropped from the dataset.

where x_{in} is the number of patents issued by firm i 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 (see Aghion et al., 2016). We provide additional empirical evidence on this in Appendix G. The weights are also remarkably persistent over time, even over a period of 20 years, see Appendix H. This suggest that time-invariant firm and country characteristics (e.g. country-specific entry costs on the supply side or idiosyncratic taste differences on the demand side) are limiting where firms export goods and file patents.

Calculating the firm-level weights based on the pre-period is done in order to minimize the risk that the weights themselves are endogenous. Nevertheless, it may be that some firms anticipate the tariff reductions taking place from 1995 and therefore adjust their 1980s exposure to different markets. We deal with this by calculating industry-specific weights that in part reflect geography, e.g. that Canadian that are more likely to be exposed to the U.S. market compared to Indian firms, see Section 5.3.1.

4.2 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 applied MFN industry-level tariff (NACE 3-digit) for the period 1992 to 2009. We use these to calculate the firm-specific weighted average tariffs, $\bar{T}_i$ which vary across firms, both because firms are exposed to different markets and because they belong to different industries. Appendix E describes the procedure followed to calculate industry-level tariffs, while Appendix F provides details about the historical background for tariff reductions during the 1990s.

Of course, the applied MFN tariff might not always be the relevant tariff facing a given firm because of regional trade agreements. We deal with this in two ways. First, in the baseline sample, the EU-15 is aggregated to one single economic unit, so that EU MFN tariff cuts do not affect the $\bar{T}_i$ of e.g a Dutch firm exporting to France. Second, we use information on regional trade agreements (RTAs) between pairs of countries in the robustness section. The information on RTAs for around 200 countries from 1948 to 2006 comes from the CEPII gravity data set.[19]

[19]See Head et al. (2010) and Head and Mayer (2013).

4.3 Final Sample of Firms

Our point of departure is a data set constructed on the basis of PATSTAT described in Section 4.1 and Appendix D, matched with the average applied MFN industry-level tariff from UNCTAD TRAINS. The final sample consists of the following firms: First, it includes firms that applied for at least one granted patent by 1992 (so that $\Delta \ln K_{it}$ is non-missing). Second, firms must be observed at least once in the pre-period (1965-1985) in order to be assigned weights ω_{in}.[20] Third, because the focus of our analysis is on foreign market access, firms need to have positive weight ω_{in} in at least one foreign country. Fourth, in some cases firms issue patents in countries with missing tariff data for their industry; these firms are dropped from the analysis.[21]

The last column of Table 1 shows that we have roughly 59,000 firms in the final sample, filing about 800,000 patents between 1992 and 2000. Our final sample consists of firms from 110 different countries, representing more than 90 percent of global gross domestic product. The initial sample of firms, which also includes firms without pre-period weights ω_{in}, consists of about 1 million firms filing about 4 million patents between 1992 and 2000. Hence, our final sample captures roughly 20 percent of global patenting over the sample period.

Note that we cannot distinguish between firm exit and zero innovation in our data. For example, if we observe zero patenting from 1995 and onwards, then p_{it} will be zero and K_{it} will be constant for the remaining years of our sample. Hence, our baseline result will capture the overall innovation effect of trade policy, including the impact on firm exit.

[20]For the pre-period we do not require patents do be granted. This is because weights reflect expectations on where future markets will be, and therefore what is relevant is the action of seeking intellectual property protection in a foreign country, rather than the final outcome of the application process.

[21]We drop firms that have positive weights ω_{in} for one or more countries with missing tariff data, i.e. if T_{jnt} is missing when calculating $\bar{T}_{it}$ from equation (4).

Table 1: Initial versus final sample.

	Initial	Final
$\sum_i \Delta K_{it}$	4,275,647	785,064
ΔK_{it}		
..mean	0.8	13.4
..median	0	1
..standard deviation	54.0	257.5
Number of firms	1,061,022	58,785

Note: The table shows the aggregate increase in the knowledge stock from 1992 to 2000 along with the mean, median and standard deviation of ΔK_{it} for the initial and the final sample of firms. The initial sample consists of all firms identified with a headquarter location and a NACE industry code that file at least one granted patent over the sample period (1992-2000).

4.4 Descriptives

Weighted average trade barriers. To illustrate our identification strategy, we take a closer look at the weighted average trade barriers, $\bar{T}_{it}$, for our sample of firms. 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 the latter half of the 1990s; the average firm experienced a decline in weighted tariffs 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 weighted tariffs ($\bar{T}_{it}$) across firms shifts markedly to the left from 1992 to 2000.

Patenting. Figure 3 shows the development in the mean number of patents per firm by year in our sample (p_{it}). We observe that the mean number of patents is steadily increasing from 1980 and onwards. Of course, these aggregate trends may not only reflect innovation, but also changes in firms' behavior, legal trends and changes in the patent systems worldwide.

Figure 4 shows the distribution of patenting firms across home countries and industries (NACE 2-digit) in our sample. We note the dominance of Japan and the US and by the industries machinery and equipment (28), computers, electronic and optical products (26), and other manufacturing (32). Tables 11 and 12 in Appendix J provide more details on patent counts and patenting firms across industries and countries.

Figure 1: Average Firm-Specific Tariffs, $\bar{T}_{it}$.

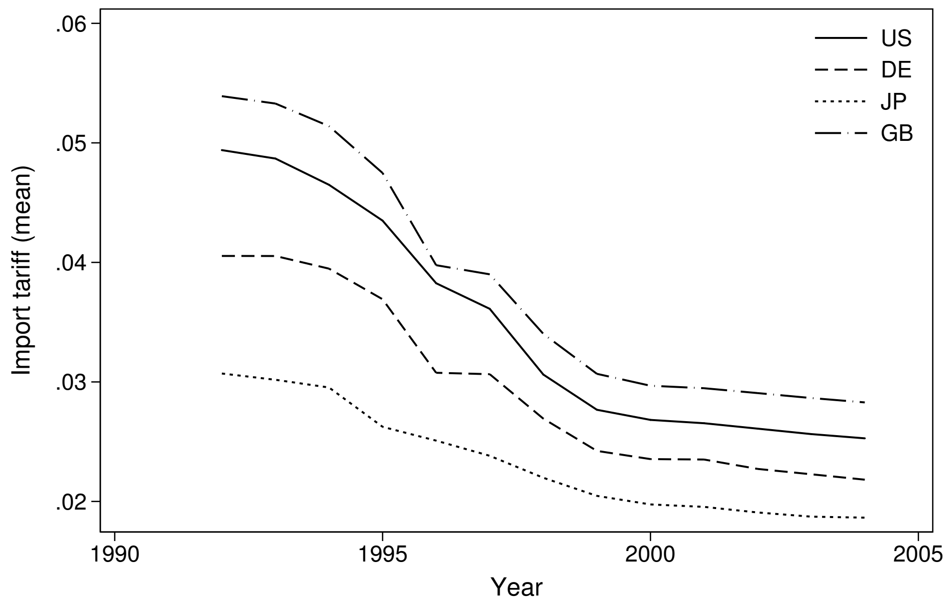

Note: The figure shows the annual average $\bar{T}_{it}$ across firms according to headquarters country.

Figure 2: Density of Firm-Specific Tariffs, $\bar{T}_{it}$, in 1992 and 2000.

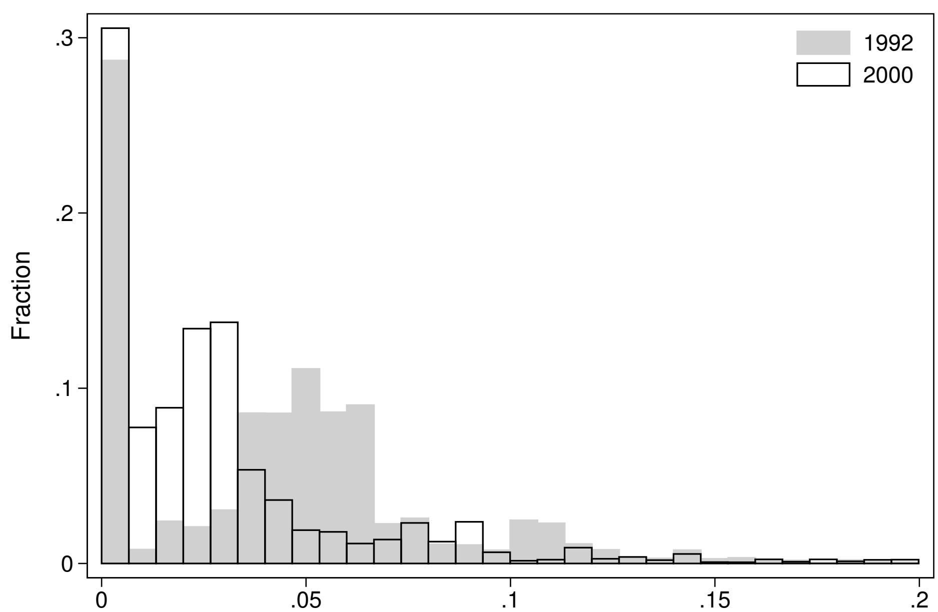

Note: $\bar{T}_{it}$ is the weighted average import tariff in firm i's markets, in 1992 and 2000. For expositional purposes the histogram is truncated at $\bar{T}_{it} = 20$.

Figure 3: Patenting per firm. 1980-2004.

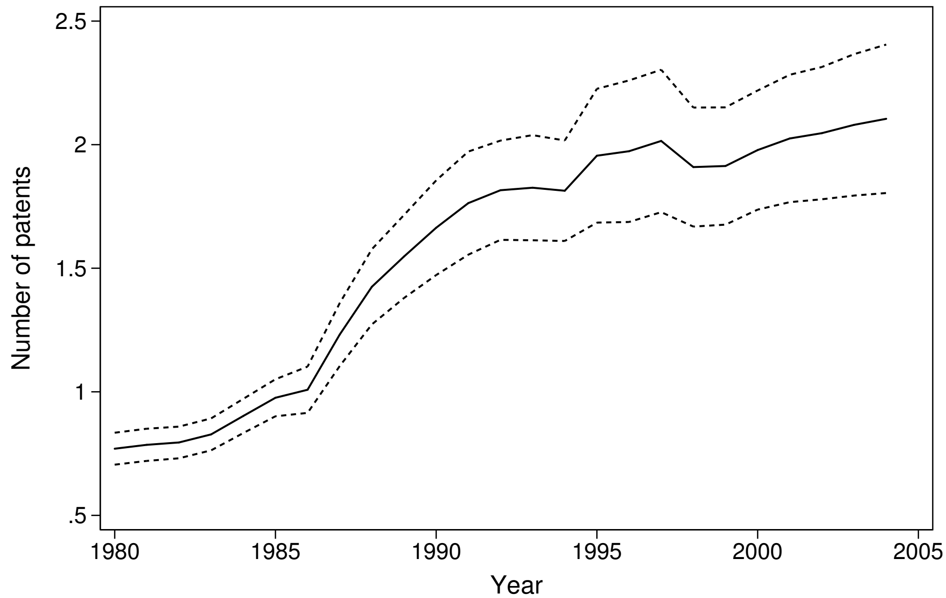

Note: The solid line shows the average number of patents filed per firm per year in the sample. The dotted lines show the 95 percent confidence interval associated with the mean..

Figure 4: Share of Firms by Country and Industry

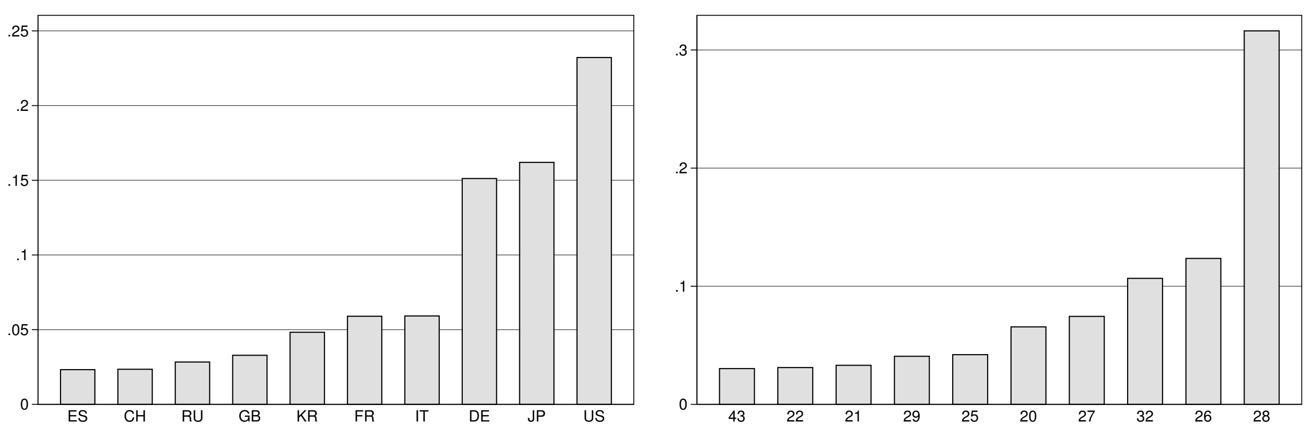

Note: The figure shows share of firms by home country and NACE Rev. 2 2-digit industry for the period 1992-2000. Only the top 10 countries/industries are shown.

5 Results

5.1 Innovation and Trade Liberalization

We proceed by estimating the model presented in equation (3) and the alternative specification outlined in equation (6). As described in Section 3, 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 2 show the results for our baseline specification with various control variables included. Column (1) has only fixed effects and column (2) adds pre-sample firm characteristics (the home weight, ω_{iH}, the number of countries the firm is patenting in during the pre-period, $n_{i,Pre}$, and 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 3. Column (4) shows the results for the model described in equation (6), where we difference out idiosyncratic firm trends. The results are highly significant across specifications, with an estimated coefficient in the range of -0.9 to -2.1. These results strongly suggest that foreign market access leads to significantly higher innovation.

A semi-log elasticity of -1.6 implies that a one percentage point reduction in tariffs causes a 1.6 percent increase in the knowledge stock of a firm over a period of eight years. As a simple back-of-the-envelope exercise, we ask how large our estimates are compared to the mean growth in the knowledge stock over the sample period. Our data shows that over the period 1992 to 2000 the mean knowledge stock globally grew by 43 percent, 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 11 ($1.6 \times 3/43$) percent of the observed increase in the knowledge stock can be explained by trade liberalization. This overall number masks considerable heterogeneity across countries and industries. For example, among developing countries firm-specific tariffs fell seven percentage points on average, while the mean knowledge stock grew by 38 percent - suggesting that trade policy explains roughly thirty percent of the increase in the knowledge stock ($1.6 \times 7/38$).[22]

[22]Developing countries are defined according to the World Bank 1995 definition of high/low income countries.

Table 2: Trade Policy and Knowledge Creation

Dep. variable:	$\Delta \ln K_i$ (1)	$\Delta \ln K_i$ (2)	$\Delta \ln K_i$ (3)	$\Delta \ln K_{i2} - \Delta \ln K_{i1}$ (4)
Change in tariff ($\Delta \bar{T}_i$)	-2.11^a	-2.14^a	-1.63^a	-.93^a
	(.35)	(.35)	(.37)	(.23)
Home country-industry FE	Yes	Yes	Yes	Yes
Firm controls	No	Yes	Yes	Yes
Destination market controls ($\tilde{\varepsilon}$)	No	No	Yes	Yes
Number of firms	58,785	58,785	58,679	58,679

Note: Standard errors clustered by home country-industry in parentheses. Firm controls are pre-sample firm characteristics: the home weight, ω_{iH}, the number of countries the firm is patenting in during the pre-period, $n_{i,Pre}$, and log knowledge stock in 1985, $\ln K_{i,Pre}$. ^{a} p< 0.01, ^{b} p< 0.05, ^{c} p< 0.1.

5.2 Is Patenting a Good Measure of Innovation?

As pointed to in Section 4.1, one may argue that patents are an imprecise measure of knowledge and innovation. Patenting is not the only way to protect innovations. Another problem is that patent quality is highly heterogeneous. According to Nagaoka et al. (2010) roughly half of the patents owned by a firm are used either by them internally or licensed to others. The remaining patents are used for strategic reasons, e.g. attempts to block inventions by competitors. Hence, it is possible that firms take out more patents, without innovating more, in response to e.g. import competition. If this were the case, one would expect that firms are taking out patents on their marginal innovations, so that the average quality of their patent stock is decreasing.

To address this issue, we use three different proxies for patent quality: the number of citations, the size of the research teams behind a patent, and the number of technology areas (IPC codes) to which a patent is attributed (patent breadth). We use citations because high value inventions are more extensively cited than low value patents (Harhoff et al., 1999). We include the size of research teams since a set of studies have associated the number of inventors listed in a patent with the economical and technological value of the patent (OECD, 2009). Finally we include number of technical classes attributed to a patent application (patent breadth) which has been found to be a measure of the value of a patent portfolio (see e.g. Lerner, 1994).

We calculate average quality of the knowledge stock as follows. Let q_p denote the number of citations three years after a patent p was filed, or the number of inventors or the number of IPC codes associated with patent p. The cumulative sum is then

$$Q_{it} = \sum_{s=1965}^{t} \sum_{p \in \Xi_{is}} q_p, \tag{9}$$

where Ξ_{is} is the set of firm i's patents filed in year s. The average quality of the knowledge stock is then calculated as $\bar{Q}_{it} = Q_{it}/K_{it}$. We proceed by using $\Delta \ln \bar{Q}_i = \ln \bar{Q}_{i2000} - \ln \bar{Q}_{i1992}$ as the dependent variable and estimate our baseline model again.

The results using all three proxies for quality are reported in columns (1)-(3) of Table 3.[23] 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 point estimates indicate that trade policy may have increased the quality of patents.

Table 3: Trade Policy and Innovation Quality.

Dep. variable: $\Delta \ln \bar{Q}_i$	Citations (1)	Research Team (2)	IPC codes (3)
Change in tariff ($\Delta \bar{T}_i$)	-1.00^a	-.22^a	-.01
	(.28)	(.06)	(.08)
Home country-industry FE	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes
Destination market controls $(\tilde{\varepsilon})$	Yes	Yes	Yes
Number of firms	37,329	56,880	58,627

Note: Standard errors clustered by home country-industry in parentheses. Firm controls are pre-sample firm characteristics: the home weight, ω_{iH}, the number of countries the firm is patenting in during the pre-period, $n_{i,Pre}$, and log knowledge stock in 1985, $\ln K_{i,Pre}$. ^{a} p< 0.01, ^{b} p< 0.05, ^{c} p< 0.1.

5.3 Robustness Checks

5.3.1 Entry to New Markets

As discussed in Section 2, firms may also innovate more because of improved market access to countries they were initially not exporting to. As outlined in Appendix B, the relevant weights ω_{in} are then based on the potential gross profits across all markets the firm may export to after the tariff cuts. We implement this by replacing firm-level weights by country-pair-industry-level ones. Specifically, we calculate aggregate weights from country m to n in

[23] The number of firms in the sample decreases when we use citations as a measure, since some firms have portfolios of patents that are never cited.

industry j as:

$$\tilde{\omega}_{mnj} = \frac{X_{mnj}}{\sum_k X_{mkj}}, \tag{10}$$

where X_{mnj} is the cumulative number of patents issued by country m in market n in industry j during the pre-period (until 1985). As shown in Appendix G, bilateral patenting adheres to a gravity model, so the share of patents from country m filed in country n are largely determined by geography and market size.

We then calculate our main independent variable as $\Delta\bar{T}_{mj} = \sum_n \tilde{\omega}_{mnj}\Delta T_{mnj}$. Note that $\Delta\bar{T}_{mj}$ is identical across firms within a country and industry. We can therefore no longer include country-industry pair fixed effects as in the main specification. Table 4 shows the results. 3-digit industry fixed effects are included in column (1), so that variation comes from comparing firms within the same industry but located in different countries. These firms happen to be exposed to different tariff cuts due to their geographical location, e.g. Canadian firms would be more exposed to U.S. tariffs and Japanese firms more exposed to Chinese ones. Column (2) also add country fixed effects. In both cases, we find a significant positive effect of market access on innovation.

Note that the alternative weights $\tilde{\omega}_{mnj}$ also address the potential concern that the firm-level weights ω_{in} are endogenous. For example, it may be that innovative firms typically choose to enter markets that subsequently cut tariffs substantially. The baseline firm-level controls and the specification with trends in equation (6) should in principle resolve this issue. Replacing firm-level weights with country-industry ones that are in part determined by gravity forces such as geographic distance and economic size is an alternative way of overcoming this issue.

Table 4: 2SLS estimates.

Dep. variable: $\Delta \ln K_{it}$	Aggregate weights		Instrumented tariffs
	(1)	(2)	(3)
Change in tariff ($\Delta \bar{T}_i$)	-1.80^a	-0.35^c	-2.06^a
	(.12)	(.18)	(.44)
Industry FE	Yes	Yes	No
Home country FE	No	Yes	No
Home country-industry FE	No	No	Yes
Number of firms	79,399	79,399	58,785
			First Stage Estimates:
$\Delta \bar{T}_i^{IV}$			-.43^a
			(.03)

Note: Standard errors clustered by industry in columns (1)-(2) and home country-industry in column (3). ^{a} p< 0.01, ^{b} p< 0.05, ^{c} p< 0.1.

5.3.2 Endogenous Tariffs

As we have argued, MFN tariff cuts in a firm's export markets are unlikely to be endogenous to a firm's innovation in the home country. Nevertheless, for robustness we proceed by instrumenting the tariff cut ΔT_{jmn} with the level of tariffs in 1992, $T_{jmn,1992}$. Specifically, we instrument $\Delta \bar{T}_i = \sum \omega_{in} \Delta T_{jmn}$ with $\Delta \bar{T}_i^{IV} = \sum \omega_{in} T_{jmn,1992}$.

Other studies typically find a strong negative correlation between initial tariff levels and subsequent tariff cuts (see Goldberg and Pavcnik, 2007, Topalova, 2010 and Loecker et al. (2016)). Also in our data, there is a strong negative correlation between the 1992 levels and subsequent changes. This occurs because tariffs in practice have a zero lower bound, so the scope for tariff cuts are by construction higher in highly protected industries and countries. The identification assumption is then that those 1992 tariff levels are not themselves a function of the knowledge growth during the 1990s (conditional on industry and country trends that are already differenced out). Column (3) in Table 4 shows the 2SLS estimate along with the first stage estimate. The 2SLS point estimate is comparable to the baseline estimate, underscoring the robustness of the results.

5.3.3 Additional Robustness checks

Falsification test. A potential concern is that firms being exposed to countries with high tariff cuts always have higher patent growth compared to other firms. To address this concern, we perform a placebo test and regress knowledge growth during the 1980s, $\ln K_{i1988} - \ln K_{i1980}$, on trade policy changes during the 1990s, $\Delta\bar{T}_{i2000} - \Delta\bar{T}_{i1992}$.[24] The results are shown in the first column of Table 5: the coefficient of interest becomes noisy and close to zero, suggesting that there are no differential pre-trends in patenting.

Country-level tariff data. Industry-level tariffs may not always be the relevant tariffs facing the firm, because it may also be exporting products associated with other 3-digit NACE industries. We therefore test the sensitivity of our results using the simple average country tariff instead of industry specific tariffs. The results, shown in the second column of Table 5, confirm our main finding that a reduction of a firm's tariffs increases innovative activity. The estimated effect is similar in magnitude to our baseline 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 the average tariff rate, $\bar{T}_{it}$, above as a weighted average of RTAs across all of firm i's markets:

$$\Delta R\bar{T}A_{it} \equiv \sum \omega_{in} \Delta RTA_{mn}, \tag{11}$$

where $\Delta RTA_{mn} = 1$ if country-pairs mn engaged in an RTA between 1992 and 2000. The results in column (3) of Table 5 show that the RTA variable is insignificant while our main variable, the change in tariffs, continues to be highly significant and negative.

Triadic patents. 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).[25] Triadic patents are commonly used in the literature to retain only highly valuable inventions and they provide a measure of innovation which is robust to administrative idiosyncrasies of the various patent offices. However, by limiting the analysis to triadic patents, the number of observations is reduced with around 94 percent. The results are shown in in column (4) of Table 5. While we observe that the sample size is reduced from around 59,000 to around 2,700 observations, our results on the impact of trade liberalization on the change in knowledge stock nevertheless remain significant and the magnitude is close to the double as we limit our analysis to these presumably highly valuable inventions.

[24] The 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 $\ln K_{i1988} - \ln K_{i1980}$.

[25] See Dernis and Khan (2004) and Martinez (2010) for additional information about how triadic patent families are constructed.

Destination country trends. The variable $\tilde{\varepsilon}_i$ was included in the regressions to capture patenting trends in destination countries. An alternative empirical strategy is to include destination country fixed effects in the regressions. Specifically, we rewrite our baseline specification to

$$\Delta \ln K_i = \beta \Delta \bar{T}_i + \sum_{n \in \Omega_i} \gamma_n + \varepsilon_i, \tag{12}$$

where γ_n is a fixed effect for destination n, and we sum over the set of countries Ω_i where the firm has non-zero weights during the pre-period. 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_{India}. 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 (5) in Table 5 shows that β is still highly significant, although the economic magnitude is somewhat lower than in the baseline specification.

Table 5: Robustness.

Dep. variable: $\Delta \ln K_{it}$	Placebo (1)	Country-level tariffs (2)	Accounting for RTAs (3)	Triadic Patents (4)	Destination trends (5)
Change in tariff ($\Delta \bar{T}_i$)	.15	-2.52^a	-1.67^a	-4.59^b	-1.38^a
	(.13)	(.44)	(.41)	(1.89)	(.30)
ΔRTA_i			.01		
			(.03)		
Home country-industry FE	Yes	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes	Yes
Destination market controls ($\tilde{\varepsilon}$)	Yes	Yes	Yes	Yes	No
Destination country trends	No	No	No	No	Yes
Number of firms	28,678	55,094	57,152	2,682	58,785

Note: Standard errors clustered by home country-industry in parentheses. Firm controls are pre-sample firm characteristics: the home weight, ω_{iH}, the number of countries the firm is patenting in during the pre-period, $n_{i,Pre}$, and log knowledge stock in 1985, $\ln K_{i,Pre}$. ^{a} p< 0.01, ^{b} p< 0.05, ^{c} p< 0.1.

6 The Mechanism

According to our model, tariff cuts in foreign markets matter for innovation becaue it raises exports. Before concluding, we want toshed some light on the economic mechanism behind our results. We therefore provide evidence that the tariff cuts of the 1990s indeed increased trade.

To investigate this question, we cannot apply the methodology developed above because our firm-level international dataset does not include trade. Therefore, this part of the paper uses variation across products and countries instead. Specifically, we test whether export-country-product combinations more exposed to tariff cuts in foreign markets have higher export growth than less exposed export-country-product combinations. We use applied MFN tariff cuts at the HS 6-digit level, as described in Section 4.2. The weighted average tariff cut is

$$\Delta \bar{T}_{jm} = \sum_{n} \omega_{jmn} \Delta T_{jnt},$$

where ω_{jmn} is the initial export share, i.e. trade from m to n of HS product j relative to total exports from m of product j, $Exports_{jmn} / \sum_o Exports_{jmo}$. Bilateral product-level trade data are gathered from BACI. Unfortunately, BACI covers the period from 1995 and forward. For this part of the paper, we therefore use the 8-year period 1995 to 2003 and construct the weights ω_{jmn} for the first year available, 1995.[26] We then estimate the regression:

$$\Delta \ln Exports_{jm} = \eta + \beta \Delta \bar{T}_{jm} + \epsilon_{jm}, \tag{13}$$

where $\Delta \ln Exports_{jm}$ is the change in log export value for HS product j and country m between 1995 and 2003 and where $\Delta \bar{T}_{jm}$ is defined over the same 8-year period.

The results are reported in Table 6. Columns (1) and (2) show results estimated by ordinary least squares, with product fixed effects in the first column and both product and country (exporter) fixed effects in the second column. These fixed effects control for different trends in exports across products and countries that are potentially correlated with $\Delta \bar{T}_{jm}$. Column (3) reports results when estimating equation (13) in levels instead of in differences and include product-country pair fixed effects. As expected, across all specifications there is a negative relationship between the two; i.e. larger tariff cuts in export markets (smaller $\Delta \bar{T}_{jm}$) lead to a greater increase in exports. According to our estimates, a 10 percentage point average tariff cut leads to roughly 5 percent increased exports.

[26] Original data are provided by the United Nations Statistical Division (COMTRADE database), see http://www.cepii.fr/cepii/en/bdd_modele/presentation.asp?id=1. As in the main part of the paper, EU-15 is aggregated to a single country.

Table 6: Tariff Cuts and Exports.

	(1)	(2)	(3)
Tariff cut ΔT_{jn}	-.49^{a}	-.36^{a}	-.98^{a}
	(.11)	(.11)	(.32)
Product (HS 6-digit) FE	Yes	Yes	No
Country FE	No	Yes	No
Product-country FE	No	No	Yes
Number of obs	99,892	99,892	926,416

Note: The dependent variable is the 1995-2003 change $\Delta \ln Exports_{jm}$ in columns (1) and (2), and the level $\ln Exports_{jm}$ in column (3). Robust standard errors clustered by product in parentheses. ^{a} p< 0.01, ^{b} p< 0.05, ^{c} p< 0.1.

7 Conclusions

We set out to analyze the impact of improved market access on firms' innovation using the global decline in tariffs during the 1990s. This is a question that so far has not been the subject of rigorous analysis despite its relevance. Our results show that the Great Liberalization of the 1990s had a large positive net impact on innovation. Our results indicate that a 1 percentage point tariff cut in export markets leads to approximately 2 percent growth in firms' knowledge stock, suggesting that trade policy was an important factor driving global innovation in the 1990s. Our findings underscore the importance of trade liberalization for firms' long term performance and for aggregate economic growth and it points to large dynamic gains from trade; gains that are typically not observed and therefore neglected in empirical analyses.

Our estimates are robust to a set of econometric issues, and in particular we provide evidence in support of patents being a useful measure of innovation. While the results are directly relevant for the analysis of trade policy, they also add to the broader literature on economic factors that govern innovation and growth..

References

Abramovsky, L., R. Griffith, G. Macartney, and H. Miller (2008). The location of innovative activity in europe. Working Paper 08/10, Institute for fiscal Studies.

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., A. Bergeaud, M. Lequien, and M. Melitz (2017). The impact of exports on innovation: Theory and evidence. Technical report, Harvard University, Mimeo.

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. V. Reenen (2016). Carbon taxes, path dependency and directed technical change: Evidence from the auto industry. *Journal of Political Economy 214*(1).

Aghion, P., C. Harris, and J. Vickers (1997). Competition and growth with step-by-step innovation: An example. *European Economic Review 41*(3–5), 771 – 782.

Altomonte, C. and T. Aquilante (2012). The EU-EFIGE/Bruegel-Unicredit dataset. (13).

Amiti, M. and J. Konings (2007). Trade liberalization, intermediate inputs, and productivity: Evidence from Indonesia. *American Economic Review 97*(5), 1611–1638.

Autor, D., D. Dorn, G. H. Hanson, G. Pisano, and P. Shu (2016). Foreign competition and domestic innovation: Evidence from u.s. patents. *MIT Working Papers*.

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.

Bloom, N. and J. V. Reenen (2002). Patents, real options and firm performance. *The Economic Journal 112*(478), C97–C116.

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), 3704–3739.

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 (2016). Environmental policy and directed technological change: Evidence from the european carbon market. *Review of Economics and statistics 98*(1).

Dekle, R., J. Eaton, and S. Kortum (2008). Global rebalancing with gravity: Measuring the burden of adjustment. *IMF Economic Review 55*(3), 511–540.

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.

Goldberg, P. K. and N. Pavcnik (2007). Distributional effects of globalization in developing countries. *Journal of economic Literature 45*(1).

Gopinath, G. and B. Neiman (2013). Trade adjustment and productivity in large crises. *American Economic Review 104*(3), 793–831.

Griffith, R., R. Harrison, and H. Simpson (2010). Product market reform and innovation in the EU. *Scandinavian Journal of Economics 112*(2), 389–415.

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.

Head, K. and T. Mayer (2013). *Handbook of International Economics*, Volume 4, Chapter Gravity Equations: Toolkit, Cookbook, Workhorse. Elsevier.

Head, K., T. Mayer, and J. Ries (2010). The erosion of colonial trade linkages after independence. *Journal of International Economics 81*(1), 1–14.

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.

Lerner, J. (1994). The importance of patent scope: An empirical analysis. *RAND Journal of Economics 25*, 319–333.

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.

Loecker, J. D., P. K. Goldberg, P. A. Khandelwal, and N. Pavcnik (2016). Prices, markups and trade reform. *Econometrica 84*(2), 445–510.

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.

Nagaoka, S., K. Motohashi, and A. Goto (2010). Patent statistics as an innovation indicator. *Handbook of the Economics of Innovation 2*, 1083–1127.

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.

Schmookler, J. (1966). *Invention and Economic Growth.* Harvard University Press.

Steinwender, C. (2015). The roles of import competition and export opportunities for technical change.

Topalova, P. (2010). Factor immobility and regional impacts of trade liberalization: Evidence on poverty from india. *American Economic Journal: Applied Economics 2*(4).

Appendix

A The Optimal Knowledge Stock

The firm then chooses the optimal K_i that maximizes global net profits, $\Pi_i - c(K_i)$, i.e.

$$\max_{z_i} \left\{ \sum_n \left[\left(\frac{\xi K_i}{\tau_{jmn}} \right)^{\sigma-1} B_{in} \right] - \psi K_i^k \right\}$$

The first order condition is

$$\begin{aligned} (\sigma-1) K^{\sigma-2} \sum_n \left[\left(\frac{\xi}{\tau_{jmn}} \right)^{\sigma-1} B_{in} \right] &= k\psi K_i^{k-1} \\ K_i &= \kappa \left(\sum_n \tau_{jmn}^{1-\sigma} B_{in} \right)^{1/[k-(\sigma-1)]}, \end{aligned}$$

where $\kappa \equiv \left[\xi^{\sigma-1} (\sigma-1) / (k\psi)\right]^{1/[k-(\sigma-1)]}$.

The second order condition is, inserting the expression for the optimal K_i,

$$\begin{aligned} (\sigma-1)(\sigma-2) K_i^{\sigma-3} \sum_n \left[\left(\frac{\xi}{\tau_{jmn}} \right)^{\sigma-1} B_{in} \right] - k(k-1)\psi K_i^{k-2} &< 0 \\ \kappa^{k-(\sigma-1)} (\sigma-2) K_i^{\sigma-3} \sum_n \left[\tau_{jmn}^{1-\sigma} B_{in} \right] - (k-1) K_i^{k-2} &< 0 \\ (\sigma-2) K_i^{\sigma-3} K_i^{k-(\sigma-1)} - (k-1) K_i^{k-2} &< 0 \\ (\sigma-1-k) K_i^{k-2} &< 0 \end{aligned}$$

which holds given that $k - (\sigma-1) > 0$.

B Fixed Exporting Costs

This section develops an extension of the benchmark model with fixed costs of exporting. In order to serve a market, a firm must incur a fixed cost $f(n)$ to market n. Net profits selling to n are then $\pi_{im}(n) = (z/\tau_{jm}(n))^{\sigma-1} B_i(n) - f(n)$. For analytical convenience, consider the case with a unit continuum of countries. Without loss of generality, countries are sorted according to their profitability, $\pi_{im}(n)$, from high to low.

The firm faces two choices: first, how much to innovate and second, where to sell. We start with the second problem. Because of the presence of market-specific fixed costs, firms will only export to countries that give them positive net profits, $\pi_{im}(n) > 0$. Label the destination with zero profits $\bar{n}_i$, i.e. $\pi_i(\bar{n}_i) = 0$. Global profits are then

$$\Pi_i = \int_0^{\bar{n}_i} \left[\left(\frac{z}{\tau_{jm}(n)} \right)^{\sigma-1} B_i(n) - f(n) \right] dn.$$

We now turn to the problem of how much to innovate. Maximizing $\Pi_i - c(K_i)$ and using Leibniz' integral rule yields

$$K_i = \kappa \left(\int_0^{\bar{n}_i} \tau_{jm}(n)^{1-\sigma} B_i(n)\, dn \right)^{1/[k-(\sigma-1)]}.$$

In changes, we obtain

$$\hat{K}_i = \left[\int_0^{\bar{n}_i{}'} \omega_{ij}(n)\, \hat{\tau}_{jm}(n)^{1-\sigma} \hat{B}_i(n)\, dn \right]^{1/[k-(\sigma-1)]},$$

where $\bar{n}'$ is the marginal destination country in the counterfactual equilibrium and $\omega_{ij}(n)$ is the *gross* profit shares in the initial equilibrium,

$$\omega_{ij}(n) = \frac{\tau_{jm}(n)^{1-\sigma} B_i(n)}{\int_0^{\bar{n}} \tau_{jm}(o)^{1-\sigma} B_i(o)\, do}.$$

We observe that when $\bar{n}_i' = \bar{n}_i$, we get a similar expression as equation (2) in the main text.

C Approximation of the Knowledge Production Function

The expression $\hat{K}_i = \left(\sum_n \omega_{in} \hat{B}_{in} \hat{\tau}_{jmn}^{1-\sigma} \right)^{1/[k-(\sigma-1)]}$ can be approximated by equation (3) in the main text, $\Delta \ln K_i = \sum_{n \in \Omega_i} \beta_n \omega_{in} \Delta T_n + \sum_{n \in \Omega_i} \omega_{in} \Delta \ln e_{in}$.

Proof. The term

$$\begin{aligned} \sum_n \omega_{in} \hat{B}_{in} \hat{\tau}_{jmn}^{1-\sigma} &= \sum_n \omega_{in} e^{(1-\sigma)\Delta \ln \tau_{jmn} + \Delta \ln B_{in}} \\ &\approx \sum_n \omega_{in} \left(1 + (1-\sigma)\Delta \ln \tau_{jmn} + \Delta \ln B_{in}\right) \\ &= 1 + \sum_n \omega_{in} \left((1-\sigma)\Delta \ln \tau_{jmn} + \Delta \ln B_{in}\right), \end{aligned}$$

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

$$\begin{aligned}
\Delta \ln K_i &= \frac{1}{k-(\sigma-1)} \ln\left[1+\sum_n \omega_{in}\left((1-\sigma)\Delta\ln\tau_{jmn}+\Delta\ln B_{in}\right)\right] \\
&\approx \frac{1}{k-(\sigma-1)} \sum_n \omega_{in}\left((1-\sigma)\Delta\ln\tau_{jmn}+\Delta\ln B_{in}\right) \\
&= \frac{1}{k-(\sigma-1)} \left(\sum_n (1-\sigma)\,\omega_{in}\Delta T_{jmn}+\sum_n \omega_{in}\Delta\ln B_{in}\right),
\end{aligned}$$

where we used $\Delta\ln\tau_{jmn} = \Delta\ln(1+T_{jmn}) \approx \Delta T_{jmn}$ for T_{jmn} close to 0. □

D PATSTAT

We use patents from PATSTAT to measure a firm's knowledge stock. To construct our data set we need to deal with a set of issues:

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 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.[27] These harmonized names have been included in PATSTAT TLS906_PERSON table since October 2011. We use the variable "HRM_L2_ID" from this table.

Patent families. To construct the knowledge stock variable 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.[28]An example can be helpful to clarify

[27]For 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).

[28]The 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).

the main idea 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 the same patent to 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.[29]

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. Firms without information on home country are dropped in what we refer to as the sample. Notice that a firm may be associated with more than one country. We have 42574 of such cases. When this is the case, we let home country be the one with the highest frequency in the data. We consider each applicant's home country as its headquarter country.

EU and the Single market. The Single market was established in 1992. To account for this we set tariffs between EU 15 members to zero.

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

[29]See also Dernis and Khan (2004) and Martinez (2010) for an overview of different types of patent families and how they are constructed.

E 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.

Details on construction. First, we convert 6-digit HS codes to 6-digit HS Combined (HSC) nomenclature using a World Bank correspondence table.[30] In some cases, a 6-digit tariff line is missing in year t, but non-missing in previous or later years; in these cases we interpolate to get a non-missing observation in year t. We also extrapolate tariffs in those cases where tariffs 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. Note, that as pointed out above, we control for the establishment of the Single market in 1992 by setting tariffs between EU 15 members, i.e. countries that were part of the EU before the Single market was established or became part of it right after, to zero. Second, we 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.[31] In cases where 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 data set contains data for 96 countries, 128 3-digit industries and 12,249 country-industry combinations. Figure 5 shows average tariffs for high- and low income countries in our final tariff data set.

[30] http://wits.worldbank.org/product_concordance.html

[31] http://unstats.un.org/unsd/cr/registry/regot.asp?Lg=1

Figure 5: Average Tariffs

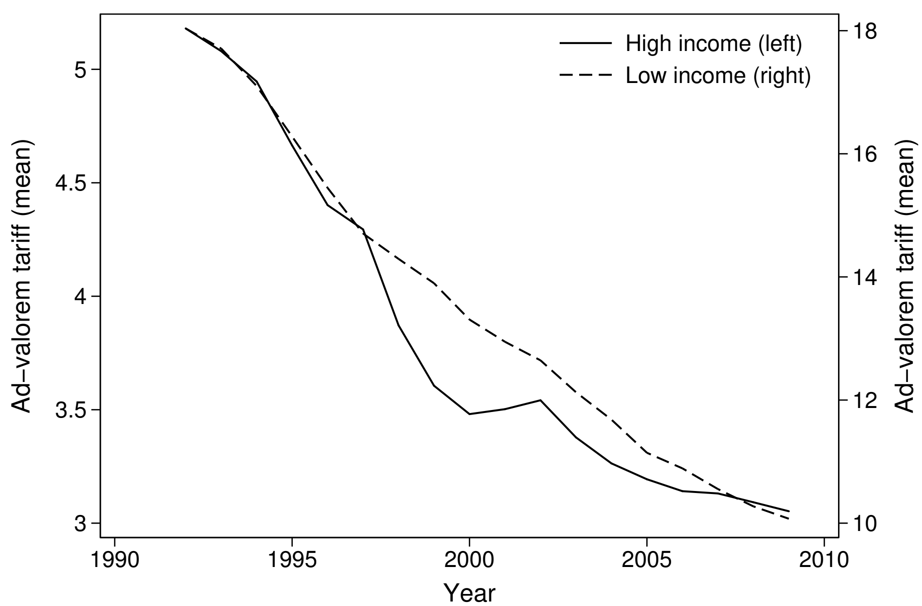

Note: The figure shows average tariffs for high- and low income countries according to the World Bank 1995 definition, using our final tariff data set. Average tariffs are calculated as the simple average across countries. 3-digit NACE tariffs are aggregated to country level tariffs using simple averages.

F 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,[32] the Uruguay Round achieved a global reduction by 40 percent of tariffs and wider market-opening agreements on goods. 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,[33] the Uruguay round achieved important tariff reductions in both developed and developing countries. 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 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 agreed on during the negotiations. This feature of the Marrakesh Agreement implies 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.[34] The first cut was made on the date of entry into force of the WTO agreement, and the following four on 1 January of each subsequent year.[35] 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 and average of 3.8%.

In addition to tariff cuts, the number of "bound" tariffs[36] increased significantly, from 78% to 99% in developed countries, from 21% to 73% in developing countries, and from 73% to 98% in transition economies.

[32] https://www.wto.org/english/docs_e/legal_e/marrakesh_decl_e.pdf

[33] Exceptions are represented by the East Asian NICs.

[34] Unless it is otherwise stated in a Member's Schedule.

[35] See Marrakesh Protocol to the General Agreement on Tariffs and Trade 1994 for additional information.

[36] Bound 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.

G Patent and Sales Weights

This section provides empirical evidence that trade and patent weights are highly correlated.

Aggregate Evidence. We aggregate the patent data to the country-pair level, where the source country is the location of the applicant firm and the destination country is the location of the patent office. 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,

$$\chi_{rst} = \frac{\text{Patents from } r \text{ to } s \text{ at time } t}{\sum_{k \neq s} \text{Patents from } k \text{ to } s \text{ at time } t} \tag{14}$$

Similarly, by using trade data from CEPII, we calculate the import share ψ_{rst} as the share of trade from r to s relative to s' total imports,

$$\psi_{rst} = \frac{\text{Import from } r \text{ to } s \text{ at time } t}{\sum_{k \neq s} \text{Imports from } k \text{ to } s \text{ at time } t} \tag{15}$$

Figure 6 shows the import and patent inflow 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 import side are also the top three countries on the patent side. In Figure 7 we plot all country pairs in our sample for the year 2000. We see that there is a strong log linear relationship between bilateral patenting and trade, with a linear regression slope of 0.80 (s.e. 0.02). Finally, we show that the patent flows adhere to a gravity model. Table 7 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 sample, 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.

Firm-Level Evidence. We use survey data for European firms from EU-EFIGE/Bruegel-UniCredit data set (henceforth EFIGE) to calculate firm specific export shares to different country groups, and compare them to patent weights from PATSTAT.[37] The EFIGE database consists of a representative sample of about 15,000 manufacturing firms (above 10 employees) across seven countries (Germany, France, Italy, Spain, United Kingdom, Austria, Hungary), and provides information on firms' international activities. We use firms' self-reported export shares for 2008 and for each firm we construct weights for market exposure based on the share of sales to eight groups of countries:[38] EU 15 countries, other EU countries, other European

[37] The EFIGE data set is described in Altomonte and Aquilante (2012).

[38] Specifically, we use the answers to two questions. D4 asks: "Which percentage of your 2008 annual

Figure 6: Import and Patents Shares.

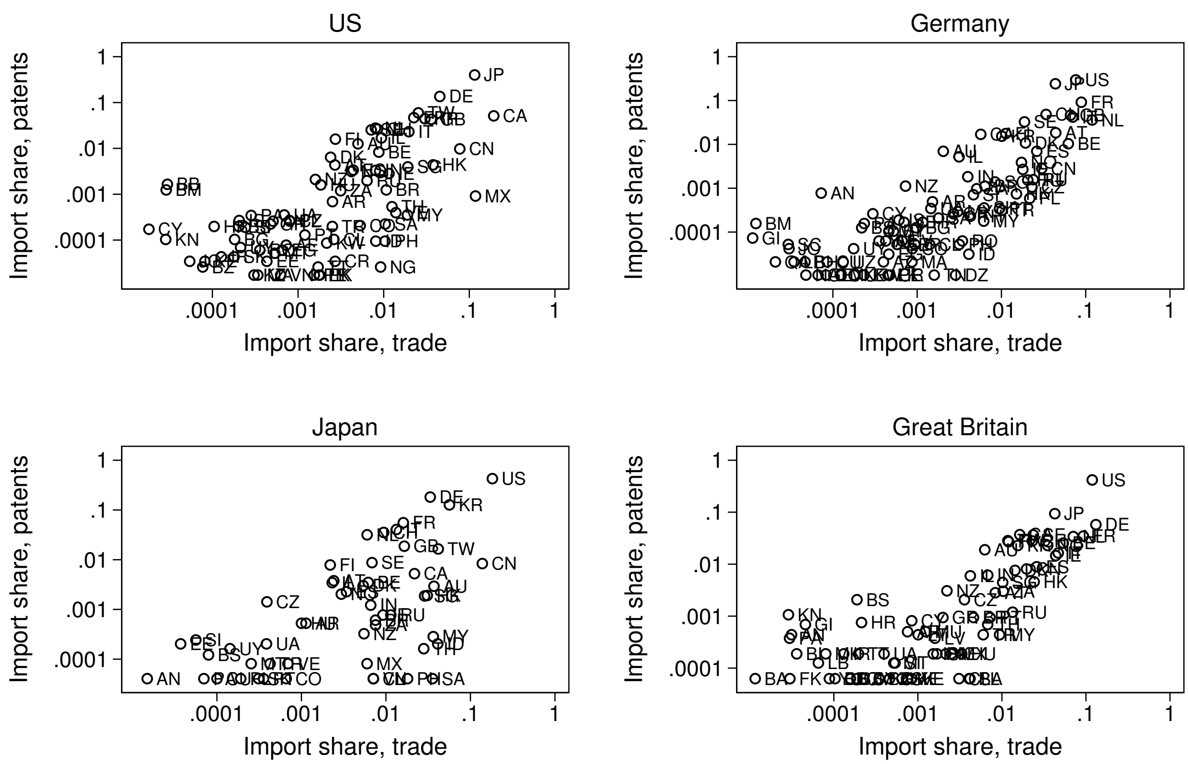

Note: The vertical axis shows the share of patents filed in U.S./Germany/Japan/Great Britain belonging to firms headquartered in source country r (log scales). The horizontal axis shows the share of total imports in U.S./Germany/Japan/Great Britain coming from source country r (log scales). Year 2000.

Figure 7: Bilateral Trade and Patenting.

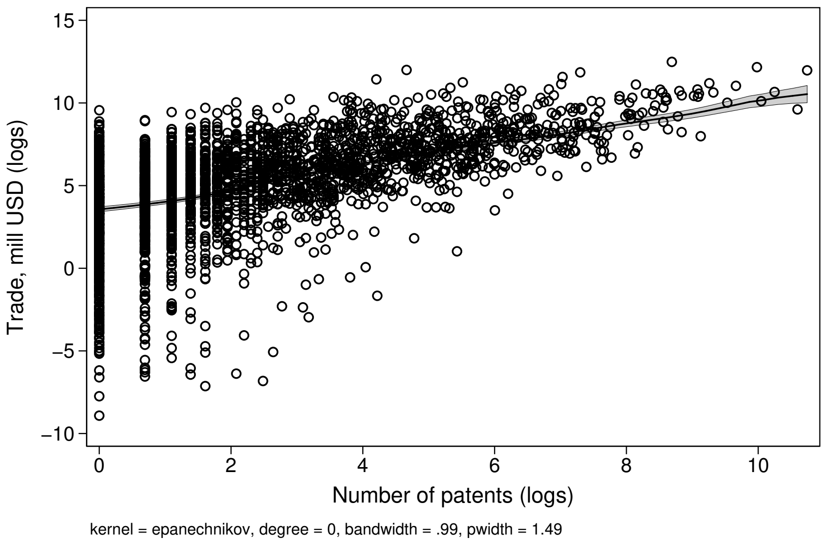

Note: The figure shows the number of patents and total trade from headquarters country r to destination country s in year 2000 (both in logs). The solid line is the local polynomial regression fit and the gray area represents the 95% confidence bands. The linear regression slope is 0.80 (s.e. 0.02).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

Table 7: Patent Flows and Gravity.

Dep. variable: $\ln Patents_{rst}$	Year 2000	1965-2006(
	(1)	(2)
Distance$_{rs}$	-.44^a	
	(.03)	
GDP$_r$	.68^a	.48^a
	(.02)	(.05)
GDP$_s$	.50^a	.27^a
	(.02)	(.04)
Year FE	No	Yes
Source-destination FE	No	Yes
R^2	0.43	0.34
Number of observations	2,558	68,447

Note: Robust standard errors in parentheses. ^{a} p< 0.01, ^{b} p< 0.05, ^{c} p< 0.1.

countries not EU, China and India, other Asian countries, USA and Canada, Central and South America, and a residual category including all remaining countries.[39]

We match the EFIGE data with firm level data from Amadeus, which in turn can be matched with PATSTAT using the patent application number of each patent owned.[40] We calculate weights for market exposure based on firms' patenting activity abroad that correspond to those we have calculated for exports using patent applications for the period of 1998 to 2008. [41] Figure 8 shows a kernel-weighted local polynomial regression of patent shares on export shares for firms with at least one patent. Again we observe that there is a strong relationship between patent and trade weights. The corresponding linear regression slope is 0.89 (s.e. 0.008).

H 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, $t = 1$: 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 same country in $t = 2$. We limit the sample to firms that filed at least one patent after 2004, which ensures that all firms exist throughout the the period in question. Table 8 reports the results. Even after 20 years, the likelihood of continuing to patent is high (44 percent). The same is true on the entry side: conditional on patenting in a market in $t = 2$, the likelihood of patenting 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

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 EU 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.

[39]The 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.

[40]Specifically, from the variable patent application number in Amadeus we are able to construct the appln_nr_epodoc in PATSTAT, and to link each patent application in Amadeus to the same patent application in PATSTAT.

[41]When 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. If a firm does not have patents, then all its weights for all groups of countries are set to zero.

that the unconditional probability is roughly 4 percent. Second, we calculate the correlation in weights conditional on patenting in that market in both t and $t+1$ (i.e., the intensive margin). Figure 9 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 between the weights.

Table 8: Persistence in Patent Weights. Extensive Margin.

	(1) $t=0$	(2) $t=1$	(3) $t=2$
Conditional Probability of continuing $P\left[p_{int} \mid p_{in0}\right]$	1	0.44	0.44
	-	(.001)	(.001)
Conditional Probability of entry $P\left[p_{int} \mid p_{in2}\right]$	0.37	0.39	1
	(.001)	(.001)	-
Unconditional Probability of patenting $P\left[p_{int}\right]$	.037	.033	.035

Note: Standard errors in parentheses.$P\left[p_{int} \mid p_{in0}\right]$) depicts the share of firm-destinations with positive patenting in $t=0$ and period t relative to all firms-destinations with positive patenting in $t=0$. $P\left[p_{int} \mid p_{in2}\right]$) depicts the share of firms-destinations with positive patenting in $t=0$ and period t relative to all firms-destinations with positive patenting in $t=2$.

Figure 9: Persistence in Patent Weights. Intensive Margin.

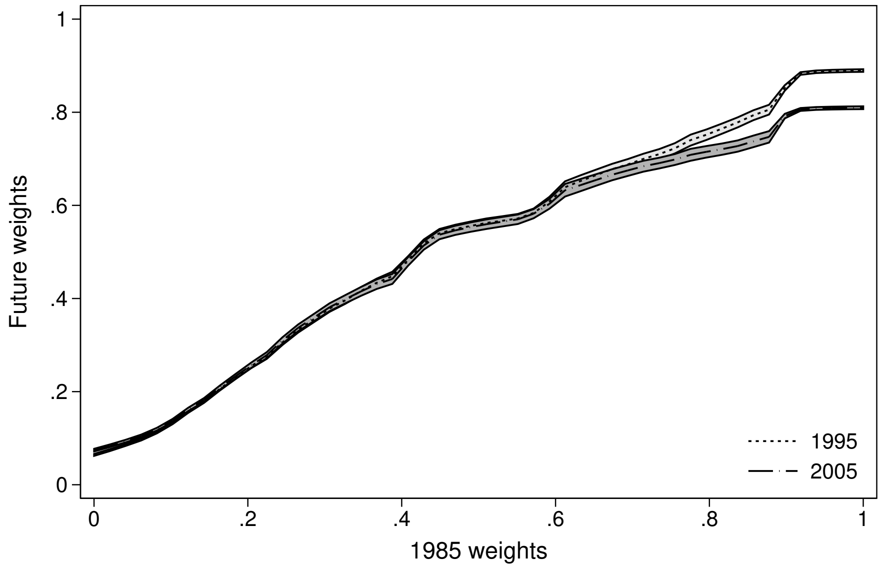

Note: The figure shows the kernel-weighted local polynomial regression of weights ω_{int} in 1995 or 2005 (vertical axis) on weights in 1985 (horizontal axis). The two lines represent two separate regressions. Gray areas denote the 95 percent confidence bands. The sample includes all pairs $(\omega_{int}, \omega_{in,t+1})$ where both values are non-zero. The population of firms is described in Section 4.3.

I Patents as a Measure of Innovation

In this Section we examine the robustness of patenting as an indicator of innovative activity by looking at the correlation between patent applications and R&D expenditures. We rely on the EFIGE survey data referred to above and match these with Amadeus and PATSTAT. This leaves us with a sample of European manufacturing firms. EFIGE contains information of firms' average investment in R&D activities as percentage of turnover for the period 2007-2009.[42] Using turnover data from Amadeus we are able to calculate average R&D expenditures for the same period.

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. 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-

[42] Calculation is based on the question C21 in EFIGE that asks: "Which percentage of the total turnover has the firm invested in R&D on average in the last three years (2007-2009)?"

Table 9: R&D expenditures and patenting

Dep. variable: Patenting	R&D expenditure	Log R&D expenditures
	(1)	(2)
Patenting	3570.16[a]	1.28[a]
	(584.17)	(0.06)
Observations	6204	6074

Note: Standard errors in parentheses. [a] $p < 0.01$, [b] $p < 0.05$, [c] $p < 0.1$.The table shows a regression of R&D expenditures on a binary variable indicating whether the firm has any patent. The population of firms is all firms in PATSTAT that can be matched with EFIGE.

2009) and the three subsequent years, until 2012. On the intensive margin, higher R&D expenditures are strongly correlated with a higher number of patent applications. Figure 10 shows a kernel-weighted local polynomial regression of firms' R&D expenditures on number of patent applications. 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 firms with 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, we find that 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 period 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 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. 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 9 respectively. In both cases we find a positive and strong correlation between R&D expenditures and patent applications.

J Additional Tables

Table 10: Patents characteristics

	Filed patents	Granted patents	Proportion granted	Citations (mean)	Inventors (mean)	IPC codes (mean)
World Total	1767861	963471	0.55	2.07	2.12	3.08
United States	258299	210294	0.81	5.00	2.03	3.43
Japan	949400	375702	0.40	1.13	2.61	3.50
Germany	111102	63292	0.57	1.54	2.06	3.08
Great Britain	34381	13910	0.41	1.66	1.82	3.21
France	42595	31737	0.75	1.71	2.00	3.27
China	20427	16048	0.79	0.12	1.55	1.49
Italy	33194	25461	0.77	0.84	1.56	2.13
Canada	14709	9947	0.68	2.96	2.08	3.25
Mexico	294	82	0.28	0.77	1.85	1.78
Brazil	4568	421	0.09	0.14	1.32	1.42

The table shows the number of filed and granted patents and the average quality of patents as proxied by the number of citations, inventors and IPC class codes in the final sample of firms for the period 1992-2000. The first row displays overall statistics, the remaining part of the table shows statistics for the ten biggest economies (nominal GDP) in 2000.

Table 11: Number of patents by industry: Final sample

Nace	World	Top 10 economies in 2000									
		United States	Japan	Germany	Great Britain	France	China	Italy	Canada	Mexico	Brazil
10	22267	4469	12686	718	353	457	45	431	311	6	21
11	1262	182	618	116	51	74	52	26	2		4
12	1127	262	202	59	18	39	122	48	3	2	16
13	5525	1055	3569	213	65	136	2	114	50		8
14	1935	210	619	80	85	118	183	166	16		19
15	2657	615	797	203	114	103	72	257	26		53
16	940	110	399	236	13	4	2	12	28		1
17	6180	253	4646	171	78	102	8	47	87		19
18	3342	400	2286	140	172	62	19	54	18		8
19	1892	259	775	89	116	129	4	3	137		10
20	232881	24106	163583	6905	3425	4043	7628	1700	1620	90	238
21	81016	34360	19260	5648	3279	2643	127	1864	1180	22	21
22	51114	4688	34138	2967	1597	1218	123	1093	362	9	213
23	38262	2146	31257	890	665	531	90	537	25	9	97
24	51482	1397	34138	858	253	418	708	143	475	63	417
25	31787	4650	10975	6778	1553	1658	276	1287	272	6	135
26	426945	86848	189669	21225	5758	10006	4225	7293	6066	21	262
27	105676	12014	62671	8712	1871	3625	1664	2632	354	9	605
28	375261	37983	236280	24437	6163	8100	2806	6828	1776	27	1278
29	140532	10128	63984	11404	1514	3061	439	2118	538	6	295
30	14080	2341	3943	1503	678	803	367	693	229	1	55
31	11749	1731	4402	1611	404	406	136	689	72		104
32	119021	26434	55573	9928	5134	3362	1101	2867	866	22	457
62	176	70	101			3					
99	9833	658	209	1474	597	116	232	2103	65	2	191

The table shows the number of patents filed during the period 1992-2000 in the world and in the 10 biggest economies in terms of nominal GDP in 2000. The population consists of the final sample of firms described in the Section 4.3. Patent count is based on filed patent applications. Patents are only counted once in each industry and country: in Column 3 (World), patents owned by more than one firm and/or filed in more than one country are not counted twice; in Column 3-12 (Top 10 economies in 2000), patents with multiple ownership in the same country are only counted once, but patents owned by firms in more than one country and/or filed in more than one country, are counted once in each destination country. Nace Rev. 2 codes: 10 Manufacture of food products; 11 Manufacture of beverages; 12 Manufacture of tobacco products; 13 Manufacture of textiles; 14 Manufacture of wearing apparel; 15 Manufacture of leather and related products; 16 Manufacture of wood and of products of wood and cork, except furniture; manufacture of articles of straw and plaiting materials; 17 Manufacture of paper and paper products; 18 Printing and reproduction of recorded media; 19 Manufacture of coke and refined petroleum products; 20 Manufacture of chemicals and chemical products; 21 Manufacture of basic pharmaceutical products and pharmaceutical preparations; 22 Manufacture of rubber and plastic products; 23 Manufacture of other non-metallic mineral products; 24 Manufacture of basic metals; 25 Manufacture of fabricated metal products, except machinery and equipment; 26 Manufacture of computer, electronic and optical products; 27 Manufacture of electrical equipment; 28 Manufacture of machinery and equipment n.e.c.; 29 Manufacture of motor vehicles, trailers and semi-trailers; 30 Manufacture of other transport equipment; 31 Manufacture of furniture; 32 Other manufacturing; 62 Computer programming, consultancy and related activities; 99 Missing in data.

Table 12: Patenting firms by industry: Final sample

Nace	World	United States	Japan	Germany	Great Britain	France	China	Italy	Canada	Mexico	Brazil
		Top 10 economies (nominal GDP) in 2000									
10	2586	374	1127	180	67	129	17	136	39	3	18
11	232	23	57	35	16	27	5	15	3		3
12	256	5	65	31	4	15	13	28	2	1	9
13	675	100	205	67	39	57	4	47	6		5
14	589	128	110	38	36	47	15	48	16		9
15	731	160	75	77	33	53	15	114	14		17
16	197	25	52	48	7	6	2	8	5		1
17	552	96	85	86	25	50	5	15	19		16
18	488	106	128	52	36	39	7	13	9	1	5
19	365	110	99	16	13	15	2	4	25		6
20	8751	1987	2326	856	377	455	127	348	174	9	94
21	4798	1325	1094	352	205	358	25	257	99	1	12
22	4719	1110	711	758	255	333	29	371	102	5	93
23	2299	383	634	283	99	152	24	163	25	5	25
24	1675	309	409	197	67	74	48	63	64	3	27
25	5091	1222	498	752	300	364	71	463	106	5	82
26	16197	5084	3241	1701	880	1074	318	698	325	8	143
27	10156	2356	1411	1337	466	714	278	871	201	5	206
28	34809	8227	6583	4722	1737	2157	508	2353	706	28	432
29	6495	1580	973	990	388	477	78	579	119	3	129
30	3363	916	368	451	195	353	74	282	81	2	37
31	2503	570	138	349	132	202	39	281	47		56
32	18162	5447	2438	2038	890	1289	324	1110	378	15	277
62	35	18	12			4					
99	4186	333	62	741	146	45	81	821	45	2	104

The table shows the number of firms patenting in the world and in the 10 biggest economies in terms of nominal GDP in 2000. The population consists of the final sample of firms described in the Section 4.3. Nace Rev. 2 codes: 10 Manufacture of food products; 11 Manufacture of beverages; 12 Manufacture of tobacco products; 13 Manufacture of textiles; 14 Manufacture of wearing apparel; 15 Manufacture of leather and related products; 16 Manufacture of wood and of products of wood and cork, except furniture; manufacture of articles of straw and plaiting materials; 17 Manufacture of paper and paper products; 18 Printing and reproduction of recorded media; 19 Manufacture of coke and refined petroleum products; 20 Manufacture of chemicals and chemical products; 21 Manufacture of basic pharmaceutical products and pharmaceutical preparations; 22 Manufacture of rubber and plastic products; 23 Manufacture of other non-metallic mineral products; 24 Manufacture of basic metals; 25 Manufacture of fabricated metal products, except machinery and equipment; 26 Manufacture of computer, electronic and optical products; 27 Manufacture of electrical equipment; 28 Manufacture of machinery and equipment n.e.c.; 29 Manufacture of motor vehicles, trailers and semi-trailers; 30 Manufacture of other transport equipment; 31 Manufacture of furniture; 32 Other manufacturing; 62 Computer programming, consultancy and related activities; 99 Missing in data.

Figure 8: Market Exposure Weights - Export and Patents.

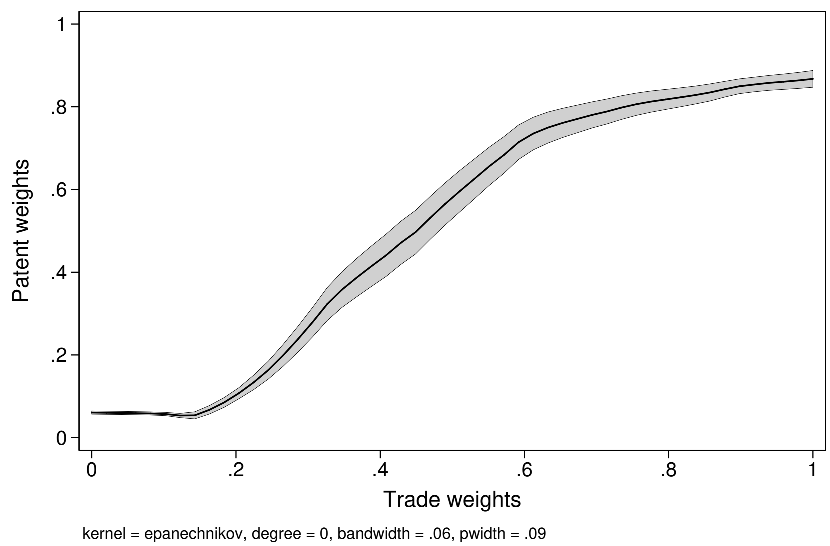

Note: The figure shows market exposure weights based on sales (2008) and patenting activity (1998 to 2008). The solid line is the local polynomial regression fit and the gray area represents the 95% confidence bands. The linear regression slope is 0.89 (s.e. 0.008). The population of firms is all firms in PATSTAT that can be matched with EFIGE.

Figure 10: R&D expenditures and patenting: Intensive margin

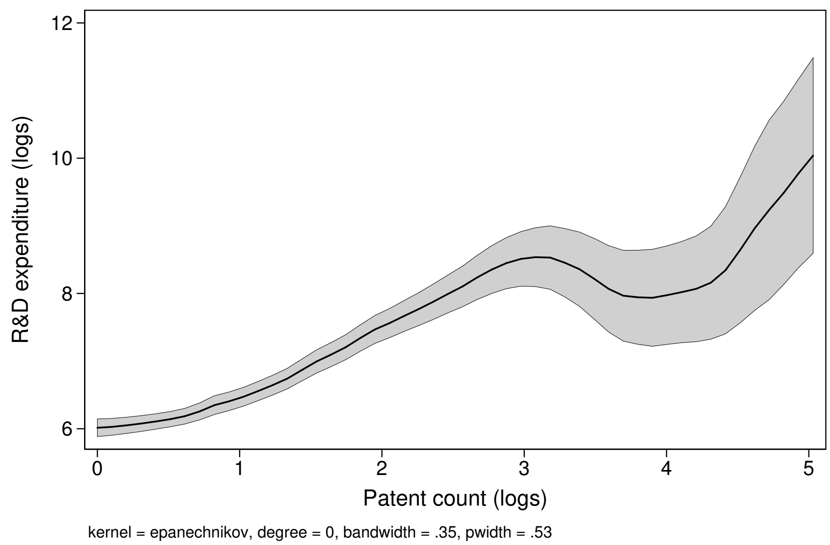

Note: The figure shows the average number of patent applications per year and average R&D expenditures per year (both in logs). R&D expenditures refer to the period 2007-2009, patent counts are calculated over a six year window, from 2007 to 2012. The solid line is the local polynomial regression fit and the gray area represents the 95% confidence bands. The linear regression slope is 0.68 (s.e. 0.05). The population of firms is all firms in PATSTAT that can be matched with EFIGE.

Figure 11: R&D Expenditures and Patenting

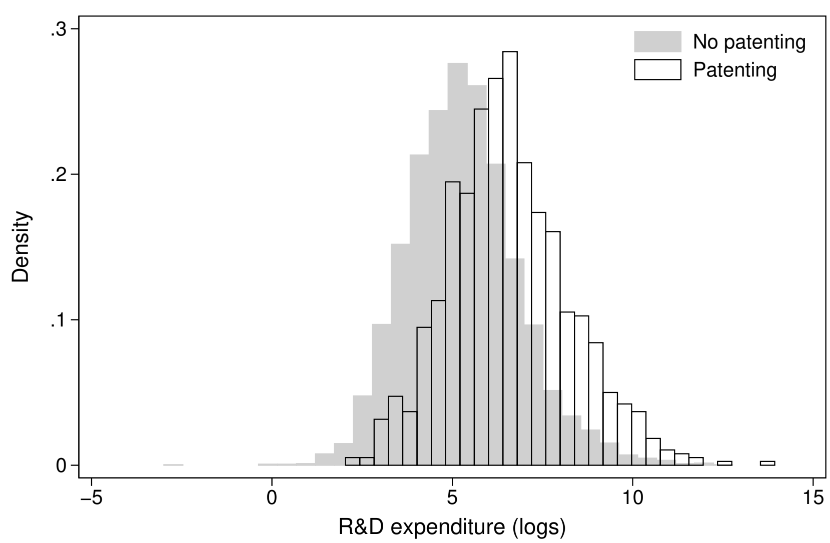

Note: The figure shows the distribution of firms' R&D expenditures (in logs) for firms with (white) and without (gray) patent applications in the period 2007-2009. The population of firms is all firms in PATSTAT that can be matched with EFIGE.