

International Patent Protection and Trade: Transaction-Level Evidence*

Gaétan de Rassenfosse¹, Marco Grazzi², Daniele Moschella³, and
Gabriele Pellegrino²

¹College of Management of Technology, Ecole polytechnique fédérale de Lausanne,
Switzerland.

²Department of Economic Policy, Università Cattolica del Sacro Cuore, Milano, Italy.

³Institute of Economics & EMbeDS, Scuola Superiore Sant'Anna, Pisa, Italy.

Abstract

We report a hitherto undocumented causal mechanism of how patent protection affects exports. The empirical analysis leverages unique data on the worldwide patenting and exporting activities at the product level for the universe of French firms. Exploiting heterogeneity of patent coverage within firm-product-country destinations, we find evidence of a patent premium. Goods protected by patents in a destination country are associated with higher export quantities, *ceteris paribus*. The effect ranges between four and eleven percent. The causality of the finding is confirmed using rejected patent applications, which are exogenous to the firm. Exports collapse when firms lose patent protection.

Keywords: export, patents, products, intellectual property rights, innovation

JEL classification: D22, F10, F14, O30, O34

**Acknowledgments.* This work has been partly supported by the European Commission under the H2020, GROWINPRO, Grant Agreement 822781. Daniele Moschella received financial support by the Italian Ministry of Education and Research under the PRIN 2017 Programme (Project code 201799ZJSN). This work is also supported by a public grant overseen by the French National Research Agency (ANR) as part of the ‘Investissements d’avenir’ program (reference: ANR-10-EQPX-17, Centre d’accès sécurisé aux données, CASD). This paper has benefited from comments of participants at several seminars and conferences: SETC, Cagliari (September 2019); Università Cattolica, Milano (October 2019); University of Maastricht (December 2019); UNU-MERIT, (June 2020); ETOS (July 2020); ESCoE (September 2020); Henley Business School (November 2020); CAED, Coimbra (November 2021); CONCORDi (November 2021); TPRI’s Brown Bag Seminar (February 2022). We are also indebted to Andrew B. Bernard, Bruno Cassiman, Paola Conconi, Bronwyn H. Hall, Alireza Naghavi, Keith Maskus, Fabio Montobbio and Nikolas Zolas for insightful comments. The editor and two anonymous reviewers provided valuable comments. Declarations of interest: none.

1 Introduction

A wealth of evidence supports the view that intellectual property (IP) rights play a central role in international trade. Contributions have focused on two main perspectives: capability and institutions. The former perspective takes patenting activity as a proxy for superior capabilities, whereas the latter seeks to document a causal effect of patenting on exports.

This paper contributes to the second perspective by documenting a hitherto ignored causal mechanism of patenting on exports. Specifically, a patent gives the right to its owner to exclude others from making, using, and selling the patented technology in the markets covered by the patent. Consequently, markets in which a firm manages to secure patent protection become more attractive to the firm, which, in all logic, should export more to these markets—a phenomenon known as the ‘patent premium’ (Jensen et al., 2011). Although this mechanism is core to the IP system, it could not be studied to date because, without exception, all studies on IP and trade miss data on whether exported goods enjoy patent protection in export markets.

The empirical analysis leverages unique data on French firms’ international patenting and exporting activities at the product-country level and exploits a novel identification strategy.¹ The analysis proceeds in two steps. First, we include a rich set of fixed effects capturing the firm, product, country, and time dimensions (as well as their interactions). We study the extent to which patent protection for a product-country-time cell affects a firm’s exports (total values, quantities, and prices). Second, acknowledging that the filing decision is endogenous to exports (*e.g.*, Zylkin and Brunel, forthcoming), we leverage information on rejected patent applications. We argue that rejection events come largely as a surprise to the firm and exploit such events to strengthen the identification of the effect of patents on trade.

In the first part of the analysis, the baseline specification shows that patent protection in a product-country destination increases total exports in that product-country destination by up to six percent. The regression model simultaneously controls for firm-level export activity in the product-country destination before patent protection, firm-level export activity of the same product in other destination countries, and export activity of the same product by other firms in the same destination country. Regression results suggest that export quantities rather than unit values drive the premium in exports associated with patents. In other words, exporting firms export larger quantities in countries with patent protection, but they do not seem to charge higher prices. We interpret this finding as evidence that firms value foreign patent protection for the legal security that it brings rather than for the possibility of setting monopoly prices. Given the rich set of fixed effects, the finding that patent protection is associated with more exports is quite surprising. Firms usually apply for patents in the largest markets, hoping that large-country protection also offers protection in the smaller, peripheral countries. However, our results suggest that patent

¹ As explained in detail in Section 2, we use the HS6 commodity classification, which relates to product categories. We sometimes use the term ‘product’ for the sake of simplification.

protection brings concrete benefits in the targeted countries. The second part of the analysis exploits rejection events. We observe massive drops in exports, up to 70 percent in value, when such cases happen, confirming the importance of patent protection for trade. The findings are robust to a range of alternative specifications including, but not limited to, changes in sample composition, changes in the set of fixed effects, a focus on product-related patents (as opposed to process patents) to increase the accuracy of the match between products and patents, and controlling for variables of patent importance.

Our analysis relates most closely to the trade literature that studies the importance of IP as an institution. This literature has traditionally shown that strengthening IP rights in a country leads to more exports to that country. It posits that an increase in the strength of IP protection facilitates technological upgrading by local firms (or provides a safer environment for foreigners to deploy technology), which, in turn, is conducive to international trade in technology-intensive goods.² Extant studies predominantly use country or industry-level export data (Maskus and Penubarti, 1995; Smith, 1999, 2001; Ivus, 2010; Palangkaraya et al., 2017), although more recent works have exploited firm-level data. Lin and Lincoln (2017) is the first study to exploit a matched firm-level dataset on exports and patents. They show that patenting firms are more likely to export to countries that strengthen IP protection. We contribute to this literature by investigating another causal mechanism of the effect of IP on trade, namely a direct ‘quantity’ effect of patent protection in an export market.

Our work also relates to the literature that has taken a capability perspective. This literature posits that patents are an inherent expression of superior capabilities that are conducive to better performance in international markets. This perspective is in tune with theoretical models of international trade that predict the self-selection of more productive firms into the export market (see Melitz, 2003; Bernard et al., 2007b, among many others). One of the first empirical works is Soete (1981), which shows a close relationship between technological performance, as proxied by patenting activity, and export performance, as proxied by export market shares, among OECD countries and across several sectors (see also Soete, 1987). A dense stream of follow-on research has investigated the nature of such technological factors both at the country level (Fagerberg, 1988; Amendola et al., 1993) and at the country-industry level (Dosi et al., 1990; Amable and Verspagen, 1995; Landesmann and Pfaffermayr, 1997; Wakelin, 1998b; Carlin et al., 2001; Laursen and Meliciani, 2010). Firm-level evidence usually relies on survey data, such as the Community Innovation Survey in Europe, in works including Wakelin (1998a), Castellani and Zanfei (2007), Cassiman et al. (2010), and Van Beveren and Vandenbussche (2010). A more recent stream of work has exploited administrative data by matching firm-level data on trade to patents including Dosi et al. (2015), Coelli et al. (forthcoming), and Aghion et al. (2018). Our results suggest that, although patents may signal superior capabil-

² This perspective is in line with studies showing that the patent system stimulates economic growth (Gould and Gruben, 1996; Hu and Png, 2013) and provides R&D incentives (Arora et al., 2008).

ities, the owning of patents delivers tangible benefits to firms, which improve their export performance.

The rest of the paper is organized as follows. Section 2 introduces the data, while Section 3 presents the core estimates of the effect of patent protection on exports. Section 4 discusses a series of robustness tests and additional results. Section 5 concludes.

2 Data and descriptive statistics

2.1 Dataset construction

We construct our dataset using two main types of firm-level data: export and patent data.

Export data

Concerning export data, we rely on transaction-level exports recorded by the French customs office (*Direction Générale des Douanes et des Droits Indirects*, DGDDI).³ The dataset contains detailed information on export flows for each year from 2002 to 2011 for all French exporters. To each exporter corresponds a unique official identification number (SIREN code). The dataset includes export value, export quantity, country of destination, and an 8-digit product code following the European Union’s Combined Nomenclature (CN8).

Because the last two digits of this code might change from year to year, we aggregate products at the 6-digit level, which corresponds to the international Harmonized System (HS6) trade classification. The HS classification is the standard for measuring trade flows at the finest level of disaggregation. Note that while the HS unequivocally identifies a given product category, more than one variety might fall within a given HS6 class. Thus, we cannot observe differences in product variety. For simplicity, we will sometimes refer to HS6 as products, but note that we are actually capturing a *category* of products. The HS classification is revised every five years. Therefore, to have a unique product classification scheme throughout the period, we have relied on a concordance table provided by the World Bank. This table maps the 2007 classification to the 2002 classification.⁴ We have used the 2002 HS classification as the main classification—this will prove helpful when linking products to patents, as explained further below.

Table 1 reports descriptive statistics for the 2011 cross-section of exporters. This data slice contains 1,955,535 observations (not reported) corresponding to different transactions at the firm-product-destination level. Column (1) shows that 92,165 firms account for these transactions, and they exported a total of 4,128 product categories (at the 6-digit level) to 230 destination countries. We discuss the rest of the table in the following sections.

³ The data are directly provided to researchers by the DGDDI upon the approval of a research proposal by the *Comité du Secret Statistique*.

⁴ Concordance tables are available at https://wits.worldbank.org/product_concordance.html

Table 1: A slice of the dataset, year 2011

	Exporters	Patentees	Exporters with patents			
			Levels of matching			
	(1)	(2)	Firm (3)	Firm-country (4)	Firm-product (5)	Firm-product-country (6)
#Firms	92,165	8,331	4,251	3,663	3,115	2,596
#Patents	.	81,774	63,401	58,979	53,519	48,954
#CPC	.	628	.	.	547	538
#Products	4,182	.	3,748	3,621	2,878	2,692
#Countries	230	86	222	72	219	72
Obs. with patents(%)	.	.	21.8	13.5	9.89	5.54
Exports with patents(%)	.	.	41.9	33	34.8	27.4

Notes: The table reports basic figures for the dataset of exporters in 2011 (column 1); the ten-year stock of patent applications worldwide (column 2); the matched observations over the two datasets at the firm level (column 3), firm-country level (column 4), firm-product level (column 5), firm-product-country level (column 6).

Patent data

Regarding patent data, we relied on two data sources. We obtained information on patent applications filed by French firms at the French patent office (*Institut National de la Propriété Industrielle*, INPI) and the European Patent Office (EPO). These data were provided to us by INPI and start in 1993, well before the beginning of the export data.⁵ Because patents are jurisdictional rights, patents granted by INPI are valid only in France, whereas patents filed at the EPO are valid potentially throughout Europe.⁶ Information on whether patent protection covers other countries comes from the official register of each national patent office across the world. This information is available in the PATSTAT database, which provides bibliographical and legal status data for patent offices worldwide (de Rassenfossé et al., 2014).⁷ Using the French patent document numbers, we could identify all the patent applications belonging to the same patent family across 90 patent offices.⁸ We have also used the PATSTAT database to obtain bibliographic data about the patents, including the priority year in the patent family and the technological class(es) to which patents pertain. We have used the Cooperative Patent Classification (CPC) codes, a classification of technologies jointly managed by the European Patent Office and the U.S. Patent and Trademark Office.

A key choice in the data collection design is the use of granted patents versus patent applications (which include granted and pending patents). Granted patents correspond to patents issued by the patent office after successfully passing the examination process, whereas pending patents are awaiting examination and, ultimately, issuance (or rejection). We have opted for the use of patent

⁵ Obtaining data from INPI required signing a license agreement to reuse the data. The data are available through FTP access in zipped XML files.

⁶ Once a patent is granted by the EPO, it must be validated in each country where protection is desired. The vast majority of EPO patents are validated in Germany, the United Kingdom, and France.

⁷ The PATSTAT database can be accessed via <https://www.epo.org/searching-for-patents/business/patstat.html>.

⁸ A family of patents is a set of patent applications covering the same invention in different countries (Martínez, 2010).

applications. Although it is only once a patent is granted that firms can seek damages in case of infringement, pending patents already offer some protection. Indeed, any infringement occurring between the patent application’s filing date and the grant date can serve as cause for a legal claim of action. Hence the phrase “patent pending” marked on products or mentioned in advertisements to provide ‘constructive notice’ to potential infringers. Many products are commercialized with a pending patent, and Section 2.3 provides evidence supporting this claim. Nevertheless, in Section 3.3, we will exploit information on refused patent applications to deal with the endogeneity of the patenting decision with respect to the export decision.

Another design choice relates to how long we should assume that the invention underlying a patent generates an economic benefit, the average ‘useful’ life for patent applications. Estimates suggest that the average lifetime of patents is about ten years (*e.g.*, de Rassenfosse and Jaffe, 2018). Accordingly, we consider a ten-year stock of patents for each exporting firm. In concrete terms, we compute the number of patents of firm f for product p in country c in year t as the sum of the number of patents filed from year t to year $t - 9$, using the priority year in the patent family, by that firm in that country for that product.

Column (2) of Table 1 shows that the ten-year stock of patent applications in 2011 consisted of 81,774 patent applications worldwide by 8,331 French firms exporting in that year. These applications cover almost all technological classes (628 CPC codes at the 4-digit level out of 641 total CPC codes) and belong to families that cover 86 different countries.

Linking export data with patent data

Export and patent data have two features in common: the SIREN identification code for firms; and the country identification code for export destinations and patent offices. Merging the datasets at the firm level and the firm-country level is thus a straightforward task. Among patenting exporters, 50 percent of firms export more than eight products, whereas the median number of products among all exporters is two. The difference is not surprising, given that the size distribution of patentees is much more skewed than the general size distribution of exporters. However, this figure suggests that it is worth considering the product dimension when analyzing the relationship between patenting and export performance. We use a concordance table between product codes and patent technological classes to merge the datasets at the product level. We have four levels of analysis: firm, firm-country, firm-product, and firm-product-country. We discuss each level in turn.

First, regarding the firm-level match, we link all patent applications from a firm (no matter in which countries they were filed) to all its transactions. This type of merge is typical in the literature on firm-level analyses of exports and patents (see, among the most recent contributions, Lin and Lincoln, 2017; Aghion et al., 2018; Autor et al., 2020). In our case, this merge leads to about 22 percent of transactions covered by a patent application. These transactions accounted for around 42 percent of total export value in 2011; see column (3) of Table 1. About five percent of firms (4251 out of 92,165) accounted for these exports, and

they jointly owned a ten-year stock of 63,401 patent applications in 2011.⁹

Second, regarding the firm-country match, we link all patent applications from a firm in a specific country to all its export transactions to that country. This match produces the figures reported in Table 1, column (4)—the share of transactions covered by a patent decreases from 22 percent to 13.5 percent. Interestingly, the share of export value accounted for by a patent decreases less than proportionally, going from 42 percent to 33 percent. Lower coverage rates with the geographical matching (72 countries) are due to the fact that firms holding a patent export to a wide range of countries (222) and file for patents in 86 over the 90 countries for which we have patent data.¹⁰ Conversely, firms sometimes file for patents in countries to which they do not export (in 2011)—this fact explains the drop in the number of applications and firms matched (from 63,401 to 58,979 and from 4,251 to 3,663, respectively). Although this may seem surprising, there are four logical explanations for such cases: the firm has not yet exported in the country; it no longer exports there; it temporarily suspended export; or it has never intended to export (but patent protection is needed, *e.g.*, to protect the manufacturing process in that country).

Third, regarding the firm-product match, we have exploited the concordance table developed by Lybbert and Zolas (2014) and Goldschlag et al. (2019). The authors propose a probabilistic crosswalk between the 2002 HS classification at the 6-digit level and CPC codes at the 4-digit level, the so-called ALP crosswalk.¹¹ We use the crosswalk to assign the corresponding CPC codes (and their probabilistic weights) to each HS6 product category. This link allows us, in turn, to match the firm-product export data with patent applications using both the SIREN and CPC codes. Given the probabilistic structure of the concordance,¹² and the fact that the same CPC code can pertain to different product categories, we end up with a dataset with many-to-many matches—*i.e.*, one patent can link to more than one product, and some products may link to more than one patent. In particular, an average of nine products (and a median of six) link to each of the 53,519 patents matched at the firm-product level (see Table 1, column 5). In all these matches, the average (median) weight is 0.31 (0.16).¹³ In the following analysis, we consider a firm-product or a firm-product-country matched to a patent using a binary indicator (0/1), where 1 denotes a probabilistic weight greater than zero. However, all the results are robust to using a more restrictive

⁹ Notice that the difference between the number of patents in columns (2) and (3) arises partly because firms applying for a patent over 2002–2011 may not be exporting in 2011 but in some previous year. If we match patents applications with firms that have exported in at least one year over the period 2002–2011, we match 73,834 patent applications.

¹⁰ There are 204 patent offices in the world, according to WIPO. PATSTAT covers 114 patent offices and contains data for all major patent offices.

Source: <https://www.wipo.int/directory/en/urls.jsp>

¹¹ The concordance tables can be downloaded at

<https://sites.google.com/site/nikolaszolas/PatentCrosswalk>.

¹² The average and the median number of CPC classes related to a single product category are 4.6 and 4.

¹³ Notice that the concordance tables based on the probability weighting structure described in Lybbert and Zolas (2014) already implement a two-percent cutoff condition for the weights, meaning that all weights below two percent are dropped, and the remaining weights are re-normalized.

threshold, *i.e.*, considering a binary indicator where 1 denotes a probabilistic weight greater than 0.10.¹⁴

Table 1, column (5), presents the descriptive statistics. The share of transactions covered by a patent drops from 22 percent in column (3) to around ten percent. The share of export value accounted for by patents decreases less markedly, going from 42 percent to 35 percent. Similarly to what we observed at the firm-country level, lower coverage rates are due to the fact that firms holding a patent export a wide range of products (3,748), but their patents' technological classes (proxied by CPC codes) are related only to a subset of such product categories (2,878)—thus 870 product categories are not covered by patents, most likely due to the fact that patent protection is not an effective means of appropriation in such cases. On the other hand, firms' patents also cover technological classes unrelated to their product portfolio.¹⁵

Fourth, regarding the firm-product-country match, we further impose that the patent application related to a specific product exported in country c is filed to patent office c . Table 1, column (6) presents the descriptive statistics. As expected, this additional matching level further reduces the share of transactions covered by a patent (around 5.5%) and the share of export value accounted for by a patent (27.4%). The number of countries and the number of products covered by a patent do not change significantly with respect to columns (4) and (5). However, there is a slight drop in the number of patent applications matched with respect to column (5). This difference means that, sometimes, a transaction is covered by a patent relevant for the country but not for the product exported or that a product covered by a patent is also exported to a country with no patent protection.

2.2 Descriptive statistics

Table 2 reports the breakdown of patents across sectors classified according to the OECD definition of technology intensity. In columns (1)–(3) (panel 'Firm'), we assign firms to sectors based on their most important product (highest export share). As expected, high-tech sectors display the highest average number of patents per firm (2.35), and the number steadily decreases with the technological intensity of sectors.

Column (3) reports the average share of patents relating to process inventions. We take advantage of the work by Ganglmair et al. (2022), who classified granted U.S. patents into product patents vs. process patents based on their textual content. A total of 55,073 patent applications in our sample are filed at the USPTO over 2002–2011, of which 35,847 have been granted. By extending the information to the patent families, we can classify 30,888 patent families out of 63,401 total patents by exporting firms. The average share of process-related patents is around 24 percent, and it is fairly stable across sectors.

¹⁴ These additional results are available upon request.

¹⁵ There are, again, four potential explanations for this pattern: the firm has not yet exported the product; it no longer exports it; it temporarily suspended export; or it has never intended to export (or even produce)—this last explanation is possibly related to firms that “know more than they do,” see Patel and Pavitt (1997) and Dosi et al. (2017).

Table 2: Breakdown of patent applications across sectors, years 2002-2011

	Firm			Firm-product			
	No obs. (1)	mean (2)	process (3)	No obs. (4)	mean (5)	weight (6)	process (7)
High-tech	12,986	2.35	0.22	111,749	1.30	0.46	0.19
Medium-high-tech	25,767	1.07	0.24	167,280	1.40	0.42	0.29
Medium-low-tech	14,446	0.49	0.26	131,026	0.67	0.18	0.25
Low-tech	38,966	0.09	0.24	232,624	0.12	0.05	0.20
Total	92,165	0.75	0.24	642,679	0.77	0.24	0.24

Notes: The table reports the breakdown of patent applications (filed in 2002–2011) across sectors classified according to the OECD definition of technology intensity. Columns (1)–(3) report: the number of firms, the average number of patents per firm, and the average share of process-related patents per firm. Columns (4)–(7) report: the number of firm-product observation, the average number of patents per firm-product, the weighted average number of patents per firm-product, and the average share of process-related patents per firm-product.

Columns (4)–(7) (panel ‘Firm-product’) report similar statistics for patents matched at the firm-product level using the ALP crosswalk. The average number of patents per firm-product in column (5) increases compared to figures in column (2) for all but the high-tech sector. This pattern is due to the fact that high-tech products are more likely to be the most important products, in terms of export shares, among multi-product exporters with patents. Around 29 percent of firm-product pairs with a patent are in high-tech sectors, but this share reaches 35 percent if we consider only the most relevant product for each firm (not reported in Table 2).

Finally, column (6) reports the weighted number of patents per firm-product, where the weights are taken from the ALP crosswalk, and column (7) reports the average proportion of process-related patents. If we consider the weighted number of patents, high-tech sectors display the highest average, whereas the medium-high-tech sectors show the highest incidence of process-related patents.

The proportions are more heterogeneous at a finer level of classification. Appendix Tables A1 and A2 present the distributions of patents per firm and firm-product across sectors defined at the two-digit level of the HS6 system. In the latter case, the highest average number of patents per firm is in Pharmaceutical Products (4.1), whereas the proportion of process-related patents is as high as 68 percent in vegetable products.

2.3 The relative timing of patenting and export

The data offer us the opportunity to investigate the relative timing of patenting with respect to exporting. They allow us to ask whether patenting tend to anticipate exporting or vice-versa.

The cross-sectional evidence for 2011 in Table 1 shows that around five percent of all trade flows match to (at least) one patent at the firm-product-country level. In absolute numbers, this figure corresponds to 106,564 transactions. If we consider the whole period for which trade data are available (2002–2011), 265,001 *unique* (*i.e.*, not repeated over time) firm-product-country transactions

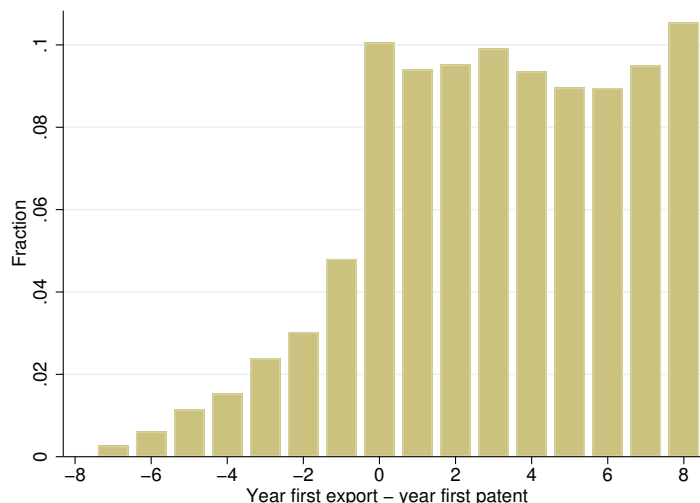


Figure 1: Distribution of the time interval between the first patent and the first export at the firm-product-country level. Patent data from 1996, export data from 2004. Source: our elaborations on the combined dataset

match to a patent. Taking advantage of the longest period available, we investigate whether the patenting activity precedes or follows exporting. To this end, we compute the difference between the year we first detect an export transaction in a firm-product-country cell and the year we first observe patenting in that cell. A positive difference provides evidence that patenting activity precedes exporting.

Figure 1 reports the distribution of the time differences between patenting and exporting. We went through the following steps to draw it. First, we set the maximum timespan between the first export and the first patent application to eight years. This limit arises because the stock of patents is computed over a ten-year time interval, and any time difference greater than nine years would be an artifact of the stock construction. Accordingly, we consider patent data starting from 1996. Second, regarding exporting activity, we only consider transactions from 2004 onward, thus allowing us to check whether a given export transaction was already taking place in the first two years (export data are available starting in 2002). These filters limit the resulting set of observations to 118,365.

Notice, first, that around 90 percent of all observations reported in Figure 1 are between zero and eight; the export activity, therefore, follows the patent application at the firm-product-country level. There is a clear jump in the distribution at zero. In around ten percent of the cases, firms export to a given product-country destination and apply for the corresponding patent in the same year. There is also a combined total of ten percent of the cases in which the patent application is subsequent to the export activity. Patent law is such that patenting needs to take place before, or very shortly after, a product is commercialized—we should thus not observe negative values. Indeed, inventions disclosed to the public cannot be patented. Therefore, inventions relating to products should first




Figure 2: Distribution of the time interval between the first patent and the first export at the firm-product-country level. Patent and export data from 2004. Source: Our elaborations on the combined dataset.

be submitted to the patent office before the product is commercialized.¹⁶ The main explanation for the negative values is that some product categories contain additional varieties that are not protected by the subsequent patent.

In Figure 2, we replicate the same exercise, but we set the same initial period for patents and exports (2004). This restriction narrows the set of observations to 64,871. The general message is unchanged: in around 75 percent of the cases, the firm-product-country export triplet is activated during or after a patent application. However, the contemporaneity between export and patent is now dominant compared to Figure 1, concerning around 18 percent of the cases. The decline of the distribution to the right of the zero is a mechanical consequence of the fact that we are leaving out older patents.

3 Patents and export value

This section investigates the effects of patenting activity on export value in the observed firm-product-country flows exploiting both cross-sectional and longitudinal variations.

First, following the trade literature (see, for instance, Bernard et al., 2007a), we decompose firm’s total exports to a product-country destination into extensive (quantity) and intensive margins (unit values):

$$\ln X_{fpct} = \ln Quantity_{fpct} + \ln UnitValue_{fpct} \quad (1)$$

where $\ln X_{fpct}$ is the log value (in euro) of exports by firm f in product p to country c in year t , which we decompose in $Quantity_{fpct}$ and $UnitValue_{fpct}$,

¹⁶Up to a one-year ‘grace period,’ in some countries. See, e.g., 35 U.S. Code §102, “Conditions for patentability; novelty.” for the United States.

Table 3: Exports and firm-product-country patents

	(1)	(2)	(3)	(4)	(5)	(6)
	$\ln X_{fcpct}$	$\ln Quantity_{fcpct}$	$\ln UnitValue_{fcpct}$	$\ln X_{fcpct}$	$\ln Quantity_{fcpct}$	$\ln UnitValue_{fcpct}$
ONLY PATENTING FIRMS						
$DPat_{fpc}$	0.063*** (0.008)	0.057*** (0.009)	0.006 (0.005)	0.064*** (0.011)	0.058*** (0.012)	0.006 (0.006)
N	3,122,592	3,122,592	3,122,592	2,540,788	2,540,788	2,540,788
adj. R^2	0.835	0.859	0.883	0.827	0.845	0.861
ONLY PATENTING FIRMS - NON-SWITCHERS + SWITCHERS FROM 0 TO 1						
$DPat_{fpc}$	0.043*** (0.011)	0.047*** (0.012)	-0.005 (0.006)	0.053*** (0.014)	0.050*** (0.017)	0.003 (0.009)
N	2,959,032	2,959,032	2,959,032	2,376,768	2,376,768	2,376,768
adj. R^2	0.838	0.861	0.885	0.832	0.848	0.864
ONLY PATENTING FIRMS - NON-SWITCHERS + SWITCHERS FROM 1 TO 0						
$DPat_{fpc}$	0.116*** (0.015)	0.090*** (0.016)	0.026*** (0.008)	0.106*** (0.020)	0.085*** (0.021)	0.021* (0.010)
N	2,809,944	2,809,944	2,809,944	2,230,901	2,230,901	2,230,901
adj. R^2	0.838	0.862	0.886	0.832	0.848	0.864
Firm-Product-Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Product-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Country-Year FE	Yes	Yes	Yes	No	No	No
Product-Country-Year FE	No	No	No	Yes	Yes	Yes

Note. The table reports estimation results from equation (2) at the firm-product-country level, using data on exports, quantity, and unit value for 2002–2011. Robust standard errors clustered at the product-year level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

capturing the physical quantity (in kilo) and the unit value of the transaction, respectively.

Then, we estimate the main equation as:

$$\ln Y_{fpc} = a + \beta DPat_{fpc} + \theta_{fpc} + \theta_{fpt} + \theta_{c(p)t} + \varepsilon_{fpc} \quad (2)$$

where $\ln Y_{fpc}$ denotes the logarithm of, alternatively, the total value, quantity, and unit value of the firm's exports in the country-product pair and $DPat_{fpc}$ is a binary variable taking the value of 1 if the firm f has a positive ten-year stock of patents in country c and product p at time t .¹⁷

Equation (2) also include a rich set of fixed effects that drive the identification of β . First, the set of destination-year fixed effects (θ_{ct}) accounts for destination-specific time-varying characteristics common to all firms. These characteristics includes, for example, ‘gravity’ variables such as the country GDP per capita and the cyclical component of the demand in the destination country, as well as other shocks and systematic changes that might correlate with the independent variable. For instance, Coelli et al. (forthcoming) show that most favored nation (MFN) tariff cuts in the aftermath of the Uruguay Round in 1994 significantly

¹⁷ The patent dummy switches from zero to one at the time of priority application. This time is identical for all destination markets in which the firm owns a patent of the same family. In practice, firms may file patent applications at different time intervals, usually up to twelve months after the priority date but sometimes up to 31 months (for PCT applications). In our dataset, the application year is contemporaneous to the priority year in around 85 percent of cases and lags by one year in the remaining 15 percent of cases. This is in agreement with Figure 2: a mode at zero suggests that second filings tend to be filed contemporaneously to priority filings on average.

affected firm-level innovation by improving market access. Similarly, Aghion et al. (2018) show that French firms respond to demand shocks in their export destinations by patenting more. The set of destination-year fixed effects also absorbs possible changes in IP institutions. To the extent that such shocks, and others, may be not just destination-specific but destination-product-specific, we also estimate equation (2) with a set of country-product-year fixed effects (θ_{cpt}).

Second, we add two more granular sets of fixed effects involving firms: θ_{fpc} and θ_{fpt} . The former, θ_{fpc} , are firm-product-country fixed effects that capture all the unobservable characteristics specific to the firm-product-country flow. They account for the fact that two different transactions, even within the same firm, might have different characteristics that the product-country fixed effects do not pick up. For example, a firm can export more expensive varieties to richer countries (Manova and Zhang, 2012). The within-transaction identification achieved through the inclusion of θ_{fpc} is, of course, not sufficient to insulate our results from endogeneity concerns related to time-varying transaction-specific shocks. In order to partially alleviate this issue, we include the latter set of fixed effects, θ_{fpt} , capturing time-varying characteristics specific to the firm-product. These fixed effects take into account, for example, the fact that firm capabilities may change over time, affecting both the propensity to patent within a product category and the expansion on international markets.

To sum up, the set of fixed effects ensure that the identification of β depends *solely* on ‘switching’ transactions, *i.e.*, firm-product-countries flows that either switch from 0 to 1 (starting to patent) or from 1 to 0 (losing patent protection—*i.e.*, no patent in the previous ten years). When a change in $DPat_{fpct}$ occurs within a transaction, the identification further requires (through the presence of θ_{fpt}) that such a change is realized only in a subset of countries toward which the firm-product pair is exported. For instance, a firm exporting a given category of motor vehicles to three different markets may decide to apply for a patent only in two destinations. Thus, we identify the effect by comparing, within a firm-product pair, the export performance over time in destination markets in which a change in $DPat_{fpct}$ occurs *versus* destination markets in which there is no change.

In this quasi difference-in-differences estimation, the set of destination-year or destination-product-year fixed effects (θ_{ct} and θ_{cpt}) accounts for time-varying characteristics common to all firms within a given market. However, as patenting and non-patenting firms might be affected differently by such common shocks, we estimate equation (2) exploiting information only on transactions made by the subset of patenting firms. In total, the estimation sample contains 4,204,015 transactions made by 7743 patenting firms.¹⁸ As the specifications involve a rich set of fixed effects, we estimated them using the high-dimensional fixed effects estimator developed by Correia (2017) and implemented in the Stata routine *reghdfe*.¹⁹

¹⁸Notice that this number is different from the 4,251 patenting firms reported in Table 1, column 3, which refers only to 2011. In our estimation sample, we consider firms with a positive stock of patents in at least one year.

¹⁹Notice that the Stata routine *reghdfe* reports the number of observations effectively used in the estimation: observations that do not contribute to identifying the coefficient are not

Table 3 reports results from estimating equation (2). Columns (1)–(3) use country-year fixed effects (θ_{ct}), and columns (4)–(6) use country-product-year fixed effects (θ_{cpt}). Robust standard errors are clustered at the product-year level to account for possible correlations across countries within the same product category. The first panel exploits the whole sample of trade flows. The sample includes the ‘switchers’ (459,780 observations) and all other observations, which fall in the following two cases: (i) observations by firms that patented in another (not exported) product-country variety; (ii) observations by firms that do have a positive patent stock in the relevant product-country throughout the whole observation period. Within firm-product-country flows, a switch in patenting status, net of any firm-product specific trend, is associated with an average increase in export value of around 6.3 percent. It is apparent that such an increase is mainly associated with a quantity effect; indeed, the coefficient on unit value is very low (0.006). Columns (4)–(6) report alternative specification results that controls for country-product-year fixed effects instead of country-year effects. This specification leads to very similar results.

The second and the third panels of Table 3 focus on specific groups of switchers. Among patenting firms, the total number of ‘switchers,’ that is, firm-product-country cells that exhibit a change in the variable $DPat_{fpct}$ in the study period, reaches 459,780. Of those, around 60 percent have a ‘first export, then patent’ dynamics, with a unique switching pattern in $DPat_{fpct}$ from 0 to 1. On the other hand, around 40 percent of observations have a ‘first patent, then export’ dynamics,²⁰ with two possible switching patterns: from 1 to 0 (around 90%) and from 1 to 0, followed by 0 to 1 (around 10%). The second panel of Table 3 restricts the set of switchers to the type ‘first export, then patent’ (*i.e.*, from 0 to 1), whereas the third panel restricts the set of switchers to the type ‘first patent, then export,’ excluding multiple switchers.²¹ In both exercises, the composition of the remaining sample is the same as in the first panel and includes cases (i) and (ii) discussed above.

Turning to the second panel, results remain very similar to those previously obtained, with the effect for export value in the range of 0.043–0.053, again driven by a quantity effect. Turning now to the third panel, the coefficients for export value increase to 0.106–0.116; moreover, we also observe a small but statistically significant effect on price of around 2.5 percent. This result implies

reported in the N statistics.

²⁰ The analysis on the exporting and patenting dynamics that we present here resonates with the evidence presented in Section 2.3. Note, however, that the two analyses rely on very different sets of observations. We recall the most relevant differences. First, we employ repeated observations over time in a given firm-product-country triplet, while in Section 2.3, we only used the *unique* occurrence of a given triplet. Second, as we have data on a transaction only when a trade occurs, this implies an over-representation of the ‘first export, then patent’ occurrences, which by construction are always switchers (from 0 to 1). On the other hand, the ‘first patent, then export’ occurrences are switchers only when they go from 1 to 0; all other cases are within-group (ii). While the first explanation is responsible for the sharp increase in the number of available observations (from 118,365 in Section 2.3 to 459,780), the second relates to the change in the proportion of the switches.

²¹ Notice that this classification includes the years 2002 and 2003, so an exporter that starts in 2002 with a patent is classified as a ‘first patent, then export.’

that for transactions already covered by a patent at the time of first export, losing patent protection comes with a decrease in export value. The decrease comes primarily from a quantity effect, even if we also detect a slight price decrease.²²

Taken together, these results suggest that international transactions covered by a patent enjoy a greater export value, which is mainly due to a greater export quantity. These results are consistent with a *market expansion* effect of patents, whereby property rights expand foreign markets by ensuring exclusive rights over inventions: firms are more likely to export when protected against the risk of imitation by foreign firms. On the other hand, the limited evidence of a unit price premium seems to rule out a *market power* mechanism, which would predict that firms securing a patent in a destination country would increase the unit price. Although this interpretation is consistent with more aggregate empirical evidence (Smith, 2001), one should still consider that the analysis is at the level of a product category. The possibility exists, for example, that a new product covered by a patent and sold with a higher unit price could compensate for a decline in the price of pre-existing products within the same product category—leading to a null aggregate price effect.

3.1 Product-related patents

Table 4: Exports and firm-product-country patents: only product-related patents

	(1)	(2)	(3)	(4)	(5)	(6)
	$\ln X_{fcpt}$	$\ln Quantity_{fcpt}$	$\ln UnitValue_{fcpt}$	$\ln X_{fcpt}$	$\ln Quantity_{fcpt}$	$\ln UnitValue_{fcpt}$
ONLY PATENTS CLASSIFIED AS PRODUCT OR PROCESS PATENTS						
$DPat_{fpet}$	0.050*** (0.011)	0.039*** (0.012)	0.011* (0.007)	0.036** (0.015)	0.032* (0.016)	0.004 (0.009)
N	2,652,500	2,652,500	2,652,500	2,101,756	2,101,756	2,101,756
adj. R^2	0.838	0.860	0.883	0.829	0.843	0.858
INCLUDING ONLY PRODUCT-RELATED PATENTS						
$DPat_{fpet}$	0.078*** (0.014)	0.059*** (0.016)	0.018** (0.009)	0.045** (0.020)	0.029 (0.021)	0.016 (0.013)
N	2,218,442	2,218,442	2,218,442	1,733,383	1,733,383	1,733,383
adj. R^2	0.833	0.851	0.877	0.822	0.832	0.852
Firm-Product-Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Product-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Country-Year FE	Yes	Yes	Yes	No	No	No
Product-Country-Year FE	No	No	No	Yes	Yes	Yes

Note. The table reports estimation results from equation (2) at the firm-product-country level, using data on exports, quantity, and unit value for 2002–2011, using the information on product-related patents. Robust standard errors clustered at the product-year level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

As mentioned before (see Section 2.1), the analysis exploits a probabilistic concordance between the 2002 HS classification and the CPC codes. As this probabilistic matching is necessarily open to measurement errors, this section proposes a robustness test that restricts the sample to transactions matched to a patent classified as a ‘product’ patent. The match between product categories and patents should be more accurate for patents marked as product patents

²² Notice that the differences between coefficients in the second and third panel should not be over-interpreted as in most cases the confidence intervals do overlap.

compared to process patents—if only because a process patent may apply potentially to a greater variety of products. Accordingly, we exclude from the sample transactions associated to a process patent using the classification by Ganglmair et al. (2022), as explained in Section 2.2. As information is available only for a subsample of patents, we first estimate equation (2) on the main sample of patenting firms but excluding firm-product pairs matched to an unclassified patent. Then, within this restricted sample, we exclude firm-product pairs matched to a process-related patent.

Table 4 reports the results. Looking at the bottom panel, excluding process-related patents from the sample confirms the central finding. The results suggest that a within-transaction change in the patenting status is associated with a significant change in export value, with a coefficient equal to 0.078, column (1). This effect is primarily driven by a quantity effect, although we note that the coefficient related to the unit value is now positive and significant, with a value of 0.018, column (3). However, this result is mostly a sample composition effect, as the coefficient is positive and significant also in the top panel, column (3). Accounting for product-country-year fixed effects (columns 4–6) yields results consistent with those presented in Table 3. The main difference is that the coefficient associated with the quantity regression loses statistical significance (also, it seems, because of a sample composition effect).

3.2 Accounting for patent importance

Table 5: Exports and firm-product-country patents: interacting with indicators of patent importance

	(1)	(2)	(3)	(4)	(5)	(6)
	$\ln X_{fcpt}$	$\ln Quantity_{fcpt}$	$\ln UnitValue_{fcpt}$	$\ln X_{fcpt}$	$\ln Quantity_{fcpt}$	$\ln UnitValue_{fcpt}$
FAMILY SIZE						
$DPat_{fpet}$	0.041** (0.017)	0.042** (0.019)	-0.001 (0.010)	0.034 (0.023)	0.046* (0.026)	-0.012 (0.013)
$DPat_{fpet} \times Family\ size$	0.003* (0.001)	0.002 (0.002)	0.001 (0.001)	0.004* (0.002)	0.002 (0.002)	0.002 (0.001)
N	2,867,657	2,867,657	2,867,657	2,285,921	2,285,921	2,285,921
adj. R^2	0.835	0.859	0.883	0.826	0.842	0.858
COUNT OF CITATIONS RECEIVED IN A 3-YEAR WINDOW						
$DPat_{fpet}$	0.043*** (0.013)	0.039*** (0.014)	0.005 (0.007)	0.030* (0.017)	0.029 (0.019)	0.001 (0.010)
$DPat_{fpet} \times \#citations$	0.020*** (0.007)	0.021*** (0.007)	-0.001 (0.004)	0.033*** (0.009)	0.033*** (0.010)	0.000 (0.005)
N	2,867,657	2,867,657	2,867,657	2,285,921	2,285,921	2,285,921
adj. R^2	0.835	0.859	0.883	0.826	0.842	0.858
Firm-Product-Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Product-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Country-Year FE	Yes	Yes	Yes	No	No	No
Product-Country-Year FE	No	No	No	Yes	Yes	Yes

Note. The table reports estimation results from equation (2) at the firm-product-country level, using data on exports, quantity, and unit value for 2002–2011, interacting with indicators of patent importance. Robust standard errors clustered at the product-year level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A further concern related to the main results of Table 3 is that the relationship between export values and patenting could depend on the importance of the

underlying patents. Testing for this possibility, we augment equation (2) with two different variables proxying for the importance of the patents associated with a transaction. In the first specification, we interact $DPat_{fpct}$ with a variable measuring the number of patent applications within the same family (variable *Family size*).²³ In the second specification, we interact $DPat_{fpct}$ with a variable measuring the number of citations received in a three-year window by all patents attached to a given transaction (variable *#citations*).

Three remarks are in order before commenting on the results presented in Table 5. First, the family size and number of citations are available only for patents filed at the USPTO. We exploited family linkage to populate values to a broader set of patents, similarly to what we did for the process information variable. Note, however, that we still end up with a slightly smaller sample than in Table 3. Second, the two patent importance variables exhibit different distributional properties. Whereas a transaction is associated with an average of 1.6 families, with a standard deviation of 4.60, it is associated with an average of 7.5 citations, with a standard deviation of 124.3. In view of this difference, we use the level of the variable *Family size* and the unit-offset logarithm of the variable *#citations*. Finally, given the interaction terms, the coefficient on $DPat_{fpct}$ corresponds to the patent effect when the importance variable takes on the baseline value, *i.e.*, 1 for *Family size* and 0 for the number of *#citations*. The coefficient associated with the importance variables measures the additional effect of having a larger family or more citations.

The top panel of Table 5, exploiting information on family size, shows that the baseline coefficient on $DPat_{fpct}$ is still positive and significant for export value and quantity in columns (1), (2), and (5). Moreover, patents with larger families are associated with larger export values. One additional protection country corresponds to an export value about 0.3–0.4 percent higher. Results using the number of citations presented in the bottom panel of Table 5 depict a similar pattern. The baseline effect remains, and a ten-percent increase in the number of citations is associated with an export value about 0.2–0.3 percent higher.

3.3 Accounting for endogeneity

Despite the high level of disaggregation of our data, which allows us to use a rich set of fixed effects, the econometric specification does not deal fully with the endogeneity between the patenting decision and the firm’s export activity. The core of the issue is that a firm that files a patent application in a specific country expects to benefit from it. This section addresses this concern by exploiting information on *rejected* patent applications.

Patent applications that compose our dataset cover particularly valuable inventions, as international patent protection implies considerable costs for the firms (Harhoff et al., 2003). Furthermore, the fact that these patent applications protect actual products sold on the market provides another indication of their importance. Evidence by Guellec and van Pottelsberghe (2000) suggests that in-

²³ We exploit the INPADOC patent family, which links together patents covering a technology. All patents in a family share common priority filings.

ternational patent applications are more likely to be granted than single-country applications. Clearly, a firm that files a patent application in an export market expects, or at least hopes, that it will receive a patent.

The grant is not automatic, however. Each patent application undergoes a substantive examination by the local patent office. Countries are sovereign in their decisions to grant patents (see, *e.g.*, de Rassenfosse and Raiteri, forthcoming), and the literature has documented heterogeneity in the grant decision across patent offices for the same invention (de Rassenfosse et al., 2021). Some of these differences in grant outcome may be due either to intrinsically different standards between patent offices or to patent offices applying their own standards inconsistently. This insight suggests that, although firms might expect that they will get a patent, the patent application may be rejected during the examination. Such rejection is plausibly an exogenous event that is, by and large, unforeseen by firms but at the same time relevant to the value of the export transaction.²⁴ Hence, our identification strategy exploits rejection events to study the effect of (a loss of) patent protection on exports.

We crawled the Google Patents website to obtain accurate data on the grant outcome of patent applications in our sample. Since a patent’s rejection is a rather rare event in our data, we limit the data collection effort to the most relevant destination countries, namely, those with the highest number of patents filed by firms in the sample and reliable data available. These countries include the United States, China, Korea, Japan, and Australia.²⁵

We start by replicating the main analysis, as presented in Table 3, to the subset of five countries. The results, reported in Table 6, top panel, are qualitatively and quantitatively similar to those obtained on the entire sample. The estimated coefficients on export value and quantity remain positive and statistically significant, although they are estimated with less precision when adding product-country-year fixed effects (columns 4–6).

Having established the stability of the results, we now study the effect of rejection events. We construct a new binary variable, $DRejection_{fpct}$, taking the value of 1 if firm f has a positive ten-year stock of rejected patents in country c and product p at time t . Since this variable is defined only for transactions with a positive value for $DPat_{fpct}$ in at least one year, we restrict the regression to those observations. Also, we focus exclusively on exogenous events and, consequently, exclude firm-product-countries with a granted patent from the estimation sample. Notice again that the set of fixed effects implies that we identify the effect of a change in the extensive margins of rejections, *i.e.*, the effect of a first rejection.

Table 6, bottom panel, reports the results. The coefficients now flip sign, suggesting that being refused a patent hits export. In column (4), export values drop by almost 70 percent, arising from a drop in export quantities and unit

²⁴ Firms may also withdraw their patent applications during the prosecution process. Given that patents in our sample are particularly valuable, withdrawals are unlikely to occur. Moreover, when they do, they may be induced by a negative search or examination report by the patent office (Lazaridis and van Pottelsberghe, 2007).

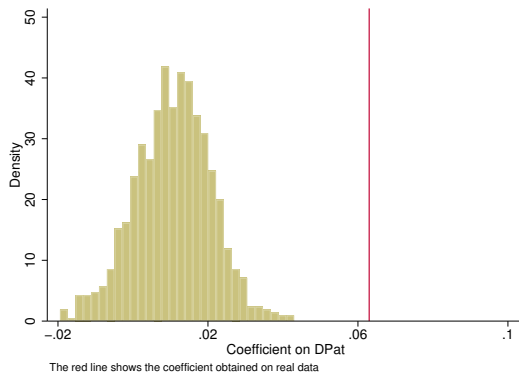
²⁵ We decided to exclude European patents, as in our sample, most of them are filed through the European patent office. It is particularly complex to correctly pinpoint the grant outcome in all the designated states.

Table 6: Exports and firm-product-country patents: United States, China, Korea, Japan, and Australia

	(1)	(2)	(3)	(4)	(5)	(6)
	$\ln X_{fcpt}$	$\ln Quantity_{fcpt}$	$\ln UnitValue_{fcpt}$	$\ln X_{fcpt}$	$\ln Quantity_{fcpt}$	$\ln UnitValue_{fcpt}$
ONLY PATENTING FIRMS						
$DPat_{fpc}$	0.096*** (0.031)	0.124*** (0.034)	-0.028 (0.018)	0.073* (0.038)	0.094* (0.047)	-0.021 (0.024)
N	156,758	156,758	156,758	121,901	121,901	121,901
adj. R^2	0.771	0.842	0.891	0.705	0.779	0.833
REJECTED PATENTS						
$DRejection_{fpc}$	-0.137 (0.086)	-0.020 (0.097)	-0.116** (0.050)	-0.693** (0.285)	-0.318 (0.325)	-0.374** (0.173)
N	13,342	13,342	13,342	3,767	3,767	3,767
adj. R^2	0.764	0.855	0.901	0.677	0.791	0.841
Firm-Product-Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-Product-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Country-Year FE	Yes	Yes	Yes	No	No	No
Product-Country-Year FE	No	No	No	Yes	Yes	Yes

Note. The table reports estimation results from equation (2) at the firm-product-country level, using data on exports, quantity, and unit value for 2002–2011. Bottom panel: only observations with positive $DPat_{fpc}$ in at least one year. Robust standard errors clustered at the product-year level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

(a) With country-year fixed effects



(b) With country-product-year fixed effects

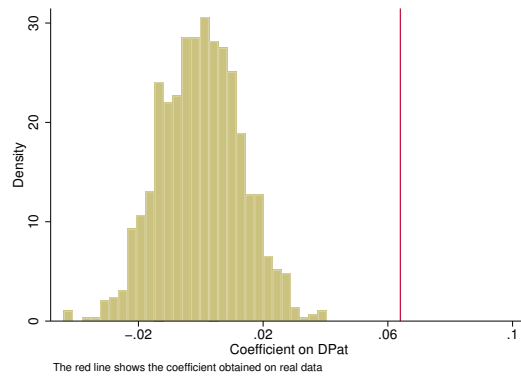


Figure 3: Distribution of coefficients on $DPat_{fpc}$ in the export value estimation of equation (2) on simulated data

values. The magnitude of the coefficient is consistent with the fact that the firm is being pushed out of the export market. Indeed, a common reason for patent rejection is lack of novelty, presumably in this case because a competitor already patented a similar invention. Hence, as the firm finds out that its patent is not novel, selling the associated product puts it at substantial risk of patent infringement litigation.

4 Placebo analysis and additional results

Equation (2) controls for common average trends between switchers and non-switchers using a full set of country-year and country-product-year fixed effects.

However, all observations may not all be equally well represented by such a common average. In particular, there could be some non-switching firm-product-countries trade flows with dynamics similar to the switchers. If this were the case, some underlying factors other than the patenting activity could affect the dynamics of the switchers and non-switchers.

In order to check that such factors do not drive the results, we run a placebo test in the spirit of Athey and Imbens (2017). We first select the set of firms that we observe for all ten years among all observations by patenting firms. This selection allows us to compare switching and non-switching observations, abstracting from possible differences in the observation period. It leaves us with 736,010 observations. To check that this particular cut of the data does not drive the placebo results, we estimate equation (2) on the (log) export value. The coefficient (clustered standard errors) on $DPat_{fpct}$ is 0.058 (0.012), which is in line with the results presented in Table 3. We then exclude observations that match to a patent (around 18%), and randomly reassign their different sequences of ‘0’ and ‘1’ to the other firm-product-country transactions, keeping constant the share of switchers within products-country pairs. After each random assignment, we estimate equation (2) on the resulting sample, both with country-year and country-product-year fixed effects.

After 1,000 simulations, we obtain the distribution of coefficients on $DPat_{fpct}$ reported in Figure 3. Panel (a) reports results for the country-year fixed effects specification and Panel (b) for the country-product-year fixed effects specification. The vertical lines indicate the value of the coefficient obtained using real data. In panel (a), the average estimated coefficient after 1,000 simulations is 0.01 (standard deviation 0.01), which is quite far from the 0.063 coefficient we obtain on real data. In panel (b), we also obtain a sizeable difference between the average coefficient, 0 (standard deviation 0.014), and the real coefficient, 0.064. Both results are reassuring. In particular, the fact that the impact of patenting is on average null when we randomly assign treatment and control with country-products-year fixed effects makes us confident that, within a product-country pair, trade flows do have relatively homogeneous dynamics—making it more likely that what happens following a real patent event is a genuine effect of the patent.

4.1 Accounting for selection in exporting

Selection in exporting is another potential source of bias that can affect our results. A non-patenting firm that exports a given product to a given destination may have an above-average export performance due to some unobserved characteristics. This selection implies that our estimates are conservative. Our main analysis already accounts for cross-sectional dimensions of selection bias usually investigated in the trade literature (Crozet et al., 2011). This section investigates the extent to which selection in the panel dimension could affect the results.

Our approach is the following. First, we expand our dataset at the firm-product-country level to include zeros when we do not observe a trade flow. Since we do not have information on firms’ domestic production, we place zeros on a given firm-product-country triplet if we observe the same firm-product variety

Table 7: Exports and firm-product-country patents: extensive and intensive margins

Dep. variable	(1) D_{exp}	(2) $\ln X_{fcpt}$	(3) X_{fcpt}
$DPat_{fpct}$	0.005*** (0.000)	0.056*** (0.007)	0.080*** (0.026)
N	23,895,179	1,436,151	23,895,179
adj. R^2	0.6712	0.835	
Pseudo R^2			0.9572
Firm-Product-Country FE	Yes	Yes	Yes
Firm-Product-Year FE	Yes	Yes	Yes
Country-Year FE	Yes	Yes	Yes

Note. The table reports results of: a selection equation with a binary dependent variable (column 1) and an export value equation with log export value (column 2) estimated via high-dimensional fixed effects; and export values with zeros (column 3) estimated via pseudo-Poisson. Robust standard errors clustered at the product-year level in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

in another country. Practically, we take the first and the final year in which we observe a firm-product variety, and we place zeros in the ‘empty’ years in which the variety is not observed—thereby eliminating gaps in the series—as well as in all countries to which the variety is not exported. In order to enhance the computational feasibility of the exercise, we restrict the dataset to firms, products, and countries that have at least one patent matched during the period, leading to a dataset of about 25 million observations. Using this expanded dataset, we first estimate the effect of a switch in the patenting status on the probability to export. Then, using the approach suggested by Silva and Tenreyro (2006), we estimate an export equation in levels with the zeros via the pseudo-maximum-likelihood Poisson technique.²⁶ Notice that this approach does not allow us to decompose the quantity and price effect.

Table (7) reports the results. Column (1) presents the result of the selection equation in which the dependent variable is a binary variable denoting whether a firm is exporting in the given product-country cell. Because of the large number of fixed effects, we estimate the equation via a simple linear probability model. Results suggest a positive and significant association between patents and the probability to export at the firm-product-country level. The coefficient reaches 0.5 percentage points, a sizeable magnitude compared to the unconditional probability to export in any given product-country cell of six percent.

Column (2) reports the standard estimation of equation (2) on log export value. The coefficient on $DPat_{fpct}$ is equal to 0.056, only slightly different from the 0.063 coefficient in Table 3, first panel. Finally, column (3) reports the

²⁶ We have used the Stata command `ppmlhdfe`, which allows to include multi-way fixed effects.

estimation results of the equation for export value in levels and with zeros. The estimated coefficient suggests that a change in $DPat_{fpc}$ comes with an increase in export value of around eight percent, confirming that our previous estimates are conservative with respect to the selection issue.

5 Concluding remarks

This paper reports novel evidence on the extent to which international trade hinges on patents. It analyzes the export and patenting activities of a panel of French exporting firms from 2002 to 2011. The noticeable feature of our study is that we observe the worldwide patenting activity at the product level, which enables a fined-grained exploration of the effect of patents on trade. Furthermore, we also exploit information on rejected patent applications to alleviate concerns about the endogeneity between the patenting and export decisions. The results of the empirical analysis provide an unprecedented perspective on the patenting and exporting behavior of firms.

First, there is overwhelming evidence suggesting that patenting activity precedes exporting in a given product-country destination. This finding is not surprising, but it establishes the validity of our empirical setup. Patent law stipulates that offering a product for sale counts as a ‘public disclosure’ of the invention—and inventions disclosed to the public cannot be patented. Therefore, inventions relating to products should first be submitted to the patent office before the product is commercialized.

Second, we assess the effect of a patent filing on exports. Controlling for a comprehensive set of fixed effects, we find that patenting is associated with an increase in the value of exports of around six percent. This result is primarily driven by greater quantities exported to the destination market rather than higher prices. We interpret it as evidence that firms value foreign patent protection for the legal security that it brings rather than for the possibility of setting monopoly prices. The findings are robust to a series of alternative specifications and a placebo test. Furthermore, acknowledging that the patenting decision is endogenous to exports, we exploit information on rejected patent applications. We find that exports collapse when patent applications are rejected by the local patent office, providing further evidence on the causal effect of patents on trade.

Taken together, the results point to the existence of a patent premium on the export values for transactions associated with a patent. The very limited evidence of a positive premium on unit values, our proxy for the price, poses some challenges to future empirical and theoretical analyses. Large-scale empirical works matching product varieties, or even actual products, to patents are needed to confirm the findings of the present work. Currently, since we work at the product class level, we cannot exclude the possibility that the introduction of patented products with a higher price (*e.g.*, a new mobile phone) is compensated by a decline in the price of the pre-existing phones in the same product category, with an average effect of zero. Thus, the finding of a lack of price effect needs to be taken with a pinch of salt. On a more theoretical viewpoint, our results suggest exploring models in which a substantial part of the value residing in

foreign patent protection is in the legal security that it offers, rather than in the possibility of setting monopoly prices as the literature traditionally assumes.

References

- AGHION, P., A. BERGEAUD, M. LEQUIEN, AND M. J. MELITZ (2018): “The Impact of Exports on Innovation: Theory and Evidence,” Working Paper 24600, National Bureau of Economic Research.
- AMABLE, B. AND B. VERSPAGEN (1995): “The role of technology in market shares dynamics,” *Applied Economics*, 27, 197–204.
- AMENDOLA, G., G. DOSI, AND E. PAPAGNI (1993): “The dynamics of international competitiveness,” *Review of World Economics (Weltwirtschaftliches Archiv)*, 129, 451–471.
- ARORA, A., M. CECCAGNOLI, AND W. COHEN (2008): “R&D and the patent premium,” *International Journal of Industrial Organization*, 26, 1153–1179.
- ATHEY, S. AND G. W. IMBENS (2017): “The state of applied econometrics: Causality and policy evaluation,” *Journal of Economic Perspectives*, 31, 3–32.
- AUTOR, D., D. DORN, G. H. HANSON, G. PISANO, AND P. SHU (2020): “Foreign Competition and Domestic Innovation: Evidence from US Patents,” *American Economic Review: Insights*, 2, 357–74.
- BERNARD, A. B., J. B. JENSEN, S. J. REDDING, AND P. K. SCHOTT (2007a): “Firms in International Trade,” *Journal of Economic Perspectives*, 21, 105–130.
- BERNARD, A. B., S. J. REDDING, AND P. K. SCHOTT (2007b): “Comparative Advantage and Heterogeneous Firms,” *The Review of Economic Studies*, 74, 31–66.
- CARLIN, W., A. GLYN, AND J. VAN REENEN (2001): “Export Market Performance of OECD Countries: An Empirical Examination of the Role of Cost Competitiveness,” *Economic Journal*, 111, 128–62.
- CASSIMAN, B., E. GOLOVKO, AND E. MARTÍNEZ-ROS (2010): “Innovation, exports and productivity,” *International Journal of Industrial Organization*, 28, 372–376.
- CASTELLANI, D. AND A. ZANFEI (2007): “Internationalisation, innovation and productivity: how do firms differ in Italy?” *The World Economy*, 30, 156–176.
- COELLI, F., A. MOXNES, AND K. H. ULLTVEIT-MOE (forthcoming): “Better, faster, stronger: Global innovation and trade liberalization,” *Review of Economics and Statistics*.

- CORREIA, S. (2017): “Linear models with high-dimensional fixed effects: An efficient and feasible estimator,” *Unpublished manuscript*, <http://scoreia.com/research/hdfe.pdf> (last accessed 25 October 2019).
- CROZET, M., K. HEAD, AND T. MAYER (2011): “Quality sorting and trade: Firm-level evidence for French wine,” *The Review of Economic Studies*, 79, 609–644.
- DE RASSENFOSSE, G., H. DERNIS, AND G. BOEDT (2014): “An introduction to the Patstat database with example queries,” *Australian Economic Review*, 47, 395–408.
- DE RASSENFOSSE, G., W. E. GRIFFITHS, A. B. JAFFE, AND E. WEBSTER (2021): “Low-quality patents in the eye of the beholder: Evidence from multiple examiners,” *The Journal of Law, Economics, and Organization*, 37, 607–636.
- DE RASSENFOSSE, G. AND A. B. JAFFE (2018): “Are patent fees effective at weeding out low-quality patents?” *Journal of Economics & Management Strategy*, 27, 134–148.
- DE RASSENFOSSE, G. AND E. RAITERI (forthcoming): “Technology protectionism and the patent system: Strategic technologies in China,” *Journal of Industrial Economics*.
- DOSI, G., M. GRAZZI, AND D. MOSCHELLA (2015): “Technology and costs in international competitiveness: from countries and sectors to firms,” *Research Policy*, 44, 1795–1814.
- (2017): “What do firms know? What do they produce? A new look at the relationship between patenting profiles and patterns of product diversification,” *Small Business Economics*, 48, 413–429.
- DOSI, G., K. PAVITT, AND L. SOETE (1990): *The Economics of Technical Change and International Trade*, New York: New York University Press.
- FAGERBERG, J. (1988): “International competitiveness,” *The Economic Journal*, 98, 355–374.
- GANGLMAIR, B., W. ROBINSON, AND M. SEELIGSON (2022): “The Rise of Process Claims: Evidence from a Century of U.S. Patents,” Tech. rep., In preparation.
- GOLDSCHLAG, N., T. J. LYBBERT, AND N. J. ZOLAS (2019): “Tracking the technological composition of industries with algorithmic patent concordances,” *Economics of Innovation and New Technology*, 0, 1–21.
- GOULD, D. M. AND W. C. GRUBEN (1996): “The role of intellectual property rights in economic growth,” *Journal of Development Economics*, 48, 323–350.

- GUELLEC, D. AND B. VAN POTTELSBERGHE (2000): “Applications, grants and the value of patent,” *Economics letters*, 69, 109–114.
- HARHOFF, D., F. M. SCHERER, AND K. VOPEL (2003): “Citations, family size, opposition and the value of patent rights,” *Research policy*, 32, 1343–1363.
- HU, A. AND I. PNG (2013): “Patent rights and economic growth: evidence from cross-country panels of manufacturing industries,” *Oxford Economic Papers*, 65, 675–698.
- IVUS, O. (2010): “Do stronger patent rights raise high-tech exports to the developing world?” *Journal of International Economics*, 81, 38 – 47.
- JENSEN, P. H., R. THOMSON, AND J. YONG (2011): “Estimating the patent premium: Evidence from the Australian Inventor Survey,” *Strategic Management Journal*, 32, 1128–1138.
- LANDESMANN, M. AND M. PFAFFERMAYR (1997): “Technological competition and trade performance,” *Applied Economics*, 29, 179–196.
- LAURSEN, K. AND V. MELICIANI (2010): “The role of ICT knowledge flows for international market share dynamics,” *Research Policy*, 39, 687–697.
- LAZARIDIS, G. AND B. VAN POTTELSBERGHE (2007): “The rigour of EPO’s patentability criteria: An insight into the “induced withdrawals”.” *World Patent Information*, 29, 317–326.
- LIN, J. X. AND W. F. LINCOLN (2017): “Pirate’s treasure,” *Journal of International Economics*, 109, 235–245.
- LYBBERT, T. J. AND N. J. ZOLAS (2014): “Getting patents and economic data to speak to each other: An ‘Algorithmic Links with Probabilities’ approach for joint analyses of patenting and economic activity,” *Research Policy*, 43, 530–542.
- MANOVA, K. AND Z. ZHANG (2012): “Export Prices Across Firms and Destinations,” *The Quarterly Journal of Economics*, 127, 379–436.
- MARTÍNEZ, C. (2010): “Patent families: When do different definitions really matter?” *Scientometrics*, 86, 39–63.
- MASKUS, K. E. AND M. PENUBARTI (1995): “How trade-related are intellectual property rights?” *Journal of International Economics*, 39, 227 – 248.
- MELITZ, M. J. (2003): “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity,” *Econometrica*, 71, 1695–1725.
- PALANGKARAYA, A., P. H. JENSEN, AND E. WEBSTER (2017): “The effect of patents on trade,” *Journal of International Economics*, 105, 1–9.

- PATEL, P. AND K. PAVITT (1997): “The technological competencies of the world’s largest firms: Complex and path-dependent, but not much variety,” *Research Policy*, 26, 141–156.
- SILVA, J. S. AND S. TENREYRO (2006): “The log of gravity,” *The Review of Economics and statistics*, 88, 641–658.
- SMITH, P. J. (1999): “Are weak patent rights a barrier to U.S. exports?” *Journal of International Economics*, 48, 151 – 177.
- (2001): “How do foreign patent rights affect US exports, affiliate sales, and licenses?” *Journal of International Economics*, 55, 411–439.
- SOETE, L. (1981): “A general test of technological gap trade theory,” *Weltwirtschaftliches Archiv*, 117, 638–660.
- (1987): “The impact of technological innovation on international trade patterns: The evidence reconsidered,” *Research Policy*, 16, 101 – 130.
- VAN BEVEREN, I. AND H. VANDENBUSSCHE (2010): “Product and process innovation and firms’ decision to export,” *Journal of Economic Policy Reform*, 13, 3–24.
- WAKELIN, K. (1998a): “Innovation and export behaviour at the firm level,” *Research policy*, 26, 829–841.
- (1998b): “The role of innovation in bilateral OECD trade performance,” *Applied Economics*, 30, 1335–1346.
- ZYLKIN, T. AND C. BRUNEL (forthcoming): “Do Cross-border Patents Promote Trade?” *Canadian Journal of Economics*.

A Sectoral distributions

Table A1: Patents' distribution across sectors, firm level, 2011

	obs.	mean	process	tech. intensity
Aircraft and Spacecraft	520	6.77	0.27	1
Animals & Animal Products	1,823	0.03	0.35	4
Chemicals & Allied Industries	5,004	2.19	0.42	2
Electrical Machinery	7,011	2.12	0.19	1
Foodstuffs	11,081	0.03	0.46	4
Footwear & Headgear	1,101	0.36	0.05	4
Metals	6,713	0.35	0.22	3
Mineral Products	658	0.68	0.39	3
Miscellaneous	7,687	0.18	0.12	4
Optical and Medical Instruments	4,816	2.11	0.20	1
Other Electrical Machinery	13,432	0.66	0.19	2
Pharmaceutical Products	638	3.11	0.41	1
Plastics and Rubber	3,534	0.76	0.25	3
Raw Hides, Skins, Leather, & Furs	1,183	0.03	0	4
Ships and Boats	326	0.33	0.03	3
Stone & Glass Products	3,214	0.45	0.36	3
Textiles	6,414	0.06	0.11	4
Transportation	7,332	1.05	0.12	2
Vegetable Products	3,322	0.06	0.68	4
Wood & Wood Products	6,356	0.14	0.27	4
Total	92,165	0.75	0.24	

Notes: The table reports the distribution of patents (applied for in the period 2002-2011) across sectors defined at the 2-digits level of the HS6 system. Columns (1)–(3) report: number of firms, average number of patents per firm, average share of process-related patents per firm. The last column classifies sectors according to their technology intensity (OECD definition): 1=High-tech; 2=Medium-high-tech; 3=Medium-low-tech; 4=Low-tech.

Table A2: Patents' distribution across sectors, firm-product level, 2011

	obs.	mean	weight	process	tech. intensity
Aircraft and Spacecraft	1,826	2.85	1.19	0.16	1
Animals & Animal Products	11,424	0.03	0.02	0.61	4
Chemicals & Allied Industries	48,123	2.58	0.86	0.44	2
Electrical Machinery	69,885	1.17	0.39	0.16	1
Foodstuffs	31,762	0.04	0.01	0.50	4
Footwear & Headgear	10,796	0.18	0.13	0.01	4
Metals	69,102	0.57	0.12	0.22	3
Mineral Products	4,976	0.75	0.33	0.49	3
Miscellaneous	35,125	0.25	0.13	0.06	4
Optical and Medical Instruments	36,784	1.23	0.45	0.22	1
Other Electrical Machinery	95,930	0.66	0.20	0.17	2
Pharmaceutical Products	3,254	4.10	1.50	0.39	1
Plastics and Rubber	36,257	0.78	0.24	0.28	3
Raw Hides, Skins, Leather, & Furs	12,797	0.040	0.02	0.01	4
Ships and Boats	771	0.57	0.38	0.08	3
Stone & Glass Products	19,920	0.78	0.20	0.33	3
Textiles	69,707	0.08	0.02	0.14	4
Transportation	23,227	1.98	0.42	0.10	2
Vegetable Products	21,870	0.08	0.02	0.66	4
Wood & Wood Products	39,143	0.22	0.05	0.17	4
Total	642,679	0.77	0.24	0.24	

Notes: The table reports the distribution of patents (applied for in the period 2002-2011) across sectors defined at the 2-digits level of the HS6 system. Columns (1)–(4) report: number of firms-products, average number of patents per firm-product, average weighted number of patents per firm-product, average share of process-related patents per firm-product. The last column classifies sectors according to their technology intensity (OECD definition): 1=High-tech; 2=Medium-high-tech; 3=Medium-low-tech; 4=Low-tech.