Do Cross-border Patents Promote Trade?

Brunel, Claire and Zylkin, Thomas

American University, University of Richmond

11 March 2021

Do Cross-border Patents Promote Trade?*

Claire Brunel[†]
American University

Thomas Zylkin[‡]
University of Richmond

March 11, 2021

Abstract

While we would expect that cross-border patents are used to protect a technology that is made available in another country, that technology could either be produced locally or imported. International patent filings could therefore be either complements or substitutes to international trade. This study combines data on patenting and trade for 149 countries and 249 industries between 1974 and 2006 with a "three-way" panel data model that addresses several biases emphasized in the trade literature in order to provide a systematic analysis of how bilateral trade responds to cross-border patent filings. We find that cross-border patents have a positive (complementary) overall effect on the patent-filing country's exports to the patent-granting country and no effect overall on imports flowing in the opposite direction. These effects vary substantially across industry groups, with patents promoting significantly more export growth in industries with a high demand elasticity and in industries that are relatively more downstream in supply chains. We also find that patents, once obtained, are associated with increased trade even in jurisdictions with weak intellectual property regimes.

JEL Classification Codes: F10; F13; F14; O34; K33

Keywords: Globalization; Intellectual Property Rights; Innovation.

*Thomas Zylkin is grateful for research support from the NUS Strategic Research Grant (WBS: R-109-000-183-646) awarded to the Global Production Networks Centre (GPN@NUS) for the project titled "Global Production Networks, Global Value Chains, and East Asian Development". We also thank Hiroyuki Kasahara, two anonymous referees, Eric Bond, David Kuenzel, Keith Maskus, and seminar participants at the International Monetary Fund, the University of Oslo, the University of Virginia, The Midwest International Economics Meetings, the Southern Economic Association Meetings, and the Western Economic Association Meetings.

[†]Email: brunel@american.edu

[‡]E-mail: tzylkin@richmond.edu

1 Introduction

As technological advancements have been a driving force behind economic growth and globalization, flows of international trade and international patents have mirrored one another in ways that hardly seem coincidental. Both have grown at rapid paces relative to world GDP since the 1970s (Fig. 1). And non-OECD countries have become substantially more involved both as recipients and originators of new patents over the last several decades (Fig. 2), just as they have also become more involved in trade in high-tech manufactured goods over this same time frame. In policy circles, the lack of patenting rights is widely cited as an important barrier to trade, and provisions on patent protections have become increasingly central to international agreements intended to promote and regulate trade (Maskus and Penubarti, 1995; McCalman, 2001; Limão, 2007; Saggi, 2016).

And yet, despite these suggestive anecdotes, much still remains unknown about how and to what extent patent protections influence trade. Through the lens of standard theory, we would generally expect that firms file patents in a foreign market to protect a technology inherent in a product they intend to sell in that market. But, empirically, we do not have evidence of how effective patents are in this regard, or when, nor can we confidently say whether the patents we observe in the data truly reflect the intent to export as opposed to other motivations, such as "strategic patenting", which aims instead to establish assets for use in potential future litigation. Indeed, when considered carefully, theory does not give unambiguous guidance on whether trade and patenting should be related positively or negatively to one another. To the extent that patent protection incentivizes technology diffusion, for example, we might expect an increase in patenting activity to coincide with reduced exports. Firms could file patents in a destination market with the intent of producing there directly, by-passing exports and instead increasing alternate modes of foreign sales such as FDI or licensing.[1] The relationship between trade and patenting could also depend on a number of factors that vary across industries and/or destination markets. For example, the same patent could have a stronger influence on trade in an market with strong protections for intellectual property rights (IPR) than in a market with only weak IPR protections. And trade and patenting could be complements in some industries and substitutes in others.

Moreover, though a vibrant empirical literature has investigated how trade responds to national patent systems, a country-level indicator, only a few studies have examined how trade is affected by actual cross-border patents, which vary by origin, destination, and the specific technology they rep-

[1] Other reasons why the relationship and trade might be theoretically ambiguous are discussed in the following section.

resent. Notably, none of these studies has investigated the trends in patenting and trade mentioned above in a panel setting with a large number of countries and for a large number of disaggregated industries. This is a conspicuous gap in the literature given the possibility that patents may be important for explaining the evolution of industry-level trade patterns over time as well as for understanding the importance of national patent systems at a more granular level. In general, the extent to which cross-border patenting has promoted trade or impeded it is a largely open question with implications for both theory and policy.

To deliver answers, we assemble an extensive database of industry-level trade flows and patent filings for 249 disaggregated industries and 149 countries over the 32 year period 1974-2006. We then implement a rigorous industry-level gravity estimation methodology that specifically assesses whether increases in bilateral patenting in a given industry are followed by increases in trade in that same industry. By pooling across the 249 industries in our study, we are able to construct econometric averages of the overall effects that patents have had on each direction of trade as well as investigate whether patents have had heterogeneous effects across different types of industries and/or across different types of markets. We are also able to exploit the panel structure of our data to demonstrate the robustness of each of our findings to natural concerns about reverse causality between trade and patenting.

By bringing together patent and trade data in this way, we are able to unpack several new empirical facts about how patents are used in conjunction with trade. For our full, unconditional sample, we find an overall complementary relationship between patents and trade: increased bilateral patenting on average tends to stimulate increased bilateral exports flowing in the same direction, with no significant effect seen for imports flowing in the opposite direction. The effect is economically significant: the difference between a "high value" flow of patents and a "low value" flow of patents—lying respectively one log-deviation above and below the log-mean of the distribution—is an average increase in exports of 8.87%, which is more than half of the average effect we find for a free trade agreement (FTA) on industry-level trade.[2]

Digging deeper, we also investigate several sources of heterogeneity in the effects of patents that are of natural interest. For example, we examine whether patents have stronger effects on trade in more differentiated industries where monopoly rights over a product's distinguishing characteristics should be more valuable. Contrary to what one might expect, we actually find overwhelming evidence for the opposite: patents have larger effects in less differentiated industries where patents are relatively less valuable for protecting monopoly rights but instead serve to protect cost and

[2]These results, explained in Section 5, are evaluated at the peak of a patent's life. We calculate patent stocks based on value-weighted patent flows, weighted by family size.

quality advantages over producers of competing products. We also test whether patents have larger trade-promoting effects when filed in countries with stronger IPR institutions. Interestingly, we do not find any overall relationship between the strength of the patent-granting country's IPR regime and the effect that patents filed in that country have on trade, indicating that patents can affect trade independently of the IPR regime.

Studying the co-evolution of patents and trade over time presents a number of challenges. This is in large part because patents are complex objects to work with empirically: they are at once "lumpy" (occurring at irregular points in time, but intended to have effects for long durations) as well as highly specific to the technology they derive from and destination market they are filed in. Furthermore, as is well documented in the trade literature, accounting for the many different dimensions of heterogeneity that naturally arise in a multilateral trade context quickly becomes computationally difficult even for estimations with aggregate trade flows (see, e.g., Glick and Rose, 2016; Larch, Wanner, Yotov, and Zylkin, 2019).

To resolve these various challenges, we draw on several recent methodological innovations that naturally complement our focus on cross-border patents. First, because of the natural lumpiness of patent flows, we construct bilateral, industry-specific stock values reflecting the gradual obsolescence of old patents filed in that market in addition to the lag in diffusion for new flows (Popp, 2002, 2003). Our analysis therefore accounts for the lasting effects of previously filed patents. Second, to address the issue that some patents are more valuable than others, we weight patents by their family size. Third, our econometric specification accounts for the industry specificity of patents by effectively pooling across 249 disaggregated industry-level gravity estimations, each with its own set of exporter-industry-time, importer-industry-time, and country pair-industry fixed effects. This "three way" fixed effects structure accounts for numerous sources of bias that have been shown to be important in the literature on trade agreements and that we similarly show to be important for assessing the effects of patents.[3] Finally, we capitalize on some recent innovations by Correia, Guimarães, and Zylkin (2019a,b) to estimate our preferred model using Poisson Psuedo-maximum Likelihood (PPML). The use of PPML not only ensures our estimates are consistent, but also ensures they are not subject to any aggregation bias from pooling across disaggregated industries (French, 2019.)

As mentioned above, the prior literature examining the empirical relationship between patenting

[3]The "three way" gravity model for panel data was first used by Baier and Bergstrand (2007) in their study of the effects of free trade agreements on trade. Pooled gravity estimators are relatively new to the literature. French (2019) describes in detail the advantages of pooled estimators such as the one described in this paper. In other recent work, Schmidt and Steingress (2018) and Kinzius, Sandkamp, and Yalcin (2018) use similar specifications to study the effects of product standard harmonization and non-tariff barriers, respectively.

and trade is limited in number and in scope. Most papers in this area have been constrained by data to examining a single flow (either imports or exports), a single bilateral relationship, a single country's trade relationships, and/or a single industry (Anderton, 1999; Greenhalgh, 1990; Antimiani and Costantini, 2013; Brunel, 2019; Chen, 2013; Chalioti, Drivas, and Katsimi, 2017) The one exception is Palangkaraya, Jensen, and Webster (2017), who also use a PPML gravity specification to examine the effects of successful patent applications on bilateral trade. Using a data set with 13 host markets, 28 industries, and a large number of exporting countries, they find, as we do, that patents generally complement bilateral exports.[4]

Our investigation goes beyond that of Palangkaraya, Jensen, and Webster (2017) in several ways. First, we investigate a wider range of effects, allowing patents to affect both directions of trade as well as allowing for heterogeneity across different industries and across different markets. Second, another key difference is that Palangkaraya, Jensen, and Webster (2017) focus on cross-sectional variation in patent application outcomes for the same invention across different patent offices with similar application processes, whereas we focus on longitudinal changes in trade within the same pair of markets as a result of successful patent filings. Our panel approach has the advantage of explicitly controlling for any omitted cross-sectional heterogeneity that might cause the propensity to file a patent in particular country to be correlated with the propensity to trade with that country. Third, our approach also facilitates our use of more expansive and varied data, with more years, more disaggregated industries, and (especially) a wider variety of OECD as well as non-OECD patent destinations.

In sum, we contribute to the literature by greatly expanding the geographical and industry-level scope of existing analyses and by using a large panel dataset which matches bilateral trade flows to bilateral patent flows at a highly detailed level of disaggregation. We notably exploit the panel dimension of this data to control for a large set of unobservable characteristics across both country-pairs and industries and simultaneously address numerous well-known sources of bias that have been articulated in the empirical trade literature and that generally cannot be addressed using firm-level data. Because we are able to surmount these many complex data and estimation challenges, our work is the first to pursue a comprehensive longitudinal analysis of the effects of cross-border patents on trade as well as the first to isolate some key sources of heterogeneity in this relationship.

[4] To be more precise, Palangkaraya, Jensen, and Webster (2017) find that the presence of either an apparent bias in the patent application process or a "blocking prior art" in the importing country deters exports.

2 Related Literatures

While only a limited amount of analysis has been devoted to the empirical relationship between patents and trade, an active related literature has looked at the "gravity"-like relationships that determine the geographical distribution of knowledge capital and patent flows. These include Keller and Yeaple (2013)'s influential work on the "gravity of knowledge" as well as Figueiredo, Guimarães, and Woodward (2015)'s study of the industry- and location-specific clustering of patent citations. Earlier contributions in this literature also include Robbins (2006) and Scherngell and Hu (2011). Keller (2004) reviews the earlier literature on technology diffusion across countries, with trade being one of the channels.

The theoretical literature on patenting and trade, meanwhile, is very rich. As also observed in Palangkaraya, Jensen, and Webster (2017), this literature does not offer a consensus view on whether patents and trade should be complements or substitutes. For example, in Lai and Qiu (2003) and Grossman and Lai (2002), stronger patenting rights in the global South could decrease trade by stimulating innovation-intensive production in the South and thereby reducing North-South comparative advantage. Yang and Maskus (2009) show that, in a model where innovating North firms can serve the South market via exports, FDI, or licensing, stronger patent rights in the South can have ambiguous effects on both directions of trade, depending on the North firms' optimal mode of foreign sales and on their strategic incentives to transfer know-how to the South.[5] And Bond and Saggi (2014) show that patent regimes that include the threat of compulsory licensing can induce an innovating foreign firm to serve a market directly rather than voluntarily license its patented product.[6]

In addition, it is likely that patent flows and trade flows we capture in our disaggregated study in many cases reflect concurrent decisions made by the same firms (or, at least, by related firms). As Saggi (2016) notes, "in a typical year", over 80% of international patent royalty payments are made between parent firms and their subsidiaries. Our work thus also broadly relates to a recent literature that studies the role that property rights play in determining how trade in intermediate inputs within a firm's supply chain is organized among related entities in different countries. Representative works in this area include Nunn and Trefler (2013), Antràs and Chor (2013), and Alfaro, Antràs, Chor, and Conconi (2015). Also related is Blonigen (2001), who finds that FDI flows and trade flows are substitutes for one another in final goods industries, but complement one another in intermediate

[5] This type of theoretical ambiguity is also discussed in McCalman (2001).

[6] Of course, the recent debates surrounding TRIPS and other similar international issues have stimulated a much wider literature studying these and other related topics; for some recent surveys, see Saggi (2016) and Maskus (2018).

goods industries. For our part, we find that cross-border patents complement exports in both intermediate goods and final goods; however, the effect for final goods is driven solely by capital goods as opposed to consumption goods. Looking at the position of each industry in the supply chain in a more textured way, we also find larger effects for industries that are relatively downstream (i.e., closer to final demand) as opposed those that are more upstream.

Finally, while the empirical literature on patents and trade is sparse, there is a rather extensive body of work relating the strength of national patent institutions to trade. See, for example, Maskus and Penubarti (1995), Briggs and Park (2014), Boring (2015), Awokuse and Yin (2010), and Ivus (2010). Saggi (2002) also provides a survey of the earlier work in this area. As also emphasized in Palangkaraya, Jensen, and Webster (2017), an empirical design based on individual patents rather than on national institutions is a natural complement to this line of research. Interestingly, we generally do not find that the relationship between patents and trade varies significantly with the strength of the host country's IPR regime, indicating that bilateral patent filings may reflect unique information about the protection of intellectual property that is not captured by national indicators.

3 Estimating the Effects of Patents on Trade

As noted in our introduction, analyzing the relationship between bilateral patents and bilateral trade requires an ambitious empirical framework. We observe a many-country world where cross-border patents and trade flows each vary by origin, destination, industry, and time—to be denoted in the analysis by i, j, k, and t, respectively. We aim to identify the effects of patents on trade by exploiting within origin-destination-industry (or "pair-industry") variation in both patents and trade. Part of what makes this task complex is the inter-connected nature of global trade. As has been well known at least since Anderson and van Wincoop (2003), correctly assessing the determinants of bilateral trade between any two countries requires first recognizing the potentially confounding role played by multilateral factors such as market size and geography with respect to third countries.[7] A further challenge is that patents are by their nature specific to the technologies they derive from. It is not necessarily the case, for example, that a patent filed for the production of aerospace equipment will affect trade in cosmetics products. We therefore require an empirical model that simultaneously accounts for both the multilateral complexity of trade and the industry-

[7]In Anderson and van Wincoop (2003) (and much of the literature following Anderson and van Wincoop 2003), each country's trade is shown to be function of the "multilateral resistance" it faces in the world trade network, such that standard controls such as GDP and other observables used in the traditional gravity literature are wholly inadequate for correctly isolating the determinants of bilateral trade linkages. Furthermore, because bilateral trade relationships and market sizes change over time, changes in multilateral resistance are themselves an important source of changes in trade over time which should be accounted for in any panel study (Baldwin and Taglioni, 2007).

level specificity of a patent, as well as the many further dimensions of complexity that arise when our object of interest simultaneously varies by industry, origin, destination, and time.

Taking all of the above concerns onboard, our proposed estimation strategy takes the form of a "three way" panel gravity model, similar to the models currently used in the structural gravity literature to study the effects of trade agreements and other bilateral policy arrangements (Baier and Bergstrand, 2007; Glick and Rose, 2016; Larch, Wanner, Yotov, and Zylkin, 2019).[8] We then interact the standard set of origin-time, destination-time, and country-pair fixed effects that have become standard in this literature with the industry dimension. That is to say, (a log-linear version of) our preferred empirical model can be expressed as

$$\ln X_{ijkt} = \delta_{ikt} + \psi_{jkt} + \eta_{ijk} + \beta \mathbf{PATENT}_{ijkt} + \delta \mathbf{Z}_{ijkt} + \varepsilon_{ijkt}, \tag{1}$$

where $\ln X_{ijkt}$ denotes log trade flows and $\mathbf{PATENT}_{ijkt}$ are our industry-level bilateral patent variables (to be described shortly). The fixed effects used are exporter-industry-time (δ_{ikt}), importer-industry-time (ψ_{jkt}), and exporter-importer-industry or "pair-industry" (η_{ijk}). $\mathbf{Z}_{ijkt}$ is a set of other time-varying controls not absorbed by η_{ijk} (such as the presence of a free trade agreement), and ε_{ijkt} is our estimation error. The loglinear model in (1) is not our preferred specification, however. As famously pointed out by Santos Silva and Tenreyro (2006), the log-transformation of the dependent variable to construct (1) is not innocuous, since the log-OLS moment condition $E[\ln X_{ijkt} - \widehat{\ln X}_{ijkt}|\cdot] = 0$ does not also imply that $E[X_{ijkt} - \widehat{X}_{ijkt}|\cdot] = 0$. Estimates of (1) should therefore be regarded as generally being inconsistent.[9] Accordingly, capitalizing on recent computational innovations by Correia, Guimarães, and Zylkin (2019a,b) that allow for all the necessary dimensions of the fixed effects in a PPML estimation, our fully specified empirical design involves estimating the above log-linear model for trade flows in levels using Poisson Pseudo-maximum Likelihood (PPML) as the underlying estimator. That is, we estimate

$$X_{ijkt} = \exp\left[\delta_{ikt} + \psi_{jkt} + \eta_{ijk} + \beta \mathbf{PATENT}_{ijkt} + \delta \mathbf{Z}_{ijkt}\right] + \nu_{ijkt}. \tag{2}$$

It can be shown that the above PPML estimator will provide consistent estimates of β and δ under standard assumptions about the unobserved heterogeneity contained in ν_{ijkt}.[10] Because it is a

[8] For more on gravity estimation, see Head and Mayer (2014) and Yotov, Piermartini, Monteiro, and Larch (2016). Yotov, Piermartini, Monteiro, and Larch (2016) provide in-depth discussion of econometric best practices related to three-way gravity models in particular.

[9] For these types of estimations, the data generating process for X_{ijt} is typically assumed to be a constant-elasticity (exponential) model as shown in (2). Under this assumption, estimating a log-linearized specification only leads to consistent estimates under the very special case where the error term ε_{ijkt} is homoscedastic, a requirement that our data strongly rejects.

[10] Weidner and Zylkin (2020) show that the three-way FE-PPML gravity estimator is consistent so long as the

pooled estimator that pools across disaggregated industries, it also benefits from the aggregation bias-minimizing properties of PPML described in French (2019).[11]

Setting aside the issues that have been raised regarding loglinear models, the rich set of fixed effects in place in (2) addresses many of the other leading econometric challenges that are commonly cited in the literature on gravity estimation. These include the need to account for multilateral inter-dependencies (*industry-level* inter-dependencies in our case, addressed by δ_{ikt} and ψ_{jkt}) as well as unobservable pair characteristics that coincide with the propensity to trade (captured by η_{ijk}), such as the trade costs currently in place for that industry-pair.[12] More generally, we also note an appealing logic to the identification based on these fixed effects that speaks to our objectives. Our question concerns how *bilateral* patenting directly affects *bilateral* trade between a given pair of countries in a given industry. By completely stripping away industry-country-specific fluctuations in trade and patenting, the country-industry-time fixed effects δ_{ikt} and ψ_{jkt} allow us to specifically isolate the pair-industry dimension of the data. The pair-industry fixed effect η_{ijk} then ensures that identification of β and δ is driven solely through observed time-variation within pairs, as in standard panel data analysis.[13]

In more intuitive terms, these fixed effects are intended to ensure that β picks up the value of protecting an invention in a particular market, rather than solely reflecting the effect of the innovation associated with that patent. That is to say, to the extent there are "direct" effects of innovation on trade irrespective of patent filing—by broadly increasing the productivity of the inventing firm, for example—we would expect these effects to influence that firm's trade with all destinations, not just the ones where it owns a patent. These direct effects would therefore be controlled for by the exporter-industry-time fixed effect δ_{ikt}. To be sure, some patents could protect

conditional mean is correctly specified and the regressors of interest are exogenous after conditioning on the fixed effects. The latter requirement is an important consideration in any panel study with a potentially endogenous regressor. To deal with this concern, we follow the advice of Wooldridge (2002) and Baier and Bergstrand (2007) in introducing "lead" variables that serve as a placebo test against the possibility that changes in patents are responding to changes in trade. This placebo approach is described further below.

[11] French (2019) shows that pooled PPML estimation of common gravity variables such as log distance can be rationalized as aggregating the information from industry-by-industry estimates in a way that preserves the same aggregate trade volume in expectation. In our setting, a similar property applies that allows us to interpret β as reflecting the partial effect of a hypothetical uniform increase in the patent stock across all industries on aggregate bilateral trade. Because $\sum_k E[\ln X_{ijkt}|\cdot] \neq \ln \sum_k E[X_{ijkt}|\cdot]$, pooled log-OLS does not have this property and is subject to aggregation bias.

[12] The observation that estimates of time-varying gravity variables tend to be biased in the absence of pair fixed effects is generally credited to Baier and Bergstrand (2007).

[13] In another recent study that uses a similar methodology, Schmidt and Steingress (2018) observe that it is also possible to include a fourth fixed effect indexed by ijt that accounts for common changes in trade across all industries within each pair. We prefer to stick with the three-way model since it conforms to the intuition of pooling our results across 249 individual industry-level gravity estimations. As we document in Section 5.3, our results are not materially affected by adding this other fixed effect; in general, we find the ikt, jkt, and ijk fixed effects are much more important for inference by comparison.

inventions aimed very specifically at a single market, in which case the direct innovation effect of the patent would not be controlled for by the fixed effects. However, in that case, we would expect that patents should have larger effects on trade in more differentiated industries where other varieties are perceived as less substitutable. As we will see in Section 5, our findings do not support that hypothesis, indicating that country-specific inventions are not driving our results. Nonetheless, we should be clear that our preferred interpretation of β depends on the assumption that, when patenting and trade both increase with one particular market relative to other markets, the increase in trade reflects the value of the increase in patent protection, which itself is a function of the value of innovations being protected. As we discuss in more detail in the next section, this logic informs our decision to weight patents by their family size, which we take to be a measure of the potential export value of an innovation.

As discussed above, theory suggests numerous reasons why the effect of patenting and trade could be ambiguous. However, to fix ideas, a reasonable prior to start with is Palangkaraya, Jensen, and Webster (2017)'s supposition that patents have a role to play in protecting the cost and/or quality advantages that exporters enjoy in foreign markets. Importantly, this view not only suggests a baseline hypothesis—that patents promote bilateral exports—but also guides our more expansive investigation of the possibly heterogeneous effects of patents on trade in two main ways. First, in contrast to the view that patents primarily protect new varieties, this perspective implies that the value of protecting a superior technology should actually be greater in less differentiated industries with a higher elasticity of substitution between products. In these industries, the value of protecting a cost and/or quality advantage is greater because the replication of the patent-holder's technology represents a more direct threat to its profitability and, similarly, because the assertion of exclusive patent rights over a disputed technology by a foreign firm should have a relatively larger impact on trade. By either of these rationales, patents should be more strongly complementary to exports in high-elasticity industries than in low-elasticity industries.[14]

Second, we will also consider interactions between our patents variables and several other variables that may reflect either costs of trading or the cost of obtaining and/or enforcing a patent. These indicators include standard gravity relationships, such as the log of distance and the sharing of a common legal system or common language, as well as a set of dummy variables reflecting the strength of the patent-granting countries patent regime ("medium" or "high", with "low" as an

[14]As noted above, we are taking the view that the patented technologies we observe are inputs that either lower the cost of producing a product or enhance its quality. A competing view would be that the patents we observe protect the distinguishing characteristics that differentiate the product in the eyes of end-users. In that case, we might expect patents to have more pronounced effects in less differentiated industries. This is not what we find.

omitted category). The inclusion of IPR levels helps us link our study to the wider literature on IPR and trade: while stronger patent protection regimes have often been found to be associated with increased trade (Maskus and Penubarti, 1995; Ivus, 2010), it is interesting to wonder whether the patent filings themselves are the main vehicle for this effect.

One remaining empirical challenge lies in the potential endogeneity of patents and trade. In general, estimates of the elasticity of trade with respect to patents will be biased if (partner-specific) changes in export demand or import supply affect the number of inventions that are subsequently patented in the partner country. Import competition, for example, has been shown by some recent papers to stifle innovation in the home country in some cases (Autor, Dorn, Hanson, Pisano, and Shu, 2016; Batrakova and Dechezleprêtre, 2013) and boost it in others (Bloom, Draca, and Reenen, 2011). Alternatively, a high level of trade in technology-intensive goods could point to a market that is highly active and innovative, or, on the contrary, to one where existing technologies are sufficient and patenting has plateaued. Moreover, multinational corporations could be influenced by strong existing trade relationships in their decisions about where to locate production. A strong bilateral trade relationship could thus be a precursor for a multinational corporation filing a patent to setup production capacity abroad.

To address these concerns, we would ideally want an instrument for bilateral patenting that is time-varying, industry- and pair-specific, and is strongly correlated with patent growth, but exogenous to changes in bilateral trade. A good instrument meeting these criteria is hard to find in this context, especially given the number of interactions we would ultimately need to instrument for. Consequently, our main strategy for assessing the direction of causality is to perform a set of standard placebo tests to determine whether the changes in trade identifying our results occur on or after the observed change in the patent stock (see Wooldridge, 2002; Baier and Bergstrand, 2007.) By constructing appropriate "lead" terms for our stock variables and their associated interactions, corresponding to the values these variables will take in 4 years time, we can effectively test whether our results may actually be picking up reverse causality flowing from changes in trade to changes in bilateral patenting decisions. We also experiment with controlling for a range of different factors that may have independently affected bilateral trade flows during the period—such as changes in the effect of geographic distance, for instance—and allow these factors to vary with the R&D intensity of the industry. As we will see, our results are generally robust to these types of tests. That said, one of the limitations of our design is that we cannot completely rule out the possibility that firms make patenting decisions today based on their expectations about future trade conditions and/or in coordination with other investments that facilitate trade. While we do offer some additional sensitivity checks that speak to the possible role of trade expectations, our findings should be interpreted with these caveats in mind.

4 Data Construction

4.1 Trade Flows and Other Gravity Variables

The trade dataset consists of bilateral trade flows between 149 developing and developed countries for 249 industries at 3-digit level based on the Standard International Trade Classification (SITC) revision 2.[15] We construct this data by combining the standard 1962-2000 NBER-UN World Trade Flows dataset of Feenstra, Lipsey, Deng, Ma, and Mo (2005) with additional data for 2000-2006 obtained from UN COMTRADE. We obtain a panel dataset covering both exports and imports between these countries from 1974-2006, using every 4 years. Both the start and the end of the time period are dictated by the patent data detailed in the next subsection. The use of every 4 years is following the recommendations of Cheng and Wall (2005) and also helps with facilitating computation.[16]

In order to investigate which types of industries drive which results, we find it useful to examine industry subgroups with similar factor intensities or usage patterns. Industry-level indicators for skill intensity, R&D intensity, and different measures of capital intensity are taken from Nunn and Trefler (2013). These data are originally provided at the Harmonized System (HS) 6 digit level, which we map to our own SITC-based industry codes using a concordance from the World Integrated Trade Systems. For data on the usage of each industry, we obtain UN Broad Economic Category (BEC) codes at the HS6 level from UN COMTRADE and concord them to SITC using the same concordance. A measure for the "upstreamness" of an HS6 category's position in the supply chain is similarly available from Antràs, Chor, Fally, and Hillberry (2012). Multiple HS codes can map to the same SITC-based industry; thus, we assign shares based on world exports of each HS category to construct usage shares for each SITC industry. We often find it useful to use these indicators to partition industries into different subcategories. For example, we divide industries into high-, medium-, and low-tech industries by establishing cutoffs at the 67th percentile and 33rd percentile of the R&D intensity distribution.

[15] Although the SITC classification technically refers to commodities, we use the term "industry" to describe our 3-digit categories since they include product groups such as motorcycles or office and stationary supplies.

[16] Specifically, Cheng and Wall (2005) note that "fixed-effects estimations are sometimes criticized when applied to data pooled over consecutive years on the grounds that dependent and independent variables cannot fully adjust in a single year's time." In our context, the use of intervals that span multiple years allows us to more cleanly observe the response of trade to changes in the patent stock. Our assumption that the effects of patents do not have their peak impact until several years after filing also assists in warding against this critique.

In addition, a particularly appealing variable for our analysis is the demand elasticity for each industry, which reflects the degree of price competition and product differentiation. We might expect patent protections to play an important role in industries characterized by a high demand elasticity, indicating a relatively low degree of product differentiation, since it is in these industries that rival innovation and/or imitation most directly threatens the market share of an innovating firm. To construct our elasticity variables, we first obtain detailed estimates of import demand elasticities at the 3 digit SITC level from Broda and Weinstein (2006). Then, to acknowledge that the point estimates of these industry-specific elasticities are measured with error, we construct an indicator that is equal to 1 if the estimated elasticity is above the median estimate across all 249 3-digit industries in our data.[17] We focus on the results for patents in "high" versus "low" elasticity industries to test our hypotheses with respect to the demand elasticity.

Finally, for data on free trade agreements (FTAs), we consult the NSF-Kellogg database on Economic Integration Agreements created by Scott Baier and Jeff Bergstrand.[18]

4.2 Bilateral Patent Flows and Stocks

The patent data are drawn from the 2013 version of the PATSTAT database, which includes all patent applications published at the national patent offices of the 149 countries in our sample as well as the European Patent Office (EPO) and from which we pull all patents filed through 2006.[19] The extraction includes all patents filed into one of our 149 countries in all industries, regardless of the origin of the inventor, and all patents filed anywhere in the world by an inventor residing in one of the 149 countries in our sample.

Since a patent grants protection for a technology only in the country where it is filed, inventors have to file in all countries in which they desire protection. In our data, there is one exception: patents filed at the European Patent Office (EPO), which covers a number of European countries. A patent filed at the EPO would only be applicable in countries designated on the patent application, but the designation data is not available in PATSTAT. Since each additional country designation

[17] The median elasticity in our data is about 2.53. Broda and Weinstein (2006) also provide their own indicators for "high", "medium", and "low" elasticities. Our results are very similar if we use these indicators instead, though we generally find the latter two categories have statistically similar effects to one another.

[18] This database is available for download on Jeff Bergstrand's website: https://www3.nd.edu/~jbergstr/.

[19] The PATSTAT 2013 database contains patents through 2012 but we truncate the data to 2006 for two reasons. First, there can be delays between the date of patent filing and the date of publication—up to 18 month for patents filed at the EPO (Lotti and Marin, 2013). Since we only observe published patents, counts in the years running up to 2012 suffer from missing observations (Hall, Jaffe, and Trajtenberg, 2001). Second, more recent patents have not had as much time to be transferred to other countries, meaning that they will have lower family sizes. As explained in this section, we weight patents by family size to account for value so patent counts in the few years preceding 2012 would appear to have low values. For both both of these reasons, we do not include the year 2010 in our analysis, though the results are qualitatively similar if we include 2010.

incurred a fee, inventors have an incentive to be selective in the countries they designate. However, designation data collected for a subset of patents in Brunel (2019) shows that the list of designated countries often covered more than half of the EPO members and often all the largest EPO member economies. Therefore, throughout our time period, we assign an EPO patent to all EPO member countries in that year.[20] Since this is not exact, we also provide robustness checks without EPO patents.

Patents are classified by the end-use of the technology based on the International Patent Classification (IPC), which can be matched to trade data using concordances developed in Lybbert and Zolas (2014).[21] We match the IPC data to SITC rev. 2 at the 3-digit level.[22]

Though patents have the advantage of being widely available and classified in a way that is linkable to trade data, several drawbacks of the data are worth discussing. First, given that filing a patent implies disclosing information on the technology, some inventors might chose to refrain from patenting their products. However, research shows that few inventors opt to protect the secrecy of their inventions, especially when the invention is more valuable (Dernis and Khan, 2004).

Second, different patents can have widely different values. A patent grants the inventor the exclusive right to use the underlying technology. Some patents are used extensively by either protecting large amounts of production, being licensed to a wide number of users, or being filed in many different countries. On the other hand, some published patents are not granted or not enforced. In accordance with the literature, we account for the value of a patent weighing each patent by its family size, i.e., the number of countries in which that invention is patented. The logic behind this weighting system is that less globally marketable inventions that are not very valuable for trade will most likely not seek wide geographical coverage. It is especially important to keep track of the value since our analysis covers a wide range of different sectors, and value is likely to differ across industries.[23]

One possible concern with weighting by family size in this way is that a more export-oriented firm may be more likely to file its patents in more destinations, thereby inducing a correlation between

[20]The list of EPO members and date of accession can be found at https://www.epo.org/about-us/foundation/member-states/date.html

[21]The concordances are available from Professor Lybbert's website: https://are.ucdavis.edu/people/faculty/travis-lybbert/research/concordances-patents-and-trademarks/

[22]Some IPC codes match to multiple SITC categories, so the flows can be non-integer values.

[23]To be sure, there are other ways of accounting for the value of patents, notably through citations data. Both have been proven to be significantly correlated with patent value and with each other. However, as discussed in De Rassenfosse (2013), the average family size is the "most internationally comparable indicator" since citation practices and the availability of citations data differ significantly across different patenting jurisdictions. In Section 5.3, we show that results computed using citation-weighted patents do not exhibit a statistically significant relationship between patenting and trade.

trade and patent value. A similar concern would also apply to value-weighting more generally, since we might expect export-oriented firms to have more valuable innovations than other firms on average. However, while it is true that this reasoning does suggest that our patent variables may be correlated with trade in the cross-section, the potential threat to identification in our panel data context is more subtle: since the country pair-industry fixed effects absorb any such correlation in the cross-section, one would need these firms to be systematically increasing their patent positions in markets they intend to export to by mere coincidence or for reasons that have nothing to do with export sales in order for there to be a bias. This perspective is consistent with the assumptions needed for identification we have described in the previous section.

We construct two measures of patents using this data. The first is a count of patents filed in each industry k in country j by an inventor residing in country i in each year weighted by family size (PAT_{ijkt}), which represents the flow of new patents by country-pair and industry in any given year. However, trade flows between two countries are not only affected by the new flows of patents, but could also be influenced by existing patents. Therefore, we also create a stock of patents in each year. In aggregating patents, we account for a diffusion time, since a brand new patent might not be known and used immediately, and a decay rate as technologies lose relevance over time. The use of the patent should thus theoretically increase and then decrease as the technology eventually becomes outdated. Therefore, as in Popp (2002), we aggregate bilateral patent stocks in the following way:

$$STOCK_{ijkt} \equiv \sum_{s=0}^{\infty} e^{-\omega_1 s}(1 - e^{-\omega_2(1+s)}) \times PAT_{ijk(t-s)} \tag{3}$$

where ω_1 is the decay rate set to 0.1, and ω_2 corresponds to the rate of diffusion and set at 0.25, following Popp (2003).[24] In line with the literature, which finds that R&D capital peaks between 3 and 5 years (Griliches, 1995), these two parameters imply the impact of a patent on trade should peak at about 4 years on average,[25] which should help us to decouple the timing of the decision to patent from the timing of trade, even before we consider any placebos. Because the life of a patent has been shown to vary by industry and originating country (Pakes and Simpson, 1989; Schankerman, 1998), we also experiment with alternative values for ω_1 and ω_2 that allow for faster and slower rates of diffusion and decay in our sensitivity analysis.

We use the subscript $ijkt$ to denote the value of the patent stock in industry k owned by filing country i in destination j at time t. Therefore, we may consider "two way" effects of bilateral

[24]Other papers which use these same parameters include Popp, Hascic, and Medhi (2011), Lovely and Popp (2011), and Brunel (2019).

[25]We do not have information on the length of validity for each patent and that length varies across patent offices. However, given our stock formula, the bulk of a patent's contribution to the stock occurs in the first 15 years, which is less than the 20 year patent term in the United States for example.

patenting on trade: sticking with the same notation as above, "$STOCK_{ijkt}$" will denote the exporter (i)'s patent stock in the importing country (j) in industry k at time t. However, we also consider the effect of "$STOCK_{jikt}$", which will be used to assess how patents affect the flow of trade in the other direction (i.e., how j's filing a patent in country i affects its imports from i).

Table 1 presents the summary statistics for the raw patent counts (unweighted), patent counts weighted by family size, as well as for stocks. On average across all industries, time, and country pairs, there are 7 new patents filed—33 on a value-weighted basis—and the stock mean is 179. For both measures, the standard deviation indicates a large amount of variation in patenting across our dataset. Stocks and flows by country pair-industry-year can often be 0 (the median is zero or small for all rows) but value-weighted stocks can also reach a maximum value of over 500,000. Unsurprisingly, both stocks and flows are larger for patents filed in OECD countries than for patents filed in non-OECD countries by an order of magnitude. Since the distribution of stock values is severely skewed, we use a log transformation of these stocks in our analysis, which means we can only identify the effect of a stock after at least one patent has been filed. Focusing on non-zero stocks, typical values remain small: the median values of bilateral industry-specific patent stocks range from 2 for non-OECD countries to 19 for the OECD group. We also provide the industries and country-pairs with the highest weighted patent growth in our data. Interestingly, the industries for which patents have grown the fastest include a diverse mix of higher-tech and lower-tech industries, with several textile and basic metal categories appearing alongside various types of machinery. The top country-pairs mainly include patent-granting countries that we traditionally think of as highly innovative and/or fast-growing during the period: the United States, Japan, Germany, China, and Great Britain.

5 Results

Having described our methods and data, we now turn to our main analysis, in which we investigate the effects of cross-border patenting on trade. We start with a broad analysis, deriving at first a set of overall, pooled estimates of how patents affect each direction of trade across all markets and industry types. Subsequent specifications then dig deeper into the heterogeneous patterns underlying these overall relationships, allowing for heterogeneity across industries and across markets. A battery of sensitivity checks then follows.

5.1 Do Cross-border Patents Promote Trade?

Our empirical investigation begins in Table 2. Columns 1-6 of this table report PPML estimates of the effects of changes in the bilateral patent stock on both directions of trade—exports as well as imports—along with the estimated effect of an FTA, a standard time-varying control from the gravity literature. Columns 7 and 8 report some additional estimates using the patent stock as the dependent variable. All columns include exporter-industry-time and importer-industry-time, and industry-country pair fixed effects. All reported standard errors are two-way clustered by country-pair and by exporter-industry. Since patents are observed at the industry-country-pair level, this should lead to highly conservative inferences that allow for correlated changes across different industries within the same pair and across different patent destinations.

The main message from Table 2 is that the effect of patenting on exports remains positive and statistically significant across a variety of different specifications (columns 1 and 3-6), while no overall effect is found for bilateral imports returning in the opposite direction (columns 2 and 3). To focus first on our main specification in column 1, the estimated coefficient for the export effect of 0.018 can be interpreted in several different ways, depending on whether one considers the filing of a single patent, multiple patent filings, or more simply a change in the stock value. But, to establish a baseline for comparison, we note that the family-weighted flow of patents filed in a particular industry and country in a particular year is well-approximated by a log-normal distribution. We define a "high value" patent flow PAT^{high} as having a value lying one log-standard deviation above the distribution's log-mean, i.e. a high value (relative to the mean) of the total of all new patents filed in a given pair-industry-year. A "low value" patent flow is thus a value one log-standard deviation below the log-mean. Following these definitions, we obtain values of $PAT^{high} = 155.774$ and $PAT^{low} = 1.125$. Applying our estimated coefficient of 0.018 to these values along with the formula we use for the stock value in (3) suggests high value patent flows increase industry-level exports at their peak by 8.87% relative to low value patent flows.[26] For comparison, this effect represents more than half of our estimated boost to industry-level trade from participating in an FTA, which can be calculated using the estimated FTA coefficient from column 1 as $e^{0.156} - 1 = 16.88\%$. Despite our conservative clustering assumptions, both coefficient estimates are statistically significant at the $p < 0.01$ significance level. As shown in columns 2 and 3, we do not observe a similar effect for bilateral imports flowing in the opposite direction.

Our next task is to determine whether we can reasonably say these estimates reflect causality.

[26] We obtain this number using $0.018 \times [\log(e^{-0.1\times5}(1-e^{-0.25\times6}) \times PAT^{high}) - \log(e^{-0.1\times5}(1-e^{-0.25\times6}) \times PAT^{low})]$. Alternatively, if we consider "very high" and "very low" value patent flows that respectively lie at the 95th percentile and 5th percentiles of the distribution, we obtain a difference in peak effects of 14.6%.

The remaining columns in Table 2 present a series of experiments that are meant to assess some possible concerns with this interpretation. First, rather than patents affecting flows of trade, causality could flow from trade to patents; thus, in column 4 of Table 2, we use "lead" stock values to test whether changes in the patent stock in a prior year anticipate changes in exports. The lead variable for the effect of the patent stock on bilateral exports is economically and statistically near-zero, indicating our results are not being driven by changes in trade patterns that precede changes in patenting. Furthermore, the original estimates are scarcely changed.[27]

Second, despite the rich set of fixed effects included, our main specification in column 1 does not control for the possibility that the changes in trade we are observing could be reflecting general changes in trade conditions across all industries rather than specific changes occurring in the industries where patents are being filed. It could be that countries provide additional protection for intellectual property while also changing market access conditions for goods at the same time, for example. Since this would presumably affect patenting and trade across all product categories simultaneously, in column 5, we experiment with adding country pair-time (ijt) fixed effects, which absorb all changes in the overall bilateral environment for both patenting and trade. This is not our preferred specification because it also rules out spillovers flowing from patents in one industry to trade in another and because the specification in column 1 is more comparable to the industry decompositions we examine later in the paper. Nonetheless, we find that the results remain stable and robust when ijt fixed effects are added.

Third, we still may be worried that our estimates fail to control for other dynamic factors that have had industry-specific impacts on trade aside from patents and FTAs—the dissemination of the internet, for example. For this reason, we next experiment with adding standard gravity controls (log distance, common language, common legal system) interacted with a time fixed effect as well as with a set of indicators reflecting the industry's degree of R&D intensity. The addition of these controls provides a reasonable test of whether our estimates are being driven by industries with different levels of patent activity responding differently to potentially omitted dynamic factors such as changes in transportation costs and infrastructure systems, ease of communication, and contracting or other similar frictions. The results, shown in column 6, are essentially unchanged.

Finally, though we have ruled out the concern that changes in trade are preceding changes in patenting, another possibility is that firms are filing patents because they expect lower trade costs

[27] What's more, recall that our stock variables assume patents only reach their peak value 5 years after the date of filing. If anything, this assumption should lend more conservatism to this type of test by making it more likely that we observe significant lead coefficients in column 2. Also notice that the lead value for "FTA" in Table 2 is actually significant, which generally goes to show that it is by no means unlikely that we should find evidence of anticipation effects and/or pre-trends in the data.

in the near future, an argument resembling that of Aw, Roberts, and Xu (2011). While we cannot claim to perfectly capture firms' expectations about future trade costs, we can at least offer a few reasonable checks in this direction. Columns 7 and 8 of Table 2 respectively regress the patent stock, as the dependent variable, on lead variables for the presence of an FTA and for log trade. The lead FTA variable plausibly captures the effect of an expected reduction of trade costs, whereas the future trade variable may convey less easily observable information about expected future trade. In both cases, we find no statistically significant relationship with the patent stock, though the effect of the lead FTA variable is positive and noisily estimated. The fact that the addition of the FTA lead variable in our earlier column 4 does not alter our main result thus adds further assurance. Altogether, the results shown in columns 4-8 of Table 2 lend significant credibility to a causal interpretation of the results shown in column 1.

5.2 Exploring the Heterogeneous Effects of Patents on Trade

While our overall results shown in Table 2 are informative, patents can be used for very different purposes (Griliches, 1998; De Rassenfosse, Palangkaraya, and Webster, 2016). The effects of patents on trade flows are therefore likely to depend in a heterogeneous way on the industry-level and market-level contexts in which patents are used. Investigating this heterogeneity is important for informing new theory as well as for understanding how the international patent data may be mapped onto existing ways of thinking about trade and innovation.

Heterogeneous effects by end-usage. Given that the cross-border patenting may be important for the spreading and deepening of global supply chains, an interesting way of splitting up industries is by end-usage. The left panel of Table 3 thus examines how our earlier results vary for intermediate goods, final goods, two different final goods subcategories ("consumption goods" and "capital goods"), different positions in the supply chain ("upstream" and "downstream"), high- and low-elasticity industries and the pharmaceutical industry. The latter warrants particular focus due to its centrality to international policy debates surrounding IPR and IPR reform. All results repeat the same specification from the first column of Table 2. Considering first the results for intermediates, final goods, capital goods, and consumption goods, we find that all of these industry subgroupings aside from consumption goods exhibit positive and significant effects for exports. Final goods are associated with modestly larger effects than intermediate goods, though this effect is really driven by capital goods as opposed to consumption goods. Looking at intermediate usage in another way, we find that industries that are relatively downstream in the supply chain exhibit positive and significant effects, whereas relatively upstream industries do not.[28]

The next source of potential ambiguity in this relationship we wish to speak to is what types of innovations the patents we observe generally represent. To fix ideas, consider a monopolistically competitive environment of a kind often used to study trade and patent protection where firms produce imperfectly substitutable varieties.[29] In this type of environment, if patents are used to patent the "variety" itself, the value of owning exclusive rights over that variety is larger when varieties are more differentiated from one another. However, if patents instead are used to protect innovations that generate cost and/or quality advantages for the producer of a particular variety over other producers, we should observe the opposite: patents should have larger effects for less differentiated industries—those with a higher elasticity of substitution between products—where cost and quality differences weigh more heavily on market share. This distinction is important for another reason as well: if patents mostly affect trade through the differentiated varieties channel, it is possible that the new varieties receiving patent protection are being specifically tailored to the markets where they are being patented. In that case, we may be worried that our estimates are picking up the value of the underlying innovations rather than the effects of patent protection.

To examine the effect of product differentiation, we include in Table 3 a subsample analysis on industries with high or low demand elasticity, where high-elasticity is defined as having an above-median demand elasticity.[30] We find that our earlier positive and significant estimates for the effects of the exporter's patent stock on trade are entirely driven by high-elasticity industries. Because the distinction between high and low elasticity industries is important for the interpretation of our results, we also find it instructive to test more formally whether this difference is statistically significant by pooling across all industries. In Table 4, we therefore interact our patent stock variable with an indicator for high-elasticity industries and consider a variety of different specifications, including with more stringent fixed effects (ijt), time-varying trade costs, and other controls. The results overwhelmingly support the interpretation that patents have larger effects on trade in less differentiated (i.e., high-elasticity) industries.[31] Focusing on column 1, which does not introduce

[28] Using regressions that pool across all industries, we have constructed formal hypothesis tests for some of these differences based on the same clustering assumptions used in the main analysis. Neither the difference between intermediate goods and final goods nor the difference between capital goods and consumption goods is statistically significant. The difference between upstream and downstream industries is significant at the 10% significance level ($p = 0.076$).

[29] For a canonical example, see Grossman and Lai (2004).

[30] Another approach would be to use the shares of products within each 3 digit SITC industry that meet Rauch's (1999)'s criteria for being "differentiated" versus "homogeneous". However, by far, most patents in our data are filed in industries that fall under the former category, whereas demand elasticities vary widely throughout the sample. Thus, we focus on the demand elasticity as our preferred measure of product differentiation.

[31] Interestingly, the distributions of patents are very similar in high- and low-elasticity industries, implying no important difference between the two industry groups in terms of the value of patents.

any added interactions or controls, the estimated coefficient of 0.034 for "log $STOCK_{ijkt}$ × high elasticity" is highly significant economically in relation to our earlier estimates in addition to being statistically significant. The sum of the coefficients on "log $STOCK_{ijkt}$" and "log $STOCK_{ijkt}$ × high elasticity" is 0.038 and is itself statistically significant ($p < 0.01$). Using our earlier calculations, this combined estimate equates to a difference in peak effects on exporting of 18.7% for high- versus low-value patent flows in these industries, comparable to the average effect of an FTA on industry-level trade.

Heterogeneous effects by factor-intensity. Continuing in a similar vein as the left-hand side panel of Table 3, the right-hand side shows further results for industry subgroupings broken out by factor intensity. Thanks to the rich data provided by Nunn and Trefler (2013), we are able to consider a wide range of factor intensities, including not only capital intensity, skill intensity, and labor intensity, but also the value added share. For most of these factor intensity measures, we take any industry with a factor intensity above the median across all industries to be intensive in that factor. The one exception is for labor-intensive industries, which we take to be any industry that is neither skill-intensive nor capital-intensive. The export effect is positive for each of the skill-, capital-, and labor-intensive categories, though none of these estimates is statistically significant. Within the capital-intensive subgroup, we also find it useful to report results that are relatively intensive in "specialized machinery", defined as machinery aside from computers or automobiles, since these industries are more likely to be reliant on proprietary capital.[32] Here, the effect is positive and significant. Turning to the results for value added, we find significant effects for low value-added industries but not for high value-added industries.

High- and medium-tech versus low-tech industries. Another natural way of examining the validity of our results is to separate out the industries where patents have traditionally been expected to play an important role. Following Palangkaraya, Jensen, and Webster (2017), we therefore show in Table 3 how our results vary across "high- and medium-tech" industries versus "low-tech" industries. To make this distinction, we use industries whose R&D intensities lie above the 33rd percentile as the "high- and medium-tech" group of industries. All other industries are considered "low-tech". Our estimates for these two categories turn out to be very similar in magnitude, though the only the result for the former is statistically significant, implying a much less noisy relationship for goods with a higher R&D intensity.[33] We also show further results for

[32]Using Nunn and Trefler (2013)'s data, we also obtained estimates in a similar fashion for automobile-, computer-, and building-intensive industries. We found significant results for automobile-intensive industries but not for the other two subcategories.

[33]One reason why patents may be associated with increased trade even for low-tech goods could be because these

"high-tech" versus "medium-tech" industries, in this case using the 67th percentile as the dividing line. Interestingly, this breakdown shows our earlier results are driven primarily by medium-tech industries, whereas the effect for high-tech industries is not significant. Looking within the high-tech category, we do find that the effect for the pharmaceutical industry, a particularly important high-tech industry in the IP context, is significant at the 10% level. Though this effect is noisily estimated, it is economically important, as it is about four-and-a-half times the magnitude of our baseline estimate.[34]

Bilateral trade cost variables. Next, we interact our patent stock variable with several further controls that could affect trade and/or patenting costs. These include the log of bilateral distance as well as indicators for whether the two countries have a common legal system or common language. In principle, the signs of these interactions could be ambiguous. On the one hand, they could mainly reflect the cost of trading and/or the cost of obtaining a patent; in which case, we would expect patents filed in markets associated with higher bilateral costs to be of higher value and thus have larger effects on trade. That is, we would expect the sign of the interaction with log distance to be positive and the signs for the interactions with common language and common legal system to be negative. On the other hand, these variables could instead mostly reflect the difficulty of enforcing a patent, in which case the expected signs would be opposite. As column 2 of Table 4 shows, the results support a mix of these two hypotheses: the interaction between log $STOCK_{ijkt}$ and the log of distance is positive but not significant, the interaction with common language is negative and significant, and the interaction with common legal system is positive and significant. These results are not affected when we add interactions with "high elasticity" or IPR regime strength (column 3), nor are they affected by the addition of our time-varying trade cost controls (column 5). However, only the estimate for the interaction with log distance remains stable when we add an ijt fixed effect (column 6), becoming marginally significant in that case.[35]

Strength of the IPR regime. Another hypothesis we wish to examine is whether patents require a strong enforcement environment in the host country in order to have an effect on trade.

goods are easier to imitate by local firms, thus necessitating more use of patent protections by patent-holding exporters. The possibility that higher-tech goods are less imitable than lower-tech goods could also explain why the effects we estimate are more concentrated in the medium-tech category than in the high-tech category.

[34] Previous studies by Delgado, Kyle, and McGahan (2013) and Maskus and Riley (2016) have indicated that strong IPR regimes boost flows of trade in the pharmaceutical industry specifically and have much weaker effects in other industries.

[35] Since log distance is a continuous variable, we use the difference between log distance and the mean log distance for these interactions. This means all other point estimates can be interpreted as occurring at the mean log distance between countries. If we do not do this, the main effect captures the mean log distance effect and changes in the mean effect across specifications become harder to interpret.

We thus use the Park (2008) index for the quality of patent protections, a standard measure in the literature, to construct indicators for whether the patent-granting countries has a "medium" or "high" level of intellectual property rights protections. Specifically, the Park (2008) index scores countries between 0 and 5. "Low", "medium", and "high" IPR correspond to scores less than 2, between 2 and 4, and above 4 respectively. Column 3 of Table 4 adds interactions between our patent stock variables and these IPR indicators. Interestingly, these IPR interactions both enter negatively and neither is significant. A possible explanation is that higher IPR protections may attract both trade and other channels for foreign sales that substitute for trade (e.g., FDI, licensing) in roughly equal measure. Alternatively, it could be that the effect of the IPR regime on trade is mainly felt through the selection of patents that are filed in different markets, rather than through how an observed patent affects trade conditional on having been filed. If lower-IPR markets are associated with higher trade costs, for example, then patents filed in those markets may need to have higher trade value in order to compensate. Our further results in columns 4-6 that add richer controls and fixed effects mostly deliver similar findings, with the interaction with "high" IPR actually becoming marginally significant when an ijt fixed effect is added in column 5.

5.3 Additional Sensitivity Checks

Because the empirical framework we have constructed requires a complex set of choices on our part, Table 5 examines the sensitivity of our analysis to various reasonable alternatives. We consider first whether our results are affected by dropping all observations for which the destination is an EPO country from our sample. Noting that our PATSTAT data does not indicate the country (or countries) where EPO patents are designated to be applied, this is an important check that ensures our results are not driven by how we have treated these patents. Though removing EPO destinations is very costly in terms of how it affects the sample (removing 50% of the original number of observations), we find that our estimated coefficient on $\log STOCK_{ijkt}$ remains similar in magnitude to our baseline estimate and retains most of its statistical significance ($p = 0.054$).

Column 2 of Table 5 experiments with dropping only a single country, the U.S. The motivation behind this check is the concern that our results could be driven by "non-practicing entities" (often referred to as "patent trolls"), which are known for asserting patent rights inventions without actually producing products associated with those innovations.[36] Anecdotally, patent trolls are

[36]This is related also to "strategic patenting", where firms acquire patents in order to raise their rivals' costs of potential litigation and as bargaining chips for settling disputes. If this type of patenting was driving our results, we should see no effect of patents filed by country i in country j on trade between that ij pair. Our results clearly show this is not the case.

known to be predominantly active in the U.S. (Fusco, 2013). Dropping the U.S. entirely is thus a conservative way of dealing with this concern. As our results show, patents continue to have a statistically significant effect on exports when the U.S. is removed, with the point estimate again being similar to our baseline.

Columns 3 and 4 then explore some different ways of weighting patents. Column 3 uses a patent stock based on unweighted patent counts, whereas the stock used in column 4 weights patents by citations, another common way of assessing patent value. Interestingly, neither result is statistically significant.[37] For unweighted patents, the lack of a result is perhaps unsurprising due to the amount of noise in the patent data and the expectation at that some patents have much more value than others. For citation-weighting, one possible explanation is that citations are known to be a less internationally comparable measure of patent value than family size (De Rassenfosse, 2013).[38] While we would not claim family size is a perfect measure of patent value either, it is noteworthy in this context that family size appears to provide a more robust measure of a patent's value for promoting exports versus these other alternatives.

Next, columns 6-8 examine the implications of varying the assumed diffusion and decay parameters from equation (3) that cause the effects of patents to phase in gradually and then fade over time. Column 6 investigates the extreme situation where the effects of patents do not decay, setting $\omega_1 = 0$. The next two columns then respectively consider the implications of doubling the diffusion parameter ($\omega_2 = 0.50$) and allowing for immediate diffusion ($\omega_2 = 1$). We find that our results remain statistically significant for either the case of no decay or if they diffuse at the faster rate implied by $\omega_2 = 0.50$. However, a comparison with our earlier results reveals that relaxing our presumed stock assumptions does weaken the fit with the data versus our baseline specification. Allowing for a gradual rate of diffusion correctly appears to be particularly important.[39] For the case of faster diffusion, the effect of the patent stock on exports is only marginally significant ($p = 0.082$), and in the extreme case of full diffusion, the export effect is insignificant.[40] It is important to note that the latter assumption implies that the life of a patent peaks within one year of the application

[37] The p-values are 0.134 and 0.125, respectively.

[38] In general, though citation counts have been found to be correlated with patent value, they are also known to be significantly influenced by idiosyncratic citation practices that reflect differences in geography, technology type, examiner effects, and strategic considerations (Jaffe and De Rassenfosse, 2019).

[39] Varying the decay rate also has a small effect: the p-value rises to 0.012 in the case no decay compared to 0.0097 using the baseline parameters.

[40] The difference in statistical significance between columns 6 and 7 in the table is easier to appreciate when the estimated values are rounded to the 4th digit. The estimated coefficient for log $STOCK_{ijt}$ in column 6 is 0.0054 with standard error of 0.0031. In column 7, the point estimate is 0.0046 and the standard error is 0.0030, implying a p value of 0.118. We have also experimented with other diffusion rates, finding that the p value for log $STOCK_{ijt}$ decreases more or less monotonically between $\omega_2 = 0.50$ and $\omega_2 = 1$, becoming statistically insignificant at about $\omega_2 = 0.75$.

being filed, which is often not enough time for the patent application to have been granted (Popp, 2003). Our default diffusion and decay parameter values imply that this peak instead occurs after about four-and-a-half years, which does still allow for the possibility that pending patents that have not yet been granted are able to offer a degree of protection.

The last column of Table 5 replaces our original patent variable with a dichotomous variable, $1[\log STOCK^*_{ijkt} > \text{median}]$, reflecting a high level of patenting. To construct this variable, we first de-mean the log of the patent stock with respect to its trade-weighted mean value for that industry-destination-year. Doing so allows us to correct for differences in market-specific conditions, which otherwise would explain a significant portion of the difference in raw stock values.[41] The dummy variable we use is set equal to 1 if the resulting de-meaned variable, $\log STOCK^*_{ijkt}$, exceeds its median. We set it to zero otherwise, including when the patent stock equals zero, thus enabling us to include additional observations. The advantage of this dummy variable approach is that it gives us a cleaner and simpler interpretation: the effect of a high level of patenting using this approach is an $\exp(0.150) - 1 = 16.2\%$ increase in trade, which we find to be statistically significant. This effect is larger than some of the estimated increases we found using a continuous stock variable, possibly because there may be bigger effects at the lower end of the stock distribution. We also include a dummy for whether a patent has ever been filed in that market by the exporting country. We find no statistically significant effect, indicating the extensive margin of patenting may matter less than the intensive margin.[42]

Finally, to highlight the validity and meaningfulness of our preferred fixed effects specification from (2), we present sensitivity checks using varying sets of fixed effects in Table 6. Column 1 starts with a relatively simple specification, with country-pair and time fixed effects only. Subsequent columns introduce industry fixed effects (column 2) and standard exporter-time and importer-time "gravity" fixed effects (column 3), before going on to demonstrate how interacting the country-pair and country-time fixed effects with the industry dimension matters for inference (columns 4-6). While the robustness of the significance of patents on export is clear, the varying of the fixed

[41] Recall that our estimations include an industry-destination-year fixed effect. Thus, simply using a dummy for whether the $\log STOCK_{ijkt}$ is above its median will not necessarily be all that informative if much of the variability in $\log STOCK_{ijkt}$ is driven by differences across destination markets, with some having much more intense patenting activity than others. Indeed, if we simply use a dummy variable for whether $\log STOCK_{ijkt}$ is above its median, we obtain results that are near zero and statistically insignificant, indicating this is not a good way of replicating the variation we observed in our prior estimates. Destination-industry-year-specific factors turn out to explain about half of the overall time variation in our original stock variable, further justifying this approach.

[42] Results in this last column are similar regardless of whether or not we include observations for which the patent stock is zero. We should also note that we exclude observations for which no patents are observed in any industry for either the exporting country or the importing country in a given year, since these observations could potentially reflect missing data. Results are again similar either way, however.

effects across columns reveals that using a more relaxed fixed effects specification introduces large biases in our estimates. When exporter-industry-time and importer-industry-time fixed effects are not used (in columns 1-3 and 5-6), the coefficient on the patent stock is always larger than 0.1. When these fixed effects are added, the estimated export effect falls by an order of magnitude. As discussed in French (2016, 2019), these *ikt*- and *jkt*-specific fixed effects can be thought of as respectively capturing the exporter's industry-level comparative advantages and the importer's pattern of relative demands. Our results therefore indicate that these industry-country-specific supply and demand shifters change substantially over the sample and should generally be regarded as an important source of bias in panel data gravity settings, consistent with the arguments of French (2016, 2019).[43] Introducing exporter-importer-pair effects also dramatically reduces the magnitude of the export effect (either compare column 2 with column 5 or compare column 4 with our baseline estimate in Table 2, column 1.) This is evidently because the propensity to trade is heterogeneous across industries within exporter-importer pairs and this heterogeneity is correlated with the propensity to file a patent in the importing country.

6 Conclusion

Cross-border patenting is regarded anecdotally and theoretically as an important marker of globalization and the spread of knowledge across borders. But what the decision to file a patent in a particular trade partner can tell us about the motivations behind cross-border patents and their implications for trade remains an open question. To provide answers, we assemble a highly disaggregated data set of industry-level trade flows and cross-border patents, which we subject to a rich set of fixed effects that allow us to specifically isolate the bilateral dimension of the data for a given industry. Notably, ours is the first study to systematically investigate what the implications of bilateral patenting are for international trade using a large and highly disaggregated data set.

Our main finding is that bilateral patents in a given industry on average promote bilateral exports in that same industry, with no significant effect seen for imports returning in the opposite direction. However, we find these overall results mask substantial sources of heterogeneity across industries. Namely, we find strong evidence that the pro-exporting effects of patents are concentrated in less differentiated industries associated with a higher demand elasticity and in industries that are relatively more downstream in supply chains. Interestingly, we also find that the effect of

[43]Notice it is these fixed effects, not the *ijk* pair-industry fixed effect, that has the largest effect on the FTA coefficient, which is not industry-specific.

patents on exports does not vary significantly with the host country's national IPR regime, indicating that bilateral patent data contains important information about the use of patent protections that cannot be captured otherwise. Taken together, these results support the interpretation that cross-border patents are used to preserve advantages in production costs and/or product quality versus foreign competitors and that they are often effective at securing these advantages irrespective of the strength of the local IPR regime. Given that technology diffusion is a significant determinant of economic growth (Comin and Hobijn, 2010), we hope these findings can inform policies intended to stimulate innovation and promote the diffusion of knowledge.

A further intended contribution of this paper is our attempt to provide a methodology that can be used to analyze bilateral patent flows in conjunction with bilateral trade flows in a panel setting. By allowing patents to have effects that grow and decay over time and by pursuing a rigorous fixed effects specification that allows us to specifically isolate the within-industry and within-pair dimension of the data, we are able to obtain estimates that we would argue are surprisingly well identified in the face of natural concerns about reverse causality. For future work, an attractive extension would be to adopt a more structural approach in order to disentangle whether foreign patents contribute to the transmission of technological comparative advantage in a manner consistent with theories of trade and growth.

References

Alfaro, L., P. Antràs, D. Chor, and P. Conconi (2015): "Internalizing Global Value Chains: A Firm-Level Analysis," Working Paper 21582, National Bureau of Economic Research.

Anderson, J. E., and E. van Wincoop (2003): "Gravity with Gravitas: A Solution to the Border Puzzle," *American Economic Review*, 93(1), 170–192.

Anderton, B. (1999): "Innovation, Product Quality, Variety, and Trade Performance: An Empirical Analysis of Germany and the UK," *Oxford Economic Papers*, 51(1), 152–167.

Antimiani, A., and V. Costantini (2013): "Trade Performances and Technology in the Enlarged European Union," *Journal of Economic Studies*, 40(3), 355–389.

Antràs, P., and D. Chor (2013): "Organizing the Global Value Chain," *Econometrica*, 81(6), 2127–2204.

Antràs, P., D. Chor, T. Fally, and R. Hillberry (2012): "Measuring the upstreamness of production and trade flows," *American Economic Review*, 102(3), 412–16.

Autor, D., D. Dorn, G. H. Hanson, G. Pisano, and P. Shu (2016): "Foreign Competition and Domestic Innovation: Evidence from U.S. Patents," Working Paper 22879, National Bureau of Economic Research, DOI: 10.3386/w22879.

Aw, B. Y., M. J. Roberts, and D. Y. Xu (2011): "R&D Investment, Exporting, and Productivity Dynamics," *American Economic Review*, 101, 1312–1344.

Awokuse, T. O., and H. Yin (2010): "Intellectual property rights protection and the surge in FDI in China," *Journal of Comparative Economics*, 38(2), 217–244.

Baier, S. L., and J. H. Bergstrand (2007): "Do Free Trade Agreements Actually Increase Members' International Trade?," *Journal of International Economics*, 71(1), 72–95.

Baldwin, R., and D. Taglioni (2007): "Trade Effects of the Euro: A Comparison of Estimators," *Journal of Economic Integration*, 22(4), 780–818.

Batrakova, S., and A. Dechezleprêtre (2013): "Is there a Trade-off between 'Dirty' Imports and 'Clean' Innovation?," Discussion paper, Mimeo.

Blonigen, B. A. (2001): "In Search of Substitution between Foreign Production and Exports," *Journal of International Economics*, 53(1), 81–104.

Bloom, N., M. Draca, and J. V. Reenen (2011): "Trade Induced Technical Change? The Impact of Chinese Imports on Innovation, IT and Productivity," Working Paper 16717, National Bureau of Economic Research, DOI: 10.3386/w16717.

Bond, E. W., and K. Saggi (2014): "Compulsory Licensing, Price Controls, and Access to Patented Foreign Products," *Journal of Development Economics*, 109, 217–228.

Boring, A. (2015): "The Impact of Patent Protection on US Pharmaceutical Exports to Developing Countries," *Applied Economics*, 47(13), 1314–1330.

Briggs, K., and W. G. Park (2014): "There will be Exports and Licensing: The Effects of Patent Rights and Innovation on Firm Sales," *The Journal of International Trade & Economic Development*, 23(8), 1112–1144.

Broda, C., and D. E. Weinstein (2006): "Globalization and the Gains from Variety," *Quarterly Journal of Economics*, 121(2), 541–585.

Brunel, C. (2019): "Green Innovation and Green Manufacturing: Links between Environmental Policies, Innovation, and Production," *Journal of Environmental Management*, 248.

Chalioti, E., S. Drivas, Kyriakos Kalyvitas, and M. Katsimi (2017): "Innovation, Patents and Trade: A Firm-level Analysis," Discussion paper, mimeo.

Chen, W.-C. (2013): "The Extensive and Intensive Margins of Exports: The Role of Innovation," *The World Economy*, 36(5), 607–635.

Cheng, I.-h., and H. J. Wall (2005): "Controlling for Heterogeneity in Gravity Models of Trade and Integration," *Federal Reserve Bank of St. Louis Review*, 87(1), 49–64.

Comin, D., and B. Hobijn (2010): "An Exploration of Technology Diffusion," *The American Economic Review*, 100(5), 2031–2059.

Correia, S. (2015): "Singletons, Cluster-robust Standard Errors and Fixed Effects: A Bad Mix," *Unpublished paper.*

Correia, S., P. Guimarães, and T. Zylkin (2019a): "PPMLHDFE: Fast Poisson Estimation with High-dimensional Data," *Unpublished manuscript.*

——— (2019b): "Verifying the Existence of Maximum Likelihood Estimates for Generalized Linear Models," *Unpublished manuscript.*

De Rassenfosse, G. (2013): "Do Firms Face a Trade-off between the Quantity and the Quality of their Inventions?," *Research Policy*, 42(5), 1072–1079.

De Rassenfosse, G., A. Palangkaraya, and E. Webster (2016): "Why do Patents Facilitate Trade in Technology? Testing the Disclosure and Appropriation Effects," *Research Policy*, 45(7), 1326–1336.

Delgado, M., M. Kyle, and A. M. McGahan (2013): "Intellectual Property Protection and the Geography of Trade," *Journal of Industrial Economics*, 61(3).

Dernis, H., and M. Khan (2004): "Triadic Patent Families Methodology," .

Feenstra, R. C., R. E. Lipsey, H. Deng, A. C. Ma, and H. Mo (2005): "World Trade Flows: 1962-2000," Working Paper 11040, National Bureau of Economic Research.

Figueiredo, O., P. Guimarães, and D. Woodward (2015): "Industry Localization, Distance Decay, and Knowledge Spillovers: Following the Patent Paper Trail," *Journal of Urban Economics*, 89(C), 21–31.

French, S. (2016): "The Composition of Trade Flows and the Aggregate Effects of Trade Barriers," *Journal of International Economics*, 98, 114–137.

——— (2019): "Comparative Advantage and Biased Gravity," *Unpublished manuscript.*

Fusco, S. (2013): "Markets and Patent Enforcement: A Comparative Investigation of Non-Practicing Entities in the United States and Europe," *Mich. Telecomm. & Tech. L. Rev.*, 20, 439.

Glick, R., and A. K. Rose (2016): "Currency Unions and Trade: A Post-EMU Reassessment," *European Economic Review*, 87, 78–91.

Greenhalgh, C. (1990): "Innovation and Trade Performance in the United Kingdom," *The Economic Journal*, 100(400), 105–118.

Griliches, Z. (1995): "R&D and Productivity: Econometric Results and Measurement Issues," in *Handbook of the Economics of Innovation and Technological Change*, pp. 52–89. Blackwell Publishers.

——— (1998): "Patent Statistics as Economic Indicators: a Survey," in *R&D and Productivity: the Econometric Evidence*, pp. 287–343. University of Chicago Press.

Grossman, G. M., and E. L.-C. Lai (2002): "International Protection of Intellectual Property," *NBER Working Paper*, 8704.

——— (2004): "International Protection of Intellectual Property," *American Economic Review*, 94(5), 1635–1653.

Hall, B. H., A. B. Jaffe, and M. Trajtenberg (2001): "The NBER patent citations data file: Lessons, insights and methodological tools," Working Paper 8498, National Bureau of Economic Research.

Head, K., and T. Mayer (2014): "Gravity Equations: Workhorse, Toolkit, and Cookbook," in *Handbook of International Economics*, ed. by G. Gopinath, E. Helpman, and K. Rogoff, vol. 4, pp. 131–195. North Holland, 4 edn.

Ivus, O. (2010): "Do Stronger Patent Rights Raise High-tech Exports to the Developing World?," *Journal of International Economics*, 81(1), 38–47.

Jaffe, A. B., and G. De Rassenfosse (2019): "Patent citation data in social science research: Overview and best practices," in *Research Handbook on the Economics of Intellectual Property Law*. Edward Elgar Publishing.

Keller, W. (2004): "International Technology Diffusion," *Journal of Economic Literature*, 42(3), 752–782.

Keller, W., and S. R. Yeaple (2013): "The Gravity of Knowledge," *The American Economic Review*, 103(4), 1414–1444.

Kinzius, L., A. Sandkamp, and E. Yalcin (2018): "Trade Protection and the Role of Non-tariff Barriers," Discussion paper.

Lai, E. L.-C., and L. D. Qiu (2003): "The North's Intellectual Property Rights Standard for the South?," *Journal of International Economics*, 59(1), 183–209.

Larch, M., J. Wanner, Y. V. Yotov, and T. Zylkin (2019): "Currency Unions and Trade: A PPML Re-assessment with High-dimensional Fixed Effects," *Oxford Bulletin of Economics and Statistics*, 81(3), 487–510.

Limão, N. (2007): "Are Preferential Trade Agreements with Non-trade Objectives a Stumbling Block for Multilateral Liberalization?," *The Review of Economic Studies*, 74(3), 821–855.

Lotti, F., and G. Marin (2013): “Matching of PATSTAT Applications to AIDA Firms: Discussion of the Methodology and Results,” Working Paper 166, Bank of Italy Occasional Paper.

Lovely, M., and D. Popp (2011): “Trade, Technology, and the Environment: Does Access to Technology Promote Environmental Regulation?,” *Journal of Environmental Economics and Management*, 61(1), 16–35.

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(3), 530–542.

Maskus, K. E. (2018): “International Agreements on Intellectual Property Rights: TRIPS and Beyond,” in *Handbook of International Trade Agreements*, pp. 59–70. Routledge.

Maskus, K. E., and M. Penubarti (1995): “How Trade-related are Intellectual Property Rights?,” *Journal of International Economics*, 39(3), 227–248.

Maskus, K. E., and W. Riley (2016): “Intellectual Property-Related Preferential Trade Agreements and the Composition of Trade,” *RCAS Working Paper*, 2016/35.

McCalman, P. (2001): “Reaping What you Sow: An Empirical Analysis of International Patent Harmonization,” *Journal of International Economics*, 55(1), 161–186.

Nunn, N., and D. Trefler (2013): “Incomplete Contracts and the Boundaries of the Multinational Firm,” *Journal of Economic Behavior & Organization*, 94(Supplement C), 330–344.

Pakes, A., and M. Simpson (1989): “Patent Renewal Data,” *Brookings Papers: Microeconomics*, pp. 331–410.

Palangkaraya, A., P. H. Jensen, and E. Webster (2017): “The Effect of Patents on Trade,” *Journal of International Economics*, 105, 1–9.

Park, W. G. (2008): “International Patent Protection: 1960–2005,” *Research Policy*, 37(4), 761–766.

Popp, D. (2002): “Induced innovation and energy prices,” *American economic review*, 92(1), 160–180.

——— (2003): “Pollution control innovations and the Clean Air Act of 1990,” *Journal of Policy Analysis and Management*, 22(4), 641–660.

Popp, D., I. Hascic, and N. Medhi (2011): “Technology and the Diffusion of Renewable Energy,” *Energy Economics*, 33(4), 648–662.

Robbins, C. A. (2006): “The Impact of Gravity-weighted Knowledge Spillovers on Productivity in Manufacturing,” *The Journal of Technology Transfer*, 31(1), 45–60.

Saggi, K. (2002): “Trade, Foreign Direct Investment, and International Technology Transfer: A

Survey," *The World Bank Research Observer*, 17(2), 191–235.

——— (2016): "Trade, Intellectual Property Rights, and the World Trade Organization," in *Handbook of Commercial Policy*, vol. 1, pp. 433–512. Elsevier.

Santos Silva, J. M. C., and S. Tenreyro (2006): "The Log of Gravity," *Review of Economics and Statistics*, 88(4), 641–658.

Schankerman, M. (1998): "How Valuavle is Patent Protection? Estimates by Technology Field," *The RAND Journal of Economics*, 29(1), 77–107.

Scherngell, T., and Y. Hu (2011): "Collaborative Knowledge Production in China: Regional Evidence from a Gravity Model Approach," *Regional Studies*, 45(6), 755–772.

Schmidt, J., and W. Steingress (2018): "No Double Standards: Quantifying the Impact of Standard Harmonization on Trade," Discussion paper, mimeo.

Weidner, M., and T. Zylkin (2020): "Bias and consistency in three-way gravity models," *arXiv preprint arXiv:1909.01327*.

Wooldridge, J. M. (2002): *Econometric Analysis of Cross Section and Panel Data*. MIT Press, Cambridge, MA.

Yang, L., and K. E. Maskus (2009): "Intellectual Property Rights, Technology Transfer and Exports in Developing Countries," *Journal of Development Economics*, 90(2), 231–236.

Yotov, Y. V., R. Piermartini, J.-A. Monteiro, and M. Larch (2016): *An Advanced Guide to Trade Policy Analysis: The Structural Gravity Model*. World Trade Organization, Geneva.

7 Figures

Figure 1: Growth in Cross-border Patenting Flows,1972-2008

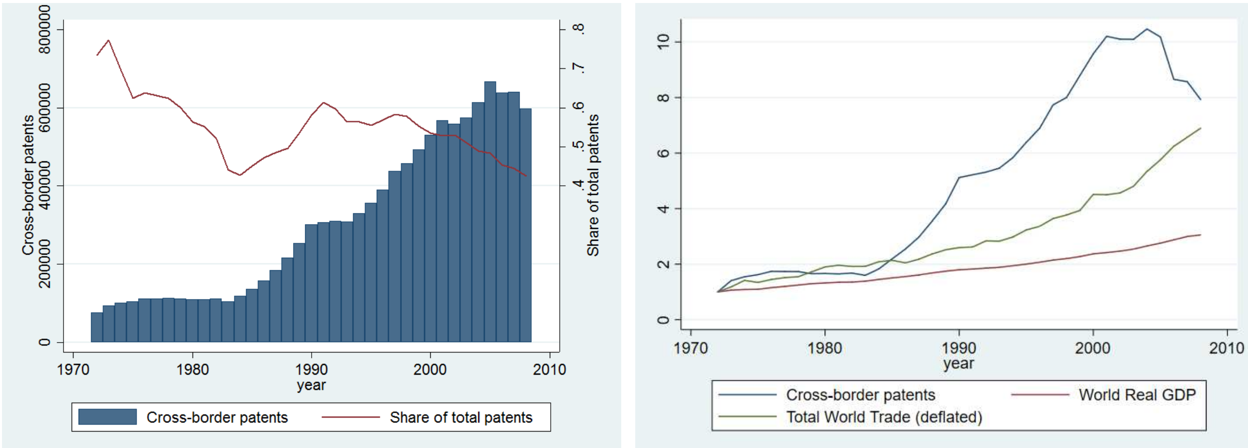

Note: the left graph shows total world cross-border patents as share of total world patents. The right graph shows growth in cross-border patenting between 1972-2008 versus growth in world trade and world real GDP. All values in the right graph are normalized to 1 in 1972. World trade values are deflated using the US GDP deflator. Patent counts are the raw number of patents filed in a given year. Sources: PATSTAT, WDI.

Figure 2: Counts of patents by origin type and destination type, 1972-2008

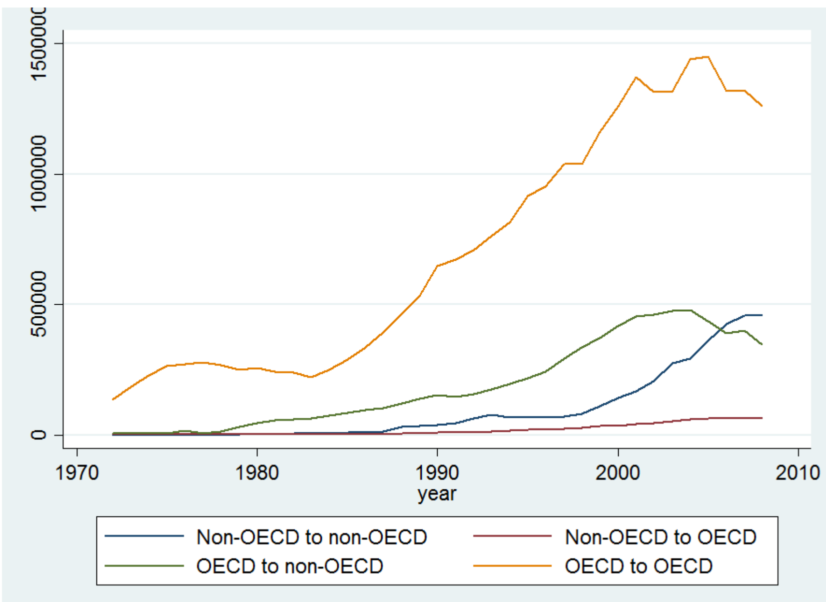

Note: Patent counts are the raw number of patents filed in a given year. We include cross-border patents only. Source: PATSTAT.

8 Tables

Table 1: Summary statistics of patent data

Variable	Mean	Median	Std. Dev.	Min.	Max.
Patent flow (unweighted)	4	0	61	0	12,829
OECD	8	0	82	0	12,829
Other	0.2	0	8	0	2,279
Patent flow (weighted)	34	0	536	0	131,409
OECD	60	0	722	0	131,409
Other	2	0	49	0	27,274
Patent stock	184	0	2954	0	593,948
OECD	326	0	3977	0	593,948
Other	11	0	190	0	34,179
Patent stock (non-zero only)	705	10	5754	0.0001	593,948
OECD	972	20	6819	0.0002	593,498
Other	64	2	457	0.0001	34,179
N=6,411,627					

Industries with highest weighted patent stock growth

759 – Parts, nes of and accessories for office machines and automatic data processing machines
971 – Gold, non-monetary (excluding gold ores and concentrates)
844 – Under garments of textile fabrics, not knitted or crocheted
846 – Under-garments, knitted or crocheted
843 – Womens, girls, infants outerwear, textile, not knitted or crocheted
716 – Rotating electric plant and parts thereof, nes
893 – Articles, nes of plastic materials
771 – Electric power machinery, and parts thereof, nes
845 – Outerwear knitted or crocheted, not elastic nor rubberized
742 – Pumps for liquids; liquid elevators; and parts thereof, nes
689 - Miscellaneous non-ferrous base metals, employed in metallurgy

Country-pairs (filer – granter) with highest weighted patent stock growth

Germany – Portugal	Japan – Poland	Japan – Portugal
Japan – Greece	United States - China	Japan - China
Great Britain - Portugal	France - Greece	Great Britain - Poland

Unweighted patent flows are simple counts of the number of patents filed. Weighted patents flows account for value by weighting each patent by its family size. Stocks are created using the value-weighted patent flows. Due to the concordance between the patent and trade datasets, flows can be non-integer values below 1.

Table 2: Do Cross-Border Patent Promote Trade?

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Trade flows						Patent stock value	
FTA	0.156***	0.155***	0.155***	0.129***		0.170***		
	(0.030)	(0.030)	(0.033)	(0.029)		(0.029)		
log $STOCK_{ijkt}$ (*export effect*)	0.018***		0.018**	0.018***	0.019***	0.016**		
	(0.007)		(0.008)	(0.007)	(0.006)	(0.007)		
log $STOCK_{jikt}$ (*import effect*)		0.001	0.000					
		(0.003)	(0.003)					
FTA_{t+4} (*lead FTA effect*)				0.065**			0.013	
				(0.031)			(0.023)	
log $STOCK_{ijt+4}$ (*export effect, lead*)				0.001				
				(0.010)				
$\ln X_{t+4}$ (*lead log trade value*)								-0.001
								(0.001)
Fixed effects and other controls:								
Exporter-industry-time	x	x	x	x	x	x	x	x
Importer-industry-time	x	x	x	x	x	x	x	x
Industry-country pair	x	x	x	x	x	x	x	x
Country pair-time					x			
Time-varying trade cost controls						x		
N	1,172,633	1,153,001	1,000,617	1,071,765	1,170,136	1,172,633	1,118,145	1,116,285

The dependent variable for columns 1-5 is industry-level (3 digit SITC rev 2) trade flows between 149 countries covering 1974-2006, using every 4 years. All estimations use PPML. Standard errors, shown in parentheses, are two-way clustered by country-pair and exporter-industry. log $STOCK_{ijt}$ refers to (the log of) the stock of patents filed by the exporting country i in the importing country j in industry k. log $STOCK_{jikt}$ refers to (the log of) the stock of patents filed by the importer j in the exporting country i in industry k. The construction of these stocks is described in Section 4.2. "Time-varying trade cost controls" refers to three-way interactions between standard pairwise gravity controls (log distance, common language, and common legal system), the year, and the degree of R&D intensity in that that industry (high vs. medium vs. low R&D intensity). Observation counts vary because we first use the algorithms of Correia (2015) and Correia, Guimarães, and Zylkin (2019b) to reduce the sample by removing singletons and perfectly predicted observations.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3: Industry-level Breakdown

	Effect of log patent stock	s.e.		Effect of log patent stock	s.e.
By end use:			**By factor intensity**:		
Intermediate Goods	0.016*	(0.009)	Low VA share	0.025***	(0.009)
Final Goods	0.022**	(0.009)	High VA share	0.011	(0.009)
Capital Goods	0.033**	(0.013)	Skill-intensive	0.010	(0.009)
Consumption Goods	0.010	(0.010)	Capital-intensive	0.014	(0.009)
Relatively Upstream	0.007	(0.009)	Specialized machinery-int.	0.021**	(0.010)
Relatively Downstream	0.027***	(0.009)	Labor-intensive	0.016	(0.010)
High Demand Elasticity	0.037***	(0.009)	High- & Med-tech	0.017**	(0.008)
Low Demand Elasticity	0.003	(0.010)	High-tech	0.009	(0.010)
Pharma Only	0.071*	(0.037)	Med-tech	0.031***	(0.009)
Without Pharma	0.018**	(0.007)	Low-tech	0.017	(0.011)

The dependent variable is industry-level (3 digit SITC rev 2) trade flows between 149 countries covering 1974-2006, using every 4 years. Standard errors are two-way clustered by country-pair and exporter-industry. All estimations use PPML with exporter-industry-time (ikt), importer-industry-time (jkt), and exporter-importer-industry (ijk) fixed effects. The estimate shown is for (the log of) the stock of patents filed by the exporting country i in the importing country j in industry k. The construction of these stocks is described in Section 4.2. Each estimate shown uses a different subset of industries. "Intermediate Goods", "Final Goods","Capital Goods", and "Consumption Goods" are determined using BEC classifications at the HS 6 digit level. More precisely, any 3 digit SITC industry in which more than one-half of world trade meets the BEC definition for one of these categories is coded as belonging to that category. Factor-intensity classifications are constructed similarly using data from Nunn and Trefler (2013). The "upstreamness" of an industry is from Antràs, Chor, Fally, and Hillberry (2012). All regressions include FTA as an additional control. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Effect of IPR, Trade Costs, and Demand Elasticity

	(1)	(2)	(3)	(4)	(5)	(6)
	All industries combined					
FTA	0.156***	0.172***	0.155***	0.170***		0.174***
	(0.030)	(0.029)	(0.030)	(0.030)		(0.029)
log $STOCK_{ijkt}$	0.003	0.015**	0.030***	0.013	0.028***	0.010
	(0.008)	(0.007)	(0.009)	(0.009)	(0.011)	(0.010)
log $STOCK_{ijkt}$ × high elasticity	0.034***			0.033***	0.023**	0.034***
	(0.010)			(0.011)	(0.010)	(0.011)
log $STOCK_{ijkt}$ × MED IPR			-0.013	-0.012	-0.009	-0.011
			(0.008)	(0.008)	(0.009)	(0.008)
log $STOCK_{ijkt}$ × HIGH IPR			-0.015	-0.014	-0.024*	-0.012
			(0.010)	(0.010)	(0.013)	(0.010)
log $STOCK_{ijkt}$ × ln DIST		0.007		0.007	0.011*	0.005
		(0.005)		(0.005)	(0.006)	(0.005)
log $STOCK_{ijkt}$ × LANG		-0.029***		-0.029***	-0.004	-0.029**
		(0.010)		(0.010)	(0.019)	(0.013)
log $STOCK_{ijkt}$ × LEGAL		0.015**		0.015**	-0.024*	0.012
		(0.006)		(0.006)	(0.012)	(0.009)
Country-pair time FE					x	
Time-varying trade cost controls						x
N	1,171,571	1,172,633	1,129,287	1,128,249	1,125,891	1,128,249

The dependent variable is industry-level (3 digit SITC rev 2) trade flows between 149 countries covering 1974-2006, using every 4 years. Standard errors, shown in parentheses, are two-way clustered by country-pair and exporter-industry. All estimations use PPML with exporter-industry-time (ikt), importer-industry-time (jkt), and exporter-importer-industry (ijk) fixed effects. log $STOCK_{ijkt}$ refers to (the log of) the stock of patents filed by the exporting country i in the importing country j in industry k. The construction of these stocks is described in Section 4.2. The interactions with ln DIST use the difference between log distance and the mean log distance. This means all other point estimates can be interpreted as occurring at the mean log distance between countries. Observation counts vary because we first use the algorithms of Correia (2015) and Correia, Guimarães, and Zylkin (2019b) to reduce the sample by removing perfectly predicted observations.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: Additional Sensitivity Checks

	(1) Drop EPO Patent Destinations	(2) Drop USA as Patenting Country	(3) Unweighted patents	(4) Citation-weighted patents	(5) No decay	(6) Faster diffusion (ω_2=0.50)	(7) Immediate diffusion (ω_2=1)	(8) Discretized patent stock
FTA	0.106***	0.175***	0.157***	0.156***	0.157***	0.156***	0.156***	0.136***
	(0.03)	(0.032)	(0.030)	(0.030)	(0.030)	(0.030)	(0.030)	(0.026)
log $STOCK_{ijkt}$	0.016*	0.016**	0.005	0.004	0.008**	0.005*	0.005	
	(0.008)	(0.007)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	
1[log$STOCK^*_{ijkt}$ > median]								0.150***
								(0.017)
1[$STOCK_{ijkt} > 0$]								0.017
								(0.020)
N	549,163	1,108,881	1,172,503	1,172,503	1,172,504	1,172,633	1,172,633	2,349,984

The dependent variable is industry-level (3 digit SITC rev 2) trade flows between 149 countries covering 1974-2006, using every 4 years. Standard errors, shown in parentheses, are clustered by country-pair and eporter-industry. All specifications use PPML with exporter-industry-time (ikt), importer-industry-time (jkt), and exporter-importer-industry (ijk) fixed effects. log $STOCK_{ikjt}$ refers to (the log of) the stock of patents filed by the exporting country i in the importing country j in industry k. The construction of these stocks is described in Section 4.2. Column 1 drops all patent destinations associated with the EPO. Column 2 drops the US as an exporting/patent-filing country. Columns 3 and 4 respectively experiment with using unweighted and citations-weighted patents instead of family size-weighted patents. Columns 5-7 experiment with different diffusion and decay assumptions for constructing patent stocks, discussed further in the text. In column 8, 1[log $STOCK^*_{ijkt}$ > median] is constructed as a dummy variable by first demeaning log $STOCK_{ijkt}$ with respect to the trade-weighted mean value for patent-granting country j in industry k at time t and then setting equal to 1 if the resulting value (log $STOCK^*_{ijkt}$) is above its median. This dummy is set equal to zero otherwise, including whenever the patent stock is zero. we first use the algorithms of Correia (2015) and Correia, Guimarães, and Zylkin (2019b) to reduce the sample by removing perfectly predicted observations.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: Varying Fixed Effects

	(1)	(2)	(3)	(4)	(5)	(6)
FTA	0.395***	0.398***	0.278***	0.161***	0.423***	0.235***
	(0.125)	(0.122)	(0.052)	(0.038)	(0.107)	(0.044)
log $STOCK_{ijkt}$	0.320***	0.302***	0.405***	0.047***	0.124***	0.108***
	(0.028)	(0.024)	(0.027)	(0.009)	(0.023)	(0.014)
Fixed effects and other controls:						
Country pair	x	x	x	x		
Industry		x	x			
Time	x	x			x	
Exporter-time			x			x
Importer-time			x			x
Exporter-industry-time				x		
Importer-industry-time				x		
Industry-country pair					x	x
N	1,292,798	1,292,798	1,292,721	1,266,741	1,199,167	1,199,069

The dependent variable is industry-level (3 digit SITC rev 2) trade flows between 149 countries covering 1974-2006, using every 4 years. All estimations use PPML. Standard errors, shown in parentheses, are two-way clustered and by country-pair and exporter-industry. log $STOCK_{jikt}$ refers to (the log of) the stock of patents filed by the importer j in the exporting country i in industry k. log $STOCK_{ijt}$ refers to (the log of) the stock of patents filed by the exporting country i in the importing country j in industry k. The construction of these stocks is described in Section 4.2. "Time-varying trade cost controls" refers to three-way interactions between standard pairwise gravity controls (log distance, common language, and common legal system), the year, and the degree of R&D intensity in that that industry (high vs. medium vs. low R&D intensity). Observation counts vary because we first use the algorithms of Correia (2015) and Correia, Guimarães, and Zylkin (2019b) to reduce the sample by removing perfectly predicted observations.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.