

Fons-Rosen, Christian; Pu, Zhaoxin

Working Paper

The transmission of sectoral shocks across the innovation network

Discussion Paper, No. 229

Provided in Cooperation with:

University of Munich (LMU) and Humboldt University Berlin, Collaborative Research Center Transregio 190: Rationality and Competition

Suggested Citation: Fons-Rosen, Christian; Pu, Zhaoxin (2020) : The transmission of sectoral shocks across the innovation network, Discussion Paper, No. 229, Ludwig-Maximilians-Universität München und Humboldt-Universität zu Berlin, Collaborative Research Center Transregio 190 - Rationality and Competition, München und Berlin

This Version is available at:

<https://hdl.handle.net/10419/222126>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

The Transmission of Sectoral Shocks Across the Innovation Network

Christian Fons-Rosen (University of California, Merced)
Zhaoxin Pu (MPI-IC Munich)

Discussion Paper No. 229

January 22, 2020

The Transmission of Sectoral Shocks Across the Innovation Network

Christian Fons-Rosen^{a,b,c}

Zhaoxin Pu^{c,d}

^a University of California, Merced

^b Centre for Economic Policy Research (CEPR), London

^c Max Planck Institute for Innovation and Competition, Munich

^d Ludwig-Maximilians-University (LMU), Munich

January 19, 2020

-

PRELIMINARY DRAFT

DO NOT CITE OR CIRCULATE WITHOUT PERMISSION

ABSTRACT

Recent innovation literature has documented the benefits of cross-pollination of ideas across a wide set of industries and technology fields in an economy. Industrial and trade policies, by contrast, tend to favor economic specialization through the promotion of selected sectors. In this paper we use a firm-level panel of 13 European countries to assess whether an industry-specific policy propagates across the network of innovating firms through technological linkages. Following the competition shock to the European textile sector, triggered by the 2001 removal of import quotas on Chinese textiles, we find that patenting and knowledge sourcing behavior of non-textile firms are negatively affected. At the aggregate regional level, this indirect effect on non-textile firms can be around three to five times larger than the direct effect.

KEYWORDS: technological linkages, spillovers, patents, knowledge sourcing, industrial policy

JEL Classification: D57, L25, L60, O33, O38

We thank participants at the CREI Workshop 2017 at the University of Bath, the IO and Trade Seminar 2018 at the LMU Munich, the MPI Junior Researcher Workshop "From Science to Innovation" 2018, and the TUM Workshop "Innovation, Technological Change, and International Trade" 2019 for valuable comments. We thank Henning Deist and Christopher Schwarting for excellent research assistance. Zhaoxin Pu gratefully acknowledges financial support by Deutsche Forschungsgemeinschaft through CRC TRR 190 (project number 280092119) and Elite Network of Bavaria through the 'Evidence-Based Economics' Doctoral Program.

1 Introduction

Recent innovation literature has documented the benefits of cross-disciplinary research and the importance of recombining ideas from diverse fields,¹ in addition to the value of inventor teams with different academic backgrounds.² At the aggregate level, these findings speak in favor of countries increasing production diversification to facilitate cross-pollination of ideas across different sectors.

One example where the recombination of ideas across industries brought groundbreaking innovation is the development of fiber optics technology. Corning, a leader in specialty glass manufacturing, was approached by the British Post Office to explore manufacturing of optical glass fiber used for light-transmission in telecommunications (Gattani, 2006). Even though Corning did not have any prior experience in that sector, soon it developed the first low-loss optical fiber. This breakthrough invention paved the way for long-distance optical communication.

Another example for cross-pollination is the cavity magnetron, originally an essential component in the development of short wavelength radar during World War II. After the war, this palm-sized gadget producing high-power microwaves proved to be the perfect source for microwave energy capable of heating up food;³ the technology is also used for drying of paper, textiles, and even pharmaceutical powders.

Governments and science funding bodies worldwide acknowledge the importance of an integrative approach in tackling challenges through innovation, explicitly supporting interdisciplinary collaboration. The US defense agency DARPA is a prominent case, for example, by funding the first materials-focused interdisciplinary laboratories in the late 1960s.⁴ Most recently, DARPA launched a purpose-built social media platform aimed at facilitating the connection among experts across disciplines.⁵ Instead of being organized by scientific disciplines, the research portfolio of the German Fraunhofer-Gesellschaft is centred on issue-oriented questions that follow an interdisciplinary approach.⁶ Also in the Chinese National Science and Technology Development Plan (2006-2020), interdisciplinary research features as an important element of the innovation strategy; strengthening the

¹See for example Acemoglu et al. (2016), Jones (2009), Uzzi et al. (2013), Weitzman (1998).

²See for example Adams et al. (2005), Jones et al. (2008), Gruber et al. (2013).

³For more details on the history of the cavity magnetron, see e.g.: <https://spectrum.ieee.org/tech-history/dawn-of-electronics/from-world-war-ii-radar-to-microwave-popcorn-the-cavity-magnetron-was-there>, last accessed on 1 Nov 2019.

⁴<https://www.darpa.mil/news-events/2015-08-14?ppl=collapse>, last accessed on 1 Nov 2019.

⁵<https://www.darpa.mil/news-events/2019-03-19>, last accessed on 1 Nov 2019.

⁶Such as health, security, mobility or communication. <https://www.fraunhofer.de/en/about-fraunhofer/profile-structure.html>, last accessed on 1 Nov 2019.

scientific cooperation across industries and disciplines.^{7 8}

By contrast, industrial and trade policy typically target the development and growth of a specific sector. Consistent with traditional trade theory,⁹ trade openness and globalization are further processes that tend to increase a country's degree of industrial specialization. By downsizing unproductive sectors and relocating factors of production towards industries which enjoy a comparative advantage, a country ends up with a higher degree of industrial production concentrated in fewer sectors. This specialization can bring about numerous upsides, for example, higher wages, lower average production costs and economies of scale. However, to gauge the aggregate impact of government interventions, one should not only consider the sectors directly targeted by a policy, but take a broader view and also evaluate the effects on the rest of the economy.

This paper aims to evaluate the indirect effects of sector-specific policies by addressing the following question: When a particular industrial sector changes its innovation strategy in response to a competition shock, does this trigger further innovation adjustments in firms belonging to non-targeted industries? How does the magnitude of this indirect effect depend on the technological distance to the sector initially targeted by the shock?

We assess this question in the context of the European textile & clothing sector, following China's accession to the WTO and its removal of import quotas on textiles in 2001. Using firm-pair level variation in technological overlap, we study how patenting decisions by non-textile firms adjust to the changes experienced by the textile industry after the competition shock. The balance sheet information for our panel of 45,012 European non-textile patenting firms comes from combining Bureau Van Dijk's Orbis and Amadeus data sets. Variables on the patent behavior come from matching PATSTAT to our firm sample and include patent filings, forward and backward citations, in addition to technological information on each patent. This allows us to calculate technological distances at the bilateral firm level, similar to Bloom, Schankerman, and Van Reenen (2013).

Our main finding is that the average non-textile firm reduces patenting by 3% of a standard deviation after the textile firms, to which it is highly technologically connected, reduce their patenting by a standard deviation. Non-textile firms react twice as strongly to reductions in patenting by textile firms that are located in the same country, suggesting that

⁷National Medium- and Long-term Plan for the Development of Science and Technology, http://www.gov.cn/jrzq/2006-02/09/content_183787.htm; PRC Scientific and Technological Progress Act 2007, http://www.gov.cn/flfg/2007-12/29/content_847331.htm; last accessed on 1 Nov 2019.

⁸The EU's Horizon 2020, the NSF in the US, the British Royal Academies as well as the German science funding body DFG have also targeted programs for the promotion of interdisciplinary research. See e.g. https://www.nsf.gov/od/oia/additional_resources/interdisciplinary_research/support.jsp, <https://royalsociety.org/grants-schemes-awards/grants/apex-awards>, https://www.dfg.de/en/research_funding/programmes/coordinated_programmes/collaborative_research_centres/index.html; last accessed on 1 Nov 2019.

⁹For example, the Ricardian model, the Heckscher-Ohlin model, or Krugman (1979).

geographical proximity is an amplifying factor. Larger non-textile firms experience greater reductions than smaller ones, and results are robust to accounting for industrial input-output relationships. Similar findings are obtained for quality-adjusted patent counts.

We then analyze if firms redirect their innovation focus or change their knowledge sourcing behavior. Non-textile firms are less likely to patent in textile-intensive technology classes and tend to diversify across a wider set of fields. Furthermore, they cite fewer textile patents and start searching for new sources of knowledge in more distant geographical and technological spaces. Main results are economically and statistically very similar when estimating the model at the 4-digit industry level.

In the final exercise, we aggregate the data to the regional NUTS3 level. In the median region, textile firms obtain 4.1% more patents after the shock (*direct effect*); at the same time, non-textile firms reduce their patenting by 1.0% (*indirect effect*). To account for the fact that there are many more non-textile than textile firms in an economy, we also look at changes in absolute terms. In the median NUTS3 region, the negative indirect effect is around three times larger than the positive direct effect on patenting.

Our paper speaks to several literature strands. First, in the literature on direct effects of competition on innovation, Bloom, Draca, and Van Reenen (2016) explore China's opening up to globalization and find a positive effect on innovation by European textile firms. By contrast, Autor et al. (2019) find a negative effect on U.S. firms across a large variety of sectors. These differing results can be reconciled by the finding of Aghion et al. (2005) that competition and innovation have an inverted-U relationship. Our paper differs from the above in that we look at the indirect effect on other sectors by highlighting the importance of knowledge linkages. Other references on the direct effect of import competition on innovation and productivity include Pavcnik (2002) using Chilean data and Amiti and Konings (2007) on Indonesia. Papers incorporating an innovation dimension but rather looking at the effects of exports include Bustos (2011) using Argentinean data and Aw, Roberts, and Xu (2011) using Taiwanese data. Following a more macro approach, Cai, Santacreu, and Li (2017) develop a theoretical model of trade, innovation and knowledge diffusion to study the role of country and sector heterogeneity on aggregate R&D and welfare. For an extensive literature review on the general relationship between trade liberalization and innovation, we refer to Shu and Steinwender (2019).

Second, our work builds on the work by Bloom, Schankerman, and Van Reenen (2013) and Lychagin, Pinkse, Slade, and Reenen (2016); we also estimate firm-pair technological distances and examine how spillover effects vary with geographical distance. We extend this research by looking at a particular policy episode (China's accession to the WTO) and by quantifying the relative importance of the indirect effects. Finally, Acemoglu, Akcigit, and Kerr (2016) analyze citation properties of 1.8 million U.S. patents to map the innovation

network. When there is more past upstream innovation for a particular technology class to build on, then that technology class innovates more.

Third, we speak to the literature on the recombination of ideas. Seminal theoretical work by Levinthal and Cohen (1989) and Bernstein and Nadiri (1989) put at the forefront how firms enjoy knowledge spillovers coming from innovation undertaken in other firms. Weitzman (1998) provides micro-foundations for the knowledge production function, modelling the production of new ideas as a function of reconfigured old ideas. On the empirical front, Jones, Wuchty, and Uzzi (2008) show that a dramatic shift in the knowledge production has occurred, as teamwork in science produces the highest-impact papers and increasingly spans university boundaries. Uzzi et al. (2013) analyze 18 million papers spanning all scientific fields and find that the highest-impact papers insert novel features into otherwise conventional combinations of prior work. Zacchia (2019) constructs a network of companies to estimate the spillover effects along the linkages between inventors of different firms. Using a game-theoretically motivated IV approach, he finds that the marginal social return of R&D performed by a firm amounts to approximately 112% of the marginal private return. Griffith, Lee, and Straathof (2017) use patent citation data to quantify the relative cost of accessing new ideas created in your own versus in another firm. They show that firm boundaries lead to slower citation times, the delay being much larger than the delay due to national borders or technological distance. They use Thomson's Derwent codes to classify the patents according to the industries in which they are used. By contrast, we use the Cooperative Patent Classification codes as we explicitly want to identify the patents according to their technological similarity and not their relatedness in the product market.

Fourth, we relate to the cumulative innovation literature, as in our case building on previous knowledge generated by other industrial sectors is pivotal. Furman and Stern (2011) investigate how institutions shape the further use of existing knowledge; exploiting variation across biological resource centers, they find that institutions amplify the cumulative impact of scientific discoveries. Galasso and Schankerman (2015) and Gaessler, Harhoff, and Sorg (2019) investigate whether patent rights facilitate or impede follow-on innovation. Using US Court of Appeals data, Galasso and Schankerman (2015) find that patent invalidation results in a 50% increase of follow-on innovation, with the effect driven by invalidation of patents owned by large patentees that trigger more follow-on innovation by small firms. Gaessler, Harhoff, and Sorg (2019) find an increase of 20% using a large-scale dataset of European patent invalidations; the effect is most pronounced where the removal of exclusive rights creates substantial freedom to operate for third parties. Sampat and Williams (2019) investigate whether patents on human genes have affected follow-on scientific research and product development. They document that patented genes appear more valuable than non-patented genes, and that on average gene patents have no quanti-

tatively important effect on follow-on innovation. In line with this literature, we highlight the notion that firms build on previous knowledge embedded in patents owned by other parties.

Fifth, we relate to the literature on industrial policy and the specialization of economic activity. In a theoretical framework, Liu (2019) analyzes industrial policy when sectors are vertically linked through an input-output network. Market imperfections in one sector compound through backward linkages to upstream sectors. He shows that the ‘centrality’ of sectors in the production network matters, when policy makers decide on which industry to target. We follow a similar idea; instead of production networks, we look at innovation networks and assign more weight to firms that are technologically more ‘central’ in our empirical analysis. Feldman and Audretsch (1999) examine the effect of the composition of economic activity on innovation. Their descriptive analysis supports the thesis that a diverse set of economic activities is more conducive to innovation than a specialization to a narrow set of areas.

The remainder of this paper is structured as follows: Section 2 describes the empirical strategy. Section 3 provides details on the dataset, the construction of variables and shows descriptive statistics. Section 4 then presents the econometric analysis and a discussion of the results. Section 5 concludes.

2 Empirical Strategy

Our empirical specification estimates within-firm changes in patenting and knowledge sourcing of non-textile firms as a function of patenting changes of textile firms, weighted by the technological distance to each of these textile firms. Consider a basic firm-level OLS equation for patents of firm i in a non-targeted manufacturing industry s , country c , and year t as:

$$\ln Pat_{isc t}^{NTXT} = \beta \sum_{j, j \neq i} tech_{ij} \ln Pat_{jt}^{TXT} + \gamma_i + \gamma_{ct} + \gamma_{st} + u_{isc t} \quad (2.1)$$

The level of observation is at the (non-textile) firm-year level. We take the first difference to sweep out firm fixed effects and estimate:

$$\Delta \ln Pat_{isc t}^{NTXT} = \beta \sum_{j, j \neq i} tech_{ij} \Delta \ln Pat_{jt}^{TXT} + \Delta \gamma_{ct} + \Delta \gamma_{st} + \Delta u_{isc t} \quad (2.2)$$

where Δ denotes the long five-year difference operator. The dependent variable is the

within-firm five-year log change in patents by non-textile firms.¹⁰ The regressor of interest is the change in patenting by textile firms, weighted by their technological distance to the respective non-textile firm ($tech_{ij}$). The variation comes from the fact that the technological distance differs for each firm-pair. The empirical specification includes country-year and industry-year fixed effects to absorb country-specific and industry-wide shocks. We use overlapping five-year differences, so to maximize the use of our data, and cluster at the four-digit industry (SIC4) level.

A possible concern is an omitted variable bias by which textile and non-textile firms are likely to face an unobserved common technology shock if their patents belong to related technology fields. In other words, intellectual proximity across a pair of textile and non-textile firms can lead to similar changes in innovation strategy as a reaction to a common technology shock. Another obvious worry is reverse causality whereby changes in non-textile firms cause a reaction by textile firms.

We hence use an instrumental variable approach to address these concerns. We use the removal of import quotas on textile and clothing from China as an instrument for changes in the innovation output of European textile firms, in the spirit of Bloom, Draca, and Van Reenen (2016) (BDVR). Following China's accession to the WTO in 2001, these import quotas were abolished and caused a competition shock for the textile firms in Europe, affecting their domestic production as well as patenting activity. The underlying quotas vary at the four-digit industry level, and reflect the toughness of quotas at their level in 2000 prior to their abolishment. Our instrument is the technology-weighted toughness of quotas across the 2,380 textile firms, where again weights differ for each of our non-textile firms.

$$\begin{aligned} \sum_{j,j \neq i} tech_{ij} \Delta \ln Pat_{jt}^{TXT} &= -\lambda \sum_{j,j \neq i} tech_{ij} \Delta QUOTA_{jt}^{TXT} + \Delta \gamma_{ct} + \Delta \gamma_{st} + \Delta u_{isct} \\ &= \lambda \sum_{j,j \neq i} tech_{ij} QUOTA_{jt}^{TXT,2000} + \Delta \gamma_{ct} + \Delta \gamma_{st} + \Delta u_{isct} \end{aligned} \quad (2.3)$$

Identification hence comes from the instrumented change in patenting of textile firms and from variation across non-textile firms in their intellectual exposure to the textile sector. The exclusion restriction here is that shocks to the patenting and knowledge sourcing of *non-textile firms* are uncorrelated with the level of *textile quotas* that were determined in the 1950s-70s. This seems plausible, especially when considering that differences in quotas across four-digit textile industries reflect historic bargaining power of the respective industry in richer western economies when the quotas were introduced. For a more detailed discussion of the quotas instrument, we refer to BDVR.

¹⁰In alternative specifications, we additionally consider dependent variables that reflect changes in knowledge sourcing and in the direction of patenting.

A priori, there are two possible outcomes for the sign of the coefficient on the instrumental variable in equation (2.3). If $\lambda > 0$, it means that the larger the quota reduction, the stronger the import competition from China and the stronger is the *increase* in domestic textile patenting. While $\lambda < 0$ would imply that the larger the quota reduction, the stronger the *drop* in domestic textile patenting.¹¹

Testing for Alternative Mechanisms

Input-Output Relationships. A concern is that textile firms and non-textile firms that are technologically close may also be closely linked through vertical input-output industrial relationships, which may confound our regression results. One could for example think of financial dependence of a non-textile firm on the textile industries if its main customers are in the textile industries. The strength of a firm’s vertical relationship with the textile industries depends on two directions: on the one hand, the share of production outputs it supplies to the textile industries; on the other hand, the share of inputs it receives from the textile industries. We therefore determine each non-textile firm’s input and output exposure to the textile industries and test whether these factors drive any of our baseline results.

Industry Level Estimations. In addition to our within-firm analysis, we use industry-level regressions for non-textile SIC4 industries to better understand the indirect effects of the textile quota removal at a more macro level. As we use an unweighted estimation specification in our firm-level estimation, all non-textile firms in the sample are basically given the same weight. This could lead to the following scenario: suppose a four-digit industry consists of one very large firm that dominates the patenting activity of the industry and many smaller patenting firms. If all the small firms reduce patenting but the one very large firm increases patenting, at the firm level our regression analysis would suggest that patenting decreases. However, at the industry level aggregate patents may actually go up.

Also here, we account for possible vertical relationships through interdependencies in sales between a non-textile industry k and a textile industry l . As an additional regressor we include changes in textile sales at the four-digit level, weighted by the industry-pair specific input link, respectively output link, from a SIC4 input-output matrix (IO_{kl}). In absence of a second instrument for changes in textile sales, and to avoid two endogenous regressors, we run equation (2.4) in an OLS model.

$$\Delta \ln Pat_{kt}^{NTXT} = \beta_1 \sum_{l, l \neq k} tech_{kl} \Delta \ln Pat_{lt}^{TXT} + \beta_2 \sum_{l, l \neq k} IO_{kl} \Delta \ln Y_{lt}^{TXT} + \Delta \gamma_{SIC2D} + \Delta u_{kt} \quad (2.4)$$

¹¹This first stage differs in multiple dimensions for the baseline estimation in BDVR. Appendix A describes these differences in great detail.

Regional Effects. For example, employees of textile firms affected by the industrial policy could reallocate to non-textile firms with a similar technological focus in the same local labor market. In order to test for potential alternative channels at the regional level, we assess if geographical distance plays a role in explaining our results. More precisely, we estimate models where we only consider textile firms located in the same country or within a 50km radius of the non-textile firm. Contrasting these estimates to model estimates including textile firms in foreign countries or further away than 50km allows us to gauge if some labor reallocation mechanism is likely driving our results.

3 Data and Construction of Variables

We use Bureau van Dijk's (BvD) ORBIS global database, linked to the PATSTAT database, for our firm-level analysis. The ORBIS database is the largest cross-country firm-level database available and includes both public and private firms from all industries. It includes firm-level data on financial accounts, industry codes, and address data among many other information. We use the 2016 Fall version of BvD ORBIS, which includes all historical ORBIS vintages from 2005-2016.¹² We complement ORBIS with the 2006 vintage of BvDs Amadeus database, which includes firm financial data from 1995-2006, in order to improve data coverage for the late 1990s. Amadeus is a similar database of the same data provider BvD, covering firms in Europe rather than globally.

As Kalemli-Ozcan et al. (2015) discuss, it is advisable to combine different BvD vintages to obtain a consistent coverage of firms over time. We link ORBIS to Amadeus at the firm-year level via BvDs unique firm identifier (BvD-ID), while accounting for duplicate accounts, different currencies and accounting standards as well as possible BvD-ID changes over time. For the harmonization and cleaning of the ORBIS and Amadeus data, we broadly follow Kalemli-Ozcan et al. (2015). In the following, we describe the sample of manufacturing firms and the construction of variables used for our econometric analysis.

Sample of European Manufacturing Firms with Patenting Activity

Consistent with the previous literature (e.g. Bloom, Draca, and Van Reenen (2016), Autor et al. (2019)) we conduct our analysis for firms in the manufacturing sector. We use the four-digit SIC industry information in ORBIS to identify all firms that belong to the manufacturing sector in any of the 13 European countries of Austria, Denmark, France, Germany, Italy, the Netherlands, Norway, Poland, Portugal, Spain, Sweden, Switzerland, and the UK. From these approximately 2 million firms, we identify those that were active as of 2001,

¹²For representativeness of ORBIS data, see Kalemli-Ozcan et al. (2015).

the treatment year when the textiles and apparels import quotas were abolished. We use the incorporation date information in ORBIS where available, and otherwise deduce from non-negative entries for revenue or employees whether a firm has been active as of 2001. Keeping only these, we are left with about 1.6 million manufacturing firms. This includes patenting and non-patenting firms.

For our empirical analysis, we are interested in those firms that innovate and undertake patenting. We use ORBIS' embedded BvD-to-PATSTAT link to merge the firm data to patent data.¹³ About five percent of the above 1.6 million manufacturing firms have a link to the PATSTAT database. In order to calculate a firm's technological proximity to other firms based on its patent filings in the pre-period, we further need to impose the condition that firms in our sample patented at least once before 2001. The above steps result in a final sample for our estimations of 45,012 non-textile and 2,380 textile firms.

Import Quotas Data for the IV Approach

For the import quotas on textiles and apparels that were abandoned following China's entry into the WTO, we use data from BDVR. The quotas variable varies at the four-digit industry level and reflects the toughness of quotas at their initial level in 2000 prior to China's WTO entry in a given four-digit industry. It is calculated as the proportion of (import value-weighted) HS6-product categories that were covered by a quota within that four-digit industry. We refer to BDVR for a more detailed discussion of the historical origin of these quotas and the quota instrument.

The removal of these quotas had a direct effect on the textile firms. Based on the four-digit industry codes of each firm, we can calculate a firm-specific measure of the intensity of the quotas reduction each textile firm was facing.¹⁴

Calculating Pairwise Technological Distance between Two Firms

We calculate the technological proximity ($tech_{ij}$) between any non-textile firm and any textile firm in the sample based on the overlap in their patent portfolio. For each firm, we determine its patent portfolio as of 2001 and construct a vector of patent shares across

¹³For the matching of ORBIS and PATSTAT, Bureau van Dijk uses string similarity matching between a company name from ORBIS and the name of the patent applicant from PATSTAT, mapping BvD-IDs to each PATSTAT person_ID. Additional information like address information is used to enhance the matching precision. Person_IDs and patent filings are linked to each other via PATSTAT tables on patent applicants and application data.

¹⁴A firm can have multiple primary and secondary SIC 4-digit industry codes. We follow the approach of BDVR, applying a two-third weight to primary codes, a one-third weight to secondary codes, and equal weighting within these groups.

patent classes, which reflects the technological profile. We then calculate the pairwise un-centred correlation between any two firms' patent portfolio vectors.¹⁵

$$tech_{ij} = \frac{\sum_c PAT_{ic} * PAT_{jc}}{\sqrt{\sum_c PAT_{ic}^2} * \sqrt{\sum_c PAT_{jc}^2}} \in (0, 1) \quad (3.1)$$

In order to determine the technological profile of a firm's patent portfolio, we need to assign each patent to a unique patent technology class. In our main specifications, we use a technology classification that builds on 34 technology areas (TF34), aggregated from the IPC codes following the proposal by Schmoch (2008), which unambiguously assigns each patent to one of these technology areas. In an alternative specification, we use the Cooperative Patent Classification (three-digit CPC codes) to classify each patent into more granular technology classes. Under the three-digit CPC scheme there are 126 distinct technology classes.¹⁶ We present our main results based on the TF34 scheme. The results also hold when we use the CPC scheme for the calculation of the technological proximity measure and are hence deferred to the Appendix.

In addition to calculating the pairwise technological proximity of two firms based on the overlap of their patent filings, we also consider calculating it based on the overlap in their patent citation behavior. Based on their respective technology classes, we can calculate the share of citations a firm makes to a certain technology class. Similar to above, we can construct a vector of relative patent citations made to the different patent technology classes for each firm, and again calculate the pairwise un-centred correlation between any two firms citation behavior vectors.

As before, we consider patents filed prior to 2001 for calculating a firms technological profile. With respect to the patent-to-patent citations, we only consider citations from EPO patents that were filed with the EPO directly or under the PCT, so to have consistent citation behavior.¹⁷ In terms of cited patents, we allow for all patent authorities, though. Due to the restriction to EPO citing patents, our sample is somewhat smaller in those specifications where we use this second variant of measuring technological proximity based on citation data.

¹⁵For similar approaches, see Jaffe (1986), Bloom, Schankerman, and Van Reenen (2013), and Lychagin, Pinkse, Slade, and Reenen (2016).

¹⁶We consider only the main technology area or main three-digit CPC code of a patent, even if patents may be assigned to multiple technology classes. CPC codes distinguish between the position and hence importance of the different codes associated with one patent (i.e. cpc-position = F first or L later). We use this information to determine the unique technology class for each patent.

¹⁷See e.g. Alcácer, Gittelman, and Sampat (2009) and Bacchiocchi and Montobbio (2010) for a discussion of different patent citation behavior between USPTO and EPO.

Patent Filings

Our main dependent variable is the change in the number of patents filed by a firm. We consider patents at the DOCDB patent family level. In our paper, we refer to patent families interchangeably as patents. Our sample of 45,012 patenting non-textile firms filed around 615,000 patent families during the years 1996-2005, while the textile firms filed approximately 10,000 patent families during the same period.

As we aggregate patent applications to the patent family level, we need to take a few decisions as to how we unify patent attributes at the patent family level. The year is determined by the filing year of the patent member that was filed first within the family. For the technology class, CPC codes distinguish between the position and hence the importance of different technology classes associated with one patent (i.e. $cpc\text{-position} = F$ first or L later). Where this information is available, we prioritise the F codes, and consider the modal code in case there are multiple F codes. When using the TF34 scheme, we consider the modal technology area. In the event of ties, we use the numerically lowest technology area.

We use this technology class information also for assessing changes in the direction and diversity of a firm's patenting activities. Focusing on the patents filed by textile firms in our sample, we identify those technology classes where textile firms are more actively patenting. For each non-textile firm, we then calculate the fraction of patents it files in these relatively textile-intensive technology classes and assess whether there is any redirection of patenting towards or away from these technology classes.¹⁸ Furthermore, we calculate a Herfindahl-Hirschman-Index that reflects the level of specialization of a firm's patent portfolio in terms of filings by technology class. We use this index to assess if firms diversify more or less as part of their broader adjustments in patenting and knowledge sourcing decisions.

Patent Citation Data

As a measure of patent value, we use the number of forward citations received from EPO patents within five years after the first filing date (Harhoff, Narin, Scherer, and Vopel, 1999). We also use patent citation data to study the knowledge sourcing behavior of firms. Based on patent-to-patent citations we form dyads consisting of the citing firm and the cited firm. As set out above, we calculate the pairwise technological distance between any two firms in our sample. Similarly, we also calculate the pairwise geographical distance between any two firms in our sample (which we describe below). We can then analyze the technological

¹⁸More precisely, we weight the number of patents a firm files in a given technology class by multiplying it with an indicator that reflects the technology class textile-intensity. The indicator value ranges between 0 and 1 and is measured as the share of patents in this technology class filed in the pre-period prior to 2001 by textile firms (vs. non-textile firms).

vs. geographical distance to the cited firms and knowledge sources, and evaluate if firms have to travel further in the knowledge or in the geographical space when a given industry becomes less pivotal as a source of knowledge.

Geocoding & Geographical Distance between Two Firms

We have detailed address data from ORBIS, which allows us to geocode the location of the firms in our sample. We use the HERE Geocoder API and, where available, a firm's street name, ZIP code, city name and NUTS codes to obtain the corresponding longitude and latitude geo-coordinates.¹⁹ We then compute the geographical distance between any non-textile firm and any textile firm in our sample, using the STATA package 'geodist'.

Vertical Industrial Input-Output Relationship

As set out in Section 2, we want to account for potential confounding factors associated with vertical linkages between textile and non-textile firms. In order to capture a non-textile firms vertical input and output exposure to the textile industries, we use a SIC4 industry-level input-output matrix, as conventionally used in the literature,²⁰ and combine it with a firm's industry profile. If a firm has multiple industry SIC codes, we weight them as before: two-third weight to primary codes, a one-third weight to secondary codes, and equal weighting within these groups. For each firm, we calculate the share in output that it supplies to the textile industries (α), as well as the share in inputs that it sources from the textile industries (σ).

Descriptive Statistics

Table 1 provides descriptive statistics for our sample of non-textile firms. The median firm employs 52.5 workers, has annual revenues of 7.4 million Euros and total assets of 4.6 million Euros. It files 0.2 patents per year and the patent stock amounted to 2.0 patents as of 2001. The sample is skewed in terms of both firm size as well as patenting activity.

¹⁹For ca. 97% of the firms in our sample, we know the address at the street level; for 1% the info is at the ZIP or city level; for 2% we have no address info.

²⁰E.g. Javorcik (2004), Liu (2019).

Table 1: Summary Statistics

	mean	p5	p25	p50	p75	p95	count
No. of Employees	538.31	2.00	16.75	52.50	169.40	1,138.00	29,510
Revenue (th.)	120,285.28	250.00	2,150.00	7,372.25	26,896.50	240,868.70	29,338
Total Assets (th.)	13,8861.30	98.85	1,177.23	4,597.00	19,871.80	240,714.80	24,389
Patent Stock	29.03	0.20	1.00	2.00	7.20	53.00	45,012
Patent Filings p.a.	1.35	0.00	0.00	0.20	0.40	3.00	45,012
No. of Primary SIC Codes	1.50	1.00	1.00	1.00	2.00	4.00	45,012
No. of Secondary SIC Codes	1.07	0.00	0.00	0.00	1.00	5.00	45,012
Observations	45,012						

Notes: This table presents summary statistics for our sample of non-textile firms. The first three financial variables are provided by Bureau Van Dijk and approximately a third of the sample has missing values. Patent stock and patent filings come from PATSTAT, the industry SIC codes come from Bureau van Dijk; these variables are available for all firms.

4 Results

4.1 Baseline Results

OLS and Reduced Form Estimations. Table 2 displays the baseline OLS estimation results. The dependent variable is the within-firm 5-year log change in patents by non-textile firms. The regressor of interest is the tech-weighted average patenting change by each of the 2,380 textile firms. As the technological distance differs for each pair of firms, the weights vary for each non-textile firm. All columns control for country-specific macro shocks by including a full set of country dummies interacted with a full set of time dummies. In Column (1) we find an elasticity of 1%, meaning that a decrease in (tech-weighted) textile patenting by one standard deviation is associated with a 1% of a standard deviation decrease in non-textile patenting. Results remains stable in Columns (2) and (3) after adding a full set of 2-digit and 4-digit industry dummies interacted with a full set of time dummies, respectively.

Table 3 presents the reduced form in which the toughness of quota removals is directly regressed on the same dependent variable as in Table 2, namely, the within-firm 5-year log change in patents by non-textile firms. Results remain stable across all three columns: a standard deviation increase in quota toughness is associated with a 2% of a standard deviation decrease in patenting by a given non-textile firm.

Instrumental Variable Estimations. Table 4 presents IV results using Chinas WTO accession. The uneven numbered columns report the first stage results and the even numbered columns report the second stage results. In this panel of non-textile firms, the endogenous variable is the tech-weighted average change in patents by textile firms, and the instrument is the tech-weighted toughness of quotas faced by textile firms in 2000. The observed negative coefficient in Column (1) implies that the removing of tougher quotas is related

Table 2: OLS

DepVar: dlnPAT of NTXT firms	(1)	(2)	(3)
dlnPAT of TXT firms	0.0102*** (0.004)	0.0112*** (0.004)	0.0098*** (0.003)
Country-Year FE	Yes	Yes	Yes
Industry(SIC-2D)-Year FE	No	Yes	No
Industry(SIC-4D)-Year FE	No	No	Yes
No. of clusters	471	471	471
Observations	225,060	225,060	225,060
Unique Firms	45012	45012	45012

Notes: ***/**/* indicate significance at the 1%/5%/10% level. Standard errors are clustered at the SIC4 industry level. This table presents the baseline OLS results for the full panel of non-textile firms. The dependent variable is the change in the log of patents by non-textile firms. The regressor of interest is the change in the log of patents by textile firms, weighted by the technological proximity between a given non-textile firm and each textile firm. To clarify, given that we have 2,380 textile firms in our sample, each regressor is a weighted average of 2,380 changes in the log of patenting, and the closer the technological proximity to the non-textile firm, the greater the weight assigned. Equation (1) is the baseline specification that controls for country-year fixed effects. Equations (2) and (3) additionally include a full set of year dummies interacted with a full set of industry 2-digit and 4-digit dummies, respectively.

Table 3: Reduced Form

DepVar: dlnPAT of NTXT firms	(1)	(2)	(3)
Toughness of quotas in 2000	-0.0196*** (0.003)	-0.0191*** (0.003)	-0.0176*** (0.003)
Country-Year FE	Yes	Yes	Yes
Industry(SIC-2D)-Year FE	No	Yes	No
Industry(SIC-4D)-Year FE	No	No	Yes
No. of clusters	471	471	471
Observations	225,060	225,060	225,060
Unique Firms	45012	45012	45012

Notes: ***/**/* indicate significance at the 1%/5%/10% level. Standard errors are clustered at the SIC4 industry level. This table presents the reduced form results for the full panel of non-textile firms. The dependent variable is the change in the log of patents by non-textile firms. The regressor of interest is the quota toughness prior to Chinas accession to the WTO for each textile firm, weighted by the technological proximity between a given non-textile firm and each textile firm. To clarify, given that we have 2,380 textile firms in our sample, each regressor is a weighted average of 2,380 quota changes, and the closer the technological proximity to the non-textile firm, the greater the weight assigned. Equation (1) is the baseline specification that controls for country-year fixed effects. Equations (2) and (3) additionally include a full set of year dummies interacted with a full set of industry 2-digit and 4-digit dummies, respectively.

to stronger reductions in textile patenting.²¹ Column (2) presents the second stage results that show a strong and significant effect of (instrumented) reductions in *textile* patenting on reductions in *non-textile* patenting with a magnitude of 3.6%. Similar to previous tables, the next columns incorporate industry-year dummies leading to only minor changes.

The test statistics for under-identification (Kleibergen-Paap rk LM statistic) and weak identification (Kleibergen-Paap rk Wald statistic) show that the first stage is very strong in all cases. The IV results in Table 4 indicate that the OLS coefficient appears downward biased. Given the robust IV results across the board, we prefer the third specification with country-year and SIC4 industry-year fixed effects and will consider the model in Column (6) as our baseline specification going forward. Appendix Table 12 shows that our baseline results also hold when we condition on a sample of firms for which we have financial data in ORBIS.

Table 4: IV - First and Second Stage

DepVar: dlnPAT of NTXT firms						
Method	(1) 1st stage	(2) IV	(3) 1st stage	(4) IV	(5) 1st stage	(6) IV
Toughness of quotas in 2000	-0.5341*** (0.004)		-0.5325*** (0.004)		-0.5340*** (0.004)	
dlnPAT of TXT firms		0.0367*** (0.006)		0.0360*** (0.006)		0.0329*** (0.006)
Country-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry(SIC-2D)-Year FE	No	No	Yes	Yes	No	No
Industry(SIC-4D)-Year FE	No	No	No	No	Yes	Yes
Underidentification test		59.8		54.1		51.9
Weak identification test		14,991.6		14,484.1		15,609.7
No. of clusters	471	471	471	471	471	471
Observations	225,060	225,060	225,060	225,060	225,060	225,060
Unique Firms	45012	45012	45012	45012	45012	45012

Notes: ***/**/* indicate significance at the 1%/5%/10% level. Standard errors are clustered at the SIC4 industry level. This table presents the instrumental variable estimation results for the full panel of non-textile firms. The dependent variable is the same as in Tables 2 and 3. The instrument is the weighted change in the quota toughness prior to Chinas accession to the WTO for each textile firm, as described in Table 3. The endogenous regression is the weighted change in the log of patents by textile firms, as described in Table 2. Columns (1), (3), and (5) presents first stage results, while columns (2), (4), and (6) present second stage results. Equations (1) and (2) are the baseline specification that control for country-year fixed effects. Equations (3) and (4) additionally include a full set of year dummies interacted with a full set of industry 2-digit dummies. Equations (5) and (6) instead include a full set of year dummies interacted with a full set of industry 4-digit dummies. The table reports test statistics for underidentification (Kleibergen-Paap rk LM statistic) and weak identification (Kleibergen-Paap rk Wald statistic).

²¹To reconcile this negative coefficient with the results in BDVR in which the removal of tougher quotas leads to more patenting by textile firms, we refer to Table 11 in the Appendix.

4.2 Heterogeneity & Robustness

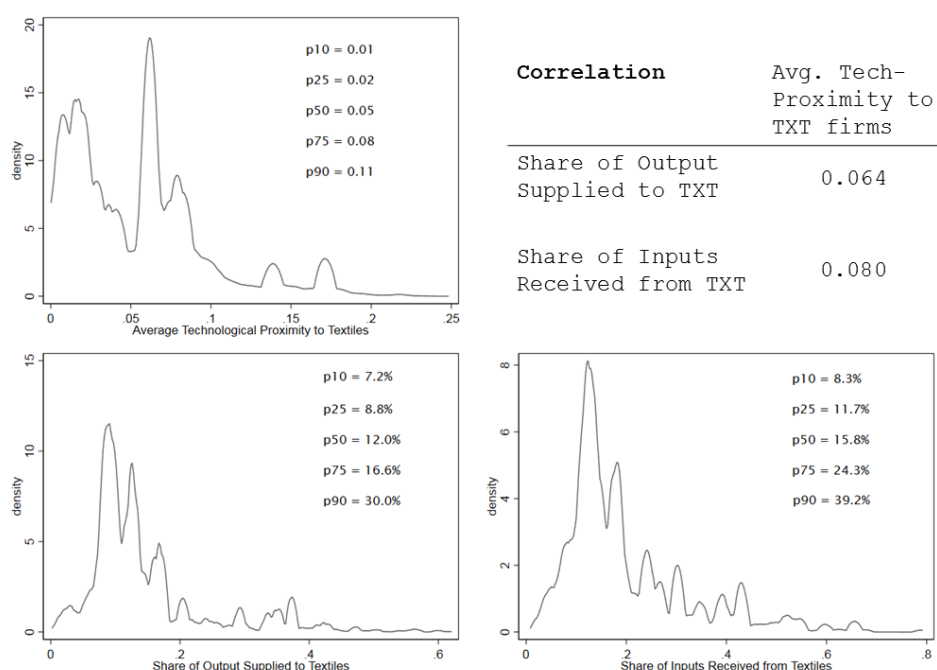
Table 5: Geographical Heterogeneity

DepVar: Reduction in PAT of NTXT firms	(1)	(2)	(3)	(4)	(5)
Reduction in PAT of TXT firms: all	0.0329*** (0.006)				
Red. in PAT of TXT firms: other countries		0.0324*** (0.006)			
Red. in PAT of TXT firms: same country			0.0575*** (0.011)		
Red. in PAT of TXT firms: same country, <50km				0.0570*** (0.019)	
Red. in PAT of TXT firms: same country, >50km					0.0516*** (0.011)
Country-year FE	Yes	Yes	Yes	Yes	Yes
Industry(SIC-4D)-year FE	Yes	Yes	Yes	Yes	Yes
Underidentification test	51.9	53.3	47.5	26.6	46.8
Weak identification test	15,609.7	6,757.8	1,177.6	203.3	1,162.5
No. of clusters	471	471	471	452	462
Observations	225,060	225,060	225,060	204,400	216,135
Unique Firms	45012	45012	45012	40880	43227

Notes: ***/**/* indicate significance at the 1%/5%/10% level. Standard errors are clustered at the SIC4 industry level. This table introduces geographical variation into the analysis. Presented are the second stages of IV estimations where each regressor differs in its geographical scope. The dependent variable is the same as in Tables 2 and 3. The endogenous regressor is the weighted change in the log of patents by textile firms, as described in Table 2. In Column (1) the regressor encompasses all textile firms at its weighted average; in Columns (2) and (3) it is limited to the sample of textile firms outside, or respectively inside, the same country as the non-textile firm. Conditional on only accounting for textile firms in the same country, Column (4) restricts the sample to the ones within a 50km radius of the non-textile firm; finally, Column (5) restricts the sample to the ones outside of a 50km radius of the non-textile firm. All specifications include country-year fixed effects as well as a full set of year dummies interacted with a full set of industry 4-digit dummies.

Geographical Heterogeneity. In this subsection we describe how the intensity of the observed reduction in non-textile patenting depends on geographical distance. For illustration purposes, Column (1) of Table 5 repeats the result of Column (6) in Table 4 - our baseline result - and is based on a tech-weighted average of *all* textile firms across our sample of European firms. In the next two columns we split these textile firms into two groups based on whether they are located in the same country as the non-textile firm or not. For the construction of the tech-weighted average textile patent changes in the regressor, Column (2) includes textile firms in other European countries and Column (3) only includes textile firms within the same country as the non-textile firm. Interestingly, the estimated coefficient almost doubles in the latter case, suggesting that geographical proximity amplifies the impact that textile firms have on non-textile firms, conditional on a given technological distance. One could expect to possibly find heterogeneous effects within the same coun-

Figure 1: Technological Proximity vs. Vertical Production Links in the Input/Output Matrix



Notes: This figure plots the distribution of the technological proximity and the vertical input & output exposure measures across the non-textile firms. The top left panel shows the histogram for the technological proximity, the bottom left panel shows the histogram for the share of production outputs supplied to textile industries, and the bottom right panel shows the histogram for the share of inputs received from textile industries. The table in the top right panel correlates our measure of technological proximity with the two measures of vertical production linkages.

try, for example of local labor markets allowing for the reallocation of affected employees across firms with a similar technological focus. Columns (4) and (5) split the textile firms into being closer or further away than 50 kilometres, respectively. Coefficients are similar and do not seem to be consistent with any major labor reallocation mechanism.

Accounting for Input-Output Relationships. A concern is that our results are driven by industrial input-output relationships that correlate with our technological distances. Figure 1 displays histograms for the share of production outputs each non-textile industry supplies to textile industries and also for the share of inputs received from textile industries. The correlation between these two measures and our technological proximity measure is positive and small (0.064 and 0.080, respectively) which mitigates the possibility that our results are driven by this alternative mechanism.

Nonetheless, Table 6 explicitly accounts for these input-output relationships in our econometric framework. As usual, Column (1) is the baseline. Columns (2) and (3) split non-textile firms by the median output exposure to the textile sector, while Columns (4) and (5) redo the same exercise for the median input exposure. We reassuringly do not

observe any substantial variation in the magnitude of the estimated coefficient, meaning that the industrial network channel seems close to orthogonal to our story. Columns (6) and (7) compare non-textile firms with weak input and output links (below 25th percentile on both dimensions) to the textile industry to the ones with strong input and output links (above 75th percentile on both dimensions). While both results remaining statistically and economically significant, if anything, the firms with the weak links to textile firms have a larger estimated coefficient.

Table 6: Accounting for Input-Output Relationships

DepVar: Reduction in PAT of NTXT firms	Base		Output (alpha)		Input (sigma)		Input & Output	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	
		<p50	>p50	<p50	>p50	<p25	>p75	
Reduction in PAT of TXT firms	0.0329*** (0.006)	0.0297*** (0.010)	0.0331*** (0.007)	0.0336*** (0.009)	0.0313*** (0.008)	0.0493** (0.022)	0.0207* (0.013)	
Country-year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
Industry(SIC-4D)-year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
Underidentification test	51.9	48.0	25.1	35.3	24.1	25.8	25.4	
Weak identification test	15,609.7	9,071.9	10,542.9	9,882.1	11,028.0	6,551.1	4,663.3	
No. of clusters	471	303	288	296	298	167	168	
Observations	225,060	110,220	107,260	109,760	108,255	25,515	41,120	
Unique Firms	45012	22044	21452	21952	21651	5103	8224	

Notes: ***/**/* indicate significance at the 1%/5%/10% level. Standard errors are clustered at the SIC4 industry level. In this table we account for the impact of industrial input-output relationships in the production network. The dependent variable and regressor of interest are as in Table 2. Column (1) repeats the baseline IV result of Table 4 Column (6). Columns (2) and (3) split the non-textile firms by the median output exposure to the textile sector, while Columns (4) and (5) split them by the median input exposure to textiles. Columns (6) and (7) estimate the effect for the non-textile firms with the weakest vs. the strongest input and output exposure to textiles (below 25th percentile vs. above 75th percentile on both dimensions).

Firm Size Heterogeneity, Patent Quality & Lag Specifications. We also check for firm size heterogeneity and find that the effect increases with the size of the firm. Appendix Table 13 reports the IV results by the quartile of size; the effect for the largest firms (4th quartile) is approximately four times larger than for the smallest firms (1st quartile). The distribution of patent quality is known to be skewed, with only a subset of patents having significant market value. We therefore re-estimate our baseline model, restricting the patent count to the subset filed with the European Patent Office (EPO) and using a cite-weighted patent count. Appendix Table 14 shows that the main result above still holds, meaning that the effect is not driven by changes in patent quality. We further test alternative models that use lagged regressors. The estimated coefficients in Appendix Table 15 show robust results, with larger effects for the two-year lag of the reduction in textile patents.

4.3 Patent Direction and Citation Behavior

Up to now, we have seen that non-textile firms reduce their patenting after an exogenous reduction in patenting by textile firms, and it does not seem to be driven by input-output relationships. The magnitude of the effect is increasing in the technological proximity to the patenting undertaken by textile firms. One mechanism consistent with this result is that non-textile firms have lost a source from which they acquired innovation expertise. We explore this possible channel in Table 7. As set out in Section 3, we identify those technology classes with the largest share of patents originating from textile firms. Column (1) is our baseline result and in Column (2) we find that the reduction in patenting by textile firms leads non-textile firms to move away from these textile-intensive technology classes and refocus towards technology areas where textile firms are less prevalent. At the same time, in Column (3) the HHI concentration index decreases, meaning that non-textile firms diversify more across a larger set of technology classes.

Table 7: Patenting Direction: 'Textile-intensity' and Diversity of Tech Classes

DepVar: Reduction in	(1) PAT filings (baseline)	(2) Share of PAT in 'txt-int.' classes	(3) HHI (1/diversity)
Reduction in PAT of TXT firms	0.0329*** (0.006)	0.1421*** (0.012)	0.0354*** (0.008)
Country-year FE	Yes	Yes	Yes
Industry(SIC-4D)-year FE	Yes	Yes	Yes
Underidentification test	51.9	51.9	51.9
Weak identification test	15,609.7	15,609.7	15,609.7
No. of clusters	471	471	471
Observations	225,060	225,060	225,060
Unique Firms	45012	45012	45012

Notes: ***/**/* indicate significance at the 1%/5%/10% level. Standard errors are clustered at the SIC4 industry level. Column (1) is our baseline result and the dependent variable in Column (2) is the share of patents granted in textile-intensive technology areas. Column (3) has the HHI concentration index as dependent variable. All specifications include country-year fixed effects as well as a full set of year dummies interacted with a full set of industry 4-digit dummies.

Table 8 also addresses the refocusing of innovation efforts by non-textile firms from a different angle. Consistent with the storyline of previous findings, the result of Column (1) tells us that non-textile firms are less likely to cite patents of textile firms after the China WTO accession. The remaining columns then analyze whether non-textile firms have to look in more distant locations (both technologically and geographically speaking) to partially substitute for the knowledge lost from textile firms. The negative estimated coefficients in Columns (2) and (3) mean that non-textile firms start citing more technologically distant patents, with results mainly driven by large firms. Columns (4) and (5) report that, for the case of large firms, non-textile firms also start citing patents from more geograph-

ically distant firms. Overall, the results of this table are consistent with non-textile firms trying to mitigate the loss of intellectual inspiration from the textile firms by exploring new sources of knowledge that are in more distant technological and geographical areas.

Table 8: Citation Behaviour

DepVar: Change in	% Cit to TXT	TechDist		GeoDist	
	(1) all	(2) all	(3) large firms	(4) all	(5) large firms
Reduction in PAT of TXT firms	0.0637** (0.031)	-0.0070 (0.030)	-0.0884** (0.041)	0.0032 (0.028)	-0.0988** (0.040)
Country-Year FE	Yes	Yes	Yes	Yes	Yes
Industry(SIC-4D)-Year FE	Yes	Yes	Yes	Yes	Yes
Underidentification test	36.5	30.7	19.4	31.3	20.4
Weak identification test	2,084.4	1,131.4	1,344.5	1,268.2	1,274.7
No. of clusters	220	186	151	182	148
Observations	11,649	8,163	4,518	7,902	4,366
Unique Firms	5567	3895	1742	3760	1684

Notes: ***/**/* indicate significance at the 1%/5%/10% level. Standard errors are clustered at the SIC4 industry level. This table addresses the refocusing of citation behavior. In Column (1), the dependent variable is the share of citations made to textile firms. Columns (2) and (3) look at the technological distance to cited patents, for the full sample and for the 25% of largest non-textile firms, respectively. Columns (4) and (5) repeat the exercise but with geographical distance to cited patents instead. All specifications include country-year fixed effects as well as a full set of year dummies interacted with a full set of industry 4-digit dummies.

4.4 Some ‘Macro’ Considerations: Industry-Level Analysis & Aggregate Regional Effects

Industry-Level Regressions Accounting for Vertical Relationships. In all the previous tables we conducted a within-firm analysis of non-textile companies in which all firms are given the same weight. In the industry regressions in Table 9, more weight is assigned to large firms. Consequently, industrial-level estimations might differ if, for example, large and small firms respond differently to reductions in textile patenting. Furthermore, from a policy perspective it is relevant to understand the evolution of non-textile patenting at the industry level in addition to the firm level. In the following table, we assign each firm to its main SIC4 industrial category, and then estimate the 5-year within-industry changes in non-textile patenting instead of the within-firm changes.²²

Table 9 presents OLS estimations for a panel of SIC4 industries over time, and all estimations include a complete set of SIC2 dummies interacted with a complete set of year

²²Additionally, by Jensen’s Inequality, the log of industry-level sum of patents is not the same as the sum of the firm-level log of patents (summed to the industry level).

Table 9: Industry-Level (SIC4) Regressions: OLS

	(1)	(2)	(3)	(4)
DepVar: \ln PAT of NTXT industries (SIC4)	Base	Output (α)	Input (σ)	Input & Output
\ln PAT of $tech_{ij}$ -weighted TXT ind.	0.0366* (0.020)	0.0360* (0.020)	0.0366* (0.020)	0.0358* (0.020)
\ln Y of $IO_{ij}\alpha$ -weighted TXT ind.		4.8856 (6.236)		7.1522 (8.039)
\ln Y of $IO_{ij}\sigma$ -weighted TXT ind.			1.0258 (6.966)	-4.0132 (8.978)
Industry(SIC-2D)-Year FE	Yes	Yes	Yes	Yes
Observations	2,065	2,065	2,065	2,065
Unique SIC4 Industries	413	413	413	413

Notes: ***/**/* indicate significance at the 1%/5%/10% level. Standard errors are clustered at the SIC4 industry level. This table presents industry-level regressions, including 2-digit industry-year fixed effects in all specifications. The dependent variable is the change in the log of patents by non-textile firms in each 4-digit industry. The regressor of interest is the change in the log of patents by textile firms, weighted by the technological proximity between a given non-textile 4-digit industry and each textile 4-digit industry. Similar to Table 7, Columns (2)-(4) control for industrial input-output relationships.

dummies. Column (1) presents the baseline result where the main regressor of interest is now the 5-year log change in tech-weighted patenting in textile industries; the new technological distances are constructed across industry pairs as opposed to firm pairs. A decrease in (tech-weighted) textile patenting by one standard deviation is associated with a 3.66% of a standard deviation decrease in non-textile patenting. This result, in which the estimated coefficient is three times larger than in our firm-level sample, is consistent with our previous finding in Table 13 that large firms reduce their non-textile patenting by a greater degree than small firms.

The remaining columns aim to account for vertical linkages. In Column (2) we add a regressor capturing changes in sales by the textile sector, weighted by how much a non-textile industry supplies its output to the textile sector. In Column (3) we generate the equivalent variable, but now weighted by how relevant textile sector products are as inputs for each of the non-textile industries. Column (4) incorporates both variables. Estimation results for our regressor of interest are unchanged, supporting the idea that industrial linkages cannot explain our finding.

Aggregate Magnitudes at the Regional Level. In Figure 2, we plot the 13 European countries in our data by the concentration of textile firms and by the relative magnitude of the direct effect in each region. Panel (a) shows the ratio of textile to non-textile firms in each NUTS2 region. The largest concentration of textile firms is in Southern Europe (Spain, Portugal and Italy) and in Poland. Nonetheless, the western coast of Scandinavia, the northern part of the United Kingdom and some areas in France and Germany also have a relatively high ratio of textile firms. Panel (b) displays the relative magnitude of the direct

effect compared to the indirect effect. While there is a positive correlation with Panel (a), strong direct effects are not confined to those regions with a high concentration of textile firms but observable across Europe. This can be explained by two factors. First, the intensity of quotas removal differed across 4-digit industries within the textile sector. Second, the various degrees of the indirect effect on non-textile firms reflect their different technological exposure to textile firms.

Table 10 compares the direct and indirect effects of China’s accession to the WTO in a five-year window. We limit our sample to NUTS3 regions with a presence of textile firms (247 regions) and condition on firms with non-missing financial data in ORBIS. In Panel (A) we start with a micro approach by documenting the predicted percentage change in annual patent filings at the firm level. To estimate the direct effect, we generate a panel of textile firms over time, as in BDVR.²³ The predicted value implies a 5.7% (4.6%) increase in patenting by the mean (median) textile firm after China’s entry into the WTO. For an estimation of the indirect effect, we again use the panel of non-textile firms to which we have been referring to throughout the paper. The predicted value of the indirect effect states that the mean (median) non-textile firm reduces its patenting by 1.0% (0.7%) after China’s entry. Overall, it is reassuring that the direct effect is much larger than the indirect effect (between six and seven times).

Table 10: Direct vs. Indirect Effects at the Regional (NUTS3) Level


Predicted change in patent filings over 5-years (\hat{y})	Mean	Median
<i>Panel A: Percentage change at the firm level</i>		
Direct effect	0.057	0.046
Indirect effect	-0.010	-0.007
<i>Panel B: Percentage change at the NUTS3 level</i>		
Direct effect	0.048	0.041
Indirect effect	-0.010	-0.010
<i>Panel C: Level change at the NUTS3 level</i>		
Direct effect	0.458	0.122
Indirect effect	-2.419	-0.341
Ratio of indirect to direct effect	5.3 x	2.8 x

In Panels (B) and (C) we aggregate the data to the regional NUTS3 level, the former displaying percentage changes and the latter presenting level changes. The mean (median) NUTS3 region experiences a 4.8% (4.1%) increase in patents by textile firms after the shock (*direct effect*) and, at the same time, a 1.0% (1.0%) reduction in non-textile patenting (*indirect effect*). Magnitudes are reassuringly similar to firm-level estimates in Panel (A), implying that our firm-level results are not driven by few very large firms or by specific

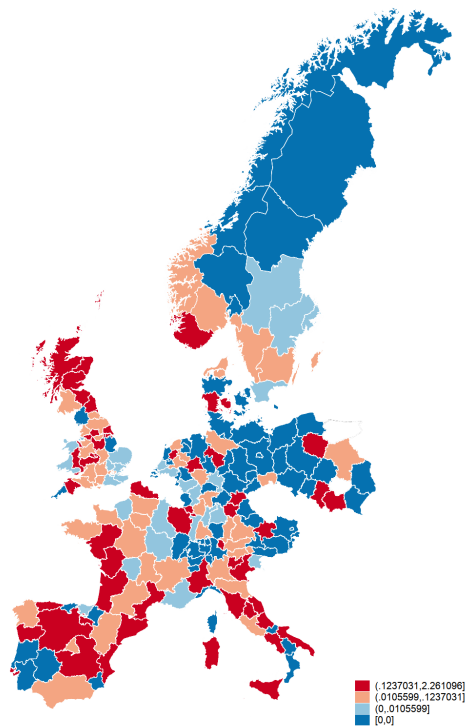
²³See Appendix Table 11 for a replication of the BDVR results with the textile firms in our sample.

Figure 2: Aggregate Effects at the Regional (NUTS2) Level

(a) Number of Textile vs. Non-Textile Firms



(b) Direct vs. Indirect Effects



Notes: The Figure maps the regions at the NUTS2 regional level for illustrative reasons, while our underlying analysis in Table 10 is undertaken at the more disaggregate NUTS3 regional level.

regions.

Both panels (A) and (B) present percentage changes, within firm and within region, respectively. In both cases, the direct effect was an order of magnitude larger than the indirect effect. To account for the fact that there are many more non-textile than textile firms in an economy, Panel (C) displays predicted level changes in the *absolute* number of patents after the China shock. In the average (median) NUTS3 region, textile firms produce an additional 0.458 (0.122) patents (*direct effect*), while non-textile firms reduce their patenting by 2.419 (0.341) units (*indirect effect*). Consequently, once we account for the relative abundance of non-textile firms in an economy, the indirect effect can be around three to five times larger than the direct effect.

5 Conclusion

Recent academic work seems to provide different viewpoints on industrial policy: while the innovation literature tends to emphasize the benefits of cross-pollination of ideas across a variety of sectors and the benefits of being exposed to other technology classes, traditional work in the trade literature rather highlights the gains from specialization. This paper addresses this debate by estimating and quantifying the impact of an industry-specific policy shock on the innovation undertaken in technologically proximate but non-targeted industries.

We construct a firm-level manufacturing panel of 13 European countries with information on innovation activity, knowledge sourcing, and technological distances across each pair of firms and sectors. We use the removal of import quotas on Chinese textiles in 2001 as an exogenous competition shock to the European textile sector to help identify the induced changes in innovation of non-textile firms through technological linkages to textile firms (Bloom, Schankerman, and Van Reenen, 2013).

Our key result is that the shock induced by the removal of Chinese import quotas propagates through technological linkages across the network of innovating non-textile firms. While the direct effect of this removal increases innovation by the average European textile firm (Bloom, Draca, and Van Reenen, 2016), the indirect effect is negative once we account for the centrality of each textile firm in the knowledge network. This negative effect increases in geographical and technological proximity to textile firms. The results are robust when accounting for vertical input-output linkages to the textile sector. The effect persists across all firm sizes and upholds when aggregating the panel to the 4-digit industry level, suggesting that results are not driven by a few large firms. Moreover, our analysis shows that non-textile firms shift their innovation away from 'textile-intensive' technology areas, cite less patents from textile firms, and instead turn to new sources of knowledge that are

further away in both the geographical and technological space.

Our results highlight the importance of accounting for indirect effects when evaluating the implications of industrial or trade policies that target specific sectors. The absolute magnitude of the indirect effects is sizeable once we account for the larger fraction of non-textile compared to textile firms. In the median NUTS3 region, the indirect effect is around three times larger than the direct effect. In terms of policy, our analysis shows that the impact of trade policy on economy-wide innovation needs to account not only for the direct effect but also for the indirect effects on non-targeted firms. By stressing that not all textile firms have the same centrality index in the network and by explicitly modelling knowledge transfers across firms and industries, we aim to incorporate some of the general equilibrium effects of an industry-specific shock.

We highlight three limitations that could be explored in future work. First, one could consider structurally estimating a general equilibrium model that explicitly accounts for knowledge transfers across industries and that incorporates technological linkages at various degrees. Second, our data do not allow us to undertake a detailed assessment of employment effects. To the extent possible, we test for potential labor reallocation within the local labor market as an alternative mechanism through our geographical analysis. Our estimated coefficients based on firms within a 50km radius (as proxy for the local labor market) versus firms outside a 50km radius are very similar to our baseline. This tentatively speaks against a major labor reallocation of inventors. Future research using worker-level data, which may account for labor mobility, could provide additional insights. Third, our empirical strategy uses a policy shock that targeted the European apparel and textile sector, which is not among the most innovative sectors. While our results can be considered a 'lower bound', further analysis exploring other types of industrial policy shocks would be valuable.

References

- Acemoglu, D., U. Akcigit, and W. R. Kerr (2016). Innovation network. *Proceedings of the National Academy of Sciences of the United States of America* 113(41), 11483–11488.
- Adams, J. D., G. C. Black, J. R. Clemmons, and P. E. Stephan (2005). Scientific teams and institutional collaborations: Evidence from U.S. universities, 1981-1999. *Research Policy* 34(3), 259–285.
- Aghion, P., N. Bloom, R. Blundell, R. Griffith, and P. Howitt (2005). Competition and Innovation: An Inverted-U Relationship. *The Quarterly Journal of Economics* 120(May), 701–728.
- Alcácer, J., M. Gittelman, and B. Sampat (2009). Applicant and examiner citations in U.S. patents: An overview and analysis. *Research Policy* 38(2), 415–427.
- Amiti, B. M. and J. Konings (2007). Trade Liberalization, Intermediate Inputs, and Productivity: Evidence from Indonesia. *American Economic Review* 97(5), 1611–1638.
- Autor, D., D. Dorn, G. H. Hanson, P. Shu, and G. Pisano (2019). Foreign Competition and Domestic Innovation: Evidence from U.S. Patents. *American Economic Review: Insights* (forthcoming).
- Aw, B., M. J. Roberts, and D. Y. Xu (2011). R&D Investment, Exporting, and Productivity Dynamics. *American Economic Review* 101(4), 1312–1344.
- Bacchiocchi, E. and F. Montobbio (2010). International Knowledge Diffusion and Home-bias Effect: Do USPTO and EPO Patent Citations Tell the Same Story? *Scandinavian Journal of Economics* 112(3), 441–470.
- Bernstein, J. I. and M. I. Nadiri (1989). Research and Development and Intra-industry Spillovers: An Empirical Application of Dynamic Duality. *The Review of Economic Studies* 56(2), 249.
- Bloom, N., M. Draca, and J. Van Reenen (2016). Trade induced technical change? The impact of chinese imports on innovation, IT and productivity. *Review of Economic Studies* 83(1), 87–117.
- Bloom, N., M. Schankerman, and J. Van Reenen (2013). Identifying Technology Spillovers and Product Market Rivalry. *Econometrica* 81(4), 1347–1393.
- Bustos, P. (2011). Trade liberalization, exports, and technology upgrading: Evidence on the impact of MERCOSUR on argentinian firms. *American Economic Review* 101(1), 304–340.

- Cai, J., A. M. Santacreu, and N. Li (2017). Knowledge Diffusion, Trade and Innovation Across Countries and Sectors. *FRB St. Louis Working Paper* (No. 2017-29).
- Cattani, G. (2006). Technological pre-adaptation, speciation, and emergence of new technologies: How Corning invented and developed fiber optics. *Industrial and Corporate Change* 15(2), 285–318.
- Feldman, M. P. and D. B. Audretsch (1999). Innovation in cities: Science-based diversity, specialization and localized competition. *European Economic Review* 43(2), 409–429.
- Furman, J. L. and S. Stern (2011). Climbing atop the shoulders of giants: The impact of institutions on cumulative research. *American Economic Review* 101(5), 1933–1963.
- Gaessler, F., D. Harhoff, and S. Sorg (2019). Bargaining Failure and Freedom to Operate: Re-Evaluating the Effect of Patents on Cumulative Innovation. *CEPR Discussion Paper Series* (DP13969).
- Galasso, A. and M. Schankerman (2015). Patents and Cumulative Innovation: Causal Evidence from the Courts. *The Quarterly Journal of Economics* 130(1), 317–369.
- Griffith, R., S. Lee, and B. Straathof (2017). Recombinant innovation and the boundaries of the firm. *International Journal of Industrial Organization* 50, 34–56.
- Gruber, M., D. Harhoff, and K. Hoisl (2013). Boundaries : Scientists vs . Engineers. *Management Science* 59(4), 837–851.
- Harhoff, D., F. Narin, F. M. Scherer, and K. Vopel (1999). Citation frequency and the value of patented inventions. *Review of Economics and Statistics* 81(3), 511–515.
- Jaffe, A. B. (1986). American Economic Association Technological Opportunity and Spillovers of R & D : Evidence from Firms ' Patents , Profits , and Market Value Author (s): Adam B . Jaffe Source : The American Economic Review , Vol . 76 , No . 5 (Dec . , 1986) , pp . 984-1. *American Economic Review* 76(5), 984–1001.
- Javorcik, B. S. (2004). Does foreign direct investment increase the productivity of domestic firms? in search of spillovers through backward linkages. *American Economic Review* 94(3), 605–627.
- Jones, B. F. (2009). The burden of knowledge and the "death of the renaissance man": Is innovation getting harder? *Review of Economic Studies* 76(1), 283–317.
- Jones, B. F., S. Wuchty, and B. Uzzi (2008). Multi-university research teams: Shifting impact, geography, and stratification in science. *Science* 322(5905), 1259–1262.

- Kalemli-Ozcan, S., B. Sørensen, C. Villegas-Sanchez, V. Volosovych, and S. Yesiltas (2015). How to construct nationally Representative Firm Level Data from the ORBIS Global Database. *National Bureau of Economic Research* (WP 21558).
- Krugman, P. R. (1979). Increasing returns, monopolistic competition, and international trade. *Journal of International Economics* 9(4), 469–479.
- Levinthal, D. A. and W. M. . Cohen (1989). Innovation and Learning: The Two Faces of R&D. *The Economic Journal* 99(397), 569–596.
- Liu, E. (2019). Industrial Policies in Production Networks. *The Quarterly Journal of Economics* 134(4), 1883–1948.
- Lychagin, S., J. Pinkse, M. E. Slade, and J. V. Reenen (2016). Spillovers in Space: Does Geography Matter? *Journal of Industrial Economics* 64(2), 295–335.
- Pavcnik, N. (2002). Trade liberalization, exit, and productivity improvements: Evidence from Chilean plants. *Review of Economic Studies* 69(1), 245–276.
- Sampat, B. and H. L. Williams (2019). How do patents affect follow-on innovation? Evidence from the human genome. *American Economic Review* 109(1), 203–236.
- Schmoch, U. (2008). Concept of a Technology Classification for Country Comparisons. *Final Report to the World Intellectual Property Organisation (WIPO)*.
- Shu, P. and C. Steinwender (2019). The impact of trade liberalization on firm productivity and innovation. *Innovation Policy and the Economy* 19(1), 39–68.
- Uzzi, B., S. Mukherjee, M. Stringer, and B. Jones (2013). Atypical combinations and scientific impact. *Science* 342(6157), 468–472.
- Weitzman, M. L. (1998). Recombinant Growth. *The Quarterly Journal of Economics* 113(2), 331–360.
- Zacchia, P. (2019). Knowledge Spillovers through Networks of Scientists. *The Review of Economic Studies* (May), 1–30.

A Appendix: Comparing Direct Effects to Bloom, Draca, van Reenen (2016)

How can we reconcile the negative coefficient of our first stage IV estimation with the positive coefficient found in BDVR? There are three key differences that can explain this easily. First, BDVR use a panel of textile firms while in this paper the main dataset is a panel of non-textile firms. In particular, as a reminder, the regressor in each observation is an average of patenting changes of all 2,380 textile firms, weighted by the technological distance to the non-textile firm. In turn, this leads to the second and third key differences: BDVR run an unweighted regression in which all textile firms are given the same weight. Contrary to their approach, the relative importance of our textile firms depends on their relative importance in the technology network and on the firms recent patent stock; textile firms with a higher degree of network centrality (i.e. understood as closer technological ties to non-textile firms) and also firms with a higher patenting propensity are more likely to drive the results.

Table 11 illustrates these differences and how to reconcile the two approaches. Column (1) transforms the data to a panel of textile firms, as in BDVR. We replicate their results as we also obtain a positive estimated coefficient in this unweighted estimation. In Column (2), we weight each textile firm by its average technological proximity to the pool of non-textile firms. The coefficient turns negative (albeit not yet statistically significant), implying that the positive coefficient on Column (1) tends to be driven by textile firms with a low technological proximity to non-textile firms. In Column (3) we additionally weight by each firms patent stock and now the estimated coefficient is both negative and statistically significant. In summary, while Columns (1)-(3) are based on exactly the same panel of textile firms, accounting for technological distance and firm size explains why we obtain a negative coefficient in our first stage. In Columns (4)-(5) we split the sample of textile firms by their average technological distance to non-textile firms. For firms with distant technological ties to non-textile firms we again replicate the positive and significant coefficient obtained in BDVR. For firms with close technological ties to non-textile firms we rather find a negative and significant coefficient, consistent with our first stage IV estimation results.

Table 11: Direct Effects on Textile Firms

DepVar: Reduction in PAT of TXT firms	Base	Weighted by		Sample split	
	(1)	(2) tech-connect.	(3) tech-connect. & pat.stock	(4) tech-connect. <p50	(5) tech-connect. >p50
Toughness of quotas in 2000	0.0401*** (0.007)	-0.0166 (0.020)	-0.1344** (0.049)	0.0977*** (0.008)	-0.0248** (0.010)
Country FE	Yes	Yes	Yes	Yes	Yes
Industry(SIC-4D) FE	Yes	Yes	Yes	Yes	Yes
No. of clusters	13	13	13	13	13
Observations	11,900	11,900	11,900	5,905	5,995
Unique Firms	2380	2380	2380	1181	1199

Notes: ***/**/* indicate significance at the 1%/5%/10% level. Standard errors are clustered at the SIC4 industry level.

B Appendix: Additional Tables

Table 12: Baseline Results
Conditional on Financial Data Availability for Sample Firms

DepVar: dlnPAT of NTXT firms				
Method	(1) IV 1st stage	(2) IV 2nd stage	(3) Reduced Form	(4) OLS
Toughness of quotas in 2000	-0.5334*** (0.005)		-0.0183*** (0.005)	
dlnPAT of TXT firms		0.0343*** (0.008)		0.0065* (0.004)
Country-Year FE	Yes	Yes	Yes	Yes
Industry(SIC-4D)-Year FE	Yes	Yes	Yes	Yes
Underidentification test		41.1		
Weak identification test		13,962.0		
No. of clusters	394	394	394	394
Observations	121,710	121,710	121,710	121,710
Unique Firms	24342	24342	24342	24342

Notes: ***/**/* indicate significance at the 1%/5%/10% level. Standard errors are clustered at the SIC4 industry level.

Table 13: Heterogeneity by Firm Size

DepVar:	(1) all (baseline)	(2) all w/ revenue info	(3) 1st quartile	(4) 2nd quartile	(5) 3rd quartile	(6) 4th quartile
Reduction in PAT of NTXT firms	0.0329*** (0.006)	0.0334*** (0.007)	0.0188** (0.007)	0.0194** (0.010)	0.0321** (0.013)	0.0636** (0.026)
Country-year FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry(SIC-4D)-year FE	Yes	Yes	Yes	Yes	Yes	Yes
Underidentification test	51.9	50.0	38.3	38.1	45.5	56.0
Weak identification test	15,609.7	14,509.7	12,323.9	9,007.0	9,237.6	11,591.7
No. of clusters	471	391	245	228	247	300
Observations	225,060	146,425	36,340	36,225	36,375	36,325
Unique Firms	45012	29285	7268	7245	7275	7265

Notes: ***/**/* indicate significance at the 1%/5%/10% level. Standard errors are clustered at the SIC4 industry level. This table shows heterogeneous effects by non-textile firm size in an IV framework, where the second stage is presented. Firm size is determined by revenue, where firms in the 1st quartile are smallest. The dependent variable and regressor of interest are as in Table 2. Column (1) repeats the baseline result of Table 4 Column (6). Column (2) conditions on non-missing revenue info in ORBIS. The remaining columns (3)-(6) split the non-textile firms by size into quartiles. Column (3) restricts the sample to the smallest non-textile firms and Column (6) does the same for the largest non-textile firms. All specifications include country-year fixed effects as well as a full set of year dummies interacted with a full set of industry 4-digit dummies.

Table 14: Accounting for Patent Quality

DepVar: Reduction in PAT of NTXT firms	(1) All patents	(2) EPO patents	(3) EPO cit.-weight
Reduction in PAT of TXT firms	0.0329*** (0.006)	0.0191*** (0.006)	0.0135** (0.005)
Country-year FE	Yes	Yes	Yes
Industry(SIC-4D)-year FE	Yes	Yes	Yes
Underidentification test	51.9	51.9	51.9
Weak identification test	15,609.7	15,609.7	15,609.7
No. of clusters	471	471	471
Observations	225,060	225,060	225,060
Unique Firms	45012	45012	45012

Notes: ***/**/* indicate significance at the 1%/5%/10% level. Standard errors are clustered at the SIC4 industry level. In this table we address patent quality by non-textile firms in an IV framework where the second stage is presented. Column (1) is our baseline result (as in Table 4 Column (6)), which includes patents filed with any patent authority. Column (2) restricts the patent count to those filed with the European Patent Office (EPO); as citations across patent authorities cannot be directly compared, we believe it is sensible to focus only on the most important patent authority for our European manufacturing firms. Finally, Column (3) weights the change in the log of patents by the number of citations received in the first five years post-grant. All specifications include country-year fixed effects as well as a full set of year dummies interacted with a full set of industry 4-digit dummies.

Table 15: Alternative Lag Specification

DepVar: Reduction in PAT of NTXT firms	(1)	(2)	(3)
Reduction in PAT of TXT firms	0.033 *** (0.006)		
L.Reduction in PAT of TXT firms		0.038 *** (0.008)	
L2.Reduction in PAT of TXT firms			0.059 *** (0.014)
Country-year FE	Yes	Yes	Yes
Industry(SIC-4D)-year FE	Yes	Yes	Yes
Underidentification test	51.9	52.4	52.4
Weak identification test	15,609.7	4,829.5	2,373.0
No. of clusters	471	471	471
Observations	225,060	180,048	135,036
Unique Firms	45012	45012	45012

Notes: ***/**/* indicate significance at the 1%/5%/10% level. Standard errors are clustered at the SIC4 industry level.