Exporting Ideas: Knowledge Flows from Expanding Trade in Goods*

Philippe Aghion Antonin Bergeaud Timothee Gigout

Matthieu Lequien Marc Melitz

December 2019

PRELIMINARY - PLEASE DO NOT CIRCULATE

Abstract

When a firm starts exporting into a new destination, its products and technologies suddenly become more visible there. Firms and inventors in that destination can then innovate building on these technologies. We combine French firm-level administrative, customs and patent data over 1995-2012 to show that entry into a new export market increases the patents' citations received from that destination. This technological spillover is concentrated in countries at intermediate levels of development and among the most productive firms.

JEL classification: O33, O34, O40, F10, F14

Keywords: International Trade, Spillover, Innovation, Patent, Competition

*Addresses - Aghion: College de France and LSE; Bergeaud: Banque de France and CEP; Gigout: College de France, CEPII and Lille University; Lequien: Banque de France and PSE; Melitz: Harvard University and NBER.

1 Introduction

Modern growth theory predicts that international trade should enhance productivity growth for several reasons. First, trade allows potential innovators to sell to a larger market; and by *increasing market size*, trade increases the size of ex post rents that accrue to successful innovators, thereby encouraging R&D investments. Second, trade *raises competition* in product markets, which in turn encourages innovation aimed at escaping competition by more advanced firms while discouraging innovation by laggard firms in the domestic economy. Third, trade *induces knowledge spillovers* which allows producers in recipient countries to catch up with the technological frontier. In previous work (see Aghion et al., 2018) we used French firm-level accounting, trade, and patent information to provide evidence on the market size and competition effects of trade expansion. In this paper, we use the same datasets to provide evidence of a knowledge spillover effect for trade expansion.

The following stylized fact motivates our analysis in this paper. In Figure 1a, we plot the long difference between the number of French exporters from 1995 to 2012 (i.e the difference between the number of French exporters in 2012 and the number in 1995) for the various geographical regions of the world. Each color corresponds to a decile in the long difference distribution across regions. Dark red corresponds to regions with the largest increase in the number of exporters from 1995 to 2012, whereas dark blue corresponds to the regions with the smallest increase in the number of exporters from 1995 to 2012. In Figure 1b, we plot the long difference between the number of citations to French patents from 1995 to 2012 for different regions worldwide; again the dark red (resp. dark red) color refers to regions lying in the highest (resp. lowest) decile in terms of long difference increases in citations. We see that those destinations experiencing the largest increase in the number of French exporters also experience the largest increase in patent citations to French innovations over the same time period. The covariance between the two long differences is equal to 1.62 (s.e.=0.22).[1]

Figure 1: Evolution of Trade and Innovation Linkages

(a) Number of French Exporters (b) Citations

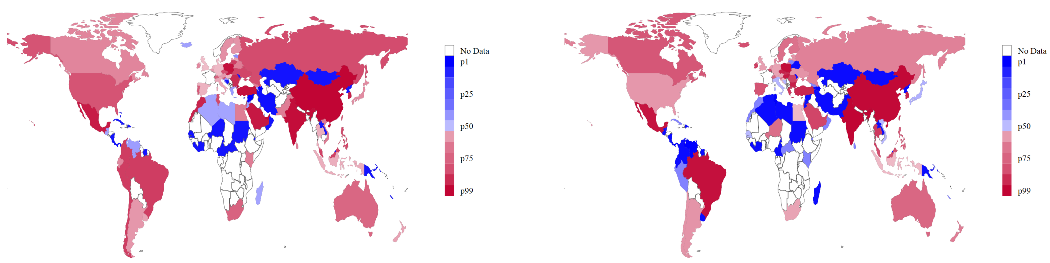

Notes: Evolution in the number of French exporters in each country (left-hand side panel) and the number of citations received from each country (right-hand side panel) between 1995 and 2012. Colors correspond to different deciles in the corresponding quantity.

[1] That is: $\texttt{LD Ln Citations}_j = 1.62(s.e. = 0.22) \times \texttt{LD Ln Exporters}_j + \upsilon_j$

We begin with a comprehensive set of patents belonging to French exporters over the 1995-2012 period. For every year and potential export destination, we construct a citation count for each exporters' patents. These citations come from new patents introduced in that year by firms or inventors operating in the destination country. We then investigate how a French firm's citation count in a destination changes whenever that firm starts exporting to that destination. Increases in a new exporter's citations represent new patents recorded in that destination subsequent to the exporter's entry into the destination. Those patents citing the French exporter represent a measure of its technological influence in that destination. We use the timing of the exporter's entry into a market and its citations in that market to infer a causal relationship between the two.

We show that exporting to a new foreign market increases the flow of citations received by the exporter from that market. The underlying idea is that entry into that new market raises the visibility of the exporter's technology to domestic firms in the market. Those domestic firms can then more easily generate further innovations that build on that technology, conditional on the host country's degree of absorptive capacity (Cohen and Levinthal, 1989).

More specifically, we use a difference-in-difference strategy to analyze the response of patent citations to a French firm's export market entry in a particular year. To deal with potential endogeneity of our results (in particular for the fact that exporting firms have better technologies or technologies that are better suited to the destination country), we leverage the high dimensionality of our dataset to condition on a large number of fixed effects. We also adapt an identification strategy from Watzinger et al. (2017, 2018), who study the knowledge spillovers induced by professor transfers across universities. This allows us to build a control group of French firms with an *ex-ante* similar probability of entry into a given market, but who did not enter in that particular year.

Our first main finding is that this impact of entry on citations (and hence knowledge flows in the destination) is positive and significant starting 3 years after export market entry, and peaking after 5 years.[2] Quantitatively, we find that export market entry induces an 38.7% increase in the exporter's mean citation rate. We also find that export market entry leads to a 3.3 pp increase in the probability of receiving citations for exporters with no citations.

We investigate whether relatively more productive firms generate more/better knowledge flows, but do not find major differences across the distribution of exporting firms. On the other hand, we find that the characteristics of the destinations do matter: the level of development (as measured by GDP per capita) strongly influences those spillovers. We find that the spillover intensity is hump-shaped with a peak around 55-60 percentile of the GDP per capita distribution across destinations. The spillover intensity steadily decreases with development for richer countries beyond that peak – but remains positive. We also find a negative and significant spillover for the poorest set of destinations. This is consistent with the view that firms in those destinations have much lower "absorptive capacity" to use the

[2]This timing lag is consistent with the time needed post-entry for new research to generate new priority patents.

knowledge spillover from the new French exporters, and mainly suffer from the increased competition effect generated by those French firms. Development then enhances a destination's ability to absorb - and build upon - the technology of the French exporters. At the other end, highly developed destinations may have already discovered the technologies that would allow them to make use of the French firm's technology. Overall, our results vindicate Cohen and Levinthal (1989)'s view stated in the following quote: "*Economists conventionally think of R&D as generating one product: new information. We suggest that R&D not only generates new information, but also enhances the firm's ability to assimilate and exploit existing information. [...] we show that, contrary to the traditional result, intra-industry spillovers may encourage equilibrium industry R&D investment*." (Cohen and Levinthal, 1989, p.569).

Our analysis relates to several other strands of literature. There is first the literature on spillovers and trade, starting with Coe and Helpman (1995a), who show that a country's TFP is positively correlated not only with domestic R&D but also with foreign R&D and to an extent which increases with the country's degree of openness to foreign trade.[3] We contribute to this literature by using firm-level data and patent citation data to identify a causal effect of export on the innovative activity in the destination country.

Second, our paper relates to the recent literature on trade and innovation, including papers on both, imports and innovation (see Bloom et al., 2016; Autor et al., 2016; Bombardini et al., 2017) and on exports and innovation (see Lileeva and Trefler, 2010; Aghion et al., 2018). Overall, this literature concentrates on the competition and market size effects of trade. We contribute to that literature by looking at the technological spillover effects of trade, and more precisely at how exporting to a destination country affects the exporting firm's patent citations by firms in that destination country.

Third is the literature on academia, scientists and citations. Thus Azoulay et al. (2010) and more recently Bell et al. (2018) analyze the impact of an inventor's death on the subsequent innovation and income patterns of the inventor's surviving coauthors. Waldinger (2011) analyzes the impact of the dismissal of Jewish scientists's by the Nazi government in Germany in the '30s. And Watzinger et al. (2017, 2018) analyze the impact of the mobility of scientists across German universities on local citations to their work. We contribute to this and the broader literature on knowledge spillovers and absorptive capacity by looking at how trade interacts with knowledge spillovers and absorptive capacity.[4]

The remaining part of the paper is organized as follows. Section 2 briefly presents the data and details our empirical strategy and section 3 shows our baseline results. We conduct further robustness tests in section 4. Section 5 concludes.

[3]See also Keller and Yeaple, 2009, Coe et al., 2009, and Keller and Yeaple, 2009.

[4]See Aghion and Jaravel (2015) for more detailed references to that literature.

2 Data and Methodology

2.1 Data

We build a database that covers all French firms and linking export, production and innovation/citation data from 1994 to 2012. Our database builds on three separate sources. First, detailed customs data provide French exports by product and country of destination for each French firm over 1993-2012. Every firm must report its exports by destination country and by very detailed product (with a classification of 10,000 different products consistent with 8-digit HS codes). From this database, we extract the date of first entry into a foreign market for each firm. Our second data source is the INSEE-DGFiP administrative fiscal dataset (FICUS-FARE), which provides extensive production and financial information for all firms operating in France. This data is drawn from compulsory reporting to fiscal authorities in France, supplemented by further census data collected by INSEE.

Our third data source is the Spring 2016 PATSTAT dataset from the European Patent Office. This contains detailed information on all patent applications from most of the patent offices around the world. We use information on the network of patent linkages via citations. Although each French firm has a unique identifying number (Siren) across all French databases, patent offices identify firms using only their name. The recording of the name is sometimes inconsistent from one patent to another, and may also contain typos. Various algorithms have been developed to harmonize assignees' names (for example this is the case of the OECD's Harmonized Assignee Name database) but none of those have been applied specifically to French firms. One notable exception is the rigorous matching algorithm developed by Lequien et al. (2019) to link each patent application with the Siren numbers of the corresponding French firms; for all firms with more than ten employees. Based on supervised learning, this new method provides significant performance improvements relative to previous methods used in the empirical patent literature: its recall rate (i.e. the share of all the true matches that are accurate) is 86.1% and its precision rate (i.e. the share of the identified matches that are accurate) is 97.0%. This is the matching procedure we use for our empirical analysis in this paper.

We seek to measure the knowledge spillovers from new French exporters-innovators (who enter an export destination during our 1995-2012 sample years) to firms/inventors in that destination. There are 7,976 such French exporters-innovators (with at least one patent that can potentially be cited) who enter one of 158 destinations during our sample years.[5] Those firms own 147,183 patents that have generated 511,326 citations (new patents) in Foreign destinations. Those citations represent 40,957 potential knowledge spillover links between a Foreign patent and a French firm. Of those potential links, 13,462 have been treated in the sense that the cited French firm has entered the corresponding export destination during our sample years. All those numbers are reported in Table 1.

[5] We do not use information on firm-destination links when we observe that firm exporting to that destination in the first two years of our customs data prior to our sample years (1993-94). We cannot infer when, before 1993, that firm started exporting to that destination.

Table 1: Descriptive Statistics

Level	N
Years	18 (1995-2012)
Destinations	158
Firms	7,976
Patent	147,183
Citations	511,326
Links (firms * destination)	40,957
$\hookrightarrow$ Ever treated	13,462

2.2 Empirical methodology

We want to estimate how a French firm's entry into a new export market affects its flow of new citations from this destination to the firm's patents.

The most natural approach is to aggregate our data at the *firm-destination* level. More specifically, we look at how a firm's entry into a new export market j affects subsequent citations received from firms located in this destination j. We thus estimate the following distributed lead/lag model:

$$Y_{f,j,t} = \sum_{\substack{k=k_{min} \\ k \neq -1}}^{k_{max}} \beta_k \times ENTRY_{f,j,t-k} + \chi_{f,j} + \chi_{f,t} + \chi_{j,s,t} + \varepsilon_{f,j,t}. \tag{1}$$

In this equation, j is a country and f a French firm operating in sector s. Our dependent variable, $Y_{f,j,t}$, measures the number of citations made at t by firms in country j to the stock of existing patents filed by firm f. $ENTRY_{f,j,t}$ is a dummy variable equal to 1 if firm f declares some positive exports to country j during year t for the first time in our sample. $\chi_{f,j}$, $\chi_{f,t}$ and $\chi_{j,s,t}$ respectively denote firm-destination, firm-year and destination-sector-year fixed effects, and are meant to capture global innovation shock in a given market ($\chi_{j,s,t}$), firm innovation intensity ($\chi_{f,t}$) and country-firm specificities. We include a relatively large set of lags and leads in order to capture the full evolution of citations made to f following its entry. Our identification relies on the fact that entry decision is not affected by future citations. We check that this is not the case by looking at the coefficient on the leads (pre-trends) and we return back to this question in the next subsection.

This specification does not allow to distinguish between the different patents owned by the firm. In particular, it does not allow to separate citations made to new patents from those made to older patents (namely to those patents that existed prior to entry). To investigate this more in detail, we disaggregate our data further and consider a panel at the *patent-destination-level*. We therefore also consider the following model:

$$Y_{p,j,t} = \sum_{\substack{k=k_{min} \\ k\neq -1}}^{k_{max}} \beta_k \times ENTRY_{f,j,t-k} + \delta \times X_{p,t_e-1} + \chi_{t_e} + \chi_{t_p,z} + \chi_{t,z} + \varepsilon_{p,j,t}, \quad (2)$$

where $Y_{p,j,t}$ is the number of priority patents by applicants in destination country j citing patent p of technology class z and filed at date t by a firm f (or said differently, the flow of destination-specific citations received by p from j during year t). As previously, $ENTRY_{f,j,t}$ is a dummy equal to one if French firm f enters destination j for the first time at date t. This level of aggregation allows us to also control for the number of citations received by patents filed at t from f prior to entry (at t=t_e) $X_{p,t_e-1} = \sum_{t=-\infty}^{t_e-1} Y_{p,f,World,t}$. We also control for the global lifecycle of an innovation within each technological field when the French and foreign (j) patent were filed by introducing the dummies $\chi_{t_p,z}$ and $\chi_{t,z}$ and add a dummy for the entry date χ_{t_e}. Lastly, we cluster the standard errors at the link-level i.e. by firm-country (f, j) pair.

2.3 Endogeneity of entry

One immediate concern is that the correlation between entry and the subsequent increase in citations may partly reflect the fact that better performing firms (with patents that are more likely to be cited) have a higher probability of entering new export markets. To deal with this selection problem, we take inspiration from Watzinger et al. (2017, 2018), who study the knowledge spillovers induced by professor transfers across universities. Those authors use administrative data from German universities. Every year a university in Germany creates a list of professors eligible for transfers. The probability of transfer within that list is as good as random. The authors then measure the effect of mobility within a list of eligible professors on the Patent-to-Article and Article-to-Article citation counts.

Similarly, we construct a control group of French firms for every French exporter observed to enter a new foreign market in a given year. Firms in this control group have a similar probability of entering that destination in that given year. All of our subsequent regressions on patent/citations flows are then reported *within* this control group. We thus start by estimating the probability that each French firm enters an export destination for the first time in each year. We then partition all those firms (by destination-year) into bins according to their predicted entry percentile. Within each bin, there are firms that enter the foreign market early, or late, or never; and there are firms that exit that foreign market early, late or never. This first-stage analysis allows us to control for the selection endogeneity by always comparing an entrant (exporter to a new destination) within its control group in our second stage.

In that second stage, we measure the impact of export entry on the knowledge flows between the entering firm and new priority patents in the destination (citing the exporter's prior patents) as in equation (2) but we now control for export entry selection by conditioning our estimates on iso-probability bin fixed effect we previously described. Thus, our results are estimated within a bin of firm-destination-year triplets with very similar probabilities of export market entry. We provide further details on this empirical methodology as well as first stage results in Appendix A.

3 Results

3.1 Baseline firm-level results

We estimate parameters β_k from equation (1) and report their values $\hat{\beta}_k$ and associated 95% confidence intervals in Figures 2. These figures vary by the set of fixed effects that are included in the estimations. More precisely, Figure 2a conditions on firm-year and destination-sector-year effects while Figure 2b conditions on destination-firm and destination-sector-year effects. Note that the former is only identified for firms that export in more than one destination. Figure 2c estimates the full model from equation (1), that is including $\chi_{f,j} + \chi_{f,t} + \chi_{j,k,t}$. In all cases, we observe a similar pattern: the marginal impact of entry on the number of citations received by the firm spikes after 3 to 5 years. We also report the same model as in Figure 2c but using a semi-dynamic approach, i.e. we impose $\beta_k = 0$ for all $k < 0$ (i.e. we set the pre-trend to exactly 0). Results are presented in Figure 2d.

In all these specifications, the dependent variable is the number of citations originated from country j divided by the previous year's stock of patents of firm f. We have also experimented with other functional forms for the dependent variable in section 3.3.

Magnitude:

TBD

3.2 Patent-level results

We now turn to our estimations at the patent-destination unit. This level of disaggregation allows us to look more precisely at how entry into a foreign market affects the standard lifecycle of the citations received by a given patents and to control more precisely for the observable characteristics of this patents (technological class, quality, age etc...).

We first run the analog of Figure 2a and 2b, i.e. conditioning on firm-year and destination-firm fixed effects. We would ideally want to control for a full *firm-destination-year* set of fixed effects in order to capture the endogenous decision of entering a given market. While this is technically feasible, in that case identification comes from comparing the flow of citations from a given in country to two different patents of the same firm, conditional on all observable characteristics and in particular their age. Instead, we prefer to implement our strategy based on the construction of destination-time specific probability of entry as described in Section 2.3.

We first verify that there is no difference between the treated group and the control group prior to entry: the regression points for the leads fluctuate around zero and are not

Figure 2: Baseline firm level results

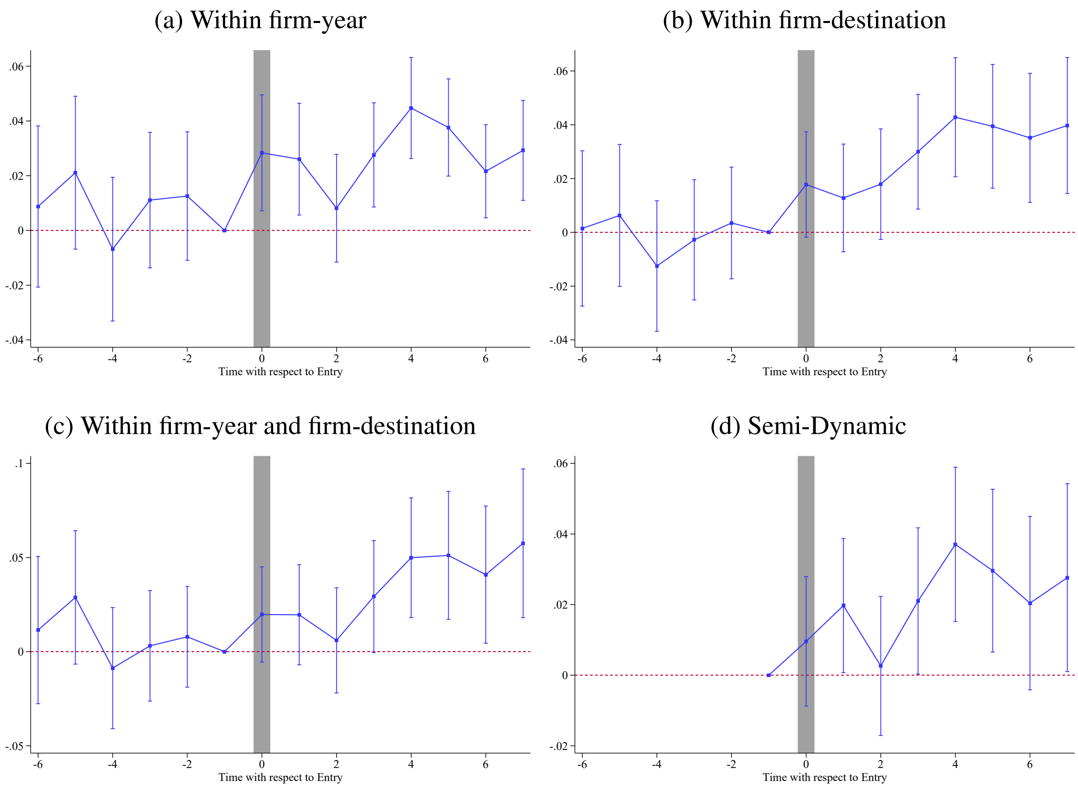

Notes: Panel 2a presents the results from estimating Equation x with firm-year and sector-destination-year fixed effects. Panel 2b presents the results from estimating Equation y with firm-destination and sector-destination-year fixed effects. Panel 2c presents the results from estimating Equation z with firm-destination, sector-destination-year and firm-year fixed effects. Panel 2d is the semi-dynamic version of Panel 2c.

significant. But entering into a market leads to a marked and significant increase in citations after 3 years – lasting for 3 years (3 to 5 years post-entry). This effect progressively dies out thereafter.

Figure 3a graphically depicts all the estimated leads and lags coefficients for entry (the main coefficients of interest $\hat{\beta}_k$), along with their 95% confidence intervals, for our fully dynamic specification with pre-entry periods, corresponding to equation (2) to which we add a probability bin fixed effect. The dependent variable this time is the flow of citations received by a patent p from a destination j.

Figure 3b repeats the same exercise as Figure 3a, but uses a semi-dynamic specification where we omit the pre-trend dummies to gain additional years of observations. This figure shows similar post-entry effects to those in the fully dynamic specification (both in magnitudes and in precision): entry increases received citations 3 to 5 years after entry, and has no significant impact at shorter or longer horizons.

Magnitude: Quantitatively, firms entering into a destination receive an additional 0.011-0.013 citations for their patents from that destination 3 to 5 years after entry, compared

Figure 3: Main Specification: Priority Citations Count

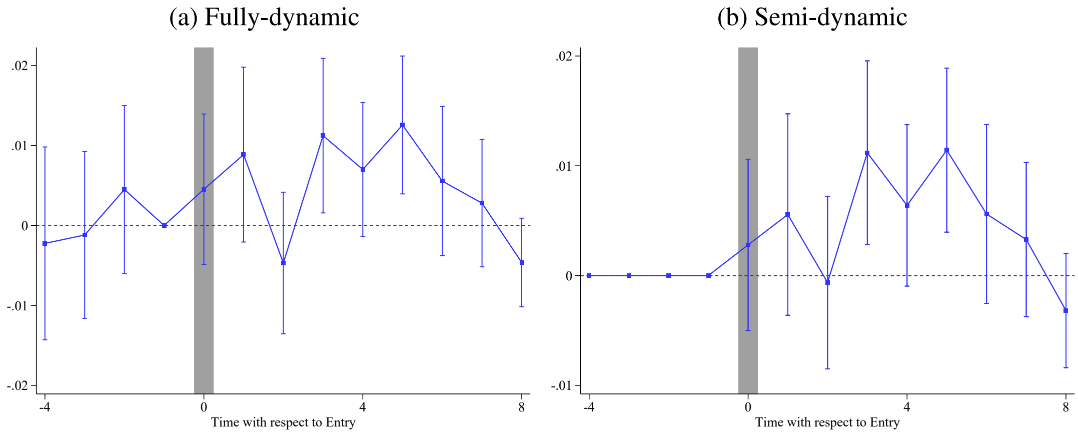

Notes: These figures show the coefficients β_k from the estimation of our baseline Equation 2. Left-hand side subfigure estimate all coefficients while right-hand side subfigure assumes that all lead coefficients are 0. 95% confidence interval are presented. Standard errors are clustered at the link (firm-destination) level.

to similar firms that had not entered that destination at that time. This corresponds to a 16-18% increase from the mean citation rate in our sample.

In order to assess the magnitude of the full treatment effect, we compute the sum of coefficients and find a total coefficient of 0.0424. Over this 8 year time window after entry, a firm receives an average of 0.51 citations whereas a firm that does not export to that destination receives an average of 0.46 citations. This corresponds to a 13.3% increase in citations from the export destination country.

3.3 Different functional forms

In the following two figures, we explore the impact of changing the functional form for the dependent variable either at the firm level or at the patent level. In all cases, we will stick with our semi-dynamic specification in order maximize the number of observations.

At the firm level: in Figure 2, we have used the number of citations per patent as a dependent variable. This is one of many functional forms that we can consider in order to capture the actual data generating process. In Figure 4a and 4b, we consider two other concave functions, respectively $ln(1 + C)$ and $\mathcal{H}[C]$[6] where C is the number of citations received by firm f from country j at time t to any patents in its portfolio. By using these functional forms, we want to allow for both the intensive and the extensive margin to be estimated (that is, we want the dependent variable to be defined if $C = 0$, we come back to this in next section).

[6] $\mathcal{H}$ is the inverse hyperbolic sin function: $\mathcal{H}(C) = \frac{1}{2}\log\left(C + \sqrt{1 + C^2}\right)$. This function gives more weight to the extensive margin than the $log(1 + C)$ and is otherwise very similar.

Figure 4: Firm level regression - other LHS variable

(a) $\log(1 + C)$ (b) $\mathcal{H}(C)$

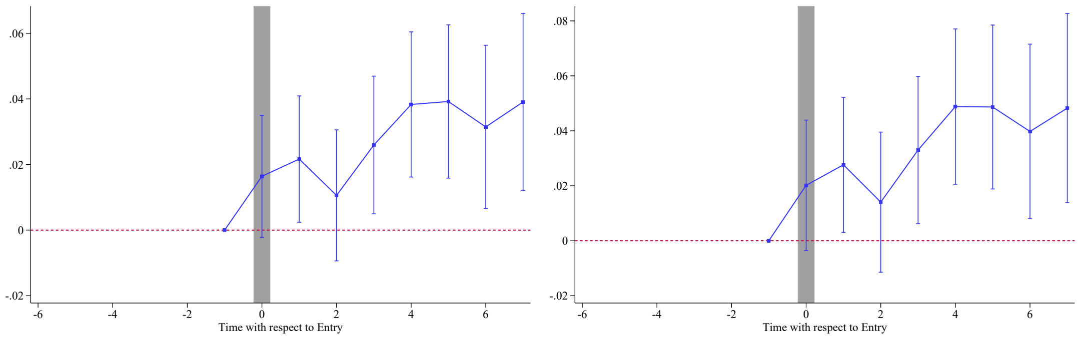

Notes: These figures show the coefficients β_k from the estimation of our baseline Equation 1. The dependent variable is defined as the log of 1 + the number of new citations made to the stock of patents in the left-hand side subfigure and the inverse hyperbolic sine function of these new citations in the right-hand side subfigure. Standard errors are clustered at the link (firm-destination) level.

At the patent level: We do the same exercises at the patent level. In Figure 5a the dependent variable is $\ln(1 + C)$, whereas in Figure 5b the dependent variable is a hyperbolic function $\mathcal{H}[C]$. These figures confirm that the pattern from our baseline Figure 3b is not particularly sensitive to changes in the functional form of the dependent variable.

Figure 5: Main Specification: alternative LHS variables

(a) $\log(1 + C)$ (b) $\mathcal{H}(C)$

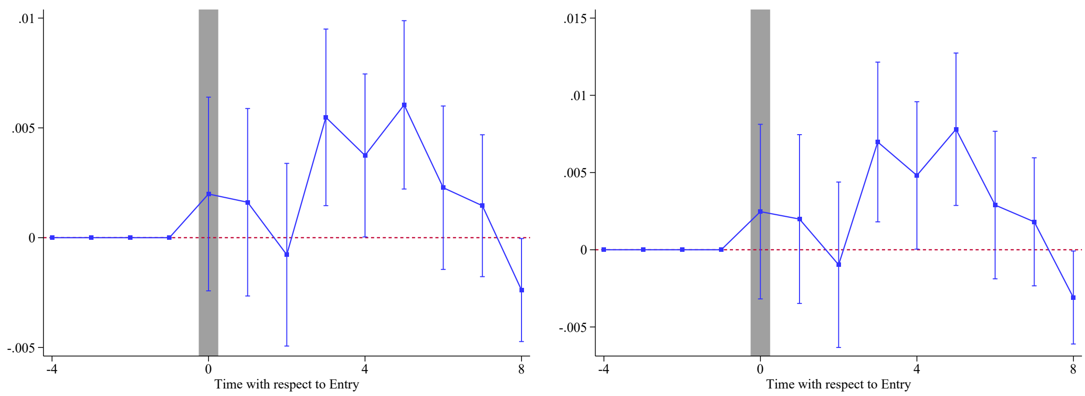

Notes: This figure shows the estimated $\hat{\beta}$ coefficients from the estimation of our baseline Equation 2. The dependant variable is $\log(1 + C)$ in the left panel, and $\mathcal{H}(C)$ in the right panel. 95% confidence interval are presented. Standard errors are clustered at the link level.

3.4 Extensive vs Intensive margin decomposition

We now decompose the overall response of citations into an extensive margin component – a binary transition from no citations to positive citations – and an intensive margin component – an increase in citations conditional on a positive number of citations. Figure 6 shows the result from the binary response regression. As we previously discussed, the results can be interpreted as a linear probability model yielding the probability that an entrant is cited in the export destination. We see that this dynamic pattern is very similar to our baseline regression, with a significant increase in the citation probability 3 to 5 years after entry. The probability of being cited increases with entry 3 to 5 years after entry. Entry increases the probability of a citation by 3.3 percentage point 3 to 5 years after entry. This implies that an entering firm is 66% more likely to obtain a citation relative to a firm that does not enter in that same year.

Figure 6: $Y = \{0, 1\}$

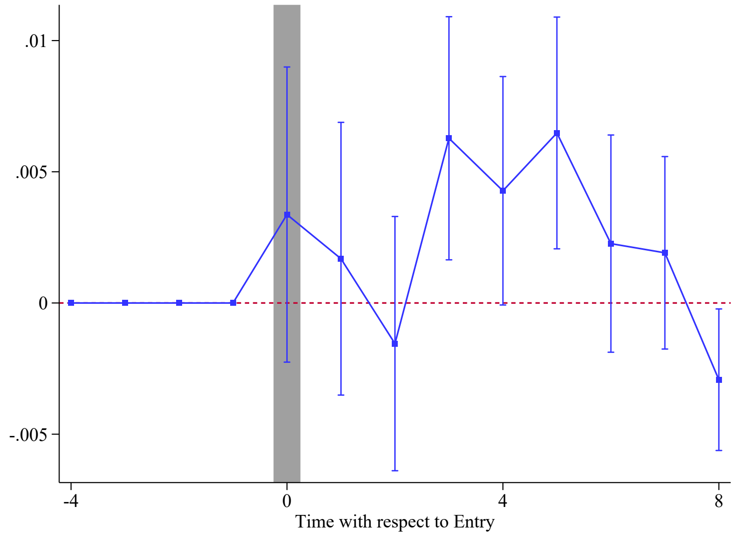

Notes: This figure shows the estimated $\hat{\beta}$ coefficients from the estimation of our baseline Equation 2. The dependant variable is the status of the technological link between the firm's applicants and the foreign country's applicants. 95% confidence interval are presented. Standard errors are clustered at the link level.

Going further, we use a distribution regression technique. We construct dummy variables that describe the cumulative distribution of citations rates. Each dummy equals one for destination that experienced a citation rate greater than x. We estimate effects on a series of dummies that move from 0 to 25 in 0.05 increments ("distribution regression" Chernozhukov et al. (2013)). See Goodman-Bacon and Cunningham (2019) for a more recent application.

The effect measured in our baseline specification is largely explained by an increase in the probability of getting a few more citations per patent. The effect on the probability of getting more than 3 citations per patent is a precisely measured at zero.

The first coefficient is a regression where the outcome variable is a dummy that takes one if the citation rate is strictly greater than zero. It is therefore the effect on the extensive margin. We also plot the fully dynamic estimation of this in Figure 8. The estimated equation is the same equation as the baseline but the outcome variable is a dummy indicating whether or not the patent receives any citations that year.

Figure 7: Distribution Regression

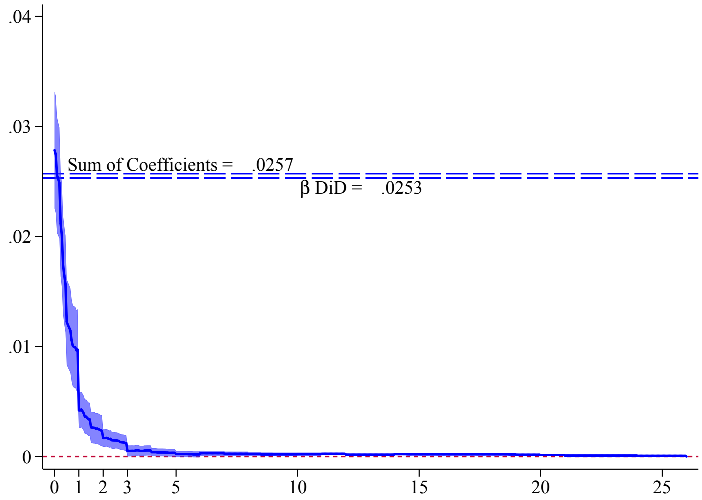

Notes: This figure plots the DiD effect of the initial entry into a foreign destination on the probability of having a citation rate greater than the amount on the x-axis. For details on "distribution regression" see Chernozhukov et al. (2013). This reflects changes in the cumulative distribution of citation rates. 95% confidence interval are presented. Standard errors are clustered at the firm and destination level.

Figure 8: Extensive margin

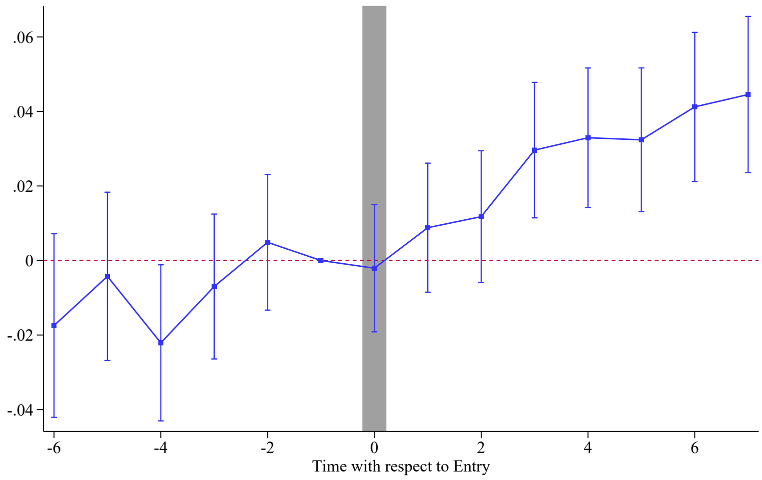

Notes: This figure plots the effect of the initial entry into a foreign destination on the probability of being cited more than once that year. 95% confidence interval are presented. Standard errors are clustered at the link level.

3.5 Heterogeneity

In this subsection, we investigate how the impact of entry on citations varies with both the exporting firm's productivity (an indication of the technology embodied in the patents), and the level of development of the destination country (which we use as a proxy for the country's degree of absorptive capacity). As we already mentioned, we measure these heterogeneous responses by moving to a static version of the treatment variable with a unique entry dummy equal to 0 before entry and to 1 thereafter. Moreover, we use a kernel reweighting scheme across the various percentiles in the variable with heterogeneous effects. The kernel approach allows us to flexibly estimate the functional form of the marginal effect of entry on patent citation across the percentiles in the heterogeneity variable. Each dot in the figure corresponds to the effect on citations estimated at a given percentile of the heterogeneous response variable (with Gaussian weights and a bandwidth of 20 percentiles).

3.5.1 Impact of the exporting firm's productivity

We first investigate the impact of a firm's productivity on the impact of the spillover entry. A more productive firm is expected to generate patents that embody better or more valuable technologies. Those patents are presumably more likely to induce follow-up innovations by other firms, and should be reflected in additional citations whenever those innovations lead to new patents. We adapt the baseline second stage regression to allow for varying β coefficients across percentiles in the distribution of French firm's productivity (at date $t-1$) at the entry stage. Productivity is measured by the firm's value added per employee. In Figure 9a, each dot corresponds to the effect of the initial entry into a foreign market estimated locally at a given percentile of the ex-ante productivity distribution. The blue band corresponds to the 90% confidence interval. We see that the effect of entry on new patent citations steadily increases over most of the productivity range (excluding firms with the lowest and the very highest levels of productivity). However, this increase is very slight.

Figure 9: Heterogeneous effects

(a) Productivity (b) GDP per capita

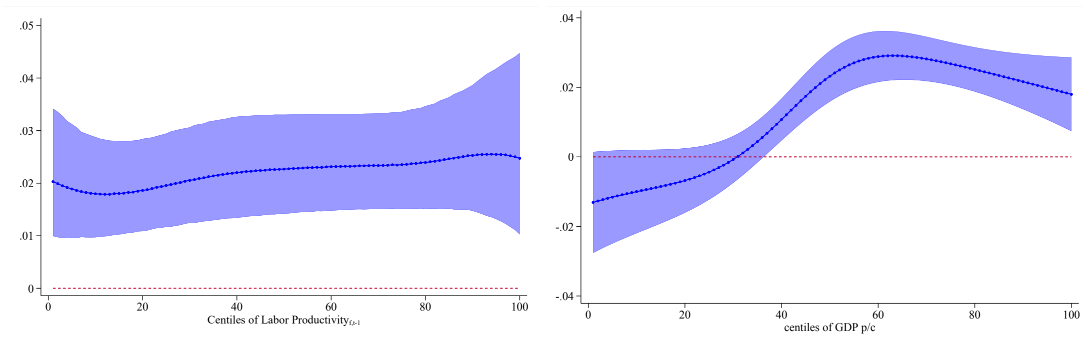

Notes: This figure plots the effect of the initial entry into a foreign destination estimated locally at a given percentile of the ex-ante distribution of firm-level productivity (LHS subfigure) and country-level GDP per capita (RHS subfigure). The dependant variable is the number of citations. We use Gaussian weights with a bandwidth set to 20 centiles. 90% confidence interval are presented. Standard errors are clustered at the link (country-firm) level.

3.5.2 Impact of a destination's development level

The transfer of knowledge from a French exporter to firms in the export destination is likely to depend upon the destination's technological development relative to the French exporter. If firms in the destination country lag far behind the French firm, then presumably these firms are not adequately equipped to build on the French firm's innovation, and therefore the French firm's entry should have limited impact on innovation in the destination country. The French firm might even deter such innovation in the destination country due to the increased competition it induces for potential innovators in that country (see Aghion et al., 2005): as a result, the impact of the French firm's entry on citations by firms in the destination country may even turn negative. On the other hand, if firms in the destination country are neck-and-neck with the French firm, then these firms can easily build upon the French firm's technology to generate new innovations: in that case entry by the French firm should increase citations by the destination country of the firm's innovations. Finally, if firms in the destination country are far ahead of the French firm's technology, then these firms will often not find it useful to develop further the French innovation as they already enjoy a better technology: entry by the French firm would then have little to no impact on its citations by firms in the destination country.[7]

To test for a differential impact of entry on citations varying with a destination's development level, we run a similar version of our static specification described above. But we now allow for our coefficient to vary across the percentiles of the destinations' GDP per capita. At low levels of GDP per capita (below the 30th percentile), entry decreases citations (Figure 9b). At intermediate-high level of GDP per capita (between the 30th and the 60th percentile), entry increases citation. The effect is then decreasing in GDP per capita for the richest countries, though it remains significantly positive.

4 Robustness and extensions

4.1 Placebos

In the baseline specification, we clustered standard errors at the firm-country link level. This provided us with standard errors that are asymptotically robust to serial auto-correlation for the error term as well as to correlations across patents within a link. Here we implement Chetty et al. (2009)'s non-parametric permutation test of $\beta_k = 0$ for $k = \{5\}$

To do so, we randomly reassign the date of entry into an export destination across links and then we re-estimate equation (1). We repeat this process 2000 times in order to obtain an empirical distribution of *placebo* coefficients $\hat{\beta}_k^p$. If entry had no effect on citations, we would expect our baseline estimate to fall somewhere in the middle of the distribution of the coefficients of the placebo coefficients $\hat{\beta}_k^p$. Since that test does not rely on any parametric assumption regarding the structure of the error term, it is immune to the over-rejection of the null hypothesis highlighted by Bertrand et al. (2004).

We plot the histogram of this distribution of placebo coefficients in Figure 10. The figure confirms that our coefficient of interest $\hat{\beta}_{k=5}$ (the solid blue line) lies on the right of the 99^{th} percentile (the red dashed lines) of the distribution of placebo coefficients. It

[7] All these developments should have different consequences for the destination firms' products as well, but the lack of data on those products prevents us from assessing such impacts. They also bring about different consequences for the French exporter's products, which we plan to investigate in future work.

confirms that initial entry into a destination leads to an increase in citations.

Figure 10: Random Permutations

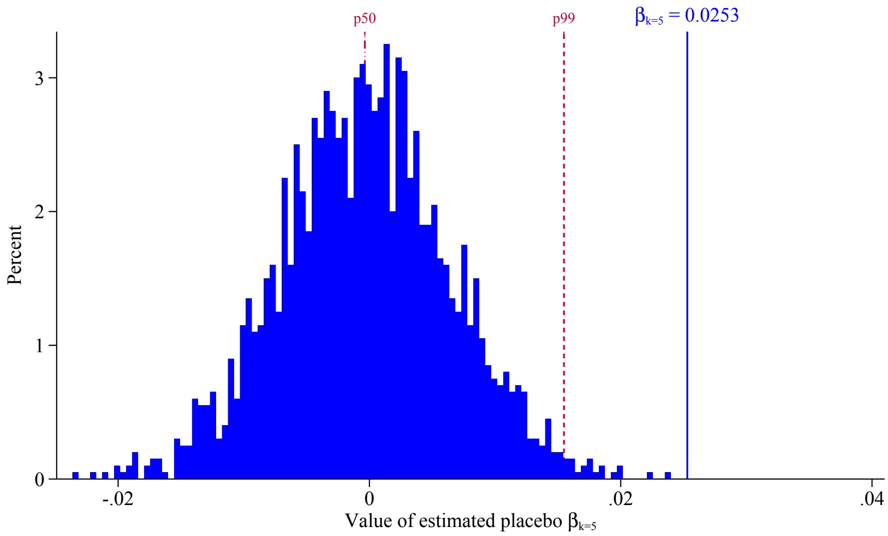

Notes: Those figures present the distribution of 2000 estimates of the coefficient $\hat{\beta}_k^p$ after performing a random permutation. The dark blue lines corresponds to the coefficient from our baseline regression. 2.

4.2 Alternative measurement of spillovers

4.2.1 Second-order citations

TBD

4.2.2 Semantic-based distance

TBD

4.3 Spillovers to the rest of the group

We change the outcome variable. It is now the citation rate of the other firms that belonged to the group at the time of the patent application. We leave the rest unchanged compared to the baseline. We find a small positive effect. include Figure 11

Figure 11: Main Specification: Triadic citations to the rest of the group

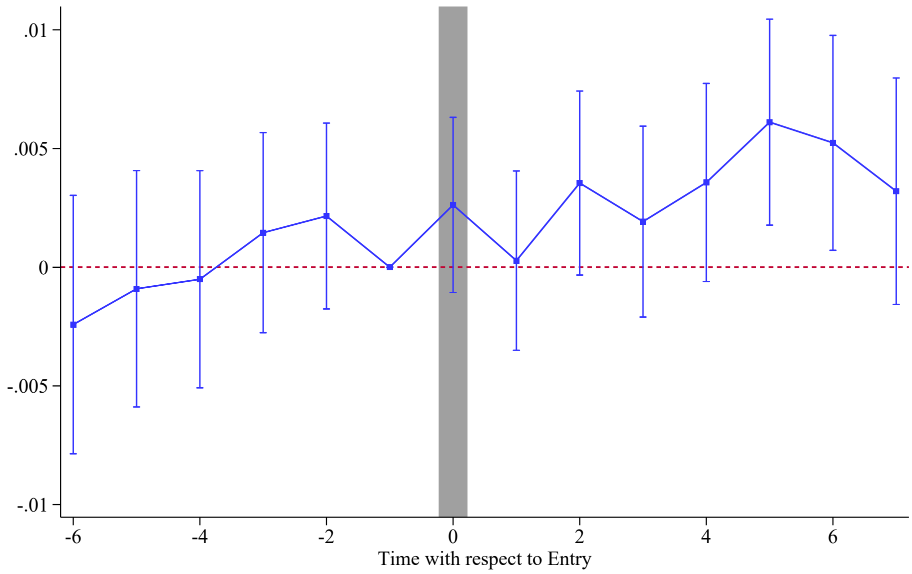

Notes: This figure shows the coefficients from the estimation of our baseline Equation 2. 95% confidence interval are presented. Standard errors are clustered at the link level.

5 Discussion and conclusion

In this paper we use French firm-level fiscal, custom, and patent citation data over the period 1995-2012 to estimate the impact of export market entry on the citations of the exporter's prior patents in the destination country. We find a positive and significant effect of entry on those citations. Overall, our results validate the notion that trade induces technological spillovers (in line with Coe and Helpman, 1995b). And the results are also consistent with Cohen and Levinthal (1989)'s view that spillovers occur conditionally upon the recipient country exhibiting sufficient *absorptive capacity*.

Our findings have several implications. First, our main findings that trade induces knowledge spillovers is in line with the notion that trade is a source of cross-country convergence. Second, fostering development in the destination country increases the country's ability to build upon the innovations brought by foreign exporters. Third, more productive firms – in addition to being more likely to export – are also more likely to induce technological spillovers (though this increase is slight).

Our analysis can be extended in several interesting directions. We have measured technological spillovers using citations of the exporter's prior patents in a destination. However, one may question whether new patents in the destination country then subsequently lead to an increase in productivity in the destination. If the answer is positive, then this should somehow be reflected in future increases in productivity growth for the affected sectors and destinations that are more highly exposed to entry by innovative firms. This and other extensions of our analysis in this paper are left for future research.

References

Aghion, Philippe and Xavier Jaravel, "Knowledge spillovers, innovation and growth," *The Economic Journal*, 2015, *125* (583), 533–573.

_ , Antonin Bergeaud, Matthieu Lequien, and Marc Melitz, "The Impact of Exports on Innovation: Theory and Evidence," NBER Working Paper 24600, National Bureau of Economic Research 2018.

_ , Nicholas Bloom, Richard Blundell, Rachel Griffith, and Peter Howitt, "Competition and Innovation: an Inverted-U Relationship," *The Quarterly Journal of Economics*, 2005, *120* (2), 701–728.

Autor, David, David Dorn, Gordon Hanson, Gary Pisano, and Pian Shu, "Foreign Competition and Domestic Innovation: Evidence from U.S. Patents," Technical Report 22879, National Bureau of Economic Research 2016. DOI: 10.3386/w22879.

Azoulay, Pierre, Joshua S Graff Zivin, and Jialan Wang, "Superstar extinction," *The Quarterly Journal of Economics*, 2010, *125* (2), 549–589.

Bell, Alex, Xavier Jaravel, and Neviana and Petkova, "Team-specific capital and innovation," *American Economic Review*, 2018, *108* (4-5), 1034–73.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, "How much should we trust differences-in-differences estimates?," *The Quarterly journal of economics*, 2004, *119* (1), 249–275.

Bloom, Nicholas, Mirko Draca, and John Van Reenen, "Trade induced technical change? The impact of Chinese imports on innovation, IT and productivity," *The Review of Economic Studies*, 2016, *83* (1), 87–117.

Bombardini, Matilde, Bingjing Li, and Ruoying Wang, "Import Competition and Innovation: Evidence from China," Technical Report 2017.

Chernozhukov, Victor, Iván Fernández-Val, and Blaise Melly, "Inference on counterfactual distributions," *Econometrica*, 2013, *81* (6), 2205–2268.

Chetty, Raj, Adam Looney, and Kory Kroft, "Salience and taxation: Theory and evidence," *American economic review*, 2009, *99* (4), 1145–77.

Coe, David T and Elhanan Helpman, "International r&d spillovers," *European economic review*, 1995, *39* (5), 859–887.

Coe, David T. and Elhanan Helpman, "International R&D spillovers," *European Economic Review*, 1995, *39* (5), 859–887.

Coe, David T, Elhanan Helpman, and Alexander W Hoffmaister, "International R&D spillovers and institutions," *European Economic Review*, 2009, *53* (7), 723–741.

Cohen, Wesley M and Daniel A Levinthal, "Innovation and learning: the two faces of R & D," *The economic journal*, 1989, *99* (397), 569–596.

Goodman-Bacon, Andrew and Jamein P Cunningham, "Changes in Family Structure and Welfare Participation Since the 1960s: The Role of Legal Services," Technical Report, National Bureau of Economic Research 2019.

Keller, Wolfgang and Stephen R Yeaple, "Multinational enterprises, international trade, and productivity growth: firm-level evidence from the United States," *The Review of Economics and Statistics*, 2009, *91* (4), 821–831.

Lequien, Matthieu, Martin Mugnier, Loriane Py, and Paul Trichelair, "Linking patents to firms: insights with French firms," 2019.

Lileeva, Alla and Daniel Trefler, "Improved Access to Foreign Markets Raises Plant-Level Productivity… for Some Plants," *Quarterly Journal of Economics*, 2010, *125* (3), 1051–1099.

Waldinger, Fabian, "Peer effects in science: Evidence from the dismissal of scientists in Nazi Germany," *The Review of Economic Studies*, 2011, *79* (2), 838–861.

Watzinger, Martin, Lukas Treber, and Monika Schnitzer, "Universities and Science-Based Innovation in the Private Sector," 2018. mimeo LMU Munich.

_ **, Markus Nagler, Lukas Treber, and Monika Schnitzer**, "Mobility of Scientists and the Spread of Ideas," Discussion paper series in econompics and management 17-22, German Economic Association of Business Administration 2017.

APPENDIX

A Endogeneity of entry

As explained in section 2.3, our first stage seeks to generate differences in the timing of entry that is *as good as random* within the iso-probability group. In Figure 12 we depict three firms with the same probability of entering a new foreign market in year t. Firm 1 enters this destination at date $t_e < t$. Firm 2 enters that same destination at date $t'_e > t$, while firm 3 never exports to that destination. Consider "treated" firms that enter this destination in year t. We estimate the average effect of entry in that year relative to those 3 firms who did not enter that destination in year t, yet have a very similar probability of having done so.

Figure 12: Exploiting random difference in timing within iso-probability bins

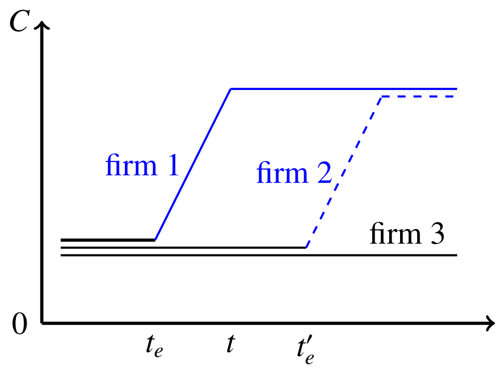

In particular, this grouping will control for two other important types of technological spillovers originating from French firms and their patents. One type does not involve any trade linkages and depends only on the fact that a French firm's technology can be observed via its patent applications (a purely "technological" link). In Figure 12, citations of firm 2's and firm 3's patents in that destination in year t must come via this link (since those firms have not exported to the destination as of year t). The other type of spillover involves a current ongoing trade relationship in year t. Citations of firm 1's patents may fall in this category as this firm is currently exporting in that destination in year t. We use the word "may" as we also measure a potential delayed impact of firm 1's entry in $t_e < t$ in year t. Our regression method allows us to separate out the impact of entry relative to the impact of a current ongoing trade relationship by using the timing of market entry and new citations (observed in new priority patents from that destination).

For each firm-destination-year, we estimate a probability of initial market entry. We estimate this first stage regression as a logit specification:

$$\Pr(ENTRY_{f,j,t}) = \alpha_G GRAVITY_{j,t} + \alpha_F FIRM_{f,t} + \varepsilon_{f,j,t}, \tag{3}$$

Table 2: Probability of First Entry

	$Pr(ENTRY_{f,j,t})$
	Probit + dummies
Ln $GDP_{j,t}$	0.127***
	(6.69)
Ln $GDPpc._{j,t}$	0.448***
	(12.71)
Ln $Employment_{f,,t-1}$	0.110***
	(72.86)
Ln $Productivity_{f,t-1}$	0.0567***
	(15.50)
$\sharp products_{f,t-1}$	0.0052***
	(19.21)
$\sharp Destinations_{f,t-1}$	0.0567***
	(153.36)
$\sharp Patents_{f,t-1}$	-0.0004***
	(-76.76)
Dummies	country + year
Destinations-Years	517025
Pseudo R2	0.202

z statistics in parentheses

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

where: (i) $ENTRY_{f,j,t}$ is a dummy variable equal to one if firm f enters destination j at date t, and is equal to zero otherwise; (ii) $GRAVITY_{j,t}$ is the usual vector of gravity variables measuring the importance of destination country j for France at date t (this includes the geographical distance between France and country j, GDP and per capita GDP of country j at date t); (iii) $FIRM_{f,t}$ includes firm-year characteristics (size, labor productivity measured as value-added per employee).

Table 2 shows the results from this first-stage regression. These results match the standard results we find in the gravity literature. In particular, French firms are less likely to enter destinations that are farther away from France, and more likely to enter bigger foreign markets. Additionally, bigger and more productive French firms are more likely to enter any given foreign market.

We present the distribution of the estimated propensity score for firms that enter and firms that stay out in Figure 13. It illustrates that there is enough overlap between the two distributions to allow the construction of equally likely bins.

We assign French firms to the same bin if their probability to enter country j at date t belongs to the same percentile of the distribution of all the probabilities to enter destination j at date t for all French firms. As a robustness test, we also run specifications with larger-

Figure 13: Distribution of Propensity Score for Entry and non Entry

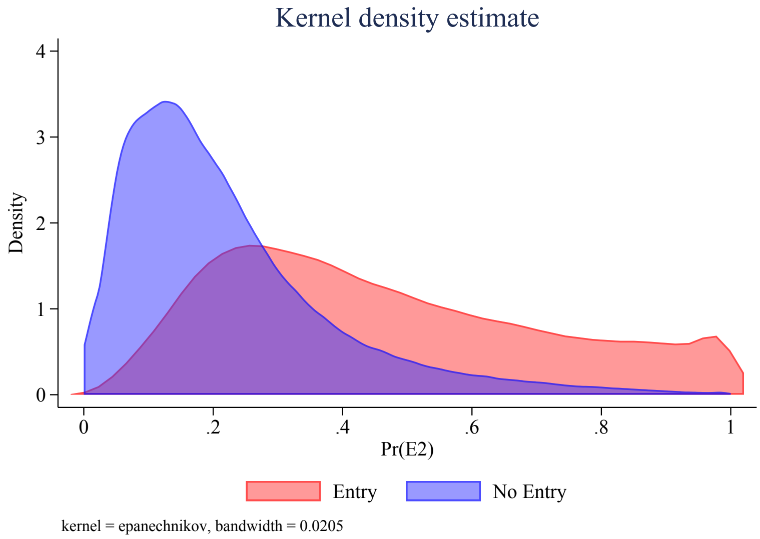

Notes: This figure shows the density function of the estimated propensity score for firms that enter and firms that do not enter.

sized bins.