Spreading Knowledge across the World: Innovation Spillover through Trade Expansion*

Philippe Aghion Antonin Bergeaud Timothee Gigout

Matthieu Lequien Marc Melitz

March 2019

Abstract

When a firm starts exporting into a new destination, its products and technologies suddenly become more visible there. Firms 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; 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 correlation coefficient between the two long differences is equal to 77%.

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 ex-

Figure 1: Evolution of Trade and Innovation Linkages

(a) Number of French Exporters (b) Citations

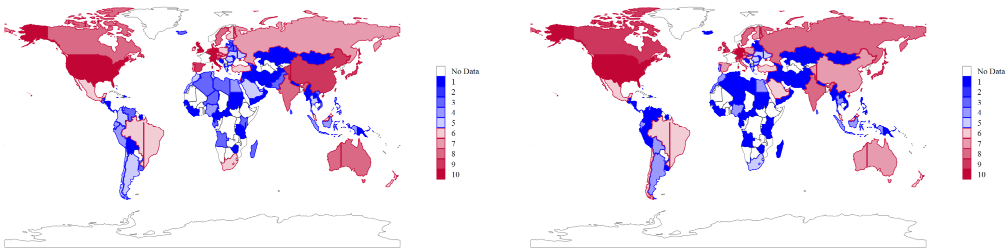

Notes: Evolution in the number of French exporters in each country (left-hand side panel) and the number of citations received from each country (right-hand side panel) between 1995 and 2012. Colors correspond to different deciles in the corresponding quantity.

porters' patents. These citations come from new patents introduced in that year by firms 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).

Our identification strategy to deal with potential selection effects (in particular for the fact that exporting firms have better technologies or technologies that are better suited to the destination country) is adapted from Watzinger et al. (2017, 2018), who study the knowledge spillovers induced by professor transfers across universities. 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: We compare this firm's citations with citations for other French firms with an *ex-ante* similar probability of entry who did not enter that market in that particular year.

We thus start by estimating the probability for each firm of entering an export destination for the first time in any given year. We then group all the French firms that belong to the same probability percentile into an "iso-probability bin" for that destination-year. 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 construct an "entry" variable which is immune to potential selection issues. In a second stage, we examine the impact of entry on the knowledge flows between the entering French exporter and the destination country – *relative* to its control group. We measure the knowledge flow using the number of new priority patents in the destination country citing the French firm's patents.

Following this event-study design, we regress the citation outcome on a set of dummy variables that indicate whether or not the firm entered the foreign market for the first time. We allow the effect to vary across time by including one dummy per year relative to the entry year. We also add an iso-probability bin fixed effect. The coefficients for this regression are thus estimated within a bin of firm-destination-year triplets with very similar entry probability: this is our control group.

We first implement this specification in a dynamic setting with a full set of leads and lags dummies to test for pre-entry effects. Once we confirm the absence of an anticipatory effect, we run the model in a semi-dynamic setting to compute the treatment/causal effect of entry on patent citation.

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.[1] Quantitatively, we find that export market entry induces an 18% increase in the exporter's mean citation rate. We also find that export market entry leads to a 1 pp increase in the probability of receiving citations for exporters with no citations.

Our second finding is that those positive effects are significant only at the top of the productivity distribution for French exporters (it is concentrated amongst the most productive exporters). This is consistent with the view that the patents owned by more productive firms embody more/better knowledge that spills over to other firms and inventors.

Our third finding regards the characteristics of destinations that receive these knowledge spillovers. We find that a destination's 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 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.[2] 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

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

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

recently Jaravel 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.[3]

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.

2 Data and Methodology

2.1 Data

We build a database covering 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

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

analysis in this paper.

We seek to measure the knowledge spillovers from French exporters to firms located in the exporters' sales destinations. Towards this goal, we count the total number of priority patents filed in each destination and year (1995-2012) that cite any patent filed by a French exporter. We restrict our count to *priority* patents as those indicate genuine innovations: Non-priority patents, by and large, reflect a geographical expansion for the protection for a priority patent.

Table 1 summarizes this data. Over our sample years, 5339 French firms have filed patent applications that have been cited at least once in a foreign destination. Across those 137 destinations reached by French exporters, 26552 priority patents have been filed citing those French firms. Of those 26552 linkages, 19691 have been "treated" in the sense that the cited French firm has entered the corresponding export destination during our sample years.

Table 1: Descriptive Statistics

Level	N
Years	18 (1995-2012)
Countries	137
Firms	5,339
Patent	114,993
Links (firms * country)	26,552
$\hookrightarrow$ Ever treated	19,691

Notes: Links to pairs of firm-country (f,j) where the stock of patents of f has received at least 1 citation from j over the observed period.

2.2 Empirical methodology

We want to estimate how a French firm's entry into a new export market affects the flow of new patents (in that destination) citing that firm's patents. 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 follow 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 (same percentile) 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 (a "diff-in-diff" approach). 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 is customary for an event-study, we regress this new patent/citation outcome on a set of dummy variables capturing the time lag (measured in years) relative to the exporter's entry into a destination. We control for export entry selection by adding the 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. In the remaining part of this section we provide further details on this empirical methodology.

2.2.1 First stage regression

As explained above, 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 2 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 2: Exploiting random difference in timing within iso-probability bins

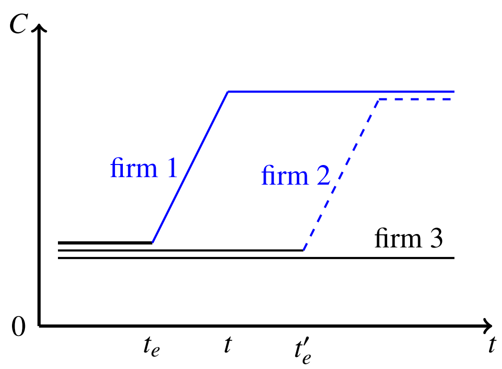

In particular, this grouping will control for two other important types of technological spillovers originating from French firms and their patents. One type does not involve any trade linkages and depends only on the fact that a French firm's technology can be observed via its patent applications

Table 2: Probability of First Entry

	$Pr(ENTRY_{f,j,t})$
Ln $GDP_{j,t}$	0.424***
	(262.50)
Ln $GDPpc._{j,t}$	-0.014***
	(-9.88)
Ln $Distance_{j,t}$	-0.132***
	(-66.61)
Ln $Employment_{f,t}$	0.630***
	(541.3)
Ln $Productivity_{f,t}$	0.217***
	(114.33)
Constant	-4.644***
	(-238.57)
Destinations-Years	452898

z statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

(a purely "technological" link). In Figure 2, 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:

$$\Pr(ENTRY_{f,j,t}) = \alpha_G GRAVITY_{j,t} + \alpha_F FIRM_{f,t} + \varepsilon_{f,j,t}, \tag{1}$$

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 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-sized bins.

2.2.2 Second stage regression

In our second stage, we estimate how the citations in destination j and year t respond to a French firm f's export market entry into j in year t_e. We estimate the following regression via OLS at the patent-level $p(f)$ (instead of the firm-level f) so that we can separately control for each patent's characteristics such as its prior citations and its filing year $t_{p(f)}$:

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

where $Y_{p(f),j,t}$ is the number of priority patents by applicants in destination country j citing patent $p(f)$ at date t; $ENTRY_{f,j,t-k}$ is a dummy equal to one if French firm f enters destination j for the first time at date $t_e = t - k$; and γ_{bin} is the *iso-probability bin* fixed effect (percentile for the firm-destination-year triplet).

We also control for the number of citations the patent has received worldwide prior to entry $X_{p(f),t_e-1} = \sum_{t=-\infty}^{t_e-1} Y_{p,f,World,t}$. We also control for the global cycle of innovation within each technological field when the French and foreign (j) patent were filed by introducing the dummies $\gamma_{t_{p(f)}} \times \gamma_z$ and $\gamma_t \times \gamma_z$, where γ_z is a two-digit technology class fixed effect. We also add a dummy for the entry date γ_{t_e}. Lastly, we cluster the standard errors at the link-level: by firm-country (f, j) pair.

We first run this specification using a fully dynamic set up: that is, we include dummies for a pre-entry effect ($k_{min} \leq -2$). Once we confirm the absence of anticipatory effects, we run the model using a semi-dynamic specification to compute the treatment effect, with $k_{min} = 0$.

We then repeat the same regression using different dependent variables in addition to the number of citations C: (i) $\log(1 + C)$; (ii) $\mathcal{H}(C)$, a hyperbolic function[4] which gives more weight to the extensive margin; (iii) a dummy variable equal to one if $C > 0$, and to zero otherwise; this is simply a linear probability model that allows us to evaluate an extensive margin effect of export on patent citation (a transition from no citations to positive citations); (iv) the log difference of the cumulative stock of citations. Since we can only compute this variable for patents that receive at least 1 citation, this specification conditions on the set of patents that are cited in the destination.

2.2.3 Introducing heterogeneity

To conclude our second-stage analysis, we estimate variants of the specification above. In particular: (i) we first use a static version of the treatment variable with a unique entry dummy equal to 0 before entry, to 1 thereafter, and then to 0 again when/if the firm exits; (ii) we introduce local effects with a kernel re-weighting scheme across the various percentiles in the variable that generates the

[4] $\mathcal{H}$ is the arsinh function: $\mathcal{H}(C) = \frac{1}{2}\log\left(C + \sqrt{1 + C^2}\right)$

heterogeneity. To do so, we follow the methodology detailed by Hainmueller et al., 2016. This kernel approach allows us to flexibly estimate the functional form of the marginal effect of entry on patent citations across the distribution of the heterogeneous variable. We focus attention on two main sources of heterogeneity: (a) the heterogeneity in French firms' productivity; and (b) the heterogeneity in the levels of development across destination countries.

3 Results

3.1 Baseline results

Figure 3 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 (with the dependent variable measured as the flow of citations C). 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 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 3: Main Specification: Priority Citations Count

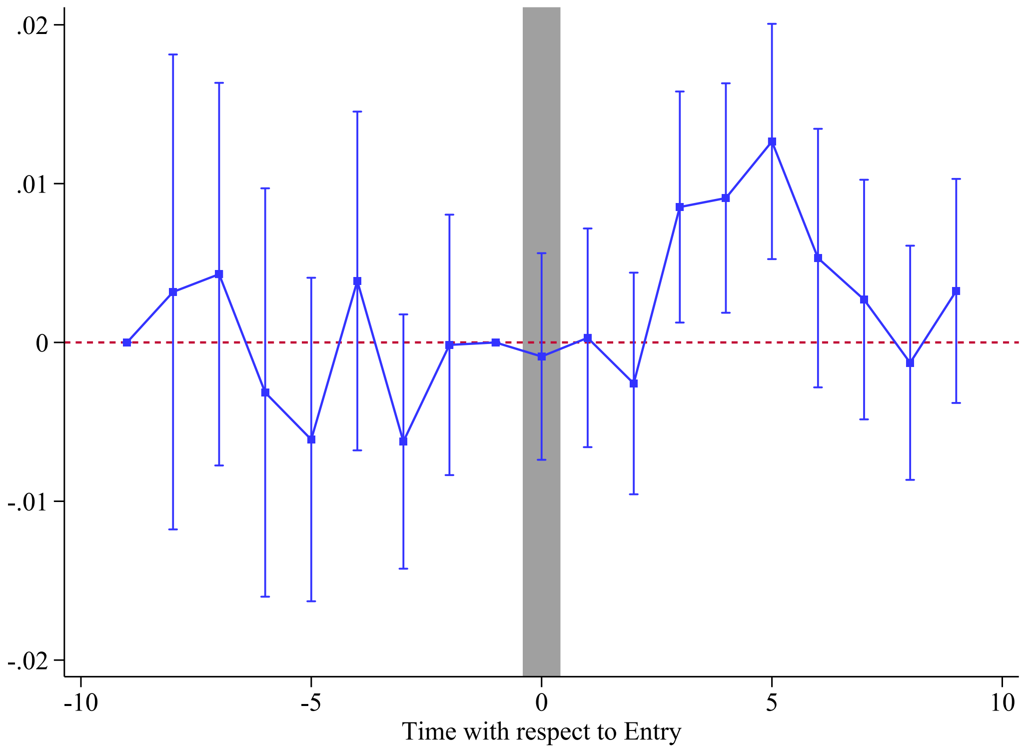

Notes: This figure shows the coefficients β_k from the estimation of our baseline Equation 2. 95% confidence interval are presented. Standard errors are clustered at the link level.

Figure 4 repeats the same exercise as Figure 3, but uses a semi-dynamic specification where we

Figure 4: Y = Priority Citations Count

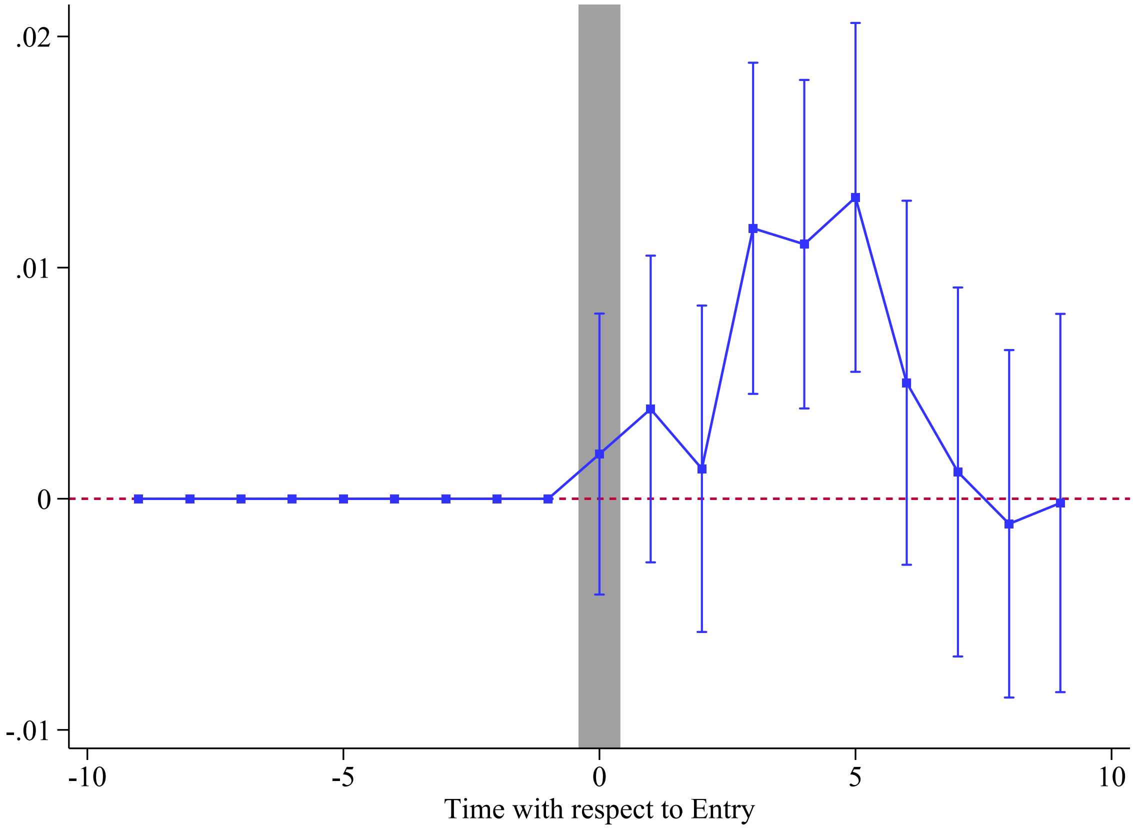

Notes: This figure shows the estimated $\hat{\beta}$ coefficients from the estimation of our baseline Equation 2 in its semi-dynamic specification. The dependant variable is the priority citations count. 95% confidence interval are presented. Standard errors are clustered at the link level.

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.

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 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 0.0497. Over this 9 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.

In the following two figures, we explore the impact of changing the functional form for the number of citations C dependent variable – sticking with our semi-dynamic specification. In Figure 5a the dependent variable is $\ln(1 + C)$, whereas in Figure 5b the dependent variable is $\mathcal{H}[h(C)]$. These

Figure 5: Main Specification: alternative LHS variables

(a) Y = log(1+C) (b) Y = arsinh(C)

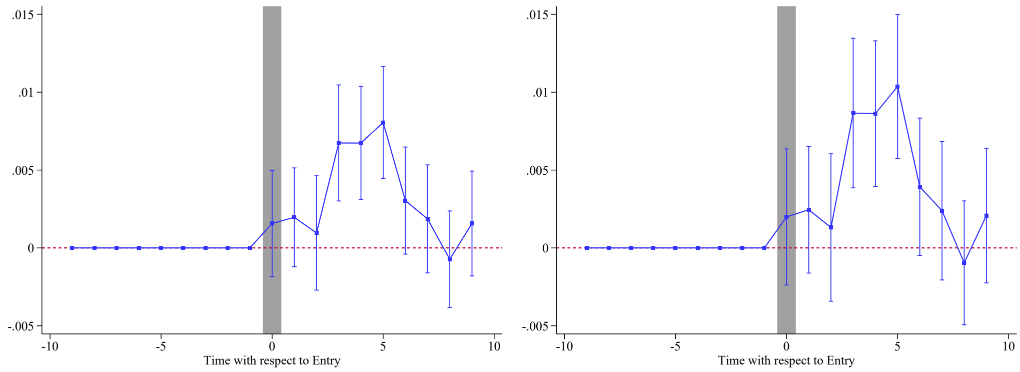

Notes: This figure shows the estimated $\hat{\beta}$ coefficients from the estimation of our baseline Equation 2. The dependant variable is $\log(1 + C)$ in the left panel, and $\text{arsinh}(C)$ in the right panel. 95% confidence interval are presented. Standard errors are clustered at the link level.

figures confirm that the pattern from out baseline Figure 4 is not particularly sensitive to changes in the functional form of the dependent variable.

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 almost 1 percentage point 3 to 5 years after entry. This implies that an entering firm is 36% more likely to obtain a citation relative to a firm that does not enter in that same year (the probability of receiving a citation for such a firm is around 4%).

In Figure 7 we condition on the subset of patents receiving at least 1 citation and measure the increase in citations with the log difference in the stock of citations. This figure shows a pattern that is slightly different from the one in the baseline Figure 4. Once again, the effect becomes significant 3 years after entry; but it reaches its maximum that same year (with a 2.23 percentage point increase in citation from the destination country) and then decays thereafter. The effect is no longer significant beyond year 5. The sum of coefficients amounts to an aggregate coefficient of 0.0689, which yields an average extra 0.7 percentage point growth rate in citations per year from the destination country over the whole time window – compared to firms that did not enter the destination country. Overall, citations to a patent of a firm that entered will have grown by 46 percent versus 36 percent for the patent of a firm that did not enter the destination country.

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

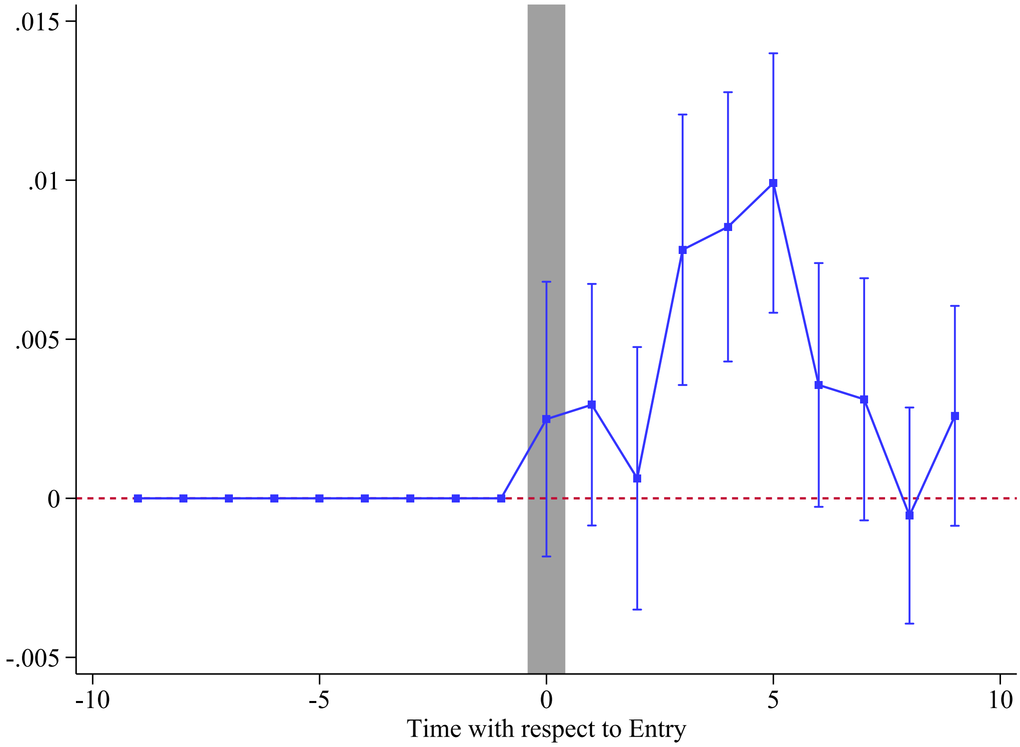

Notes: This figure shows the estimated $\hat{\beta}$ coefficients from the estimation of our baseline Equation 2. The dependant variable is the status of the technological link between the firm's applicants and the foreign country's applicants. 95% confidence interval are presented. Standard errors are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

3.2 Heterogeneous effects

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, to 1 thereafter, and then to 0 again when/if the firm exits. Moreover, we use a kernel re-weighting 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 15 percentiles).

3.2.1 Impact of the exporting firm's productivity

A more productive firm is expected to generate patents that embodies better/more valuable technologies. Those patents are presumably more likely to induce follow-up innovations by other firms,

Figure 7: Y = Δlog Cumulative Citation Stock

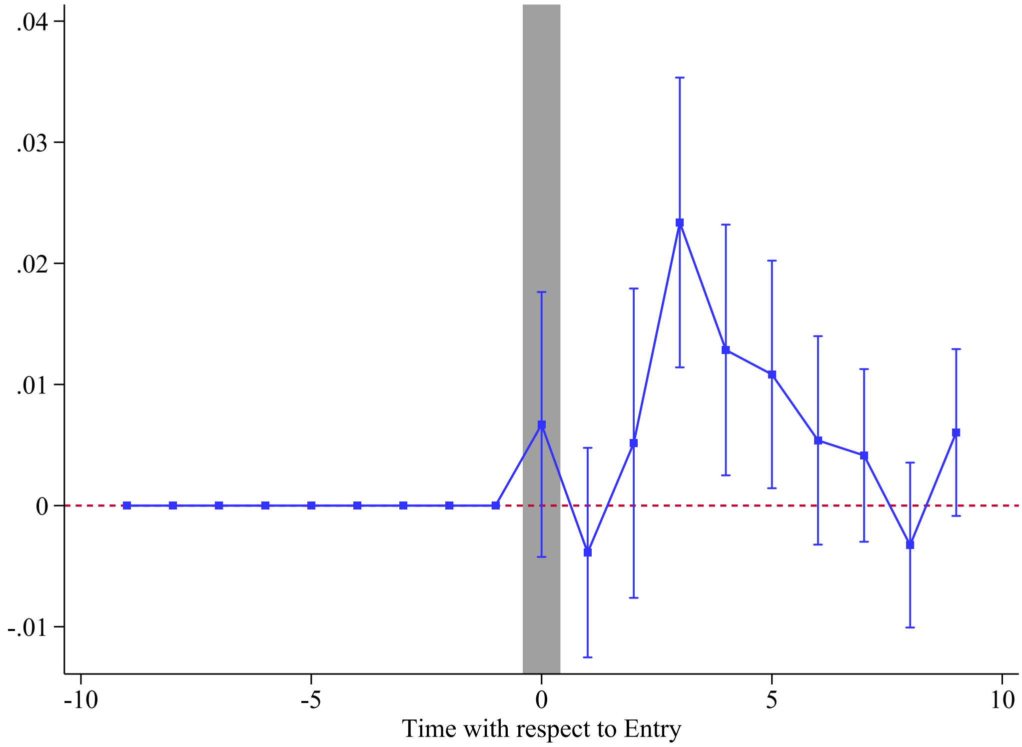

Notes: This figure shows the estimated $\hat{\beta}$ coefficients from the estimation of our baseline Equation 2. The dependant variable is Δlog Cumulative Citation Stock. 95% confidence interval are presented. Standard errors are clustered at the link level.

and should be reflected in additional citations whenever those innovations lead to new patents. To test this prediction, 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 8, 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 citation is linearly increasing in productivity and that spillovers becomes significantly different from zero above the $20th$ percentile.

3.2.2 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 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

Figure 8: Productivity and Spillovers

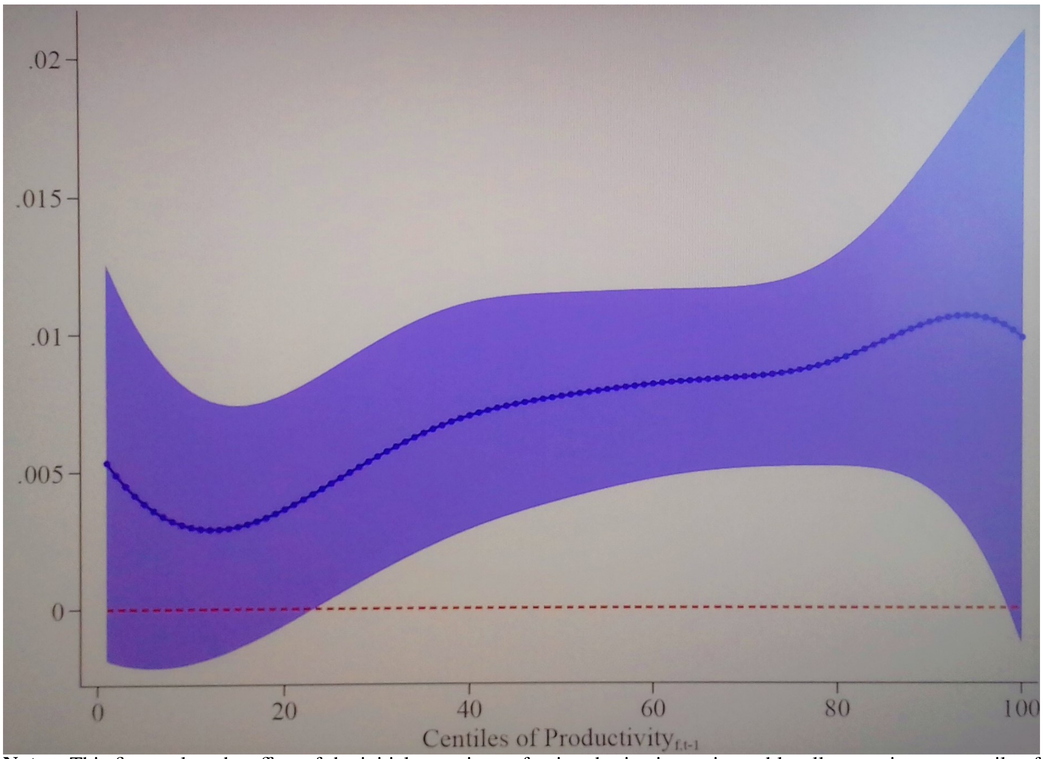

Notes: This figure plots the effect of the initial entry into a foreign destination estimated locally at a given percentile of the ex-ante productivity distribution. The dependant variable is the number of citations. We use Gaussian weights with a bandwidth set to 15 centiles. 90% confidence interval are presented. Standard errors are clustered at the link level.

entry on citations by firms in the destination country may even turn negative. On the other hand, if firms in the destination country are close to 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.[5]

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 40th percentile), entry decreases citations (Figure 9). At intermediate-high level of GDP per capita (between the 40th and the 90th percentile), entry increases citation. And the effect dissipates at higher levels of GDP per capita.

[5]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.

Figure 9: Development and Spillovers

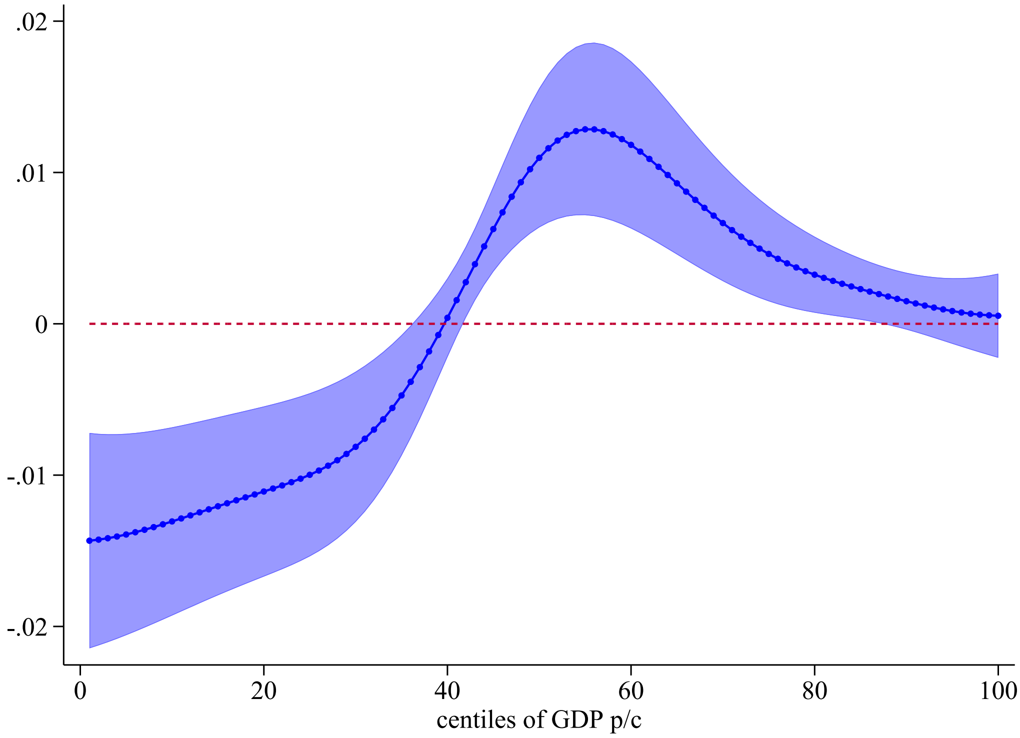

Notes: This figure plots the effect of the initial entry into a foreign destination estimated locally at a given percentile of the per capita GDP. The dependant variable is the number of citations. We use Gaussian weights with a bandwidth set to 15 centiles. 90% confidence interval are presented. Standard errors are clustered at the link level.

4 Robustness

We now report several robustness tests. We first consider departures from our baseline specification, before turning to placebo tests.

4.1 Departing from our baseline specification

4.1.1 Size of the bins

When constructing the iso-probability bins, we face a trade-off between increasing the number of observations per bin and a better approximation of randomness for the timing of entry. Choosing a larger bandwidth for the entry probability provides a higher number of observations per bin but makes each entry within a bin less random (because we cannot control for selection differences within bins). In the baseline specification, we used relatively narrow iso-probability bins representing centiles. As a robustness test, we replicate our regression analysis using bins based on deciles. We find that our results remains qualitatively unchanged (cf. Figure 10).

Figure 10: Different Size of Bins

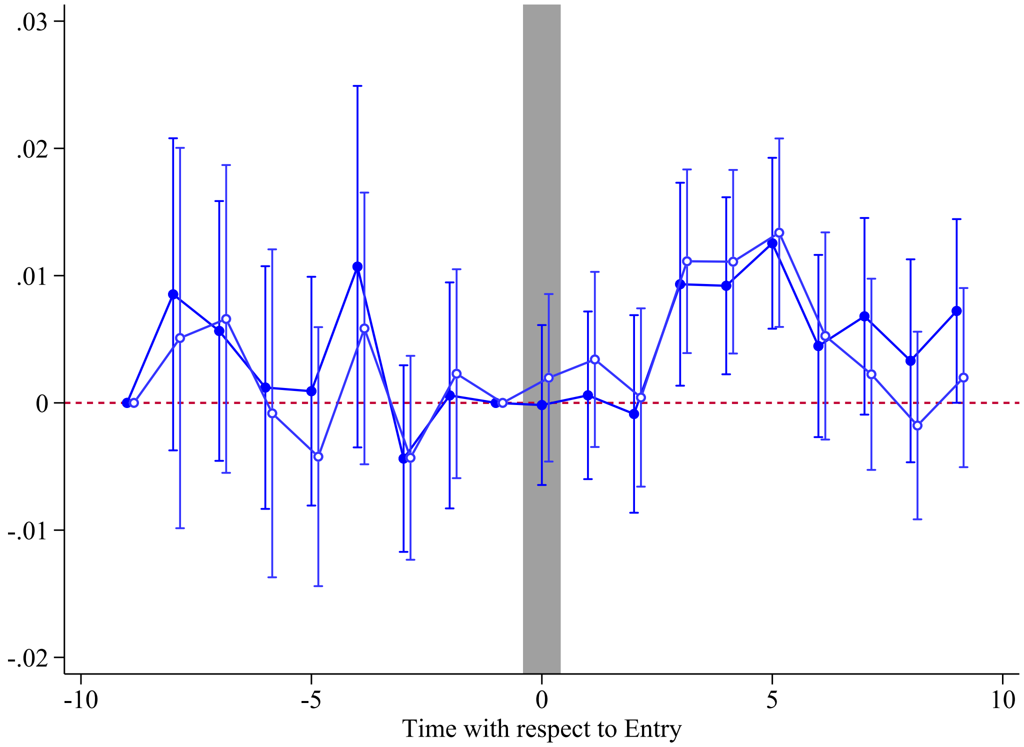

Notes: This figure shows the coefficients from the estimation of our baseline Equation 2 with two different size of iso-probability bins. Here the dependant variable is count of priority citations. 95% confidence interval are presented. Standard errors are clustered at the link level.

4.1.2 Treatment

In the baseline specification, we consider a firm to be "treated" so long as it continues to export to the same destination. Alternatively, we can consider that a firm remains "treated" even after it exits the export destination. We find that our results are robust to this alternative definition of the treatment group.

Additionally, in the baseline specification, once a firm exits the export destination, it is assigned back to the control group. Instead, we could drop that firm from the sample altogether. We find that doing so does not affect our qualitative findings on the dynamic effect of entry on citations (Figure 11)

4.1.3 Sample

A large number of patents receive only very few citations over their life-cycle. We find that our main findings are robust to dropping these patents from our sample (see Figure 12)

Figure 11: Different Definition of Entry

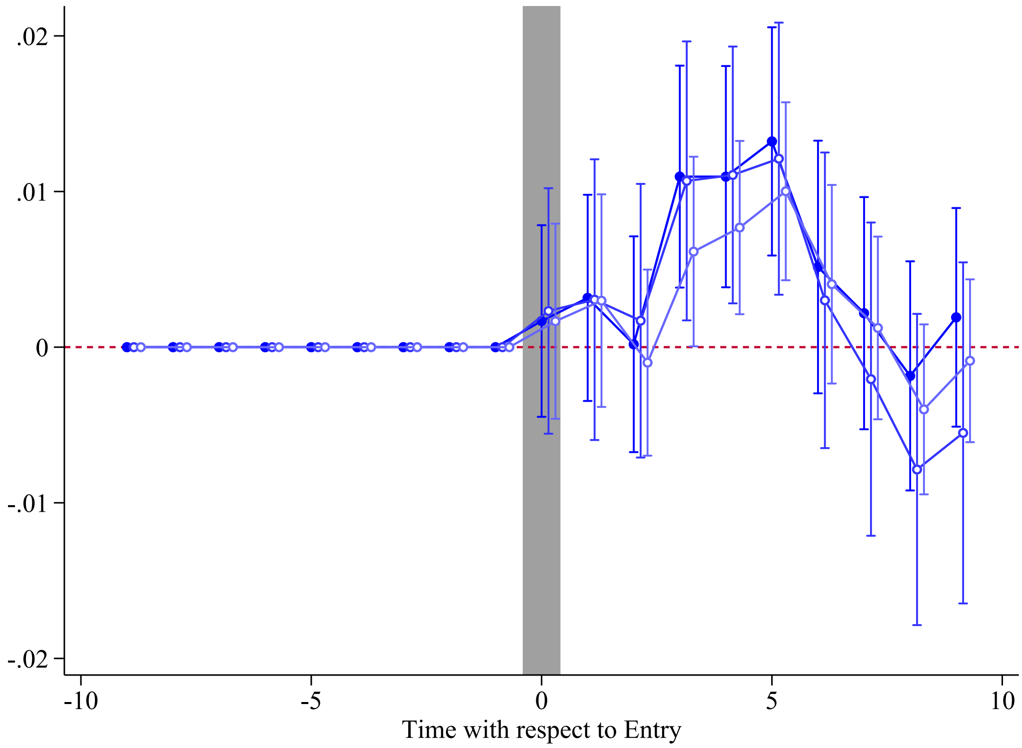

Notes: This figure shows the coefficients from the estimation of our baseline Equation 2 with three different definition of entry. The dark blue line is the baseline definition, the medium blue one corresponds to the case where we drop formerly treated links, the light blue one corresponds to the case where links remain treated even after the firm exit from the market Here the dependant variable is count of priority citations. 95% confidence interval are presented. Standard errors are clustered at the link level.

4.1.4 Other control variables

We estimate the baseline semi-dynamic specification with different set of control variables (see Figure 13). Each line is a specification that includes a different set of control variables. It includes at most: the lagged number of citations of the patent, a dummy variable indicating the year the patent was filled, the natural logarithm of the lagged number of employees of the firm and the natural logarithm of the lagged labor productivity. We find that the pattern and overall treatment effect remains stable regardless of the control variables used. The median value of the overall estimated treatment effect across the different specifications is 0.0471 with a minimum of 0.0385 and a maximum of 0.0477).

4.2 Placebo tests

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 $k = \{5\}$

Figure 12: Removing zombie patents

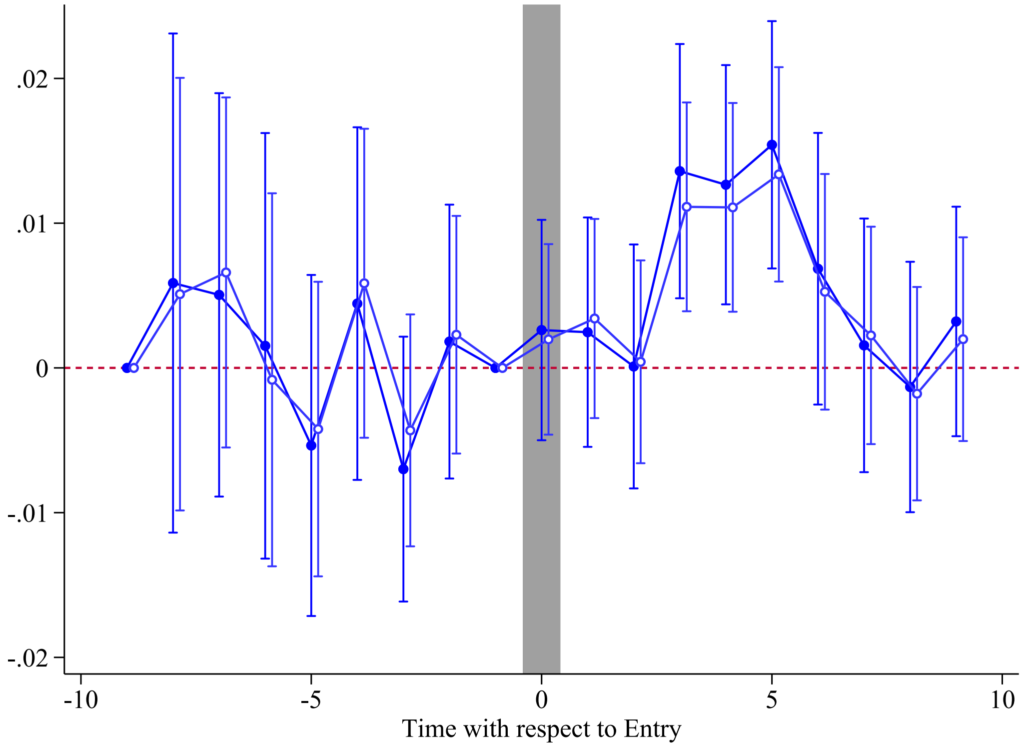

Notes: This figure shows the coefficients from the estimation of our baseline Equation 2. The dark blue line corresponds to the estimation without zombie patents, the light blue line to the baseline sample. Here the dependant variable is count of priority citations. 95% confidence interval are presented. Standard errors are clustered at the link level.

To do so, we randomly reassign the date of entry into an export destination across links and then we re-estimate the second-stage regression. 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 14. The figure confirms that our coefficient of interest $\hat{\beta}_{d=5}$ (the solid blue line) lies on the right of the [p0.5,p99.5] interval (the red dashed lines) of the distribution of placebo coefficients. It confirms that initial entry into a destination leads to an increase in citations.

4.3 Alternative windows of estimation

So far, we have assumed that the date of the first entry as observed in the custom data is the first true year of export into that destination. Here, we relax this assumption. To do so, we define new sample periods and different windows of lags and leads.

The first test corresponds to a fully dynamic specification on a sample between 1999 and 2010 with

Figure 13: Sensitivity to different control variables

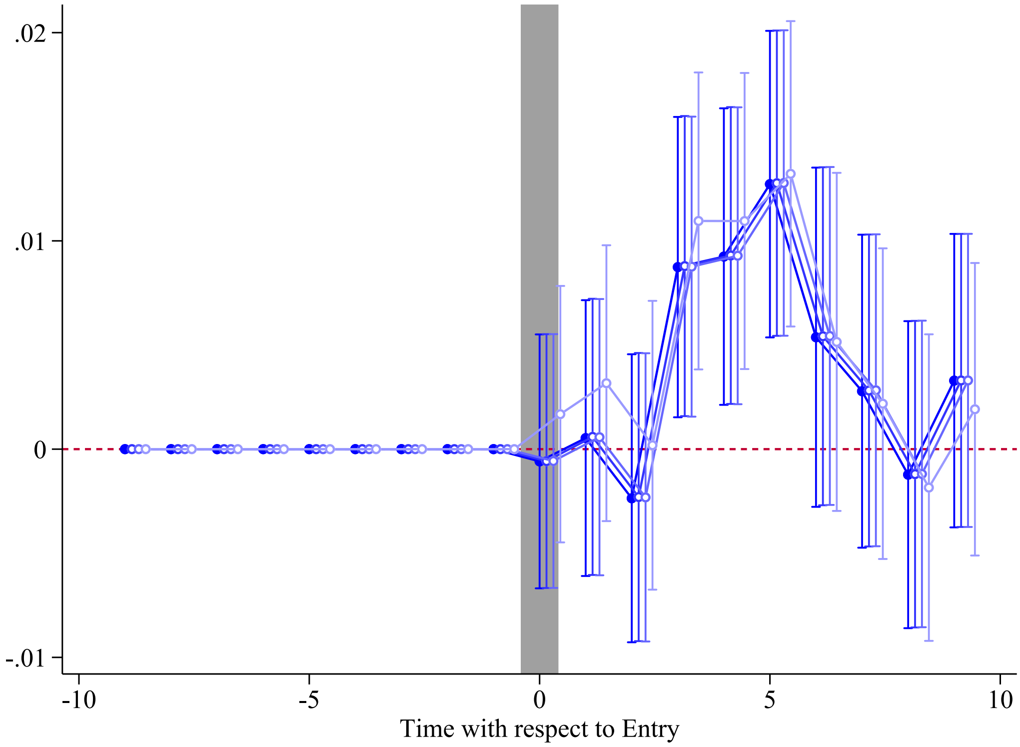

Notes: This figure shows the coefficients from the estimation of our baseline Equation 2. The dark blue line corresponds to the estimation without zombie patents, the light blue line to the baseline sample. Here the dependant variable is count of priority citations. 95% confidence interval are presented. Standard errors are clustered at the link level.

3 periods pre-entry, 1 omitted year just before the entry and 8 periods after the entry. We present the results in Figure 15a. This figure shows no evidence of a pre-trend. The response function follows a pattern somewhat similar to our baseline with a comparable order of magnitude . The main difference is a slightly sharper initial increase during the two years following entry, followed by a negative coefficient three years after entry. The effect then increases to a level near identical to that of the baseline regression.

We then test a semi-dynamic specification in order to include 2011 and 2012 in our sample and reduce the number of estimated coefficients. We omit any pre-entry indicators and keep 8 post-entry indicators. We present the results in Figure 15b. The negative coefficient for k = 2 is no longer statistically significant. The other coefficients are somewhat smaller than in the baseline but the overall response remains persistently positive and significant for two more years. The overall treatment effect is 0.0424 compared to 0.0497 in the baseline (see section 3.1).

We repeat this test for all dependent variables used in the baseline specification and find that our results are robust to the change in specification (see fig 15c, 15d, 15e, 15f). The only substantial difference is the intensive margin response (see 15f). Its response is much more persistent than in the baseline specification.

We then repeat the last test with different combinations of the control variables. We present the

Figure 14: Main Specification: Priority Citations Count

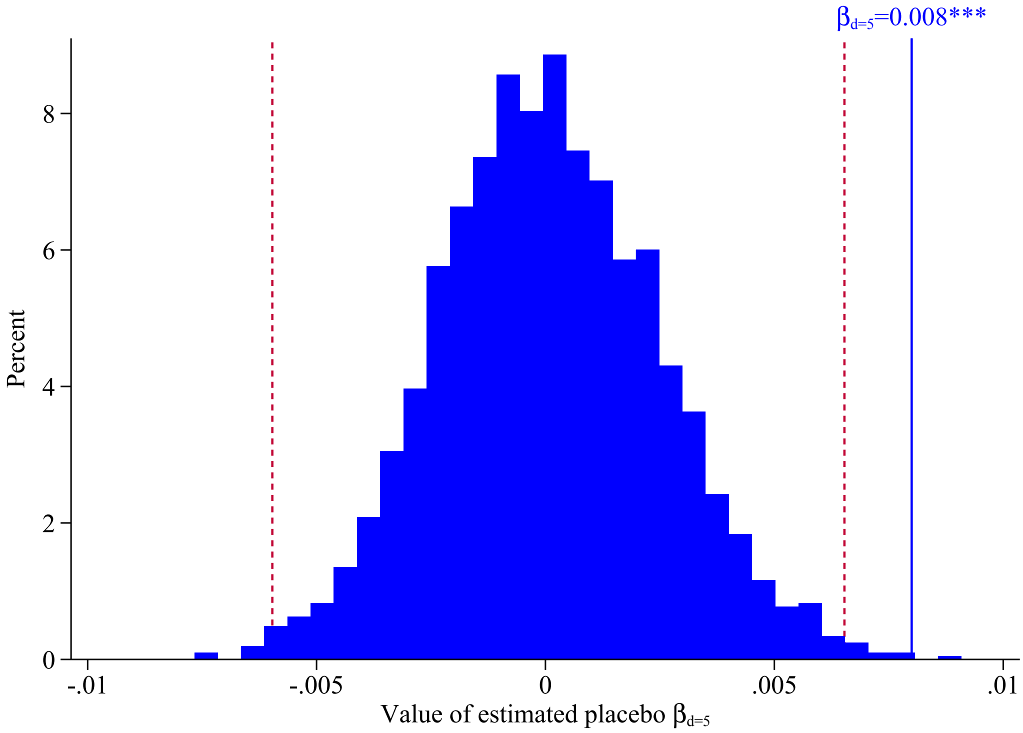

Notes: This figure shows the coefficients from the estimation of our baseline Equation 2. 95% confidence interval are presented. Standard errors are clustered at the link level.

results in Figure 16. We find that the estimated coefficients remains stable to all these changes in the regression specification.

Figure 15: Alternative Window Specification with alternative LHS variables

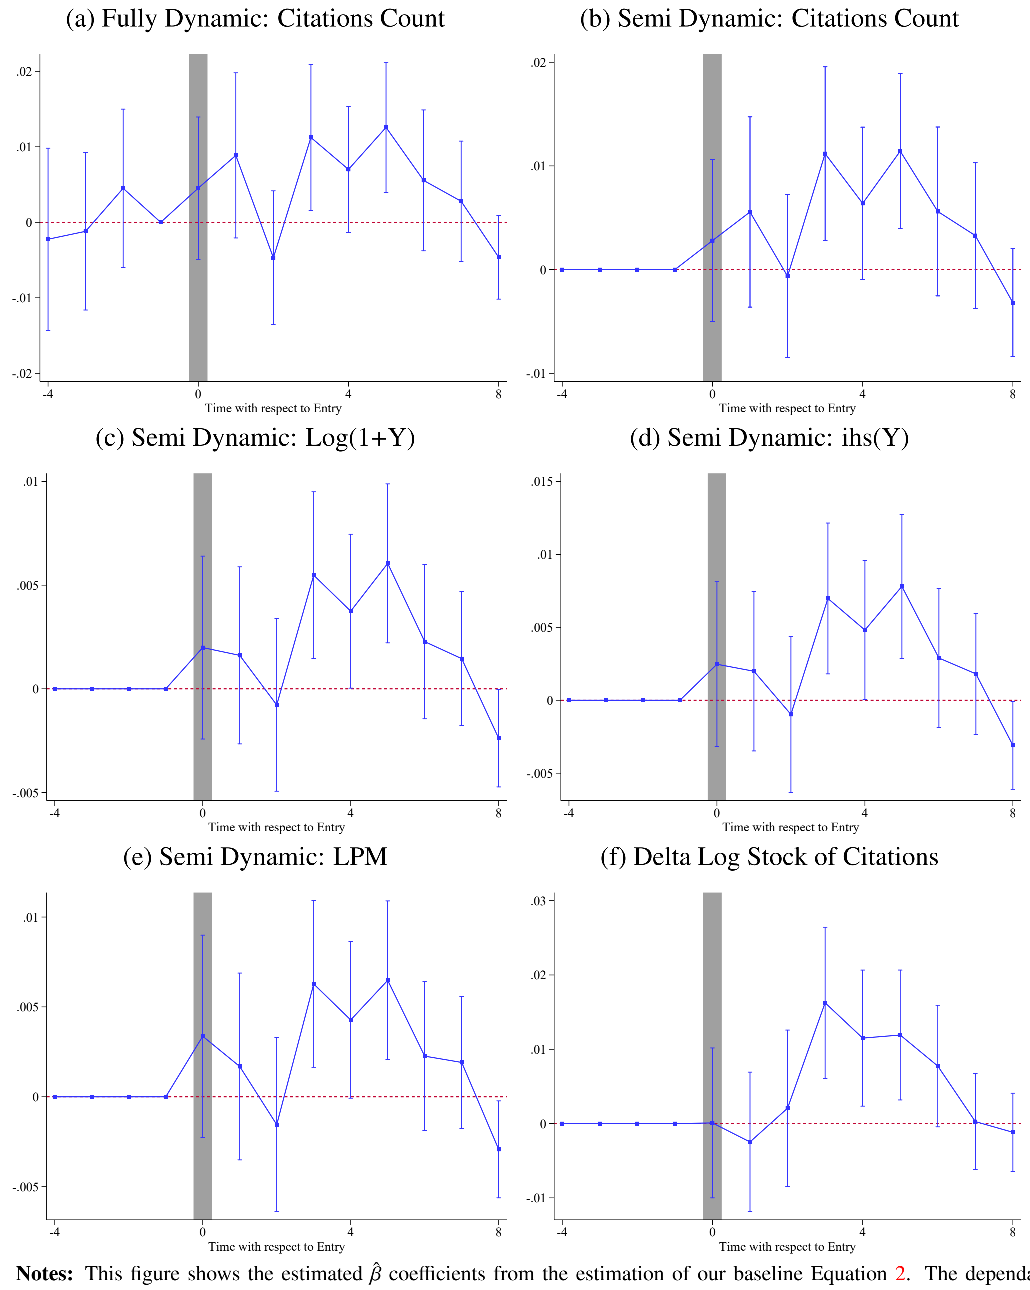

Notes: This figure shows the estimated $\hat{\beta}$ coefficients from the estimation of our baseline Equation 2. The dependant variable is either the count of priority citations, the log of 1+Citations, the inverse hyperbolic sine of Citations, a dummy variable indicating 1 if the patent is receiving citations, and the delta log of the stock of citations. 95% confidence interval are presented. Standard errors are clustered at the link level.

Figure 16: Priority Citations Count with different control variables

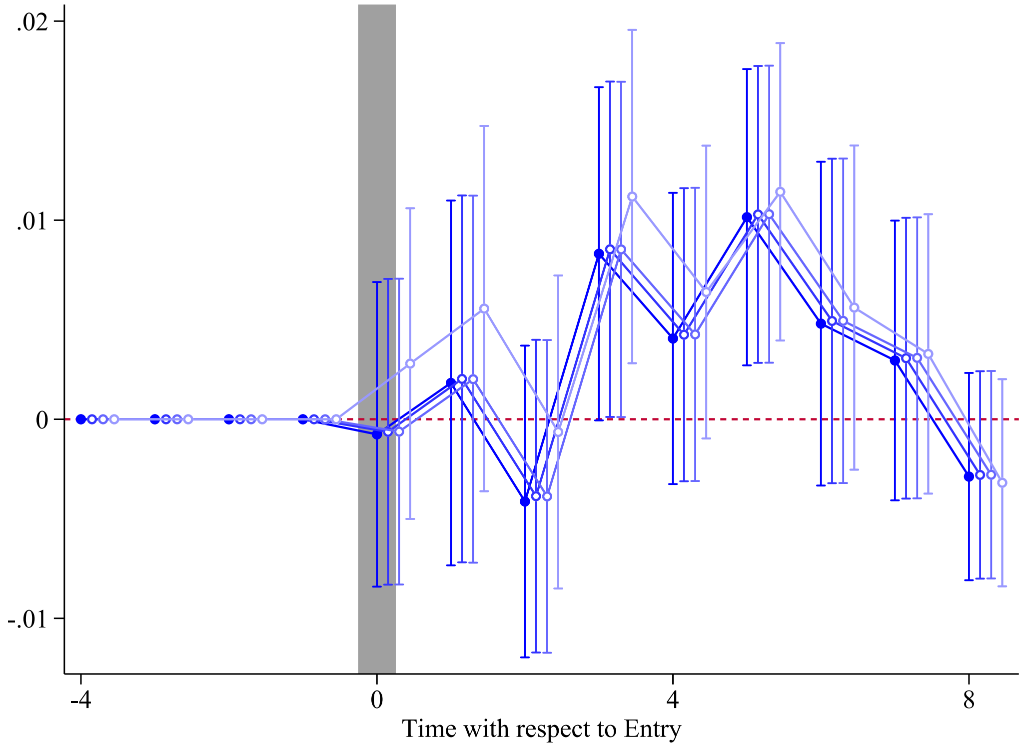

Notes: This figure shows the coefficients from the estimation of our baseline Equation 2. The dependant variable is the count of priority citations. Each line represent the coefficients from a regression with a different vector of control variables. 95% confidence interval are presented. Standard errors 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. Moreover, we find that this effect is concentrated among the most productive French exporters and in destinations at intermediate levels of development. 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.

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.

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.

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.

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.

Hainmueller, Jens, Jonathan Mummolo, and Yiqing Xu, "How much should we trust estimates from multiplicative interaction models? Simple tools to improve empirical practice," *Political Analysis*, 2016, pp. 1–30.

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

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.