The Heterogeneous Impact of Market Size on Innovation: Evidence from French Firm-Level Exports*

Philippe Aghion   Antonin Bergeaud

Matthieu Lequien   Marc J. Melitz

March 2022

Abstract

We analyze how demand conditions faced by a firm in its export markets affect its innovation decisions. We exploit exogenous firm-level export demand shocks and find that firms respond by patenting more; furthermore this response is driven by the subset of initially more productive firms. The patent response arises 2 to 5 years after the shock, highlighting the time required to innovate. In contrast, the demand shock raises contemporaneous sales and employment for all firms regardless of their productivity. This skewed innovation response to common demand shocks arises naturally from a model of endogenous innovation and competition with firm heterogeneity.

JEL codes: D21, F13, F14, F41, O30, O47

Keywords: Innovation, export, demand shocks, patents.

*Addresses - Aghion: Collège de France, LSE and PSE; Bergeaud: Banque de France and CEP; Lequien: Insee, Banque de France and PSE; Melitz: Harvard University and NBER.

1 Introduction

Beyond their immediate impacts, trade shocks can have long-run consequences, in particular on firms' innovation, one of the main driver of long-run economic growth. The economic magnitude of this link is substantial. In our more conservative specification, we find that a 1 percent expansion/contraction in export demand leads to 62 additional/fewer priority patents (corresponding to the first patent publication for an invention) in the French manufacturing sector – a .81 aggregate elasticity. We analyze how the quantity and quality of this innovation response unfold over time and vary across firms with different initial levels of productivity.

In order to study those patenting responses, we merge comprehensive patent records with exhaustive firm-level production and customs data, which cover the whole population of French manufacturing firms. The combined use of these datasets has been made possible by a new algorithm developed in Lequien et al. (2019) that matches a French firm's name with its unique identifier (*Siren*) used in all French administrative business records and allows us to link the innovation activities of a firm with the other firm data sources.

We measure innovation by the flow of *priority* patent applications. All subsequent filings of the same intellectual property (in particular if they are filed at patent authorities in other countries) are secondary filings. We focus on priority patents for two reasons. First because our goal is not to measure a response in patenting but a response in innovation. By focusing on priority patents, we concentrate on patents that correspond to new inventions. Second because we want to avoid capturing the fact that firms that are more involved in international trade are more likely to patent many secondary filings so as to protect their invention in their sales' destinations.

Our first finding is that on average firms respond to a positive export demand shock by innovating more. In other words, we find a significant *market size effect* of export demand shocks on French firms' innovation. Since our specifications always control for sector-year effects, this innovation response must be driven by differences in firm-level innovation responses to demand shocks within each sector. This stands in sharp contrast to the literature measuring sector-wide innovation responses – whether across sectors or for a given sector over time.[1]

Our second finding is that the innovation response to a positive export demand shock takes 2 to 5 years to materialize. In contrast, we find that the response of sales and employment is immediate. We interpret this difference as a confirmation that the response to export demand shocks captures a market size effect.

Our third finding is that the impact of a positive export demand shock on innovation is entirely driven by French firms with above median productivity levels (in an initial period prior to the demand shocks). This heterogeneous response could simply reflect the fact that the demand shock only affects the most productive firms. We check that this is not the case by allowing for a different impact of the export demand shocks on sales or employment depending upon initial productivity levels. We find that in contrast to what we observe for innovation, there is no heterogeneous response of sales or employment to a demand shock for low versus high productivity firms. Thus, similar demand shocks only lead to future innovation responses by relatively more productive firms.

These results provide some additional context to the recent literature documenting the rise of superstar firms: the skewed innovation response is likely to generate further increases in market share for the best performing firms leading to increases in market concentration. Indeed, Autor et al. (2020b) document that this growth in concentration is most apparent in industries with above average growth in patent-intensity.

Our identification strategy relies on the construction of firm-level demand shocks that respond to aggregate conditions in a firm's export destinations but are independent of firm-level decisions (including the concurrent decisions for export-market participation and the forward looking innovation response). Following Hummels et al. (2014), this type of export demand shock has been used extensively in the recent empirical trade

[1] In an influential study, Acemoglu and Linn (2004) measure the sector-wide innovation response of the pharmaceutical industry to changes in demand over time.

literature. It leverages detailed information on the set of products exported to specific destinations by a firm at a prior given date (prior to any changes in innovation that we analyze in our sample). Focusing on this export-driven measure of market size means that we are abstracting from the potential effects of domestic market demand on firms' innovation. For this market, we cannot separate out the causal effects of domestic market size on innovation from the reverse effect of innovation on domestic demand and market size.

We show that our results using this identification strategy are robust to many different specifications including variations in the measure of and functional forms for innovation. We also perform placebo tests that independently confirm that our causation inference from increases in market size to innovation are well founded.

While several explanations might be entertained to explain why the effect of export on innovation should be skewed towards more frontier firms, we show that this outcome arises naturally from a model of exports and innovation with endogenous innovation and markups. In this setting, a positive export demand shock induces not only a direct market size effect – which increases innovation for all firms – but also a competition effect. The idea is that an increase in market size in any export destination will attract new firms into the export market as more firms find it profitable to sell there. And indeed we find a positive correlation between our export demand shocks and various measures of firm entry into the corresponding destination markets. With endogenous markups (linked to endogenous price elasticities), this competition effect associated with entry impinges disproportionately on the market share of the less productive firms, reducing their incentives to innovate. Overall, this combination of the direct market size effect and of its induced competition effect leads to a skewed innovation response between more and less productive firms. Firms closest to the technological frontier increase innovation the most, while the combined effect can even be negative for the least productive firms.

Our analysis relates to several strands of literature. There is first the theoretical literature on trade, innovation and growth (see Grossman and Helpman, 1991a,b, Aghion and Howitt, 2009, chapter 13, and more recently Akcigit et al., 2018).[2] Our paper also relates to the recent empirical literature on firm-level trade and innovation. In particular Lileeva and Trefler (2010); Damijan et al. (2010); Aw et al. (2011) and Bustos (2011) highlight a clear relationship between R&D efforts and export status. Our analysis contributes to this literature in two main respects: (i) this literature focuses on the *extensive* margin of export markets (i.e. whether a firm exports or not to a particular market or set of markets) whereas we consider instead the effect of the *intensive* margin of exports (i.e. of the size of export markets) on innovation;[3] (ii) we use innovation outcomes - the flow of priority patent filings - instead of R&D spending as our main measure of innovation, whereas these papers consider the causal impact of new export markets on R&D spending.[4]

There is also a recent literature on trade and innovation that focuses on the impact of import competition on domestic firms (see Bloom et al., 2016; Iacovone et al., 2011; Autor et al., 2020a; Bombardini et al., 2017). These papers investigate whether import competition induces firms to innovate more in order to escape competition as in Aghion et al. (2005). Empirically, our work is quite distinct as we examine the market expansion channel related to exports. Our theoretical model therefore does not feature an escape competition channel: reductions in market share generate reductions in innovations, though disproportionately so for low productivity firms.

Finally, our work contributes to the empirical literature on market size and innova-

[2] Akcigit et al. (2018) develop and calibrate a new dynamic trade model where firms from different countries compete strategically for leadership in domestic import and export markets. Their model predicts that trade openness encourages innovation in advanced sectors and discourages it in backward sectors. Dhingra (2013) and Impullitti and Licandro (2018) also develop theoretical models with endogenous firm innovation and endogenous competition (via endogenous markups). Dhingra (2013) focuses on the firm-level trade-offs between innovation and product variety, whereas Impullitti and Licandro (2018) focuses on the consequences of innovation for growth and welfare.

[3] Restricting attention to the extensive margin makes it somewhat more difficult to analyze the details of how the market size channel operates: one reason being that several aspects are changing for a firm as it makes the big step of becoming an exporter.

[4] In related work, Coelli et al. (2020) document the patenting response of firms in response to the Uruguay round of tariff levels.

tion, starting with Acemoglu and Linn (2004). We add to this literature in three main respects: (i) by providing evidence of a widespread (manufacturing) firm-level market size effect that is not driven by any sector-level dynamics; (ii) by showing that this market size effect is skewed and mainly driven by the most productive French firms; (iii) by looking at the time dynamics of the market size effect of expanded export markets on firm-level innovation: in particular we show that while a positive export demand shock immediately increases the firm's sales, the innovation response takes several additional years to materialize in new patents. However, one should keep in mind that our analysis is grounded in the market size variations arising from export destinations, which means that we do not use variations coming from the domestic markets. Thus we leave open the question as to whether the domestic market size affects firms in a similar way as the export market size effect that we document.

The remaining part of the paper is organized as follows. Section 2 presents the data and shows some descriptive statistics on export and innovation. Section 3 describes our estimation methodology. Sections 4 and 5 present our empirical results respectively regarding the effect of market size on innovation and its heterogeneous impact with productivity. Section 6 develops a model of export and innovation featuring both a direct market size and an induced competition effect, which predicts that the innovation response to a positive export shock is skewed towards the more productive firms. Section 7 concludes.

2 Exporters and innovators: data and descriptive statistics

In this section, we briefly present our datasets and show some descriptive evidence. Further details about data construction can be found in Appendix A.

2.1 Data sources

Our goal is to explore information on French firms' exports to capture variations in their market size that we can connect to innovation (patenting) outcomes. We also want to look at how this relationship varies across firms with different levels of productivity. Toward this goal, we build a database covering all French firms by linking export, production and innovation data from 1994 to 2012. Our database draws from three sources: (i) French customs, which reports yearly export flows at a very disaggregated product level (representing over 10,000 manufacturing products) by destination; (ii) administrative fiscal datasets (`FICUS` and `FARE` from Insee-DGFiP), which provide extensive production and financial information for all firms operating in France; (iii) the Spring 2016 vintage of `PATSTAT` patent dataset from the European Patent Office, which contains detailed information on all patent applications from many patent offices in the world. In our analysis we will focus on patent applications by French firms, regardless of the origin of the patent office (see below and Appendix A for details).

Matching patents and firms: Although each French firm has a unique identifying number (*Siren*) across all French databases, patent offices do not identify firms applying for patents using this number but instead use the firm's name. This name may sometime carry inconsistencies from one patent to another and/or can contain typos. Various algorithms have been developed to harmonize assignees' names (see Morrison et al., 2017 for a review) but none of those have been applied specifically to French firms. One notable exception is the rigorous matching algorithm developed in Lequien et al. (2019) to link each patent application with the corresponding French firms' Siren numbers, for all firms with more than 10 employees. This new method, based on supervised learning and described in Appendix A.4, provides significant performance improvements relative to previous methods used in the empirical patent literature: based on a verification sample similar to the learning sample, its recall rate (share of all the true matchings that are accurate) is 86.1% and its precision rate (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.

Measure of innovation: Our main measure of innovation consists of a count of *priority* patent applications. This corresponds to the first patent publication that describes an invention. All subsequent filings of the same intellectual property in other jurisdictions (for example in order to extend the geographical coverage of the protection) are secondary filings. We make this restriction for two reasons. First because our goal is not to measure a response in patenting but a response in innovation. By focusing on priority patents, we concentrate directly on patents that correspond to new inventions. Second because we want to avoid capturing the fact that firms that are more involved in international trade are more likely to patent many secondary filings so as to protect their invention in the markets they export to. Priority patents correspond to 35% of the total set of patents but 95% of innovative firms (firms that hold any patent, whether a priority or a secondary filing) in our sample hold at least one priority patent. This suggests that most of the patents we observe in the data are successive secondary filings of the same innovation by the same firm, and legitimate the use of priority applications as our main measure of innovation. Appendix A provides additional details on the construction of our patent measures. For robustness, we report all of our main results using an alternative patent measure based on citation weights for all patent applications by a firm (citations received within a 5 year window). Following Hall et al. (2005), this measure has been widely used in the literature to more accurately capture the innovative relevance of patents. We have also confirmed that our results are robust to a much wider set of patent measures in Appendix C (see in particular Figures C1).

Capturing variations in market size: Finally, to capture variations in firms' market size, we use CEPII's `BACI` database of bilateral trade flows at the HS6 product level (covering more than 5,000 manufacturing products, see Gaulier and Zignago, 2010) to construct measures of demand shocks across export destinations. These data cover the period 1995-2012.

Sample restrictions: Although our main firm-level administrative data source is comprehensive, with more than 46.8 million observations spanning nearly 7.5 million different firms from 1995 to 2012, we restrict our data sample for several reasons. First, we restrict our attention to private business corporations (legal category 5 in the INSEE classification). We thus drop state-owned firms, self-employed businesses, and non-profit organizations as we focus on profit-maximizing firms. Second, we drop firms with less than 10 employees since our matching to the patent data is substantially less complete for those firms (as we previously described). These two restrictions substantially reduce the number of firms in our sample. Yet, the bulk of aggregate employment (77%), sales (79%), and exports (92%) remain in our sample. Those firms are matched with 505,000 patents in `PATSTAT`, including 183,000 priority patents. Lastly, since our detailed customs trade data only covers goods trade (and not services), we will further restrict our sample to the manufacturing sector.[5] This reduces our working sample to 67,684 firms. Nevertheless French aggregate exports and innovation are still concentrated in manufacturing covering 54% of aggregate exports and 43% of patents. Table 1 summarizes these successive sample restrictions and also shows the average number of firms operating in any given year of our sample. For our manufacturing sample, we see that this represents 43,318 firms on average per year between 1995-2012.

Table 1: Successive restriction of the sample

	Total Firms	Firms per Year	Employment	Sales	Exports	Patents
Full	7,474,147	2,597,852	100	100	100	
Private business Corp.	2,888,647	1,114,651	88	90	96	
More than 10 emp	400,662	260,386	77	79	92	100
Manufacturing	67,684	43,318	19	20	54	43

Notes: This Table gives the number of distinct firms and average number of firms per year as well as the share of employment, sales, exports and patents in each sample as compared to the Full (raw) firm level dataset (in %). All columns except the first consider yearly average over the period 1995-2012. Full correspond to our complete sample of firms based on administrative data (see Section 2). "Private Business Corp." corresponds to this sample restricted to firms that are in Legal category ("*catégorie juridique*") number 5. "More than 10 emp" further reduces the sample to firms that are at least once over 10 employees over the period of observation. "Manufacturing" restricts to firms that are always classified in a manufacturing sector.

The case of multinational groups: Our dataset does not allow us to properly take into account the case of multinational groups, an issue which often arises when dealing with national firm level data. The presence of multinational groups tends to break the relationship between export shocks and patenting since these groups may locate their

[5] Although the customs data also covers the wholesale sector, we also exclude those firms as they do not produce the goods that they export.

R&D activities in different countries than the location of production. In particular, the R&D activity for production based in France may be located elsewhere under a different entity of a multinational's group. Conversely, the R&D activity can be located in France whereas a substantial share of production is outsourced in other countries. In these cases, we will not record the appropriate link between the export shocks for this producer and an induced innovation (patents). This measurement issue works against our obtained results of a positive response of patenting to export shocks that is increasing with a firm's proximity to its industry frontier. Thus, we conjecture that our results would be strengthened if we had the needed information to exclude broken production/R&D links amongst the multinational groups in our sample.

2.2 Sector breakdown and skewness

Starting from our sample of manufacturing firms from Table 1, Table 2 shows how those firms are distributed across sectors (using the NACE 2-digit classification), along with their average employment and sales per firm over our sample period from 1995-2012 – shown as yearly averages.[6] Table 2 also shows the proportion of exporters and innovators (firms with at least one patent) in each sector (again, averaged over our sample years) – along with the average exports per exporter (firms with positive exports) and the average number of patents and priority patents per innovator. We clearly see that innovators represent a small minority of manufacturing firms. Only 2.6% of firms introduce any new patents in any given year (on average). Looking across years, 10.7% of firms have at least one patent in one of those years. This is the set of firms we will classify as innovators in our ensuing analysis. Although a minority of firms, they nevertheless represent 39% of employment, 46% of sales, and 58% of exports for the manufacturing sector. In Table 3, we report the same statistics for employment, sales, exports, and patents as sector-level shares. We see that priority patents are concentrated in the computer and electronic,

[6] Throughout, we define sectors at the 2-digit level of the European NACE rev2 classification. We also eliminate the tobacco sector (sector code 12) as it only contains two firms.

machinery and equipment, and motor vehicles sectors, jointly accounting for 44.4% of the priority patents in manufacturing.

Table 2 reveals that the number of patents introduced each year by innovators can be substantial – especially in some sectors. There is a huge amount of dispersion underlying that average number of patents. To highlight this skewness, we show the Lorenz curve for the distribution of those patents in Figure 1, along with the Lorenz curves for exports, sales, and employment in one of our sample years (2007). Figure 1 confirms the previously reported finding that firm-level exports are significantly more skewed than sales and employment (e.g. see Mayer and Ottaviano, 2008 and Bernard et al., 2018): 1% of firms account for 69% of aggregate exports in 2007, whereas the top 1% of firms based on total size account for 51% of sales (ranked by sales) and 33% of employment (ranked by employment). But Figure 1 also shows that patenting is even significantly more skewed than exporting: 1% of all firms account for 92% of priority patents in 2007. Yet, these univariate statistics for patenting and exporting do not capture the massive overlap between these two activities across firms – which we investigate in more detail below.

Table 2: Exports and innovation in the manufacturing sector

			Mean per Firm		%	Mean per Exporter	%	Mean per Innovator	
Sector	Description	Firms	Employment	Sales	Exporter	Exports	Innov.	Patents	Priority
10	Food products	6,775	48	12.8	26	6.8	0.4	5.3	1.5
11	Beverages	606	60	27.9	64	11.6	*	*	*
13	Textiles	1,691	43	5.9	64	3.2	2.2	3.5	1.3
14	Wearing apparel	1,764	41	4.3	53	2.4	0.5	1.8	0.9
15	Leather	514	61	7.3	60	4.1	1.0	1.7	0.8
16	Wood	1,932	30	4.3	42	1.6	0.7	1.5	0.9
17	Paper	2,818	48	9.0	45	4.2	1.3	4.5	1.6
18	Printing	1,530	26	3.6	27	0.7	0.5	3.4	1.7
19	Coke	132	309	469.7	75	106.5	9.1	40.9	10.4
20	Chemicals	980	105	37.1	80	17.4	6.0	10.0	2.9
21	Basic pharmaceutical	311	216	83.8	78	37.3	11.3	15.5	2.0
22	Rubber and plastic	2,374	78	13.0	64	4.9	5.0	5.2	2.1
23	Other non-metallic	1,633	67	13.4	42	5.1	2.6	11.4	2.6
24	Basic metal	1,229	86	20.5	50	15.5	2.8	5.1	1.5
25	Fabricated metal	7,660	34	4.7	39	2.1	1.7	3.2	1.5
26	Computer and electronic	2,613	85	15.7	56	10.3	6.9	8.5	3.2
27	Electrical equipment	621	144	31.3	68	16.2	8.5	15.4	6.5
28	Machinery and equipment	3,376	96	27.6	62	10.0	7.2	5.8	2.5
29	Motor vehicles	1,033	113	32.7	56	21.9	4.2	18.3	13.6
30	Other transport equipment	436	185	56.5	58	40.6	7.2	18.6	7.1
31	Furniture	990	38	4.8	42	1.4	1.1	1.9	1.2
32	Other manufacturing	1,084	47	8.4	53	6.9	4.2	19.0	7.7
33	Repair of machinery	3,027	28	3.4	23	1.3	0.9	4.0	1.5
	All Manufacturing	43,634	58	13.9	45	7.9	2.6	8.0	3.1

Notes: This table presents the number of firms, average employment, sales, employment and exports (sales and exports are in million of Euros, employment in number of employees), the share of exporters, the total number of patents and of priority patents in the sector and the share of innovators.Cells with too few observations to ensure data confidentiality are replaced with ∗. Sector codes correspond to the 2 digit NACE classification. The data presented represents the yearly averages from 1995 to 2012.

Table 3: Relative importance of each sector

NAF	Description	Share of total (in %)					
		Firms	Employment	Sales	Exports	Patents	Priority
10	Food products	15.8	13.1	14.3	7.7	1.5	1.1
11	Beverages	1.3	1.3	2.6	2.8	*	*
13	Textiles	3.7	2.7	1.6	2.3	1.4	1.4
14	Wearing apparel	3.6	2.4	1.2	1.4	0.1	0.2
15	Leather	1.1	1.2	0.6	0.8	0.1	0.1
16	Wood	4.4	2.2	1.4	0.8	0.2	0.3
17	Paper	5.3	4.4	3.6	3.2	1.5	1.4
18	Printing	3.2	1.4	0.8	0.2	0.2	0.3
19	Coke	0.3	1.3	8.6	5.5	4.0	2.8
20	Chemicals	2.3	4.2	6.0	8.8	6.6	4.7
21	Basic pharmaceutical	0.7	2.6	4.1	5.7	5.6	1.9
22	Rubber and plastic	5.5	7.4	5.1	4.9	6.6	7.1
23	Other non-metallic	3.7	4.3	3.6	2.3	5.1	3.1
24	Basic metal	2.6	3.8	3.8	5.9	1.8	1.4
25	Fabricated metal	17.8	10.5	5.9	4.0	4.6	5.4
26	Computer and electronic	5.3	7.7	6.1	9.1	15.2	15.1
27	Electrical equipment	1.3	3.1	2.8	4.1	7.7	8.3
28	Machinery and equipment	7.5	12.1	15.0	13.5	14.5	16.4
29	Motor vehicles	2.2	4.5	5.0	6.6	7.5	12.9
30	Other transport equipment	1.0	3.1	4.0	6.9	7.0	6.8
31	Furniture	2.3	1.5	0.8	0.4	0.2	0.4
32	Other manufacturing	2.4	1.9	1.4	2.5	7.2	7.5
33	Repair of machinery	7.0	3.4	1.7	0.6	1.3	1.2

Notes: This table presents the share of value added, employment, export and patents (all patents and priority patents) accounted for by each 2-digit manufacturing sector as well as the share of firms in each sector. Cells with too few observations to ensure data confidentiality are replaced with $*$. Data are averaged over the period 1995-2012.

Figure 1: Lorenz curves for priority patents, exports, sales and employment

(a) Whole distribution (b) Top 5 percentiles

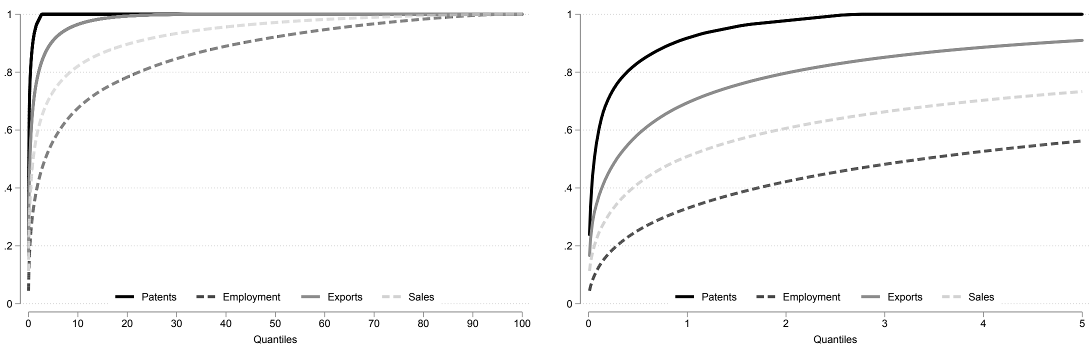

Notes: Lorenz curves plot cumulative distribution function for priority patents, employment, export and sales. Data are for manufacturing firms and for the year 2007.

2.3 The nexus between innovation and exports

Looking across our sample years (1995-2012), Table 4 reports different size-related performance measures (averages per firm) based on their exporter and innovator classification.

As we previously discussed, we classify firms as innovators if they introduced at least one patent during those sample years. From here on out, we classify exporters in a similar way as a firm with positive exports in at least one of our sample years. This raises the proportion of exporting firms to 62% of our manufacturing sample (45% of firms export on average in any given year, c.f. Table 2). Table 4 confirms the well-documented size differential in favor of exporters. However, several new salient features regarding innovators pop-out from this table. First, innovating firms are massively concentrated among exporters: only 4% of innovators do not report any exporting. Second, non-exporting innovators do not look very different from non-exporting non-innovators, and the various measures of firm size (employment, sales, value-added) respectively for innovators and non-innovators among non-exporters remain close to each other;[7] and third, these same measures of firm size differ markedly between innovators and non-innovators among exporters: innovators employ on average 4.2 times more workers and produce 5-6 times more output and value-added than non-innovating exporters. They export almost 10 times more than non-innovators and reach more than three times the number of export destinations. These size differentials are several times larger than those between exporters and non-exporters. In the aggregate, this small subset of innovators accounts for over half of French manufacturing exports.

In order to compare exporters to non-exporters and innovators to non-innovators, within specific groups, we compute export and innovation premia (in log points). Consider first the exporter premia reported in the top panel of Table 5. These premia are generated by regressing the performance measure of interest (listed in the rows) on our exporter indicator – with each cell representing a separate regression. Column 1 includes no other controls; Column 2 adds a 2-digit sector fixed effect (see Table 2); and Column 3 controls for firm employment, in addition to the sector fixed effect. Since we are using a broad

[7]This is not the case outside of the manufacturing sector. In those other sectors, non-exporting innovators are substantially bigger than their non-exporting and non-innovating counterparts. We conjecture that this is driven by the fact that exporting no longer serves the same performance screening function outside of manufacturing.

Table 4: Exporters and innovators are bigger

	Non-exporter		Exporter		Total
	Non-innovator	Innovator	Non-innovator	Innovator	
Firms	13,411	173	25,292	4,442	43,318
Employment	20	19	51	215	58
Sales	4.1	2.4	10.8	62.3	14.1
Value Added	0.8	0.9	2.7	14.7	3.3
Export	0.0	0.0	2.4	20.7	3.6
Countries	0.0	0.0	4.8	17.3	4.6
Products	0.0	0.0	5.0	16.1	4.6
Patents	0.0	0.1	0.0	0.8	0.1

Notes: This table presents basic descriptive statistics across four categories of manufacturing firms whether they innovate, export, both or none. Employment is given in full-time equivalent on average over the year and exports, sales and value added are in million euros. Countries is the number of destination countries for exports. Employment, Sales, Value Added, Age, Exports, Countries and Patents are taken as a yearly average over the whole period 1995-2012.

Table 5: Export and innovation premia

Panel 1: Premium for being an exporter (among all manufacturing firms)					
	(1)	(2)	(3)	Obs.	Firms
log Employment	0.860	0.835		760,700	67,563
log Sales	1.354	1.339	0.468	771,306	67,601
log Wage	0.122	0.101	0.114	759,453	67,547
log Value Added per Worker	0.210	0.185	0.184	750,495	67,094
Panel 2: Premium for being an innovator (among all exporting manufacturing firms)					
	(1)	(2)	(3)	Obs.	Firms
log Employment	1.002	0.977		523,714	42,582
log Sales	1.265	1.239	0.206	530,076	42,602
log Wage	0.119	0.096	0.111	522,923	42,577
log Value Added per Worker	0.207	0.184	0.185	516,099	42,341
log Export Sales (Current period exporters)	2.010	1.897	0.795	349,222	42,201
Number of destination countries	12.54	11.48	6.97	535,210	42,644

Notes: This table presents results from an OLS regression of firm characteristics (rows) on a dummy variable for exporting (upper table) or patenting (lower table) from 1994 to 2012. Column 1 uses no additional covariate, column 2 adds a 2-digit sector fixed effect, column 3 adds a control for the log of employment to column 2. Wage corresponds to total payroll divided by employment. All firm characteristic variables are taken in logs. All results are significant at the 1 percent level. Upper table uses all manufacturing firms whereas lower table focuses on exporting manufacturing firms.

cross-year definition for exporter status, we expect these premia to be lower than measures based on current-year exporter status since firms who drop in and out of export markets tend to be substantially smaller than year in year out exporters. This is the case for the premia in column 1 compared to similar numbers reported by Bernard et al. (2018) for U.S. firms in 2007. Yet, once we control for sectors in column 2, the reported premia become much more similar. In particular, we find that even within sectors, exporters are substantially larger than non-exporters. And we also find that large differences in productivity and wages in favor of exporters persist even after further controlling for firm employment.

In the bottom panel, we focus on the subset of exporters from the top panel, and report the *additional* premia in favor of innovators within this subset. As with the top panel, those premia are calculated by running separate regressions on our innovator indicator. Even within this subset of bigger and better performing firms, innovators stand out: they are substantially bigger, more productive, and have larger total wage bill. They also export substantially more (and to more destinations) than non-innovative exporters. All these differences persist within sectors and controlling for firm employment.

Even these large premia do not fully reflect the concentration of innovative and exporting activities within the more restricted subset of firms that are both exporters and innovators. Figure 2 plots the share of innovating firms for each percentile of the firm export distribution. We see that the innovative firms are highly concentrated within the top percentiles of the export distribution. At the 80^{th} percentile of the export distribution, 29% of the firms have some patenting experience. And the increase in the share of innovative firms with the percentile of the export distribution is highly convex. Above the 95^{th} percentile of the export distribution, a majority of firms are innovators; in the top percentile, 66% of the firms are innovators. Those firms in the top export percentile account for 41% of the aggregate share of French patents.[8]

[8] Of course part of the relationship in Figure 2 could be driven by a scale effects: large firms tend to export more and are more likely to innovate. When we rank firms in percentile of export intensity (instead of absolute export) we still find a near monotonic increase in the share of innovators for export intensity in the 5-95% range. After this threshold, the relationship becomes negative as the last 5 percentiles of export intensity are dominated by unusual small firms that export virtually all of their sales.

Figure 2: THE SHARE OF INNOVATORS JUMPS AT THE TOP OF THE EXPORT DISTRIBUTION

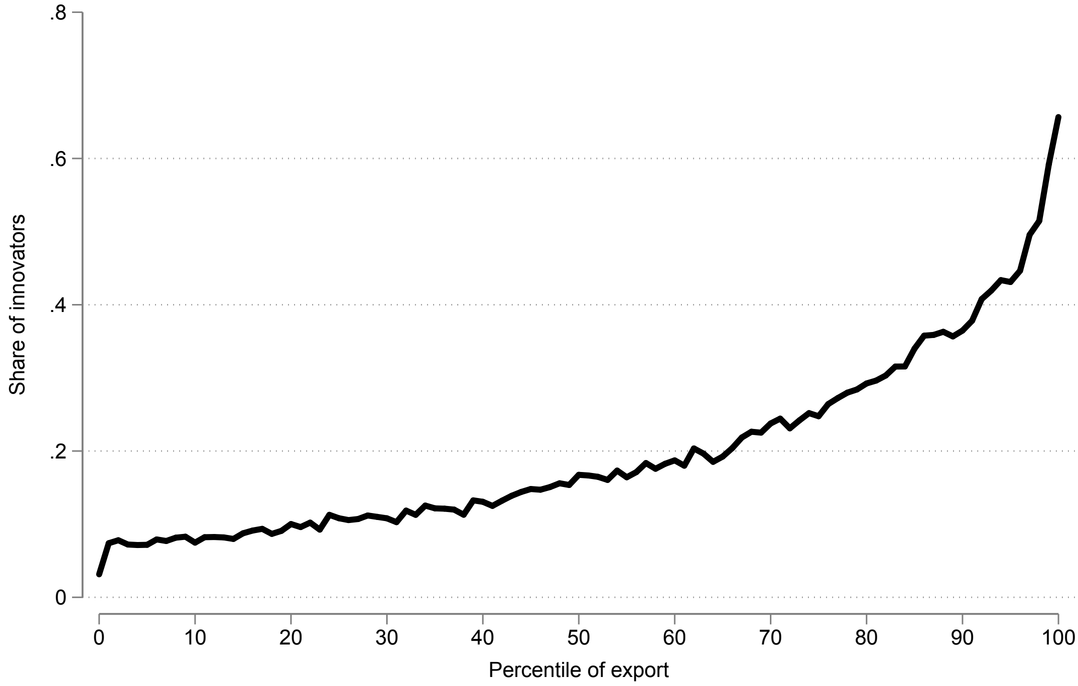

Notes: Percentiles of exports are computed each year from 1995 to 2012 separately and then pooled together. For each percentile, we compute the share of innovators. Each percentile contains the same number of firms, except for percentile 0 that contains all the firms with no export. Manufacturing firms only.

3 Empirical Framework

3.1 Firm level export demand shocks

We have just documented a strong correlation between exports and innovation in the cross-section of French manufacturing firms. However, this correlation does not say much about the direction of causation: from innovation to exports (a major innovation leads to growth in export demand and entry into new export markets), or from exports to innovation. Moreover, other firm-level changes could generate concurrent changes in both innovation and exports (for example, a new management team). Thus, to identify the causal relationship from exports to innovation, we need to identify a source of variation in firm exports that is exogenous to changes within the firm (and in particular to the innovation activity of the firm). We follow Mayer et al. (2020) in building an exogenous firm-level measure of export demand shocks.

To construct these export demand shocks, consider a French exporter f who exports a product s (measured at the HS6 level) to destination j at an initial date t_0. Let $M_{j,s,t}$ denote the aggregate import flow in product s into country j from all countries except France at time $t > t_0$. $M_{j,s,t}$ reflects the size of the (s, j) export market at time t. We then sum over the $M_{j,s,t}$ across destinations j and products s weighted by the relative importance of each market (s, j) in firm f's exports at the initial date t_0. The underlying idea is that subsequent changes in destination j's imports of product s from the world (excluding France) will be a good proxy for the change in export demand faced by this firm. By excluding French exports to this destination, we seek to exclude sources of variation that originate in France and may be correlated with changes for the firm.[9]

We then scale the weighted export demand variable by the firm's initial export intensity (at t_0) so that our demand shock scales proportionately with a firm's total production (as a firm's export intensity goes to zero, so does the impact of any export shock on total production).

Formally, t_0 is the first year with positive exports in both customs (to compute destination market shares) and production data (to compute export intensity).[10] X_{f,j,s,t_0} denotes firm f's export flow to market (j, s) at time t_0. The export demand shock for firm f between t and $t-1$ is then constructed as:

$$\Delta D_{f,t} = \sum_{j,s} w_{f,j,s,t_0} \left(\frac{M_{j,s,t} - M_{j,s,t-1}}{\frac{1}{2}(M_{j,s,t} + M_{j,s,t-1})} \right), \tag{1}$$

where the weight $w_{f,j,s,t_0} \equiv (X^*_{f,t_0}/S^*_{f,t_0})(X_{f,j,s,t_0}/X_{f,t_0})$ represents firm f's initial share of sales of product s, at the HS6 level, to destination j and $X_{f,t_0} = \sum_{j,s} X_{f,j,s,t_0}$ represents the firm's total exports at date t_0. The asterisks on firm f's initial export intensity

[9]One potential source of endogeneity may arise in markets where a French firm has a dominant position. We check that our results are robust to dropping firm-destination pairs whenever the firm's market share in the destination exceeds 10%. See Figure C3 in Appendix C.

[10]This year is 1994 for about half of the firms and is used as a reference year in most of our analysis.

$X^*_{f,t_0}/S^*_{f,t_0}$ indicate that the underlying data for total exports X^*_{f,t_0} and sales S^*_{f,t_0} come from the production data (as opposed to customs data which we use to calculate the destination/product specific market shares).[11]

There are some clear outliers in the distribution of this demand shock $\Delta D_{f,t}$ across firms. They typically involve firms that export a small number of often highly specialized products to small destinations (such as yachts to Seychelles and Maldives). In order to deal with these outliers in a consistent way, we trim our demand shock $\Delta D_{f,t}$ at 2.5% (eliminating those trade shocks below/above the 2.5th and 97.5th percentiles in each year). We report our main results on the response of innovation to this trade shock using trimming thresholds between 0-5% in Appendix C (Figures C4).

Demand shock as a shift share instrument: We note that the time variation in our demand shock $\Delta D_{f,t}$ only stems from the variation in the world export flow $M_{j,s,t}$ and not in the firm-level weights, which are fixed at their value in the initial export period t_0. We expect that a firm's innovation response at time $t > t_0$ will induce changes to its pattern of exports at time t and beyond, including both intensive margin responses (changes in exports for a previously exported product s to a destination j) and extensive margin responses (changes in the set of products s sold across destinations j). By fixing the firm-level weights in the initial period t_0 (including the extensive margin set of products and destinations), we exclude this subsequent endogenous variation in exports from our demand shock. This is quite similar to a standard shift-share or "Bartik" (Bartik, 1991) setting in which aggregate shocks are combined with measures of shock exposure. In our case the sum of exposure weights w_{f,j,s,t_0} across (s, j)'s is different from 1 and varies across firm. We follow Borusyak et al. (2021) who argue that in such "incomplete shift-share" case with panel data, one needs to control for this sum interacted with a time dummy in

[11] Total exports reported by customs and in the production data do not always exactly match, though they are highly correlated. One potential source of difference comes from small exports towards other European Union countries which are not reported in customs data (see Appendix A for more details).

our regressions.[12]

3.2 Estimation strategy

Here we spell out the baseline regression equations of French firm's innovation on the export demand shock variables $\Delta D_{f,t}$. Our identifying assumption is that after controlling for any sector-level variation by year and firm characteristics at and prior to t_0, subsequent variations in the firm-level export demand shock are uncorrelated with firm-specific shocks to innovation.

As we have no presumption regarding the timing of this innovation response to demand shocks, we include a full set of lags and leads for the demand shock $\Delta D_{f,t}$ in our regressions. Our identification strategy nevertheless relies on the fact that our shock is independent of previous innovation decisions and we will check that the response of innovation to future shocks remains insignificant – in other words, no pre-trends.[13] We restrict our analysis to the subset of innovating firms (i.e. firms with at least one patent between 1985 and 2012), and check that entry into innovation subsequent to 1994 does not bias our sample.[14] Out of our sample of 67,684 manufacturing firms (see Table 1), there are 5,308 such innovators. Not all of them are active throughout our sample period.

[12] Borusyak et al. (2021) discuss the possibility of purging the shift-share instrument from a specific dimension. In our case, we have experimented using the residual of $M_{j,s,t}$ on different sets of fixed effects (product and time, country and time) in the construction of the demand shock (using the same weights as in the baseline). Our results are qualitatively unchanged. Even with the unresidualized shock, the effective number of shocks, measured by the inverse concentration index of the exposure share, is also much smaller than the number of export destination markets (s, j) (see Borusyak et al., 2021 and Appendix E for more details).

[13] A simpler and more naive approach would have been to estimate a static model in which the dependent variable would be directly regressed on the demand shock lagged appropriately. Although such a static model delivers predictions consistent with our results, the coefficients would most likely be biased in the sense that the firm's response is also affected by subsequent and previous shocks.

[14] In Appendix C, Figure C5 shows that our main results are essentially unchanged when we further restrict the sample to firms who innovated before 1994. Our sample also includes firms for which we can define a t_0, i.e. firms that exported at least once since 1994. t_0 is used as a reference year and can be any year from 1994. Figure C6 shows that our results hold if we restrict to firms for which $t_0 = 1994$.

On average across those years there are 1,140 innovators in our sample (2.6% of 43,318 manufacturing firms operating in a given year).

Our main estimation strategy is described by:

$$\begin{aligned}\Delta Y_{f,t} &= \left(\sum_{\tau=-k'}^{k} \alpha_\tau \Delta D_{f,t-\tau}\right) + \boldsymbol{\gamma} \cdot \boldsymbol{Z}_{f,t_0} + \tilde{\boldsymbol{\gamma}} \cdot \left(\tilde{\boldsymbol{Z}}_{f,t_0} \times \chi_t\right) + \varepsilon_{f,t} \\ &= \boldsymbol{\alpha} \cdot \Delta_k \boldsymbol{D}_{f,t} + \boldsymbol{\gamma} \cdot \boldsymbol{Z}_{f,t_0} + \tilde{\boldsymbol{\gamma}} \cdot \left(\tilde{\boldsymbol{Z}}_{f,t_0} \times \chi_t\right) + \varepsilon_{f,t},\end{aligned} \tag{2}$$

where $\Delta Y_{f,t}$ is firm f's outcome of interest between t and $t-1$; $\boldsymbol{Z}_{f,t_0}$ is a vector of controls for firm f at t_0; and $\tilde{\boldsymbol{Z}}_{f,t_0}$ is a subset of that vector, which is interacted with year interval fixed-effects χ_t. The second equation uses the vector notation $\Delta_k \boldsymbol{D}_{f,t} = [\Delta D_{f,t+k'}, \Delta D_{f,t+k'-1}, ..., \Delta D_{f,t}, ..., \Delta D_{f,t-k}]$ and $\boldsymbol{\alpha} = [\alpha_{-k'}, ..., \alpha_k]$. As we previously discussed, we include a sector indicator and the firm's prior export intensity (at t_0) in the subset $\tilde{\boldsymbol{Z}}_{f,t_0}$ of $\boldsymbol{Z}_{f,t_0}$, so those are also interacted with the year dummies.

Our specification in first-difference eliminates any bias that would be generated by a correlation between non time-varying firm characteristics (likely to affect current and future innovation) and the *level* of the demand shock shock $D_{f,t}$.[15] We additionally want to control for a potential correlation between those firm characteristics and future *changes* in the demand shock $\Delta D_{f,t}$. Following Blundell et al. (1999) and Blundell et al. (2002), we use a control function approach based on firm performance variables measured at t_0. We use the levels and growth rates of sales and employment as controls, which we include in the vector $\boldsymbol{Z}_{f,t_0}$. In addition, we include controls for the firm's past and current rate of innovation at t_0 whenever we use an innovation measure as the dependent outcome. We describe the functional form for those additional controls in more detail in the following section. We note that this type of correlation between changes in the demand shock $\Delta D_{f,t}$ and firm characteristics is substantially less likely than a correlation with the level of the demand shock $D_{f,t}$. We have checked that there is indeed a strong correlation between that demand shock in levels and the firm characteristics in our control function

[15] As discussed in Borusyak et al. (2021), this would require a firm fixed-effect control for a specification in levels.

(better performing firms tend to export to destinations with higher levels of demand). However, there is no correlation between those variables and changes in demand $\Delta D_{f,t}$.

Lastly, Borusyak et al. (2021) and Goldsmith-Pinkham et al. (2020) point out that even when such a correlation between firm characteristics and future demand shocks remains, the induced bias disappears as the number of shocks (our combination of destination-product pairs) grows large. Additionally, we discuss in Appendix E the validity of our identification following Borusyak et al. (2021).

4 Market Size and Innovation

We first show that our constructed export demand shock has a strong and contemporaneous impact on a firm's market size. We thus run our estimating equation (2) using the growth rate of sales and employment as our outcome variable $\Delta Y_{f,t}$ on the left-hand-side. We compute the average growth rate $\Delta Y_{f,t} = (Y_{f,t} - Y_{f,t-1})/[.5(Y_{f,t} + Y_{f,t-1})]$ in the same way that we constructed the export demand shock $\Delta D_{f,t}$.[16] The results for our key estimated coefficients α_τ (large darker dot) and their confidence intervals (95% as bar and 99% as dots) are represented graphically in Figure 3 for $\tau = -4, \ldots, 5$. The α_τ coefficients for $\tau > 0$ represent a response of the outcome variable $\Delta Y_{f,t}$ to a demand shock $\Delta_{f,t-\tau}$ τ years earlier; and conversely the coefficients for $\tau < 0$ represent a response of the outcome variable to a demand shock $-\tau$ years later.[17] It clearly shows a strong and contemporaneous response in both sales and employment to the export demand shock. As one would expect, the contemporaneous ($\tau = 0$) employment elasticity is lower than the one for sales; but it nevertheless becomes strongly positive (and significant beyond

[16]Using this average growth rate computation is important for the trade shock in order to accommodate the substantial number of import flow changes to/from zero. It is inconsequential for our measurement of the growth rate of sales and employment: our results are nearly identical when we compute the growth rate using the log difference instead.

[17]From here on out, we set this timing window for the demand shock $\Delta D_{f,t}$ to 4 leads and 5 lags. We have experimented with longer and shorter windows; this does not qualitatively affect our results. See Figures C7 for a longer window and C8 for a semi-dynamic specification without pre-trends.

the 1% level) in the same time interval as the demand shock. This highlights that this shock induces "real" growth for the firm (and that the increase in sales is not just associated with higher prices). As is also expected given the sluggish nature of employment adjustments, the response is longer-lasting than the one for sales and still significant one year following the demand shock. None of the pre-trend coefficients ($\tau < 0$) are significant except for the response of sales one year prior to the demand shock. This is entirely explained by the reporting lag between the booking of an order (when it shows up in the firm's sales accounting data) and the delivery of the exported goods (when it shows up in the export customs data) – that can potentially occur in different calendar years.[18]

Figure 3: OLS: Average response to a demand shock

(a) Sales (b) Employment

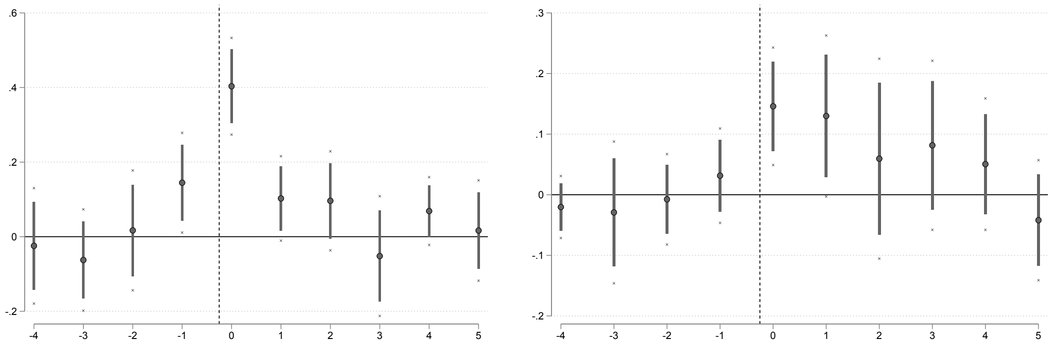

Notes: Estimates of coefficients α_τ for $\tau = -4 \ldots 5$ from equation (2) are reported graphically with the growth rate of sales (left-hand panel) and employment (right-hand panel) as the dependent variable. The x-axis represents the value of τ, the darker dots the point estimates of α_τ, the bar the 95% confidence intervals and the smaller dots the boundaries of the 99% confidence intervals. These estimations are obtained from an OLS regression with standard errors clustered at the NACE 2-digit sector level and robust to heteroskedasticity. Number of observations: 22,231. Time period for t: 2000-2008.

We now investigate how the firm's innovation responds to the same export demand shock using the same estimation strategy. We are left with a choice of functional form for a firm's patent response $\Delta Y_{f,t}$ between t and $t-1$. We do not think that the growth rate of a firm's full (over time) patent stock $P_{f,t}$ would be appropriate – because this puts too much weight on patents that may have been accumulated very far in the past and may

[18] In Appendix B, we use the monthly customs export data to show that this discrepancy is explained by shipments that arrive at the beginning of a new calendar year. It also mostly affects firms with volatile sales: the significant pre-trend coefficient disappears when we exclude those firms with sales growth rates above ±50%.

not be relevant for more recent patents (reflecting current innovation success). Instead of dividing the change in patent stock – new patents introduced between t and $t-1$ – by the average stock in those 2 periods (the Davis-Haltiwanger growth rate), we directly control for the average number of new patent introductions $\Delta P_{f,0}$ during our pre-sample time interval from 1985-1994 (prior to t_0). Given the very large dispersion across firms in new patents $\Delta P_{f,t}$ filed during year t, including the prevalence of zeros in many years (and for many firms, most years), we use the functional form $\log(1 + \Delta P_{f,t})$ with $\log(1 + \Delta P_{f,0})$ in our control vector $\boldsymbol{Z}_{f,t_0}$ for our OLS specification (2). We also address the zeros and over-dispersion in $\Delta P_{f,t}$ using a negative binomial specification where we can then use $\Delta P_{f,t}$ directly on the left-hand-side:[19]

$$\mathbb{E}_{\boldsymbol{Z}}\left[\Delta P_{f,t}\right] = \exp\left[\boldsymbol{\alpha} \cdot \Delta_k \boldsymbol{D}_{f,t} + \boldsymbol{\gamma} \cdot \boldsymbol{Z}_{f,t_0} + \tilde{\boldsymbol{\gamma}} \cdot \left(\tilde{\boldsymbol{Z}}_{f,t_0} \times \chi_t\right)\right], \tag{3}$$

where the expectation $\mathbb{E}_{\boldsymbol{Z}}$ is taken conditional on $\boldsymbol{Z}_{f,t}$ and on past and future values of $\Delta D_{f,t}$. We keep the same functional form $\log(1 + \Delta P_{f,0})$ in $\boldsymbol{Z}_{f,t_0}$ to control for the average rate of new patent introductions during our pre-sample years.[20] We choose a negative binomial (NB) specification as it is best suited (especially compared to Poisson) for the over-dispersion in the empirical distribution of new patents $\Delta P_{f,t}$, which standard deviation is 10.9, an order of magnitude higher than the 0.9 mean.

The graphical results for our OLS specification with the $\log(1 + \Delta P_{f,t})$ functional form are presented in Figure 4 with the innovation response $\Delta P_{f,t}$ measured both as new priority patents as well as our alternative measure based on citations received within five years. The graphical results for our negative binomial specification (3) are presented in

[19]Figure C2 and Tables C1 and C2 in the Appendix C also show that our results hold using the inverse hyperbolic sine transformation of the dependent variable: $\log\left(y + \sqrt{1+y^2}\right)$.

[20]This control is then defined for firms with zero new patents during some pre-sample years. We have also experimented with using $\log \Delta P_{f,0}$ directly in $\boldsymbol{Z}_{f,t}$ – hence a control for $\Delta P_{f,0}$ outside of the exponential in (3) – along with an indicator variable when $P_{f,0}$ is zero. This does not qualitatively affect our results. See Blundell et al. (1999) and Aghion et al. (2016) for a use of this type of control function in a similar specification.

Figure 4: OLS: AVERAGE RESPONSE TO A DEMAND SHOCK

(a) Priority patents (b) Citations within 5 years

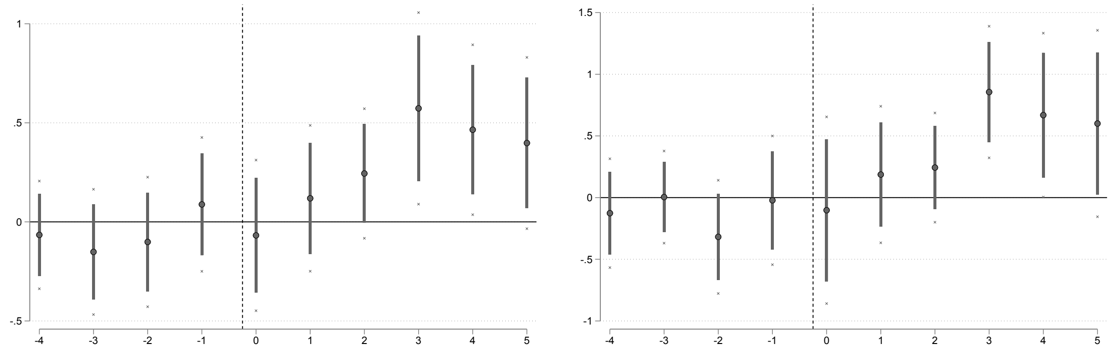

Notes: Estimates of coefficients α_τ for $\tau = -4 \ldots 5$ from equation (2) are reported graphically. The two panels differ in the dependent variable: the left-hand side panel considers the log of the number of new priority patents + 1 and the right-hand side panel considers the log of the number of accumulated citations received within 5 years + 1. The x-axis represents the value of τ, the darker dots the point estimates of α_τ, the bar the 95% confidence intervals and the smaller dots the boundaries of the 99% confidence intervals. These estimations are obtained from an OLS regression with standard errors clustered at the NACE 2-digit sector level and robust to heteroskedasticity. Number of observations: 22,237. Time period for t: 2000-2008.

Figure 5: NEGATIVE BINOMIAL: AVERAGE RESPONSE TO A DEMAND SHOCK

(a) Priority patents (b) Citations within 5 years

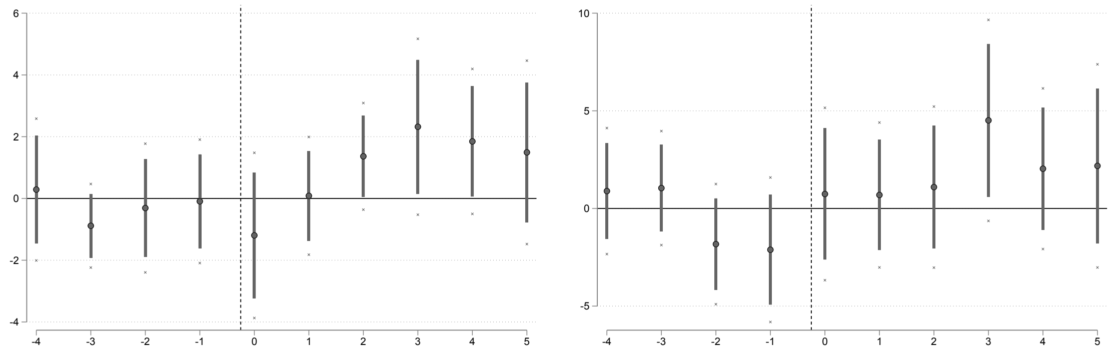

Notes: Estimates of coefficients α_τ for $\tau = -4 \ldots 5$ from equation (3) are reported graphically. The two panels differ in the dependent variable: the left-hand side panel considers the number of new priority patents and the right-hand side panel considers the number of accumulated citations received within 5 years. The x-axis represents the value of τ, the darker dots the point estimates of α_τ, the bar the 95% confidence intervals and the smaller dots the boundaries of the 99% confidence intervals. These estimations are obtained from a negative binomial regression with standard errors clustered at the NACE 2-digit sector level and robust to heteroskedasticity. Number of observations: 22,237. Time period for t: 2000-2008.

Figure 5 with the same two options for the innovation response $\Delta P_{f,t}$.

All four figures (across different functional form specifications and new patent measures) show a significant and sustained response of patenting activity starting 3 years after the export shock. The pre-trends are centered around zero and do not show any sign that the patenting activity precedes the change in export demand. We thus find a significant aggregate *market size* effect of export demand shocks on French firms' innovation. Since our specifications include sector-year fixed effects, this innovation response cannot be explained by any sector-wide innovation changes. Rather, it must be driven by the firm-level innovation responses to demand shocks.

Table 6 summarizes our results from Figures 3-5 for the response of both market size (scale) and innovation to the export demand shock. The dynamic leads and lags are cumulated (the coefficients are summed) into a pre-period (1 to 4 years prior to the shock), a current period (concurrent and 1 year after the shock), and a future period (2 to 5 years after the shock). Even when cumulated, there is no evidence of pre-trends for either scale (sales and employment) or innovation. Table 6 also highlights how the response of scale occurs concurrently with the shock while the response of innovation is delayed to the future period. This cumulative response is significant beyond the 1% level in our OLS specifications, and significant around the 5% level (a bit stronger for the patents; and weaker when measured as citations) in our negative binomial specification.[21]

The economic magnitude of those cumulated innovation responses are substantial.[22] On average, there are 3,116 firms (the innovators in our sample) operating in the future period 2-5 years following a demand shock in 1999, 2000, ..., 2003. Those firms introduced 7,637 priority patents (on average, in that same future period), which generated 25,346 citations. The future period coefficients for innovation in Table 6 imply that a 1 point export demand shock would induce 62 (OLS) - 171 (NB) new priority patents associated with 98 (OLS) - 1,149 (NB) citations during that same future period (again, on average for demand shocks in 1999, 2000, ... , 2003). This represents an aggregate (macro) elasticity of .81-2.2 for patents to an aggregate export demand shock; and an elasticity

[21] As the discussion of the economic magnitudes below makes clear, this is due to very large but imprecisely estimated coefficients in the negative binomial specification. The results from Tables 6 and 7, which use 2-digit NACE controls, are robust to using 5-digit NACE controls.

[22] The demand shock has a mean of 0.011 with a standard deviation of 0.043. The 25^{th} percentile is -0.002 and the 75^{th} is 0.025. Its yearly mean changes in line with international trade. A one point change in the export demand shock corresponds to a fourth of a standard deviation and is a small move in the distribution (moving from the bottom 10% to the top 90% corresponds to an increase of 9pp).

of .39-4.5 in terms of citations. The economic magnitude of that innovation response to demand shocks in export markets is therefore substantial.[23]

Table 6: CUMULATIVE RESPONSE TO DEMAND SHOCK

	Scale		Innovation			
	Sales	Employment	Priority Patents		Citations	
	OLS	OLS	OLS	NB	OLS	NB
Pre-Trend	0.074 (0.132)	-0.026 (0.063)	-0.230 (0.339)	-1.000 (1.599)	-0.462 (0.478)	-2.005 (2.729)
Current	0.506*** (0.072)	0.276*** (0.064)	0.051 (0.253)	-1.111 (1.566)	0.085 (0.469)	1.431 (2.851)
Future	0.129 (0.092)	0.149 (0.103)	1.680*** (0.532)	7.026** (3.117)	2.368*** (0.738)	9.819* (5.300)

Notes: This table reports point estimate and standard errors (under parentheses) for different linear combinations of coefficients from various estimations of equations (2) and (3). Pre-Trend corresponds to the estimate of $\alpha_{-4} + \alpha_{-3} + \alpha_{-2} + \alpha_{-1}$, current to $\alpha_0 + \alpha_1$ and Future to $\alpha_2 + \alpha_3 + \alpha_4 + \alpha_5$. Column 1 corresponds to the results displayed in Figure 3a, column 2 to Figure 3b, column 3 to Figure 4a, column 4 to Figure 5a, column 5 to Figure 4b and column 6 to Figure 5b. ***, ** and * indicate p-value below 0.01, 0.05 and 0.1 respectively.

5 Heterogeneous Impact: Distance to Frontier

We now investigate whether this innovation response varies across firms based on their distance to their sector's frontier. We use labor productivity (value-added per worker) as our metric for this distance. Just as we did with the firm-level export shares, we use the initial year t_0 to generate a distance measure that does not subsequently vary over time $t > t_0$. We partition firms into those with productivity above their 2-digit sector median (in year t_0), $a_{f,t_0} \geq \overline{a}_{t_0}$ (represented by indicator dummy $\mathbf{1}_a^+$), and those with productivity below the sector median, $a_{f,t_0} < \overline{a}_{t_0}$ (represented by indicator dummy $\mathbf{1}_a^-$). More specifically, we consider the following regression equation:

$$\Delta Y_{f,t} = \boldsymbol{\alpha}_H \cdot \left(\Delta_k \boldsymbol{D}_{f,t} \times \mathbf{1}_a^+\right) + \boldsymbol{\alpha}_L \cdot \left(\Delta_k \boldsymbol{D}_{f,t} \times \mathbf{1}_a^-\right) + \boldsymbol{\gamma} \cdot \boldsymbol{Z}_{f,t_0} + \tilde{\boldsymbol{\gamma}} \cdot \left(\tilde{\boldsymbol{Z}}_{f,t_0} \times \chi_t\right) + \varepsilon_{f,t}. \quad (4)$$

Since the firm's initial productivity level a_{f,t_0} is now used to construct our two different trade shocks on the right-hand-side, we add that variable to the control vectors $\boldsymbol{Z}_{f,t_0}$ and

[23] We have chosen throughout to report the magnitudes of the innovation responses in terms of the export demand shock. We could alternatively consider an instrumental variable specification in order to report those innovation magnitudes in terms of a shock to scale (market size or employment), using our scale regression as a first stage. Our innovation regressions can be viewed as the reduced form for that instrumental variable specification. Since the magnitude of those reduced form coefficients have a natural and direct interpretation, we stick to this specification.

$\tilde{\boldsymbol{Z}}_{f,t_0}$. We use the same functional form $\Delta Y_{f,t} = \log(1 + \Delta P_{f,t})$ for our OLS specification (adding $\log(1 + \Delta P_{f,0})$ to our control vector $\boldsymbol{Z}_{f,t_0}$). And we also estimate a negative binomial specification with the 'untransformed' new patent measure $\Delta P_{f,t}$ on the left-hand side, along with a control for $\Delta P_{f,0}$ in $\boldsymbol{Z}_{f,t_0}$:

$$\mathbb{E}_{\boldsymbol{Z}}\left[\Delta P_{f,t}\right] = \exp\left[\boldsymbol{\alpha}_H \cdot (\Delta_k \boldsymbol{D}_{f,t} \times \mathbf{1}_a^+) + \boldsymbol{\alpha}_L \cdot (\Delta_k \boldsymbol{D}_{f,t} \times \mathbf{1}_a^-) + \boldsymbol{\gamma} \cdot \boldsymbol{Z}_{f,t_0} + \tilde{\boldsymbol{\gamma}} \cdot \left(\tilde{\boldsymbol{Z}}_{f,t_0} \times \chi_t\right)\right], \quad (5)$$

where the expectation $\mathbb{E}_{\boldsymbol{Z}}$ is again taken conditional on $\boldsymbol{Z}_{f,t}$ and on past and future values of $\Delta D_{f,t}$.

The graphical results for both our OLS and negative binomial specifications are presented in Figures 6 and 7, once again using both priority patents and the accumulated citations as our measure of new patent activity $\Delta P_{f,t}$. All four figures show a significant and sustained response of patenting activity starting 3 years after the export shock – but *only* for firms that are initially closer to their sector's frontier (with labor productivity above the median level). In Appendix C, we return to the full battery of robustness checks that we previously described for the analysis of the non-heterogeneous responses. The main messages from Figures 6 and 7 remain unchanged (See Figures C1-C8). We also provide the estimates of all the coefficients from the regressions underlying Figures 4a, 5a, 6a and 7a (Table C3).

Could this heterogeneous response simply reflect the fact that the demand shock only affects the most productive firms? To check that this is not the case, we replicate the results shown in Figures 3a and 3b: that is, we allow for a different impact of the export demand shocks on sales or employment depending upon initial productivity levels. Looking at Figures 8a and 8b, we see that in contrast to what we observe for innovation, there is no heterogeneous response of sales or employment to a demand shock for low versus high productivity firms.[24] The responses for both sets of firms match the magnitudes of

[24] This instantaneous response of sales and employment to the demand shock is the short-term adaptation of firms to an increase in demand: they produce more to satisfy this increase in demand, and this response is homogeneous. However, following the conclusions of the model in section 6, this increased market size will trigger new entry

the average response that we previously documented.[25]

We summarize once again our dynamic results in Figures 6-8 by cumulating the coefficients into pre-trend, current, and future periods just as we previously reported in Table 6 for the case without the heterogeneous impact by productivity: pre-trend for 1-4 years prior to the export demand shock; current for 0-1 year following the shock; and future for 2-5 years following the shock. Those coefficients are reported in Table 7 for the above and below median firm productivity groups. In addition, we now report a significance test for their difference across those two groups. Appendix D describes falsification tests that highlight that those heterogeneous innovation responses cannot be explained by a firm-level trend or a firm's prior exposure to export markets.

Figure 6: OLS: Heterogeneous response to a demand shock

(a) Priority Patents (b) Citations within 5 years

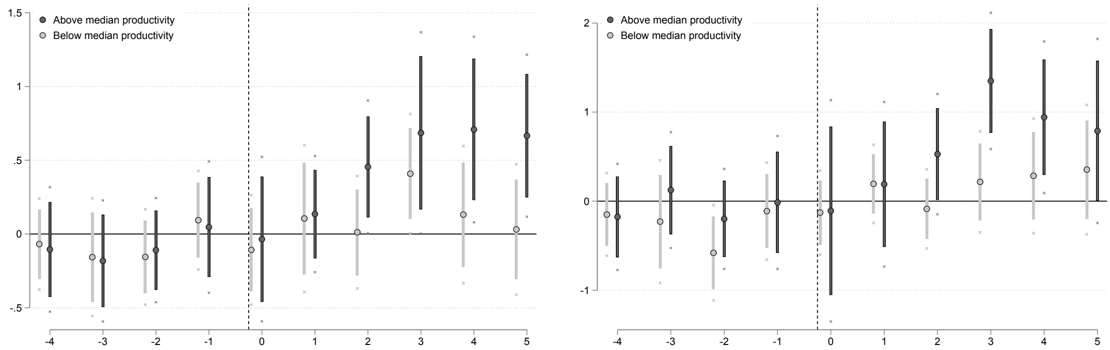

Notes: Estimates of coefficients $\alpha_{H,\tau}$ and $\alpha_{L,\tau}$ for $\tau = -4 \ldots 5$ from equation (4) are presented graphically. The two panels differ in the dependent variable: the left-hand side panel considers the log of the number of new priority patents + 1 and the right-hand side panel considers the log of the number of accumulated citations received within 5 years + 1. The x-axis represents the value of τ, the darker dots the point estimates of α_τ, the bar the 95% confidence intervals and the smaller dots the boundaries of the 99% confidence intervals. These estimations are obtained from an OLS regression with standard errors clustered at the NACE 2-digit sector level and robust to heteroskedasticity. Number of observations: 22,237. Time period for t: 2000-2008.

As can also be seen in Figures 6b and 7b, the response of innovation in terms of

in the long-run, with an impact on sales and employment that varies with productivity. We do not have a long-enough time-span to capture that long-run response. However, since innovation is a forward-looking variable (decided based on expected future profits), it captures those longer-run expectations.

[25]As can be seen in Figure 8a, the growth rate of the sales response for the below median firms fluctuates up and down following the trade shock. This effect is driven by firms with volatile sales: it disappears when we exclude those firms with sales growth rates above ±50%.

Figure 7: Negative Binomial: Heterogeneous response to a demand shock

(a) Priority Patents (b) Citations within 5 years

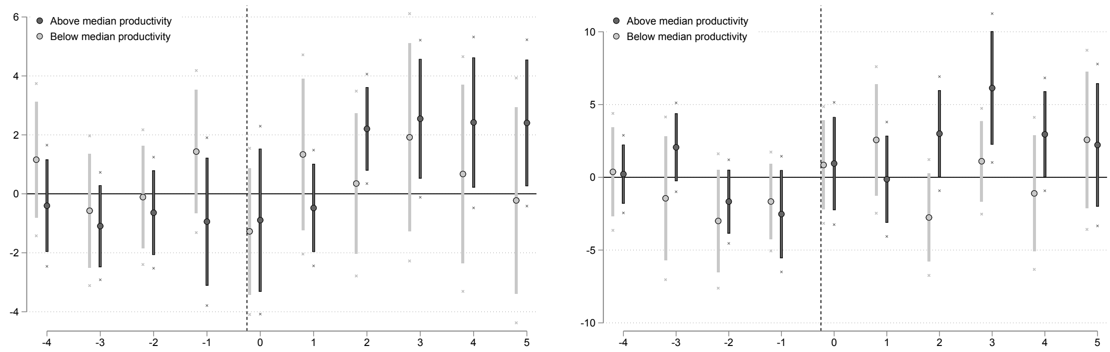

Notes: Estimates of coefficients $\alpha_{H,\tau}$ and $\alpha_{L,\tau}$ for $\tau = -4 \dots 5$ from equation (5) are presented graphically. The two panels differ in the dependent variable: the left-hand side panel considers the number of new priority patents and the right-hand side panel considers the number of accumulated citations received within 5 years. The x-axis represents the value of τ, the darker dots the point estimates of α_{τ}, the bar the 95% confidence intervals and the smaller dots the boundaries of the 99% confidence intervals. These estimations are obtained from a negative binomial regression with standard errors clustered at the NACE 2-digit sector times productivity group level and robust to heteroskedasticity. Number of observations: 22,237. Time period for t: 2000-2008.

Figure 8: Heterogeneous response to a demand shock - sales and employment

(a) Sales (b) Employment

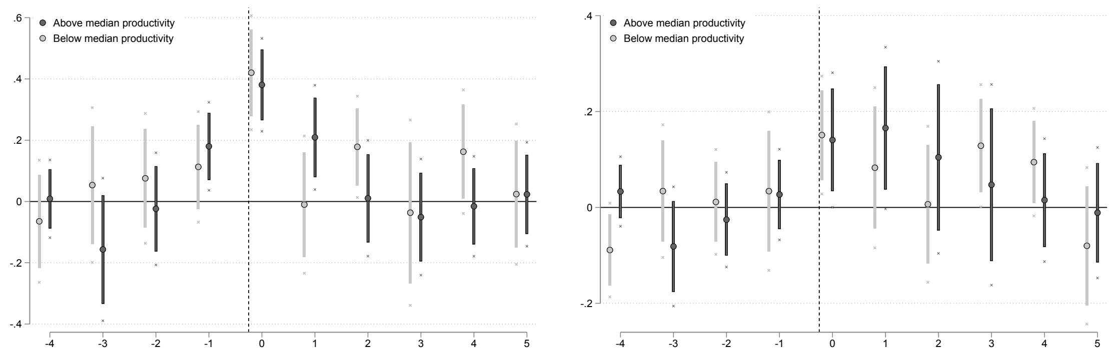

Notes: This Figure replicates Figure 3 but allowing for heterogeneity based on the initial productivity level as described in equation (4). Number of observations: 22,231. Time period for t: 2000-2008.

citations is negative in the pre-trend period for below-median productivity firms. There is no evidence of pre-trends in any of the other dependent variables for either scale of innovation among either subset of more or less productive firms. Table 7 also highlights the strong contemporaneous response of the scale variables to the demand shock that we previously emphasized. Importantly, there is no evidence of any significant differences in the responses of those scale variables across the two productivity groups. On the other hand, Table 7 makes clear that the strong future response of innovation for the above-

Table 7: Cumulative Heterogeneous Response to Demand Shock

	Scale		Innovation			
	Sales	Employment	Priority Patents		Citations	
	OLS	OLS	OLS	NB	OLS	NB
Pre-Trend						
Below Median	0.177 (0.153)	-0.009 (0.082)	-0.288 (0.339)	1.907 (2.496)	-1.071** (0.499)	-5.730 (4.510)
Above Median	0.009 (0.173)	-0.047 (0.087)	-0.350 (0.394)	-3.082* (1.641)	-0.265 (0.709)	-1.910 (2.926)
Difference	-0.169 (0.201)	-0.037 (0.084)	-0.062 (0.435)	-4.988 (3.225)	0.806 (0.714)	3.819 (5.337)
Current						
Below Median	0.410*** (0.108)	0.234*** (0.077)	-0.003 (0.249)	0.063 (1.460)	0.066 (0.296)	3.416 (2.684)
Above Median	0.590*** (00.079)	0.306*** (00.101)	0.1 (0.297)	-1.372 (1.768)	0.081 (0.778)	0.818 (2.833)
Difference	0.180 (0.116)	0.073 (0.118)	0.103 (0.272)	-1.435 (1.455)	0.015 (0.705)	-2.598 (2.822)
Future						
Below Median	0.328** (0.143)	0.150 (0.093)	0.583* (0.336)	2.717 (3.113)	0.769 (0.674)	-0.184 (4.148)
Above Median	-0.032 (0.155)	0.156 (0.136)	2.515*** (0.718)	9.582*** (3.066)	3.609*** (0.909)	14.312*** (5.241)
Difference	-0.361* (0.201)	0.006 (0.151)	1.932*** (0.591)	6.865** (3.123)	2.841*** (0.967)	14.496*** (4.832)

Notes: This table reports point estimate and standard errors (under parentheses) for different linear combinations of coefficients from various estimations of equations (4) and (5). Pre-Trend corresponds to the estimate of $\alpha_{X,-4}+\alpha_{X,-3}+\alpha_{X,-2}+\alpha_{X,-1}$, current to $\alpha_{X,0}+\alpha_{X,1}$ and Future to $\alpha_{X,2}+\alpha_{X,3}+\alpha_{X,4}+\alpha_{X,5}$ where $X=H$ for lines "Above Median" and $X=L$ for lines "Below Median". The lines "Difference" corresponds to the difference between the corresponding above and below median linear combinations. Column 1 corresponds to the results displayed in Figure 8a, column 2 to Figure 8b, column 3 to Figure 6a, column 4 to Figure 7a, column 5 to Figure 6b and column 6 to Figure 7b. ***, ** and * indicate p-value below 0.01, 0.05 and 0.1 respectively.

median firms is statistically distinguishable from the response for the below-median firms. Not only are the future response coefficients for those below-median firms insignificant, the coefficients difference in favor of the relatively more productive firms is statistically significant beyond the 1% level in all of our specifications. This strongly supports our main finding that the innovation response is entirely concentrated within the subset of relatively more productive firms.

The economic magnitude of the innovation response for those above-median firms corresponds roughly to a similar aggregate response as the one we reported for the case without firm heterogeneity – except that this response is now concentrated more intensely and exclusively within the top-half of relatively more productive firms. The OLS coefficients in Table 7 imply that an aggregate 1 point increase in the export demand shock would induce 63 patents associated with 107 citations amongst the above-median firms. Those numbers are similar than the 63 patents and 98 citations we previously recorded for the aggregate response without firm heterogeneity. The NB coefficient for the patent response implies a higher number of patents from the 1 point increase in the export demand shock: 245 patents relative to 171 for the case without firm heterogeneity. The NB coefficient for the citation response implies a substantially larger response relative to the case without heterogeneity (1,624 versus 1,149). But both NB coefficients for citations (and especially the one for the above-median firms) have a large standard error, so there is still a wide overlap between our predictions for the aggregate response in terms of citations with and without firm heterogeneity.

Once again, we find that the economic magnitude of the innovation response – concentrated within the subset of relatively more productive firms – is substantial.

6 A model

6.1 Presentation

In this section, we show that our main finding of a skewed innovation response to common demand shocks arises naturally from a model of endogenous innovation and competition with firm heterogeneity. Our model features a "standard" market size effect that increases innovation for all firms. But it also embodies an endogenous competition effect that discourages innovation by low productivity firms. This skewed induced competition effect captures the idea that the expanded market for exports will attract new firms into the export market as more firms find it profitable to sell their products there; this in turn will raise competition for exporters into that market. Due to the nature of competition between firms – featuring endogenous markups – this effect gradually dissipates as productivity (and resulting market share) increases. This competition effect is thus more salient for smaller French firms with initially lower productivity, as they lose market share to larger more productive firms.

The model develops a very simple (and specific) functional form for simplicity, but our result that the increase in market size triggers a skewed competition effect extends to much more general specifications, in particular to the broad class of preferences under monopolistic competition that satisfy Marshall's Second Law of Demand (MSLD), i.e. lead to residual (firm-level) inverse demands that become more inelastic as consumption increases. Instead, a model with monopolistic competition and CES preferences (and hence exogenous markups) would not generate a skewed induced competition effect of increased market size. The recent empirical trade literature provides mounting evidence for the relevance of endogenous markups associated with MSLD demand.[26]

Finally, we stress that our empirical work and results in the previous sections are not meant to specifically test whether the heterogeneous impact of increased market size on innovation is due to the skewed competition effect with endogenous markups that we model in this section. We are just showing that this evidence is consistent with – and easily explained by – a competition channel highlighted by our model. Our model also illustrates the fact that very few assumptions are needed beyond MSLD demand to generate a skewed innovation response to increased market size.

6.2 Basic setup

French firms exporting to some export market destination D are competing with local firms producing in D. We let L denote the number of consumers in that destination. This indexes market size. These consumers have preferences over all varieties available in D. There is a continuum of differentiated varieties indexed by $i \in [0, M]$, where M is the measure of available products. Suppose that the demand for variety q_i is generated by a representative consumer in country D with additively separable preferences with sub-utility:[27]

$$u(q_i) = \alpha q_i - \frac{\beta q_i^2}{2}, \qquad \alpha > 0, \quad \beta > 0.$$

Those preferences do not differentiate between French or locally produced varieties. Thus, the output, profit and revenues for the French exporters and local producers have the same expression. For simplicity, we assume that both types of firms have access to

[26] See Aghion et al. (2018) for a more general model and Melitz (2018) for a summary of this evidence and how it is connected to endogenous markups and MSLD demand. This evidence for endogenous markups adjustments would also be consistent with oligopoly models where the elasticity of substitution between products remains constant. Such a model would nevertheless feature endogenous price elasticities that respond in a very similar way to those in a model of monopolistic competition with MSLD demand.

[27] As we previously discussed, our analysis can be extended to a broader class of preferences that satisfy Marshall's Second Law of Demand (such that residual demand becomes more inelastic as consumption increases).

the same innovation technology, which leads to similar innovation decisions.

6.2.1 Consumer optimization

This representative consumer facing prices p_i solves:

$$\max_{q_i \geq 0} \int_0^M u(q_i)di \;\; \text{s.t.} \;\; \int_0^M p_i q_i di = 1.$$

This yields the inverse residual demand function (per consumer):

$$p(q_i) = \frac{u'(q_i)}{\lambda} = \frac{\alpha - \beta q_i}{\lambda}, \tag{6}$$

where $\lambda = \displaystyle\int_0^M u'(q_i)q_i di > 0$ is the corresponding Lagrange multiplier, also equal to the marginal utility of income. Given the assumption of separable preferences, this marginal utility of income λ is the unique endogenous aggregate demand shifter. Higher λ shifts all residual demand curves downwards; we thus interpret this as an increase in competition for a given exogenous level of market size L.

6.2.2 Firm optimization

Consider a (French or domestic) firm with marginal cost c facing competition λ. This firm chooses the output per consumer $q(c;\lambda)$ to maximize operating profits $L\left[p(q)q - cq\right]$. The corresponding first order condition yields

$$q(c;\lambda) = \frac{\alpha - c\lambda}{2\beta}, \tag{7}$$

so long as the firm's cost is below α/λ; the remaining firms with higher cost do not produce. This output choice in turn leads to the maximized profit per consumer

$$\pi(c;\lambda) = \frac{(\alpha - c\lambda)^2}{4\beta\lambda}.$$

In particular, we see that both output and profit are decreasing in both firm level cost

c and the endogenous competition measure λ. More productive firms (with lower cost c) are larger and earn higher profits than their less productive counterparts; and an increase in competition λ lowers production levels and profits for all firms.

6.2.3 Innovation choice

A firm is characterized by its baseline cost $\tilde{c}$. It can reduce its marginal cost of production c below its baseline cost by investing in innovation. More formally, we assume that

$$c = \tilde{c} - \varepsilon k,$$

where k is the firm's investment in innovation and $\varepsilon > 0$; and we assume that the cost of innovation is quadratic in k, equal to $c_I k + \frac{1}{2} c_{I2} k^2$.[28]

Thus a firm with baseline cost $\tilde{c}$ will choose its optimal R&D investment $k(\tilde{c}; \lambda)$ so as to maximize total profit:

$$\Pi(\tilde{c}, k; \lambda) = L\pi(\tilde{c} - \varepsilon k; \lambda) - c_I k - \frac{1}{2} c_{I2} k^2.$$

The optimal R&D investment $k(\tilde{c}; \lambda)$, if positive, satisfies the first order condition:

$$\varepsilon Q(\tilde{c}, k; \lambda) = c_{I2} k + c_I, \tag{FOC}$$

where

$$Q(\tilde{c}, k; \lambda) \equiv Lq(\tilde{c} - \varepsilon k; \lambda) = L[\alpha - (\tilde{c} - \varepsilon k)\lambda]/2\beta$$

is the total firm output (across consumers) produced by a firm with baseline cost $\tilde{c}$ and innovation k. We assume that the baseline cost $\tilde{c}$ is bounded below by $\tilde{c}_{\min}$ such that

[28] Since we only consider a single sale destination D for our firms, we are implicitly assuming that the innovation is directed at the delivered cost to consumers in D. We should thus think of innovation as specific to the appeal/cost trade-off to consumers in D. Our companion paper describes how our main skewness result holds for more general functional forms for the cost and return to innovation.

$\tilde{c}_{\min} - \varepsilon k(\tilde{c}_{\min}; \lambda) = 0$, or equivalently

$$\tilde{c}_{\min} = \frac{\varepsilon}{c_{I2}} \left(\frac{\varepsilon L \alpha}{2\beta} - c_I \right).$$

This in turn ensures that the post-innovation marginal cost is bounded away from zero, even for the most productive firms.

Figure 9a depicts the optimal innovation choice at the intersection between the marginal cost (MC, right-hand side of FOC) and the marginal benefit of innovation (MB, left-hand side of FOC). As long as the marginal benefit is above the marginal cost of investing in R&D, the firm wants to increase innovation, because the marginal profit made by investing one more unit of R&D exceeds its marginal cost. We assume that the second order condition holds so that the slope of the marginal cost is strictly larger than the slope of the marginal gain:

$$c_{I2} > \varepsilon \frac{\partial Q}{\partial k} = \frac{\varepsilon^2 \lambda L}{2\beta}. \tag{SOC}$$

This ensures a smooth innovation response to productivity differences.

When comparing a more productive firm (with lower baseline cost, depicted by the blue curve) and a less productive firm (with higher baseline cost, depicted by the red curve), we see that both firms face the same marginal cost curve and their marginal gain curves have the same slope. Only the zero intercepts of the two marginal gain curves are different: the lower $\tilde{c}$ firms have a higher intercept, thus a higher marginal gain, and therefore invest more in R&D. Firms with sufficiently high baseline costs do not innovate, as the zero intercept of their marginal gain curves falls below c_I, so that even their first innovation unit would not be worth its cost. These are firms with baseline costs above the baseline cost of the marginal innovator, which is equal to:

$$\widehat{C}_I = \frac{1}{\lambda} \left(\alpha - \frac{2\beta c_I}{\varepsilon L} \right). \tag{8}$$

In the next subsection we analyze how the optimal innovation choice $k(\tilde{c}; \lambda)$ responds to a positive demand shock, i.e. to an increase in market size L.

6.3 The market size and competition effects

We first analyze the direct effect of an increase in L, holding the competition level λ constant. At each firm's current innovation choice $k(\tilde{c};\lambda)$, this triggers a proportional increase in firm output, and an upward shift in the marginal benefit of innovation, inducing all firms to increase innovation.

Figure 9b shows this innovation response for firms with different baseline costs. Both the intercept and the slope of the marginal gain curve increase. We see how this leads to higher innovation for all firms. Given our assumptions on the benefits and costs of innovations, this leads to higher innovation responses for more productive firms:

$$\frac{\partial^2 k}{\partial L \partial \tilde{c}} < 0.$$

This increase in market size also induces some firms to begin R&D (higher $\widehat{C}_I$, see (8)).

We now consider the effect of an increase in competition λ, holding market size L constant. At each firm's current innovation choice $k(\tilde{c};\lambda)$, this triggers a decrease in firm output (see equation (7)). However, unlike the case of a change in market size L, this output response is no longer proportional across firms: high cost firms bear the brunt of the competition increase and disproportionately lose market share. Even though all firms respond by reducing innovation, this reduction in innovation is most pronounced (larger) for those high cost firms:

$$\frac{\partial^2 k}{\partial \lambda \partial \tilde{c}} < 0.$$

This contrasts with the case of a market size decrease (leading to proportional output decreases), which would lead to bigger innovation reductions for low cost firms instead. In the limit for the most efficient firms (with baseline cost approaching $\tilde{c}_{\min}$), the negative impact of increased competition on innovation dissipates completely (see (FOC)).

Figure 9c shows this innovation response for firms with different baseline costs. The increase in competition decreases the marginal benefit of innovation, but substantially more for the high cost firm – because the intercept decrease is larger (recall that the slope of the marginal benefit curve does not change with the firm's baseline cost).[29] Thus, the high cost firm's reduction in innovation is most pronounced. The competition increase also induces some firms to stop R&D (lower $\widehat{C}_I$, see (8)).

6.4 The heterogeneous innovation response to an export shock

How can our model generate the skewness we observed in firms' innovation response to a positive export demand shock? In Appendix F we endogenize the equilibrium competition level λ in country D and we show that it increases with L. The intuition is that an increase in market size L induces entry on the export market D by new firms; this in turn increases the elasticity of the inverse demand curve faced by each French exporter to D and an increase in λ. It then follows that an increase in market size L will have two effects on firms' innovation incentives: (a) a *direct* - positive - *market size effect*, whereby the increase in L induces all firms to increase innovation; this effect was shown above to be more positive for more frontier firms (i.e. for firms with lower initial production cost $\tilde{c}$); (b) an *induced* - negative -*competition effect* whereby the increase in L increases competition λ which in turns reduces firms' innovation incentives; as we saw above, the effect of an increase in λ on firms' innovation is more negative for less productive firms (i.e. for firms with higher initial production cost $\tilde{c}$). The overall effect of an increase in market size L on innovation – which combines the direct market size effect and the induced competition effect – will be unambiguously more positive for more frontier firms; moreover, this overall effect can turn out to be negative for the least productive firms - depending on the relative magnitude of the direct and indirect impacts. This heterogeneous response is fully consistent with our empirical analysis: we showed that the most productive half of the firms increase their innovation when their market size expands, while the response for the least productive half of the firms is essentially muted.

[29] The new dotted marginal benefit curve remains below the old one at least until it meets the marginal cost curve, even though an increase in competition increases the slope of the marginal benefit curve.

Figure 9: Theoretical relationships

(a) Optimal innovation

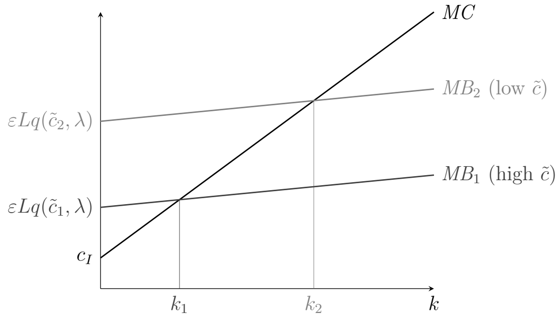

(b) Direct market size effect (increase in L)

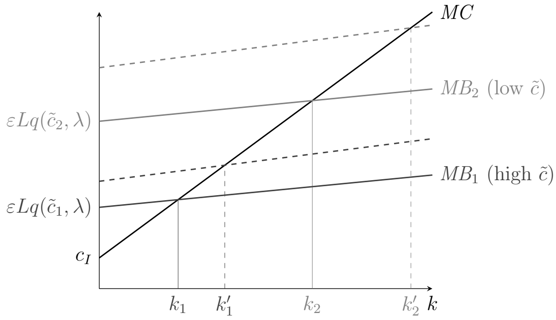

(c) Competition effect (increase in λ)

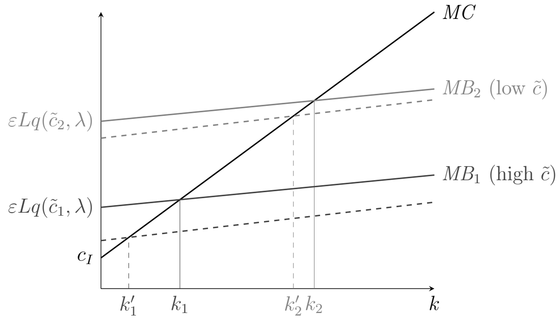

7 Conclusion

In this paper we used exhaustive data covering the French manufacturing sector to analyze the impact of export demand shocks on patenting by French exporting firms. To disentangle the direction of causality between export demand and innovation, we constructed a firm-level export demand shock which responds to aggregate conditions in a firm's export destinations but is exogenous to firm-level decisions.

We first showed that French firms respond to exogenous growth shocks in their export destinations by patenting more. Second, we showed that this positive impact of market size on innovation is skewed and entirely driven by French firms with above-median initial labor productivity within their sector. Third, we showed that the innovation response arises 2 to 5 years after a demand shock, whereas the same demand shock raises contemporaneous sales and employment for all firms. And lastly, we developed a simple theoretical model with endogenous innovation and endogenous markups which rationalizes the skewed innovation response to increases in export demand.

Our paper contributes to the existing literature on innovation and market size in several respects: To our knowledge, we are the first to identify a causal impact of firm-level market size on innovation that is independent of any sector-level dynamics (controlling for arbitrary sector level year-on-year changes) and widespread across the entire manufacturing sector. Given the detailed timing of the changes in demand, we are also able to precisely measure the time-lag required before the ensuing patenting activity is recorded. And lastly, we have showed that this innovation response is highly skewed and dominated by relatively more productive firms within each sector.

Acknowledgment: A previous version of this paper was circulated under the title: "The Impact of Export on Innovation: Theory and Evidence". The authors would like to thank, without implicating Daron Acemoglu, Ufuk Akcigit, Pol Antras, Kirill Borusyak, Elhanan Helpman, Giammario Impullitti, Thierry Mayer, Isabelle Mejean, Torsten Persson, Stephen Redding, Thomas Sampson, Dan Trefler, John Van Reenen, Daniel Xu, and participants at the 2017 AEA meeting, the NBER Summer Institute 2017, CIFAR and numerous seminars for useful discussions and advises. We are grateful to Connie Lee for excellent research assistance. This work is supported by a public grant overseen by the French National Research Agency (ANR) as part of the "Investissements d'Avenir" program (reference: ANR-10-EQPX-17 - Centre d'accès securisé aux données – CASD), and has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 786587). Philippe Aghion and Matthieu Lequien thank the support of the EUR grant ANR-17-EURE-0001.

References

Acemoglu, Daron and Joshua Linn, "Market size in innovation: theory and evidence from the pharmaceutical industry," *The Quarterly journal of economics*, 2004, *119* (3), 1049–1090.

Adao, Rodrigo, Michal Kolesár, and Eduardo Morales, "Shift-share designs: Theory and inference," *The Quarterly Journal of Economics*, 2019, *134* (4), 1949–2010.

Aghion, Philippe and Peter Howitt, *The economics of growth*, Cambridge, Massachusetts. London, England: MIT Press, 2009. OCLC: 610950320.

_ , **Antoine Dechezleprêtre, David Hemous, Ralf Martin, and John Van Reenen**, "Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry," *Journal of Political Economy*, 2016, *124* (1), 1–51.

— , **Antonin Bergeaud, Matthieu Lequien, and Marc J Melitz**, "The Impact of Exports on Innovation: Theory and Evidence," Working Paper 24600, National Bureau of Economic Research May 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.

Akcigit, Ufuk, Sina T. Ates, and Giammario Impullitti, "Innovation and Trade Policy in a Globalized World," NBER Working Paper 24543, National Bureau of Economic Research 2018.

Autor, David, David Dorn, Gordon Hanson, Gary Pisano, and Pian Shu, "Foreign competition and domestic innovation: Evidence from US patents," *American Economic Review: Insights*, 2020, *2* (3), 357–74.

— , — , **Lawrence F Katz, Christina Patterson, and John Van Reenen**, "The fall of the labor share and the rise of superstar firms," *The Quarterly Journal of Economics*, 2020, *135* (2), 645–709.

Autor, David H., David Dorn, and Gordon H. Hanson, "The China syndrome: Local labor market effects of import competition in the United States," *American Economic Review*, 2013, *103* (6), 2121–68.

Aw, Bee Yan, Mark J Roberts, and Daniel Yi Xu, "R&D investment, exporting, and productivity dynamics," *American Economic Review*, 2011, *101* (4), 1312–44.

Bartik, Timothy J., *Who Benefits from State and Local Economic Development Policies?* number wbsle. In 'Books from Upjohn Press.', W.E. Upjohn Institute for Employment Research, 1991.

Bernard, Andrew B, J Bradford Jensen, Stephen J Redding, and Peter K Schott, "Global firms," *Journal of Economic Literature*, 2018, *56* (2), 565–619.

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.

Blundell, Richard, Rachel Griffith, and Frank Windmeijer, "Individual effects and dynamics in count data models," *Journal of econometrics*, 2002, *108* (1), 113–131.

_ , _ , and John van Reenen, "Market Share, Market Value and Innovation in a Panel of British Manufacturing Firms," *Review of Economic Studies*, 1999, *66* (3), 529–554.

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

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "Quasi-Experimental Shift-Share Research Designs," *The Review of Economic Studies*, 06 2021. rdab030.

Bustos, Paula, "Trade Liberalization, Exports, and Technology Upgrading: Evidence on the Impact of MERCOSUR on Argentinian Firms," *American Economic Review*, 2011, *101* (1), 304–340.

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

Coelli, Federica, Andreas Moxnes, and Karen Helene Ulltveit-Moe, "Better, faster, stronger: Global innovation and trade liberalization," *Review of Economics and Statistics*, 2020, pp. 1–42. forthcoming.

Damijan, Jože P, Črt Kostevc, and Sašo Polanec, "From innovation to exporting or vice versa?," *World Economy*, 2010, *33* (3), 374–398.

de Rassenfosse, Gaetan, Helene Dernis, Dominique Guellec, Lucio Picci, and Bruno van Pottelsberghe de la Potterie, "The worldwide count of priority patents: A new indicator of inventive activity," *Research Policy*, 2013, *42* (3), 720–737.

Dhingra, Swati, "Trading Away Wide Brands for Cheap Brands," *American Economic Review*, October 2013, *103* (6), 2554–2584.

Gaulier, Guillaume and Soledad Zignago, "BACI: International Trade Database at the Product-Level. The 1994-2007 Version," Working Papers 2010-23, CEPII research center October 2010.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik instruments: What, when, why, and how," *American Economic Review*, 2020, *110* (8), 2586–2624.

Grossman, Gene M. and Elhanan Helpman, *Innovation and growth in the global economy*, MIT Press, 1991.

— **and** — , "Quality Ladders in the Theory of Growth," *The Review of Economic Studies*, 1991, *58* (1), 43.

Hall, Bronwyn, Adam Jaffe, and Manuel Trajtenberg, "Market Value and Patent Citations," *RAND Journal of Economics*, 2005, *36* (1), 16–38.

Hummels, David, Rasmus Jørgensen, Jakob Munch, and Chong Xiang, "The wage effects of offshoring: Evidence from Danish matched worker-firm data," *American Economic Review*, 2014, *104* (6), 1597–1629.

Iacovone, Leonardo, Wolfgang Keller, and Ferdinand Rauch, "Innovation Responses to Import Competition," 2011. Unpublished manuscript.

Impullitti, Giammario and Omar Licandro, "Trade, firm selection, and innovation: the competition channel," *The Economic Journal*, February 2018, *128*, 189–229.

Insee, "NAF REV. 2 et CPF REV. 2 : Guide d'utilisation," 2016.

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

Lileeva, Alla and Daniel Trefler, "Improved access to foreign markets raises plant-level productivity... for some plants," *The Quarterly journal of economics*, 2010, *125* (3), 1051–1099.

Malgouyres, Clément, Thierry Mayer, and Clément Mazet-Sonilhac, "Technology-induced Trade Shocks? Evidence from Broadband Expansion in France," CEPR Discussion Papers 13847, C.E.P.R. Discussion Papers 2019.

Mayer, Thierry and Gianmarco I. P. Ottaviano, "The Happy Few: The Internationalisation of European Firms," *Intereconomics - Review of European Economic Policy*, 2008, *43* (3), 135–148.

_ , Marc J. Melitz, and Gianmarco I. P. Ottaviano, "Market Size, Competition, and the Product Mix of Exporters," *American Economic Review*, 2014, *104* (2), 495–536.

_ , Marc J Melitz, and Gianmarco IP Ottaviano, "Product mix and firm productivity responses to trade competition," *Review of Economics and Statistics*, 2020, pp. 1–59. forthcoming.

Melitz, Marc J., "Competitive effects of trade: theory and measurement," *Review of World Economics*, February 2018, *154* (1), 1–13.

Morrison, Greg, Massimo Riccaboni, and Fabio Pammolli, "Disambiguation of patent inventors and assignees using high-resolution geolocation data," *Nature Scientific Data*, 2017, *4*.

ONLINE APPENDIX

A Data description

A.1 Patent data

Our first database is PATSTAT Spring 2016 which contains detailed information about patent applications from every patent office in the world. Each patent can be exactly dated to the day of application, which is sometimes referred to as the "filing date".

Counting patent applications Each French firm is associated with a number of patent applications by that firm each year (see section A.4). If the firm shares a patent with some other firms, then we only allocate a corresponding share of this patent to the firm. This raises the well-documented issue of truncation bias (Hall et al., 2005). Indeed as we come closer to the end of the sample, we observe a smaller fraction of all patents since many of them are not yet granted.[30] In addition, there is a legal obligation to wait 18 months before publication in PATSTAT. With our version of Spring 2016 this implies that we can assume the data to be reasonably complete up to 2012. The sector-time fixed effects also deal with the truncation bias in our regressions. An alternative solution could be to use the year of granting instead of the year of application. However, the former is less relevant than the latter as it is affected by administrative concerns and also by potential lobbying activities that have little to do with the innovation itself. In order to be as close to the time of the innovation as possible, we follow the literature and consider the filing date. We consider every patent owned by a French firm, regardless of the patent office that granted the patent rights. Here we need to be aware of the differences in regulations across intellectual property offices. Some patent offices, especially those of Japan and Korea, offer less breadth to a patent, which implies that more patents are needed to protect a given invention than in other patent offices (see de Rassenfosse et al., 2013). Since we only consider French firms, this would become an issue only if some French firms patent relatively more in countries like Japan or Korea, which would induce an upward bias in the number of patents held by those firms. However, we use a count of priority patent applications only, which is immune to this potential bias.

[30]The time between patent application and patent granting is a little more than 2 years on average but the distribution of this lag is very skewed with few patent applications still waiting for patent granting many years after the application.

Priority patent applications The fact that an inventor might want to patent its invention in different countries (or through supranational patent offices like PCT or EPO) makes it impossible to consider that one patent is equal to one invention. For this reason, patents are associated with a family which gather different patents which are more or less related to the same invention. More precisely, during a 12-month period following the filing of an application, the applicant has the *right of priority*. During this period, the applicant can file a similar patent in a different patent office and *claim the priority* of the first application when filing this subsequent application. If the priority claim is valid, the date of filing of the first application is considered to be the effective date of filing for the subsequent applications. This first application corresponds to the priority patent. All subsequent filings of the same intellectual property (in particular if they are in other countries) are secondary filings.

Citations We also use `PATSTAT` information on citations received by patents owned by French firms. Citations are often used to address the problem that all patents are not of equal quality and that simply counting the number of patent applications provides a noisy measure of the true innovation performance of a firm. However, the truncation bias issue is even worse with citations than with patent count. Patents from say 2010 have less time to be cited than patents from 1980 regardless of their respective qualities. Comparing different cohorts of patents can thus lead to misinterpreting what is reflected by the total number of citations received by a firm. To address this problem, we consider the number of citations received within a certain time window after the application date (usually 3 or 5 years). Using sector times year fixed effects in the regressions also helps to alleviate this concern.

A.2 Firm-level accounting data

Our second data source, provided by the DGFiP-Insee and called `FICUS` and `FARE`, provides us with accounting data for French firms. The corresponding data are drawn from compulsory reporting of firms and income statements to fiscal authorities in France. Since every firm needs to report every year to the tax authorities, the coverage of the data is all French firms from 1994 to 2012 with no limiting threshold in terms of firm size or sales. This dataset provides us with information on the turnover, employment, value-added, the four-digit sector the firm belongs to ... This corresponds to around 47 million observations and the number of observations per year increases from 1.9m to 3.9m over the period we consider.

The manufacturing sector is defined as category C of the first level of the NAF (*Nomenclature d'Activités Françise*), the first two digits of which are common to both NACE (Statistical Classification of Economic Activities in the European Community) and ISIC (International Standard Industrial Classification of All Economic Activities). Insee provides each firm with a detailed principal activity code (APE) with a top-down approach: it identifies the 1-digit section with the largest value added. Among this section, it identifies the 2-digit division with the largest value-added share, and so on until the most detailed 5-digit APE code (Insee, 2016). It is therefore possible that another 5-digit code shows a larger value-added share than the APE identified, but one can be sure that the manufacturing firms identified produce a larger value-added in the manufacturing section than in any other 1-digit section, which is precisely what we rely on to select the sample of most of our regressions. The 2-digit NAF sector, which we rely intensively on for our fixed effects, then represents the most important activity among the main section of the firm. Employment each year is measured on average within the year and may therefore be a non-integer number.

A unique 9-digit identifier called *Siren number* is associated to each firm, this number is given to the firm until it disappears and cannot be assigned to another firm in the future. When a firm merges with another firm, or is acquired by another firm, or makes significant changes in its organization, this number may change over time. Hence, new entrant *Sirens* in our database do not necessary correspond to new firms.

A.3 Trade data

Customs data for French firms Detailed data on French exports by product and country of destination for each French firm are provided by the French Customs. These are the same data as in Mayer et al. (2014) but extended to the whole 1994-2012 period. Every firm must report its exports by destination country and by very detailed product (at a level finer than HS6). However administrative simplifications for intra-EU trade have been implemented since the Single Market, so that when a firm annually exports inside the EU less than a given threshold, these intra-EU flows are not reported and therefore not in our dataset. The threshold stood at 250 000 francs in 1993, and has been periodically reevaluated (650 000 francs in 2001, 100 000 euros in 2002, 150 000 euros in 2006, 460 000 euros in 2011). Furthermore flows outside the EU both lower than 1 000 euros in value and 1 000 kg in weight are also excluded until 2009, but this exclusion was deleted in 2010.

Country-product bilateral trade flows CEPII's database `BACI`, based on the UN database `COMTRADE`, provides bilateral trade flows in value and quantity for each pair of countries from 1995 to 2015 at the HS6 product level, which covers more than 5,000 products.

A.4 Matching

Our paper is the first to merge those three very large - patent, administrative, and customs - datasets covering exporting French firms. Merging administrative firm-level data from FICUS/FARE and Customs data is fairly straightforward as a firm can be identified by its *Siren* identifier in both datasets.[31] Thus the main challenge is to match either of these two datasets with PATSTAT. Indeed, PATSTAT only reports the name of the patent applicant(s). Not only can this name be slightly different from the name reported in the other two databases, but it may also change over time, for example because of spelling mistakes. We thus relied on the work of Lequien et al. (2019) who developed a matching algorithm to map patents with the corresponding French firms.

Lequien et al. (2019) proceed in three main steps to merge PATSTAT and SIRENE:

1. For each *Siren* number from SIRENE, find a small subset of applicant firms in Patstat with phonetic similarities:
    - perform cleaning, splitting and phonetic encoding on firms' name in both databases. Too common words are deleted (THE, AND, CO, FRANCAISE ...).
    - sort each name by least frequent encoding in SIRENE. The more often a word appears in the database, the less information it can convey to identify firms.
    - for each SIRENE firm, the first (ie least frequent) cleaned word of the firm's name is compared with every PATSTAT name. All the PATSTAT names containing this word form a first subset of possible matches. Then the second word of the firm's name is compared with every name in this subset, reducing it further. This procedure stops before arriving at a null subset, and yields a set of likely PATSTAT matches for each SIRENE name. Very often this set is null because the majority of firms do not patent. On average, this subset contains 10 applicants, reducing a lot the computationally intensive comparisons.
2. Computation of parameters on these possible matches
    - Comparison of the names (raw names, and cleaned names), using Levenshtein distances and an inclusion parameter (all the words in one name are included in the name from the other database)
    - zip code comparison (*code postal*)
    - date comparisons (a firm cannot have patented before its creation)

[31] Although one must keep track of the different definitions of firms across these two datasets.

3. Matching with supervised learning

    - Sample from INPI (*Institut National de la Propriété Intellectuelle*) with 15,000 true matches between *Siren* number and `PATSTAT` *person id* (and in total 170,000 pairs, with the corresponding known mismatches).
    - This sample is randomly split into a learning sample and a verification sample (this procedure is repeated 10 times, and the recall and precision measures are averaged over them, so that the choice of the sample does not alter the results). This allows to choose the relevant variables and estimate the parameters.
    - apply this model on all the possible matches identified in the previous step.
    - in 90% of cases, unique matching. In the remaining 10% of cases, filter further with a decision tree (is the date of creation of the firm lower than the first filing of the applicant?, which couple has the minimum Levenshtein distance between raw names, between cleaned names, is one of the names included in the other?, which firm has the maximum number of employees?)

Based on the (rotating) verification sample taken from INPI data, the recall rate (share of all the true matchings that are accurate) is at 86.1% and the precision rate (share of the identified matches that are accurate) is at 97.0%.

B Time lag in exports reporting between production and customs data

The different timing for recording the export transaction between tax and customs authorities materializes in the annual data in particular when the transaction occurs at the end of a year t – it is recorded in the tax data for year t – but the shipment occurs at the beginning of the following year, in which case it is recorded in the customs data in year $t+1$. Because part of the January $t+1$ (customs) exports is recorded as year t (tax) exports, a firm with larger (customs) exports in January of year $t+1$ is expected to show a larger discrepancy between tax and customs exports in year t. Figure B1b reports the bin-scatter of the ratio of customs over production exports in year t (y axis) versus the share of January $t+1$ exports over exports in year t (both from the customs data, x axis), absorbing firm fixed effects. It shows that when January $t+1$ (customs) exports represent a bigger share of year t (customs) exports, then the customs data falls shorter than the production data for year t.

We extend this analysis over the last months of year t and the first months of year $t+1$ with the following regression:

$$X_{f,t}/X^*_{f,t} = \sum_{m=-6}^{6} \alpha_m \frac{X_{f,m}}{X_{f,t}} + \mu_{s(f,t),t} + \nu_f + \varepsilon_{f,t} \tag{B1}$$

where $X_{f,m}$ is (customs) exports for month m of year $t+1$ if $m > 0$, or month $12+m$ of year t if $m \leq 0$. 0 corresponds to December t, 1 to January $t+1$. We control for firm fixed effects and the sector of the firm. We keep in the regression only observations where $X_{f,t}/X^*_{f,t} \leq 10$ and where each share $\frac{X_{f,m}}{X_{f,t}} \leq 1$. Figure B1a reports the coefficients α_m along with their 95 and 99 confidence intervals (standard errors are clustered at the firm level). Everything else equal, if the first months of year $t+1$ represent a larger share of year t exports, then the ratio of yearly exports from customs to production data is smaller. Conversely if the last months of year t represent a larger share of yearly exports, then the customs yearly figure is bigger relative to the production figure.

Figure B1: Customs/production discrepancy in year t versus $t+1$ January share of year t exports

(a) Binscatter (b) Event Study

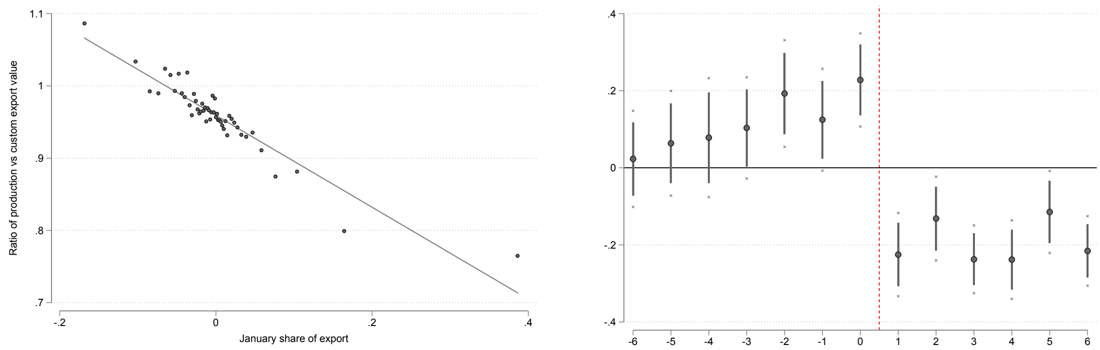

Notes: Left-hand side Figure reports the bin-scatter of the ratio of customs over production exports for year t (y axis) against the ratio of January $t+1$ exports over year t exports (both taken from customs data). Firm fixed effects are absorbed. Right-hand side Figure reports the coefficients α_m and corresponding 95% and 99% confidence intervals from equation (B1). Number of observations: 53,186 (Fig B1a) and 57,927 (Fig. B1b). Years: 1994-2012

C Additional Empirical results

Figure C1: OLS: Other variables

(a) Citations within 3 years

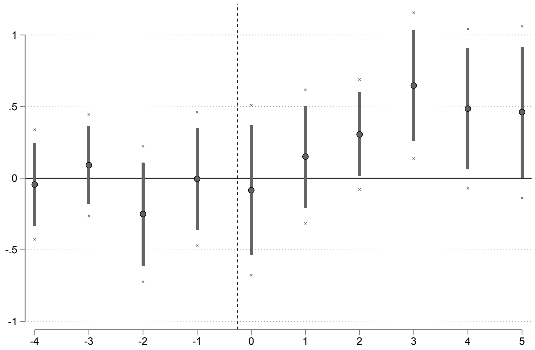

(b) Citations within 3 years

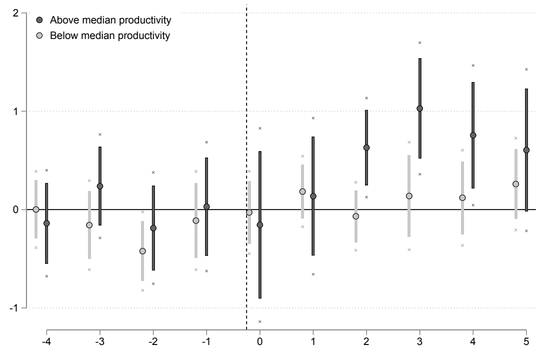

(c) Top 10% patents

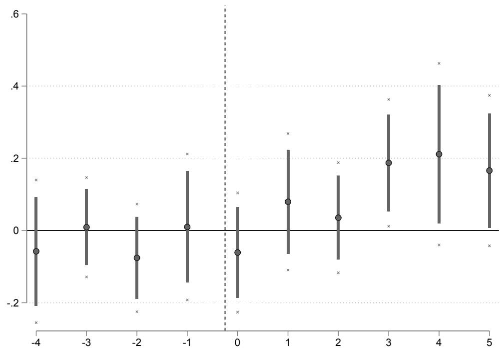

(d) Top 10% patents

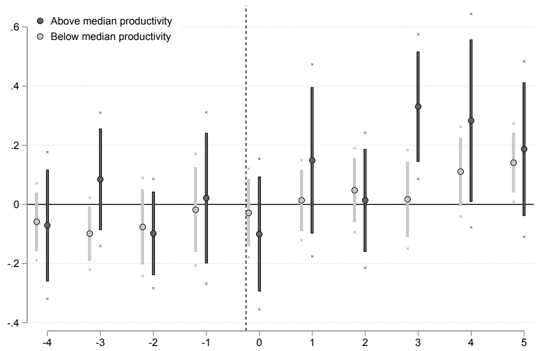

(e) All patents

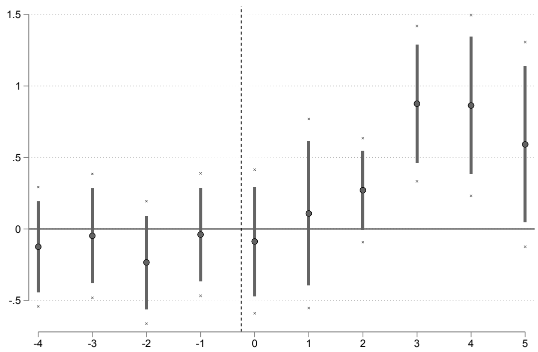

(f) All patents

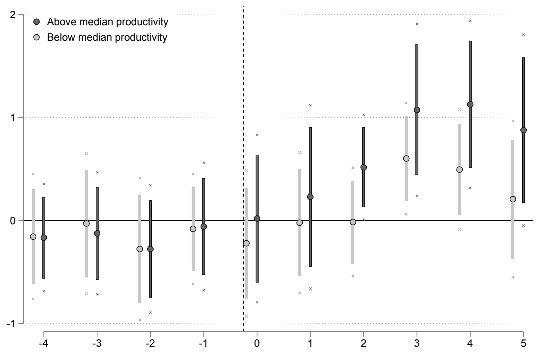

Notes: These Figures replicate Figures 4a and 6a but using different measures of innovation as the dependent variable: respectively counting citations received within a 3 year window, counting the number of patents among the 10% most cited in the year and counting any patent (whether priority or secondary filing). Number of observations: 22,237.

Figure C2: Inverse Hyperbolic Sine transformation of the dependent variable

(a) Priority patents

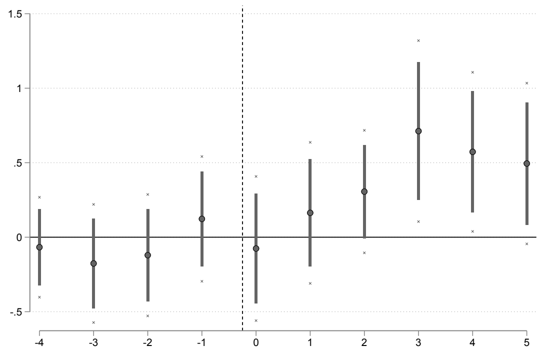

(b) Priority patents

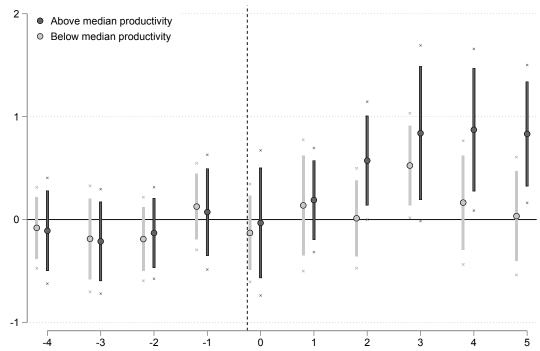

(c) Citations within 5 years

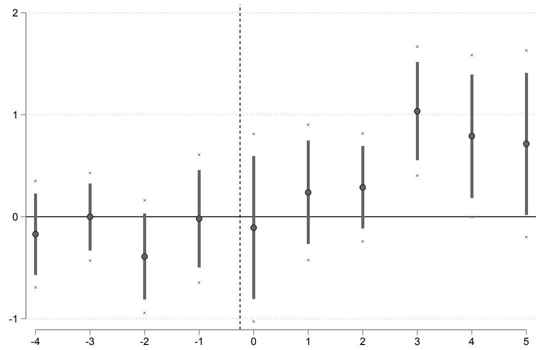

(d) Citations within 5 years

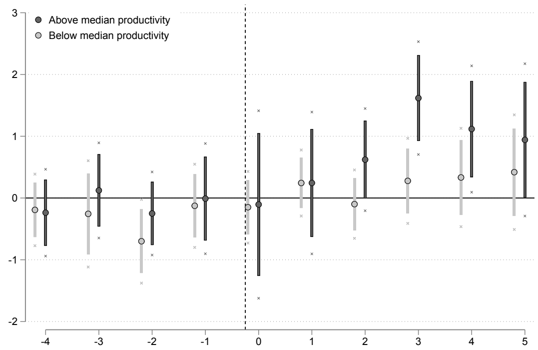

Notes: These Figures replicate Figures 4 and 6 but use the inverse hyperbolic sine transformation (asinh) of the number of priority patents or of 5-year citations as the dependent variable. Number of observations: 22,237

Table C1: Cumulative Response to Demand Shock - hyperbolic sine

	Priority Patents		Citations	
	Baseline	Hyperbolic Sine	Baseline	Hyperbolic Sine
Pre-Trend	-0.230 (0.339)	-0.241 (0.418)	-0.462 (0.478)	-0.583 (0.556)
Current	0.051 (0.253)	0.088 (0.323)	0.085 (0.469)	0.130 (0.564)
Future	1.680*** (0.532)	2.086*** (0.658)	2.368*** (0.738)	2.827*** (0.864)

Notes: Columns 2 and 4 of this table replicate columns 3 and 5 of Table 6 but use an inverse hyperbolic sine function ($H(x) = \log(x + \sqrt{x^2+1})$ instead of a logarithm function as the dependent variable. The corresponding columns 3 and 5 of Table 6 are replicated in columns 1 and 3. ***, ** and * indicate p-value below 0.01, 0.05 and 0.1 respectively.

Table C2: Cumulative Heterogeneous Response to Demand Shock

	Priority Patents		Citations	
	Baseline	Hyperbolic Sine	Baseline	Hyperbolic Sine
Pre-Trend				
Below Median	-0.288 (0.339)	-0.330 (0.424)	-1.071** (0.499)	-1.276** (.613)
Above Median	-0.350 (0.394)	-0.379 (0.485)	-0.265 (0.709)	-0.375 (0.840)
Difference	-0.062 (0.435)	-0.048 (0.546)	0.806 (0.714)	0.901 (0.867)
Current Below Median	-0.003 (0.249)	0.008 (0.317)	0.066 (0.296)	0.094 (0.359)
Above Median	0.100 (0.297)	0.157 (0.381)	0.081 (0.778)	0.139 (0.953)
Difference	0.103 (0.272)	0.149 (0.353)	0.015 (0.705)	0.045 (0.862)
Future				
Below Median	0.583* (0.336)	0.736* (0.431)	0.769 (0.674)	0.929 (0.830)
Above Median	2.515*** (0.718)	3.117*** (0.891)	3.609*** (0.909)	4.298*** (1.074)
Difference	1.932*** (.591)	2.380*** (.736)	2.841*** (0.967)	3.369*** (1.170)

Notes: Columns 2 and 4 of this table replicate columns 3 and 5 of Table 7 in the paper but use an inverse hyperbolic sine function ($H(x) = \log(x + \sqrt{x^2+1})$ instead of a logarithm function as the dependent variable. The corresponding columns 3 and 5 of Table 6 are replicated in columns 1 and 3. ***, ** and * indicate p-value below 0.01, 0.05 and 0.1 respectively.

Figure C3: OLS: removing the markets where a French firm is a leader

(a) Priority patents

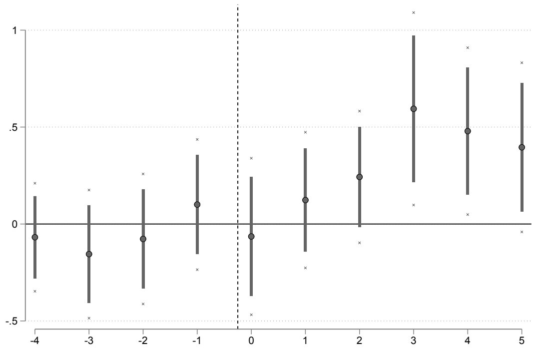

(b) Priority patents

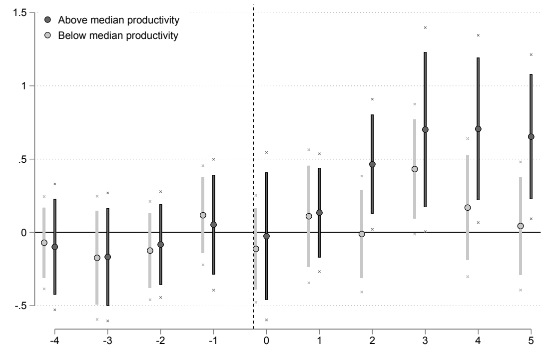

(c) Citations within 5 years

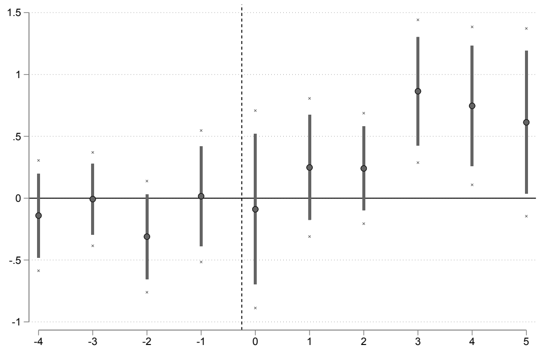

(d) Citations within 5 years

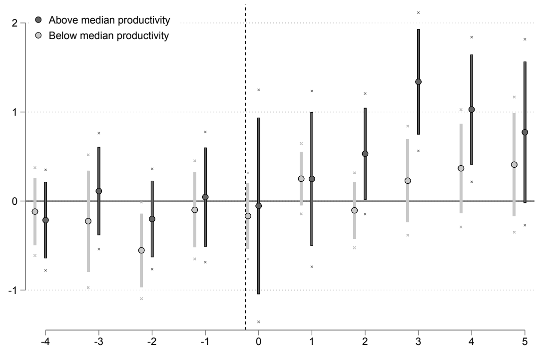

Notes: These Figures replicate Figures 4 and 6 but dropping firm-destination pairs whenever the firm's market share in the destination exceeds 10% in the construction of the demand shock variable. Number of observations: 22,156

Figure C4: Different trimming thresholds - OLS: Priority patents

(a) No Trimming

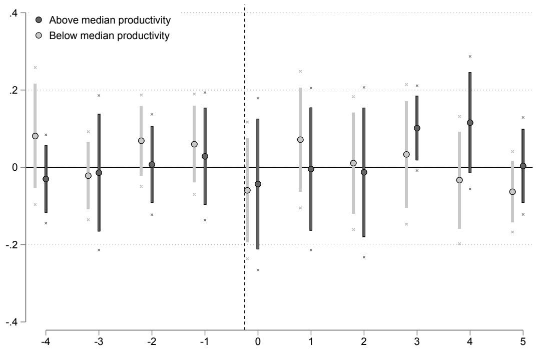

(b) Trimming 1%

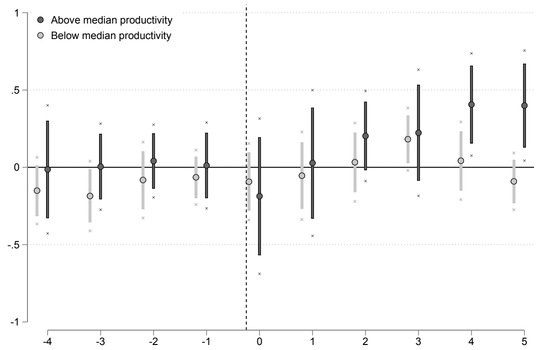

(c) Trimming 2%

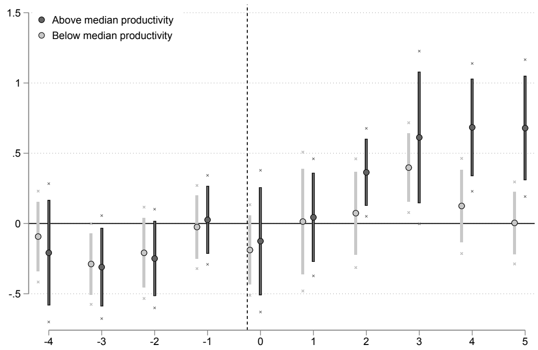

(d) Trimming 3%

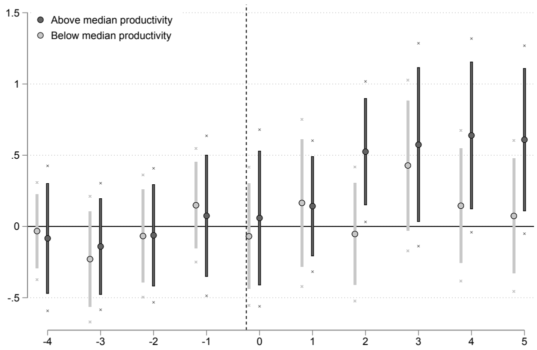

(e) Trimming 4%

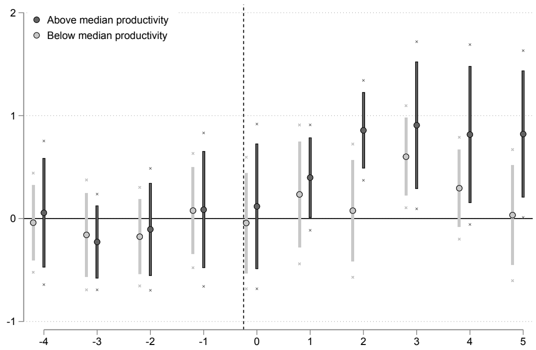

(f) Trimming 5%

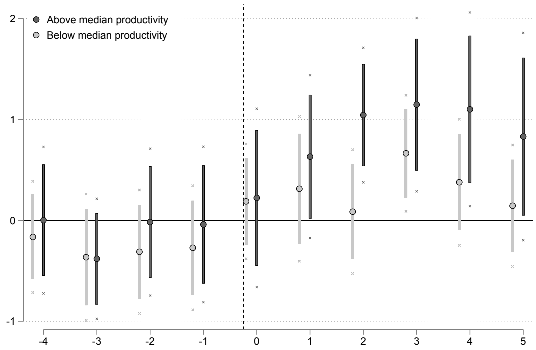

Notes: These Figures replicate Figure 6a but use a different trimming of the demand shocks, respectively: 0, 1, 2, 3, 4 and 5%. Number of observations: 27,036, 24,946, 23,149, 21,442, 19,887 and 18,522 respectively.

Figure C5: OLS: Sample of firms that innovated before 1994

(a) Priority patents

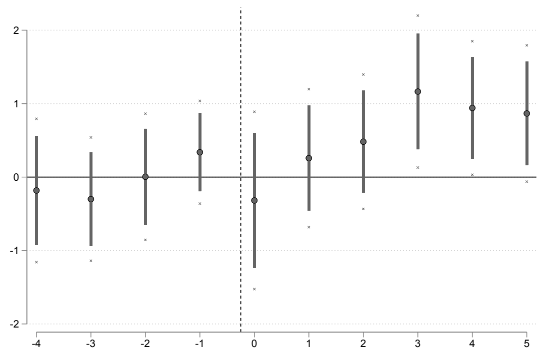

(b) Priority patents

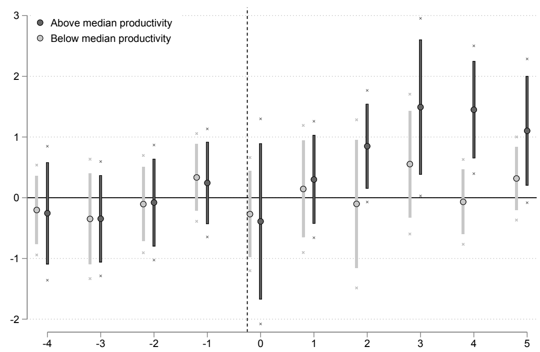

(c) Citations within 5 years

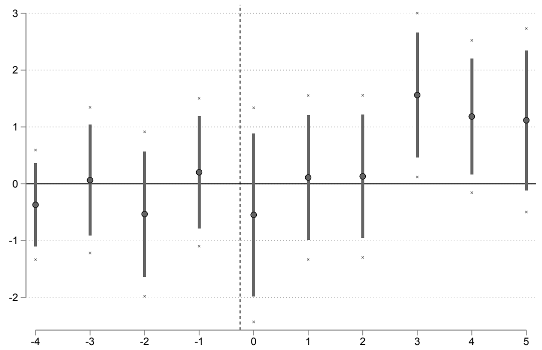

(d) Citations within 5 years

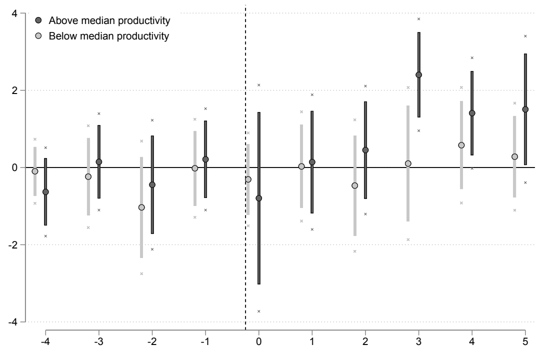

Notes: These Figures replicate Figures 4 and 6 but restricting to firms that innovated before 1994 (i.e. with a patent application with filing date between 1985 and 1994). Number of observations: 7,209

Figure C6: OLS: Sample of firms with $t_0 = 1994$

(a) Priority patents

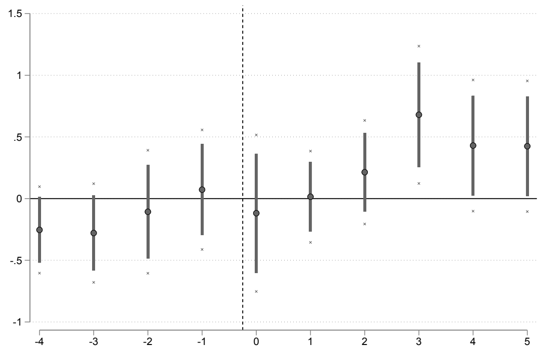

(b) Priority patents

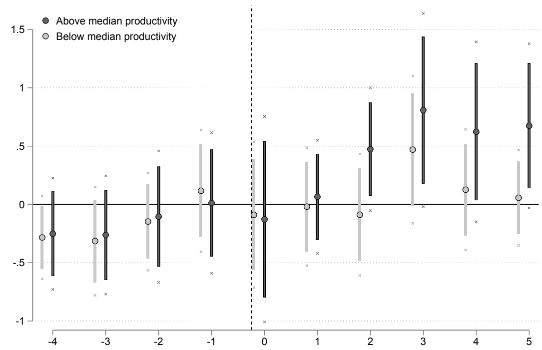

(c) Citations within 5 years

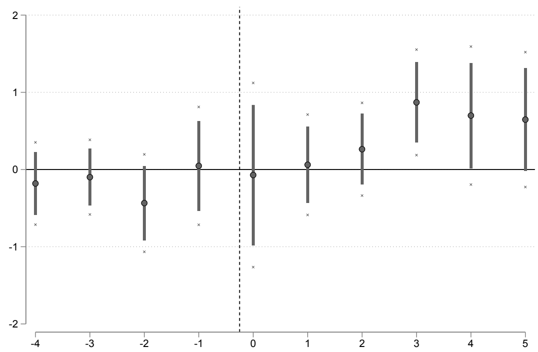

(d) Citations within 5 years

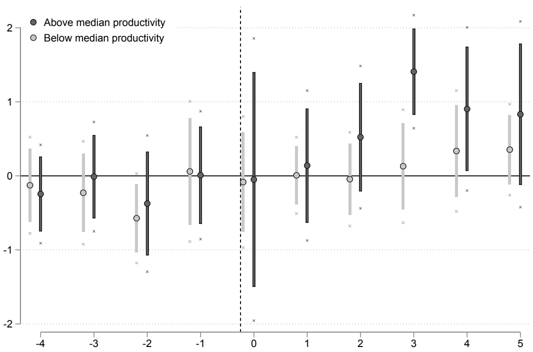

Notes: These Figures replicate Figures 4 and 6 but restricting to firms that first exported before 1994 ($t0 = 1994$ in equations (2) and (4)). Number of observations: 14,165

Figure C7: OLS: 6 LAGS

(a) Priority Patents

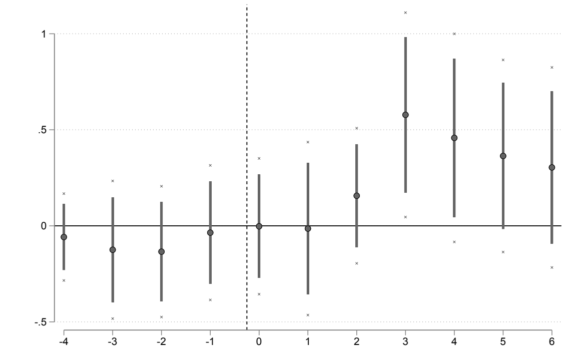

(b) Priority Patents

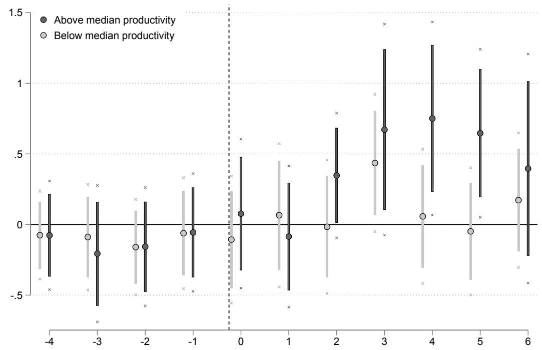

(c) Citations within 5 years

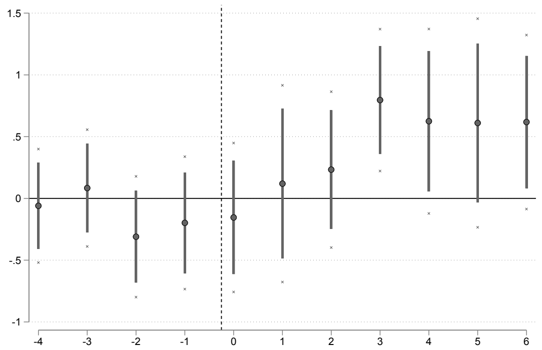

(d) Citations within 5 years

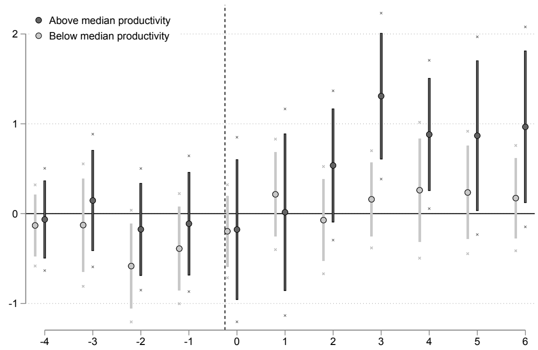

Notes: These Figures replicate Figures 4 and 6 but using 6 lags and 4 leads (therefore defining $k = 6$ and $k' = 4$ in equations (2) and (4)). Number of observations: 18,759

Figure C8: OLS: NO PRE-TREND

(a) Priority Patents

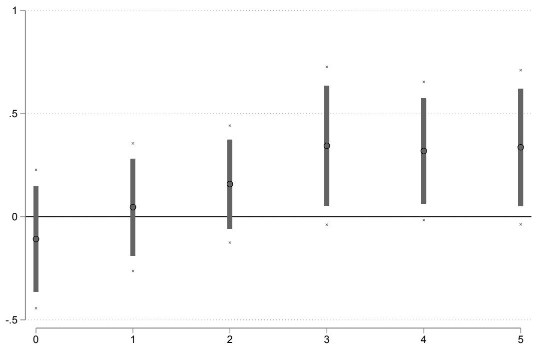

(b) Priority Patents

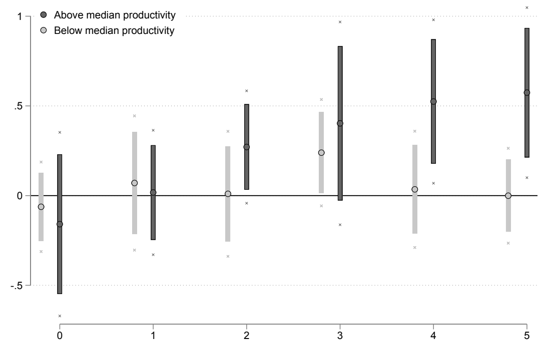

(c) Citations within 5 years

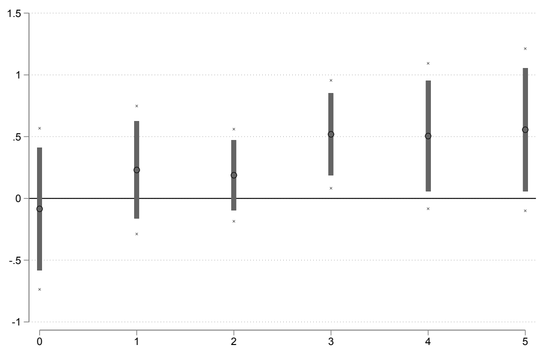

(d) Citations within 5 years

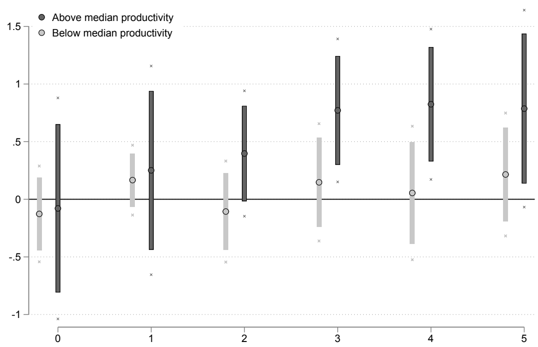

Notes: These Figures replicate Figures 4 and 6 but without using any lead (therefore defining $k = 5$ and $k' = 0$ in equations (2) and (4)). Number of observations: 25,315

Table C3: Detailed regression results

	One group		Two groups	
	OLS	NB	OLS	NB
Shock $D_{f,t}$				
at t+4	-0.066 (0.106)	0.288 (0.891)		
Below Median			-0.068 (0.120)	1.158 (1.003)
Above Median			-0.105 (0.164)	-0.404 (0.798)
at t+3	-0.152 (0.123)	-0.885* (0.525)		
Below Median			-0.157 (0.155)	-0.573 (0.986)
Above Median			-0.183 (0.159)	-1.094 (0.707)
at t+2	-0.101 (0.127)	-0.310 (0.809)		
Below Median			-0.155 (0.126)	-0.112 (0.887)
Above Median			-0.110 (0.137)	-0.641 (0.732)
at t+1	0.088 (0.131)	-0.093 (0.776)		
Below Median			0.093 (0.130)	1.433 (1.067)
Above Median			0.047 (0.173)	-0.943 (1.105)
Total pre-trend	-0.230 (0.339)	-1.000 (1.599)		
Below Median			-0.288 (0.339)	1.907 (2.496)
Above Median			-0.350 (0.394)	-3.082* (1.641)
at t	-0.068 (0.148)	-1.196 (1.038)		
Above Median			-0.108 (0.144)	-1.274 (1.094)
Above Median			-0.035 (0.216)	-0.892 (1.237)
at t-1	0.119 (0.143)	0.085 (0.740)		
Below Median			0.105 (0.193)	1.336 (1.312)
Above Median			0.135 (0.153)	-0.480 (0.762)
Total current	0.051 (0.253)	-1.111 (1.566)		
Below Median			-0.003 (0.249)	
Above Median			0.100 (0.297)	-1.372 (1.768)
at t-2	0.244* (0.127)	1.364** (0.670)		
Below Median			0.012 (0.148)	0.350 (1.217)
Above Median			0.455*** (0.175)	2.205*** (0.721)
at t-3	0.573*** (0.188)	2.322** (1.106)		
Below Median			0.409*** (0.157)	1.918 (1.629)
Above Median			0.685*** (0.265)	2.548** (1.035)
at t-4	0.465*** (0.166)	1.847* (0.911)		
Below Median			0.131 (0.180)	0.673 (1.546)
Above Median			0.709 (0.244)	2.422 (1.125)
at t-5	0.398** (0.168)	1.493 (1.153)		
Below Median			0.031 (0.172)	-0.224 (1.612)
Above Median			0.666*** (0.213)	2.406** (1.095)
Total Future	1.680*** (0.532)	7.026** (3.117)		
Below Median			0.583 (0.336)	2.717 (3.113)
Below Median			2.515*** (0.718)	9.582*** (3.066)
Initial Stock of patent (log)	0.744*** (0.078)	1.274*** (0.108)	0.735*** (0.068)	1.245*** (0.119)
Initial Employment (log)	0.052 (0.038)	0.240 (0.209)	0.047 (0.038)	0.191 (0.212)
Initial Sales (log)	0.042** (0.018)	0.350*** (0.112)	0.044** (0.015)	0.369*** (0.097)
Initial Employment variation (log)	-0.002 (0.018)	-0.043 (0.100)	0.020 (0.018)	0.080 (0.115)
Initial Sales variation (log)	0.048*** (0.016)	0.425*** (0.063)	0.027* (0.015)	0.298*** (0.082)
Fixed Effects				
Sector × year	✓	✓	✓	✓
Productivity group × year			✓	✓
Initial export intensity × year	✓	✓	✓	✓
Obs.	22,237	22,237	22,237	22,237
R2	0.294		0.296	

Notes: This table reports point estimates and standard errors (under parentheses) of all coefficients underlying Figure 4a (equation (2), col 1), Figure 5a (equation (3), col 2), Figure 6a (equation (4), col 3) and Figure 7a (equation (5), col 4). Firms are either taken all together (columns 1 and 2) or grouped according to their initial level of productivity (Above Median and Below Median, columns 3 and 4). The dependent variable is the log of 1 + the number of priority patent applications (col 1 and 3) and the number of priority patent applications in columns 2 and 4. In addition to the estimation of each coefficient, the Table also reports the Wald test on different linear combinations of coefficients (see Tables 6 and 7). ***, ** and * indicate p-value below 0.01, 0.05 and 0.1 respectively. Standard errors are clustered by groups of sector (2 digit) and group of productivity.

D Falsification Tests

In order to reinforce our finding of a causal impact for market size (via our demand shocks) on innovation by above-median productivity firms, we develop a falsification test that highlights that those innovation responses cannot be explained by a firm-level trend: that is, that those firms observed to increase innovation would have done so anyway absent an increase in export demand.[32] This test also provides a further check on our control that the innovation response is not explained by a firm's prior exposure to export markets (since we use prior export intensity to construct our trade shocks).[33] Our test involves the construction of placebo demand shocks for each firm and then showing that firm innovation does not respond to this placebo shock.

In our first placebo construction, we allocate products to firms randomly (based on their empirical distribution across firms) and compute the demand shocks that each firm would have experienced had it actually exported those products at t_0. In our second placebo construction, we instead allocate the export destinations randomly across firms (again, based on the empirical distribution of destinations across firms).[34]

We construct 2000 different placebo demand shocks using both methods, and then estimate our baseline OLS specification (4) each time with the response of priority patents on the left-hand side. Figure D1 shows the cumulative distribution for the coefficient $\alpha_{H,4}$ and its t-statistic for the response by firms with above-median productivity 4 years after the shock. Against those distributions, we show (red vertical line) the coefficient value and t-statistic for $\alpha_{H,4}$ that we reported in Figure 6a using the 'true' demand shocks. We immediately see that the value and significance of the demand shock coefficient we previously obtained are clear outliers in those distributions (well beyond the 100th percentile for the coefficient values; and at the 98.6 and 95.6 percentiles for the associated t-statistics). We can thus easily reject the hypothesis that a similar innovation response by the above-median productivity firms would have been observed absent the impact of the "true" demand shock. We have repeated this falsification test summing the coefficients representing 2 to 5 years after the shock (instead of just year 4), along with its associated t-statistic. In all those cases, our reported coefficients (and their t-statistics) are again clear outliers in the simulated cdf: above the 95th percentile of the distribution

[32] A similar approach has been implemented by Chetty et al. (2009) and Malgouyres et al. (2019).

[33] Our main check is to add export intensity interacted with the year fixed effects as controls.

[34] To be more precise, each placebo demand shock is the outcome of a random permutation across firms from either the empirical distribution of products, or the empirical distribution of destinations.

in all cases (and above the 100th percentile in a few). We have also repeated this exercise, with similar results, with citations as the dependent variable.

Figure D1: OLS: Falsification Tests

(a) Randomly Switch Products

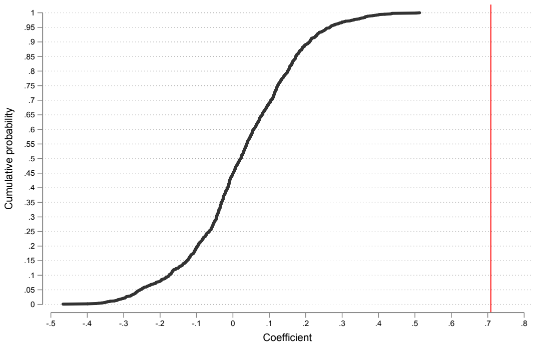

(b) Randomly Switch Destinations

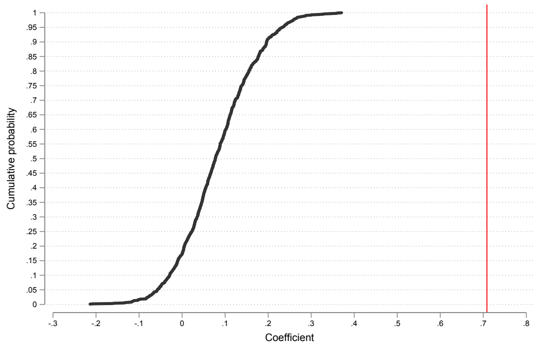

(c) Randomly Switch Products

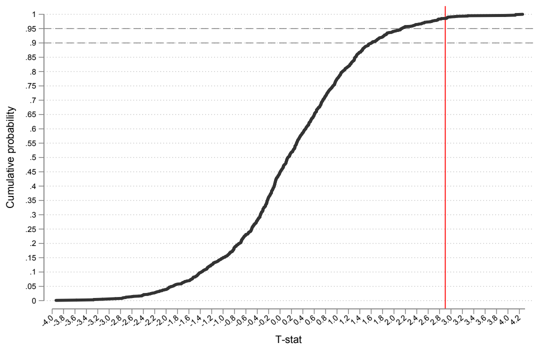

(d) Randomly Switch Destinations

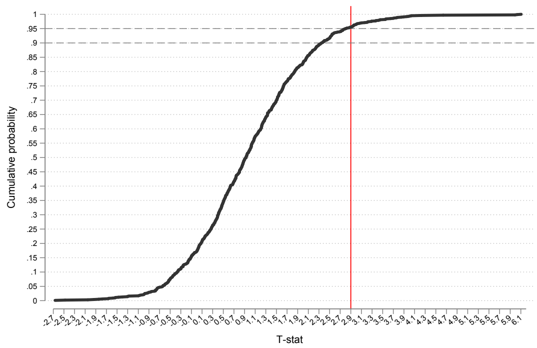

Notes: This figure plots the cumulative distribution of the point estimates (top panels) and the associated t-stat (bottom panels) and for the $\alpha_{H,4}$ coefficient when equation (4) is estimated 2000 times with a placebo shock, randomly switching the products exported at t_0 (left panel) or randomly switching the export countries at t_0 (right panel). $\alpha_{H,4}$ coefficient and t-stat from Figure 6a in red line.

E Shift share

Our shock belongs to the family of shift-share ("Bartik") instruments which have been the subject of a recent active literature (e.g. Borusyak et al., 2021; Goldsmith-Pinkham et al., 2020; Adao et al., 2019). This literature in particular offers a number of diagnostics and tests to assess the plausibility of the empirical strategy's validity. We follow Borusyak et al. (2021) and consider an approach based on the "exogenous shock assumption", in contrast the paper by Goldsmith-Pinkham et al. (2020) takes the "exogenous share assumption".

Formally, a shift share variable can be seen as the dot product of a vector of exposure shares (in our case measures at the firm level and taken at t_0) and of a vector of shocks (in our case, the growth rate of import demand for a given product in a given country). The framework of Goldsmith-Pinkham et al. (2020) assumes exogeneity of the exposure shares and imposes no restriction regarding the exogeneity of the shocks while Borusyak et al. (2021) explores a framework based on the assumptions that shocks are quasi random. In our setting, we argue that the shocks are idiosyncratic in the sense that they are viewed as independent from any country-product unobservables that could affect French firms' innovation.

To summarise their approach, we denote our shift share measures as:

$$z_{f,t} = \sum_{o \in \Omega} w_{o,f} g_{o,t}$$

We choose these notations to match Borusyak et al. (2021)'s, f is a firm, o an export market composed of a country j and a product s, w is the share of market o in the basket of export markets of firm f in the first year we observe exports for f and g is the growth rate of import demand for a good s by a country j. From this measure, we construct the average shock exposure w_o as:

$$w_o = \sum_{f \in \mathcal{F}} w_{o,f} e_f,$$

where $\mathcal{F}$ denotes the set of French firms. These exposure shares help to assess the plausibility of their main assumptions. In our setting, we are not concerned that there are too few shocks as the number of export markets is very large (more than 100,000) but these shocks can be too concentrated which could alter the finite sample properties of the estimator in a shift-share setting. Moreover, analysing the variation of these shocks, weighting by exposure shares w_o across different groups (country, product, time) is helpful to think about potential omitted variables in our main regressions. Finally, these exposure shares are useful to estimate an equivalent export market level regression which allows to (1) construct exposure-robust standard errors and (2) control for export market specific heterogeneity.

In Table E1, we replicate the analysis presented in Borusyak et al. (2021, Table 1) in the context of Autor et al. (2013). This Table presents the anatomy of our shocks across export markets and time. Column 1 includes the "missing" export market. This arises because firm-level weights do not sum up to one due to multiplication by the export intensity. Hence, the "missing" export market in our case corresponds to the part of the production that is sold domestically. In column 1, the average value of the shock is 0.016 which is significantly different from the average presented in column 2 which excludes this market. The effective number of shocks, measured by the inverse HHI of the exposure share is also significantly larger once domestic sales are excluded (22 against 50,939). The effective number of markets decreases to 158 when the shock is averaged using 2 digit products (instead of 6 digits, that is moving from about 5,000 categories to 100). This suggests that we can reduce the variation coming from the product side of our shocks when they are clustered by larger product categories.
Digging further into the shift share, we present the average exposure share by country in Table E2. In this table, we show the top 10 countries in terms of the average value of w_o. We clearly see that (1) the list of countries unsurprisingly follows the list of top trade partners with France and (2) no country accounts for a disproportionately large share.

Table E1: Shock Summary Statistics

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mean	0.0116	0.0485	-0.0117	0.0138	-0.0018	-0.0123	0.0132	-0.0018
Standard Deviation	0.1960	0.3985	0.3954	0.3982	0.3951	0.3501	0.3799	0.3400
Interquartile Range	0.0000	0.2854	0.2861	0.2851	0.2855	0.2563	0.2388	0.2362
Specification								
Excluding missing market		✓	✓	✓	✓	✓	✓	✓
Residualized on product			✓		✓			
Residualized on country				✓	✓			
Residualized on product-year						✓		✓
Residualized on country-year							✓	✓
Observation counts								
Nb of markets	105,003	105,002	105,002	105,002	105,002	105,002	105,002	105,002
Nb of countries	199	198	198	198	198	198	198	198
Nb of products	4,596	4,595	4,595	4,595	4,595	4,595	4,595	4,595
Nb of markets at product 2 digit	8,111	8,110	8,110	8,110	8,110	8,110	8,110	8,110
Largest share w_o								
Largest share across o	0.0482	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002
Larges share across markets at product 2 digit	0.7606	0.0425	0.0425	0.0425	0.0425	0.0425	0.0425	0.0425
Effective sample size								
1/HHI	22,3	50,939	50,939	50,939	50,939	50,939	50,939	50,939
1/HHI across markets at product 2 digit	1.73	158	158	158	158	158	158	158

Notes: This table summarises the shocks $g_{o,t}$ for each export market o (a country and a product at the 6 digit level) in the spirit of Borusyak et al. (2021, Table 1). Column 1 includes the "missing market" which corresponds to domestic sales. Mean, Standard Deviation and Interquartile range are computed after a residualization of the demand vector $g_{o,t}$ on a set of fixed effects and weighted by the exposure share w_o. Markets at product 2 digit correspond to an association of a country and a product taken at the 2 digit level.

Finally, we present results from the equivalent regression at the shock level, again following Borusyak et al. (2021). Table E3 presents the results when the dependent variable is the logarithm of the number of priority patents + 1, and the regressors include different leads and lags of the shocks. One advantage of this equivalent regression is that (1) we can cluster the standard errors by 2 digit product category and (2) we can include

Table E2: Equivalent export market regressions results

Rank	Country	Average Exposure share	Rank	Country	Average Exposure share
1	DE	0.066	6	BE	0.033
2	US	0.053	7	JP	0.028
3	CH	0.042	8	ES	0.027
4	GB	0.039	9	NL	0.022
5	IT	0.037	10	CN	0.021

Notes: This table presents the average value of exposure share for the top 10 countries. Countries with less than 1000 observations are excluded.

Table E3: Equivalent export market regressions results

	(1)	(2)	(3)	(4)	(5)	(6)
Pre-Trend	-0.120 (0.143)	-0.023 (0.150)	-0.339** (0.155)	-0.143 (0.179)	-0.094 (0.148)	-0.300* (0.158)
Current	-0.113 (0.078)	0.006 (0.077)	-0.134 (0.093)	0.001 (0.119)	-0.079 (0.082)	-0.172* (0.096)
Future	0.663*** (0.144)	0.208* (0.126)	0.871*** (0.190)	0.623*** (0.198)	0.503*** (0.157)	0.891*** (0.227)

Notes: This table run the equivalent shift share regression at the shock level based on Borusyak et al. (2021). Column 1 does not use any fixed effect, column 2 include a product and country fixed effects, column 3 includes a country-year fixed effect, column 4 uses a product-year fixed effects, column 5 use a country-2digit-product fixed effects and column 6 includes a country-2digit-product-year fixed effect. Number of observations: 538,249. Standard errors are clustered at the 2 digit product level.

product and or country fixed effects. This is what we do in the different columns. By using this lead/lag specification, we can directly assess the validity of the shock exogeneity assumption by considering the significativity of the sum of pre-trend coefficients. This assumption is verified in column 1, where we do not include any fixed effect and in other columns as long as product level fixed effects are included (except for column 6 which uses an interacted country - 2 digit product category - year fixed effects, i.e. leveraging only variations within the same product bundle from a same destination).

F Theoretical appendix

We describe how the equilibrium competition level λ in destination D is endogenously determined and show that λ increases with L. Although this equilibrium involves all the firms operating in D, including both the French exporters to D along with the domestic producers in D, we show that the equilibrium competition level λ is determined independently of the export supply to D (which then only impacts the number of domestic entrants and producers).

Let $\Gamma_D(\tilde{c})$ denote the cumulative distribution of baseline costs $\tilde{c}$ among domestic producers in D. We assume that $\Gamma_D(\tilde{c})$ has support on $[\tilde{c}_{0D}, +\infty)$ with $\tilde{c}_{0D} > \tilde{c}_{\min}$. Let F_D denote the fixed production cost faced by those domestic firms in D. Since a firm's operating profit is monotonic in its baseline cost $\tilde{c}$, producing for the domestic market D is profitable only for domestic firms with a baseline cost $\tilde{c}$ below a cutoff value $\widehat{C}_D$ defined by the zero profit condition:

$$\Pi(\widehat{C}_D, 0; \lambda) = F_D, \tag{ZCP}$$

where we have assumed that $\widehat{C}_D > \widehat{C}_I$ so that the firm with the cutoff cost $\widehat{C}_D$ does not innovate (and hence does not incur any innovation cost). Entry is unrestricted subject to a sunk entry cost F_D^E. In equilibrium, the expected profit of a prospective entrant will be equalized with this cost, yielding the free-entry condition:

$$\int_{\tilde{c}_{0D}}^{\widehat{C}_D} [\Pi(\tilde{c}, k(\tilde{c}; \lambda); \lambda) - F_D] \, d\Gamma_D(\tilde{c}) = F_D^E. \tag{FE}$$

Proposition 1 *The two conditions (ZCP) and (FE) jointly determine a unique pair $(\lambda, \widehat{C}_D)$.*

Proof *Uniqueness*: in $(\widehat{C}_D, \lambda)$ space, the (ZCP) condition is strictly downward-sloping while the (FE) condition is strictly upward-sloping, ensuring uniqueness of the equilibrium if such an equilibrium exists. More precisely: (a) an increase in competition from λ to $\lambda + d\lambda$ reduces the profit of firms with baseline cost $\widehat{C}_D(\lambda)$, so that those firms no longer operate; this means that $\widehat{C}_D(\lambda + d\lambda) < \widehat{C}_D(\lambda)$, which proves that the (ZCP) curve is strictly downward-sloping; (b) an increase in competition from λ to $\lambda + d\lambda$ reduces the profit of all firms (the envelope theorem ensures that at the optimal innovation level $\frac{\partial \Pi}{\partial k} = 0$ so that $\frac{d\Pi}{d\lambda} = \frac{\partial \Pi}{\partial \lambda} < 0$); this in turn means that $\widehat{C}_D$ has to strictly increase for the (FE) condition to hold, which proves that the (FE) curve is strictly upward-sloping.

Existence: We show that the (FE) curve lies below the (ZCP) curve for values of $\widehat{C}_D$ close to $\tilde{c}_{0D}$, and that the (FE) curve ends up above the (ZCP) curve for high values of $\widehat{C}_D$. As $\widehat{C}_D$ becomes close to $\tilde{c}_{0D}$, (ZCP) implies a value for λ which is positive and bounded away from zero, whereas (FE) requires λ to become arbitrarily small, because the integrand must go to $+\infty$ for the integral over a very small interval to remain equal to F_D^E. Next, recall that the (ZCP) curve must remain below the $\lambda = \frac{\alpha}{\widehat{C}_D}$ curve. Given that $\frac{\alpha}{\widehat{C}_D} \longrightarrow 0$ when $\widehat{C}_D \longrightarrow +\infty$, the $\frac{\alpha}{\widehat{C}_D}$ curve must cross the (FE) curve at some point. At this point, the (ZCP) curve lies below the (FE) curve.

For simplicity, we have abstracted from any export profits for the domestic firms. This is inconsequential for our prediction that the equilibrium competition level λ increases with market size L, so long as destination D is small relative to the size of the global export market.[35]

Proposition 2 *An increase in market size L in D leads to an increase in competition λ.*

Proof We prove this proposition by contradiction. If λ were to decrease, then the cutoff $\widehat{C}_D$ would have to increase (see (ZCP)). Since $\pi(c;\lambda)$ is decreasing in λ, then $\Pi(\tilde{c}, k; \lambda)$ must also increase for any given innovation level k when λ decreases. Given the optimization principle, $\Pi(\tilde{c}, k(\tilde{c};\lambda); \lambda)$ must also increase. This, together with an increase in the cutoff $\widehat{C}_D$, represents a violation of the (FE) condition. Thus competition λ must increase when L increases.

[35]More precisely, the free entry condition can be extended to incorporate the (net) export profits Π_{-D} earned in other destinations:

$$\int_{\tilde{c}_{0D}}^{\widehat{C}_D} \left[\Pi(\tilde{c}, k(\tilde{c};\lambda); \lambda) - F_D + \Pi_{-D}(\tilde{c}, k; \{\lambda_{-D}\})\right] d\Gamma_D(\tilde{c}) = F_D^E,$$

where $\{\lambda_{-D}\}$ denotes the vector of competition levels in countries other than D. So long as these competition levels $\{\lambda_{-D}\}$ do not respond to changes in D, the export profits shift up the marginal benefit of innovation in (FOC) by an amount that does not depend on λ or L. This marginal benefit curve will remain an increasing function of innovation k and will shift up with any market-wide change in D that increases firm output $Q(\tilde{c}, k; \lambda)$ at fixed innovation k.