WORKING paper

The Impact of Exports on Innovation: Theory and Evidence

Philippe Aghion[1] & Antonin Bergeaud[2] &

Matthieu Lequien[3] & Marc J. Melitz[4]

April. 2018, WP #678

ABSTRACT

This paper investigates the effect of export shocks on innovation. On the one hand a positive shock increases market size and therefore innovation incentives for all firms. On the other hand it increases competition as more firms enter the export market. This in turn reduces profits and therefore innovation incentives particularly for firms with low productivity. Overall the positive impact of the export shock on innovation is magnified for high productivity firms, whereas it may negatively affect innovation in low productivity firms. We test this prediction with patent, customs and production data covering all French manufacturing firms. To address potential endogeneity issues, we construct firm-level export proxies which respond to aggregate conditions in a firm's export destinations but are exogenous to firm-level decisions. We show that patenting robustly increases more with export demand for initially more productive firms. This effect is reversed for the least productive firms as the negative competition effect dominates.[5]

Keywords: Innovation, trade, export, demand shocks, patents

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

[1] Collège de France, London School of Economics and Paris School of Economics, p.aghion@lse.ac.uk

[2] Banque de France, antonin.bergeaud@banque-france.fr

[3] Banque de France and Paris School of Economics, matthieu.lequien@banque-france.fr

[4] Harvard University and NBER, mmelitz@fas.harvard.edu

NON-TECHNICAL SUMMARY

We combine modeling insights from the literature on firm heterogeneity and trade into the new growth theory. The former literature has focused on the causality link from productivity to trade whereas the latter literature has focused on the reverse link from trade to productivity by investigating the various channels whereby trade liberalization affects innovation-led productivity growth across firms. Our model derives testable predictions on how firms' access to export markets affects innovation; and how this link will vary across exporters. We then take advantage of the availability of exhaustive firm-level data on productivity, trade, and patenting in France to test these predictions.

The main insight from the model is that export has two impacts on firm's innovation response: a market size effect and a competition effect. Think of French firms that export to China. An increase in Chinese demand for products produced by these firms will first have a direct market size effect: namely, the expanded market for exports will increase the size of innovation rents and thereby increase those firms' incentives to invest more in innovation. But it will also have a competition effect: namely, the expanded market for exports will attract new firms into the Chinese market as more firms find it profitable to sell there their product; this in turn will raise competition for exporters into that market. Due to the nature of competition between firms in the model this effect dissipates the larger is the firm's productivity. This competition effect is therefore most salient for French firms with initially lower market shares and higher production costs (these firms will suffer more than -or at the expense of- more efficient exporting firms). Hence our prediction that a positive export shock should raise innovation more in more frontier firms; and that it may induce less innovation for those firms that are far from the frontier.

To test this prediction, we use three exhaustive datasets on French firms: patenting, production, and customs data. To disentangle the direction of causality between innovation and export performance, we construct a firm-level export demand variable following Mayer et al. (2016). This variable responds to aggregate conditions in a firm's export destinations but is exogenous to firm-level decisions (including the concurrent decisions for export-market participation and the forward looking innovation response). We show that: (i) firms that are initially more productive (closer to their sector's technology frontier) strongly respond to a positive export demand shock by patenting more; (ii) this effect dissipates for firms further from the "frontier" and is reversed for a subset of initially less productive firms (see Figure). These results confirm the predictions of the model for both a market size and a competition effect of the export shock. They are robust to (i) considering other patent indicators; (ii) controlling for firm specific characteristics; (iii) excluding dominant firms in a destination market; (iv) considering alternative measures and heterogeneous effects for a firm's proximity to frontier; and (v) controlling for pre-trends with sector-decile specific time trends. We also show that the competition effect is most salient in high competition export destinations. Finally we obtain similar results with an alternate specification based on long time-differences.

Figure: Patenting increases more with demand for initially more efficient firms
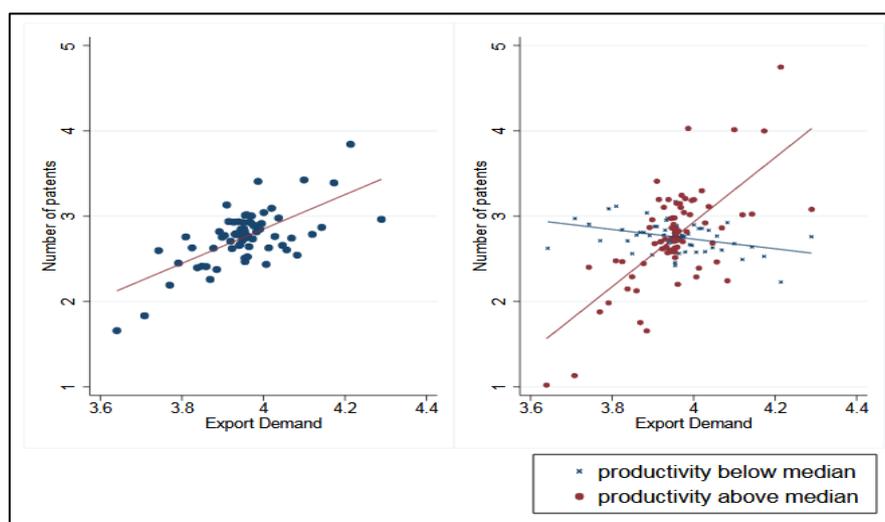
Notes: Each dot represents the quantile-average of the residual of the number of patent on a firm fixed effect (y-axis) against the quantile-average of the residual of the demand variable on a firm fixed effect (x-axis), for all firms (left-hand side panel) or for both firms below (blue) or above (red) the productivity median at initial date (right-hand side panel).

L'effet du commerce international sur l'innovation, théorie et résultats empiriques

Cet article examine l'effet des chocs de demande à l'exportation sur l'innovation. D'une part, un choc positif augmente la taille du marché et donc les incitations à l'innovation pour toutes les entreprises. D'autre part, un tel choc accroît la concurrence à mesure qu'un plus grand nombre d'entreprises entrent sur le marché d'exportation. Cela a pour effet de réduire les bénéfices et donc les incitations à l'innovation, en particulier pour les entreprises dont la productivité est faible. Dans l'ensemble, l'impact positif du choc à l'exportation sur l'innovation est amplifié pour les entreprises à productivité élevée, alors qu'il peut affecter négativement l'innovation dans les entreprises à faible productivité. Nous testons cette prédiction avec des données de brevets, de douanes et de production couvrant l'ensemble des entreprises manufacturières françaises. Pour traiter les questions d'endogénéité potentielles, nous construisons des indicateurs d'exportation au niveau de l'entreprise qui répondent aux conditions globales des destinations d'exportation d'une entreprise, mais qui sont exogènes aux décisions prises au niveau de l'entreprise. Nous montrons que le nombre de brevets augmente fortement avec la demande à l'exportation pour des entreprises initialement plus productives. Cet effet est inversé pour les entreprises les moins productives, car l'effet négatif sur la concurrence domine.

1   Introduction

In this paper we combine modeling insights from the literature on firm heterogeneity and trade (see Melitz and Redding, 2014 for a survey) into the new growth theory (e.g., see Aghion and Howitt, 2009). The former literature has focused on the causality link from productivity to trade whereas the latter literature has focused on the reverse link from trade to productivity by investigating the various channels whereby trade liberalization affects innovation-led productivity growth across firms. Our model derives testable predictions on how firms' access to export markets affects innovation; and how this link will vary across exporters. We then take advantage of the availability of exhaustive firm-level data on productivity, trade, and patenting in France to test these predictions.

Figure 1 below motivates our analysis. The curve depicts, for each percentile in exports, the share of French innovative firms within that percentile. It clearly shows a positive relationship between exports and innovation.[1] One of the most striking features that emerges from our merged production-export-innovation dataset is a massive correlation between export and innovation performance across firms. This holds both at the extensive margin (exporters are substantially more likely to innovate, and innovators are more likely to export) as well as the intensive margin (large exporters tend to be big innovators and vice-versa). We describe these relationships in much more detail in Section 3. Does this correlation reflect a causal effect of export on innovation, or the effect of innovation on exports, or both? How does the innovation behavior of a firm react to its export markets' conditions? Our paper is a first attempt at understanding these firm-level patterns connecting innovation and trade using the matching between patenting, balance sheet, and customs exhaustive datasets.

In the first part of the paper we develop a simple model of trade and innovation with heterogeneous firms. The model is one with monopolistic competition and heterogeneous firm capabilities, but adds the innovation dimension of new growth theory to it. It features a continuous set of firms indexed by their heterogeneous production costs. Innovation allows firms to reduce their production costs by an amount that increases with the size of the innovation investment. Think of French firms that export to China. An increase in Chinese demand for products produced by these firms will have two main effects on their innovation incentives. First, a direct market size effect: namely, the expanded market for exports will increase the size of innovation rents and thereby

---

[1]Obviously the relationship shown in Figure 1 is partly driven by a scale effect: large exporters are larger firms and larger firms are more likely to innovate. However, as we will see below, the distribution of export is more skewed than the distribution of sales or value added. In addition, the positive relationship still exists when centiles of exports intensity (exports divided by sales) are used.

Figure 1: The share of innovators jumps at the top of the export distribution
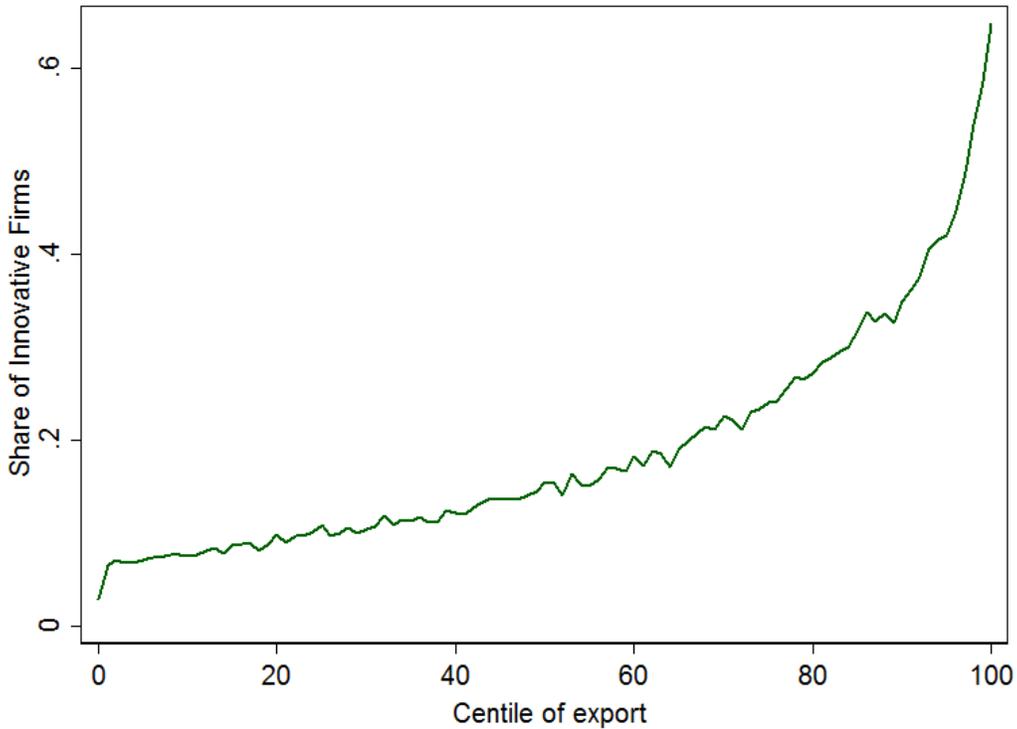
increase those firms' incentives to invest more in innovation. Second, a competition effect: namely, the expanded market for exports will attract new firms into the Chinese market as more firms find it profitable to sell there their product; this in turn will raise competition for exporters into that market. Due to the nature of competition between firms – featuring endogenous markups – this effect dissipates the larger is the firm's market share (and hence its productivity). This competition effect is therefore most salient for French firms with initially lower market shares and higher production costs (these firms will suffer more than -or at the expense of- more efficient exporting firms). Hence our prediction that a positive export shock should raise innovation more in more frontier firms; and that it may induce less innovation for those firms that are far from the frontier.

In the second part of the paper we take this prediction to the data. More specifically, we merge three exhaustive firm-level datasets – patenting, production, and customs data –, which cover the whole population of French firms to analyze how the access to export markets affects the quantity and quality of patents generated by these firms. The patent data are drawn from PATSTAT (Spring

2016 version) and contain detailed information on all patents and patent applications from the main patent offices in the world. We use an algorithm developed in Lequien et al. (in progress) that matches a French firm's name with its unique administrative identifier. This allows us to link the innovation activities of a firm with all other firm data sources. The production datasets FICUS and FARE, from INSEE/DGFiP, contain balance sheet information for each firm registered in France from 1993 to 2012 (total and export sales, number of employees, sector, etc.). French customs trade data (1993-2014) cover nearly comprehensive export flows by firm and destination at a very detailed level of product disaggregation (over 10,000 product categories). We complement these firm-level data sets with bilateral trade data from BACI (Gaulier and Zignago, 2010, updated to cover the period 1995-2013) at the product level (at a slightly higher level of aggregation than our French firm-level export data); and with country-level data (primarily GDP).

To disentangle the direction of causality between innovation and export performance, we construct a firm-level export demand variable following Mayer et al. (2016). This variable responds to aggregate conditions in a firm's export destinations but is exogenous to firm-level decisions (including the concurrent decisions for export-market participation and the forward looking innovation response). We show that: (i) firms that are initially more productive (closer to their sector's technology "frontier") strongly respond to a positive export demand shock by patenting more; (ii) this effect dissipates for firms further from the "frontier" and is reversed for a subset of initially less productive firms. These results confirm the predictions of the model for both a *market size* and a *competition effect* of the export shock. Our theoretical model highlights how an industry equilibrium with endogenous markups is key for this type of competition, which induces a reversal in the innovation response across 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., 2014)[2]. We contribute to this literature by uncovering a new -indirect- effect of market size on innovation working through competition differences within sectors;[3] and by testing the overall effect of export expansion on innovation using exhaustive firm-level data. Second, our paper relates to recent papers on import competition, innovation and productivity growth (see Bloom et al., 2016; Iacovone et al., 2011; Autor et al., 2016; Bombardini

---

[2]Akcigit et al. (2014) 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.

[3]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.

et al., 2017). These papers show that increased import competition induces firms to innovate more in order to escape competition as in Aghion et al. (2005).[4] Instead we look at how the export side of trade affects innovation and analyze a competition effect that varies across producers within a sector. We find very strong empirical confirmation for this type of differential response across firms.

Our paper is most closely related to Lileeva and Trefler (2010) and Bustos (2011), and the ensuing empirical literature connecting exports to innovation at the firm-level.[5] We add to those analyses in three main respects: first, by uncovering an indirect- competition-enhancing effect of increased export market size on innovation; second, by showing that this effect leads to heterogeneous innovation responses to the same market-size shock across firms (strongest for the frontier firms and turning negative for the firms furthest from the frontier); and third, by documenting this type of heterogeneous innovation response – including an innovation reversal for the least efficient firms in a sector.

The remaining part of the paper is organized as follows. Section 2 develops our model of export and innovation, and generates the prediction that the market size effect of a positive export shock, is stronger for more frontier firms. Section 3 briefly presents the data and shows some descriptive statistics on export and innovation. Section 4 describes our estimation methodology and presents our empirical results and Section 5 concludes.

2   Theory

We start with a highly parametrized version of our model in order to highlight the key interactions between market size and competition – leading to innovation reversals for the least productive firms. We then discuss how this result extends to much more general functional forms.

This model is essentially an open economy long-run version of the model in Mayer et al. (2014), augmented with innovation. 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, and 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

---

[4]Interestingly, in this paper we use *firm-level* competition data, whereas Aghion et al. (2005) as well as previous papers by Nickell (1996) and Blundell et al. (1999) regress innovation and/or productivity growth on *sectoral* measures of product market competition.

[5]In related work, Coelli et al. (2016) document the patenting response of firms in response to the Uruguay round of tariff levels.

representative consumer in country D with additively separable preferences with sub-utility:[6]

$$u(q_i) = \alpha q_i - \frac{\beta q_i^2}{2},$$

where $\alpha > 0$ and $\beta > 0$.

As this consumer makes no difference between a French or a locally produced variety, 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 the same innovation technology, which leads to similar innovation decisions, but the results can be easily extended without this assumption.

2.1 Consumer optimization

This representative consumer facing prices p_i solves:

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

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

$$p(q_i) = \frac{u'(q_i)}{\lambda} = \frac{\alpha - \beta q_i}{\lambda}, \tag{1}$$

where $\lambda = \int_0^M u'(q_i)q_i di > 0$ 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.

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 $L[p(q)q - cq]$. The corre-

---

[6]As we argue below, 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).

sponding first order condition yields

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

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.

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 $c_I k + \frac{1}{2} c_{I2} k^2$.[7] 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:

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

---

[7]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. As we discuss in further detail later, our analysis extends to the more general case where innovation affects (and responds to) changes in other destinations, as long as competition in other destinations does not respond to changes in market D (namely market size).

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 $\tilde{c}_{\min} - \varepsilon k(\tilde{c}_{\min}; \lambda) = 0$, 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 2 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, which ensures that the slope of the marginal cost is strictly larger than the slope of the marginal gain (otherwise firms may end up with infinite innovation), namely:

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

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:

$$\widehat{C}_I = \frac{1}{\lambda} \left(\alpha - \frac{2\beta c_I}{\varepsilon L} \right). \tag{3}$$

2.4 The direct impact of an increase in market size or competition on innovation

In the next section, we describe how an increase in market size L induces an endogenous increase in competition λ in the destination market D. To build intuition on the combined impact of these

Figure 2: Optimal innovation is higher for more efficient firms
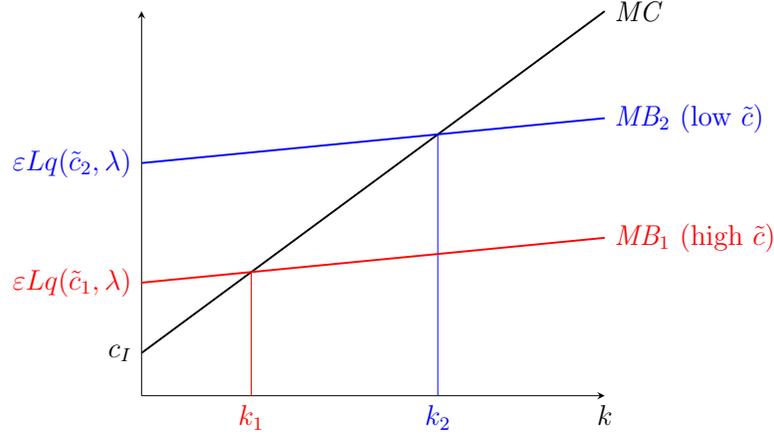
demand-side changes for innovation, we first consider these effects separately. We first analyze the direct effect of an increase in market size L, holding the competition level λ constant. At each firm's current innovation choice $k(\tilde{c}; \lambda)$, this triggers a proportional increase in firm output, and an upward shift in the marginal benefit of innovation, inducing all firms to increase innovation.

Figure 3 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 our unambiguous prediction of 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 3).

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 2). 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

Figure 3: Direct market size effect (increase in L)
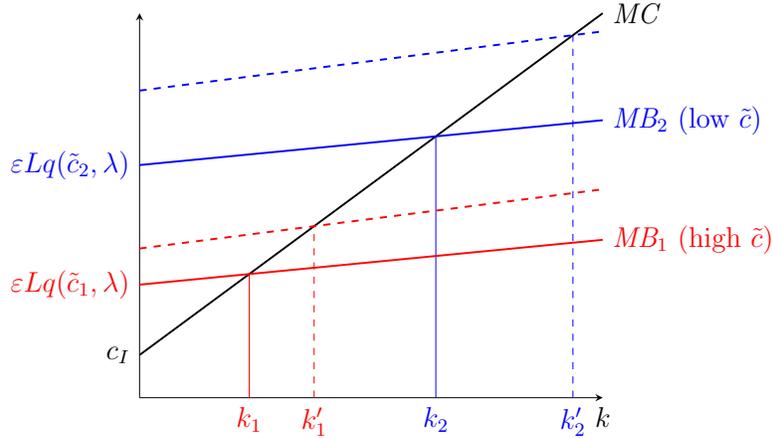
Figure 4: Competition effect (increase in λ)
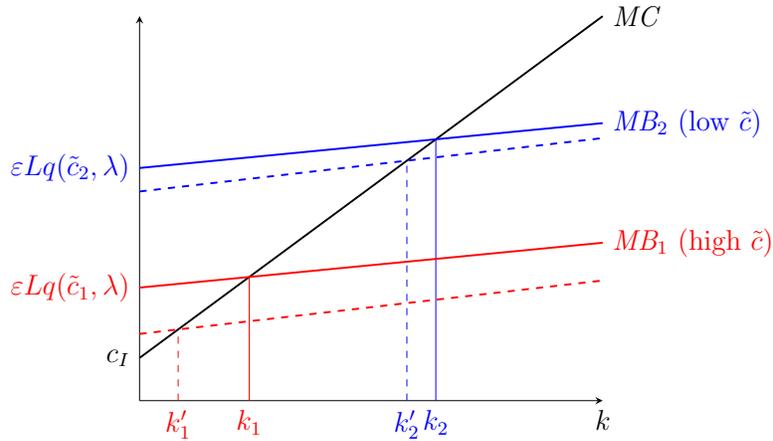
the most efficient firms (with baseline cost approaching $\tilde{c}_{\min}$), the negative impact of increased competition on innovation dissipates completely (see FOC).

Figure 4 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).[8] 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 3).

In the next subsection we endogenize the competition variable λ by introducing a free entry

---

[8]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.

condition.

2.5   The induced competition effect of increased market size

We now describe how the equilibrium competition level λ in country D is endogenously determined and show that λ increases with L. This equilibrium involves all the firms operating in D, including both the French exporters to D along with the domestic producers in D. However, in the long-run with entry of domestic producers into 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).[9] Since innovation is inherently forward looking, we focus on these long-run implications for competition.

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). In the long-run, 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* $(\lambda, \widehat{C}_D)$.

The proof is developed in A.1. It relies on the fact that (ZCP) is downward-sloping whereas (FE) is upward sloping in $(\widehat{C}_D, \lambda)$ space. For simplicity, we have abstracted from any export profits for the domestic firms. This is inconsequential for the qualitative predictions of our model that we emphasize below – so long as changes in country D (in particular its market size) do not affect

---

[9]This is not the case in the short-run, where the export supply to D affects both the competition level λ as well as the number of domestic producers.

the equilibrium competition levels in those export markets. Essentially, the key assumption is that market D is small relative to the size of the global export market.[10]

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

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. ■

2.6 The overall effect of increased market size on innovation

We now analyze the overall effect of an increase in market size L on innovation resulting from both the direct market size effect and the induced competition effect (increase in λ). Based on our previous discussion, we already know that the overall innovation response for the most efficient firms (those firms with the lowest baseline costs) will be positive. The reason is that the negative impact of the induced increase in competition on innovation dissipates for the most efficient firms, so that the positive direct market size effect must prevail for those firms.

As we move away from the technology frontier towards less efficient firms, the negative impact of increased competition on innovation strengthens and is no longer negligible. This leads to lower and lower innovation responses for firms as we move away from the technology frontier up the baseline cost scale. The only remaining question is whether the overall impact of an increase in L on innovation can indeed become negative for operating firms with sufficiently high baseline costs. In order to show that this is the case, we first point out that – in response to any changes in market conditions λ and L – innovation co-moves monotonically with firm output as well as its output at a *given* innovation level k.

---

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

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

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. This is the key property that we rely on in our comparative statics for the impact of market size and competition.

Proposition 3 *For a given innovating firm, innovation, output and output holding k fixed all move comonotonically in response to a change in the market size L.*

The proof is a straightforward extension of the first and second order conditions for optimal innovation depicted in Figure 2. Any change in market conditions L or λ leading to an increase in firm output $Q(\tilde{c}, k; \lambda)$ at constant innovation k results in an upward shift of the marginal benefit curve. This induces an increase in the optimal level of innovation (given the second order condition ensuring that the marginal benefit curve cuts the marginal cost curve from above). This innovation response further re-enforces the increase in firm output $Q(\tilde{c}, k(\tilde{c}; \lambda); \lambda)$.

We have just shown that any changes in market size or competition that increase a firm's output choice (at its current innovation level) will induce this firm to increase innovation – and vice-versa. In order to show that high cost firms reduce innovation in response to an increase in market size, we need only show that this market size increase induces those firms to reduce output at their current innovation levels (due to the increase in competition λ). This output response is given by:

$$\left. \frac{dQ}{dL} \right|_{k \text{ fixed}} = \frac{\partial Q}{\partial L} + \frac{\partial Q}{\partial \lambda} \frac{d\lambda}{dL} = q(\tilde{c} - \varepsilon k; \lambda) - L \frac{\tilde{c} - \varepsilon k}{2\beta} \frac{d\lambda}{dL}. \tag{4}$$

As a firm's baseline cost $\tilde{c}$ increases, the direct positive impact of the market size increase diminishes, whereas the negative impact due to increased competition strengthens. The positive impact shrinks to zero in the limit as a firm's marginal cost approaches the choke price α/λ that induces no consumer demand $q(\alpha/\lambda; \lambda) = 0$. Thus, the total output response must be negative for at least some high cost firms. So long as those firms innovate, their innovation response to a market size increase will be negative. Hence, whenever the set of innovating firms is broad enough – the innovation threshold $\widehat{C}_I$ is high enough – then there will be an innovation reversal within this set of firms. This will be ensured so long as the innovation cost c_I is low enough, along with a low fixed cost F to ensure that production is profitable for all innovating firms. More generally, in the Appendix we establish:

Proposition 4   a) *For any cost distributions Γ and Γ_D, there exists values of c_I and F such that some high cost firms reduce their innovation when the market size L increases.*

b) *For any values of c_I and F, there exists cost distributions Γ and Γ_D and a range of L such that some high cost firms reduce their innovation when the market size L increases within that range.*

Figure 5 illustrates this parameter configuration leading to an innovation reversal for the high cost firm in response to an increase in market size L. Lastly, we note that this innovation reversal

13

Figure 5: Overall response of innovation to an increase in market size L
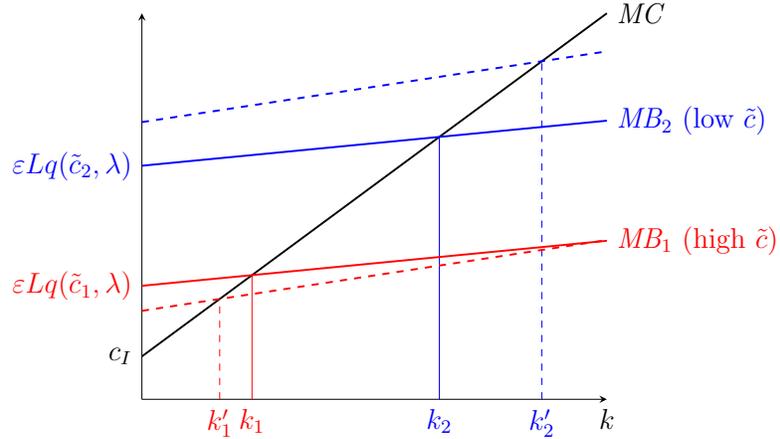
applies to French exporters in the same way that it does for domestic producers in D. It also applies to a much more general version of our model without specific functional form representations for the cost-saving benefit of innovation (replacing ε with a function of the baseline cost $\tilde{c}$ and innovation k); the cost of innovation (a convex function of innovation k); and for preferences. The key feature of the optimal innovation choice – linking the direction of the innovation response with the direction of the change in firm output at a *constant* innovation level – would continue to hold. Thus, an innovation reversal will occur for high baseline cost firms so long as a market size increase induces a reduction in output for those firms. As discussed in Mayer et al. (2014), this output reversal in response to market size increases will occur whenever residual demand satisfies Marshall's Second Law of Demand – whereby demand becomes more inelastic with consumption. Importantly, such a reversal cannot occur under C.E.S. preferences and exogenous markups. The competition increase induced by the endogenous response of the markups is a critical necessary ingredient for the prediction of both output – and therefore innovation – reversals.

3 Exporters and Innovators: data and descriptive statistics

In this section, we briefly present our datasets and show some descriptive evidence on the link between firms' innovation and exports. Further details about data construction can be found in Appendix B.

14

3.1 Data sources

We build a database covering all French firms and 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 HS8 product level (representing over 10,000 manufacturing products) by destination; (ii) Insee-DGFiP administrative fiscal datasets (`FICUS` and `FARE`), 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 all patent applications and on patents filed in some specific patent offices (see section 4.2 and Appendix B for details).

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 using 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 (for example this is the case of the OECD Harmonized Assignee Name database, 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. (in progress) 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 B.4, provides significant performance improvements relative to previous methods used in the empirical patent literature: 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.

Finally, 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 over 1995-2012.

Sample restrictions

Although our main firm-level administrative data source is comprehensive, with more than 47.1 million observations spanning over 7.3 million different firms from 1995 to 2012, we restrict our data sample for several reasons. The first is due to the matching with patent data mentioned above, which is most complete for firms above 10 employees. We therefore impose this size restriction, which drops a large number of firms but a relatively small share of aggregate French production:

17.1% of employment, 15.6% of sales, and 13.6% of exports (predominantly within EU exports). Second, 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. This further reduces our sample from 1.7 million to 835,000 firms. Yet, the bulk of aggregate employment (74.2%), sales (77.7%), and exports (77.2%) remain in our dataset after imposing these restrictions. These remaining firms are matched with an average of 27,640 patents per year in PATSTAT. Lastly, since our detailed customs trade data only covers goods trade (and not services), we will further restrict our sample to the manufacturing sector for most of our analysis.[11] This reduces our working sample to 105,000 firms. Nevertheless the bulk of French aggregate exports and innovation are still concentrated in manufacturing as only 20.6% of aggregate exports and 33% of patents are recorded outside of this sector.

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. Multinational groups tend to break the relationship between export shocks and patenting since these groups may locate their R&D activities in different countries from 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. In this case, 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.

3.2   Sector breakdown and skewness

Table 1 shows the breakdown of those firms across sectors, along with their average employment, exports, and patents (per firm) for 2007.[12] As has been widely reported in the empirical literature on micro-level trade patterns, many firms are only occasional exporters: they export in some years, but not in others. This pattern is even more pronounced for innovation: even firms with substantial ongoing R&D operations do not typically file patent applications year in and year out. We therefore use the broadest possible cross-year definition to classify firms as exporters and innovators. We label a firm as an *exporter* if it has exported at least once between 1993-2012;

---

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

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

and as an *innovator* if it has filed at least one patent application between 1995-2012.[13] Thus, our reported export participation rates in Table 1 are higher than in other studies. However, even with this broadest classification, innovators represent only a small minority of manufacturing firms. For comparison, Table 1 also reports the share of exporters and innovators based on the more standard definition of current year (2007 for this table) exporting or patenting activity – shown in parentheses.

Table 1: Exports and innovation in the manufacturing sector

Sector	Description	Firms	Emp	Export	% Exporter	Patents	% Innov.
10	Food products	8,814	43	1,847	41 (26)	20	3 (0)
11	Beverages	1,463	47	5,974	80 (59)	11	2 (0)
13	Textiles	1,802	37	2,335	86 (63)	91	12 (2)
14	Wearing apparel	1,558	39	2,577	80 (59)	49	5 (1)
15	Leather	492	56	2,566	85 (59)	38	9 (2)
16	Wood	2,432	29	790	64 (36)	16	5 (1)
17	Paper	2,950	44	2,056	79 (44)	86	7 (1)
18	Printing	842	24	167	53 (20)	5	3 (0)
19	Coke	171	225	75,957	92 (69)	3,061	24 (9)
20	Chemicals	1,229	116	17,607	94 (79)	1,992	21 (6)
21	Basic pharmaceutical	357	288	42,065	96 (82)	2,808	35 (13)
22	Rubber and plastic	2,745	80	3,820	86 (64)	339	21 (6)
23	Other non-metallic	2,158	63	2,320	65 (38)	272	11 (2)
24	Basic metals	1,648	80	12,487	66 (44)	147	12 (3)
25	Fabricated metal	8,392	36	1,125	67 (40)	82	9 (2)
26	Computer and electronic	3,511	85	7,620	73 (54)	769	23 (8)
27	Electrical equipment	447	106	8,812	91 (70)	1,764	26 (8)
28	Machinery and equipment	4,668	80	8,252	79 (58)	558	24 (7)
29	Motor vehicles	791	61	2,549	79 (47)	173	15 (3)
30	Other transport equipment	558	215	54,911	83 (56)	2,293	18 (7)
31	Furniture	1,146	34	598	67 (36)	14	7 (1)
32	Other manufacturing	1,017	41	2,472	82 (58)	321	12 (3)
33	Repair of machinery	3,430	28	302	54 (23)	25	6 (1)
	Aggregate manufacturing	52,621	57	4,654	68 (44)	288	11 (3)

Notes: This table presents the number of firms, average employment, average export (in thousands of Euros), average number of patents (in thousands), and the shares of exporters and innovators (cross-year definitions). The shares in parentheses are calculated based on current year export participation or patent filing. Data are for 2007.

Even within the minority set of innovators, patenting activity is extremely skewed. This is

---

[13]The initial year for both ranges do not coincide in order to reflect our subsequent empirical analyses. We will use prior years of export data to construct exogenous export share weights (see section 4.1 for more details).

Figure 6: Lorenz curves - patents are more concentrated than exports, sales and employment

(a) Top 5 percentiles　　　　　　　　　　(b) Whole distribution
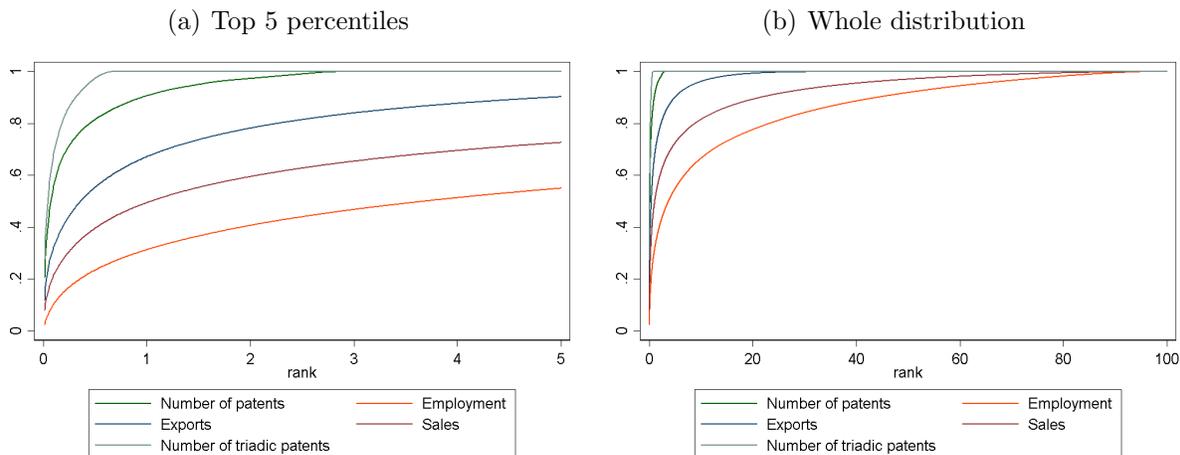
Notes: Lorenz curves plot cumulative distribution function for patents, triadic patents, employment, export and sales. Data are for manufacturing firms.

clearly visible in Figure 6, which plots the Lorenz curve for patents and triadic patent families at manufacturing firms in 2007, along with the Lorenz curves for exports, sales, and employment. Figure 6 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., 2016): 1% of firms account for 67% of aggregate exports in 2007, whereas the top 1% of firms based on total size account for 50% of sales (ranked by sales) and 31% of employment (ranked by employment). But Figure 6 also shows that patenting is even significantly more skewed than exporting: 1% of all firms account for 91% of patents in 2007. And less than 1% of firms own all the triadic families - i.e. patent families which include patents filed in Asia, Europe and in the USA, see Section 4.2). Indeed fewer than 2.9% of manufacturing firms have patented in 2007. This fraction is significantly smaller than our previously reported 11% share of innovators in Table 1 measured across our full sample years. Similarly, only 44% of manufacturing firms report any exporting activity for 2007 compared to a 68% share when exporting is measured across our full sample years.

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.

18

3.3 The innovation-export nexus

Looking across our sample years (1995-2012), Table 2 reports different size-related performance measures (averages per firm) based on their exporter and innovator classification. This table confirms the well-documented size differential in favor of exporters. However, several new salient features regarding innovators pop-out from this table: 1) Innovating firms are massively concentrated among exporters: only 5% of innovators do not report any exporting; 2) non-exporting innovators do not look very different than 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;[14] 3) these same measures of firm size differ markedly between innovators and non-innovators among exporters: innovators employ on average 4.5 times more workers and produce 7-8 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.

Table 2: Exporters and innovators are bigger

	Non-exporter		Exporter		Total
	Non-innovator	Innovator	Non-innovator	Innovator	
Firms	45,789	392	52,043	6,893	105,117
Employment	17	20	51	233	59
Sales	2,147	2,460	11,499	69,200	13,982
Value Added	639	883	2,733	16,091	3,332
Age	14	15	20	22	18
Export	0	0	2,463	22,875	3,622
Countries	0	0	5	18	5
Patents	0	0.2	0	2.6	0.2

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 thousand of 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.

In order to compare exporters to non-exporters and innovators to non-innovators, within specific

---

[14]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.

groups, we compute export and innovation premia. Consider first the exporter premia reported in the top panel of Table 3. 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 1); and Column 3 controls for firm employment, in addition to the sector fixed effect. Since we are using a broad 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. (2016) 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 after controlling for firm employment (within sectors).

Table 3: Export and innovation premia

Panel 1: Premium for being an exporter (among all manufacturing firms)

	(1)	(2)	(3)	Obs.	Firms
log Employment	0.854	0.767		937,050	91,563
log Sales	1.616	1.478	0.417	979,413	104,368
log Wage	0.132	0.097	0.109	935,489	91,525
log Value Added Per Worker	0.217	0.171	0.175	923,535	90,876

Panel 2: Premium for being an innovator (among all exporting manufacturing firms)

	(1)	(2)	(3)	Obs.	Firms
log Employment	1.043	0.997		645,522	58,121
log Sales	1.284	1.239	0.197	656,218	58,803
log Wage	0.125	0.095	0.109	644,533	58,104
log Value Added Per Worker	0.203	0.173	0.179	635,144	57,720
log Export Sales (Current period exporters)	2.018	1.911	0.806	433,456	56,509
Number of destination countries	13	11	7	663,004	58,936

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 3-digit sector fixed effect, column 3 adds a control for the log of employment to column 2. 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.

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 1 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 80th percentile of the export distribution, 30% of the firms have some patenting experience. And the increase in the share of innovative firms with the percentile of the export distribution

21

is highly convex. Above the 95th percentile of the export distribution, a majority of firms are innovators; in the top percentile, 68% of the firms are innovators. Those firms in the top export percentile account for 41% of the aggregate share of French patents.

4 Empirical framework and results

4.1 Identification strategy: firm-level export demand shocks

We have just documented the strong correlation between exports and innovation in the cross-section of French manufacturing firms. However, this correlation does not shed light on 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 (as we emphasize in our theoretical comparative statics). 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. (2016) in building such a measure of exogenous export demand.

To construct these export demand shocks, consider a French exporter f who exports a product s to destination j at an initial date t_0. Let M_{jst} denote the aggregate import flow in product s into country j from all countries except France at time $t > t_0$. M_{jst} reflects the size of the (s, j) export market at time t. We then sum M_{jst} across destinations j and products s weighted by the relative importance of 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.[15] 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 on any export shock on total production.)

More precisely, let X_{fjst_0} denote firm f's export flow to market (j, s) at time t_0. This is the firm's first observed export year in our sample.[16] The export demand shock for firm f at time t is

---

[15]Another distinct potential source of endogeneity may arise in markets where a French firm has a dominant position. In this case, imports into those markets may respond to this firm's decisions (including innovation). We address this issue in Section 4.5.3.

[16]We consider this firm to be an exporter only if we observe positive exports in both customs data (so we can

then:

$$D_{ft}^{M_s} = \frac{X_{ft_0}^*}{S_{ft_0}^*} \sum_{j,s} \frac{X_{fjst_0}}{X_{ft_0}} \log M_{jst}, \tag{5}$$

where the weight $(X_{ft_0}^*/S_{ft_0}^*)(X_{fjst_0}/X_{ft_0})$ represents firm f's initial share of sales of product s to destination j and $X_{ft_0} = \sum_{j,s} X_{fjst_0}$ represents the firm's total exports. The asterisks on firm f's initial export intensity $X_{ft_0}^*/S_{ft_0}^*$ indicate that the underlying data for total exports $X_{ft_0}^*$ and sales $S_{ft_0}^*$ come from the production data (as opposed to customs data which we use to calculate the destination/product specific market shares).[17]

We note that the time variation in our demand shock $D_{ft}^{M_s}$ only stems from the world export flow M_{jst} and not the firm-level weights, which are fixed 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 from our demand shock.

We will also experiment with an alternate measure of this demand shock using more aggregated data (across products). We thus aggregate both the world and the firm's export shares at the 3-digit ISIC level:

$$D_{ft}^{M_I} = \frac{X_{ft_0}^*}{S_{ft_0}^*} \sum_{j,I} \frac{X_{fjIt_0}}{X_{ft_0}} \log M_{jIt},$$

where $M_{jIt} = \sum_{s \in I} M_{jst}$ measures aggregate imports (excluding France) in destination j for industry I, and $X_{fjIt_0} = \sum_{s \in I} X_{fjst_0}$ is the associated firm-level exports for that industry-destination pair in the initial year t_0. This measure will no longer reflect the cross-firm variation at the detailed product level. However, it captures some potential spillovers across related products in the construction of the demand shock (an increase in export demand for closely related products may induce a firm to direct innovation towards these related products).

Constructing these export demand shocks generates outliers for a few firms that export a small set of products (often highly specialized) to small destinations (such as yachts to Seychelles and Maldives). We therefore trim our sample by removing firms with extreme changes in our con-

---

calculate destination market shares) as well as production data (so we can calculate export intensity).

[17]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 custom data (see Appendix B for more details).

structed export demand. We regress this demand shock on a firm fixed effect and trim observations with a residual that is above/below the 97.5th/2.5th percentile. That is, observations with the largest variations in their export demand shock (relative to their firm mean) are eliminated from our sample.[18]

Before turning to the impact of the export demand shock on the firms' innovation response, we note that this demand shock has a very significant explanatory power for a firm's total export response (see Table B1, which uses a similar estimating strategy as the one we develop in the next section for innovation).

4.2 Innovation measures

We consider three main measures of innovation for each French firm f in any year t. The first measure counts all patent applications filed by the firm during year t. To better reflect the firm's individual contribution, we use fractional counts for patents shared with other firms (so that a patent filed with 2 other applicants counts as a third). The second measure selects higher quality innovations by counting triadic families of patents: when the same innovation (within a patent family) is filed at three different patent offices in Europe (the European Patent Office, EPO), the United States (U.S. Patent Office, USPTO) and at least one of the major Asian economies (Japan, China or Korea).[19] The rationale behind this measure is that the best ideas are more likely to be protected in the three main economic regions worldwide. Another feature of this triadic patent measure is that it is more immune to geographical and institutional biases, i.e. to the possibility that different patent offices would differ with regard to quality standards and/or the enforcement of Intellectual Property Rights (see Park, 2008). When aggregating triadic patent families by firm, we also use fractional counts to reflect a given firm's contribution to the patent family. The third measure counts (fractional) patent applications in Europe (at the EPO) – the main domestic market for French firms. This measure is thus less restrictive than the triadic measure (which requires filing in 3 regions including Europe); but it provides a homogeneous institutional framework for the assessment of intellectual property. Other innovation measures yield similar results (see section 4.5.1). Appendix B provides additional details on the construction of our patent measures.

---

[18]The incidence of these outliers decreases as we aggregate the trade flows from products to industries. We have experimented with different threshold cutoffs in the 1-5% range. Our qualitative results are robust to these changes (see Table B3 in Appendix C).

[19]This definition is slightly broader than the definition of triadic patent families used elsewhere in the literature as we consider China and Korea in addition to Japan.

4.3 Main estimation strategy

Our baseline regression seeks to capture the impact of the exogenous demand shock on a firm's innovation response. We expect this innovation response to incorporate the cumulated effects of the trade shocks over time. We thus do not use year-to-year differences in the demand shock (first differences) and use that shock as constructed (in "levels") along with a firm fixed-effect in order to capture its within-firm variation (relative to the firm mean over time). We also add sector-time fixed-effects to remove any time variation that is common to the firm's sector. We restrict our analysis to the subset of innovating firms (i.e. firms with at least one patent since 1994).[20]

To capture the indirect competition effect of an export demand shock (which varies with a firm's initial productivity level), we add an interaction between the demand shock and the firm's initial productivity. Just as we did with the firm-level export shares, we only use our initial year t_0 to generate a productivity measure that does not subsequently vary over time $t > t_0$.[21] We assign a 0-9 productivity index d_f to all firms based on their labor productivity (value-added per worker) decile in year t_0 within their 2-digit sector.[22]

The left-hand panel of Figure 7 shows a bin-scatter of our main patent measure (all applications in year t) against the firm's export demand $D_{ft}^{M_s}$ for the same year – absorbing a firm fixed-effect. This clearly shows that there is a very strong correlation between changes in export demand and changes in patent flows (within firms over time); and that a linear relationship provides a very good functional form fit for that correlation. We thus use a linear OLS specification as our baseline regression equation to quantitatively assess this relationship:

$$Y_{ft} = \alpha\ D_{ft} + \beta\ D_{ft} * d_f + \chi_{s,t} + \chi_f + \varepsilon_{ft}, \tag{6}$$

where Y_{ft} is one of our innovation measures based on the flow of patent applications during year t by firm f and D_{ft} is one of our export demand proxies ($D_{ft}^{M_s}$ or $D_{ft}^{M_I}$) for that same year t. Those proxies are constructed so that they are exogenous to the firm's decision in year $t > t_0$.[23]

---

[20]Investigating the entry margin into - or the exit margin out of - the set of innovating firms is also an important topic, but we leave it for further research.

[21]Recall that t_0 is the first year since 1994 in which the firm reports positive exports. This year is equal to 1994 for about 50% of the firms and is always removed from the estimation.

[22]When a firm belongs to the manufacturing sector for a subset of our sample years, we only use those years in our estimation. For a firm not in our manufacturing sample at t_0, we compute its productivity decile within its previous sector at t_0.

[23]Serial correlation in the innovation shocks could induce some correlation between a firm's export structure a time t_0 (which we use to construct our export demand shock for year t) and the subsequent innovation shock in year t. First, we note that a time-persistent effect in the innovation shock will be captured by the firm fixed-effect. To ensure that our results are not driven by transitory serial correlation, we also experiment with dropping additional

The $\chi_{s,t}$ and χ_f capture the sector-time and firm fixed effects and ε_{ft} is an error term. In most of our regressions, we estimate coefficients and standard errors using a Newey-West estimator for the covariance matrix which allow for heteroskedasticity and autocorrelation in the error terms with a maximum lag of 5 years (see Newey and West, 1987 and Wooldridge, 2010).[24]

While the linear panel fixed effect model is our preferred specification, we also show results from a Poisson model in Table B6 using the following specification:

$$\ln \left(\mathbb{E} \left[\tilde{Y}_{ft} \right] \right) = a \ D_{ft} + b \ D_{ft} * d_f + \chi_{s,t} + \chi_f, \tag{7}$$

where $\tilde{Y}_{ft}$ is a measure of total patent applications that does not use fractional counts and is therefore an integer. We estimate a and b and corresponding standard errors using maximum likelihood.

4.4 Baseline results

In the model, an increase in market size induces both a positive market size effect and a counteracting competition effect which is most pronounced (and potentially dominant) for the least productive firms (Section 2.6). In the data this feature clearly stands out. Graphically the right-hand side panel in Figure 7 shows that the number of patents increases much more with export demand in firms initially more productive (above the median productivity). Quantitatively Table 4 reports the results from the baseline regression (6). A positive export shock reduces innovation for firms in the lowest productivity decile ($d_f = 0$); the export shock's impact on patenting increases with initial productivity and turns positive at productivity around the third or fourth decile: firms initially more efficient increase more their innovation when they face a positive export shock. This pattern holds for all three patent and two demand shock measures.

Our coefficients from column 1 using the product-level demand shock imply the following quantitative response in the number of patents to the average demand shock: for a firm in the lowest productivity decile, the number of patents (relative to the firm average) is 3.3 patents lower than the sector average; and each additional productivity decile increases this patent response by 0.96 patent. Thus the response of a firm in the top productivity decile amounts to 8.7 more

---

years of data following t_0. Instead of starting our sample at $t_0 + 1$, we have tried starting at $t_0 + 2$ or $t_0 + 3$. We report those alternative specifications in Table B4, which are qualitatively very similar to our baseline results (even though the sample size is reduced in our critical time dimension).

[24]We also show robustness of our baseline results using standard errors clustered by firm (see Table B5). However, given the relatively long time dimension of our sample, we find this latter specification too conservative and prefer the Newey-West estimator. As argued by Abadie et al. (2017), in such fixed-effect models, clustering standard error only matters if we expect heterogeneity in the treatment effect.

Figure 7: Patenting increases more with demand for initially more efficient firms
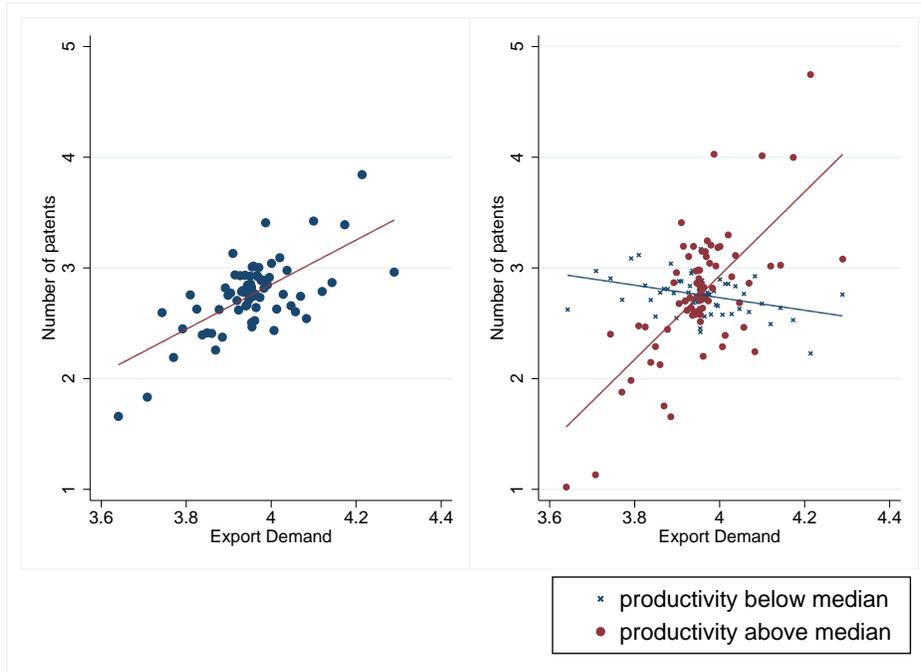
Notes: Each dot represents the quantile-average of the residual of the number of patent on a firm fixed effect (y-axis) against the quantile-average of the residual of the demand variable $D_{ft}^{M_s}$ on a firm fixed effect (x-axis), for all firms (left-hand side panel) or for both firms below (blue) or above (red) the productivity median at t_0 (right-hand side panel).

patents than the sector average (still relative to the firm average). Our coefficients when using the industry demand shocks (column 2) yield very similar results.

4.5 Robustness analysis

Our main finding – that initially more productive firms respond to an export demand shock by innovating relatively more – is robust to various alternative specifications. In this subsection, we show the robustness of our main results to: (i) considering other patent indicators; (ii) controlling for firm specific characteristics; (iii) excluding dominant firms in a destination market; (iv) considering alternative measures and heterogeneous effects for a firm's proximity to frontier; and (v) controlling for pre-trends with sector-decile specific time trends.

4.5.1 Other patent indicators

There are many alternative ways of aggregating patent counts, which yield different measurements of a firm's innovation output. In Table 5 we consider 6 alternative measures of Y_{ft} and

Table 4: BASELINE RESULTS

Dependent variable	All patents		Triadic patents		EPO patents	
Demand measure	$D_{ft}^{M_s}$ (1)	$D_{ft}^{M_I}$ (2)	$D_{ft}^{M_s}$ (3)	$D_{ft}^{M_I}$ (4)	$D_{ft}^{M_s}$ (5)	$D_{ft}^{M_I}$ (6)
Demand	-3.260***	-2.578**	-0.265***	-0.224***	-0.368***	-0.447***
	(1.014)	(1.056)	(0.0786)	(0.0820)	(0.123)	(0.143)
Decile × Demand	0.960***	0.909***	0.0859***	0.0862***	0.125***	0.114***
	(0.255)	(0.304)	(0.0195)	(0.024)	(0.029)	(0.039)
Cutoff decile	4	3	4	3	3	4
Nb of observation	77,901	77,918	77,901	77,918	77,901	77,918
R^2	0.897	0.888	0.759	0.747	0.849	0.836

Notes: This table presents regression results of an OLS estimation of equation 6. Sample includes manufacturing firms with at least one patent in 1995-2012 and for which we can compute the export demand shock (see section 4.1). Cutoff decile corresponds to the first value of d_f for which the overall effect becomes positive. Coefficients and standard errors are obtained using a panel fixed-effect estimator. Autocorrelation and heteroskedasticity robust standard errors using the Newey-West variance estimator with a bandwidth set to 5 years. ***, ** and * respectively indicate 0.01, 0.05 and 0.1 levels of significance.

run our baseline model (6) with these new dependent variables. (1) We first restrict our triadic measure to a dyadic one by only considering the number of patent families at the USPTO and at the EPO (dropping Asia); (2) We then expand the definition by counting all patent families (instead of individual patents as in our baseline); (3) We return to a count of individual patents but restrict this to EPO patent applications; (4) We count only the number of priority patent applications; (5) We drop the construction of fractional patent counts ("raw" number of patent); and (6) We only count the number of patent applications that will ultimately be granted in any patent office.[25] The results for all these alternative patent measures are similar to those in our baseline Table 4: the export demand shock has a positive effect on the corresponding measure of innovation in frontier firms but has a negative effect on innovation in lagging firms.

4.5.2 Direct control for firm size

A firm experiencing an increase (decrease) in market size which is not initially related to innovation, may still respond to it by innovating and exporting more (less). We do not control for firm size directly in our baseline as our theoretical model suggests that changes in size that are driven by the export shocks should be incorporated into our measure of the impact of exports on innovation. However, in order to eliminate a direct impact of firm scale on our estimates, we

---

[25]See Appendix B for additional details on the construction of all these indicators.

Table 5: ALTERNATIVE WAYS OF COUNTING PATENTS

Dependent variable	Dyadic	Families	EPO*	Prior	Nb Appln	Granted
Demand Measure	$D_{ft}^{M_s}$	$D_{ft}^{M_s}$	$D_{ft}^{M_s}$	$D_{ft}^{M_s}$	$D_{ft}^{M_s}$	$D_{ft}^{M_s}$
	(1)	(2)	(3)	(4)	(5)	(6)
Demand	-0.233**	-0.936***	-0.883***	-1.364***	-4.732***	-1.197*
	(0.098)	(0.330)	(0.222)	(0.502)	(1.253)	(0.629)
Decile × Demand	0.072***	0.271***	0.245***	0.408***	1.375***	0.363**
	(0.024)	(0.087)	(0.050)	(0.133)	(0.315)	(0.154)
Nb of observation	77,901	77,901	77,901	77,901	77,901	77,901
R^2	0.819	0.802	0.848	0.830	0.881	0.900

Notes: This table presents regression results of an OLS estimation of equation 6. Sample includes manufacturing firms with at least one patent in 1995-2012 and for which we can compute the export demand shock (see section 4.1). See Appendix B for a complete definition of the different indicators used in this Table. Coefficients and standard errors are obtained using a panel fixed-effect estimator. Autocorrelation and heteroskedasticity robust standard errors using the Newey-West variance estimator with a bandwidth set to 5 years. ***, ** and * respectively indicate 0.01, 0.05 and 0.1 levels of significance.

now include such a control for firm size (at time t). We select different empirical measures of size from the production data: employment, raw materials, net and gross capital stock, and sales. The corresponding regression results are reported in Table 6. They clearly show that a direct control for size does not affect our previously reported baseline coefficients (reported again in column 1): the coefficients remain virtually unchanged. Similar results are obtained when using our two other main measures of innovation (triadic and EPO families) and are available upon request.

4.5.3 Excluding markets where a firm is a leader

When a French firm has a dominant market share in a market (j, s), then the world exports M_{jst} may be correlated with the firm's exports X_{fjst} (even though French exports are excluded from the construction of the world exports M_{jst}) as those other Foreign exporters may respond to actions taken by this French firm. To investigate this further, we drop from our dataset the markets (j, s) (in all years) for a firm f whenever its export sales in market (j, s) are above 10% of world exports (including France) into this market for any given year. These instances represent 6.7% of the customs data observations and predominantly reflect firms exporting to African destinations. The results are reported in Table B7 in the Online Appendix; and once again leave our baseline results virtually unchanged.

Table 6: CONTROL FOR FIRM SIZE

Dependent variable	Number of patents					
Demand measure	$D_{ft}^{M_s}$ (1)	$D_{ft}^{M_s}$ (2)	$D_{ft}^{M_s}$ (3)	$D_{ft}^{M_s}$ (4)	$D_{ft}^{M_s}$ (5)	$D_{ft}^{M_s}$ (6)
Demand	-3.260***	-3.434***	-3.285***	-3.218***	-2.799***	-3.280***
	(1.014)	(1.039)	(1.024)	(1.017)	(1.004)	(1.012)
Decile × Demand	0.960***	0.970***	0.954***	0.960***	0.874***	0.976***
	(0.255)	(0.263)	(0.257)	(0.256)	(0.258)	(0.255)
Size		0.696***	1.257***	2.007***	2.007***	1.227***
		(0.114)	(0.096)	(0.201)	(0.350)	(0.187)
Nb of observation	77,901	76,236	76,678	76,860	77,240	77,605
R^2	0.897	0.900	0.898	0.898	0.898	0.898

Notes: This table presents regression results of an OLS estimation of equation 6 where we add a control for firm size. Column 1 uses no control, column 2 controls for the log of raw material inputs, column 3 (resp. 4) controls for the log of net (resp. gross) capital stock, column 5 controls for the log of employment and column 6 controls for the log of sales (we obtain similar results when we control jointly by any subset of these covariates). Sample includes manufacturing firms with at least one patent in 1995-2012 and for which we can compute the export demand shock (see section 4.1). Coefficients and standard errors are obtained using a panel fixed-effect estimator. Autocorrelation and heteroskedasticity robust standard errors using the Newey-West variance estimator with a bandwidth set to 5 years. ***, ** and * respectively indicate 0.01, 0.05 and 0.1 levels of significance.

4.5.4 Other measures of proximity to frontier

So far, our model uses initial productivity deciles to measure a firm's proximity to its sector's technology frontier. We now consider alternative measures for this proximity. Table 7 shows that our baseline results (1) are robust to (2) measuring productivity deciles using sales instead of value added per worker; as well as (3)-(6) measuring proximity to frontier using a binary threshold for initial productivity within sector set at 50%, 75%, 90% and 95%. Those results highlight how the impact of the export shock is magnified for firms very close to the frontier (at the very top of the distribution of initial productivity).[26]

In light of these results, which suggest that the positive effect of export demand on innovation may be concentrated among the most productive firms, we allow our key interaction coefficient β in equation (6) to vary with the productivity decile d_f. Our estimating equation then becomes:

$$Y_{ft} = \sum_{d=0}^{9} \left[\beta_d\ D_{ft} * \mathbb{1}_{d_f=d} \right] + \chi_{s,t} + \chi_f + \varepsilon_{ft}, \tag{8}$$

---

[26]Similar results are obtained when using Triadic and EPO families. Tables are available upon request to the authors.

Table 7: ALTERNATIVE DEFINITION OF FRONTIER

Dependent variable	Number of patents					
Demand measure	$D_{ft}^{M_s}$ (1)	$D_{ft}^{M_s}$ (2)	$D_{ft}^{M_s}$ (3)	$D_{ft}^{M_s}$ (4)	$D_{ft}^{M_s}$ (5)	$D_{ft}^{M_s}$ (6)
Demand	-3.260***	-1.316	-1.042**	-0.0904	0.520	0.578
	(1.014)	(0.890)	(0.501)	(0.680)	(0.699)	(0.671)
Interaction	0.960***	0.578**	4.438***	5.456***	8.375**	16.63**
	(0.255)	(0.235)	(1.117)	(1.835)	(3.570)	(7.276)
Nb of observation	77,901	77,901	77,901	77,901	77,901	77,901
R^2	0.905	0.905	0.905	0.905	0.905	0.906

Notes: This table presents regression results of an OLS estimation of equation 6. Sample includes manufacturing firms with at least one patent in 1995-2012 and for which we can compute the export demand shock (see section 4.1) Column (1) is our baseline model, column (2) defines productivity using sales instead of value added, columns (3) to (6) no longer construct decile groups but use a dummy variable for being above the sectoral 50^{th}, 75^{th}, 90^{th} and 95^{th} percentile of the initial productivity distribution. This dummy is interacted with the demand variable. Autocorrelation and heteroskedasticity robust standard errors using the Newey-West variance estimator with a bandwidth set to 5 years. ***, ** and * respectively indicate 0.01, 0.05 and 0.1 levels of significance.

where $\mathbb{1}_{d_f=d}$ are indicator dummies for each productivity decile. This new specification allows us to relax the assumption that there is a constant slope shift across decile groups and to account for potential non linear effects of our export demand shock variable for different levels of productivity. Coefficients and corresponding confidence intervals are graphically reported in Figure 8. They show that the assumption of a constant effect across decile group is a good approximation – with the possible exception of the top-decile where the effect is magnified relative to the linear trend (this confirms the results from Table 7). This figure also highlights that the effect of the demand shock is clearly negative for some of the lowest productivity deciles; and that this effect turns positive for all deciles above the median (deciles 5 through 9). The confidence intervals remain relatively wide since those decile indicators induce substantial collinearity.

4.5.5 Controlling for sector-decile specific time trends

To deal with the possibility that firms in different productivity deciles and different sectors may evolve differently over time and in particular may follow different innovation trends independently from the export demand shock, we add dummies for all year-productivity decile bins. Those results are reported in the odd-numbered columns (1), (3), (5) in Table 8. In addition, we also consider dummies for the triple interaction of all year-sector-productivity decile bins. Those results are reported in the even-numbered columns (2), (4), (6). This sectoral classification is the same

Figure 8: Response to Demand by decile of productivity
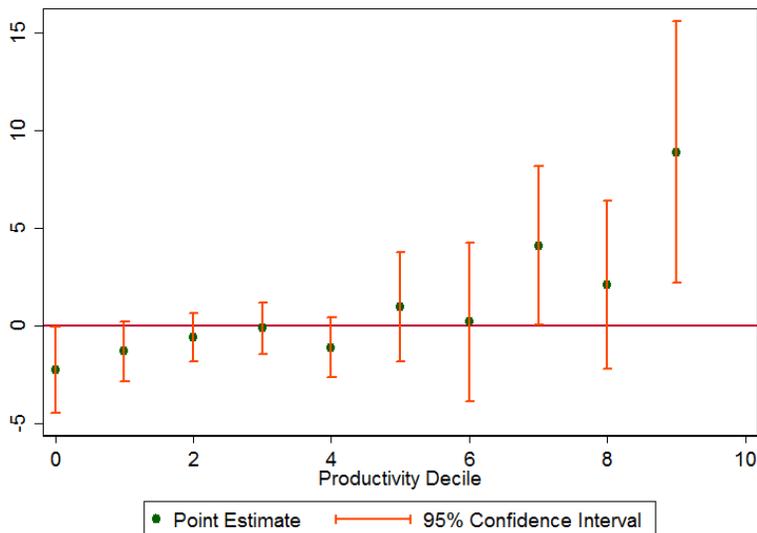
that we use to construct the productivity deciles – which are now orthogonal to any sector-level changes over time. Table 8 highlights that these controls do not change the message from our baseline results (though the negative impact for low productivity firms is no longer significant for the case of EPO patent families).

4.6 Extensions

In this last section, we extend our empirical analysis in different directions. First, we categorize export destinations based on a separate measure of competition in those destinations and show that the competition effect is most salient in high competition export destinations. We then show how we obtain similar results with an alternate specification based on long time-differences (splitting our sample years into two intervals). This provides an alternate way of capturing the slow-moving changes in the variables of interest (changes in export demand and the associated innovation response). And third, we contrast our main results using measures of innovation output (patents) with results obtained with measures of innovation inputs (R&D inputs) – which are available for a subsample of firms in our sample.

Table 8: Decile group specific evolution

Dependent variable	All patents		Triadic patents		EPO patents	
Demand measure	$D_{ft}^{M_s}$ (1)	$D_{ft}^{M_s}$ (2)	$D_{ft}^{M_s}$ (3)	$D_{ft}^{M_s}$ (4)	$D_{ft}^{M_s}$ (5)	$D_{ft}^{M_s}$ (6)
Demand	-2.876***	-2.135**	-0.179**	-0.135**	-0.168	-0.0755
	(0.975)	(0.868)	(0.0716)	(0.0638)	(0.121)	(0.120)
Decile × Demand	0.900***	0.727***	0.0689***	0.0515***	0.0833***	0.0614**
	(0.267)	(0.251)	(0.0200)	(0.0188)	(0.0316)	(0.0309)
Nb of observations	77,901	77,744	77,901	77,744	77,901	77,744
R^2	0.897	0.899	0.759	0.763	0.849	0.849

Notes: This table presents regression results of an OLS estimation of equation 6, to which a productivity decile × year fixed effect is added (columns 1, 3 and 5) or replacing the sector × time fixed effect with a sector × productivity decile × time fixed effect (columns 2, 4 and 6). Sample includes manufacturing firms with at least one patent in 1995-2012 and for which we can compute the export demand shock (see section 4.1). Coefficients and standard errors are obtained using a panel fixed-effect estimator. Autocorrelation and heteroskedasticity robust standard errors using the Newey-West variance estimator with a bandwidth set to 5 years. ***, ** and * respectively indicate 0.01, 0.05 and 0.1 levels of significance.

4.6.1 Direct competition effect

We now highlight how the skewed response of innovation to the export shock is driven by the induced competition effect (a demand-side explanation) – as opposed to supply-side effects (such as skewness in the costs or returns to R&D). Towards this end, we use an index of market competition from Djankov et al. (2002) to separate all French export market destinations into high- and low- competition categories. These data on competition levels across countries are now regularly updated and reported in the World Bank's "Doing Business" database. There are several different measures for competition; we use the index reflecting the ease of opening up a business in a country. This generates a time-invariant index by destination on a 0-100 scale.[27]

We then separate destinations into high (H, above median) and low (L, below median) competition according to this index and construct two separate export demand shock measures for those two categories:

$$D_{ft,H}^{M_s} = \frac{X_{ft_0}^*}{S_{ft_0}^*} \sum_{j,s} \mathbb{1}_{C_j > \hat{C}_t} \frac{X_{fjst_0}}{X_{ft_0}} \log M_{jst},$$

and

$$D_{ft,L}^{M_s} = \frac{X_{ft_0}^*}{S_{ft_0}^*} \sum_{j,s} \mathbb{1}_{C_j \leq \hat{C}_t} \frac{X_{fjst_0}}{X_{ft_0}} \log M_{jst},$$

---

[27]See Appendix B for more details about these data and for an explanation on why we can only construct a time-invariant measure of competition by country.

where C_j denotes the country-specific competition index (ease of doing business) and $\hat{C}_t$ is the median of this value in year t. Hence $\mathbb{1}_{C_j \leq \hat{C}_t}$ is equal to 1 if country j is less competitive than the median country. These two new demand shocks sum up to our baseline measure $D_{ft}^{M_s}$ and capture separate export demand proxies for destinations with high/low competition.

We then estimate the following model:

$$Y_{ft} = \alpha_H \ D_{ft,H}^{M_s} + \beta_H \ D_{ft,H}^{M_s} * d_f + \alpha_L \ D_{ft,L}^{M_s} + \beta_L \ D_{ft,L}^{M_s} * d_f + \chi_{s,t} + \chi_f + \varepsilon_{ft}. \tag{9}$$

Our theory predicts that we should observe the skewness impact of the demand shocks more (or entirely) for the high-competition destinations. The results reported in Table 9 strongly confirm this prediction. This table considers two separate ways of measuring the threshold value $\hat{C}$. In the odd-numbered columns (1), (3), (5), $\hat{C}$ is the yearly median value once the measure of competition has been aggregated by product (i.e. on a sample containing one observation per product/firm). In the even-numbered columns (2), (4), (6), we use the threshold value when we keep only one observation per country. Both threshold measures for high/low competition confirm that the skewness effect is predominantly driven by the impact of export demand in high competition destinations.

4.6.2 Regression in long differences

In this section, we explore an alternate estimation strategy based on long differences over time. We decompose our full 1995-2012 sample into two periods $p \in \{p_0, p_1\}$ of equal length. Our demand variable is then measured in log differences, at the product (6 digit HS) or industry (3 digit ISIC) level, as:

$$\Delta D_f^{M_s} = \frac{X_{fp_0}^*}{S_{fp_0}^*} \sum_{j,s} \frac{X_{fjsp_0}}{X_{fp_0}} \log \frac{M_{jsp_1}}{M_{jsp_0}},$$

$$\Delta D_f^{M_I} = \frac{X_{fp_0}^*}{S_{fp_0}^*} \sum_{j,I} \frac{X_{fjIp_0}}{X_{fp_0}} \log \frac{M_{jIp_1}}{M_{jIp_0}},$$

where all trade flows are aggregated over each period p_0 and p_1.[28] Similarly, we measure innovation output ΔY_f as the difference in patent introductions between both periods (same measures as for our baseline analysis).

---

[28] $X_{fjsp_0} = \sum_{t \in p_0} X_{fjst}$, and $X_{fp_0} = \sum_{j,s} X_{fjsp_0}$. $X_{fp_0}^*$ and $S_{fp_0}^*$ are the average over period p_0 of yearly (Ficus) exports and sales. M_{jsp} is the average over period p of M_{jst}. If $M_{jsp} = 0$, we replace it with 1 euro. The construction is similar at the industry level. Finally, d_{fp_0} is the decile of the average (within sector) productivity decile over period p_0. A firm's sector is its most representative one during period p_0.

Table 9: More direct estimation of the competition effect

Dependent variable	All patents		Triadic patents		EPO patents	
Demand Measure	$D_{ft,H}^{M_s}, D_{ft,L}^{M_s}$		$D_{ft,H}^{M_s}, D_{ft,L}^{M_s}$		$D_{ft,H}^{M_s}, D_{ft,L}^{M_s}$	
	(1)	(2)	(3)	(4)	(5)	(6)
Low Competition	-0.030	-0.113	-0.026	-0.041	-0.147	-0.140
	(0.457)	(0.465)	(0.031)	(0.032)	(0.092)	(0.092)
High Competition	-5.686***	-5.275***	-0.373**	-0.387**	-0.691***	-0.902***
	(1.980)	(1.769)	(0.153)	(0.151)	(0.222)	(0.221)
Interact. Low	0.086	0.008	0.012	0.013**	0.024	0.001
	(0.144)	(0.086)	(0.009)	(0.006)	(0.029)	(0.014)
Interact. High	2.134***	1.338**	0.182***	0.103**	0.306***	0.179***
	(0.652)	(0.522)	(0.052)	(0.043)	(0.084)	(0.059)
Nb of observation	76,821	76,821	76,821	76,821	76,821	76,821
R^2	0.892	0.896	0.836	0.756	0.861	0.843

Notes: This table presents regression results of an OLS estimation of equation (9). Demand (low comp.) corresponds to $D_{ft,B}^{M_s}$ as defined in section 4.6.1 and Demand (high comp.) to $D_{ft,A}^{M_s}$. Interaction low comp. (resp high comp.) is defined as the interaction between the productivity decile of the firms and $D_{ft,B}^{M_s}$ (resp $D_{ft,A}^{M_s}$). Columns 1, 3 and 5 define the competition median at the firm × product level each year to compute demand shocks while columns 2, 4 and 6 compute the median at the country level (see section 4.6.1). Sample includes manufacturing firms with at least one patent in 1995-2012 and for which we can compute the export demand shock (see section 4.1). Coefficients and standard errors are obtained using a panel fixed-effect estimator. Autocorrelation and heteroskedasticity robust standard errors using the Newey-West variance estimator with a bandwidth set to 5 years. ***, ** and * respectively indicate 0.01, 0.05 and 0.1 levels of significance.

The firm fixed-effect is differenced-out but we keep the sector fixed effect; and we add the firm's productivity decile as an additional (pre-trend) control. Our estimating equation then becomes:

$$\Delta Y_f = \alpha \; \Delta D_f^{M_s} + \beta \; \Delta D_f^{M_s} * d_{fp_0} + \gamma \; d_{fp_0} + \chi_s + \varepsilon_f, \tag{10}$$

The results are reported in Table 10. The impact of the demand shock for the lowest productivity decile is still negative though barely significant. However, the interaction between the demand shock and the firm's initial productivity decile is positive and much more strongly significant. These results confirm our previous baseline findings.

4.6.3 Measures of innovation inputs: R&D investment

Up to now, we have used measures for the output of innovation based on patents. Another commonly used measure for innovation is based on R&D investment (the inputs for the innovation

Dependent variable	Nb patents	Triadic families	EPO families
Demand Shock	$\Delta D_f^{M_s}$	$\Delta D_f^{M_s}$	$\Delta D_f^{M_s}$
	(1)	(2)	(3)
Demand	-5.382**	-0.314*	-0.514*
	(2.431)	(0.191)	(0.302)
Decile × demand	1.260**	0.108**	0.140**
	(0.562)	(0.0461)	(0.0612)
Decile	-0.0345	-0.00493	0.00832
	(0.0554)	(0.00356)	(0.00661)
Nb of observation	4,707	4,707	4,707
R^2	0.0197	0.0171	0.0138

Notes: This table presents regression results of an OLS estimation. Sample includes one observation per manufacturing firm with at least one patent in 1995-2012 and for which we can compute the export demand shock (see section 4.1). 1995-2012 is broken down in 2 periods (1995-2003 and 2004-2012), over which trade flows and firm characteristics are aggregated. ***, ** and * respectively indicate 0.01, 0.05 and 0.1 levels of significance.

process). As with any input-based measure, the latter generates biases against firms that use those inputs more efficiently to generate innovation. A separate issue is that this measurement of R&D inputs is only available based on survey responses covering a subsample of firms. The sample is exhaustive for the largest innovators, but the sampling frequency decreases steeply with firm size. Thus, smaller firms are not consistently surveyed over time, thwarting the construction of time-varying measures of innovation (see more details in Appendix B.6).

Table 11 reports the correlations between our three main patent measures and a firm's total R&D budget and the associated number of R&D researchers – whenever this survey data is available. The left-hand side panel reports the between-firm correlations based on firm averages across years. Although the correlations across innovation inputs (R&D) and outputs (patents) in the bottom-left rectangle are weaker than the correlations within a set of inputs or outputs, the between firm correlations are nevertheless substantial and highly significant. However, those correlations between inputs and outputs drop precipitously when focusing on within-firm variations – whereas the correlations within the set of either inputs or outputs remain strong. Those correlations are reported in the middle and right-hand side panels in Table 11. (The middle panel reports within-firm correlations across all years after absorbing a firm fixed-effect; while the right-hand side panel reports the within-firm correlation between periods p_0 and p_1 as defined in our

Table 11: Correlations between R&D and patent measures of innovation
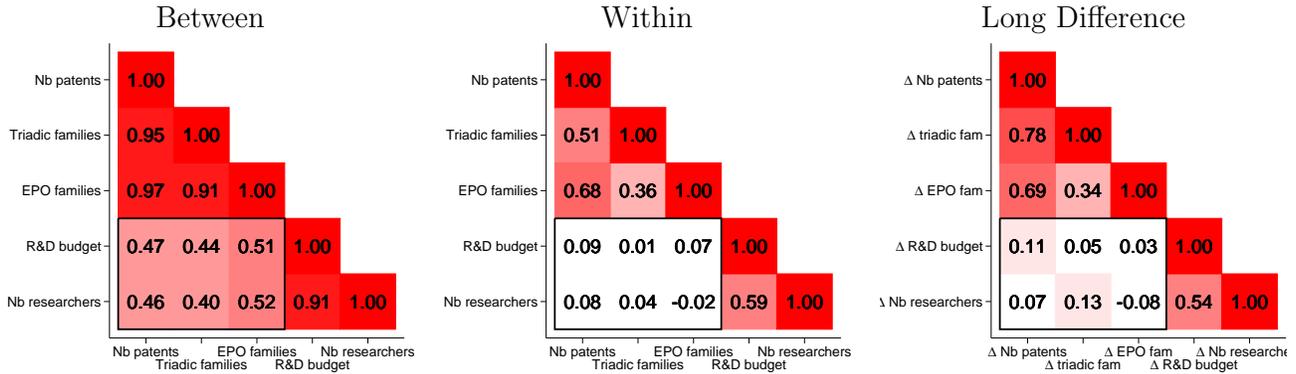
Notes: This table presents the pairwise correlations between innovation measures from PATSTAT and from the R&D survey. It is based on the sample of manufacturing firms with at least one patent in 1995-2012, for which we can compute the export demand shock (see section 4.1). Between correlations are the correlations between the firms' averages over the period. Absorbing a year or year x sector fixed effect prior to taking the firms' average leaves the correlations and their p-values virtually unchanged. Within correlations are the correlations after taking out the firm fixed effect. The long difference correlations are the correlations between the period 1 innovation measures. Using Sidak-adjusted p-values, all the correlations are significant at 5%, except the within and LD correlations between triadic families and total R&D budget and the LD correlation between EPO families and total R&D budget.

previous long-difference regressions.) Those very low correlations could be driven by the fact that R&D investments within a firm occur at discrete time intervals and slowly translate into increased patents – along with unmeasured changes in the efficiency/utilization of those R&D inputs.

In Table 12, we report the regression results for both our baseline specification as well as the long-difference one using both R&D input measures. In order to separate out the impact of the reduction in sample size associated with the availability of R&D data, we report results using our main patent innovation output variable with the same subsample of firm-years. In the left-hand columns reporting the level regressions, we see that the reduction in sample size does not affect our main results for the skewness impact of the export demand shock on the patent response. The coefficients using the R&D inputs have the same signs, but are not significant. We conjecture that this is due to the fact that the patent measure better captures the within-firm changes in innovation intensity at a yearly frequency. This is confirmed by the results for long-differences, where the coefficients for the R&D inputs are now substantially stronger and significant for the case of the number of researchers. As was the case in our full sample, the significance of the patent response is reduced when moving to the long-difference specification. In this case with a much smaller sample, those coefficients for the patent response are no longer significant.

Table 12: R&D SURVEY

Dependent variable	Regression in levels			Regression in long differences		
Innovation measure	R&D budget (1)	Nb researchers (2)	Nb patents (3)	ΔR&D budget (4)	ΔNb researchers (5)	ΔNb patents (6)
Demand	-1441.1 (3015.6)	-17.67* (9.531)	-8.859*** (3.340)	-4228.2 (4279.9)	-34.01** (15.03)	-6.506 (4.242)
Decile × Demand	1072.7 (720.5)	4.181* (2.321)	2.219*** (0.714)	1541.8 (1048.4)	8.554** (3.438)	1.560* (0.823)
Decile				-97.65 (173.3)	0.657 (0.568)	0.0835 (0.102)
Nb of observations	20,030	20,662	21,480	1,746	1,713	1,827
R^2	0.918	0.844	0.896	0.0112	0.0288	0.0301

Notes: This table presents regression results of an OLS estimation of equation 6 (columns 1-3) and 10 (columns 4-6). Sample includes manufacturing firms with at least one patent in 1995-2012 and for which we can compute the export demand shock (see section 4.1); the sample is further reduced to pairs (firm, year) in the R&D survey (columns 1-3) or to firms surveyed at least once in p_0 and p_1 in the R&D survey (columns 4-6). For the regressions in levels, coefficients and standard errors are obtained using a panel fixed-effect estimator. Autocorrelation and heteroskedasticity robust standard errors using the Newey-West variance estimator with a bandwidth set to 5 years. ***, ** and * respectively indicate 0.01, 0.05 and 0.1 levels of significance.

5 Conclusion

In this paper we analyzed the impact of export shocks on innovation for French firms. On the one hand those shocks increase market size and therefore innovation incentives for all firms. On the other hand they increase competition as more firms enter the export market. This in turn reduces profits and therefore innovation incentives particularly for firms with low initial productivity. Overall an export demand shock has a more positive effect on innovation in high productivity firms, whereas it may negatively affect innovation in low productivity firms. We tested this prediction with patent, customs and production data covering all French firms. To address potential endogeneity issues, we constructed firm-level variables which respond to aggregate conditions in a firm's export destinations but are exogenous to firm-level decisions. We showed that patenting robustly increases more with demand for initially more productive firms. This effect is reversed for the least productive firms as the negative competition effect dominates. Moreover, we showed that the positive interaction between a firm's initial productivity and the export demand shock is primarily driven by those export destinations where product market competition is highest. This further confirms the fact that export demand shocks involve both, a market size and a competition effect for French manufacturing innovators.

Our analysis can be extended in several directions. A first direction will be to use the same data to explore the effect of imports on innovation, using the same comprehensive databases. This would allow us to better understand why Bloom et al. (2016) and Autor et al. (2016) get

opposite conclusions. A second direction would be to look at the impact of exports on the citations to previous innovations, thereby shedding new light on the knowledge spillover effects of trade. These await future research.

References

Abadie, Alberto, Susan Athey, Guido W. Imbens, and Jeffrey Wooldridge, "When Should You Adjust Standard Errors for Clustering?," Working Paper 24003, National Bureau of Economic Research November 2017.

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

_ , **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, "Technology Dynamics, Innovation Policy and the Gains from Globalization," 2014. Unpublished manuscript.

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

Bernard, Andrew B., Bradford J. Jensen, Stephen J. Redding, and Peter K. Schott, "Global Firms," NBER Working Paper 22727, National Bureau of Economic Research 2016.

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 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.

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.

Coelli, Federica, Andreas Moxnes, and Karen-Helene Ulltveit-Moe, "Better, Faster, Stronger: Global Innovation and Trade Liberalization," *CEPR Discussion Paper 11506*, 2016.

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.

Djankov, Simeon, Rafael La Porta, Florencio Lopez de Silanes, and Andrei Shleifer, "The Regulation of Entry," *The Quarterly Journal of Economics*, 2002, *117* (1), 1–37.

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.

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.

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, Loriane Py, and Paul Trichelair, "Linking patents to firms: insights with French firms," Working Paper, Banque de France in progress.

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

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.

_ , _ , **and** _ , "Product Mix and Firm Productivity Responses to Trade Competition," NBER Working Papers 22433, National Bureau of Economic Research, Inc July 2016.

Melitz, Marc J. and Stephen J. Redding, "Heterogeneous Firms and Trade," in Elhanan Helpman, Kenneth Rogoff, and Gita Gopinath, eds., *Handbook of International Economics*, Vol. 4 of Handbook of International Economics, Elsevier, 2014, pp. 1–54.

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

Newey, Whitney K and Kenneth D West, "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, May 1987, *55* (3), 703–708.

Nickell, Stephen J, "Competition and Corporate Performance," *Journal of Political Economy*, August 1996, *104* (4), 724–746.

Park, Walter G, "International patent protection: 1960–2005," *Research policy*, 2008, *37* (4), 761–766.

Wooldridge, Jeffrey, *Econometric Analysis of Cross Section and Panel Data* Econometric Analysis of Cross Section and Panel Data, MIT Press, 2010.

A   Theoretical Appendix

A.1   Proof of Proposition 1

Uniqueness: in the $(\widehat{C}_D, \lambda)$ space, one can show that 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: to prove the existence of an equilibrium, 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.

A.2   Proof of Proposition 4

Proof.   The first part of proposition 4 follows directly: for a given distribution of French firms Γ and a given market size L in destination D, one can find c_I and F small enough such that $\overline{C} < \widehat{C}_I < \widehat{C}$, with $\overline{C}$ such that $\frac{dQ}{dL}(\overline{C}) = 0$. Importantly, the limit cases with $c_I = 0$ or $F = 0$ have a solution, which ensures that we can choose c_I and F as small as possible to obtain a solution verifying $\overline{C} < \widehat{C}_I < \widehat{C}$. When $c_I = 0$, all firms innovate, because the first unit of innovation has a zero cost. When $F = 0$ there is no operating cost, so that all firms with positive output actually produce.

Formally and after some algebra, we have:

$$\overline{C} < \widehat{C}_I \quad \Leftrightarrow \quad \frac{2\beta c_I}{\varepsilon\alpha L} < \frac{\frac{L}{\lambda}\frac{d\lambda}{dL}}{1 + \frac{L}{\lambda}\frac{d\lambda}{dL}}$$

$$\widehat{C}_I < \widehat{C} \quad \Leftrightarrow \quad F < \frac{c_I^2}{\beta\varepsilon^2 L\lambda},$$

Because the elasticity of competition with respect to the market size exists when $c_I = 0$ and is positive (it is the ratio of the average profits over the average revenues, see below), we can choose c_I small enough that $\overline{C} < \widehat{C}_I$. Then and because λ exists and is strictly positive when $F = 0$, we can choose F small enough that $\widehat{C}_I < \widehat{C}$.

The second part is detailed below.

Differencing (FE) with respect to L yields:

$$\int_{\tilde{c}_{0D}}^{\widehat{C}_D} \pi d\Gamma_D(\tilde{c}) = -L \int_{\tilde{c}_{0D}}^{\widehat{C}_D} \frac{\partial\pi}{\partial\lambda}\frac{d\lambda}{dL}d\Gamma_D(\tilde{c})$$

$$\Rightarrow \frac{d\lambda}{dL}\frac{L}{\lambda} = \frac{\displaystyle\int_{\tilde{c}_{0D}}^{\widehat{C}_D} \pi d\Gamma_D(\tilde{c})}{\displaystyle\int_{\tilde{c}_{0D}}^{\widehat{C}_D} r d\Gamma_D(\tilde{c})} = \beta\frac{\displaystyle\int_{\tilde{c}_{0D}}^{\widehat{C}_D} q^2(\tilde{c}, k; \lambda)d\Gamma_D(\tilde{c})}{\displaystyle\int_{\tilde{c}_{0D}}^{\widehat{C}_D} q(\tilde{c}, k; \lambda)\left[\alpha - \beta q(\tilde{c}, k; \lambda)\right] d\Gamma_D(\tilde{c})}.$$

Innovation – or equivalently output or output holding k fixed – decreases for the marginal innovator if and only if $\frac{d\lambda}{dL}\frac{L}{\lambda} > \frac{\alpha - \widehat{C}_I\lambda}{\widehat{C}_I\lambda}$ (see equation (4)). Using the above value for the elasticity of competition with respect to market size and after some algebra, output decreases for the marginal innovator if and only if:

$$\int_{\tilde{c}_{0D}}^{\widehat{c}_D} \left[q(\tilde{c}, k; \lambda) - \frac{2c_I}{\varepsilon L}\right] q(\tilde{c}, k; \lambda)d\Gamma_D(\tilde{c}) > 0. \tag{11}$$

Given that $q(\widehat{C}_I, k; \lambda) = \frac{c_I}{\varepsilon L}$, the distribution Γ_D must put enough weight on local firms with output above twice the output of the marginal innovator. ∎

B Data description

B.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 B.4). If the firm shares a patent with another firm, 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 [1] 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 a lot in countries like Japan or Korea, in which case the number of patents by such firms would be artificially large. To check that this problem does not drive our results, we build different measures of patent counts as detailed below.

Different counts of patents The various indicators from `PATSTAT` used in the regressions are described in detail below. All these indicators, based on different ways of counting or selecting patents, have pros and cons and shed a different light on our analysis. As stated by de Rassenfosse et al. (2013), it is virtually impossible to define a measure of innovation based on patents that is immune to the various biases that are associated with such data.

Following the innovation literature, we always only select patents of invention (the bulk of patents), thus dropping utility model and design patents.

- `Number of patents`: Each year, we sum over the patents filed by a firm f. When a patent has other applicants than f, we only count the share that f represents among all the co-

---

[1]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.

applicants (one third if f has 2 other applicants). This variable thus is a fractional count, as most of the variables shown in the regressions.

- `Triadic families`: when the same invention is filed in different patent offices, in practice the firm typically files for a different patent at each office, each referring to the first it has filed (called priority patent): these patents relate to each other, they belong to the same (DOCDB) family.[2]. Triadic families refer to such families with at least one patent filed at the EPO, one patent filed at the USPTO, and one patent filed at either the Japanese, the Chinese or the Korean Patent Office. We want to select innovations filed in the 3 main economic regions worldwide (Europe, USA and Asia). We depart slightly from the literature regarding the treatment of Asia: we do not want to consider Japan as the only relevant country, but instead add the two main other innovating countries, China and Korea. Finally the family is weighted with how much f contributes to it: $\frac{\sum_{k \text{ patents} \in \text{ family}} \frac{\mathbb{1}_{f \text{ is applicant of } k}}{\text{nb applicants}_k}}{\text{nb patents in the family}}$. The date of the family corresponds to the earliest filing year of the patents in this family.

- `EPO families`: The construction is very similar to that of the triadic families, except that the family will be taken into account if there is at least one patent in it filed at the EPO.

- `Dyadic families`: The construction is very similar to that of the triadic families, except that the family will be taken into account if there is at least one patent in it filed at the EPO, and another filed at the USPTO.

- `Families`: The construction is very similar to that of the triadic families, except that we take into account all the families containing a patent applied for by f.

- `EPO*`: We use the fractional count of the patents filed by firm f at the EPO.

- `Raw number of patents`: we use the (non-fractional) count of the number of patents filed by firm f.

- `Only Granted`: We use the fractional count of the granted patents filed by firm f.

---

[2]The `PATSTAT` data catalog states that "a large DOCDB family might indicate that the applicant seeks a wide geographical protection for the invention", and that "if two applications claim exactly the same prior applications as priorities (these can be e. g. Paris Convention priorities or technical relation priorities [...], then they are defined by the EPO as belonging to the same DOCDB simple family."

B.2 Firm-level accounting data

Our second data source provides us with accounting data for French firms from the DGFiP-INSEE, this data source is called `FICUS` and `FARE`. 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'Activits Franaise*), 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.

B.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. To convert HS products into ISIC industries we use a United Nations correspondence table (when 1 HS code corresponds to 2 ISIC codes, we split the HS flow in half into each ISIC code).

B.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[3] as a firm can be identified by its *Siren* identifier in both datasets. Thus the main challenge is to match either of these two datasets with `PATSTAT`. Indeed, `PATSTAT` only reports the name of the patent owner. 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. (in progress) who developed a matching algorithm to map patents with the corresponding French firms. The advanced methodology, described below, is a leap forward compared with other methods proposed by the literature.

Lequien et al. (in progress) 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 ...).

---

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

- 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)

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?)

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.5   Other data

We also use additional databases at the country level for our analysis. First we use the October 2015 vintage of the IMF's `World Economic Outlook` which provides country information such as GDP and population with a coverage as wide as possible. Second, to measure the level of competition for each country, we use the "Doing Business" project, based on the work of Djankov et al. (2002) and updated by the World Bank. Among all the available information, we consider the "ease of starting a business" which is the variable with the largest spatial coverage. This is a rating of all country for 0 to 100 that measures the constraints when one want to open a new company in the country. Because most countries are not surveyed each year, we choose to take a time invariant average value of this measure as our competition indicator.

B.6   R&D survey variables and sample

B.6.1   The survey

The annual survey on R&D expenses in firms exists since 1963. It describes the private sector R&D in terms of financial means (spending and financing) and mobilized workers. It covers firms established in France and doing R&D, and gathers information on previous year R&D activities. Usually surveys on firms are sampled with the *rpertoire Sirene*, but this database has no information on R&D activities to select firms that one would like to cover. firms with R&D activities represent 1/200 among active firms in *Sirene*.

The Ministry of Higher Education and Research therefore selects firm according to the following procedure:

1. **The historical repertoire**: All units having had a R&D activity are considered. This repertoire is updated with the newest information from the previous survey: takeover, absorption ... Firms answering they do not do R&D the year of survey but they might do some the following year are kept in this sample.

2. **External sources**: Administrative files and surveys allow to detect new firms possibly doing R&D: firms receiving the *Crdit Impt Recherche*, having the *young innovating firm* status, receiving help from firms incubators, firms reporting R&D activities in other surveys (Community Innovation Survey ...)

3. **Updating with Sirene cessations**: Firms known as having shut down in *Sirene* are eliminated.

4. **Stratification**:

- Firms with internal spending of R&D over 400 000 are exhaustively interrogated (and above 2 million , they fill a bigger questionnaire).

- New firms (CIR, JEI . . .) are exhaustively interrogated as well (but only since 2001, see below).

- The rest stays only two years in a row in the panel: firms interrogated in N-1 and N-2 are excluded, those interrogated in N-1 are kept, and some newly selected firms are drawn.

Main changes in the survey methodology

- 1992: reform leading to broadly the survey as it exists today. Most variables exist since 1992 or 1993.

- 2000: Increase of the threshold separating the simplified from the general questionnaire, from 5 million francs to 10 million francs. Some variables therefore are missing for firms that filed the simplified questionnaire in 2000.

- 2001: New firms are all interrogated in the first year. Until 2000, only 1 in 2 new firm was interrogated, the rest was kept for the following year.

- Change in units: in 1998, the answer is in francs and not in thousands francs anymore, because many errors were seen. In 2004, the answer is in thousands euros and not in euros anymore. After 2008, the answer is again asked in euros.

Some firms provide a "group" answer. Indeed for larger firms, R&D activity is more often organized at the group level than at the legal unit level. A variable lists the legal units concerned, but only after 2009.

B.6.2 The variables

- Total R&D Budget : total spending of a firm on R&D activities. It is the sum of internal and external spending. One has to be careful, this variable can count twice contracts made between two firms of the same group, once in internal spending and a second time in external spending.

- *of which* Current spending : gross wages of R&D workers and general expenses (spending in capital excluded), such as small tools, raw materials, administrative costs . . .

- *of which* Gross wages of R&D workers. It includes all fiscal and social contributions.

- R&D workers : Researchers and technicians (support). in full time equivalent (prorata of time spent in R&D activities, with a minimum of 10%).

- *among which* Nb of researchers : scientists and engineers working at creating knowedge, products, processes, methods or new systems. It includes PhDs paid by the firm or high-level staff responsible with animating the researchers' teams. In full time equivalent.

C   Additional Empirical Results

Table B1 presents regression results of an OLS estimation of equation 6 replacing the innovation left-hand-side with (the log of) total firm exports and dropping the export intensity from the demand shock computation. We thus use the following unadjusted export demand variables (at both the product M_s and industry M_I level):

$$\tilde{D}_{ft}^{M_s} = \sum_{j,s} \frac{X_{fjst_0}}{X_{ft_0}} \log M_{jst}$$

$$\tilde{D}_{ft}^{M_I} = \sum_{j,I} \frac{X_{fjIt_0}}{X_{ft_0}} \log M_{jIt}.$$

We see that those unadjusted export demand variables very strongly predict a firm's export response (first two columns). On the other hand, there is no evidence for a skewness effect for that export response – according to a firm's proximity to frontier. Our theoretical model predicts that the skewness effect evolves slowly over time as competition increases. The innovation response is forward looking and captures this anticipated effect, whereas the export measure does not.

Table B1: IMPACT OF THE DEMAND SHOCK ON FIRM'S EXPORTS

Dependent variable	log(Exports)	log(Exports)	log(Exports)	log(Exports)
Demand Shock	$\tilde{D}_f^{M_s}$	$\tilde{D}_f^{M_I}$	$\tilde{D}_f^{M_s}$	$\tilde{D}_f^{M_I}$
	(1)	(2)	(3)	(4)
Demand	0.0419***	0.0592***	0.0334	0.0537**
	(0.0135)	(0.0164)	(0.0228)	(0.0260)
Decile × demand			0.00203	0.00135
			(0.00396)	(0.00453)
Nb of observation	72,380	72,416	72,380	72,416
R^2	0.855	0.856	0.855	0.856

Notes: Sample includes manufacturing firms with at least one patent in 1995-2012 and for which we can compute an the export demand shocks $\tilde{D}_{ft}^{M_s}$ and $\tilde{D}_{ft}^{M_I}$. Coefficients and standard errors are obtained using a panel fixed-effect estimator. Autocorrelation and heteroskedasticity robust standard errors using the Newey-West variance estimator with a bandwidth set to 5 years. ***, ** and * respectively indicate 0.01, 0.05 and 0.1 levels of significance.

Table B2 considers more aggregated (across industries and products) demand shock measure using the GDP of destination j at t instead of world imports (excluding France) for a particular

industry or product. Namely:

$$D_{ft}^G = \frac{X_{ft_0}^*}{S_{ft_0}^*} \sum_j \frac{X_{jt_0}}{X_{ft_0}} \log GDP_{jt}.$$

Table B2: OTHER DEMAND SHOCK

Dependent variable	All patents	Triadic patents	EPO patents
Demand measure	$D_{ft}^{M_s}$ (1)	$D_{ft}^{M_s}$ (2)	$D_{ft}^{M_s}$ (3)
Demand	-3.037**	-0.217*	-0.448**
	(1.466)	(0.119)	(0.195)
Decile × Demand	0.852**	0.103***	0.149***
	(0.394)	(0.0328)	(0.0544)
Nb of observations	77,002	77,002	77,002
R^2	0.891	0.753	0.838

Notes: This table presents regression results of an OLS estimation of equation 6. Sample includes manufacturing firms with at least one patent in 1995-2012 and for which we can compute an export demand shock D_{ft}^G. Coefficients and standard errors are obtained using a panel fixed-effect estimator. Autocorrelation and heteroskedasticity robust standard errors using the Newey-West variance estimator with a bandwidth set to 5 years. ***, ** and * respectively indicate 0.01, 0.05 and 0.1 levels of significance.

Table B3: DIFFERENT TRIMMING

Dependent variable	All patents					
Demand Measure Trimming	$D_{ft}^{M_s}$ 2.5% (baseline) (1)	$D_{ft}^{M_s}$ 0% (2)	$D_{ft}^{M_s}$ 1% (3)	$D_{ft}^{M_s}$ 2% (4)	$D_{ft}^{M_s}$ 3% (5)	$D_{ft}^{M_s}$ 5% (6)
Demand	-3.260***	-0.751**	-1.994***	-3.052***	-3.491***	-3.474***
	(1.014)	(0.305)	(0.745)	(0.925)	(1.119)	(1.256)
Decile × Demand	0.960***	0.166**	0.510***	0.842***	1.065***	1.086***
	(0.255)	(0.0816)	(0.169)	(0.227)	(0.285)	(0.327)
Nb of observation	77,901	82,043	80,378	78,722	77,077	73,784
R^2	0.897	0.881	0.892	0.896	0.898	0.908

Notes: This table presents regression results of an OLS estimation of equation 6. Sample includes manufacturing firms with at least one patent in 1995-2012 and for which we can compute the export demand shock (see section 4.1). Different trimming on extreme variations of the Demand variable are done. Coefficients and standard errors are obtained using a panel fixed-effect estimator. Autocorrelation and heteroskedasticity robust standard errors using the Newey-West variance estimator with a bandwidth set to 5 years. ***, ** and * respectively indicate 0.01, 0.05 and 0.1 levels of significance.significance.

Table B4: Removing first years

Dependent variable	All patents		Triadic patents		EPO patents	
Demand Measure	$D_{ft}^{M_s}$	$D_{ft}^{M_s}$	$D_{ft}^{M_s}$	$D_{ft}^{M_s}$	$D_{ft}^{M_s}$	$D_{ft}^{M_s}$
Years removed	$t < t_0 + 2$	$t < t_0 + 3$	$t < t_0 + 2$	$t < t_0 + 3$	$t < t_0 + 2$	$t < t_0 + 3$
	(1)	(2)	(3)	(4)	(5)	(6)
Demand	-2.703***	-2.378**	-0.233***	-0.203***	-0.339***	-0.321**
	(1.003)	(0.938)	(0.0788)	(0.0733)	(0.129)	(0.129)
Decile × Demand	0.811***	0.654***	0.0752***	0.0620***	0.112***	0.0896***
	(0.249)	(0.229)	(0.0194)	(0.0172)	(0.0298)	(0.0292)
Nb of observation	72,265	66,684	72,265	66,684	72,265	66,684
R^2	0.911	0.920	0.763	0.763	0.870	0.885

Notes: This table presents regression results of an OLS estimation of equation 6. First years following t_0 are excluded from the estimation. Sample includes manufacturing firms with at least one patent in 1995-2012 and for which we can compute the export demand shock (see section 4.1). Coefficients and standard errors are obtained using a panel fixed-effect estimator. Autocorrelation and heteroskedasticity robust standard errors using the Newey-West variance estimator with a bandwidth set to 5 years. ***, ** and * respectively indicate 0.01, 0.05 and 0.1 levels of significance.

Table B5: Baseline results - Clustered standard errors

Dependent variable	All patents		Triadic patents		EPO patents	
Demand measure	$D_{ft}^{M_s}$	$D_{ft}^{M_I}$	$D_{ft}^{M_s}$	$D_{ft}^{M_I}$	$D_{ft}^{M_s}$	$D_{ft}^{M_I}$
	(1)	(2)	(3)	(4)	(5)	(6)
Demand	-3.260**	-2.578*	-0.265**	-0.224**	-0.368**	-0.447**
	(1.475)	(1.530)	(0.115)	(0.112)	(0.168)	(0.200)
Decile × Demand	0.960***	0.909**	0.0859***	0.0862***	0.125***	0.114**
	(0.372)	(0.444)	(0.0287)	(0.0327)	(0.0397)	(0.0554)
Nb of observation	77,901	77,918	77,901	77,918	77,901	77,918
R^2	0.897	0.888	0.759	0.747	0.849	0.836

Notes: This table presents regression results of an OLS estimation of equation 6. Sample includes manufacturing firms with at least one patent in 1995-2012 and for which we can compute the export demand shock (see section 4.1). Coefficients and standard errors are obtained using a panel fixed-effect estimator. Heteroskedasticity robust standard errors clustered at the firm level ***, ** and * respectively indicate 0.01, 0.05 and 0.1 levels of significance.

Table B6: Poisson regressions

Dependent variable	All patents	Triadic patents	EPO patents
Demand measure	$D_{ft}^{M_s}$ (1)	$D_{ft}^{M_s}$ (2)	$D_{ft}^{M_s}$ (3)
Demand	-0.652***	-0.937***	-0.466*
	(0.208)	(0.360)	(0.246)
Decile × Demand	0.127***	0.164***	0.103***
	(0.0292)	(0.0481)	(0.0336)
Nb of observations	73,488	18,410	47,648

Notes: This table presents regression results of a Poisson estimation of equation 7. To obtain integer dependent variables, we do not use the fractional count (using these integer number with OLS has negligible effects) . Coefficients and standard errors are obtained using a maximum likelihood estimator. Autocorrelation and heteroskedasticity robust standard errors using the Newey-West variance estimator with a bandwidth set to 5 years. Sample includes manufacturing firms with at least one patent in 1995-2012 and for which we can compute the export demand shock (see section 4.1). ***, ** and * respectively indicate 0.01, 0.05 and 0.1 levels of significance.

Table B7: Excluding leaders

Dependent variable	All patents		Triadic patents		EPO patents	
Demand measure	$D_{ft}^{M_s}$ (1)	$D_{ft}^{M_I}$ (2)	$D_{ft}^{M_s}$ (3)	$D_{ft}^{M_I}$ (4)	$D_{ft}^{M_s}$ (5)	$D_{ft}^{M_I}$ (6)
Demand	-2.820***	-2.730**	-0.251***	-0.219***	-0.357***	-0.482***
	(0.970)	(1.074)	(0.0838)	(0.0849)	(0.122)	(0.145)
interac	0.766***	1.246***	0.0790***	0.104***	0.122***	0.181***
	(0.239)	(0.318)	(0.0199)	(0.0244)	(0.0296)	(0.0405)
Nb of observation	77790	77806	77790	77806	77790	77806
R^2	0.897	0.887	0.751	0.748	0.846	0.835

Notes: This table presents regression results of an OLS estimation of equation 6. Sample includes manufacturing firms with at least one patent in 1995-2012 and for which we can compute the export demand shock (see section 4.1). The Demand variable does not include country j and products s for a firm f with a market share above 10% for the pair (j, s). Coefficients and standard errors are obtained using a panel fixed-effect estimator. Autocorrelation and heteroskedasticity robust standard errors using the Newey-West variance estimator with a bandwidth set to 5 years. ***, ** and * respectively indicate 0.01, 0.05 and 0.1 levels of significance.