Estimating firm product quality using trade data*

Paul Piveteau[†]

Gabriel Smagghue[‡]

February 2019

Abstract

We propose a new instrumental variable strategy to estimate product quality at the firm-level, using trade data. Interacting firm importing shares by country with real exchange rates (RER), we obtain a cost shifter that varies across firms and is arguably orthogonal to product quality. We use this import weighted RER as an instrument for export prices and we identify firm-level quality from residual export variations, after controlling for prices. Our quality estimates correlate to firm characteristics (e.g. wages) and to alternative measures of quality available for some rare sectors. Moreover, we document cases in which our estimates more adequately characterize quality compared to prices, a popular proxy for quality. We show for instance that firms add products to their export portfolio when their quality increases, as expected, while simultaneously their prices decrease. This suggests that our empirical strategy, by delivering quality estimates which, unlike prices, are not polluted with productivity variations, should contribute to future research on the link between firm-level product quality and globalization.

*First version: July 2013. We are especially grateful to the editor, two anonymous referees, Maria Bas, Tibor Besedes, Arnaud Costinot, Jonathan Dingel, Gilles Duranton, Juan Carlos Hallak, James Harrigan, Amit Khandelwal, Thierry Mayer, Julien Martin, Marc Melitz, Eric Verhoogen and David Weinstein for helpful remarks and suggestions. The paper was previously circulated under the title "A new method for quality estimation using trade data: an application to French firms". We acknowledge the financial support of the Spanish Ministry of Science and Innovation under grants ECO2011-27014. We thank CNIS and French customs for confidential data access. We are also grateful to audiences of the International Trade colloquium at Columbia University, the Sciences Po lunch seminar, the MIT International Tea Seminar, the LSE trade seminar, EEA, ETSG and FREIT-EIIT.

[†]School of Advanced International Studies, Johns Hopkins University. 1717 Massachusetts Avenue NW. Washington DC, 20036. Email: ppiveteau@jhu.edu

[‡]University Carlos III of Madrid. Email: gsmagghu@eco.uc3m.es

1 Introduction

Trade economists have long investigated the role played by product quality in shaping the patterns of trade at the macroeconomic level. A more recent literature has shown the importance of product quality at the microeconomic level: in addition to being one of the main sources of firm heterogeneity,[1] the quality supplied by firms impacts the relative demand for inputs, which makes it decisive to understand the link between globalization and inequalities.[2] These findings triggered an increasing demand from trade economists for disaggregated data on product quality. However, despite this need, estimating firm-level quality on trade data remains an empirical challenge: traditional techniques developed in Industrial Organization cannot be applied to datasets in which product characteristics are not observed,[3] which is typically the case with international trade data.[4]

In this paper, we propose and implement a new empirical methodology to estimate product quality at the firm level. We create a new instrument for prices, based on exchange rate variations interacted with firm-specific importing shares, that allows us to consistently estimate demand equations in the absence of observable product characteristics. Implementing this methodology using customs data from France, we first document the reliability of our estimation: we compare the estimated price elasticities and measures of quality with industry and firm characteristics as well as alternative measures of quality. Then, we employ the obtained quality measures to study the link between export performance and quality: we show that firms which add products or destinations in their portfolio simultaneously exhibit an increasing quality. Importantly, using prices, a common proxy for quality, to study this question leads to a different conclusion.

The main contribution of this paper is to provide a new method to estimate quality using trade data. We estimate quality from the demand side. The main challenge one faces when estimating demand functions is to deal with the endogeneity of prices: prices are likely to be correlated to demand shocks, because quality is costly to produce.[5] Consequently, researchers have used unit values or prices as proxies for quality, or have estimated demand equations in contexts where unobserved vertical differentiation is limited.[6] To address this endogeneity issue, we construct a novel instrument for prices, exploiting fluctuations in exchange rates. These fluctuations, interacted with firm-specific import shares, shift a firm's costs of importing goods. As the firm passes importing cost variations to its consumers, the instrument generates firm-specific export price and sales variations. These variations are arguably exogenous to unobserved demand shocks (e.g., quality shocks) and allow us to identify the price-elasticity of exports. Quality is then identified at the firm, destination, product, year level, from the residual variations of

[1]See Roberts, Yi Xu, Fan, and Zhang (2017) and Hottman, Redding, and Weinstein (2016) for empirical quantifications of the relative importance of different sources of heterogeneity at the firm level.

[2]Verhoogen (2008) and Brambilla et al. (2012) document the consequences of trade openness on wage inequality.

[3]Industrial Organization has developed strategies to back out quality by estimating a demand equation. In this approach, the presence of omitted product characteristics challenges the identification as these characteristics are likely to be correlated with the price of the product which induces an endogeneity bias.

[4]Exceptions include Crozet et al. (2012) and Garcia-Marin (2014) who use expert ratings of quality of Champagne and wine, as quality measures.

[5]See, e.g., Hallak and Sivadasan (2013), Johnson (2012) and Kugler and Verhoogen (2012) for trade models where quality is costly and endogenous at the firm-level.

[6]Broda and Weinstein (2010) and Handbury (2012) use barcode-level data, that features no quality variation within barcode across time, whereas Foster, Haltiwanger, and Syverson (2008) restrict their analysis to homogeneous products.

demand once price variations have been controlled for; a strategy that is present throughout the literature.

The implementation of this method using customs data from France, supports the validity of the procedure. First, we find that the import-weighted exchange rate, our instrument, is strongly and positively correlated to export prices charged by firms. This is consistent with the assumption we make to motivate the instrumentation, namely that exchange rates shift a firm's production costs. Second, in order to evaluate the ability of our instrument to correct for the endogeneity of prices, we estimate the demand equation using both ordinary least squares and instrumental techniques. Our instrumental variable procedure affects the estimates of price-elasticities consistently with a correction of an omitted variable bias: while ordinary least squares estimates deliver a low (in absolute value) price-elasticity (-0.8), the instrumental variable approach estimates an average price-elasticity of demand of -4.3, consistent with existing studies in the industrial organization literature. Moreover, elasticities estimated at a more disaggregated level are positively correlated with existing estimates from Broda and Weinstein (2006) and Soderbery (2015), and, as expected, negatively correlated with a measure of vertical differentiation from Sutton (2001).[7]

We then investigate the properties of the quality estimates obtained from the procedure. We show that the dispersion of these estimates within a market is positively correlated with existing measures from Khandelwal (2010). Moreover, we directly relate our quality measure to quality measures at the firm-level. A natural benchmark is provided by Crozet et al. (2012) who use one of the very few "direct" measure of firm-specific quality present in the literature, by relying on ratings attributed by an expert to a sample of French Champagne producers. We compare these ratings with our estimated quality of exported Champagne and find a positive and strongly significant correlation. Similarly, we find that the obtained quality measures are intuitively correlated with firms characteristics and in particular the average wage paid by firms.

Finally, we compare our estimated quality measures with export prices, the most commonly employed proxy for quality. We show that prices and quality are positively correlated in the cross-section of firms, as well as over time within a firm. However, this correlation is significantly stronger for vertically differentiated markets. In other words, prices are informative on quality, but less so in more homogeneous sectors. Then, we show that this imperfect correlation between prices and quality can be misleading when studying the role of quality in explaining export performance. In particular, we show that firms adding destinations or varieties to their portfolio do so as they experience an increase in the estimated quality of their products. On the contrary, using prices to study this question leads to contradictory answers as prices tend to increase with the addition of a destination, and decrease with the addition of a product. We argue these results highlight the superiority of our measure over prices that conflate many factors other than vertical differentiation.

This paper is directly related to the literature aiming to measure quality using trade data. Most of the literature back up quality measures from the estimation of a demand system, following the tradition in Industrial Organization.[8] In particular, we can cite Hallak and Schott (2011)

[7]However, given the low number of observations (14), these correlations are not statistically different from zero.

[8]Most notable contributions in IO include Berry, Levinsohn, and Pakes (1995) and Berry (1994). These papers

and Khandelwal (2010) who rely on an instrumental variable approach to identify quality at the country-product level using trade data. To be applied at the firm-product level, their methods require an instrument for prices which varies across firms. We provide such an instrument. Gervais (2015) and Roberts et al. (2017) also estimate quality at the firm level by instrumenting prices. However, these studies use instruments, respectively physical productivity and wages, which are questionable if quality varies over time, within the firm. More recently, Fontagné, Martin, and Orefice (2018) develop a strategy using variations in electricity prices across French firms. By contrast, our strategy relies on an instrument available to most trade economists as it can be constructed solely from customs data.

Because of the difficulty of estimating demand equations at the firm level, in the absence of product characteristics, researchers have relied on alternative strategies: Khandelwal, Schott, and Wei (2013) construct quality by calibrating price-elasticity with estimates from Broda and Weinstein (2006). The relevance of these price-elasticities estimates is open to question as they are obtained from country-level data. Alternatively, demand equations have been estimated in contexts where unobserved vertical differentiation is limited: for instance, Broda and Weinstein (2010) and Handbury (2012) use barcode-level data, whereas Foster et al. (2008) restrict their analysis to homogeneous products. Finally, another strand of the literature has relied on structural models to overcome the endogeneity problem when estimating demand equations.[9] In comparison to these methods, our paper provides an instrument for export prices based on trade data, which allows the consistent estimation of demand functions for potentially all industries and under weaker assumptions.

A number of papers have used prices to investigate the role played by quality in explaining export performance across firms. Most of these papers used output and input prices as proxy for quality: we can cite for instance Kugler and Verhoogen (2012) and Manova and Zhang (2012) that document quality variations across firms, and within firm across destinations, using firm-level or customs data. While the use of prices is appropriate in the context of their studies,[10] we believe the use of prices as proxy for quality can be problematic in other situations. Indeed, while product quality usually increases the cost of a good, many other factors determine the price charged by a firm for its product. Moreover, the presence of multi-products firms makes the use of prices even more so challenging since firms self-select their set of products based on their quality. Manova and Yu (2017) studies the complexity of the relationship between prices and product quality in the context of multi-product firms.

Finally, the use of exchange rates as an instrument for prices links our paper to Berman, Martin, and Mayer (2012) and Amiti, Itskhoki, and Konings (2014). These studies empirically analyze the firm-level pass-through from exchange rates to export prices. However, while both

have contributed to the estimation of structural demand parameters by introducing demand systems exhibiting more sophisticated substitution patterns. However, the structure included in these papers does not solve the issue that prices are endogenous to quality in the demand equation. Therefore, these structural empirical models do not dispense from finding an instrument for prices, but can usually rely on product characteristics that control for mosts of the variation in quality across goods.

[9]See Redding and Weinstein (2016) and Redding and Weinstein (2017) who estimate demand systems to recover an aggregate price index and aggregate trade patterns that are consistent with micro data.

[10]The positive correlation between prices and export performance they show clearly points toward a positive effect of product quality on firms' performance. However, a negative relationship would not imply a negative role for product quality.

works are interested in the heterogeneity of the pass-through across firms, we only use the effect of exchange rates on export prices as a first stage to a demand function estimation. More recently, Amiti, Itskhoki, and Konings (2016) studies the price setting of firms in response to shocks on their costs and the prices of their competitors. In this context, they also use exchange rates to obtain exogenous variations in the cost of imported inputs.

This paper is structured as follows. In the next section, we derive a simple model of demand with vertically-differentiated goods and present our identification strategy to consistently estimate demand equations using trade data. In section 3, we describe the French customs data used for the implementation and show the results of the estimation. Section 4 describes the relevance of the quality estimates we obtain by relating them to existing measures. Moreover, we explore the link between these measures and prices to show that using prices as proxy can be misleading in some contexts. Finally, section 5 concludes.

2 Quality Estimation Strategy

In this section, we present a novel strategy to estimate the quality of exports at the firm-product-destination-year level, using customs data. Since we identify quality from the demand side, this section introduces a demand system with constant elasticity of substitution (CES) in which the quality of a product acts as a utility shifter. This implies that variations in the quality of exported goods over time and across firms will be revealed from variations in sales that cannot be explained by price movements.

In order to identify the demand system and infer product quality measures, we then present a new instrument for the price of firms' exports. This instrument is obtained by interacting firm-specific importing shares with real exchange rates. We argue that this instrument is exogenous to quality choices made by firms and measurement errors on prices, which constitutes an improvement relative to existing instruments in the literature, allowing us to consistently estimate demand functions using trade data.

2.1 An Empirical Model of Demand

Let us consider a global economy composed of a collection of destination markets d. In each market, the representative consumer allocates her revenue over the different varieties of each product g. Our definition of product categories follows the structure of French customs data, namely a product corresponds to a 8 digit position of the Combined Nomenclature (CN). A variety is defined as an unique combination of a destination market d, a producing firm f and a product g.

Representative consumers have two tier preferences. The lower level of the utility function aggregates consumptions of varieties by product. The upper level aggregates consumptions across products. We assume that the lower part of the utility function displays a constant elasticity of substitution (CES) across varieties, while we do not need to impose any functional form on the upper level. It follows that an expression of the utility of the representative consumer in market d at year t is

$$U_{dt} = U\left(C_{1dt}, .., C_{Gdt}\right),$$
$$C_{gdt} = \left[\sum_{f \in \Omega_{gdt}} \left(\lambda_{fgdt}\, q_{fgdt}\right)^{\frac{\sigma_j - 1}{\sigma_j}}\right]^{\frac{\sigma_j}{\sigma_j - 1}} \qquad \text{for each } g = 1..G, \tag{1}$$

with $U(.)$ a well-behaved utility function, C_{gdt} the CES aggregate consumption of product g in destination d at year t, Ω_{gdt} the set of varieties of good g available to consumers, and σ_j the elasticity of substitution across varieties within a product category, that varies across industries j.[11] Moreover, q_{fgdt} and λ_{fgdt} are respectively the aggregate physical consumption and the quality of variety fgd at year t.

Utility function (1) imposes that varieties are equally substitutable within product categories.[12] In equation (1), quality is modeled as a utility shifter, i.e. a number of units of utility per physical unit of good. This implicitly defines quality as an index containing any characteristic of a variety which raises consumers' valuation of it. These characteristics may be tangible (e.g. size, color) as well as intangible (e.g. reputation, quality of the customer service, brand name). This broad definition is consistent with most of the literature in international trade and quality.[13]

The representative consumer allocates her total expenditure, E_{dt}, across goods and varieties, in order to maximize her utility (1). This behavior results in the following aggregate demand function for variety fgd:

$$q_{fgdt} = p_{fgdt}^{*-\sigma_j} \lambda_{fgdt}^{\sigma_j - 1} P_{gdt}^{\sigma_j - 1} E_{gdt}, \tag{2}$$

with E_{gdt} the expenditure optimally allocated to good g. p^*_{fgdt} is the price of variety fgd faced by consumers in destination d, labeled in market d's currency. P_{gdt} is the price index of good g in market d at year t.[14] In order to properly grasp the properties of demand function (2), it is worth noting that $-\sigma_j$ is not the own price elasticity of variety fgd's demand. It is rather the own price elasticity *keeping constant the price index P_{gdt} and the aggregate expenditure E_{gdt}*. In a monopolistic competition setting, firms are atomistic and their individual decisions do not influence these aggregate variables. However, with non-atomistic firms, individual prices

[11] In the empirical application, goods g will based on the 8-digit product classification, while we estimate different elasticities of substitution for 15 different industries j. Therefore, each product category g is nested within an industry j.

[12] This feature is shared by most estimations of demand systems with vertically differentiated goods using trade data. In the nested logit specification of Khandelwal (2010) for instance, the cross price elasticity is the same for any two varieties within a nest, irrespective of their quality, after controlling for their market shares.

[13] Because of the wide range of product attributes potentially captured by our concept of "quality", some papers have adopted a more conservative terminology. For instance, Roberts et al. (2017) refer to the variety-specific utility shifter as a "demand index", Foster et al. (2008) to "demand fundamental" and Hottman et al. (2016) to "product appeal".

[14] The price index is

$$P_{gdt} = \left(\sum_{f \in \Omega_{gdt}} \left(\frac{p^*_{fgdt}}{\lambda_{fgdt}}\right)^{1-\sigma_j}\right)^{\frac{1}{1-\sigma_j}}.$$

may have an aggregate impact, and thus the own price elasticity may differ from $-\sigma_j$ and be heterogeneous across firms.

Producing firms are located in different countries and we assume that exporting involves iceberg trade costs. Let "Home" be the country from which firms export in the data (France in our application). Domestic firms need to ship $\tau_{gdt} \geq 1$ units of good g for one unit to reach the consumer in market d at year t. Therefore, for varieties exported from home to market d, the customer price in d currency (p^*_{fgdt}) is linked to the FOB (Free on Board) price in home currency (p_{fgdt}) by the following relationship:

$$p^*_{fgdt} = \frac{\tau_{gdt}}{e_{dt}} p_{fgdt}, \tag{3}$$

with e_{dt} the direct nominal exchange rate from home currency (Euro in the application) to market d's, i.e. that one unit of d currency buys e_{dt} units of home currency. Plugging (3) and log-linearizing, we can re-express demand function (2) for domestic firms as follows:

$$\log q_{fgdt} = -\sigma_j \log p_{fgdt} + \tilde{\lambda}_{fgdt} + \mu_{gdt} \tag{4}$$
$$\text{with} \quad \begin{cases} \tilde{\lambda}_{fgdt} \equiv (\sigma_j - 1) \left(\log \lambda_{fgdt} - \overline{\log \lambda}_{gdt} \right) \\ \mu_{gdt} \equiv -\sigma_j \log \left(\frac{\tau_{gdt}}{e_{gdt}} \right) + (1 - \sigma_j) \log P_{gdt} + \log E_{gdt} + (\sigma_j - 1) \overline{\log \lambda}_{gdt} \end{cases}$$

and $\overline{\log \lambda}_{gdt} \equiv \frac{1}{N_{\mathcal{H}_{gdt}}} \sum_{f \in \mathcal{H}_{gdt}} \log \lambda_{fgdt}$ the average log-quality of good g supplied by domestic firms to market d at year t, $N_{\mathcal{H}_{gdt}}$ being the number of firms in the set $\mathcal{H}_{gdt}$ of firms exporting good g from home to country d at year t.

Equation (4) is the one that we bring to the data. In (4), $\log q_{fgdt}$ and $\log p_{fgdt}$ are observable to the econometrician while σ_j, $\tilde{\lambda}_{fgdt}$ and μ_{gdt} have to be estimated. One can see from (4) that the demand shifter of a firm contains a variety-specific as well as a market-specific term (respectively $\tilde{\lambda}_{fgdt}$ and μ_{gdt}). The latter term will be estimated by including destination-product-year fixed effects in the regression. This term is not informative on quality as it conflates the average quality of domestic exports with other aggregate variables. Thus, the estimation developed in this paper identifies quality from $\tilde{\lambda}_{fgdt}$, the variety-specific part of the demand shifter. Incidentally, the presence of quality in the demand shifter also causes the potential endogeneity of prices, as we discuss further below.

From the structural expression of $\tilde{\lambda}_{fgdt}$ in (4), one can see that our strategy does not deliver an absolute measure of quality. Instead, we obtain a measure of quality which is relative to the average quality supplied by domestic firms to a market. A corollary is that $\tilde{\lambda}_{fgdt}$ will not be suited to analyze variations in the aggregate quality of domestic exports, but rather how firms move relative to each other along the quality ladder across markets and over time. Moreover, because we assume that all firms have the same elasticity within a category, any deviation in the price elasticity across firms will be attributed to our quality measure. Therefore, our quality measure can also capture the relative market power of firms. As long as this market power is monotonically increasing with quality, this confounding effect will not affect the ranking of quality across firms.

The next subsection describes the estimation of demand function (4) with a focus on our treatment of the endogeneity of prices.

2.2 Dealing with Price Endogeneity

In our setup, the endogeneity of prices comes from two mechanisms. First, we face a well-known simultaneity problem as prices are likely to be correlated to quality, which is in the residual of the demand function. Assuming that high quality varieties are more costly to produce, this correlation would result from firms passing on the cost of quality to consumers. Similarly, firms with higher ability are likely to exert market power that will result in higher mark-up. In both scenarios, final prices charged by firms are correlated with demand, which leads ordinary least squares to underestimate the price-elasticity of demand, σ. Indeed, when a firm increases the quality of its products, the effect of prices on demand is compensated with the greater appeal of the good to consumers.

A second source of endogeneity, more specific to international trade data, comes from the construction of prices. Because prices are not directly observed, we follow the standard practice and use unit values as a proxy for prices. Unit values are obtained by dividing the value of a shipment by the physical quantity shipped. The use of this proxy may generate an attenuation bias due to the measurement error contained in the price variable.[15]

Existing Methods Existing literature has used different empirical strategies to deal with price endogeneity. In particular, the literature in Industrial Organization has developed estimation procedures with instruments for prices. For instance, Berry et al. (1995) use competitors' product characteristics, Hausman (1996) and Nevo (2000) use product's price on other markets, while Foster et al. (2008) rely on estimated physical productivities. However, these instruments are not valid in the presence of *unobserved* vertical differentiation.[16] As a consequence, these instruments cannot be used in our context. Indeed, trade data contain no product characteristics, except for the classification in product categories. Despite a narrow definition of these categories (8-digit CN classification present in our data has around 8,000 positions), there is still a wide scope for (unobserved) vertical differentiation within each category.[17]

Some strategies for demand estimation with trade data exist at the country level. Khandelwal (2010) and Hallak and Schott (2011) use instrumental variables approaches. Their strategy are not suited to firm-level demand estimation as their instruments vary at the market level, not across firms within a market. Feenstra (1994) and Broda and Weinstein (2010) respectively develop and refine a very influential demand estimation using country-level trade data. Their identification exploits the heteroskedasticity of supply and demand shocks. Although there strat-

[15]This attenuation bias will certainly be magnified by the flow fixed effects we use in our estimation. In fact, in the time series of a trade flow, the measurement error may represent a larger share of the variation of unit values than in the cross-section.

[16]Berry et al. (1995), Hausman (1996) and Nevo (2000) all study specific markets, for which they clearly observe different varieties of a good, as well as their characteristics, reducing the possibility for unobserved quality differences. In a different setup, Foster et al. (2008) and Handbury (2012) estimate demand functions for a wide range of products, but either restrict their analysis to homogeneous products or use barcode-level data, which rule out the possibility of unobserved quality differences.

[17] Consider cars, for instance. This product category contains multiple cn8 position, among which position 8703 21 10 'new and used vehicles, with spark-ignition internal combustion reciprocating piston engine'. There is clearly room for vertical differentiation across different exporting firms within this position.

egy could be applied to firm-level trade data, it involves an orthogonality assumption between demand and supply shocks which is likely to be violated in the presence of vertical differentiation (e.g., if quality is costly).

Literature on demand estimation with trade data is scarcer at the firm-level. Roberts et al. (2017) and Gervais (2015) use firms' wages and physical productivities as instruments for prices. These instruments are only valid if product quality is constant over time within the firm. For instance, if a firm upgrades its quality, it might need more workers per physical unit of output. In that case physical productivity is (negatively) correlated to quality and IV estimates of σ would be biased downward. Khandelwal et al. (2013) construct a firm-level quality measure by calibrating a CES demand system with price-elasticity estimates from Broda and Weinstein (2006). Conceptually, this approach raises two concerns. First, it implicitly inherits the identifying assumptions from Broda and Weinstein (2006). We explained above that these assumptions are problematic in the presence of vertical differentiation. Second, Broda and Weinstein (2006) estimates are obtained from country-level data. Elasticity may differ at the micro and the macro level,[18] which would generate biases in estimated firm-level quality.

Because existing methods do not lend themselves to our exercise, we develop a new instrumental strategy, robust to unobserved and time-varying quality differences within product categories.

A New Instrument for Prices at the Firm-level The approach developed in this paper takes advantage of the information coming from the importing activity of exporters. We use real exchange rates fluctuations faced by importing firms to instrument prices of exported goods. The basic idea is that real exchange rate shocks on a firm's imports are cost shocks. As the firm passes these cost shocks through to its export prices, sales adjust and the demand function is identified. In order to generate firm-specific exchange rate shocks, we take advantage of the fact that the spatial structure of imports varies across firms.

To gain insight into the identification, let us study the example of two firms selling in a same market. One firm imports from the United States, while the other imports from Europe. An appreciation of the dollar would induce an increase of the export price of the former, leaving unchanged the price of the latter. The response of these firms' relative sales to the change in their relative prices identifies the price-elasticity of demand. Importantly, these relative real exchange rate shocks across firms are likely to be exogenous to relative demand shocks. As such, they are not correlated with changes in quality or mark-up that are the two sources of endogeneity in prices. Therefore, exchange rates movements on imports constitutes a valid instrument to estimate the price-elasticity of demand. The next subsection discusses this assumption. It acknowledges situations where it is likely to be violated and adjusts the econometric specification accordingly.

To construct this instrument, we take advantage of two sources of variations at the firm-level: the set of countries a firm imports from and the share of these imports in the production cost of

[18] See Imbs and Méjean (2015) or Chetty (2012) for instances where the price elasticity depends on the level of aggregation considered.

the firm. First, we construct an import-weighted log real exchange rates defined as

$$\overline{\log rer}_{ft_0t} = \sum_c \omega_{cft_0} \times \log(\widetilde{\text{rer}}_{ct}), \tag{5}$$

where $\omega_{cft_0} = \frac{m^{(d)}_{cft_0}}{\sum_{c=1}^{C} m_{cft_0}}$ is the share of imports of differentiated goods of firm f from source country c, and $\widetilde{\text{rer}}_{ct}$ is the real exchange rate from home (France in our application) to country c at time t, deviated from its trend. The import weights are constructed at the initial date t_0, and measure the share of imports classified as 'differentiated' in the classification from Rauch (1999), $m^{(d)}_{cft_0}$, in the total imports of the firm. This restriction allows us to exclude homogeneous goods traded on organized exchanges for which firms can easily find substitutes when facing adverse exchange rates movements. On the contrary, imports of differentiated goods are more exposed to these movements because it is more difficult for firms to switch to a different supplier. Moreover, we construct $\widetilde{\text{rer}}_{ct}$ as deviations of the real exchange rates from country-specific trends. Specifically, we define $\log(\widetilde{\text{rer}}_{ct})$ using direct quotation and CPI indices as

$$\log(\widetilde{\text{rer}}_{ct}) = \log(\text{rer}_{ct}) - \hat{\rho}_c \times t \qquad \text{with} \quad \text{rer}_{ct} = \text{er}_{ct} \frac{\text{CPI}_{ct}}{\text{CPI}_{France,t}}$$

and $\hat{\rho}_c$ is estimated from regressing $\log(\text{rer}_{ct})$ on a set of country-specific trends. We decide to account for these trends to eliminate exchange rate variations driven by long-run movements. These trends can be problematic for two reasons: first, firms can design contracts to protect themselves against these changes. Second, these movements can be anticipated by firms and therefore are not orthogonal to initial sourcing decisions.

To obtain our final instrument, we interact this import-weighted exchange rate with the share of these imports in the operating costs of the firm. Formally, we define our instrument as

$$\overline{RER}_{ft_0t} = \overline{\log rer}_{ft_0t} \times \frac{\sum_t m_{ft}}{\sum_t OC_{ft}}, \tag{6}$$

with m_{ft} and OC_{ft} respectively the total imports and the operating costs of firm f at date t.[19] The motivation for interacting the RER shocks on imports with the import share is simply to get a well-defined cost shifter at the firm-level. If we were to omit the import share from the formula of $\overline{RER}_{ft_0t}$, our instrument would not capture the fact that two firms facing a given RER shock may experience different cost shocks depending on the role of their imports in the production process.

This instrument generates a cost shifter at the firm level that will identify the elasticity of demand, as firms pass this cost shock into their export prices. In a CES demand system where all firms are atomistic, this pass-through should be equal to one as firms maintain a constant mark-up over their marginal costs. However, if individual firms have an effect on the price index of the nest in which they are operating, this mark-up is not constant and firms feature heterogeneous pass-through. In order to capture this potential heterogeneity, we create an additional instrument

[19] Note that we compute the import share in the operating costs from the entire sample period available, rather than a specific reference date such as t_0. We made this decision because of discrepancies in the coverage of the two datasets we have access to. This methods avoids loosing firms that do not appear in the firm-level dataset in later years.

by interacting $\overline{RER}_{ft_0t}$ with ms_{fgdt_0}, defined as the market share of exporter f in its HS6 product category in market d at the initial date t_0.[20] Our motivation for doing so follows Amiti et al. (2014) which shows that in an oligopolistic model with nested-CES preferences, this market share is a sufficient statistic to capture this heterogeneity in pass-through. Specifically, in presence of these nested preferences, export prices of firms with larger market should respond less strongly to our instrument.

Finally, we create a third instrument based on the lagged real exchange rates faced by firms. The production of many goods span more than a year. As a consequence, we expect that cost shocks on imports purchased in the previous year might also generate an increase in the current price charged by an exporter. This instrument used a similar set of weights than our main instrument, but relies on real exchange rates at time $t-1$.

We conclude the presentation of the instruments with three remarks. First, the instrument is orthogonal to measurement errors on unit values as its construction does not involve information on exports. Therefore, our instrumental strategy deals with the measurement errors problem existing when estimating demand functions using unit values. Second, similar instruments have been used in a series of recent international trade contributions (see Brambilla et al. (2012) or Bastos et al. (2018)). In these papers, the export-weighted exchange rate generates exogenous change in firms' destination portfolio. In our case, the import-weighted average exchange rate creates exogenous firm-specific cost shifters due to the mechanical increase of the price of imported inputs. Lastly, we are not the first paper looking at the pass-through from the cost of imported input to export prices. Amiti et al. (2014) and Berman et al. (2012) run the same type of regression using respectively Belgian and French customs data. However, the motivation for their analysis differs greatly from ours. While, they are interested in the heterogeneity of the pass-through across firms, we only use the effect of exchange rates on export prices as a first stage to a demand function estimation.

2.3 Discussion of the Identification

There are a few mechanisms that could affect the exogeneity of the instrument. First of all, the instrument is constructed from import shares, which are potentially endogenous to quality. Put simply, higher quality firms most likely import from countries with a stronger currency, from where they can source higher quality inputs. Therefore, we expect the instrument to be positively correlated to quality in the cross-section of firms. If not controlled for, this correlation would induce the price elasticity of demand (which is negative) to be biased upward.[21] To address this problem, we decide to introduce fixed effects to capture time-invariant differences across firms. As a result, the identification of the parameters is in the time series of export prices and export volumes.

Specifically, we introduce exporting spell fixed effects in our empirical specification. We define a spell as a sequence of consecutive years during which a firm-product-destination triplet

[20] We compute the market share at the HS6 level, rather than HS8, because this is the most disaggregated level available for a large range of country. We use the BACI dataset to define the total size of the market.

[21] In the cross-section of firms, the instrument is likely to be positively correlated to quality. So, provided that higher quality goods are more expensive, an increase in the value of the instrument is associated to an increase in both prices and the demand shifter. Hence the upward bias.

is exported. Moreover, we use the first year of this spell as the reference year t_0 from which we construct the import weights that form the basis of our instrument. For instance, if Renault exports cars to Argentina from 2000 to 2003, stops exporting in 2004, and resumes from 2005 to 2007, then the exports of the variety Renault-Car-Argentina will correspond to two spells whose initial dates are respectively $t_0 = 2000$ and $t_0 = 2005$. Therefore, the instrument for the first exporting spell will be based on importing weights from the year 2000 while the second will use 2005 as reference year. Since the instrument is constructed using time-invariant import shares, its time series variations are fully driven by firm-specific exchange rates dynamics and not contaminated by (endogenous) import share dynamics. Moreover, it allows us to create an instrument with time-invariant weights that are closer to the current importing weights. This leads to stronger instruments while using weights that are plausibly exogenous to any change in quality decisions made by the firm during the exporting spell.

Beside the endogeneity of the import shares, another potential threat to the identification comes from the dual impact of exchange rates variations on firm performance. While a change in exchange rates can increase input prices, it can also affect the competitiveness of firms on foreign markets. This is a concern to us as it suggests that our instrument could be correlated to a firm's demand shifter. In reality, this is not an issue with the structural demand equation we consider. As one can see from the demand function (4), the competitiveness effect will be fully captured by destination-product-year fixed effect μ_{gdt}.

Moreover, exchange rate variations can directly cause quality adjustments. Bastos, Silva, and Verhoogen (2018) show that an exchange rate shock may induce a firm to upgrade its quality if it improves its competitiveness in rich destination markets. Similarly, Bas and Strauss-Kahn (2015) show that a change in tariffs or exchange rates on imported goods can lead firms to adjust their product quality. This import side effect is based on the premise that source countries produce inputs of different qualities. When an exchange rate shock makes imports from high (low) input quality countries more affordable, a firm upgrades (downgrades) the quality of its imported inputs, and output quality adjusts accordingly.

However, even if firm-level quality adjustments actually take place as the real exchange rate fluctuates, it is not clear what would be the sign, if any, of the resulting correlation between quality and our instrument. An increase in $\overline{RER}_{ft}$ can equally result from the appreciation of the currency of a rich source country as of the currency of a poor source country. So the existence of a bias on price-elasticity is unclear. Nonetheless, we take a conservative approach and neutralize the effect of exchange rates on quality by adding controls to the estimation. Namely, we incorporate the import weighted average GDP per capita of the firm as well as the export weighted average GDP per capita to the demand equation. The formula of these controls is:

$$\begin{cases} \overline{gdpc}_{ft}^{\text{exp}} & = \sum_c \omega_{cft}^{\text{exp}} \times \log(\text{gdpc}_{ct}) \\ \overline{gdpc}_{ft}^{\text{imp}} & = \sum_c \omega_{cft}^{\text{imp}} \times \log(\text{gdpc}_{ct}) \end{cases}. \tag{7}$$

These terms aim to capture quality adjustments following changes in the set of countries the firm imports from and exports to. The implicit assumption here is that GDP per Capita proxies the quality of inputs supplied by a country.[22] In the mechanism described above, exchange rates are suspected to affect quality only through an impact on a firm's spatial structure of imports. Controlling for that structure of exports thus makes the instrument orthogonal to the demand residual.

One last threat to the identification can come from the endogenous selection in trade flows. It has been extensively documented that trade data are very sparse.[23] If firms decide to stop exporting when they face an adverse shock, our estimation procedure will underestimate the price and export adjustments to exchange rate movements. To account for this selection bias, we follow Fitzgerald and Haller (2018) and Fontagné et al. (2018) by limiting our sample to long-lasting exporting spells. The idea is that firms with long exporting spells are far away from exit thresholds. As a consequence, their exit probability is close to zero such that they are unlikely to generate a selection bias. Therefore, we use exporting spells lasting more than 6 years to measure the full effect of our instrument.[24] Even though this adjustment does not account for the extensive margin response to cost shocks, it allows us to consistently estimate the demand equation faced by these firms.

Finally, we include an additional control to the specification to account for the partial-year effect that might contaminate our quality measures. Recent papers such as Berthou and Vicard (2015) and Bernard, Boler, Massari, Reyes, and Taglioni (2017) have documented that the construction of trade statistics in calendar year leads to systematic lower sales when a firm enters a market. This effect comes from the fact that firms are likely to enter at any time, leading to partial calendar years. To account for these systematic deviations, we add a specific dummies, $entry_{fst}$, equal to one in the first year of an export spell.

2.4 Econometric specification

Consistently with the above discussion, our econometric specification will proceed in two steps. In a first step, we regress the exported price of the firm on the instrument, RER_{ft_0t}, spell and market fixed effects, and the controls defined in equation (7). Bearing in mind that the reference year t_0 is the initial year of an export spell and that the index s characterizes a spell number for a firm f, destination d, and product g triplet, the formal expression of the first stage is

$$\log p_{fgdt} = \eta \overline{RER}_{ft_0t} + \beta_1 \overline{gdpc}_{ft} + \beta_2 entry_{fgdt} + \delta_{fgds} + \delta_{gdt} + u_{fgdt} \tag{8}$$

with $\overline{RER}_{ft_0t}$ our main instrument, $\overline{gdpc}_{ft}$ a vector containing the two controls defined in equation (7) and δ_{gdt} are market-year fixed effects. δ_{fgds} is a full set of export spell fixed effects. If a variety f-g-d is not exported continuously over the period but rather in 2 spells s and s' for instance, then two fixed effects δ_{fgds} and $\delta_{fgds'}$ are estimated for that variety.

Using the predicted values of exporting prices from this first stage, we can then estimate the structural equation (4) in a second stage:

$$\log q_{fgdt} = -\sigma_j \widehat{\log p}_{fgdt} + \alpha_1 \overline{gdpc}_{ft} + \alpha_2 entry_{fgdt} + \gamma_{fgds} + \gamma_{gdt} + \varepsilon_{fgdt} \tag{9}$$

[22] In line with this assumption, Schott (2004) shows evidence that richer countries specialize in the export of higher quality goods.

[23] See Blum et al. (2013) for instance.

[24] Figure 3 in appendix C shows that it is the minimum number of years needed to eliminate the selection bias.

in which γ_{fgds} and γ_{gdt} are spell and market-year fixed effects. The estimation of this equation is consistent if the structural error ε is orthogonal to our set of instruments. As we argue in the previous paragraphs, we believe this condition is reasonable given our specification. In equation (9), demand equation is identical to structural demand equation (4) except that we now impose our measure of quality, $\tilde{\lambda}_{fgdt}$, to take following form:

$$\tilde{\lambda}_{fgdt} = \hat{\alpha}_1 \overline{gdpc}_{ft} + \hat{\alpha}_2 entry_{fgdt} + \hat{\gamma}_{fgds} + \hat{\varepsilon}_{fgdt}. \tag{10}$$

In the next section, we implement this methodology using French customs data. Then, we assess its effectiveness by comparing our estimates of the elasticity of demand, and of product quality to existing measures.

3 Data and Demand Estimation Results

In this section, we apply our estimation strategy to French exporting firms using customs data. We start by describing the data we use, and provide descriptive statistics showing that they suit our exercise. Then, we report our results on the price elasticity and show that these estimates are almost systematically larger, in absolute values, than corresponding OLS estimates. This is strongly suggestive that the use of our IV estimation corrects the endogeneity bias described in section 2.2. Finally, we estimate product quality by separately estimating demand function (4) for different product categories.

3.1 Data

We exploit two sources of data. Our main source is firm-level trade data collected by French customs administration. These data provide a comprehensive record of the yearly values and quantities exported and imported by French firms from 1997 to 2010. Trade flows are disaggregated at the firm, country and eight-digit product category of the combined nomenclature.[25] Imports and exports are reported separately.

Our second dataset is the BRN ("Bénéfice réel net"). It covers all French firms with revenue larger than 763,000 euros, and is constructed from reports of French firms to the tax administration. This dataset has been widely used in the literature (see Eaton et al. 2011 or Berman et al. 2012 for instance). We use it mainly for two purposes: constructing the share of imports in firm total costs and correlating our quality estimates with firm-level characteristics such as wages or the number of employees.

Before implementing the estimation, we perform a series of operations to clean the data. In particular, information on quantities is known to be noisy in trade data. To mitigate this issue, we drop observations that display large variations in unit values from year to year.[26] Moreover, because of changes in the HS classification across years, we apply the algorithm described in Pierce and Schott (2012) in order to obtain well-defined and time invariant product categories.

[25] Only annual values which exceeds a legal threshold are included in the dataset. For instance, in 2002, this threshold was 100,000 euros. This cutoff is unlikely to significantly affect our study since, this same year, the total values of flows contained in the dataset represented roughly 98 percent of the aggregated estimates of French international trade.

[26] Appendix A provides the details of the cleaning procedure.

Descriptive Statistics The empirical strategy described in the previous section requires large variations in the data. First, our set of instruments relies on variations across firms in the set of countries they import from. Second, the large number of fixed effects included in the regression requires enough observations to identify variations across varieties within markets and across time within varieties. Table 1 provides statistics regarding the amount of variation contained in the data. Due to the large number of fixed effects required for the estimation procedure, many observations will not be helpful in identifying the estimated demand elasticity.[27] Therefore, we separately report statistics for the full sample and the sample that contributes to the parameters identification. For the latter, it is important to note that our preferred specification restricts our sample to exporting spells lasting more than six years.[28]

Table 1: Descriptive Statistics

	p5	p25	p50	p75	p95	mean
Full sample:			N = 21 624 643			
# **Source countries by firm**	0	0	0	3	16	3.14
# **Observations by exporting spell**	1	1	1	2	5	1.75
# **Varieties by export market**	1	1	2	5	21	5.87
Final estimating sample:			N = 3 481 154			
# **Source countries by firm**	0	2	9	17	36	11.86
# **Observations by exporting spell**	4	7	9	12	14	9.71
# **Varieties by export market**	2	2	3	6	19	6.33

Notes: An observation is an export flow at the firm, nc8 product, destination, year level. An exporting spell is a set of consecutive export flows for a firm-destination-nc8 product triplet. An export market is a nc8 product-destination-year triplet, and a variety a firm-nc8 product pair. Final estimating sample excludes exporting spells shorter than 7 years.

First, table 1 reports the number of source countries per firm over the period: focusing on the estimating sample, more than 50 percent of firms import from at least 9 sources and the average number of source countries per firm is equal to 11.9. This is reassuring that there is substantial variation across firms regarding their exposure to exchanges rates movements. Moreover, note that a significant share of exporters do not import and therefore will not be affected by variations in foreign exchange rates. Second, rows two and three report the numbers of observations by market and varieties. Even though many observations will not contribute to the identification, more than 50% of the destination product year fixed effects are identified and a significant share of exporting spells are long enough to identify variations across years between firms: selecting exporting spells lasting more than 6 years and eliminating observations that are absorbed by market-specific fixed effects, we still have more than three millions observations to identify our parameters.

The instrument crosses two informational sources: import shares and real exchange rates. Figure 1 provides information on the latter source by reporting the 1997-2010 evolution of real

[27] In particular, this is the case of exporting spells that only last one year.

[28] Note that some spells have less than 7 observations in the second panel of table 1. Those are cases where the exporting spell lasts more than 6 years but for which some observations are absorbed by market-specific fixed effects, and therefore are excluded from the estimating sample.

exchange rates for the top 5 exporters to France over the period. Even though the real exchange rate movements of Euro zone countries are solely due to inflation after 1999, this figure shows large and non-monotonic movements in exchange rates: this is likely to affect firms that import relatively more from these specific countries.

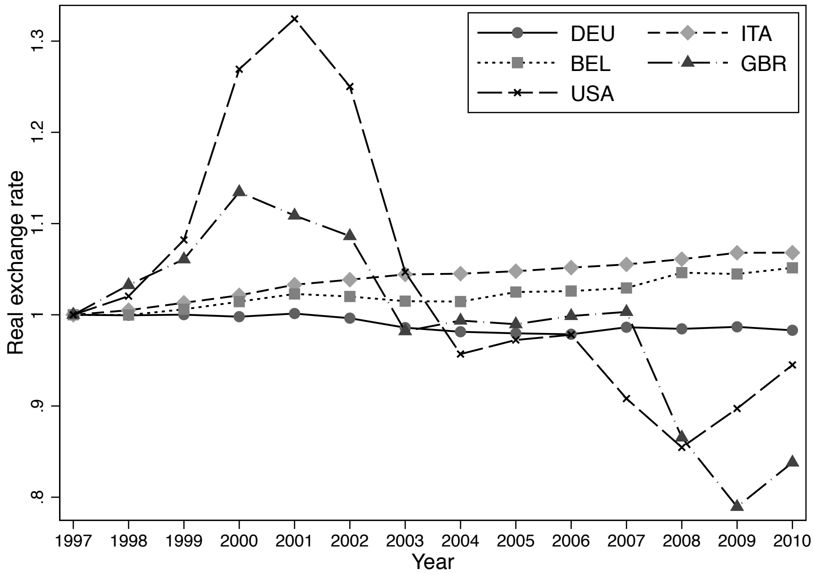

FIGURE 1: RER (1997-2010) - Top Source Countries

Notes: Real exchange rates are calculated as $e_{Euro,st} \times \frac{\text{CPI}_{st}}{\text{CPI}_{France,t}}$ where $e_{Euro,st}$ is the direct nominal exchange rate from Euro to j's currency at date t. CPI is the consumer price index. After 1999, Real-exchange-rate movements of Euro zone countries are solely due to inflation. 1997 real exchange rates are normalized to one.

3.2 Pooled Industries Results

In order to describe the effectiveness of our instrumental strategy, we first present results obtained by pooling the data, before moving on to separately estimating the model on different product categories. The pooled results are reported in table 2. Panel A and panel B respectively contain first stage and second stage results. All regressions in this table are obtained including firm-destination-cn8 product-spell fixed effects and destination-cn8 product-year fixed effects.[29] Moreover, since our instrument is defined at the firm-year level, standard errors are clustered at the firm level to account for the Moulton factor and potential autocorrelation within the panel dimension.

In Panel A, we report the first stage of the 2SLS procedure and the reduced form effect of the instruments on export volumes. The main instrument, $\overline{RER}_{ft_0t}$, has a positive and significant effect on the export price charged by firms: using the full sample, column (2) shows that a firm's export prices increase by 0.16 percent on average when its real exchange rates on imports increase

[29]Estimation of linear equations with two sets of high-dimensional fixed effects and unbalanced panel, as is the case in our estimation, is cumbersome. To perform the estimation, we rely on the algorithm developed in Correia et al. (2016). This algorithm first demeans the variables along the two sets of fixed effects. Parameters of interest are then estimated using demeaned variables.

Table 2: Results on Pooled Data

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	OLS		**2SLS**			**Reduced Form**	
Panel A (1^{st} STAGE)		*log price*				*log qty*	
$\overline{RER}_{ft_0t}$		0.16***	0.25***	0.23***	0.25***	-1.06***	-1.16***
		(0.041)	(0.065)	(0.076)	(0.067)	(0.30)	(0.34)
$\overline{RER}_{ft_0t-1}$				0.029			0.16
				(0.068)			(0.29)
$\overline{RER}_{ft_0t} \times ms_{fpdt_0}$					-0.41		
					(0.83)		
$Entry_{fpdt}$		0.0023***	0.0015	0.0014	0.0015	-0.95***	-0.95***
		(0.00070)	(0.0020)	(0.0020)	(0.0020)	(0.012)	(0.012)
$\overline{gpdc}_{ft}^{\mathbf{imp}}$		0.0037***	0.0060**	0.0060**	0.0061**	0.020*	0.020*
		(0.0014)	(0.0025)	(0.0025)	(0.0025)	(0.012)	(0.012)
$\overline{gpdc}_{ft}^{\mathbf{exp}}$		0.0033*	0.0053	0.0053	0.0054	0.24***	0.24***
		(0.0017)	(0.0041)	(0.0042)	(0.0042)	(0.021)	(0.021)
Panel B (2^{nd} STAGE)	*log qty*						
Log price ($-\hat{\sigma}$)	**-0.78*****	**-3.03****	**-4.26*****	**-4.18****	**-4.39*****		
	(0.0080)	(1.39)	(1.65)	(1.65)	(1.69)		
$Entry_{fpdt}$	-0.32***	-0.31***	-0.95***	-0.95***	-0.94***		
	(0.0041)	(0.0053)	(0.014)	(0.014)	(0.015)		
$\overline{gpdc}_{ft}^{\mathbf{imp}}$	0.012*	0.020**	0.045***	0.045***	0.046***		
	(0.0066)	(0.0090)	(0.017)	(0.017)	(0.018)		
$\overline{gpdc}_{ft}^{\mathbf{exp}}$	0.15***	0.16***	0.26***	0.26***	0.27***		
	(0.010)	(0.012)	(0.028)	(0.028)	(0.029)		
Sample	Full	Full	$>$ 6 yrs	$>$ 6 yrs	$>$ 6 yrs	$>$ 6 yrs	$>$ 6 yrs
Kleibergen-Paap F-stat		14.3	14.5	7.3	7.6		
Hansen p-value				0.4	0.00		

Notes: The full sample contains 10 762 689 observations while the restricted sample contains 3 481 154 observations. Firm×prod×dest×spell and prod×dest×year fixed effects are included in all regressions. Standard errors in parentheses are clustered at the firm level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

by 1 percent. When limiting our sample to long export spells in column (3), which mitigates selection bias, we obtain a larger response with a pass-through of 25 percent. This imperfect pass-through of imported exchange rates to export prices can be explained by two main reasons: first, a large literature documents the imperfect pass-through of exchange rates movements to import prices. Therefore, we can expect import prices paid by French firms to not entirely follow exchange rates changes. Second, this instrument only measures with error the import costs faced by French firms. In particular, the use of import weights from the initial period creates measurement error in the true cost of importing, hence driving the coefficient toward zero. As a consequence, even if French firms fully pass-through import costs on their export prices, the estimated coefficient of this first stage is likely to be lower than one. Importantly, this is not an issue for our empirical strategy: even with this incomplete pass-through, the instrument generates exogenous variations in the export prices of French firms, which is sufficient to identify the price elasticity of demand.

In columns (4) and (5), we test the relevance of additional instruments. We see from these columns that the instrument constructed using the lag real exchange rate does not significantly explain export prices. Similarly, the interaction of our instrument with the market share of the firm in the export market does not significantly affect the degree of pass-through of French exporters. This absence of results could be explained by the fact that French firms are relatively small in foreign markets. Therefore, they do not feature market power that would lead them to strategically adjust their mark-up. Finally, columns (6) and (7) directly look at the reduced form impact of our instruments on export volumes. We show that a positive import cost shock significantly reduces the volume exported by a firm. However, the lagged cost shock once again does not appear to have a significant effect on export volume.

The relative strength of these different sets of instruments is measured by the Kleibergen-Paap F statistic reported at the end of table 2. The F-statistic of 14.5 is above the thresholds commonly used to detect weak instruments when we only used our main instrument $\overline{RER}_{ft_0t}$. However, including the lagged instrument or the interaction between our instrument and the market share significantly reduces this F-statistic: introducing more instrument mechanically reduces the F-statistic by decreasing the number of degree-of-freedom, without bringing enough additional explanatory power to compensate.

Turning to the second stage, panel B of table 2 reports the demand elasticity estimates using our different specifications. We start by reporting the estimation of the demand equation using ordinary least squares (OLS). The purpose of this specification is to serve as a reference point to assess the impact of our instrumentation on the estimates. With OLS, we obtain an estimated price-elasticity of demand of $\hat{\sigma}^{OLS} = 0.78$, well below the usual estimates found in the literature. This is not surprising as this estimates is polluted by measurement errors and endogeneity between demand and supply shocks.

By contrast, all the specifications using instrumental variables lead to a larger and consistent elasticity in absolute values. In column (2), we report the result of the 2SLS estimation on the full sample. We find a large negative estimates of -3.03, much more in line with estimates found in the literature. However, this estimate is likely to be biased from endogenous selection, since many firms respond to cost shocks by entering or exiting foreign markets. Specification (3) reports our preferred specification, which uses our instrument and controls for selection bias by reducing the sample to long exporting spells. We obtain a large price-elasticity, equal to 4.26. Adding more instruments in specifications (4) and (5) does not significantly change this estimates as these additional instruments brings no explanatory power in our first stage.

Overall, these estimated parameters are consistent with other elasticities found in the literature, even though we obtain larger estimates than studies employing datasets with a wide range of products. For instance, Foster, Haltiwanger, and Syverson (2008) obtains a mean estimate of 2.41 with eleven homogeneous industries, Handbury (2012) finds a mean of 1.97 with 149 industries, and Gervais (2015) a median of 2.11 with 504 products.[30] More recently, Fontagné et al. (2018) estimates an elasticity around 5 when accounting for selection bias using long exporting

[30] Other papers estimating firm-level demand functions include Nevo (2000), who finds estimates between 2.2 and 4.2 in the cereal industry, Dubé (2004) who gets estimates between 2.11 and 3.61 in the soft drinks industry.

spells.

Finally, all our control variables play a role in the estimation. Firms entering the export markets record much lower export volume in their first year due to the partial-year effect. Second, firms who export and import from richer countries charge higher prices and export larger volumes. This is consistent with Bastos, Silva, and Verhoogen (2018), which predicts that, following an increase in the average GDP per capita of its destinations, a firm should upgrade its product, generating a positive impact on prices and on sales. Similarly, the average GDP per capita of source countries is positively correlated with output prices and sales, suggesting that $\overline{gdpc}_{ft}^{\text{imp}}$ actually proxy for the quality of imported inputs. These results imply that exchange rates could affect the quality choices made by firms, by changing their average destination or origin countries. However, this variation is orthogonal to our instrument which maintains import shares constant within a spell.

Having demonstrated the relevance of our empirical strategy, we now turn to a more disaggregated estimation of demand elasticities, taking into account heterogeneity across product categories.

3.3 Demand Estimation by Product category

In this section, we describe the results obtained when replicating the instrumentation strategy for fifteen product categories. To perform the estimation, we employ specification (3) from table 2 that uses a single instrument, $\overline{RER}_{ft_0t}$, the three controls ($\overline{gpdc}_{ft}^{\textbf{imp}}$, $\overline{gpdc}_{ft}^{\textbf{exp}}$ and $Entry_{fpdt}$), and the two sets of fixed effects.[31] Moreover, we restrict our sample to export flows that last more than 6 years, to account for endogenous selection. In table 3, we report three estimates for each product category. First, we estimate the elasticity of substitution using OLS (columns 1 and 2) to obtain a benchmark to which to compare our parameters estimated with instrumental variables. Second, we present the results of the 2SLS estimation performed separately for each product category: column 3 reports the estimated coefficient, column 4 the standard error of the parameter and column 5 the F-stat of the first stage describing the strength of the instrument. Finally, columns 6 and 7 report the point estimate when the first stage is common across all product categories, but the price elasticity is allowed to vary across product group.

First of all, estimates obtained with OLS display the same issue observed in the aggregate data: due to measurement errors and endogeneity between demand and supply shocks, the parameters are biased toward zero. On the contrary, the estimated elasticities in the IV specification tend to be larger, in absolute values, relative to the OLS. This confirms that the instrument does correct for endogeneity as expected. However, due to the reduction in the number of observations, the first stage of the 2SLS does not appear strong enough, as illustrated by the F-statistic that does not exceed the critical value conventionally adopted of 10 to reject weak instruments. As a consequence, the IV estimates display a large variance which sometimes leads to unrealistic values of the price elasticity. In particular, industries relying on homogeneous inputs such as 'Mineral products' or 'Stone, Glass', have F-statistic close to zero, which generates very noisy estimates of their price elasticity. By contrast, industries with a larger F-stat, hence a stronger

[31] We choose this specification because it delivers the highest F statistic in the first stage and appears more robust in our results.

TABLE 3: Price-elasticity estimates ($-\hat{\sigma}$) for different product categories

Product categories	**OLS**		**IV**			**IV (single FS)**		
	Coef	SE	Coef	SE	F-stat	Coef	SE	N
Animal Products	-0.88***	(0.042)	1.82	(5.17)	1.21	**-10.1***	[5.63]	190 887
Vegetable Products	-0.75***	(0.028)	10.4	(24.2)	0.29	**-3.28**	[5.01]	195 596
Foodstuffs	-0.94***	(0.018)	-1.03	(4.24)	0.81	**-1.70**	[5.34]	409 242
Mineral Products	-0.84***	(0.083)	-171.6	(5.8e3)	0.00	**-3.78**	[7.79]	24 125
Chemicals and Allied	-0.90***	(0.021)	-1.34	(1.61)	2.11	**-4.12**	[2.97]	374 169
Plastics, Rubbers	-0.92***	(0.025)	-1.26	(1.20)	7.78	**-2.43**	[2.95]	227 886
Skins, Leather	-0.74***	(0.042)	-18.8	(41.9)	0.20	**-5.90****	[3.00]	56 251
Wood, Wood products	-0.86***	(0.023)	-3.06*	(1.75)	5.52	**-1.47**	[2.83]	178 783
Textiles	-0.70***	(0.038)	-5.82	(4.45)	3.91	**-4.42**	[2.72]	663 856
Footwear, Headgear	-0.68***	(0.061)	-7.09	(5.21)	2.30	**-6.79****	[3.37]	65 454
Stone, Glass	-0.84***	(0.034)	-1837.6	(5.5e5)	0.00	**-4.97**	[3.03]	80 316
Metals	-0.81***	(0.025)	-2.43	(3.03)	1.82	**-3.01**	[2.69]	260 784
Machinery, Electrical	-0.87***	(0.021)	-2.87**	(1.32)	6.14	**-3.78**	[2.57]	392 429
Transportation	-0.79***	(0.031)	-6.02	(5.65)	1.26	**-8.95****	[4.41]	113 832
Miscellaneous	-0.79***	(0.023)	-4.96	(3.52)	2.87	**-4.02***	[2.41]	247 544

Notes: Estimates in columns "OLS" and "IV" are obtained by estimating equation (4) separately for each industry, respectively by OLS and 2SLS. Estimates in column "IV (single FS)" is obtained by estimating a single first stage and a second stage where the price-elasticity is allowed to vary across industries. Controls for GDP per capita ($\overline{gpdc}_{ft}^{\mathbf{exp}}$ and $\overline{gpdc}_{ft}^{\mathbf{imp}}$) and for partial-year effect ($Entry_{fpdt}$) are included in all regressions. Firm×Prod×Dest×Spell and Prod×Dest×Year fixed effects are included in all regressions. IV specifications use $\overline{RER}_{ft_0t}$ as instrument. Standard errors are clustered at the firm level and standard errors for the "IV (Single FS)" specification are obtained through 1000 bootstrap replications using firm as the sampling unit. Column "F-stat" reports the value of the Kleibergen-Paap F-stat. * p<0.1, ** p<0.05, *** p<0.01

first stage, display more realistic price elasticity estimates.

To circumvent the weakness of this first stage, we also present estimates based on a single first stage for all industries. This first stage is identical to table 2, except that we allow demand elasticities to vary across product categories in the second stage. To estimate these elasticities, we obtain predicted export prices from specification (3) in table 2, and then regress export volumes on these predicted prices, interacted with industry dummies. The standard errors of these estimates are computed from 1000 bootstrap samples, to take into account the variability of the first stage predictions. In order for this empirical strategy to be valid, we need to assume that the pass-through of import exchange rates to export prices is similar across industries. In particular, it requires that firms do not adjust mark-ups differently across sectors in response to a cost shock. In a CES demand system without oligopolistic power, this assumption is satisfied since mark-ups do not respond to supply or demand shocks. However, in a model such as Atkeson and Burstein (2008) where mark-ups are a function of the nested market share of a firm, this assumption would be rejected. Since we did not find evidence of heterogeneous pass-through in the aggregate data, we believe this assumption is reasonable in our context.

This procedure with a single first stage allows us to obtain much more robust estimates of the price elasticities. Using this specification, all product categories display an elasticity of substitution larger than one, ranging from 1.47 to 10.1. Moreover, they appear relatively consistent with 2SLS estimates for industries with a moderate F-statistic, which reassures us regarding the validity of the specification using a single first stage. Because of the robustness of this specification, we use those estimates to construct our quality measures in the rest of the paper.[32] However, note that these estimates still display large standard errors: using exporting spells that last more than 6 years strongly reduce the size of the sample, which ultimately results in less precision for our estimates.

In order to make sense of the price-elasticity variation across sectors, we compare our estimates to industry-level characteristics that are related to the degree of substitution across varieties. First, we compare the elasticities to those obtained by Soderbery (2015) in a paper that refines the estimation strategy of Broda and Weinstein (2006).[33] Similarly, we relate our estimates with the Sutton measure, that characterizes the scope for vertical differentiation of an industry.[34] We expect industries with a large scope for vertical differentiation to have lower price elasticity of their demand.

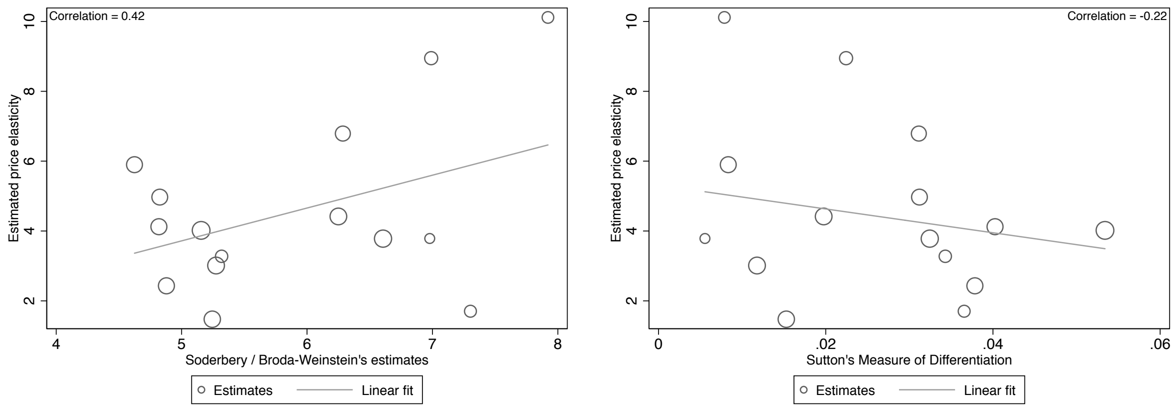

FIGURE 2: Estimated price elasticity and existing estimates.

Notes: Each circle corresponds to a product category in Table 3. The size of a circle is proportional to the inverse of the standard errors of the price elasticity estimate. The vertical axis the estimated price-elasticities while the horizontal axis is the demand elasticity estimates from Soderbery (2015), improving on Broda and Weinstein (2006) on the left figure and the measure of vertical differentiation from Sutton (2001) on the right figure. The lines are the predicted values of the OLS regressions using the inverse of the standard errors as weights.

We can see in the left quadrant of figure 2 that the correlation between the two sets of price elasticities is positive, although non-significant, which should not come as a surprise given the small number of data points (15). Similarly, we do find in the right quadrant a negative relationship between the estimated price elasticities and the measure of vertical differentiation from Sutton (2001): we estimate a lower price elasticity of demand for industries with a larger scope for vertical differentiation. Even though these relationships are not statistically significant, they are reassuring regarding the relevance of our estimates.

[32]Moreover, we show in appendix C that most of our findings hold when using an aggregated price elasticity of 4.26, as estimated in specification (3) of table 2.

[33]Soderbery (2015) estimates are demand elasticities faced by countries (not firms) on their exports and are defined at the 4-digit level. We therefore aggregate them up to the 1-digit level using a simple arithmetic mean.

[34]The Sutton measure, initially at the ISIC rev. 2 level is converted to the 1-digit level to be compared with our price elasticities.

4 Analysis of Estimated Quality

In this section, we document features of $\tilde{\lambda}_{fpdt}$, our quality measure obtained from the demand estimation. Specifically, we obtain our quality measure from equation (10), using estimates from column 6 of table 3.[35] We start by briefly decomposing the variance of $\tilde{\lambda}_{fpdt}$ along different dimensions. Then, in order to assess the relevance of our measure, we document its correlation with existing but sporadic measures of quality, as well as with firm-level data and industry-level measures of vertical differentiation. Finally, we show contexts in which our measure might be preferred to other variables commonly used to proxy product quality.

As a first way to describe our estimates of quality, we provide a variance decomposition in table 4. Here, it is important to remember that the quality measure is obtained at the firm $\times$ product category $\times$ destination $\times$ year level. Moreover, quality is defined relatively to the average quality in the market. Therefore, it defines a position over the quality ladder in a market, rather than an absolute quality which can be compared across markets.

A first observation about table 4 is that firm-specific variation in quality only explains 20 percent of total variation. The dispersion of quality is significantly better predicted by variety-specific effects. Indeed, 52 percent of this quality dispersion is captured by time-invariant variety-specific effects, and 60 percent by time-variant variety fixed effect. Therefore, there is substantial variation in quality across products within firms, and this variety-specific quality is fairly constant over time. Table 4 is also suggestive of the presence of important market-specific tastes, or of the fact that firms adjust the quality to their product depending on the country they serve, which explains that the R^2 increases from 52 percent to 78 percent once we allow for quality to vary across destinations within a given firm-product variety.

TABLE 4: Variance decomposition of the quality measure ($\tilde{\lambda}_{fpdt}$)

Set of Fixed Effects	R^2
Firm	0.20
Firm $\times$ Prod	0.52
Firm $\times$ Prod $\times$ Dest	0.78
Firm $\times$ Year	0.25
Firm $\times$ Prod $\times$ Year	0.60

Notes: Each R^2 is obtained from the separate regression of the quality measures on fixed effects only.

Having briefly described the sources of variation of this measure, we next document its consistency with existing measures of vertical differentiation.

[35] In appendix C, we show that the results of this section are robust when constructing our quality measure using the aggregate estimate of the price-elasticity.

4.1 Consistency tests

Comparison with expert assessed quality First, we relate the estimated quality to one of the only objective product quality measure existing in the literature. Crozet et al. (2012) take advantage of expert ratings for Champagne to analyze the importance of quality in explaining international trade flows at the firm level. These expert assessed ratings (initially from Juhlin (2008)) are expressed in number of stars ranging from 1 to 5, one being the lowest quality. To test the relevance of our quality measure, we non-parametrically regress our measure $\tilde{\lambda}_{fpdt}$ for Champagne exports over the number of stars assigned by experts.[36]

Table 5: Correlation with ratings of Champagne

	Dep. variable: estimated quality $\tilde{\lambda}_{fpdt}$ Coef.	se
2 Stars	0.32**	(0.15)
3 Stars	0.49**	(0.19)
4 Stars	1.38***	(0.22)
5 Stars	1.77***	(0.19)
N	32 448	
R^2	0.096	

Notes: Champagne ratings from Juhlin (2008). A larger number of star means a higher expert assessed quality. We drop non-Champagne exports of Champagne producers. Standard errors in parentheses are clustered at the firm level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

From table 5, it appears that our measure of quality is monotonically increasing with the number of stars assigned by Juhlin (2008). Even though Champagne is a specific good in many dimensions, this case study provides compelling evidence of the relevance of our measure of quality.

Correlation with firms' characteristics In order to further assess the validity of our quality measure, we relate the measure $\tilde{\lambda}_{fpdt}$ to firms' characteristics obtained from the BRN dataset. In particular, this allows us to inspect how $\tilde{\lambda}_{fpdt}$ is related to the average wage paid by a firm, a measure commonly used to proxy the qualitative of a firm's workforce. Table 6 reports these correlations.

It appears from table 6 that quality is strongly correlated with the average wage of the firm. In order to control for the size of the firm, we also add as regressors the number of employees and the total stock of capital employed by the firm. Adding these controls does not affect the correlation between quality and average wage. Moreover, this link is even stronger when we include destinations, product and year fixed effects, such that firms with higher wages

[36] We thank the authors for sharing their data.

Table 6: Correlation with firms' characteristics

	(1)	(2)	(3)	(4)	(5)	(6)
	Dependent variable: estimated quality ($\tilde{\lambda}_{fpdt}$)					
	No fixed effects		Dest, prod and year FE		Dest×prod×year FE	
log(wage)	0.83***	0.76***	0.97***	0.94***	1.09***	1.06***
	(0.020)	(0.018)	(0.021)	(0.020)	(0.023)	(0.022)
log(employment)		0.067***		0.10***		0.12***
		(0.0098)		(0.010)		(0.012)
log(capital)		0.053***		0.061***		0.070***
		(0.0069)		(0.0073)		(0.0081)
N	13 391 548	13 201 466	13 391 528	13 201 442	13 297 957	13 096 862
R^2	0.0068	0.012	0.0094	0.016	0.022	0.032

Notes: The variable log(*wage*) is obtained by taking the logarithm of the total wage bill divided by the number of employees. Specifications (1) and (2) have a non-reported constant. Standard errors in parentheses are clustered at the firm-year level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

systematically have higher product quality, relative to other exporting firms in the same market. These results provide further evidence that our measure captures heterogeneity across firms that is related to vertical differentiation and product quality differences.

Length of quality ladders and vertical differentiation As a final test of our quality estimation, we construct a market specific measure of the "length" of the quality ladder. Following Khandelwal (2010), this length is obtained by taking the difference, for any product × destination × year combination, between the 95th and the 5th percentile of the quality distribution. This quantity may be interpreted as a revealed measure of the degree of vertical differentiation in a market. We start by verifying that this measure is correlated with the quality ladders obtained by Khandelwal (2010). In his work, these measures are obtained at the industry level by comparing exporting countries' qualities in the US market. In contrast, we obtain this measure by comparing French exporters' qualities in different industries and destinations.

Table 7 shows the positive link between the quality ladders constructed from our quality measures, and the ones from Khandelwal (2010). This positive correlation is not significant when we include all markets, yet it appears positive and robust once we exclude markets in which the number of firms is too small to reliably compare the 5th and the 95th percentile. This correlation remains stable as we control for market destinations and time fixed effects such that the identification is obtained across product categories in the same country at the same time.

These different tests demonstrate the pertinence of $\tilde{\lambda}_{fpdt}$ to describe the quality of the good produced by the firm, and the vertical differentiation across firms. In order to further establish the relevance of our measure, we show in the next subsection why it may be preferred to prices, a popular proxy for quality.

Table 7: Length of quality ladders and vertical differentiation

	(1)	(2)	(3)	(4)	(5)	(6)
		Dependent variable: quality ladders				
	All markets		More than 5 firms		More than 20 firms	
Khandelwal (2010)'s ladders	0.062	0.063	0.48***	0.47***	0.65***	0.63***
	(0.12)	(0.12)	(0.17)	(0.17)	(0.24)	(0.24)
Dest FE	Yes	No	Yes	No	Yes	No
Year FE	Yes	No	Yes	No	Yes	No
Dest×year FE	No	Yes	No	Yes	No	Yes
N	1 452 329	1 452 226	540 705	540 456	137 068	136 865
R^2	0.055	0.058	0.041	0.049	0.085	0.10

Notes: Each coefficient in the table comes from separate regressions. Khandelwal (2010)'s measure is averaged from 10-digit product categories to 6-digit categories. Standard errors in parentheses are clustered at the 6-digit product level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

4.2 How well do Prices proxy for Quality?

As highlighted in the introduction, the absence of product quality measures has led researchers to use proxies, directly available in many datasets, to measure product quality. Among these proxies, the price of a good is probably the most common in microeconomic studies. However, an important drawback of this proxy is that it conflates many factors that are not related to product quality, and ignores characteristics that are not accounted in the price of a product but still enter the consumers' valuation of that product. In particular, goods produced by firms with low productivity, and therefore high price, would be assigned a high quality according to prices.

In this subsection, we confirm this imperfect relationship between prices and quality. First, we show that prices and $\tilde{\lambda}_{fpdt}$ are more correlated in industries with larger vertical differentiation. In these industries, variations in prices are more driven by quality variations than by cost variations, which generates a stronger correlation with our quality measures. Second, we describe a situation in which using prices as proxy for quality can be misleading: we show that, unlike prices, the quality measure of a firm increases as this firm increases its number of destinations or varieties.

Prices and vertical differentiation In table 8, we display the correlation between prices and quality $\tilde{\lambda}_{fpdt}$ across firms. All regressions include destination×product×year fixed effects such that the relation is identified within a market. One can see from this table that export prices and quality are positively correlated, justifying the use of prices as a proxy for quality. This positive correlation is significant in the cross-section of firms within a market, but also when tracking firms over time: firms which move their prices over time simultaneously move their quality in the same direction (specifications (2), (4) and (6)).

If prices and quality measures are correlated across and within firms, table 8 also shows that this correlation is stronger in markets with a larger degree of differentiation. The coefficient on the interaction between quality and quality ladder shows that markets with a long quality

TABLE 8: Correlation between prices and quality

	(1)	(2)	(3)	(4)	(5)	(6)
	Dependent variable: log price					
	All markets		More than 5 firms		More than 20 firms	
Quality	0.67***	0.59***	0.66***	0.59***	0.65***	0.58***
	(0.00063)	(0.00071)	(0.00077)	(0.00083)	(0.0011)	(0.0012)
Quality × quality lad.	0.16***	0.16***	0.17***	0.15***	0.18***	0.18***
	(0.00088)	(0.00064)	(0.0011)	(0.00084)	(0.0020)	(0.0016)
Firm×Prod×year FE	No	Yes	No	Yes	No	Yes
N	20 048 513	13 830 187	16 147 196	11 683 523	10 225 420	7 608 688
R^2	0.49	0.88	0.50	0.88	0.50	0.89

Notes: Quality ladder is the difference between the 95th and 5th percentiles of the quality distribution within a market, normalized to have a mean of zero and a variance of one. Quality measures and prices are also normalized to have zero mean and a standard deviation of one within markets. Each regression includes product×dest×year fixed effects. Standard errors in parentheses are clustered at the product×dest×year level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

ladder also display a tighter link between prices and product quality.[37] As a consequence, in markets with little vertical differentiation, prices might contain little information about quality, but instead, be mostly driven by productivity.

Prices, quality and the extensive margins Despite the positive relationship between price and product quality measure, the use of prices as proxy for quality can lead to erroneous conclusions. In this section, we show this is the case when studying the link between the extensive margins of a firms (its number of destinations and products) and its product quality.

Recent research in international trade has studied the link between quality and export performance. While prices are often positively correlated with export performance in developing countries (see Kugler and Verhoogen (2012) or Manova and Zhang (2012) for instance), the positive correlation in developed economies is less established. The main reason for this ambivalence comes from the fact that prices conflate productivity and quality, with opposite predictions on the link between prices and export performance. On the contrary, a measure of product quality should be positively correlated with measures of performance.

We test this relationship by looking at the link between prices, quality and the number of destinations or varieties of a firm. Table 9 shows regressions of the quality $\tilde{\lambda}_{fpdt}$ and the price of a good on the number of destinations or varieties a firm is exporting. Importantly, we include market fixed effects to identify this relationship within a specific market. Moreover, an additional complication comes from the fact that adding a new destination or a new variety might lead to a decrease of the average quality exported by the firm.[38] To account for this endogenous sorting,

[37]In these regressions, we normalize prices and quality such that, within markets, they are centered around zero and have a standard deviation of one. This normalization aims at avoiding the dispersion of qualities to affect the slope between prices and qualities.

[38]This can be easily seen in a model in which firms have a core-product of higher quality. In this framework, adding a product generates a reduction in the average quality produced because this additional variety (or

we identify the link between quality and the extensive margin within a specific firm-destination-product dimension: we use both fixed effects and first differences to look at the change in quality for existing destinations or products. In other words, we capture the change in quality of existing destinations or varieties when new destinations or products are added: if quality is positively correlated with export performance, the quality of existing varieties should increase as the firm expands its scope.

TABLE 9: Quality and extensive margins

	(1)	(2)	(3)	(4)
Dependent variable:	*estimated quality* $\tilde{\lambda}_{fpdt}$		*log price*	
log(# destinations)	0.66***	0.57***	0.0027***	0.0047***
	(0.0047)	(0.0048)	(0.00075)	(0.00048)
log(# products)	0.40***	0.37***	-0.0060***	-0.0017***
	(0.0041)	(0.0038)	(0.00071)	(0.00043)
Firm-dest-product FE	Yes	No	Yes	No
First difference	No	Yes	No	Yes
N	13 830 187	8 298 223	13 830 187	8 298 223

Notes: Each coefficient is obtained from a separate regression. Each regression includes product×dest×year fixed effects. Standard errors in parentheses are clustered at the firm×year level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 9 shows that this correlation is indeed positive when using our measure of quality. The addition of a destination or a variety takes place simultaneously to an increase in the quality of existing destinations and varieties. However, this prediction is mixed when looking at prices: while expanding to a new destination is associated with a price increase in existing destinations, existing products see a price reduction when a new product is introduced. This mixed result highlights the drawbacks of prices: by conflating many factors, they might lead to somewhat conflicting conclusions. In this specific context, the addition of a new destination in a firm portfolio seems to be mostly driven by demand factors, while the addition of new goods is correlated with price reductions at the firm level. By contrast, the use of estimated quality measures clearly shows that product quality is positively associated with the export performance of the firm.

5 Conclusion

A recent literature has evidenced that product quality has implications for key economic outcomes such as firms' profitability or welfare inequalities. These findings make it crucial to understand the determinants of quality at the firm-level. In this paper, we have provided a necessary tool to

destination) is of lower quality. Manova and Yu (2017) develops such a model.

pursue this research agenda. Namely, we have proposed a novel strategy to estimate time-varying quality at the firm-level. Our strategy is robust to unobserved vertical differentiation and only requires firm-product level information on prices, sales and imports by country.

We first show that the measures of quality obtained from this method are consistently related to a range of measures: estimated quality is positively correlated with the average wage paid by firms, with direct measures of product quality from outside sources. Moreover, we use our firm-level quality estimates to build an index of vertical differentiation and we find a strong correlation with comparable indices in the literature.

To highlight the relevance of this work, we then study the link at the firm level between product quality and export performance. We show that the most common proxy for quality, export prices, displays a mixed relationship between quality and exporters' scope. Instead, when using the estimated quality measures, we show that firms adding varieties to a market or destinations to their portfolio do so as the quality of their existing varieties increase, implying a positive link between export performance and quality. In light of these results, we believe that the methodology developed in this paper could help exploring existing and new questions in which the quality of the good produced by exporters plays an important role.

References

Amiti, M., O. Itskhoki, and J. Konings (2014): "Importers, Exporters, and Exchange Rate Disconnect," *American Economic Review*, 104, 1942–1978.

——— (2016): "International Shocks and Domestic Prices: How Large Are Strategic Complementarities?" Working Paper 22119, National Bureau of Economic Research.

Atkeson, A. and A. Burstein (2008): "Pricing-to-market, Trade Costs, and International Relative Prices," *American Economic Review*, 98, 1998–2031.

Bas, M. and V. Strauss-Kahn (2015): "Input-trade liberalization, export prices and quality upgrading," *Journal of International Economics*, 95, 250–262.

Bastos, P., J. Silva, and E. Verhoogen (2018): "Export destinations and input prices," *American Economic Review*, 108, 353–92.

Berman, N., P. Martin, and T. Mayer (2012): "How Do Different Exporters React to Exchange Rate Changes?" *Quarterly Journal of Economics*, 127, 437–492.

Bernard, A. B., E. A. Boler, R. Massari, J.-D. Reyes, and D. Taglioni (2017): "Exporter dynamics and partial-year effects," *American Economic Review*, 107, 3211–28.

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63, 841–90.

Berry, S. T. (1994): "Estimating Discrete-Choice Models of Product Differentiation," *RAND Journal of Economics*, 25, 242–262.

Berthou, A. and V. Vicard (2015): "Firms' Export Dynamics: Experience versus Size," *The World Economy*, 38, 1130–1158.

Blum, B. S., S. Claro, and I. J. Horstmann (2013): "Occasional and perennial exporters," *Journal of International Economics*, 90, 65–74.

Brambilla, I., D. Lederman, and G. Porto (2012): "Exports, Export Destinations, and Skills," *American Economic Review*, 102, 3406–38.

Broda, C. and D. E. Weinstein (2006): "Globalization and the Gains from Variety," *Quarterly Journal of Economics*, 121, 541–585.

——— (2010): "Product Creation and Destruction: Evidence and Price Implications," *American Economic Review*, 100, 691–723.

Chetty, R. (2012): "Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply," *Econometrica*, 80, 969–1018.

Correia, S. et al. (2016): "REGHDFE: Stata module to perform linear or instrumental-variable regression absorbing any number of high-dimensional fixed effects," *Statistical Software Components*.

Crozet, M., K. Head, and T. Mayer (2012): "Quality Sorting and Trade: Firm-level Evidence for French Wine," *Review of Economic Studies*, 79, 609–644.

Dubé, J.-P. (2004): "Multiple Discreteness and Product Differentiation: Demand for Carbonated Soft Drinks," *Marketing Science*, 23, 66–81.

Eaton, J., S. Kortum, and F. Kramarz (2011): "An Anatomy of International Trade: Evidence from French Firms," *Econometrica*, 79, 1453–1498.

Feenstra, R. C. (1994): "New Product Varieties and the Measurement of International Prices," *American Economic Review*, 84, 157–177.

Fitzgerald, D. and S. Haller (2018): "Exporters and shocks," *Journal of International Economics*, 113, 154–171.

Fontagné, L., P. Martin, and G. Orefice (2018): "The international elasticity puzzle is worse than you think," *Journal of International Economics*, 115, 115–129.

Foster, L., J. Haltiwanger, and C. Syverson (2008): "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?" *American Economic Review*, 98, 394–425.

Garcia-Marin, A. (2014): "Income Distribution, Quality Sorting and Trade," *Manuscript*.

Gervais, A. (2015): "Product Quality and Firm Heterogeneity in International Trade," *Canadian Journal of Economics*, 48, 1152–1174.

Hallak, J. and P. Schott (2011): "Estimating Cross-Country Differences in Product Quality," *Quarterly Journal of Economics*, 126, 417–474.

Hallak, J. C. and J. Sivadasan (2013): "Product and Process Productivity: Implications for Quality Choice and Conditional Exporter Premia," *Journal of International Economics*, 91, 53–67.

Handbury, J. (2012): "Are Poor Cities Cheap for Everyone? Non-Homotheticity and the Cost of Living Across US Cities," *Manuscript*.

Hausman, J. A. (1996): "Valuation of New Goods under Perfect and Imperfect Competition," in *The Economics of New Goods*, University of Chicago Press, 207–248.

Hottman, C., S. J. Redding, and D. E. Weinstein (2016): "Quantifying the Sources of Firm Heterogeneity," *Quarterly Journal of Economics*.

Imbs, J. and I. Méjean (2015): "Elasticity Optimism," *American Economic Journal: Macroeconomics*, 7, 43–83.

Johnson, R. C. (2012): "Trade and Prices with Heterogeneous Firms," *Journal of International Economics*, 86, 43–56.

Juhlin, R. (2008): *Champagne Guide*, Richard Juhlin Publishing AB.

KHANDELWAL, A. (2010): "The Long and Short (of) Quality Ladders," *Review of Economic Studies*, 77, 1450–1476.

KHANDELWAL, A., P. SCHOTT, AND S. WEI (2013): "Trade Liberalization and Embedded Institutional Reform: Evidence from Chinese Exporters," *American Economic Review*, 103, 2169–2195.

KUGLER, M. AND E. VERHOOGEN (2012): "Prices, Plant Size, and Product Quality," *Review of Economic Studies*, 79, 307–339.

MANOVA, K. AND Z. YU (2017): "Multi-product firms and product quality," *Journal of International Economics*, 109, 116–137.

MANOVA, K. AND Z. ZHANG (2012): "Export Prices across Firms and Destinations," *Quarterly Journal of Economics*, 127, 379–436.

NEVO, A. (2000): "Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry," *RAND Journal of Economics*, 31, 395–421.

PIERCE, J. R. AND P. K. SCHOTT (2012): "Concording US Harmonized System Codes over Time," *Journal of Official Statistics*, 28, 53–68.

RAUCH, J. E. (1999): "Networks versus Markets in International Trade," *Journal of International Economics*, 48, 7–35.

REDDING, S. J. AND D. E. WEINSTEIN (2016): "Measuring Aggregate Price Indexes with Demand Shocks: Theory and Evidence for CES Preferences," Working Paper 22479, National Bureau of Economic Research.

——— (2017): "Aggregating from Micro to Macro Patterns of Trade," Working Paper 24051, National Bureau of Economic Research.

ROBERTS, M. J., D. YI XU, X. FAN, AND S. ZHANG (2017): "The Role of Firm Factors in Demand, Cost, and Export Market Selection for Chinese Footwear Producers," *The Review of Economic Studies*, 85, 2429–2461.

SCHOTT, P. K. (2004): "Across-product versus Within-product Specialization in International Trade," *Quarterly Journal of Economics*, 119, 647–678.

SODERBERY, A. (2015): "Estimating import supply and demand elasticities: Analysis and implications," *Journal of International Economics*, 96, 1–17.

SUTTON, J. (2001): *Technology and Market Structure: Theory and History*, The MIT Press.

VAN BEVEREN, I., A. B. BERNARD, AND H. VANDENBUSSCHE (2012): "Concording EU Trade and Production Data over Time," Working Paper 18604, National Bureau of Economic Research.

VERHOOGEN, E. (2008): "Trade, Quality Upgrading, and Wage Inequality in the Mexican Manufacturing Sector," *Quarterly Journal of Economics*, 123, 489–530.

Appendices

A Data preparation

We perform two main operations to prepare the final sample. First, we harmonize the product codes to obtain consistent categories across time. Then, we clean the dataset to take into account the existence of measurement errors in trade data.

Harmonization of product codes The product classification used by custom authorities is regularly updated to follow changes in product characteristics. We need to account for these changes to maintain a coherent set of product categories across time. To achieve this, we follow the procedure from Van Beveren et al. (2012) who apply the methodology from Pierce and Schott (2012) to European statistics. This allows us to obtain consistent product categories from 1997 to 2010.

Choice of units for quantity information Data on quantities are known to be subject to measurement errors, which could lead to spurious relationships between quantities and prices (computed by dividing values with quantities). Moreover, the customs statistics from France allows exporters to declare the quantities in two different units: the weight or a supplementary unit that is product specific and more relevant to describe the quantities of certain types of goods. Therefore, we decide to use the supplementary unit when at least 80 percent of firms in the category are providing this unit. Otherwise, we use the weight of the good as quantity.

Data cleaning After harmonizing quantities within product categories, we can compute prices as the export value divided by quantity. Then, because of the potential measurement errors in prices, we drop prices that display large variations from one year to another. In particular, given our identification strategy, we perform the following procedure:

- We declare a price p_{fpgt} as abnormal when $\log p_{fpdt} - \log p_{fpdt-1}$ is larger than one or lower than minus one.
- We declare a price p_{fpgt} as missing when the quantity for that observation is missing.
- We drop from the sample the entirety of an exporting spell that contains at least one abnormal or missing price.

By performing this cleaning procedure, we ensure that each exporting spell contained in our sample displays reasonable price changes across years.

B Correspondence with existing datasets

In this section, we specify how we merged existing datasets with the French customs to run the analysis in sections 3 and 4.

Elasticities of substitution from Soderbery (2015) Elasticity estimates are available from Anson Soderbery's personal website. From the LIML estimates available at the HS10 level, we compute the arithmetic average for each one of the 15 broad product categories, as presented in table 3, to be compared with our estimates at the product level.

Measure of vertical differentiation from Sutton (2001) This measure of vertical differentiation is defined as the ratio of R&D and advertising expenditures in an industry to the total sales of the industry, contained in the U.S Federal Trade Commission (TC) 1975 Line of Business Survey. We manually copied the information from Kugler and Verhoogen (2012) which provides information at the ISIC rev.2 4-digit industry level (column 1 of table A3 in the online appendix).

From this level of aggregation, we match the 4-digit industry level to the HS6 code level from the WITS concordance tables and compute the arithmetic average by 1-digit product categories, to be able to compare it with our elasticity estimates.

Champagne ratings from Juhlin (2008) The dataset of ratings for Champagne producers has been compiled by Crozet et al. (2012). The authors matched ratings contained in Juhlin (2008) to a unique firm identifier that allows concordance with administrative French datasets. We thank the authors for providing us the complete data with this identifier, allowing an easy match at the firm level, with the custom dataset. With this information, we regress these expert-assessed quality measures to our estimated quality of each Champagne producers. Importantly, in the regression reported in table 5, we only use estimated qualities for product code 22041011 that is specifically corresponds to sparkling wine from the Champagne region.

Firm-level dataset (BRN) The BRN dataset ("Bénéfice réel net") provides information at the firm-year level for all French firms with a revenue larger than 763,000 euros. This dataset can easily be matched with the custom data through a common firm identifier. From this dataset, we compute the average wage by dividing the total wage bill by the number of employees. By construction, the regressions in table 6 are based on the sample of firms present in both datasets.

Quality ladders from Khandelwal (2010) We obtained measures of quality ladders from Amit Khandelwal's personal website. We download estimates available at the HS10 level. Because hs codes can vary across countries beyond the 6-digit levels, we compute the average ladder length's by computing the arithmetic average by hs6 product group. From this level of aggregation, we can merge these quality ladders to our measures at the hs8 level. To account for the mismatch in the level of variation between the two measures, standard errors in table 7 are clustered at the hs6 level.

C Additional results

C.1 Quality results - Robustness

In this section, we report alternative results using quality measures constructed from a single aggregate demand elasticity estimated in table 2. Specifically, we use $\sigma = 4.26$ as estimated in

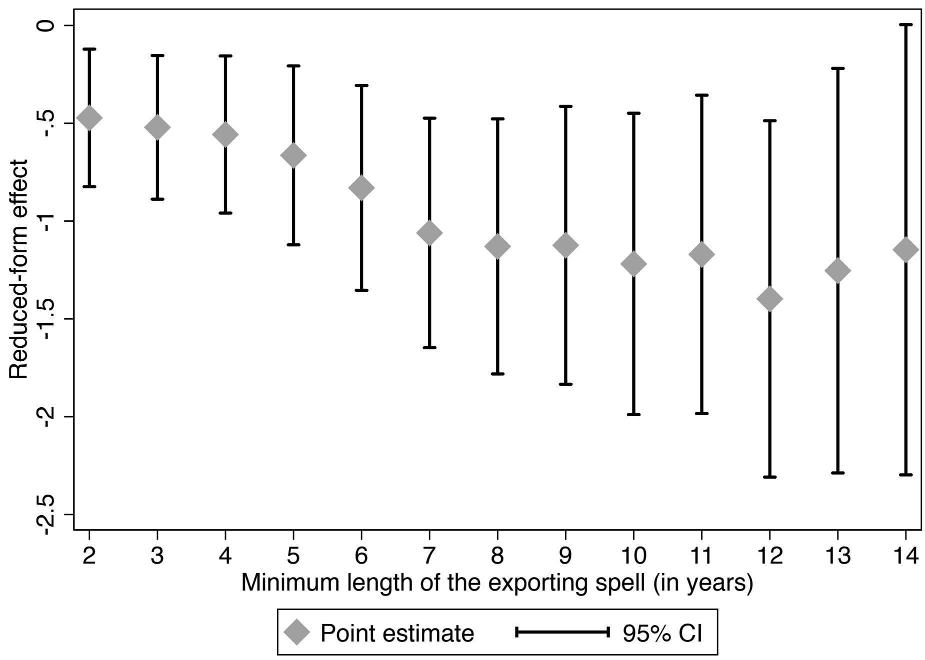

FIGURE 3: Reduced form impact across exporting spell length

Notes: This figure reports the regression of the instrument on the logarithm of export volumes, following specifications (6) in table 2. For each spell length in the figure, we estimate the reduced-form relationship including all exporting flows that are longer than this length.

column (4) of table 2. These results show that all the findings hold when using this alternative estimate of the price-elasticity.

TABLE 10: Correlation with ratings of Champagne

	Dep. variable: estimated quality ($\tilde{\lambda}_{fpdt}$)	
	Coef.	se
2 Stars	0.44***	(0.15)
3 Stars	0.75**	(0.18)
4 Stars	1.99***	(0.25)
5 Stars	3.33***	(0.59)
N	32 448	
R^2	0.19	

Notes: Champagne ratings from Juhlin (2008). A larger number of star means a higher expert assessed quality. We drop non-Champagne exports of Champagne producers. Estimated quality are constructed with the estimate from column 4 of table 2. Standard errors in parentheses are clustered at the firm level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

TABLE 11: Correlation with firms' characteristics

	(1)	(2)	(3)	(4)	(5)	(6)
	Dependent variable: estimated quality ($\tilde{\lambda}_{fpdt}$)					
	No fixed effects		Dest FE, prod FE, year FE		Dest×prod×year FE	
log(wage)	0.85***	0.80***	0.99***	0.98***	1.12***	1.11***
	(0.019)	(0.018)	(0.020)	(0.020)	(0.022)	(0.022)
log(employment)		0.088***		0.12***		0.14***
		(0.0095)		(0.010)		(0.011)
log(capital)		0.032***		0.037***		0.044***
		(0.0066)		(0.0070)		(0.0077)
N	13 391 548	13 201 466	13 391 528	13 201 442	13 297 957	13 096 862
R^2	0.0072	0.012	0.0095	0.016	0.021	0.030

Notes: The variable log(*wage*) is obtained by taking the logarithm of the total wage bill divided by the number of employees. Specifications (1) and (2) have a non-reported constant. Estimated quality are constructed with the estimate from column 4 of table 2. Standard errors in parentheses are clustered at the firm-year level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

TABLE 12: Length of quality ladders and vertical differentiation

	(1)	(2)	(3)	(4)	(5)	(6)
	Dependent variable: quality ladders					
	All markets		More than 5 firms		More than 20 firms	
Khandelwal (2010)'s ladders	-0.11	-0.11	0.26**	0.25**	0.43***	0.42***
	(0.094)	(0.094)	(0.11)	(0.11)	(0.14)	(0.14)
Dest FE	Yes	No	Yes	No	Yes	No
Year FE	Yes	No	Yes	No	Yes	No
Dest×year FE	No	Yes	No	Yes	No	Yes
N	1 452 329	1 452 226	540 705	540 456	137 068	136 865
R^2	0.064	0.067	0.077	0.088	0.15	0.18

Notes: Each coefficient in the table comes from separate regressions. Khandelwal (2010)'s measure is averaged from 10-digit product categories to 6-digit categories. Quality ladders are constructed with the estimate from column 4 of table 2. Standard errors in parentheses are clustered at the 6-digit product level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 13: Correlation between prices and quality

	(1)	(2)	(3)	(4)	(5)	(6)
	Dependent variable: log price					
	All markets		More than 5 firms		More than 20 firms	
Quality	0.75***	0.65***	0.74***	0.66***	0.75***	0.66***
	(0.00030)	(0.00047)	(0.00038)	(0.00052)	(0.00056)	(0.00069)
Quality × quality lad.	0.12***	0.12***	0.12***	0.097***	0.13***	0.10***
	(0.00026)	(0.00035)	(0.00035)	(0.00041)	(0.00063)	(0.00068)
Firm×Prod×year FE	No	Yes	No	Yes	No	Yes
N	20 048 513	13 830 187	16 147 196	11 683 523	10 225 420	7 608 688
R^2	0.94	0.99	0.93	0.99	0.93	0.99

Notes: Quality ladder is the difference between the 95th and 5th percentiles of the quality distribution within a market, normalized to have a mean of zero and a variance of one. Each regression includes product×dest×year fixed effects. Estimated quality are constructed with the estimate from column 4 of table 2. Standard errors in parentheses are clustered at the product×dest×year level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 14: Quality and extensive margins

	(1)	(2)	(3)	(4)
Dependent variable:	*estimated quality*		*log price*	
log(# destinations)	0.66***	0.57***	0.0027***	0.0047***
	(0.0047)	(0.0047)	(0.00075)	(0.00048)
log(# products)	0.39***	0.37***	-0.0060***	-0.0017***
	(0.0041)	(0.0038)	(0.00071)	(0.00043)
Firm-dest-product FE	Yes	No	Yes	No
First difference	No	Yes	No	Yes
N	13 830 187	8 298 223	13 830 187	8 298 223

Notes: Each coefficient is obtained from a separate regression. Each regression includes product×dest×year fixed effects. Estimated quality are constructed with the estimate from column 4 of table 2. Standard errors in parentheses are clustered at the firm×year level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$