Heterogeneous firms and asymmetric product differentiation*

Carlo Altomonte (Bocconi University)
Italo Colantone (Bocconi University)
Enrico Pennings† (Erasmus University Rotterdam)

March 2015

Abstract

We introduce asymmetric product differentiation in a model characterized by a linear demand system, endogenous markups and heterogeneous firms (as in Melitz-Ottaviano, 2008). In particular, a single industry is divided into a number of market segments, each characterized by a different degree of horizontal product differentiation. Such a setup allows us to explain, within a single theoretical framework, the non-linear relations between firm productivity, size and exporting behavior that have been documented by the empirical literature. The theoretical results are tested empirically by examining the performance of French wine producers operating in market segments characterized by different levels of horizontal product differentiation. Such segments are identified using the official classification of French wines based upon the controlled denomination of origin, i.e. the "Appellation d'Origine Contrôlée" (AOC) system.

JEL classification: L10, F10
Keywords: Firm heterogeneity, product differentiation, productivity, exports

* **Acknowledgements**: The paper has greatly benefited from the comments of the editor and an anonymous referee. We wish to thank seminar participants at the ERWIT-CEPR meeting in Rome, the ETSG Conference in Lausanne, the University of Brescia, and Politecnico of Milan for their very useful comments. The authors retain the sole responsibility for their errors and omissions.

† *Corresponding author:* Department of Applied Economics, Erasmus University Rotterdam, Burg. Oudlaan 50, 3062PA Rotterdam, Netherlands. Phone +31 10 4082166. Fax +31 10 4089149. pennings@ese.eur.nl

1 Introduction

The literature on heterogeneous firms has pointed at productivity as being the main driver of multiple dimensions of firm performance, e.g. size, survival and export status. In particular, it has been shown that more productive firms tend to be larger, less likely to exit and more likely to export.[1] And yet, more recent evidence has also shown some limitations of the predictive power of productivity in explaining cross-firm differences in performance. For instance, a firm's total factor productivity (TFP) does not unambiguously determine its export status: contrary to the standard theoretical predictions, there are many small (low productivity) firms that export, while some large (high productivity) firms only sell domestically (Hallak and Sivadasan, 2013). The same one-to-one relation between productivity and firm size, often taken for granted in this literature, has also been found to be less straightforward than originally thought (Brooks, 2006; Foster et al., 2008).

In this paper, we explain the fact that equally productive firms operating in the same industry display differences in their size, or exporting status, by modeling the fact that firms compete in different market segments within the same industry, and these segments are characterized by asymmetric degrees of substitutability (i.e. horizontal product differentiation) across varieties. The idea is that equally productive firms operating in different segments face different price elasticities of demand, which allows them to charge different markups and to achieve a different size and exporting status in equilibrium, thus leading to a closer matching of theory and empirical evidence.

We first describe a model that captures the key theoretical idea, building upon the Melitz-Ottaviano (2008) framework, and then test the model's main predictions using data from the French wine industry. In particular, we compare the performance of French wine producers operating in market segments that are characterized by different levels of horizontal product differentiation. Market segments are identified using the official classification of French wines according to the controlled denomination of origin, that is the world-famous "Appellation d'Origine Contrôlée" (AOC) system. Within this system, which has been introduced by law in 1935, a wine possesses a certain AOC geographical label 'if and only if' it is produced within a well delimited regional area. Based on official data from the French Institute of Origin and Quality (INAO), we are able to identify ten AOC geographical areas, mapping the classical subdivision of the French wine industry: Alsace, Bordeaux, Bourgogne, Champagne, Jura-Savoie, Languedoc-Roussillon, Loire, Provence, Rhône and South-West.

The AOC classification, being geographically based and well rooted in the history of France, allows for an exogenous partition of the industry into market segments according to the location of the producers. In particular, we rely on a number of earlier studies that have shown how the AOC system is essentially associated with a higher level of horizontal differentiation, while it

[1] Melitz (2003) and Bernard et al. (2003) provide the seminal theoretical contributions in a standard CES demand setting and in a Ricardian framework, respectively, while Melitz and Ottaviano (2008) extend the framework to linear demand systems.

cannot guarantee a higher level of quality.[2] That is, while AOC-denominated wines are not necessarily better than others, they stand out in terms of typicity and distinct taste, as determined by the specific characteristics of the "terroir" (from *terre*, i.e. land in French) where they are produced. As such, they are found to be less substitutable among each other, as compared to non-AOC wines. Our main empirical strategy exploits this characteristic of the French wine industry, treating all the AOC producers as competing in a single high-differentiation segment, as opposed to the low-differentiation segment of non-AOC producers (i.e. firms producing outside of the AOC areas identified above). In a refinement of the analysis, we then treat each AOC region as a separate market segment, each characterized by a measurable level of horizontal differentiation that is given by the number of "sub-appellations", i.e. smaller regional appellations within each of the ten larger AOC regions.

The empirical analysis focuses on around 1,000 wine producers, observed over the time span 1999-2008, and is based on balance sheet data from AMADEUS, including information on firms' geographical location and export activities. We find evidence that, at low levels of productivity, AOC firms are larger than non-AOC firms, while this relation is inverted as TFP grows beyond a certain productivity threshold. The latter evidence confirms our key theoretical finding on the relation between productivity and size across market segments: for a relatively inefficient firm, it is easier to attain a relatively higher level of sales if varieties are more differentiated (AOC segment), while a very productive firm can leverage upon its efficiency to a larger extent if varieties are less differentiated from each other (non-AOC segment). This complex relation between productivity and size across market segments, confirmed by the data, is our key novel and distinctive finding, as alternative approaches based on vertical differentiation arguments (i.e. quality differences across varieties) cannot lead to the same result.

Our empirical findings are robust to the inclusion of several controls, to the use of different TFP estimates, and cannot be obtained in the case of alternative food industries, where an equivalent geographical-based market segmentation is not present, i.e. meat and bread products. Moreover, results also hold as we differentiate across AOC areas, using the number of sub-appellations codified within each area as a regional-specific measure of horizontal differentiation. As a corollary result, we also present evidence in line with a second theoretical prediction: conditioning for the level of productivity, AOC firms are more likely to export than non-AOC firms. Thus, also the relation between firm productivity and export engagement within an industry is crucially moderated by the segment-specific (asymmetric) degree of product differentiation. The latter constitutes an additional contribution of this paper.

Our theoretical approach fits a variety of industries, such as food or beverages, where geographical designations create an asymmetry in the degree of horizontal product differentiation across varieties; or the clothing industry, where some producers sell relatively undifferentiated clothes through standard outlets (e.g. hypermarkets), while trendy firms like Zara or H&M

[2] See Gergaud and Ginsburgh (2008), Cross et al. (2011), Barham (2003), Veale and Quester (2008), Lecocq and Visser (2006).

produce and sell clothes of a similar quality, but invest in differentiating their products from competitors, e.g. through dedicated distribution chains. Similar examples would apply to industries where design and brand are in general important drivers of prices once controlling for quality (e.g. electronics, furniture, leather products).

This paper is related to a large body of recent research on heterogeneous firms, in which several authors have started to combine differences in productivity with additional dimensions of heterogeneity. In particular, a number of papers have explored the role of product quality, in order to explain the variation in firms' size, prices and export performance conditional on productivity.[3] In these studies, quality is in general modeled as a demand shifter, thus introducing vertical differentiation over the set of product varieties produced by firms.

Our paper is complementary with respect to the latter stream of studies, in that we explicitly abstract from quality differences in order to uncover the specific role of asymmetric horizontal differentiation. While modeling heterogeneity in terms of quality certainly broadens the scope of analysis with respect to earlier productivity-based studies, we show that differences in horizontal differentiation have distinctive and empirically verifiable implications on the relation between productivity and firm size that cannot be reproduced through a vertical differentiation (quality) argument. This allows us to go one step further with respect to the current attempts at reconciling theory and empirical evidence. Other recent papers have combined horizontal product differentiation with differences in product quality, in order to study the relation between productivity, prices and sales (Antoniades, 2012; Eckel et al., 2011; Di Comite et al., 2014). However, these papers have not modeled an asymmetric degree of horizontal differentiation across market segments, as all varieties are assumed symmetric from a demand perspective. Breaking such a symmetry constitutes the main novel feature of our approach.

The paper is organized as follows. In Section 2 we introduce the theoretical model. In Section 3 we derive the main predictions of our model. Section 4 presents the French wine industry and our data. Section 5 discusses the empirical evidence. Finally, Section 6 concludes.

2 Theoretical model

2.1 Model setup: consumers

We start with a model characterized by linear demand systems, endogenous markups and heterogeneous firms à la Melitz-Ottaviano (2008), in which we introduce some asymmetry in the degree of horizontal product differentiation across varieties. In particular, in our framework consumers choose between a homogeneous good and a continuum of differentiated varieties, indexed by $i \in \Omega$. However, Ω is now split in two separate subsets: Ω^l and Ω^h, where Ω^l is assumed to be the subset of varieties characterized by a "low" degree of horizontal differentiation (i.e. high

[3] See Johnson (2012), Verhoogen (2008), Kneller and Yu (2008), Hallak and Sivadasan (2013), Khandelwal (2010), Baldwin and Harrigan (2011), Kugler and Verhoogen (2012), Crozet et al. (2012), Spearot (2013), Crino' and Epifani (2012).

substitutability across varieties), while Ω^h contains those varieties characterized by a "high" degree of horizontal differentiation (i.e. low substitutability across varieties). Ω^l and Ω^h can thus be thought of as being two distinct market-segments within a narrowly defined industry.[4] For instance, in line with our empirical application, we may think of Ω^h and Ω^l as two different segments of the wine industry, e.g. highly differentiated varieties of AOC-denominated wines (Ω^h) vs. relatively less differentiated non-AOC wines (Ω^l).

Formally, considering an economy with L consumers, each supplying one unit of labour, the utility function of a representative consumer can be written as:

$$U = q_o + \sum_{\xi=l,h} \left[\alpha \int_{i\in\Omega^\xi} q_i di - \frac{1}{2}\gamma^\xi \int_{i\in\Omega^\xi} q_i{}^2 di - \frac{1}{2}\eta^\xi \left(\int_{i\in\Omega^\xi} q_i di \right)^2 \right] - \beta \int_{z\in\Omega^l} \int_{j\in\Omega^h} q_z q_j \; d_z d_j \qquad (1)$$

where $\xi = l, h$ indicates the low vs. high differentiation market segment, q_o stands for the consumption level of a homogeneous good (taken as a numeraire) and q_i represents the consumption level for each variety $i \in \Omega^\xi$. The parameter α captures the degree of vertical differentiation (i.e. quality) with respect to the numeraire, and is assumed to be constant across market segments. This allows us to focus solely on the implications of horizontal differentiation, abstracting from quality considerations. The parameters γ^ξ measure the segment-specific degree of product differentiation across varieties *within* each market segment. Consistent with the partition of Ω described above, we assume that $\gamma^h > \gamma^l$. The η^ξ parameters measure the degree of substitution of the differentiated varieties in each market segment with respect to the homogeneous good.

Parameter β indexes the substitutability pattern across varieties belonging to different market segments. Its inclusion constitutes the main departure from the Melitz-Ottaviano (2008) framework, where there is no partition of the industry in multiple segments, and thus no need to allow for substitutability of varieties between segments. And yet, our utility function is still consistent with the general formulation of the quasi-linear utility function presented in Ottaviano, Tabuchi and Thisse (1998), and it shares the same properties. Varieties are assumed to be more substitutable between the different market segments than with respect to the numeraire, in line with a general love for variety behavior, and thus we have that $\eta^\xi > \beta > 0$.[5]

A recent paper by Di Comite et al. (2014) provides a complementary extension of the Melitz-Ottaviano (2008) utility function. In particular, they allow for variety-specific vertical and horizontal differentiation, but the applicability of their model is limited to a single market segment, where varieties are characterized by the same substitutability. On the contrary, the primary focus of our paper is on the cross-segments implications of different degrees of horizontal

[4] The theoretical results are independent on the level of disaggregation of the industry in which firms are observed, as long as different market segments / degrees of product differentiation can be identified within the same industry.

[5] If β would be equal to zero, the consumers' choice problem would boil down to the less interesting case of two independent maximizations for the two market segments. Anyway, all the predictions of the model would still hold.

differentiation.

By solving the consumer problem, and assuming a positive demand for the numeraire good (q_0), we can obtain the inverse demand function for each variety in each of the two subsets Ω^l and Ω^h:

$$p_i^l = \alpha - \gamma^l q_i^l - \eta^l Q^l - \beta Q^h \tag{2}$$

$$p_i^h = \alpha - \gamma^h q_i^h - \eta^h Q^h - \beta Q^l \tag{3}$$

where $Q^{l,h} = \int_{i\in\Omega^{l,h}} q_i di$. By taking $\gamma^\xi q_i^\xi$ to the left-hand side and p_i^ξ to the right-hand side in both equations and subsequently integrating over all $i \in \Omega^\xi$, we get the following equations:

$$\gamma^l Q^l = N^l\alpha - N^l\overline{p}^l - N^l\eta^l Q^l - N^l\beta Q^h \tag{4}$$

$$\gamma^h Q^h = N^h\alpha - N^h\overline{p}^h - N^h\eta^h Q^h - N^h\beta Q^l \tag{5}$$

where N^l and N^h are the number of consumed varieties in the subsets Ω^l and Ω^h respectively, $\overline{p}^\xi = \frac{1}{N^\xi}\int_{i\in\Omega_c^\xi} p_i d_i$, and Ω_c^ξ is the subset of consumed varieties within Ω^ξ. The solution to the system of equations is:

$$\begin{bmatrix} Q^l \\ Q^h \end{bmatrix} = D \begin{bmatrix} \gamma^h + N^h\eta^h & -N^l\beta \\ -N^h\beta & \gamma^l + N^l\eta^l \end{bmatrix} \begin{bmatrix} N^l(\alpha - \overline{p}^l) \\ N^h(\alpha - \overline{p}^h) \end{bmatrix} \tag{6}$$

where $D = \left[\left(\gamma^l + N^l\eta^l\right)\left(\gamma^h + N^h\eta^h\right) - N^l N^h\beta^2\right]^{-1}$.

By substituting the aggregate demand for both segments in eq. 2 and 3, the inverse demand functions can be written in the following compact way:

$$p_i^l = f_1^l - \gamma^l q_i^l + f_2^l\overline{p}^l + f_3^l\overline{p}^h \tag{7}$$

$$p_i^h = f_1^h - \gamma^h q_i^h + f_2^h\overline{p}^l + f_3^h\overline{p}^h \tag{8}$$

where f_1^ξ, f_2^ξ and f_3^ξ are functions of N^l, N^h, β, γ^l, γ^h, η^l and η^h.[6] We can now derive the price conditions in order for a variety in segment ξ to display a positive consumption level (i.e. $q_i^\xi > 0$):

$$p_i^l < f_1^l + f_2^l\overline{p}^l + f_3^l\overline{p}^h \equiv p_{\max}^l \tag{9}$$

$$p_i^h < f_1^h + f_2^h\overline{p}^l + f_3^h\overline{p}^h \equiv p_{\max}^h \tag{10}$$

where $p_{\max}^\xi$ stands for the price level at which demand in segment ξ is just equal to zero. The total demand for goods in each market segment can then be expressed as:

$$Lq_{i,c}^\xi = \frac{L}{\gamma^\xi}(f_1^\xi - p_i^\xi + f_2^\xi\overline{p}^l + f_3^\xi\overline{p}^h) \tag{11}$$

where $q_{i,c}^\xi$ stands for the consumption of variety i in segment ξ. From here, exploiting the

[6]The expressions of these parameters are reported in Appendix 1.

property that $q^{\xi}_{i,c}(p^{\xi}_{\max}) = 0$, we can express the price elasticity of demand for the two subsets of varieties as follows:

$$\varepsilon^{\xi}_{i} = \left(\frac{p^{\xi}_{\max}}{p^{\xi}_{i}} - 1\right)^{-1} \qquad \forall i \in \Omega^{\xi}_{c}, \xi \in \{l, h\} \tag{12}$$

Assuming without loss of generality that $\eta^{l} = \eta^{h}$, it can be shown from eq. 2, 3 and 6 that, if $\gamma^{h} > \gamma^{l}$, then $p^{h}_{\max} > p^{l}_{\max}$, which in turn implies $\varepsilon^{h}_{i} < \varepsilon^{l}_{i}$, for any given price p_i.[7] Hence we have that consumers are willing to pay a higher maximum price for varieties in the high-differentiation sector, as compared to the low-differentiation one. Consistently, the price elasticity of demand is lower for the highly differentiated varieties in Ω^{h} than for their counterparts in Ω^{l}, where substitutability is higher. This result is intuitive and in line with previous empirical evidence, e.g. Goldberg (1995) for the car industry. The same argument can be extended to the differentiated products of virtually all other industries in which it can be assumed that different segments of the market do vary in demand structure and density of products, thus providing a rationale for a partition of the differentiated goods' set (Ω) such as ours.

2.2 Model setup: firms

We maintain the same assumptions as in Melitz-Ottaviano (2008). In particular, labor is the only factor of production and is inelastically supplied in a competitive market. The production of the homogeneous good requires one unit of labor, under constant returns to scale. Since this numeraire good is sold in a competitive market, a unit wage is implied.

Entry in the differentiated sector involves a sunk cost, which is related to product development and start-up investments. An entrepreneur decides *ex-ante* whether to enter in the low-differentiation market segment (Ω^{l}), paying a sunk cost f^{l}_{E}, or in the high-differentiation one (Ω^{h}), at cost f^{h}_{E}.[8] We assume that the choice of the market segment is exclusive, that is, once the sunk cost for one segment is paid, the firm cannot compete also in the other segment. One way of interpreting this assumption is that, even though they are active in the same industry, firms operating in different market segments are rather distinct in terms of organization of the value chain (e.g., they have different suppliers and distribution channels). Alternatively, to the extent that product design and branding play a role within an industry, firms may be thought of as making initial investments that are segment-specific. These arguments imply that a diversification of activities from one segment to the other would require additional sunk investments in order to produce and market the alternative variety.

Having paid the sunk cost, each firm draws, independently on the chosen market segment, an inverse productivity parameter c from the same industry-wide common distribution $G(c)$, with support $[0, c_M]$. Here c represents the firm-level marginal cost (in terms of units of labour) for the production of the differentiated good. There are no fixed costs of production, so the

[7] The formal proof is provided in Appendix 2.

[8] We do not assume *a priori* any ranking in the two sunk costs.

technology is characterized by constant returns to scale. Hence, those firms that can cover the marginal cost start producing, while the others exit.

Competition in each of the differentiated market segments is of a monopolistic nature, with each firm in Ω^ξ facing a residual demand function as expressed in inverse form in eq. 7 and 8. As the choice of the high vs. low differentiation segment in which to produce is exclusive (a firm cannot produce in both), we can optimize the firm decision within each market segment. In particular, optimum price $p(c)$ and output $q(c)$ must satisfy the following condition:

$$q_i^\xi(c) = \frac{L}{\gamma^\xi}[p_i(c) - c] \qquad \forall i \in \Omega_c^\xi,\ \xi \in \{l, h\} \tag{13}$$

If the profit maximizing price is above the relevant $p_{\max}^\xi$ the firm exits. Thus the marginal firm (indifferent between staying and exiting) in each market segment is characterized by a cut-off cost level c_D^ξ such that its price is driven down to the marginal cost ($p(c_D^\xi) = c_D^\xi = p_{\max}^\xi$), and the demand goes to zero. We assume that both cut-offs c_D^l and c_D^h are lower than the (common) upper bound of costs c_M, which implies that those firms with a cost draw between the segment-specific cut-off level and c_M do exit, while the others stay in the market and earn positive profits.

In the previous section we have shown that, if $\gamma^h > \gamma^l$, then $p_{\max}^h > p_{\max}^l$. This implies that $c_D^h > c_D^l$, i.e. the cost cut-off for survival is higher in Ω^h than in Ω^l. As a result, some less productive firms (with costs ranging between c_D^l and c_D^h) can survive in the high-differentiation market segment, while they would exit in the low-differentiation one.

2.3 Equilibrium in the closed economy

From the profit maximizing price, and using the expressions for the segment-specific cut-offs derived in eq. 9 and 10, we can solve for the optimal price $p^\xi(c)$:

$$p^\xi(c) = \frac{1}{2}\left(c_D^\xi + c\right) \qquad \text{for } \xi = l, h \tag{14}$$

and from here for the optimal produced quantity $q^\xi(c)$ and markup $\mu^\xi(c)$ in each market segment:

$$q^\xi(c) = \frac{L}{2\gamma^\xi}\left(c_D^\xi - c\right) \tag{15}$$

$$\mu^\xi(c) = p^\xi(c) - c = \frac{1}{2}\left(c_D^\xi - c\right) \tag{16}$$

Analogously, it is then possible to obtain firm-specific revenues and profits:

$$r^\xi(c) = \frac{L}{4\gamma^\xi}\left[(c_D^\xi)^2 - c^2\right] \tag{17}$$

$$\pi^\xi(c) = \frac{L}{4\gamma^\xi}(c_D^\xi - c)^2 \tag{18}$$

The model can then be solved for the closed-economy free entry equilibrium. Since the expected profits, net of entry costs, are zero in both Ω^l and Ω^h, an entrepreneur is *ex-ante* indifferent between entering one or the other segment. In Appendix 3 we also derive the average firm performance measures for a Pareto parameterization of the cost distribution $G(c)$, as well as closed form solutions for the segment-specific cut-offs (c_D^h and c_D^l).

From eq. 14 and 16, since $c_D^h > c_D^l$ we have that, for any given productivity level (i.e. c draw) a firm in the high-differentiation segment (Ω^h) charges a higher price and obtains a higher markup than an equally productive firm operating in the low-differentiation segment (Ω^l). The latter provides an explanation for the survival of relatively less productive firms in the high-differentiation market segment as compared to the low-differentiation one, and is consistent with available empirical evidence (e.g. Goldberg, 1995). The pattern of firm sizes (eq. 15 and 17) for equally productive firms operating in different market segments is instead more complex. As it entails the main testable prediction of the model it will be discussed separately in the next section.

3 Predictions

The crucial implication of our model is that the productivity of a firm does not unambiguously determine its size. Indeed, given a cost draw c, the level of firm output and revenues will depend on the market segment in which the firm is competing. Since we only observe revenues and not physical output in our firm-level data, in order to link theory and empirics in the closest possible way we formulate our proposition in terms of revenues.[9] In particular, from equation 17 it is possible to prove the following:

Proposition 1 *The ratio of firm revenues in Ω^h over Ω^l is > 1 for high levels of c (low productivity), and decreases for decreasing levels of the cost draw (increasing productivity), becoming < 1 after a threshold level c^T.*

Proof. First, by equating the optimal revenues $r^l(c)$ and $r^h(c)$ from eq. 17 we can derive the threshold cost level $c^T = \sqrt{\frac{\gamma^h (c_D^l)^2 - \gamma^l (c_D^h)^2}{\gamma^h - \gamma^l}}$, with $c^T > 0$ as long as $\gamma^h (c_D^l)^2 > \gamma^l (c_D^h)^2$. If the latter holds, it is straightforward to prove that $c^T < c_D^l < c_D^h$ and hence that a level of the cost draw exists, at which a firm operates in either the high or the low differentiation market segment (as the threshold is smaller than both cut-offs) with the same optimal size. In order to study the variation in the optimal size around the threshold, for any cost level below c^T, say $c^T - \epsilon$, we would have from eq. 17 that the optimal revenues in the two market segments are equal up to a term $\frac{L}{4\gamma^\xi}(2c^T\epsilon - \epsilon^2)$. From here, since $(2c^T\epsilon - \epsilon^2) > 0$ for any $\epsilon < 2c^T$ (i.e. for all levels of c between 0 and c^T), $\gamma^h > \gamma^l$ implies that $r^l(c) > r^h(c)$. Symmetrically, for any cost level $c^T + \epsilon$ the optimal revenues would be equal up to a term $\frac{L}{4\gamma^\xi}(-2c^T\epsilon - \epsilon^2)$. Since $(-2c^T\epsilon - \epsilon^2) < 0$, $\gamma^h > \gamma^l$ implies that $r^l(c) < r^h(c)$. ■

[9] We have proven and tested the same proposition also in terms of physical output (as proxied by deflated revenues), starting from the expression in (15). Results are available upon request.

The above proposition thus explains the coexistence of equally productive firms displaying different sizes within the same industry. Importantly, this result does not require the introduction of quality differences across varieties, but relies only on the existence of an asymmetric degree of horizontal product differentiation across market segments. The result is described graphically in Figure 1: for relatively low levels of productivity, larger firms operate in the market segment characterized by a higher degree of product differentiation; the opposite holds true for relatively high levels of productivity, where larger firms operate in the low-differentiation market segment. Notice that the difference between firms' sizes in the two segments (in absolute value) is directly proportional to the distance of each firm from the threshold $1/c^T$, and to the difference in the degrees of product differentiation $(\gamma^h - \gamma^l)$.

Figure 1: The productivity - size relation across different market segments

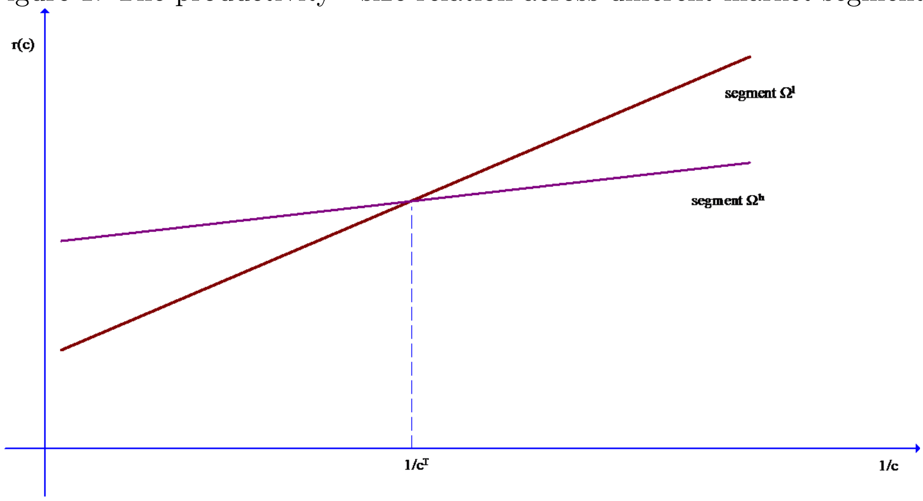

The result of Proposition 1 has an intuitive explanation: if two firms are very productive (low c), then the firm in the low-differentiation market segment will realize greater revenues with respect to the firm operating in the high-differentiation segment. In fact, the former can leverage upon the favorable cost draw to a larger extent, thanks to the high substitutability across varieties. The reverse will be true if the two firms have a low productivity (high c), for exactly the same reason: for a relatively inefficient firm it will be easier to attain a relatively higher level of revenues if varieties are less substitutable for each other.

In our analysis we have explicitly ruled out a vertical-differentiation dimension, that is, all varieties in our model are assumed to share the same level of quality α. The question is then whether the same result could be obtained by focusing instead on quality differences across segments. One could assume, for instance, that varieties in Ω^h are characterized by a superior quality with respect to varieties in Ω^l. Would this lead to the same pattern as described in Figure 1? The answer is no. Indeed, for that result to be obtained in a vertical differentiation framework, one should build up a model such that a superior level of quality (in Ω^h) results in a size premium up to a certain level of productivity, but the premium is then reverted to a a size penalty for the most productive firms, which is quite implausible. Indeed, in a quality-augmented extension of the Melitz-Ottaviano (2008) model, Kneller and Yu (2008) show that firms producing higher quality products always earn greater revenues for any given productivity level, in stark contrast with our findings.

In light of the above considerations, the non-linear result in the relation between size and productivity obtained in Proposition 1 constitutes the main distinctive contribution of our approach with respect to the literature focusing on the implications of quality differences across producers.

A second testable prediction can be derived from an open-economy version of our model, concerning the relation between productivity and exporting status across market segments. In particular:

Proposition 2 *Self-selection into exporting in* Ω^h *requires a relatively smaller productivity premium than in* Ω^l*.*

That is, relatively less productive firms are able to export in the high-differentiation market segment, while they would not be exporters in the low-differentiation segment. This is due to the fact that, for any given cost c, firms in Ω^h charge a higher markup, and thus are able to cover the transport costs and break-even on the export market already at a lower productivity level. A discussion of the open-economy version of the model, as well as a formal proof of Proposition 2, are provided in Appendix 4.

The latter finding extends the open-economy results of Melitz-Ottaviano (2008) to a demand system with asymmetric differentiation, and conveys a novel message: while it holds true that the relatively more productive firms within each market segment do export, the minimum productivity level (the productivity premium) which is required for becoming an exporter is inversely proportional to the segment-specific level of product differentiation. This result is another implication of the lower substitutability across varieties experienced in the high-differentiation segment, resulting in a lower price elasticity of demand for firms in Ω^h than in Ω^l.

Contrary to Proposition 1, Proposition 2 can however be obtained also by modeling vertical differentiation (i.e. quality differences) across varieties. For instance, assuming that varieties in Ω^h are characterized by a superior average quality, and assuming the existence of quality constraints on the export market (i.e. a minimum quality threshold for exporting, as in Hallak and Sivadasan, 2013), it would immediately follow that firms in Ω^h are more likely to export than firms in Ω^l, given the same level of productivity. Nevertheless, we believe it is important to show how differences in horizontal differentiation may also break the linear relation between productivity and exporting status, without having to rely necessarily on quality arguments. In the next sections, we will provide empirical evidence in support of both our predictions.

4 The French wine industry

The empirical test of our model is conducted by analyzing firm-level performance measures in the French wine-making industry (NACE-Rev.1.1 code 15.93). There are two main reasons for such a choice. The most important reason is the divisibility of the industry in multiple segments, as defined by a specific system of controlled denomination of origin: the ‘Appellation d’Origine Contrôlée’ (AOC). The AOC classification, introduced in 1935, is geographically based and follows the historical distribution of wine production areas in France, which goes back centuries. Hence, it allows for an exogenous partition of the industry which is well known to consumers, and can be directly associated to the level of horizontal differentiation across wines. The second reason for choosing the French wine industry is the availability of balance sheet data on the firms operating within each market segment over time, including information on export activities (i.e. exports as a share of turnover).

In this section, we first present the French wine industry and discuss its partition into market segments, explaining what are the differences between AOC and non-AOC wines, and why such differences relate to horizontal rather than vertical differentiation. We discuss as well how the degree of horizontal differentiation varies also across the various AOC regions. We then move to the description of the firm-level dataset, and present our estimation of TFP.

4.1 The AOC system and market segments

The French wine-making industry has been historically characterized by a strong geographically based partition. In the Middle Ages, wine consumers in Paris or London were already aware of the differences between wines produced in different areas of France (Wilson, 1998). Indeed, besides the natural influence of different soil and climate conditions across regions, wine-making activities followed regional-specific practices, which jointly determined some recognizable characteristics of the final product. Based upon this tradition, a law decree of 1935 has codified an official national system of controlled denomination of origin: the ‘Appellation d’Origine Contrôlée’ (AOC). Within this system, a given AOC-denominated wine can be labelled and sold as such ‘if and only if’ the production takes place within a specific geographic area. Ten "macro" AOC regional areas are codified in the system, mapping the historical subdivision of the French wine industry which is well known to consumers: Alsace, Bordeaux, Bourgogne, Champagne, Jura-Savoie, Languedoc-Roussillon, Loire, Provence, Rhône and South-West. The French AOC system is the oldest of the European label of origin systems, and involves a great deal of administration and control for its maintenance. In particular, monitoring activities are related to the specific methodological practices that are codified for each AOC product (through official documents called "cahiers de charges" in French), and must be closely followed by the producers. The "cahiers de charges" regulate many aspects of wine production, such as grape varieties that can be used and in what percentage, minimum level of alcohol content, timing and procedures for planting, harvesting and pruning of vines etc. The AOC system is widely regarded as being a key factor for the global success of the French wine industry, and it is perceived as a reference model worldwide (Barham, 2003).

At the "micro" level, the AOC classification includes around 380 official wine denominations (also called appellations). For each of them, the French Institute of Origin and Quality (INAO), which is the public agency responsible for the management of the AOC system, publishes the list of municipalities that are included in the specific production area. Each municipality (a French 'commune') is identified by a unique INSEE code, which allows us to identify the AOC wine producers based on their geographical location (as the INSEE code is also available for each producer in the firm-level dataset).[10] After downloading and merging all these lists, we have noticed that the same municipality can appear in the list of several denominations. The most evident case is that of Vosne-Romanée, in the Burgundy region, a small town where wines can be produced with 15 different AOC denominations, from the standard 'Bourgogne' to the exclusive 'Romanée-Conti'. We have then aggregated the different denominations in the ten non-overlapping "macro" AOC areas mentioned above, in such a way that each AOC municipality is assigned to a unique general AOC appellation. Thus, our industry partition maps exactly the classical subdivision of the French wine industry as recognized by consumers.[11] Table 1 reports the number of "sub-appellations" aggregated within each of the ten general AOC appellations. The highest number of sub-appellations is witnessed by Bourgogne, 100, while there is only one appellation for Champagne. Such a peculiarity of Champagne has been discussed and exploited in a recent paper on quality sorting and trade by Crozet et al. (2012), as we discuss later.

Table 1: AOC areas and their sub-appellations

AOC	number of sub-appellations
Alsace	63
Bordeaux	47
Bourgogne	100
Champagne	1
Jura-Savoie	9
Languedoc-Roussillon	19
Loire	69
Provence	9
Rhône	45
South-West	17
Total	379

We proceed by explaining the differences between AOC-denominated wines and non-AOC wines. According to the French regulatory system, the AOC labels are assigned to products that derive their authenticity and typicity from their geographical origin. In the wording of INAO, an AOC label is thus the expression of an "intimate link" between a product and its "terroir", i.e. a distinct combination of particular climatic, agronomic and geological conditions. These environmental factors, along with specific and codified methodological practices, make AOC wines unique and not replicable outside of their production area.[12] A distinct taste and wine

[10] INSEE codes are used by the French National Institute for Statistics and Economic Studies for identifying geographical entities. These codes allow for deeper territorial disaggregation than zipcodes. In fact, several small municipalities often share the same zipcode. Instead, INSEE codes are always specific to a single municipality.

[11] A distinct 'Cognac' area has also been identified. However, we have decided not to consider it, given the particular nature of this product, which is a spirit rather than a wine. Accordingly, producers located in the Cognac area are dropped from the analysis. See the next section for more details.

[12] In a widely cited book, Wilson (1998) analyses the central concept of terroir in wine production, by which

identity are clearly horizontal attributes, for which consumers who love variety are willing to pay (Cross et al., 2011; Lecocq and Visser, 2006; Quandt, 2007).[13]

Consistent with the above discussion, and important for our purposes, several studies on the economic implications of AOC denominations have provided evidence that "the complex and costly French AOC system seems unable to produce more than just horizontal differentiation (typicity). As a matter of fact, it cannot guarantee a high level of quality (vertical differentiation)" [Gergaud and Ginsburgh, 2008, p. 150]. Indeed, many low-quality wines are also produced in municipalities that are eligible for an AOC label, and respecting the methodological practices specified in the "cahiers de charges". Gergaud and Ginsburgh (2008) even report the extreme case of the famous wine-maker Didier Daguenau, who has produced an admittedly very bad AOC-denominated wine, making the point that an AOC label is not per se a guarantee for high-quality, a remark which is also made by Crozet et al. (2012). Overall, what is unique about French AOC wines is thus their distinctive character (horizontal differentiation), rather than their quality (vertical differentiation).

Wines produced within a single AOC macro region, say Bordeaux, will share some fundamental elements, conferred especially by the types of grapes employed in production (which will distinguish them from other AOC wines), but will also witness a substantial degree of horizontal heterogeneity across different products of similar quality (Gergaud and Ginsburgh, 2008; Cross et al., 2011; Parker, 1985; Wilson, 1998), as reflected by the large number of sub-appellations that have been codified through history (Barham, 2003).[14] Indeed, the number of sub-appellations within each AOC macro area is another element that we are going to use in the empirical analysis, to investigate the variation in horizontal differentiation across regions.

The role of the AOC system has also been extensively studied from a consumers' behavior perspective. In particular, as discussed by Veale and Quester (2008), wine is a type of product for which extrinsic cues tend to have a stronger influence on consumers' choices than intrinsic ones, which are difficult to be assessed (Dodds, 1991; Kardes et al., 2004; Monroe, 1976; Verdù Jover et al., 2004).[15] A fundamental extrinsic cue of wine is the geographical origin, as indicated by the AOC system on a bottle's label. Not surprisingly then, AOC labels have been found to be a more important driver of wine prices as compared to sensory variables related to quality (Lecocq and Visser, 2006; Combris et al., 1997; Combris et al., 2000).

different terrains in the AOC areas determine distinctive characteristics of the final products. In turn, Barham (2003) discusses how the AOC system has been a crucial success factor for the French wine industry, by translating the role of natural endowments and human know-how in economic value recognized by consumers.

[13] There is plenty of evidence that the wine drinking experience is very subjective (Lecocq and Visser, 2006), and that even experts' ratings of the same wines do not strongly correlate with each other (Quandt, 2007). There are even concerns that experts' descriptions may not tell much to professional tasters, let alone average drinkers (Goldstein et al. 2008; Weil, 2001, 2005, 2007).

[14] Talking about Bordeaux, the famous wine critic Robert Parker (1985) notices how "subtle differences in soil may lead to very different styles."

[15] Experiments have in fact shown how average consumers do feel somewhat intimidated rather than confident of correctly evaluating different wine products, and so prefer to rely on extrinsic cues (Verdù Jover et al., 2004).

In light of the above discussion, in the empirical strategy we will start by considering all the AOC producers as competing in a single high-differentiation segment (Ω^h), as opposed to the low-differentiation segment of non-AOC producers (Ω^l), with each producer assigned to one of the two segments based on its geographical location (i.e. in municipalities within or outside of any AOC area). Going back to our model, we are thus assuming that consumers attach a greater utility penalty to uneven consumption of distinctive AOC wines (a high γ in the model's notation), as compared to uneven consumption of different varieties of non-AOC wine, a market segment in which one bottle is likely to be relatively less different from the other (i.e. a low γ).[16] In a refinement of the analysis, we will also treat each of the ten "macro" AOC areas as separate market segments, each characterized by a measurable degree of horizontal differentiation that is given by the corresponding number of sub-appellations, as reported in Table 1.

4.2 The French wine producers

Firm-level data are obtained from AMADEUS. This is a commercial database produced by Bureau Van Dijk, containing annual balance sheet data for over 14 million companies across all European countries, spanning the period 1999-2008 in the release we have used. In general, for each firm, information is available on turnover, value added, capital, number of employees, materials, labor costs and other financial indicators. In the case of France, the official data source embedded in AMADEUS is constituted by the balance sheet information that most French firms (and foreign multinationals located in France) are obliged to deposit each year to the 'Tribunaux de Commerce'.[17]

Firms in the dataset are classified according to the NACE (Rev.1.1) classification of industries. In particular, we focus on the 4-digit industry 15.93, named 'Manufacture of wines'. For this industry, the release of AMADEUS we have used contains data on 1,124 French firms.[18] These firms report wine production as being their primary activity, so they cannot be simple wholesalers of wine, in line with Crozet et al. (2012). For all firms, we do have information on the municipality where they are located. This allows us to identify all the producers being active in one of the AOC areas, as well as firms located outside of these areas (non-AOC). After dropping a handful of clearly problematic observations (e.g. obvious mistakes in the data input process), as well as those firms located in the 'Cognac' AOC area (as the latter is a spirit, not a wine), we are left with 1,052 firms. In terms of representativeness, these firms account for around 88% of the official total turnover reported by Eurostat for the French wine industry

[16] And yet, the consumption patterns across segments cannot be assumed to be uncorrelated, as there would be in general a certain positive degree of substitution between AOC and non-AOC wines (a positive β in our model).

[17] Firms exempted from presenting their balance sheets to the *Tribunaux* are some of the partnerships ('Sociétés de personnes'), and some of the cooperative companies ('Sociétés coopérative et unions'), under specific conditions. No censoring in terms of size is used in the data, i.e. also small firms are represented.

[18] In the vast majority of cases, data refer to single-establishment firms. Only in 17 cases our observational units are part of a broader group. In these cases we use unconsolidated data.

(NACE 15.93), on average over the time span.

Table 2 reports the distribution of firms in our sample across AOC vs. non-AOC areas. The largest group of AOC producers is located in the Champagne region, with 387 firms, followed by Languedoc-Roussillon, with 166 companies. 117 firms in our dataset are instead wine makers located outside of any AOC area. The proportion of AOC vs. non-AOC firms in our sample is in line with official statistics (INAO), according to which AOC areas account for about 85% of total turnover in the French wine industry. The large majority of non-AOC firms (around 80%) are still located in administrative counties ('*Départements*', defined at the NUTS-3 level of disaggregation[19]) which also contain some municipalities belonging to an AOC area. We can thus rule out a geographical bias in our sample composition.

Table 2: Firms' distribution and share of exporters across AOC areas

AOC	Number of Firms		Share of Exporters
non - AOC	117		0,13
Alsace	23		0,43
Bordeaux	46		0,26
Bourgogne	65		0,40
Champagne	387		0,25
Jura-Savoie	18		0,10
Languedoc-Roussillon	166		0,19
Loire	54		0,27
Provence	55		0,19
Rhône	92		0,21
South-West	29		0,27
Total	1.052	Avg.	0,23

As previously anticipated, our firm-level data include information on exports. According to the French law in fact, each company reporting to the 'Tribunaux de Commerce' is mandatorily required to provide the figure on turnover accounted for by exports, i.e. zero or positive. The latter feature results in a broad coverage of firms' export activities across all categories of size, a characteristic of the French data already exploited in the empirical trade literature (e.g. Konings and Vandenbussche, 2013). In particular, the reported export figures refer to the direct exports of each firm, and do not include those exports taking place through domestic intermediaries. Hence, our analysis focuses on the self-selection of producers into direct exporting activities, in line with Crozet et al. (2012). Such a selection has been shown to be stronger than for indirect exporters (Ahn et al., 2011), consistent with higher fixed costs for direct exports than for indirect ones.

The second column of Table 2 presents descriptive figures on the extensive margin of trade, i.e. the share of firms reporting a positive value of exports, for each group of producers, on

[19] The NUTS classification partitions the territory of the EU Member States into administrative regions at different levels of aggregation. The NUTS-3 level corresponds to the level of provinces or counties, with population generally comprised between 150,000 and 800,000 inhabitants. With exclusion of the overseas regions, there are 96 NUTS-3 Départements in France.

average across years. Overall, 23% of firms are direct exporters. With the exception of Jura-Savoie, the share of direct exporters is higher for AOC firms than for non-AOC producers. And yet, there seems to be a substantial degree of heterogeneity across different AOC areas. For instance, the share of direct exporters for Bourgogne is much higher than for Champagne: 40 vs. 25%. This is consistent with the fact that a large share of direct exports in Bourgogne is accounted for by small firms with less than 10 employees, as reported by Crozet et al. (2012): up to 80%, against 3% for Champagne.

Our panel is unbalanced, as not all firms are observed throughout the time-span. Moreover, missing observations for some of the variables of interest do not allow us to estimate firms' total factor productivity (TFP) in all cases. Focusing only on those firm-year observations with complete information, Table 3 reports some descriptive statistics for different categories of firms. It can be inferred that, on average, non-AOC firms are not systematically different from AOC firms.

Table 3: Firm-level data: descriptive statistics

Variable	Obs	Mean	Std. Dev.	Min	Max
a) Overall Sample					
employees	2894	27.7	83.6	1	1274
turnover (thousand euros)	2894	15090.7	52613.5	27	918375
materials (thousand euros)	2894	7947.6	23116.7	1	347686
tangible fixed assets (thousand euros)	2894	2843.4	8295.5	1	133466
b) non-AOC firms					
employees	252	38.1	95.1	1	436
turnover (thousand euros)	252	14580.1	49984.4	70	325672
materials (thousand euros)	252	7829.6	26430.7	2	167457
tangible fixed assets (thousand euros)	252	1721.3	2640.3	1	12851

4.3 TFP estimation in the wine industry

We first estimate total factor productivity through a simple OLS procedure, by regressing value added (output minus materials) over capital and labor inputs (employment). Physical output is proxied by deflated turnover. The employed deflator is specific to the 4-digit industry, and is published by the French National Statistical Institute. Materials' costs are deflated using a specific input deflator obtained from the EU-KLEMS database.[20] Capital is proxied by tangible fixed assets, deflated using the GDP deflator.

Considering the specificities of the wine industry, it may be useful to give an idea of what is behind the balance sheet figures for materials and capital, even though an official breakdown is not provided in AMADEUS. A close look at the published informative notes accompanying the balance sheets of some producers (e.g. Pommery) reveals that the tangible fixed assets are constituted by land, vines, machines, barrels, vehicles, and buildings for production and ageing.

[20] The EU KLEMS database is the outcome of a project financed by the European Commission for the analysis of productivity and growth. More details are available on the EU KLEMS website: http://www.euklems.net/index.html

Materials include standard items such as purchases of services, electricity and fuel, along with fertilizers, bottles, corks, and all that is needed for packaging.

Given the well-known simultaneity problems of OLS productivity estimates (see Van Beveren, 2012, for a review), we have also estimated TFP by applying the value added version of the Levinsohn-Petrin (2003) algorithm (Lev-Pet from now on). The results from both estimations can be compared in Table 4. In line with expectations, the labor coefficient is significantly reduced in the Lev-Pet estimation, from 0.69 to 0.51. We will then consider the Lev-Pet estimated figures as our first benchmark measures of TFP. Although we have not found any earlier TFP estimates for wine producers to perform a comparison with respect to our capital and labor coefficients, their sum being <1 is consistent with the fact that wine producers charge relatively high markups. In particular, in our data we observe an average markup of 0.33, a figure consistent with earlier findings by Cranfield (2002).[21] In line with what one would expect from the theoretical model (i.e. higher markup for given productivity in case of higher differentiation), the price-cost margin is somewhat higher for AOC firms than for non-AOC firms: 0.33 vs 0.30 on average, a statistically significant difference.

A potential problem with our Lev-Pet estimates is related to the fact that we do not observe physical output nor firm-level prices, and thus we have to rely on a common deflator for the revenues of all firms. This may bias upwards the productivity estimates for those firms charging higher markups, i.e. the AOC producers. If this is the case, note however how the latter bias would clearly work against us in finding supportive evidence for Proposition 1, which predicts that, at high TFP levels, AOC firms should display *smaller* nominal revenues than equally productive non-AOC firms. Moreover, as discussed by Van Beveren (2012), to the extent that input prices and output prices are positively correlated (which is arguably the case for wine, according to Crozet et al. 2012), the biases induced by omitted prices are likely to net-out each other, as both inputs and output would be over-estimated (or under-estimated) by using a single deflator at the industry level, depending on whether a firm is charging higher (lower) prices than the average and facing higher (lower) input costs than the average.

In a robustness check discussed later, we will also employ TFP estimates obtained through an augmented version of the Lev-Pet productivity estimation, which incorporates year dummies in the first stage. This is meant to account for the impact of yearly-specific transitory shocks which could be determined, for instance, by weather conditions.

[21] Markups are computed as price-cost margin according to the following formula: (turnover-labor costs-material costs)/turnover.

Table 4: TFP estimation - OLS vs. Levinsohn-Petrin

Dependent Variable: ln(value added)

	OLS	Lev-Pet
	(1)	(2)
ln(labor)	0.688***	0.514***
	(0.012)	(0.023)
ln(capital)	0.332***	0.126**
	(0.009)	(0.055)

A more general concern one might have is whether it is appropriate to estimate a single production function for the pooled sample of AOC and non-AOC producers. Empirical researchers have normally estimated TFP at the 3-digit industry level, and often a 2-digit approach has been adopted, while in this case we only look at one single 4-digit industry. The production technology can thus be expected to be the same for all firms in our sample, as wine production remains essentially the same in different geographical locations, even though the types of grape that are used do change, along with some methodological practices foreseen by the different AOC regulations (the "Cahiers de charges").

As an alternative to Levinsohn and Petrin (2003), we also estimate TFP through Index Numbers, as in Aw et al. (2001). Given the peculiarity of the production function characterizing the wine industry (where an important part of the capital stock is constituted by immobile factors such as land), the idea is to compare our results against a fully non-parametric methodology, in which the efficiency of each firm in each year is computed relative to a hypothetical firm operating in the base year. The hypothetical firm has input revenue shares and log input levels equal to the arithmetic means of the revenue shares and log inputs observed across all observations in the base year. More in detail, the TFP index for a firm f in year t is defined as:

$$\ln TFP_{ft} = (\ln Y_{ft} - \overline{\ln Y_t}) + \sum_{s=2}^{t} (\overline{\ln Y_s} - \overline{\ln Y_{s-1}}) -$$
$$- \left[\sum_{i=1}^{n} \frac{1}{2}(s_{ift} + \overline{s_{it}})(\ln X_{ift} - \overline{\ln X_{it}}) + \sum_{s=2}^{t} \sum_{i=1}^{n} \frac{1}{2}(\overline{s_{is}} + \overline{s_{is-1}})(\overline{\ln X_{is}} - \overline{\ln X_{is-1}}) \right] \quad (19)$$

where Y_{ft} stands for the output and X_{ift} is the level of each employed input $i = 1, .., n$. The term s_{ift} is the share of firm's expenditure for input i out of total revenues, while $\overline{\ln Y_t}$, $\overline{\ln X_{it}}$ and $\overline{s_{it}}$ stand for the corresponding arithmetic means over all firms in year t.

In our Index Numbers estimation, the base year is 1999, while output, capital, labor and material inputs are proxied as above. The revenue shares of materials and labor are computed by taking the ratio of materials and labor costs over turnover, in nominal terms. The capital share is instead computed as a residual, by relying on the product-exhaustion theorem, which entails the assumption of constant returns to scale (in line with our theoretical model).

Table 5 reports some descriptive statistics for the different measures of productivity. As expected, these measures are all positively and significantly correlated, although with some differences. In particular, while both Levinsohn-Petrin and Index Numbers estimates are highly correlated with OLS ones (about 0.75), the correlation between them is somewhat lower, about 0.44, thus providing room for a solid robustness check.

Table 5: TFP - descriptive statistics

Variable: ln(TFP)	Obs.	Mean	Std. Dev.	Min	Max
Lev-Pet	2894	5.154	0.890	1.658	8.090
Index Numbers	2893	0.178	0.414	-1.935	2.555
OLS	2894	3.460	0.680	0.127	7.286

5 Empirical results

5.1 TFP, size and horizontal differentiation

Our model predicts size to be an increasing linear function of TFP within each market segment. However, such a linear function should have both a different intercept and a different slope for market segments characterized by heterogeneous degrees of product differentiation (eq. 17). More specifically, Proposition 1 states that, for low levels of productivity, firms in the high-differentiation segment earn higher revenues than their counterparts in the low-differentiation segment, while this relation is inverted as productivity increases, since revenues grow faster with productivity in the low-differentiation segment. We test this prediction by comparing the performance of AOC producers (high-differentiation) versus non-AOC producers (low-differentiation). Empirically, the linear relation between revenues and TFP should thus present a higher intercept but a smaller slope for AOC firms relatively to non-AOC firms. In turn, this should determine an inversion in the size-ratios for equally productive firms belonging to the two different segments as productivity grows, in line with Figure 1. In what follows we test for these predictions.

Table 6: Proposition 1 - econometric test

Dependent variable: ln(size)

TFP estimated through:	OLS	Lev-Pet	Index Numbers	Augmented Lev-Pet		Lev-Pet	Index Numbers
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
ln(TFP)	1.047***	1.656***	1.088***	1.728***	1.723***	1.696***	1.095***
	[0.000]	[0.000]	[0.000]	[0.000]	[0.003]	[0.005]	[0.007]
AOC Overall Dummy	1.161***	0.666***	0.427***	0.727***	0.705**	0.658**	0.420***
	[0.000]	[0.000]	[0.000]	[0.000]	[0.015]	[0.014]	[0.002]
*ln(TFP)*AOC Overall Dummy*	-0.293***	-0.147***	-0.840***	-0.181***	-0.176**	-0.153**	-0.853***
	[0.000]	[0.000]	[0.000]	[0.000]	[0.004]	[0.003]	[0.003]
Constant	4.309***	-0.390***	7.623***	-0.064***	-0,033	-0.016**	7.888***
	[0.000]	[0.000]	[0.000]	[0.000]	[0.006]	[0.001]	[0.034]
Year Dummies	no	no	no	no	yes	yes	yes
N. of obs.	2.894	2.894	2.893	2.894	2.894	2.894	2.893
R-sq	0,11	0,67	0,02	0,67	0,67	0,68	0,02

Standard errors are clustered within the AOC and the control group. ***, **, * indicate significance at the 1, 5 and 10% level, respectively.

Table 6 reports the outcome of a first econometric test of Proposition 1, in which we perform a pooled regression of firm-level nominal revenues on TFP (in logs), the AOC "overall" dummy, and the interaction of the two, where the AOC "overall" dummy identifies all firms located in either one of the ten AOC areas. Their size/TFP relation is then compared against the control group of non-AOC producers, employing the three different measures of productivity discussed above: OLS, Lev-Pet and Index Numbers. The baseline results, reported in the first three columns, are in line with Proposition 1. In particular, for low levels of productivity firms tend to be larger in the high-differentiation market segment, i.e. a higher intercept in the linear relation, as measured by the positive and significant AOC overall dummy; however, as TFP increases, size grows slower in the high-differentiation segment than in the low-differentiation one, i.e. a smaller slope, as measured by the negative and significant interaction term between TFP and the AOC overall dummy. These baseline results are robust to using the three different measures of productivity. As the OLS estimates of TFP are known to be biased, for convenience of exposition we will focus on Lev-Pet and Index Numbers estimates in the rest of the analysis.[22] To clarify the interpretation of these results, a simple back-of-the-envelope calculation based on column 2, for example, reveals that the inversion of size ratios across segments should be observed around a level of log TFP equal to 4.5 (i.e. nominal revenues equal to around 1.2 mn. euros).

A potential concern in estimating an econometric relation between size (measured by revenues) and TFP is related to a possible mechanical endogeneity of the latter term, as TFP is essentially the share of value added (i.e., revenues - materials) left unexplained by capital and labor in the production function. By regressing revenues over TFP one could thus fear that a spurious positive correlation might arise in a mechanical way between the two terms. In practice this is not the case, as TFP is not a size-related but a *relative* variable.[23]

The remaining columns of Table 6 report the results of four robustness checks on the test of Proposition 1, aimed at controlling for the potential role of transitory shocks, which could spuriously induce a positive correlation between TFP estimates and revenues. First, we have employed TFP estimates obtained through an augmented version of the Lev-Pet procedure, incorporating year dummies in the first stage of the semi-parametric estimation (column 4). This does not have any notable impact on our results; if anything, the positive correlation between productivity and size is even stronger. Second, we have included year dummies in the regression, using at the same time the Lev-Pet estimates obtained by including year dummies in the first-stage (column 5). Also in this case the results are virtually unaffected. Third, we have included year dummies in the regressions while using either the original Lev-Pet estimates (not

[22] We have also run all the other regressions employing the OLS productivity estimates. Results are always consistent with those reported in the paper, and are available upon request.

[23] In particular, in our econometric model we regress the log of revenues on the log of TFP. Analytically, starting from a production function in revenues (Y), materials (M), capital (K) and labor (L), one can write TFP in log form as: $\ln(TFP) = \ln(Y) - \alpha \ln(L) - \beta \ln(K) - \gamma \ln(M)$. The test of Proposition 1 involves the following type of estimation: $\ln(Y) = \delta_0 + \delta_1 \ln(TFP) + \varepsilon$, which could be easily rewritten by substituting the TFP term as: $\ln(Y) = \delta_0/(1-\delta_1) - \alpha\delta_1/(1-\delta_1)\ln(L) - \beta\delta_1/(1-\delta_1)\ln(K) - \gamma\delta_1/(1-\delta_1)\ln(M) + \varepsilon$. The latter shows that size, as measured by revenues, only impacts one side of the equation, not both, thus ruling out a mechanical relation between size and TFP in our regressions.

augmented with year dummies in the first stage), or the Index Numbers measures of productivity (columns 6 and 7, respectively), always finding the same results.

Table 7: Proposition 1 - additional controls

Dependent variable: ln(size)

TFP estimated through:	Lev-Pet	Lev-Pet	Lev-Pet	Lev-Pet	Index Numbers	Index Numbers	Index Numbers	Index Numbers
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ln(TFP)	1.656***	1.645***	1.639***	1.614***	1.088***	1.930**	2.460**	2.370**
	[0.000]	[0.000]	[0.013]	[0.009]	[0.000]	[0.052]	[0.061]	[0.078]
AOC Overall Dummy	0.666***	0.629***	0.636**	0.593**	0.427***	0.411***	0.254*	0.283*
	[0.000]	[0.001]	[0.014]	[0.019]	[0.000]	[0.001]	[0.022]	[0.026]
*ln(TFP)*AOC Overall Dummy*	-0.147***	-0.131***	-0.132***	-0.121**	-0.840***	-1.067***	-0.525*	-0.519*
	[0.000]	[0.000]	[0.002]	[0.003]	[0.000]	[0.014]	[0.081]	[0.082]
ln(materials/turnover)		0.681**	0.678***	0.647**		0.851**	0.910***	0.822***
		[0.013]	[0.005]	[0.012]		[0.053]	[0.006]	[0.007]
ln(capital/employees)			0.017	-0.003			0.535*	0.466
			[0.036]	[0.041]			[0.076]	[0.087]
ln(firm age)				0.099*				0.247*
				[0.015]				[0.025]
Constant	-0.390***	0.201**	0.153	-0.006	7.623***	8.246***	5.919**	5.293**
	[0.000]	[0.011]	[0.092]	[0.070]	[0.000]	[0.039]	[0.343]	[0.293]
N. of obs.	2,894	2,894	2,894	2,881	2,893	2,893	2,893	2,880
R-sq	0.67	0.78	0.78	0.79	0.02	0.17	0.32	0.34

Standard errors are clustered within the AOC and the control group. ***, **, * indicate significance at the 1, 5 and 10% level, respectively.

In Table 7, we add sequentially to the baseline specifications of Table 6 three firm-level controls: material intensity (i.e. log of materials costs over turnover), capital intensity (log of capital over employees), and the logarithm of firm age. The inclusion of such controls, aimed at controlling for systematic differences across AOC vs. non-AOC firms, does not change our main results, both when using Lev-Pet and Index Numbers measures of productivity. Size appears to be positively correlated with material intensity and firm age, in line with traditional stylized facts in the firm dynamics literature (e.g. Evans, 1987), while a positive correlation with capital intensity is less robust. We shall stress that the three additional controls do not show statistically significant differences between AOC and non-AOC firms, as shown in Table 8. As already mentioned, there is instead evidence of a statistically significant difference in price cost margins (PCM), with AOC firms charging a somewhat higher PCM on average. In two unreported regressions, we have included the firm-level PCM as an additional control, both with Lev-Pet and Index Numbers, with no notable differences in our results with respect to those reported in columns 4 and 8 of Table 7.

Table 8: Firm-level controls - descriptives

Variable	Obs.	Mean	Std. Dev.	Min	Max
a) AOC firms					
material intensity	2642	0.52	0.23	0.00	0.97
capital intensity	2642	4.36	1.42	-1.39	8.91
firm age	2637	45.54	37.32	0.00	197.00
b) non-AOC firms					
material intensity	252	0.53	0.21	0.00	0.88
capital intensity	252	4.38	1.37	-1.10	8.13
firm age	252	50.59	37.62	1.00	108.00

A further concern one may have with our empirical analysis is that we are not controlling for possible quality differences across market segments. The latter do not play any role in obtaining the theoretical result, but may still be empirically relevant. To that extent, Crozet et al. (2012) have studied the implications of quality differences across Champagne producers only, relying on published quality ratings from wine critics. A similar approach is not viable in our case, as quality ratings (retrieved from publications such as Parker's "Wine Buyer's Guide") tend to cover only a limited number of solely AOC producers, so we would miss the information for many firms, especially in the crucial non-AOC control group. Besides the issue of data availability, Crozet et al. (2012) discuss in detail why Champagne is a quite unique case for the purposes of quality analysis. Specifically, within the single and homogeneous Champagne appellation, 95% of Champagne production is "non-vintage", i.e. firms can blend wine produced in different years in order to keep the characteristics of the final product stable over time. As a result, in most cases, the label of a bottle of Champagne does not report the vintage (i.e. year of production) on it. The opposite is true for all the other AOC wines, where quality may change a lot from one year to the other even for the same producer. This implies that, for a meaningful control, one would need to know which vintage accounts for which share of turnover in each year, something we cannot observe; let alone the fact that the same firm normally produces multiple varieties (e.g. a white Chardonnay and a red Pinot Noir in Burgundy), with potentially different associated qualities.

Quality is therefore not the focus of our paper, but of course we need to make sure that omitted quality is not driving our results. To that extent, if one believes that wines produced by AOC firms possess, on average, a higher quality than non-AOC ones, the latter would actually work against us in finding supporting evidence for Proposition 1. In fact, higher-quality firms would always tend to be larger for any given productivity level, as shown by Kneller and Yu (2008) in a quality-augmented version of Melitz-Ottaviano (2008). Hence, not controlling for quality would make it more difficult to find, as we do, evidence of an inversion in the size-productivity pattern for high levels of TFP, in which non-AOC firms grow larger than AOC ones after a certain productivity threshold. Kneller and Yu (2008) also show that, intuitively, higher-quality firms do charge higher markups for any given productivity level. And still, as discussed above, our results are unchanged when controlling for the firm-level price cost margin, another hint that omitted quality is unlikely to be driving our results.

Beyond the econometric test discussed so far, Proposition 1 can also be tested in a non-parametric way, by directly showing how the initial size-relation across segments is inverted after a productivity threshold within our sample. That is, AOC firms are larger than non-AOC firms at low productivity levels, while the opposite holds true at high productivity levels. Table 9 reports this straightforward test. In detail, we have allocated our firms to ten different classes based on their productivity level (cut-off points for each class are the deciles of the TFP distribution, computed separately for Lev-Pet and Index Numbers estimates). Then, we have computed the average firm size within each class, separately for non-AOC and AOC firms. As a measure of size we have employed nominal turnover, as in the previous regressions. Finally, we have calculated the ratio between non-AOC and AOC average-sizes for each of the ten classes. The findings are clear: AOC firms are on average larger than non-AOC firms for the lowest classes of TFP; this relation is then inverted as productivity grows, thus providing a direct confirmation of Proposition 1. This result holds both when considering Lev-Pet and Index Numbers estimates of TFP, with the only notable difference being that the inversion of size ratios for Index Numbers is observed at a higher decile of productivity.

Table 9: Proposition 1 - non-parametric test

	Average firm size: (turnover, 000s eur)				Average firm size: (turnover, 000s eur)		
Deciles of ln(TFP) Lev-Pet	non -AOC	Within AOCs	Ratio	Deciles of ln(TFP) Index Numbers	non -AOC	Within AOCs	Ratio
1	306.7	365.7	0.84	1	724.0	1360.1	0.53
2	716.5	811.4	0.88	2	3081.0	3655.2	0.84
3	1390.1	1453.3	0.96	3	1906.3	3256.1	0.59
4	1494.7	1769.4	0.84	4	1616.8	3593.6	0.45
5	3110.1	2399.7	1.30	5	1833.3	3812.3	0.48
6	3369.1	3223.4	1.05	6	2260.4	4024.3	0.56
7	4962.2	4366.0	1.14	7	2365.2	3690.1	0.64
8	7611.4	6740.6	1.13	8	3537.6	4304.0	0.82
9	14259.8	11337.8	1.26	9	6475.0	4309.6	1.50
10	150512.3	47978.1	3.14	10	7797.0	2291.4	3.40

5.2 Other robustness checks

It could well be the case that the AOC "overall" dummy employed in our empirical analysis is just capturing generic regional effects, instead of any specific role of the distinct market segments, as our model postulates. If that would be the case, then we would expect to find similar patterns in the size/TFP relation also for other comparable industries, when considering firms that are located within the same geographical areas.

In order to rule out this possibility, we have repeated the test of Proposition 1 for two alternative French 4-digit food industries: "Production of meat and poultrymeat products" (NACE Rev.1.1 code 15.13) and "Manufacture of bread; manufacture of fresh pastry goods and cakes" (NACE Rev.1.1 code 15.81). The reason for selecting these two particular industries as a counterfactual is twofold. First, these industries are still part of the food sector, but they produce goods which are less differentiated than wine, at least on a geographical base. Second, these industries display a significant coverage over the entire French territory in terms of number of firms, as reported by AMADEUS. As a result, when restricting ourselves to the same municipalities covered by the wine-producers database (i.e. where at least one wine producer is located), we could still rely upon a relatively high number of observations for our tests.

Table 10: Proposition 1 - econometric test for alternative industries

Dependent variable: ln(size)

Industry:	NACE 1593 (wine)		NACE 1513 (meat products)		NACE 1581 (bread products)	
TFP estimated through:	Lev-Pet	Index Numbers	Lev-Pet	Index Numbers	Lev-Pet	Index Numbers
	(1)	(2)	(3)	(4)	(5)	(6)
ln(TFP)	1.656***	1.088***	2.209***	1.607***	1.705***	1.045***
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
AOC Overall Dummy	0.666***	0.427***	1.453	-0.09	0.822	-0.231***
	[0.000]	[0.000]	[1.605]	[0.099]	[0.587]	[0.065]
*ln(TFP)*AOC Overall Dummy*	-0.147***	-0.840***	-0.81	-0.51	-0.535	-0.273*
	[0.000]	[0.000]	[0.857]	[0.837]	[0.301]	[0.149]
Constant	-0.390***	7.623***	2.148***	6.310***	2.613***	5.941***
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
N. of obs.	2,894	2,893	1,005	1,002	5,709	5,698
R-sq	0.67	0.02	0.09	0.06	0.13	0.08

Standard errors are clustered within the AOC and the control group. ***, **, * indicate significance at the 1, 5 and 10% level, respectively.

Table 10 replicates the baseline econometric test of Proposition 1 presented in Table 6 for the meat and bread industries, focusing on Lev-Pet and Index Numbers measures of productivity, as in the previous robustness checks. For convenience of exposition, we report again the results for the wine industry in the first two columns. In the remaining columns, the AOC "overall" dummy is now identifying all the bread and meat producers that are located in municipalities where AOC wine is produced. Essentially, we are thus imposing on these alternative industries the same regional partition of the wine industry. If our results for wine would be driven by generic regional effects, rather than horizontal differentiation, then we would expect to find similar results for meat and bread firms. As it can be seen, the results are instead much different. Only in column 6, for bread producers, we find a negative and mildly significant interaction term between the AOC "overall" dummy and productivity, but the latter is associated with a negative and significant AOC dummy, contrary to the results for wine. Overall, this evidence suggests that our findings for wine producers are not driven by generic regional effects, but are indeed associated with the specific geographical-based segmentation of the wine industry, which reflects differences in the degree of horizontal differentiation.

A similar counterfactual analysis is performed in Table 11 with respect to the non-parametric test of Proposition 1, as presented in Table 9. That is, we have split meat and bread producers in ten different classes of productivity, based both on Lev-Pet and Index Numbers TFP estimates. We have then calculated the average firm size (nominal turnover) within each class, and we have taken the ratio of average-size figures between non-AOC and AOC firms, separately for the two industries. Also in this case, the results for meat and bread producers are very different than those reported in Table 9 for the wine industry. In particular, for the meat industry the average-size relation between AOC and non-AOC firms looks pretty erratic across different deciles of productivity. Indeed, while average size grows with productivity for both groups of producers (consistent with the theory of heterogeneous firms), there is not a clear pattern concerning the ratio of average-size figures across groups. For instance, non-AOC producers are on average smaller than AOC producers both in the first and last classes of productivity, while they are larger in the sixth class. For the bread industry, instead, we find that firms located outside any AOC area are larger than AOC firms, on average, at all levels of productivity. Overall, for both the alternative industries, there is no inversion of the average-size relation across groups as observed for the 'AOC-segmented' wine industry. From this we can further infer that our segmentation of the wine industry based on the AOC dummy is not capturing generic regional effects, but rather a market segmentation that is specific to the wine industry, consistent with our theoretical framework.

Table 11: Proposition 1 - non-parametric test for alternative industries

NACE 1513 (meat products)				**NACE 1581 (bread products)**			
	Average firm size: (turnover, 000s eur)				Average firm size: (turnover, 000s eur)		
Deciles of ln(TFP) Lev-Pet	non -AOC	Within AOCs	Ratio	Deciles of ln(TFP) Lev-Pet	non -AOC	Within AOCs	Ratio
1	163.6	299.8	0.55	1	175.4	164.1	1.07
2	422.8	354.7	1.19	2	273.1	247.7	1.10
3	444.7	473.0	0.94	3	297.2	281.4	1.06
4	542.0	601.4	0.90	4	358.7	288.1	1.24
5	628.7	737.2	0.85	5	364.8	313.0	1.17
6	879.9	519.4	1.69	6	402.1	326.6	1.23
7	783.5	439.3	1.78	7	478.0	346.9	1.38
8	705.4	537.4	1.31	8	520.3	332.8	1.56
9	927.0	634.5	1.46	9	496.2	378.4	1.31
10	654.6	846.9	0.77	10	591.6	415.6	1.42
Deciles of ln(TFP) Index Numbers	non -AOC	Within AOCs	Ratio	Deciles of ln(TFP) Index Numbers	non -AOC	Within AOCs	Ratio
1	187.7	285.3	0.66	1	175.6	173.6	1.01
2	617.0	383.8	1.61	2	274.7	252.8	1.09
3	498.6	573.3	0.87	3	377.4	275.3	1.37
4	508.4	669.5	0.76	4	379.5	312.0	1.22
5	859.5	588.5	1.46	5	455.2	319.7	1.42
6	956.7	544.2	1.76	6	430.4	324.6	1.33
7	799.3	472.4	1.69	7	405.5	378.5	1.07
8	670.0	456.0	1.47	8	449.2	314.9	1.43
9	645.5	593.4	1.09	9	529.5	343.7	1.54
10	630.9	779.0	0.81	10	419.5	361.9	1.16

5.3 Refinements and additional results

Up until now, consistent with our theoretical model, we have performed the empirical tests of Proposition 1 by partitioning the wine industry in just two segments. In particular, we have treated all the AOC firms as belonging to a single high-differentiation market segment, as opposed to the low-differentiation segment of non-AOC producers. This is the most clear-cut subdivision of the industry that can be made. As a refinement of the analysis, one could also try to explore the role of variation in horizontal differentiation across the ten AOC "macro" areas, which could be treated as being separate market segments. For this purpose, Table 12 presents the results of two regressions where we have included separate dummies for each AOC area, and all their interactions with TFP, thus allowing for both intercepts and slopes in the size/TFP relation to vary across ten different market segments, while keeping the non-AOC producers as the control group. Results are in line with the idea that different AOC areas may constitute different market segments: both the AOC dummies and their interactions are jointly different from zero, and statistically different from each other.

Table 12: Proposition 1 - a refinement of the analysis

Dep. Variable: ln(size)	Lev-Pet	Index Numbers
	(1)	(2)
ln(TFP)	1.656***	1.088***
	(0.000)	(0.000)
Constant	-0.390***	7.623***
	(0.000)	(0.000)
AOC dummies	yes	yes
*ln(TFP) * AOC dummies*	yes	yes
H0: All intercepts equal to zero, F-stat	24.06	13.94
p-value	(0.000)	(0.000)
H0: All intercepts equal, F-stat	26.38	13.85
p-value	(0.000)	(0.000)
H0: All interactions equal to zero, F-stat	25.67	4.57
p-value	(0.000)	(0.000)
H0: All interactions equal, F-stat	28.09	3.69
p-value	(0.000)	(0.000)
N. of obs.	2894	2893
R-sq	0.70	0.08

Standard errors are clustered within AOC areas and within the control group. ***,**,* indicate significance at the 1, 5 and 10% level, respectively.

A relevant question is if we can relate these significant differences in intercepts and slopes across AOC areas to a varying degree of horizontal differentiation. In Table 13 we focus on AOC producers only, and we employ the number of sub-appellations (in logs) within each "macro" AOC area as a proxy for the specific level of horizontal differentiation of each segment, with a higher number of sub-appellations capturing a higher level of horizontal differentiation. As reported in Table 1, such a number goes from 1, in the case of Champagne, up to 100 in the case of Bourgogne. This empirical approach is motivated by the fact that, as discussed in Section 4.1, a large number of sub-appellations is the result of a historical recognition of significant differences across distinct wine varieties within a region (Barham, 2003; Crozet et al., 2012). In line with the previous econometric tests, we then include in the regressions both the linear term of the number of sub-appellations and its interaction with TFP. The results are once again consistent with our theoretical model, as we find a positive and significant coefficient for the linear term, and a negative and significant one for the interaction term. Hence, an increasing degree of horizontal differentiation, as proxied by a larger number of sub-appellations, is found to be associated with a higher intercept and a lower slope of the size/productivity relation also across AOC areas.[24]

Table 13: Proposition 1 - a refinement of the analysis

Dependent variable: ln(size)

TFP estimated through:	Lev-Pet	Index Numbers
	(1)	(2)
ln(TFP)	1.628***	0.761***
	[0.029]	[0.101]
ln(number of sub-appellations)	0.348***	0.051***
	[0.062]	[0.018]
*ln(TFP)*ln(number of sub-appellations)*	-0.057***	-0.397***
	[0.012]	[0.041]
Constant	-0.477***	7.861***
	[0.158]	[0.053]
N. of obs.	2642	2491
R-sq	0.66	0.04

Standard errors are clustered within AOC areas. ***, **, * indicate significance at the 1, 5 and 10% level, respectively.

Finally, we report an econometric test of Proposition 2, concerning the productivity premium of exporters. Consistent with Melitz-Ottaviano (2008), as well as any other model of international trade with heterogeneous firms, our model predicts that the most productive firms self-select into exporting within each market segment. However, when introducing asymmetric product differentiation, the cut-off levels of productivity inducing self-selection are heterogeneous across different segments. That is, firms' exporting behavior is systematically related to the segment-specific level of horizontal differentiation. More specifically, Proposition 2 states that selection into exporting in the high-differentiation segment requires a smaller productivity advantage than in the low-differentiation segment. In other words, equally productive firms are more likely to export if operating in the high-differentiation segment.

Table 14 presents an econometric test of Proposition 2. In particular, we run a probit analysis on the probability that a firm reports a positive value of exports in a given year, including the AOC "overall" dummy, along with the lagged export status and TFP level, as it is standard in this type of analysis. Columns 1 and 2 display the results for the wine industry. Probit marginal effects (evaluated at the mean) are reported. Consistent with earlier findings in the literature, exporting is found to be persistent and positively associated with firm productivity, however measured. In addition, in line with our model, the AOC "overall" dummy is estimated to be

[24] The lower number of observations in column 2 of Table 13 is due to the omission of a number of outliers (some 15 firms), detected when looking at the distribution of Index Numbers TFP across the sub-appellations.

positive and statistically different from zero. In particular, the magnitude of the coefficients indicate that, ceteris paribus, the probability that an AOC firm exports is about 12% higher than for an equally productive non-AOC firm, both when using Lev-Pet and Index Numbers estimates of TFP.

Table 14: Proposition 2 - probit marginal effects

Dependent variable: export status (binary)

Industry:	NACE 1593 (wine)		NACE 1513 (meat products)		NACE 1581 (bread products)	
TFP estimated through:	Lev-Pet	Index Numbers	Lev-Pet	Index Numbers	Lev-Pet	Index Numbers
	(1)	(2)	(3)	(4)	(5)	(6)
ln(TFP) t-1	0.065***	0.061***	0.038	0.028	-0.014	-0.005
	[0.001]	[0.016]	[0.027]	[0.028]	[0.011]	[0.008]
Export Status (t-1)	0.802***	0.810***	0.673***	0.682***	0.628***	0.629***
	[0.014]	[0.015]	[0.081]	[0.079]	[0.055]	[0.055]
AOC Overall Dummy	0.117***	0.124***	0.012	0.011	0.002	0.003
	[0.004]	[0.001]	[0.010]	[0.010]	[0.004]	[0.004]
N. of obs.	2052	2052	836	833	4470	4466
Pseudo R-sq	0.57	0.56	0.48	0.48	0.40	0.40

Standard errors are clustered within the AOC and the control group. ***, **, * indicate significance at the 1, 5 and 10% level, respectively.

The latter finding can also be validated against generic regional effects that could operate across industries and spuriously drive our results (Koenig et al., 2010). To that extent, Columns 3 to 6 of Table 14 report the results of the same probit regression for the meat and bread industries. Also in this case, data for the latter industries refer to firms which are located in the same municipalities covered by the wine-producers database, and the same AOC "overall" dummy is considered as a regional effect. As it can be seen, in both cases the AOC dummy is not significantly different from zero. This reinforces our claim that selection into exporting in the high-differentiation segment is milder than in the low-differentiation segment, consistent with our model.

Still, the above evidence, although in line with Proposition 2, does not allow us to single out horizontal differentiation as the only driver of our results. In fact, in this case unobserved quality factors may potentially provide an alternative explanation for the uncovered relation between TFP and export status across market segments. In particular, assuming that AOC wines are of higher average quality, and assuming the existence of quality constraints on the export markets, one could obtain similar results as in the first two columns of Table 14. For this reason, the empirical evidence on exporting shall only be seen as suggestive and complementary with respect to the previous analyses. The real test of the model has been carried on Proposition 1 (i.e. on the relation between productivity and firm size), which is the distinctive implication of our approach based on differences in horizontal differentiation across market segments.

6 Conclusion

In this paper we have introduced an asymmetrically differentiated demand system across multiple market segments within a single industry, in the context of the productivity-based literature on firm heterogeneity. In particular, we have extended the Melitz-Ottaviano (2008) framework by modeling asymmetric horizontal product differentiation across market segments. In so doing, we have been able to derive more complex relations between productivity, size and firms' export engagement, all crucially moderated by the degree of product differentiation within each segment. Such an enriched theoretical framework has proved useful in explaining the behavior of French firms in the wine industry, which is characterized by officially defined exogenous market segments.

Horizontal product differentiation can thus be seen as a second relevant dimension of heterogeneity, besides productivity, allowing us to enhance the richness of theoretical results with respect to earlier models. Other papers have identified quality (vertical differentiation) as an additional relevant factor of firm heterogeneity. In particular, quality has been modeled to explain the non obvious empirical relation between productivity and exporting. With respect to this literature, our setup explicitly abstracts from quality differences across producers, in order to isolate the role of horizontal product differentiation. The advantage of such a choice is that it allows to account not only for non-linearities in the relation between productivity and export status, but also, within the same theoretical framework, for the imperfectly linear relation between productivity and size observed in the data. The latter is a distinctive feature of our model, as it cannot be obtained in a quality-augmented firm heterogeneity setting.

Our results should thus be seen as complementary to the ones obtained insofar by the quality literature, and open the way to new promising lines of research attempting to combine the two approaches. In particular, efforts shall be devoted to studying the implications of the interplay between heterogeneous levels of horizontal differentiation and firms' quality decisions, across different market segments within an industry.

References

Ahn, J., Khandelwal, A. and S.J. Wei, 2011. The role of intermediaries in facilitating trade, Journal of International Economics 84: 73-85.

Antoniades, A., 2012. Heterogeneous Firms, Quality, and Trade, mimeo, Princeton University and Georgetown University.

Aw, B.Y., Xiaomin, C. and M.J. Roberts, 2001. Firm-Level Evidence on Productivity Differentials and Turnover in Taiwanese Manufacturing, Journal of Development Economics 66: 51-86.

Baldwin, R. and J. Harrigan, 2011. Zeros, Quality and Space: Trade Theory and Trade Evidence, American Economic Journal: Microeconomics 3: 60-88.

Barham, E., 2003. Translating Terroir: the Global Challenge of French AOC labelling, Journal of Rural Studies 19: 127-138.

Bernard, A.B., Eaton, J., Jensen, J.B. and S.S. Kortum, 2003. Plants and Productivity in International Trade, American Economic Review 93: 1268-1290.

Brooks, E.L., 2006. Why don't firms export more? Product quality and Colombian plants, Journal of Development Economics 80: 160-178.

Combris, P., Lecocq, S. and M. Visser, 1997. Estimation of a Hedonic Price Equation for Bordeaux Wine: Does Quality Matter?, Economic Journal 107: 390-402.

Combris, P., Lecocq, S., and M. Visser, 2000. Estimation of a Hedonic Price Equation for Burgundy Wine, Applied Economics 32: 961-967.

Cranfield, J.A.L., 2002. Persistence of Price-Cost Margins in the U.S. Food and Tobacco Manufacturing Industries: A Dynamic Single Index Model Approach, Journal of Food Distribution Research 33: 35-49.

Crino', R. and P. Epifani, 2012. Productivity, Quality and Export Behavior, Economic Journal 122: 1206-1243.

Cross, R., A.J. Plantinga and R.N. Stavins, 2011. What Is the Value of Terroir?, American Economic Review - Papers and Proceedings 101: 152-156.

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.

Di Comite, F., Thisse, J.F. and H. Vandenbussche, 2014. Verti-zontal Differentiation in Export Markets, Journal of International Economics 93: 50-66.

Dodds, W., 1991. In Search of Value: How Price and Store Name Information Influence Buyers' Product Perceptions, The Journal of Consumer Marketing 8: 15-24.

Eckel, C., Iacovone, L., Javorcik, B. and J.P. Neary, 2011. Multi-Product Firms at Home and Away: Cost-versus Quality-based Competence, CEPR Discussion Paper No. 8186.

Evans, D.S., 1987. The Relationship between Firm Growth, Size, and Age: Estimates for 100 Manufacturing Industries, Journal of Industrial Economics 35: 567-581.

Foster, L., Haltiwanger, J. and C. Syverson, 2008. Reallocation, firm turnover, and efficiency: selection on productivity or profitability?, American Economic Review 22: 527-534.

Gergaud, O. and V. Ginsburgh, 2008. Natural Endowments, Production Technologies and the Quality of Wines in Bordeaux. Does Terroir Matter?, Economic Journal 118: 142-157.

Goldberg, P. K., 1995. Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry, Econometrica 63: 891-952.

Goldstein, R., Almenberg, J., Dreber, A., Emerson, J.W., Herschkowitsch, A. and J. Katza, 2008. Do More Expensive Wines Taste Better? Evidence from a Large Sample of Blind Tastings, Journal of Wine Economics 3: 1-9.

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.

Johnson, R., 2012. Trade and Prices with Heterogeneous Firms, Journal of International Economics 86: 43-56.

Kardes, F.R., Cronley, M.L., Kellaris, J.J. and S.S. Posavac, 2004. The Role of Selective Information Processing in Price-Quality Inference, Journal of Consumer Research, 31: 368-374.

Khandelwal, A., 2010. The Long and Short (of) Quality Ladders, Review of Economic Studies 77: 1450-1476.

Kneller, R. and Z. Yu, 2008. Quality Selection, Chinese Exports and Theories of Heterogeneous Firm Trade, GEP Discussion Paper 2008/44, Nottingham.

Koenig, P., Mayneris, F. and S. Poncet, 2010. Local Export Spillovers in France, European Economic Review 54: 622-641.

Konings, J. and H. Vandenbussche, 2013. Antidumping Protection Hurts Exporters: Firm-level Evidence, Review of World Economics 149: 295-320.

Kugler, M. and E. Verhoogen, 2012. Prices, Plant Size, and Product Quality, Review of Economic Studies 79: 307-339.

Lecocq, S. and M. Visser, 2006. What Determines Wine Prices: Objective vs. Sensory Characteristics, Journal of Wine Economics 1: 42-56.

Levinsohn, J. and A. Petrin, 2003. Estimating Production Functions Using Inputs to Control for Unobservables, Review of Economic Studies 70: 317-341.

Melitz, M.J., 2003. The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity, Econometrica 71: 1695-1725

Melitz, M.J. and G.I.P. Ottaviano, 2008. Market Size, Trade, and Productivity, Review of Economic Studies 75: 295–316.

Monroe, K. B., 1976. The Influence of Price Differences and Brand Familiarity on Brand Preferences, Journal of Consumer Research 3: 42-49.

Ottaviano, G. I. P., Tabuchi, T. and J.F. Thisse, 1998. Agglomeration and trade revisited, CEPR Discussion Paper 1903, London.

Parker, R.M., 1985. Bordeaux, The Definitive Guide for the Wines Produced since 1961, New-York: Simon and Schuster.

Quandt, R.E., 2007. On wine bullshit: some new software, Journal of Wine Economics 2: 129-135.

Spearot, A. C., 2013. Variable Demand Elasticities and Tariff Liberalization, Journal of International Economics 89: 26-41.

Van Beveren, I., 2012. Total factor productivity estimation: A practical review, Journal of Economic Surveys 26: 98-128.

Veale, R. and P. Quester, 2008. Consumer Sensory Evaluations of Wine Quality: The Respective Influence of Price and Country of Origin, Journal of Wine Economics 3: 10-29.

Verdù Jover, A., Lloréns Montes, F.J. and M.M. Fuentes Fuentes, 2004. Measuring Perceptions of Quality in Food Products: the Case of Red Wine, Food Quality and Preference 15: 453-469.

Verhoogen, E., 2008. Trade, Quality Upgrading and Wage Inequality in the Mexican Manufacturing Sector, Quarterly Journal of Economics 123: 489-530.

Weil, R.L., 2001. Parker vs. Prial: The Death of the Vintage Chart, Chance 14: 27-31.

Weil, R.L., 2005. Analysis of Reserve and Regular Bottlings: Why Pay For a Difference Only the Critics Claim to Notice?, Chance 18: 9-15.

Weil, R.L., 2007. Debunking Critics' Wine Words: Can Amateurs Distinguish the Smell of Asphalt From the Taste of Cherries?, Journal of Wine Economics 2: 136-144.

Wilson, J., 1998. Terroir: The Role of Geology, Climate and Culture in the Making of French Wines, Berkeley: University of California Press.

Appendix 1: Expressions of the f-parameters

$$f_1^l = \alpha - \eta^l D\left[\left(\gamma^h + N^h\eta^h\right)N^l\alpha - N^lN^h\alpha\beta\right] - \beta D\left[\left(\gamma^l + N^l\eta^l\right)N^h\alpha - N^lN^h\alpha\beta\right] \tag{20}$$

$$f_1^h = \alpha - \eta^h D\left[\left(\gamma^l + N^l\eta^l\right)N^h\alpha - N^lN^h\alpha\beta\right] - \beta D\left[\left(\gamma^h + N^h\eta^h\right)N^l\alpha - N^lN^h\alpha\beta\right] \tag{21}$$

$$f_2^l = D\eta^l\left(\gamma^h + N^h\eta^h\right)N^l - D\beta^2N^lN^h \tag{22}$$

$$f_2^h = D\beta\left(\gamma^h + N^h\eta^h\right)N^l - D\eta^h\beta N^lN^h \tag{23}$$

$$f_3^l = D\beta(\gamma^l + N^l\eta^l)N^h - D\eta^l\beta N^lN^h \tag{24}$$

$$f_3^h = D\eta^h(\gamma^l + N^l\eta^l)N^h - D\beta^2N^lN^h \tag{25}$$

Appendix 2: Market segment-specific cut-offs

From equations 2, 7 and 9, and equations 3, 8 and 10, respectively, we have that:

$$p^l_{\max} = f_1^l + f_2^l\overline{p}^l + f_3^l\overline{p}^h = \alpha - \eta^lQ^l - \beta Q^h \tag{26}$$

$$p^h_{\max} = f_1^h + f_2^h\overline{p}^l + f_3^h\overline{p}^h = \alpha - \eta^hQ^h - \beta Q^l \tag{27}$$

Hence, we can prove that $p^h_{\max} > p^l_{\max}$, and thus market segment-specific cut-offs are a direct function of the degree of product differentiation, if:

$$\alpha - \eta^hQ^h - \beta Q^l > \alpha - \eta^lQ^l - \beta Q^h \tag{28}$$

Assuming $\eta^l = \eta^h = \eta > \beta$, then $p^h_{\max} > p^l_{\max}$ if $Q^l > Q^h$. From eq. 6, where $D > 0$, we then have that $Q^l > Q^h$ if:

$$(\alpha - \overline{p}^l)(N^lN^h\eta + N^lN^h\beta + \gamma^hN^l) > (\alpha - \overline{p}^h)(N^lN^h\eta + N^lN^h\beta + \gamma^lN^h) \tag{29}$$

We can define: $K = N^lN^h\eta + N^lN^h\beta$. Then, through some simple algebra, we have that $Q^l > Q^h$ (and thus $p^h_{\max} > p^l_{\max}$) if:

$$N^l > \left(\frac{\alpha - \overline{p}^h}{\alpha - \overline{p}^l}\right)\frac{\gamma^l}{\gamma^h}N^h + \left(\frac{\overline{p}^l - \overline{p}^h}{\alpha - \overline{p}^l}\right)\frac{K}{\gamma^h} \tag{30}$$

The latter inequality always holds, for any value of N^l, as long as its right hand side is negative. Assuming that N^l and N^h are large enough, then K is an order of magnitude greater than N^h. The latter implies that the sign of the right hand side essentially depends on the relation between $\overline{p}^h$ and $\overline{p}^l$. Now, if $p^h_{\max} < p^l_{\max}$, then $c^h_D < c^l_D$ and thus, from eq. 14, the distribution of prices in Ω^h would be stochastically dominated by the one in Ω^l. This in turn would imply that $\overline{p}^h < \overline{p}^l$, which would prove the inequality false. By contradiction, then, it has to be that $p^h_{\max} > p^l_{\max}$.

Appendix 3: Pareto parametrization

We follow Melitz-Ottaviano (2008) and assume that the productivity draws $(1/c)$ follow a Pareto distribution with lower bound $(1/c_M)$ and shape parameter $k \geq 1$. The cumulative distribution function for the cost draws (the inverse of productivity) can then be written as $G(c) = \left(\frac{c}{c_M}\right)^k$ with $c \in [0, c_M]$.

The distribution of cost draws for the two sets of surviving firms (in Ω^l and Ω^h) is a truncation of $G(c)$, with upper bound c^ξ_D. Since a truncated Pareto distribution is still Pareto distributed with the same shape parameter, we have that:

$$G^\xi_D(c) = \left(\frac{c}{c^\xi_D}\right)^k \ , \ c \in [0, c^\xi_D] \ , \ \ \xi = l, h \tag{31}$$

In equilibrium, the expected firm profits (net of sunk entry costs) for a potential entrant need to be equal to zero, for both market segments. Hence, from equation 18, we can write:

$$\int_0^{c^\xi_D} \pi^\xi(c) dG(c) = \frac{L}{4\gamma^\xi} \int_0^{c^\xi_D} \left(c^\xi_D - c\right)^2 dG(c) = f^\xi_E \quad \text{with } \xi = l, h \tag{32}$$

Using the fact that $\frac{dG(c)}{dc} = g(c)$, we can write $g(c) = \frac{kc^{k-1}}{c^k_M}$, solve the Riemann–Stieltjes integral in eq. 32 and derive the following parametric expression for the cut-offs:

$$c^\xi_D = \left[\frac{2(k+1)(k+2)(c_M)^k f^\xi_E \gamma^\xi}{L}\right]^{\frac{1}{(k+2)}} \qquad \text{with } \xi = l, h \tag{33}$$

Assuming for simplicity that $f^l_E = f^h_E$, then $\gamma^h > \gamma^l$ implies once again $c^h_D > c^l_D$[25]. The Pareto parametrization then allows us to obtain average measures of firm performance in terms of the cost cut-off c^ξ_D. In particular, for $\xi = l, h$ we have[26]:

$$\overline{c}^\xi = \frac{k}{k+1} c^\xi_D \tag{34}$$

$$\overline{p}^\xi = \frac{2k+1}{2k+2} c^\xi_D \tag{35}$$

[25] This result would be even stronger when assuming $f^h_E > f^l_E$, in line with the plausible idea that pre-entry product development costs are larger for the highly differentiated product varieties in Ω^h than for the more standardized ones in Ω^l.

[26] The average figure for the generic performance measure z has been obtained as follows: $\overline{z} = [\int_0^{c^\xi_D} z(c) dG(c)]/G(c^\xi_D)$, starting from the firm level performance measures defined in eq. 14-18.

$$\overline{\mu}^{\xi} = \frac{1}{2}\frac{1}{k+1}c_D^{\xi} \tag{36}$$

$$\overline{q}^{\xi} = \frac{L}{2\gamma^{\xi}}\frac{1}{k+1}c_D^{\xi} = \frac{(k+2)(c_M)^k}{(c_D^{\xi})^{k+1}}f_E^{\xi} \tag{37}$$

$$\overline{r}^{\xi} = \frac{L}{2\gamma^{\xi}}\frac{1}{k+2}(c_D^{\xi})^2 = \frac{(k+1)(c_M)^k}{(c_D^{\xi})^{k}}f_E^{\xi} \tag{38}$$

$$\overline{\pi}^{\xi} = f_E^{\xi}\frac{(c_M)^k}{(c_D^{\xi})^k} \tag{39}$$

Having shown that $c_D^h > c_D^l$ (given $f_E^l = f_E^h$), it follows that firms in Ω^h are on average less productive (higher $\overline{c}$), they charge higher average prices and earn higher average mark-ups. However, notwithstanding such higher prices and mark-ups, firms in Ω^h are on average smaller in terms of produced output, and thus earn on average less (total) revenues and profits.

Appendix 4: Productivity and export across market segments

We follow Melitz-Ottaviano (2008) and consider two countries: H and F, with L^H and L^F consumers respectively. Consumers in the two countries share the same preferences, resulting in the same inverse demand functions as in eq. 7 and 8. In both countries we have the same market segmentation as before (Ω^h and Ω^l) in terms of product differentiation. Firms operating in one (and only one) market segment can produce in one country and sell in the other by incurring an iceberg-type per-unit trade cost $\tau^{\delta} > 1$, where δ indexes the destination country H or F. There are no fixed-costs of exporting, and the per-unit iceberg trade cost $\tau > 1$ is assumed to be the same for both goods in Ω^h and Ω^l, for each country.

For each market segment ξ we now have a δ country-specific maximum price denoted $p_{\max}^{\delta\xi}$, such that a variety displays a positive consumption level. Since national markets are segmented and production is characterized by constant returns to scale, each firm in country δ solves two distinct profit maximization problems, one for the domestic and one for the export market, within each and the same market-segment ξ. Solving within each market segment, we can derive:

$$q_D^{\delta}(c) = \frac{L^{\delta}}{\gamma^{\xi}}[p_D^{\delta}(c) - c] \quad , \quad \delta = H, F \text{ and } \xi = l, h \tag{40}$$

$$q_X^{\delta}(c) = \frac{L^{\psi}}{\gamma^{\xi}}[p_X^{\delta}(c) - \tau^{\psi}c] \quad , \quad \delta = H, F \quad , \psi \neq \delta \text{ and } \xi = l, h \tag{41}$$

where $p_D^{\delta}(c)$ and $q_D^{\delta}(c)$ are the domestic profit maximizing price and quantity, while $p_X^{\delta}(c)$ and $q_X^{\delta}(c)$ are the profit maximizing delivered price and quantity for the export market, denoted with $\psi \neq \delta$.

As only firms earning non-negative profits in a certain market (domestic vs. foreign) will decide to sell in that market, this determines the existence of two different cost cut-offs for domestic versus foreign sales in each country-market segment pair. We call $c_D^{\delta\xi}$ the upper bound cost for firms in market segment ξ selling in their domestic market (country δ). The upper bound cost for exporters to country ψ is instead $c_X^{\delta\xi}$. These cut-offs must satisfy:

$$c_D^{\delta\xi} = \sup\left\{c : \pi_D^{\delta\xi}(c) > 0\right\} = p_{\max}^{\delta\xi} \tag{42}$$

$$c_X^{\delta\xi} = \sup\left\{c : \pi_X^{\delta\xi}(c) > 0\right\} = \frac{p_{\max}^{\psi\xi}}{\tau^{\psi}} \tag{43}$$

The last equation clearly shows how trade costs make it harder for exporters to break even relative to domestic producers selling in their home market, in line with Melitz-Ottaviano (2008) and with all the other models of international trade where heterogeneous firms self-select into exporting. However, in our model the cost cut-off for exporting is also market segment-specific. In particular, it is easy to prove Proposition 2, stating that "self-selection into exporting in Ω^h requires a relatively smaller productivity premium than in Ω^l". Indeed, from equations 42 and 43 we have that $c_X^{\delta\xi} = c_D^{\psi\xi}/\tau^{\psi}$. Since $c_D^{\psi l} < c_D^{\psi h}$ and $\tau^{\psi} > 1$ is the same for both market segments, it follows that $c_X^{\delta l} < c_X^{\delta h}$.