

Matthieu Crozet, Keith Head and
Thierry Mayer

Quality sorting and trade: firm-level evidence for French wine

EFIGE working paper 07
April 2009

Funded under the
Socio-economic
Sciences and
Humanities
Programme of the
Seventh
Framework
Programme of the
European Union.

LEGAL NOTICE: The
research leading to these
results has received
funding from the
European Community's
Seventh Framework
Programme (FP7/2007-
2013) under grant
agreement n° 225551.
The views expressed in
this publication are the
sole responsibility of the
authors and do not
necessarily reflect the
views of the European
Commission.



The EFIGE project is coordinated by Bruegel and involves the following partner organisations: Universidad Carlos III de Madrid, Centre for Economic Policy Research (CEPR), Institute of Economics Hungarian Academy of Sciences (IEHAS), Institut für Angewandte Wirtschaftsforschung (IAW), Centro Studi Luca D'Agliano (Ld'A), Unitcredit Group, Centre d'Etudes Prospectives et d'Informations Internationales (CEPII).
The EFIGE partners also work together with the following associate partners: Banque de France, Banco de España, Banca d'Italia, Deutsche Bundesbank, National Bank of Belgium, OECD Economics Department.

Quality sorting and trade: Firm-level evidence for French wine*

Matthieu Crozet[†] Keith Head[‡] Thierry Mayer[§]

April 30, 2009

Abstract

Investigations of the effect of quality differences on heterogeneous performance in exporting have been limited by lack of direct measures of quality. We examine exports of French wine, matching the exporting firms to producer ratings from two wine guides. We show that high quality producers export to more markets, charge higher prices, and sell more in each market. More attractive markets are served by exporters that, on average, make lower rated Champagne. Market attractiveness has a weakly negative effect on prices and a strongly positive effect on quantities, confirming the sign predictions of a simple quality sorting model.

Methodologically, we make several contributions to the literature. First, we propose an estimation method for regressions of firm-level exports on ability measures and use Monte Carlo simulations to show that it corrects a severe selection bias present in OLS estimates. Second, we show how the means of quality, price, and quantity for exporters to a given market can be used to recover estimates of core parameters (which we compare with firm-level estimates) and discriminate between productivity and quality-sorting versions of the Melitz model. Our new method regresses country means on an index of each country's attractiveness and the fixed costs of entering it. We compare our method, which utilizes explanatory variables estimated in the firm-level regressions, to the conventional approach that relies on a reduced-form relationship with proxies for attractiveness and fixed costs.

*This paper is produced as part of the project European Firms in a Global Economy: Internal policies for external competitiveness (EFIGE), a collaborative project funded by the European Commissions Seventh Research Framework Programme, Contract # 225551. The views expressed in this publication are the sole responsibility of the authors and do not necessarily reflect the views of the European Commission. We thank Isaac Holloway for research assistance. Participants in conferences and seminars at Louvain-la-Neuve, Nottingham, NOITS, ERWIT, LSE, George Washington University, and Kyoto University provided helpful suggestions.

[†]CEPII and University of Reims.

[‡]Sauder School of Business, University of British Columbia, and CEPR

[§]Paris School of Economics (Université de Paris 1), CEPII, and CEPR.

1 Introduction

Since firm-level data on trade have become available, researchers have documented overwhelming evidence of dramatic differences in export performance. Most firms do not export; the few that do tend to export relatively small shares of their output and export to only a handful of destinations.¹ Only the highest performing firms export substantial amounts to large sets of destinations. While the fact of performance differences is well-established, the source of this heterogeneity remains unclear.

Theoretical papers following the seminal work of Melitz (2003) mainly assume that the sorting of firms into export markets depends upon individual *productivity* draws. However, the proxies used for measuring productivity differences, such as value-added per worker (Bernard and Jensen, 1999) or sales in the home market (Eaton et al., 2008, and Yeaple, 2009) could be driven by primitives other than physical output per unit of input. Casual observation suggests that product *quality* differences are important in many industries. Presence and performance in foreign markets could therefore be driven by quality sorting, productivity sorting, or a combination of the two. The precise quantification of the role of quality in explaining trade outcomes has been hindered by the lack of direct measures of quality, forcing reliance on proxies such as unit values.

This paper studies the exports of Champagne producers, where firm-destination export flows can be matched to firm quality ratings from wine guides. Firm-level regressions illustrate how directly measured quality affects the prices firms charge, the set of countries to which they export, and the amounts they export to each destination. The firm-level regressions show that there is a payoff to quality in terms of greater presence in export markets. Since direct measures of quality are only available for particular products, we also consider tests of the quality sorting hypothesis using indirect evidence from the average prices and quantities of Champagne exported to different destinations. Under standard theoretical assumptions (namely Pareto distributed heterogeneity), there are discriminating predictions for both average price and quantity. We find that indirect tests corroborate the direct evidence for the hypothesis that quality sorting is important for the Champagne industry. Since our model and estimation methods were not tailored for application to this industry, we believe they can be usefully applied in other settings.

Work on quality and trade began with the question of what makes a country export higher quality goods—as inferred from unit values. Schott (2004) finds that within goods categories, unit values tend to increase with the exporters’ per capita income, capital to labor ratio, skill ratio, and the capital intensity of production. Hummels and Klenow (2005) find that, within categories, price and quantity indexes rise with origin-country income per capita. Economists have also investigated which countries tend to *import* high quality goods. Hallak (2006) finds some evidence that richer countries have relatively greater demand for high quality, again measured by unit values. Hummels and Skiba (2004) find that average FOB export prices rise with freight costs to a destination market. They interpret this as a confirmation of the Alchian-Allen (1964) effect (“shipping the

¹Bernard et al. (2007) summarize the evidence on US firms and Mayer and Ottaviano (2007) compile evidence for firms from seven European nations.

good apples out”).²

A more recent set of papers builds upon Melitz (2003) to consider the implications of heterogeneity in firm quality for patterns of trade at the industry level. Baldwin and Harrigan (2007) propose a model where lower productivity is more than compensated by higher product quality. Using product-level export data from the US, they confirm their model’s prediction that average prices are higher for long distances but decrease with destination GDP. With data for a wider set of developed countries, Baldwin and Ito (2007) corroborate the positive distance effect on export prices increase for a small, but significant, number of products. Moreover, they find that countries with a comparative advantage in raw materials exhibit less evidence of quality sorting. Johnson (2009) relates export prices to quality-adjusted price thresholds for exporting to different destinations. For the majority of sectors, export prices tend to be higher when markets are inferred to require greater ability for profitable entry. This is inconsistent with a homogeneous quality model in which high ability firms charge *low* prices. Echoing Schott (2004), Johnson also finds a home-country component of export prices that is highly correlated with per capita income.

The next step taken by the quality and trade literature confronts firm-level theories in which product quality drives exporter performance with *firm-level data*. Manova and Zhang (2009) analyze Chinese firm-level export prices to distinguish between several models of trade with heterogeneous firms. They find that none of the existing models can explain all aspects of exporter behavior, but still present evidence of quality sorting since firms that export to more destinations charge higher prices. Verhoogen (2008) hypothesizes that higher quality goods require higher quality workers and finds supportive evidence in his study of the performance of Mexican firms during the 1994 Peso crisis. Kugler and Verhoogen (2008) show that Colombian firms’ size and export propensities are positively correlated with input and output prices, corroborating the linkage between the quality of inputs and outputs. Hallak and Sivadasan (2008) also find a positive relationship between exporting and output prices using data from India, the United States, Chile and Colombia.³ Furthermore, Iacovone and Javorcik (2008) find that Mexican exporting firms charge higher prices, and that firms experience an increase in their price two years before they start exporting.

Our paper contributes to the quality and trade literature in terms of data and method. Contrasting with the existing literature, we focus on a particular industry, Champagne, where we can obtain direct quality measures from wine guides. This allows us to assess how well the quality inferred from prices corresponds to directly measured quality. One important advantage of the Champagne industry is that we are able to match firm-level quality measures with firm-level destination-specific exports obtained from customs dec-

²The Alchian-Allen effect relies upon freight costs that are less than proportional to product value. An increase in freight costs therefore lowers the relative price of high quality goods leading to an increase in their relative demand.

³Their model attributes firm-level heterogeneity to two separate draws: A classical productivity draw determines marginal cost and a “caliber” draw determines fixed costs. Unlike other models, the two dimensions of heterogeneity have independent impacts on the endogenous quality choice when combined with a minimum export quality requirement.

larations. Our theoretical model combines the Hallak (2006) specification of preferences for quality with the Baldwin and Harrigan (2007) assumption that higher quality entails higher marginal costs.

Methodologically, we make several contributions to the literature. We identify an important selection bias that is generic to firm-level regression of export outcomes on observed measures of firm quality (or productivity). A firm with low observed quality that manages to export to a difficult market must have an above-average realization of some unobserved determinant of export profitability. Our proposed solution is easy to implement and our Monte Carlo simulations indicate that it corrects almost all the selection bias. The firm-level regressions allow us to estimate the structural parameters of the model in terms of consumers' marginal valuation of quality and of producers' marginal cost of quality. We also develop new predictions for the heterogeneous quality model, relating conditional means of quality, price, and quantity to an index of market attractiveness and to the fixed costs of entering a market. We show how to estimate both of these explanatory variables using the firm-level regressions, rather than approximating cross-country variation in attractiveness and fixed costs with a set of gravity variables, as done in the existing literature. These regressions of country means on country characteristics generate an additional set of estimates of the core parameters of the model which can be compared to those estimated in the firm-level regressions.

The paper proceeds as follows. We first derive testable predictions from a model of firm-level heterogeneity in quality in section 2. Section 3 then proceeds to explain why applying this model to Champagne producers makes sense, and details the sources and main features of the data we use. The firm-level equations of the model are estimated in section 4.1, where we also back out the implied values of the key structural parameters. Section 4.2 estimates three sets of country mean relationships implied by the model. Our conclusion summarizes our results and outlines some desirable generalizations to the quality-sorting model.

2 Theory

The theory examined in this paper is based on work by Baldwin and Harrigan (2007), who introduce a cost-quality tradeoff in the model of Melitz (2003). We also draw on Eaton, Kortum, and Kramarz (2008). Wherever practical, we adopt the notation from prior work.

2.1 General Set-up

Consider a category of goods with a sub-utility function that is assumed to have a constant elasticity of substitution (CES), $\sigma > 1$, over the set, \mathcal{B}_d , of all varieties, j , available in country d :

$$U_d = \left(\int_{j \in \mathcal{B}_d} [a_d(j)s(j)^\gamma q(j)]^{\frac{\sigma-1}{\sigma}} dj \right)^{\frac{\sigma}{\sigma-1}}. \quad (1)$$

In this expression $q(j)$ denotes quantity of variety j consumed and $s(j)$ denotes its measured quality.⁴ Following Hallak (2006), the intensity of the consumers’ desire for quality is captured in parameter γ .

The $a_d(j)$ are destination d -specific demand parameters, a feature that Eaton, Kortum and Kramarz (2008) added to the Melitz model. Heterogeneity in the $a_d(j)$ provides a structural error term for firm-level regressions, as shown in subsection 2.6. There are a variety of possible interpretations for $a_d(j)$. In addition to cross-country variation in the tastes for the good made by firm j , it could also represent a firm’s network of connections with purchasers in each market. Foster, Haltiwanger, and Syverson (2008) argue that firm-level demand shocks—which they attribute in part to “webs of history-laden relationships between particular consumers and producers”—are important even for suppliers of the nearly homogenous goods they study. Firm-destination demand shocks allow the model to accommodate the fact that two firms with the same observed quality, s , differ in the amounts exported to the same country.⁵

The sub-utility enters full utility with a Cobb-Douglas parameter determining budget shares denoted b_d . The foreign country comprises M_d individuals with y_d income per capita. Aggregate expenditures on all \mathcal{B}_d varieties are given by $b_d y_d M_d$.

We assume that, within a detailed product classification, each firm exports a single variety.⁶ The solution to the consumers’ utility maximization is usually expressed in terms of trade-cost inclusive export values. Since the export values in our data set are reported on an FOB basis, we divide the destination d consumers’ desired expenditures on firm j by a trade cost factor, τ_d , to obtain FOB exports. Using $x_d(j)$ to denote FOB exports, we obtain

$$x_d(j) = \frac{[p_d(j)\tau_d/(a_d(j)s(j)^\gamma)]^{1-\sigma}}{\int_{i \in \mathcal{B}_d} [p_d(i)\tau_d/(a_d(i)s(i)^\gamma)]^{1-\sigma} di} b_d y_d M_d / \tau_d. \quad (2)$$

In this expression, the prices paid by consumers in d are given by $p_d(j)\tau_d$, where $\tau_d - 1$ is the *ad valorem* tariff equivalent of all trade costs incurred by firm j to sell in destination d .

Using $w(j)$ to denote a factor price index and $z(j)$ to denote factor productivity, a firm’s unit costs of production are given by $w(j)/z(j)$. This specification allows factor prices to vary across firms to take into account the idea (supported by Kugler and Verhoogen, 2009) that a firm making high-quality output might need more expensive inputs. The model entails a constant mark-up, $\sigma/(\sigma - 1)$, which can be factored out of the numerator and the denominator of equation (2). Taken together, these assumptions imply

⁴In our application “star” ratings in wine guides provide $s(j)$.

⁵Firm-destination demand shocks are one of several dimensions of flexibility that can be added to heterogeneous firm models to make them more consistent with actual trade patterns. Without a market-specific component to firm-level performance, all French firms that serve Thailand for instance, a remote and relatively small market, should also export to *all* “easier” countries. This is not the case for Champagne and Eaton et al. (2008) show that such “hierarchy” relationships do not hold strictly for French exports in general.

⁶The single-variety assumption is standard in the theory but is counterfactual for wine exporters, who often export multiple varieties. We will return to this issue in the empirical section.

export revenue from destination d is given by

$$x_d(j) = \left(\frac{w(j)/z(j)}{a_d(j)s(j)^\gamma} \right)^{1-\sigma} b_d y_d M_d \tau_d^{-\sigma} P_d^{\sigma-1}, \quad (3)$$

where the price index is defined in terms of quality-adjusted costs,

$$P_d \equiv \left(\int_{i \in \mathcal{B}_d} \left[\frac{\tau_d(i)w(i)/z(i)}{a_d(i)s(i)^\gamma} \right]^{1-\sigma} di \right)^{1/(1-\sigma)}.$$

We collect all country-specific determinants of exports into a single factor, A_d , defined as

$$A_d \equiv b_d y_d M_d \tau_d^{-\sigma} P_d^{\sigma-1}.$$

We refer to A_d as the “attractiveness” of a destination market. It depends positively on the size ($b_d y_d M_d$) and relative accessibility ($\tau_d^{-\sigma} P_d^{\sigma-1}$) of the market.

The net contribution to firm profits of destination d is given by

$$\pi_d(j) = x_d(j)/\sigma - F_d = ([z(j)/w(j)]s(j)^\gamma)^{\sigma-1} A_d \alpha_d(j)/\sigma - F_d, \quad (4)$$

where F_d is a destination-specific fixed cost for exporters, and $\alpha_d(j) \equiv a_d(j)^{\sigma-1}$ is a monotonic transformation of the utility parameter which captures the idiosyncratic firm-destination demand shock.⁷

2.2 The cost-quality tradeoff

We have so far allowed for heterogeneity in productivity (the standard approach following Melitz (2003)), factor prices, quality, and preferences. Four sources of heterogeneity is too many for a tractable model. One option is to hold w/z constant and have a model of pure (costless) quality variation. The problem with this is that in the Dixit-Stiglitz framework, mark-ups do not vary across firms and so quality has no independent effect on price. In this framework, the only way to make prices depend on quality is to stipulate a relationship between costs and quality. Thus, we imagine that when firms draw a “recipe” for quality level s , this entails a set of inputs and production methods such that marginal costs are increasing in s . Like Baldwin and Harrigan (2007), Mandel (2008), and Johnson (2009) we assume that this relationship takes the form of a power-function:

$$\frac{w(j)}{z(j)} = \omega s(j)^\lambda, \quad (5)$$

with $\lambda \geq 0$.⁸ Mandel (2008) and Johnson (2009) show derivations for power function tradeoffs between cost and quality in models where firms choose quality subject to a cost

⁷This transformation of a allows us to use the α notation of Eaton et al. (2008).

⁸Baldwin and Harrigan (2007) assume that firms draw a unit labour requirement ($1/z$) which is related to quality with an elasticity of θ . Our cost-quality parameter, λ , is related to their θ as follows: $\lambda = 1/(1 + \theta)$. Johnson (2009) assumes firms take a draw on something called ability, which is defined as quality divided by cost. Quality is a power function of ability with elasticity ϕ . Our λ corresponds to $(\phi - 1)/\phi$ in Johnson (2009).

of upgrading. We discuss the reasons to expect a positive relationship between costs and quality for Champagne in section 3.1.

Using equation (5), we can express the key equations of the model in terms of the remaining sources of heterogeneity. Export values are given by

$$x_d(j) = \omega^{1-\sigma} s(j)^\beta A_d \alpha_d(j) \quad (6)$$

where we use $\beta \equiv (\gamma - \lambda)(\sigma - 1)$ as an abbreviation for the elasticity of firm-level exports with respect to quality. Since $\sigma > 1$, a positive value of β implies that “quality pays,” i.e. consumer’s marginal valuation of quality exceeds the marginal cost to producers. The next two subsections show that this parameter determines how entry and means of firm-level quality, quantity, and price vary with destination attractiveness, A_d . In the empirical section we estimate β using firm-level data.

Individual firms charge FOB prices of

$$p_d(j) = \frac{\sigma}{\sigma - 1} \omega s(j)^\lambda. \quad (7)$$

Thus, the model predicts that firms charge the same FOB prices to all destinations and that these prices increase in quality with elasticity λ . A more general model might incorporate pricing-to-market via cross-country differences in σ , and hence the mark-up.

The parameterization of our model in terms of γ and λ is useful because when we set $\gamma = 0$ and $\lambda = -1$, utility does not depend directly on s , and costs are inversely related to s , i.e. $w(j)/z(j) = \omega/s(j)$. Hence, we can reinterpret $s(j)$ as a *productivity* draw. This allows us to compare the results of the quality-sorting model ($\gamma > \lambda \geq 0$) with the original Melitzian productivity-sorting model ($\gamma = 0, \lambda = -1$). As can be seen in equation (6), both models predict that higher s leads to higher export values. However, the price effects differ in sign. When s is interpreted as costly quality, it causes higher prices, but when s is a productivity draw, it causes low prices. This distinction is important later because it leads to contrasting predictions for conditional mean prices and quantities.

2.3 Entry threshold quality

The next step is to determine which firms export to a given destination. Substituting (6) into (4), we obtain

$$\pi_d(j) = \omega^{1-\sigma} s(j)^\beta A_d \alpha_d(j) / \sigma - F_d, \quad (8)$$

For any given value of α , the *zero-profit* quality level is

$$\hat{s}_d(\alpha) = \left(\frac{A_d \alpha}{\sigma F_d \omega^{\sigma-1}} \right)^{\frac{-1}{\beta}}. \quad (9)$$

Equation (9) shows the minimum quality needed to export profitably to destination, $\hat{s}_d(\alpha)$ (the ring on top is a mnemonic for zero profit), is a decreasing power function of how “easy” this market is for the average exporter (A_d). On the flip-side, higher fixed (F_d) or

variable (ω) costs increase the quality cut-off. In contrast to Baldwin and Harrigan (2007) and Johnson (2009), the quality threshold is not *only* country-specific. Rather, it depends on the individual firm's realization of its market-specific demand shifter, captured in α . This means even the lowest quality producer can enter any market as long as it obtains a sufficiently high α draw.

We assume that $s(j)$ and $\alpha_d(j)$ are independently distributed with probability density functions denoted $g(s)$ and $h(\alpha)$, respectively. The probability of entering is given by

$$\Pr[\pi_d(j) > 0] = \int_0^\infty \Pr[s(j) > \hat{s}_d(\alpha)]h(\alpha)d\alpha = \int_0^\infty (1 - G[\hat{s}_d(\alpha)])h(\alpha)d\alpha. \quad (10)$$

The precise functional form of $h(\alpha)$ can be left unspecified until we need to estimate the market-entry probability. However, to obtain closed-form relationships between conditional means and A_d , we need a tractable distribution for s . Following the recent literature (Chaney, 2008, Eaton et al., 2008, Helpman et al., 2008) we assume a Pareto distribution for firm-level heterogeneity. Letting \underline{s} denote the lower support of s , the CDF, $G(s)$, and PDF, $g(s)$, take the forms

$$G(s) = 1 - (s/\underline{s})^{-\kappa}, \quad \text{and} \quad g(s) = \kappa \underline{s}^\kappa s^{-\kappa-1}. \quad (11)$$

Plugging the Pareto CDF into equation (10) we express the probability of exporting to a market as

$$\Pr[\pi_d(j) > 0] = \hat{s}(1)^{-\kappa} \underline{s}^\kappa \mu_1 \quad \text{where} \quad \mu_1 \equiv \int_0^\infty \alpha^{\frac{\kappa}{\beta}} h(\alpha) d\alpha,$$

and $\hat{s}(1)$ is the zero-profit quality for a firm with $\alpha = 1$.

The continuum of firms assumption used in the monopolistic competition model allows us to equate $\Pr[\pi_d(j) > 0]$ with the fraction of firms that actually export to the market, which we denote N_d/N . Making this substitution and expressing $\hat{s}(1)$ in terms of its determinants from equation (9), we obtain

$$N_d/N = \left(\frac{A_d}{\sigma F_d \omega^{\sigma-1}} \right)^{\frac{\kappa}{\beta}} \underline{s}^\kappa \mu_1. \quad (12)$$

We refer to N_d as the ‘‘popularity’’ of market d and treat the set of firms at risk of exporting, N , as exogenous. Equation (12) shows that popularity is predicted to be a power function of the attractiveness of the market, A_d .

We now proceed to specify predictions for measurable aggregate statistics: the average quality, average price, and average quantity for each destination d . We show the relationship between the conditional expected values of these variables and both attractiveness, A_d , and popularity, N_d .

2.4 Conditional expectations of quality, price and quantity

The general form for the expected value of quality conditional on being a profitable exporter to some destination is

$$E[s \mid \pi_d(j) > 0] = \frac{\int_0^\infty \int_{\hat{s}_d(\alpha)}^\infty s g(s) h(\alpha) ds d\alpha}{\Pr[\pi_d(j) > 0]}. \quad (13)$$

A derivation shown in the appendix shows that the expected value of quality exported to a given market is

$$E[s \mid \pi_d(j) > 0] = \left(\frac{A_d}{\sigma F_d \omega^{\sigma-1}} \right)^{\frac{-1}{\beta}} \frac{\kappa}{\kappa - 1} (\mu_2 / \mu_1), \quad (14)$$

where μ_2 is another moment of the α distribution.⁹ Since $\beta > 0$ under both quality and productivity sorting, the expected ability (s) of exporters to d is always decreasing in the attractiveness of that market.

Although A_d is not directly observable, its chief determinants (population, per capita income) are measured and reasonable proxies are available for the others as we discuss below. Alternatively, we take advantage of the fact that the number of firms exporting to a market is directly observed and the model implies that this number is determined by the attractiveness index. Inverting equation (12) to obtain A_d as a function of N_d/N and substituting this value into equation (14) yields

$$E[s \mid \pi_d(j) > 0] = \left(\frac{N_d}{N} \right)^{-1/\kappa} \frac{\kappa}{\kappa - 1} \underline{s} \mu_2 \mu_1^{(1-\kappa)/\kappa}. \quad (15)$$

Equation (15) implies that the elasticity of expected quality with respect to the number of firms that export to the market is negative. Intuitively, if more firms make it in, the marginal entrants are worse, bringing down the average.

Prices are a power function of s (with parameter λ) given by equation (7). Derivations in the appendix reveal the expected price conditional on exporting to be

$$E[p \mid \pi_d(j) > 0] = \left(\frac{A_d}{\sigma F_d} \right)^{\frac{-\lambda}{\beta}} \omega^{\frac{\gamma}{(\gamma-\lambda)}} \frac{\sigma \kappa}{(\sigma - 1)(\kappa - \lambda)} (\mu_3 / \mu_1), \quad (16)$$

where μ_3 is another moment of the α distribution shown in the appendix. The elasticity with respect to A_d is negative under quality sorting since $\lambda > 0$ and $\beta > 0$. Under productivity sorting, attractive destinations have higher expected prices since $-\lambda/\beta = 1/(\sigma - 1) > 0$, when $\gamma = 0$ and $\lambda = -1$.

As with the conditional expectation of quality, we can express the conditional expectation of price in terms of the fraction of firms that enter the market:

$$E[p \mid \pi_d(j) > 0] = \left(\frac{N_d}{N} \right)^{-\lambda/\kappa} \frac{\omega \sigma \kappa}{(\sigma - 1)(\kappa - \lambda)} \underline{s}^\lambda \mu_3 \mu_1^{\lambda/\kappa - 1}. \quad (17)$$

As with expected quality, the expected price conditional on exporting is decreasing in the fraction of firms that export to market d for the quality-sorting model ($\lambda > 0$). However, the prediction is opposite for the productivity-sorting model (where $\lambda = -1$). When quality sorting takes place, only high quality varieties are exported to difficult

⁹Each conditional expectation depends on a different moment of the α distribution. All are shown explicitly in the Appendix.

countries, and those are high price varieties, because high quality is associated with high costs. When productivity sorting drives firms' selection into export markets, only the most productive firms with low marginal costs make it to difficult markets, and—with a constant markup—the selected firms charge low prices.

The model also makes predictions about the expected quantity shipped by a firm to a given market. Firm-level quantity ($q_d(j)$) is obtained by dividing (6) by the FOB price equation, (7), yielding

$$q_d(j) = \frac{\sigma - 1}{\sigma} \omega^{-\sigma} s(j)^{(\sigma-1)\gamma - \lambda\sigma} A_d \alpha_d(j). \quad (18)$$

A derivation shown in the appendix yields the expected quantity conditional on being a profitable exporter to a market as

$$E[q \mid \pi_d(j) > 0] = \left(\frac{A_d}{\sigma} \right)^{\frac{\lambda}{\beta}} F_d^{\frac{\beta-\lambda}{\beta}} \omega^{\frac{-\gamma}{(\gamma-\lambda)}} \frac{\kappa(\sigma - 1)}{\kappa - (\sigma - 1)\gamma + \sigma\lambda} (\mu_4/\mu_1), \quad (19)$$

where μ_4 is another moment of the α distribution. The power on A_d is positive as long as quality is costly and “worthwhile,” i.e. $\gamma > \lambda > 0$. Under the parameterization corresponding to productivity sorting ($\gamma = 0$, $\lambda = -1$), the power on A_d is negative since $\lambda/\beta = -1/(\sigma - 1) < 0$. Under productivity sorting, expected quantity should be increasing in fixed costs, with an elasticity, $\sigma/(\sigma - 1)$, equal to the mark-up factor. The sign is positive under quality sorting if and only if γ/λ exceeds $\sigma/(\sigma - 1)$.

As with quality and price, one can obtain an expression for the conditional expectation of quantity as a function of the probability that a firm exports to that market:

$$E[q \mid \pi_d(j) > 0] = \left(\frac{N_d}{N} \right)^{\lambda/\kappa} \frac{F_d}{\omega} \frac{\underline{s}^{-\lambda} \kappa (\sigma - 1)}{\kappa - (\sigma - 1)\gamma + \sigma\lambda} \mu_4 \mu_1^{-\lambda/\kappa - 1}. \quad (20)$$

Average quantity exported to d is a positive function of the fraction of firms that enter that market in the quality-sorting model. As with price, the sign of the relationship is reversed in the productivity-sorting model. The drawback of this expression compared to those obtained for quality and price is that it is not a simple bivariate relationship between observables. Indeed fixed costs are expected to enter with a unit elasticity.

2.5 Country mean predictions

We now show how to transform the relationships between conditional expectations and model parameters into relationships that can be estimated using observables. We can estimate the expected value of quality, prices, and quantities using the observed average level of these variables for exporters to a given destination. Thus for firm-level variable v —which can represent s , p , or q —we have

$$\tilde{v}_d \equiv (1/N_d) \sum_{j \in \mathcal{H}_d} v_d(j), \text{ where } \mathcal{H}_d \text{ is the set of French exporters to } d. \quad (21)$$

There are three types of empirical relationships between the country means— \tilde{s}_d , \tilde{p}_d , and \tilde{q}_d —and country characteristics that we can estimate as a way of establishing the relevance of quality sorting and backing out implied model parameters.

The relationship requiring the least data is the *mean-popularity regression*, obtained by taking logs of equations (15), (17), and (20):

$$\ln \tilde{v}_d = \text{constant} + \eta_{vN} \ln N_d + \text{error}_d.$$

The error term in this and subsequent country-mean specifications is statistical (unlike the structural error term in the firm-level regressions discussed in the next subsection). It captures the fact that in finite samples means do not equal expected values.

Table 1 provides the predicted values for η_{vN} in the first and fourth rows. This specification omits F_d which is justified for quality and price. In the case of quantity we will control for F_d , expecting an elasticity of one. Under quality sorting and the Melitzian productivity-sorting parameterization, $\eta_{pN} = -\eta_{qN}$. However, the signs flip based on the type of sorting: η_{pN} is negative under quality-sorting and positive under productivity-sorting.

Table 1: Predicted elasticities in two sorting models

Sorting model:	Explanatory variable:	Dependent variable:			
		N_d	\tilde{s}_d	\tilde{p}_d	\tilde{q}_d
$s = \mathbf{quality}$ ($\gamma > \lambda \geq 0$)	“popularity” (N_d)		$\frac{-1}{\kappa}$	$\frac{-\lambda}{\kappa}$	$\frac{\lambda}{\kappa}$
	“attractiveness” (\hat{A}_d)	$\frac{\kappa}{\beta}$	$\frac{-1}{\beta}$	$\frac{-\lambda}{\beta}$	$\frac{\lambda}{\beta}$
	“entry threshold” ($\widehat{F}_d\sigma$)	$\frac{-\kappa}{\beta}$	$\frac{1}{\beta}$	$\frac{\lambda}{\beta}$	$\frac{\beta-\lambda}{\beta}$
$s = \mathbf{productivity}$ ($\gamma = 0, \lambda = -1$)	“popularity” (N_d)		$\frac{-1}{\kappa}$	$\frac{1}{\kappa}$	$\frac{-1}{\kappa}$
	“attractiveness” (\hat{A}_d)	$\frac{\kappa}{\sigma-1}$	$\frac{-1}{\sigma-1}$	$\frac{1}{\sigma-1}$	$\frac{-1}{\sigma-1}$
	“entry threshold” ($\widehat{F}_d\sigma$)	$\frac{-\kappa}{\sigma-1}$	$\frac{1}{\sigma-1}$	$\frac{-1}{\sigma-1}$	$\frac{\sigma}{\sigma-1}$

Note: $\beta \equiv (\sigma - 1)(\gamma - \lambda) > 0$ is the elasticity of firm-level exports ($x_d(j)$) with respect to quality $s(j)$.

The second exercise is the *mean-attractiveness regression*, based on taking logs of equations (14), (16), and (19):

$$\ln \tilde{v}_d = \text{constant} + \eta_{vA} \ln \hat{A}_d + \eta_{vF} \ln \widehat{F}_d\sigma + \text{error}.$$

The model's predicted elasticities for each variable (s , p , and q) with respect to A_d and $F_d\sigma$ are shown in rows 2 and 3 for quality sorting and rows 5 and 6 for productivity sorting. Estimation of this relationship can provide estimates of β , the elasticity of firm export values with respect to quality, and λ , the elasticity of cost with respect to quality. These can be compared with more direct estimates from the firm-level regressions described in the next section. One can also discriminate between quality- and productivity-sorting based on the sign pattern of the η_{vA} and η_{vF} coefficients for price and quantity.

A third exercise replaces $\ln \hat{A}_d$ with its determinants, which are just the standard set of explanatory variables in a gravity equation. We therefore refer to this approach as the mean-gravity regressions.

$$\ln \tilde{v}_d = \text{constant} + \boldsymbol{\rho}_v \mathbf{R}_d + \text{error},$$

where $\mathbf{R}_d = [\ln b_d \quad \ln y_d \quad \ln M_d \quad \ln \text{Dist}_j \quad \text{French}_d \quad \ln \text{Prod}_d]$. Of these variables, the first three follow directly from the theory. Distance and speaking a common language are standard proxies for trade costs (τ_d). The variable Prod_d measures production of wine in country d . The idea is that for a given amount of consumption ($b_d y_d M_d$), more domestic production in d tends to crowd out imports from France. In the model this occurs through a reduction in P_d , the price index. This specification is the one that is most closely related to exercises conducted in other papers and we therefore show results for it prior to the other two which we believe to be original to this paper.

The relationship between the reduced-form coefficients ($\boldsymbol{\rho}_v$) and the structural parameters depends on what assumption we make regarding F_d , the fixed costs. If they are constant or orthogonal to the determinants of A_d , then $\boldsymbol{\rho}_v$ is given by the product of η_{vA} and the derivative of $\ln A_d$ with respect to each column i of \mathbf{R}_d . Thus, we can use the estimated $\boldsymbol{\rho}_v$ to discriminate between models since the signs of $d \ln A_d / d \mathbf{R}_d^i$ are known. Next, consider the more general case where the fixed costs depend on at least some of the determinants of A_d :

$$\boldsymbol{\rho}_v^i = \eta_{vA} \frac{d \ln A_d}{d \mathbf{R}_d^i} + \eta_{vF} \frac{d \ln F_d}{d \mathbf{R}_d^i}.$$

Since theory gives little guidance as to the determinants of fixed costs, this general case renders the estimates of $\boldsymbol{\rho}$ of little use for discriminating between models. We instead consider a more restrictive case where we can make predictions. Suppose fixed costs are a power function of the attractiveness index: $F_d \propto A_d^\phi$, where ϕ can be positive or negative. In that case, the coefficients on gravity variables correspond to

$$\boldsymbol{\rho}_v^i = (\eta_{vA} + \phi \eta_{vF}) \frac{d \ln A_d}{d \mathbf{R}_d^i}.$$

In the case of $v = s$ and $v = p$, we have $\eta_{vA} = -\eta_{vF}$. This implies

$$\boldsymbol{\rho}_v^i = \eta_{vA} (1 - \phi) \frac{d \ln A_d}{d \mathbf{R}_d^i}.$$

For $\phi < 1$, the sign of $\boldsymbol{\rho}_v^i$ provides the sign of η_{vs} and η_{vp} .

There is one additional technical point to be noted in the estimation of those conditional mean regressions. Because of the very nature of our model where firms self-select in exporting to different destination countries, markets d have a different population size when computing the averages. This generates heteroskedasticity since the variance of the error term will be inversely proportionate to N_d , the number of exporters constituting the mean. We therefore weight by the number of exporters in generalized least squares (GLS) regressions.

Table 1 also shows, in its first column, the predicted elasticities derived from equation (12), where the probability of entry is replaced with N_d , the actual number of entrants. Comparing the elasticities shown in the first two columns of Table 1, we see $\eta_{NA} = -\kappa\eta_{sA}$. This implies that any variable that increases N_d should have the opposite effect on \tilde{s}_d since κ , the Pareto shape parameter, is strictly greater than zero.

Yeaple (2009) is the first paper we know of to use this implication of heterogeneous firms models to test for productivity sorting. Yeaple regresses the number of American multinationals with affiliates in country d on a set of country-level profitability determinants (which closely resembles our \mathbf{R}_d). He compares the estimates to the regression of the media sales of those firms in the US on the same set of profit determinants. For every country characteristic except distance, he finds the predicted pattern of opposite signs. Since Yeaple’s proxy for productivity, home sales, could also be a proxy for quality in our model (size is increasing in s as long as $\gamma > \lambda$), we view Yeaple’s regressions as a test for “ability sorting” that does not discriminate between sorting by physical productivity and sorting by product quality.

2.6 Firm-level predictions

We can estimate the model using firm-level data for three different dependent variables: the probability of exporting, the FOB price, and exported value. Taking logs of equations (8), (7), and (6) we obtain the estimating equations.

The probability of exporting is given by

$$\Pr[x_d(j) > 0] = \Pr[\pi_d(j) > 0] = \Pr[\ln x_d(j) - \ln \sigma - \ln F_d > 0].$$

Using equation (6) to express $x_d(j)$ in terms of its determinants, firm j will export to d with probability

$$\Pr[x_d(j) > 0] = \Pr[\beta \ln s(j) - (\sigma - 1) \ln \omega - \ln(F_d \sigma) + \ln A_d + \ln \alpha_d(j) > 0]. \quad (22)$$

The parameters can be estimated using a binary choice model whose form depends on the assumption made on the distribution of the unobserved heterogeneity term $\alpha_d(j)$. Assuming log-normality for $\alpha_d(j)$ implies a probit form. This is the assumption made in Helpman, Melitz, and Rubinstein (2008). The logged attractiveness of country d and its fixed export costs, $\ln A_d$ and $\ln(F_d \sigma)$ appear on the right hand side of the export probability. Rather than attempt to estimate these terms as a parametric function of country d primitives, we absorb them with country-year-specific fixed effects.

From (7), the price charged by firm j takes the following estimable form:

$$\ln p_d(j) = \lambda \ln s(j) + \ln \omega + \ln[\sigma/(\sigma - 1)]. \quad (23)$$

The price equation in this model lacks an error term. However, for parallelism with the export probability and value equations, we add a normally distributed (in log scale) error with country-year specific fixed effects. One possible interpretation of the fixed effects are that σ varies across markets and the fixed effects estimate $\ln[\sigma_d/(\sigma_d - 1)]$. However, this is unattractive because it implies country-specific elasticities with respect to quality in the export value equation.

From (6), the log of firm-level exports (for firms that export positive values) is

$$\ln x_d(j) = \beta \ln s(j) - (\sigma - 1) \ln \omega + \ln A_d + \ln \alpha_d(j). \quad (24)$$

Country-year-specific fixed effects will capture the $\ln A_d$. We then collect those fixed effects and re-use them later in the country-mean regressions.

Assuming log-normal $\alpha_d(j)$ implies that OLS would be the maximum likelihood estimator for equation (24)—*if* we observed positive exports to all markets. In fact most Champagne exporters have positive exports to only a small number of destinations. This zero problem is predicted by the model unless fixed costs of exporting are negligible. The zero problem implies that OLS (with the dependent variable $\ln x_d(j)$ set as missing for $x_d(j) = 0$) would yield inconsistent estimates of the quality effect on exports.

Inspecting equation (22) reveals a negative relationship between the quality observed among exporters, $s(j)$, and the unobserved idiosyncratic shock that firms experience when considering exports to each market, $\alpha_d(j)$. For firms with identical quality levels, the probability of passing the cutoff and exporting increases with $\alpha_d(j)$. It follows that firms with high observed quality will become exporters even with relatively low draws of $\alpha_d(j)$, while low quality firms need high draws of $\alpha_d(j)$ to be observed as positive exporters. This negative correlation will tend to bias estimates of the effect of $s(j)$ on exports toward zero, since low quality firms will tend to do better than expected.

Helpman et al. (2008) use a Heckman correction to address the sample selection issue.¹⁰ The Heckman approach can be identified off functional form but it is generally recognized that we can only have confidence in the results if we have a variable explaining the firm-level decision to export to individual markets that is excludable from equation (24). According to the theory, it should be a variable that influences firm-level fixed costs. Helpman et al (2008) use overlap in religion in trade partners, as well as measures of entry costs based on World Bank data. They make this data dyadic by interacting indicators for the exporting and importing country. This will not work in our context because our country fixed effects are *de facto* dyadic fixed effects given that all our exports originate in only one region. We would therefore need an additional firm-level dimension here. The problem is that it is very difficult to conceive of a variable that

¹⁰For analysis of aggregated trade flows, Helpman et al. (2008) show that an additional, non-linear, correction term is needed to account for the fact that “a larger fraction of firms export to more ‘attractive’ export destinations.”

would affect one firm’s country-level fixed costs but not affect its variable costs of trade or its individual demand shock.

We pursue an alternative method that adheres closely to theory. Firm j exports to destination d if and only if $\pi_d(j) > 0$. From (4) it can be seen that exports are profitable if and only if $x_d(j) > F_d\sigma$. If equation (6) predicts $x_d(j) < F_d\sigma$ then $\pi_d(j) < 0$ and we would observe $x_d(j) = 0$. Thus we can define $\hat{x}_d = F_d\sigma$ as the minimum observed value of $x_d(j)$, that is the zero-profit export level for destination d . Assuming a log-normal distribution for $\alpha_d(j)$, we have a Tobit structure. The problem is that we do not observe $F_d\sigma$. Fortunately, Eaton and Kortum (2001) suggest that a maximum likelihood estimate of the censoring point, \hat{x}_d , can be obtained from the minimum observed positive value of $x_d(j)$. Thus we set

$$\ln \widehat{F_d\sigma} = \ln(\min_{j \in \mathcal{B}_d} x_d(j)). \quad (25)$$

For $\ln x_d(j) > \ln \widehat{F_d\sigma}$ the likelihood is based on the continuous $\ln x_d(j)$ from equation (24). For $\ln x_d(j) < \ln \widehat{F_d\sigma}$ the likelihood is the probability that $\ln x_d(j) \leq \ln \widehat{F_d\sigma}$.

To assess the reliability of this estimation approach, under the assumptions of our model, we conducted Monte Carlo simulations using equations (4) and (6) (profits and export value, which are the core of our selection problem). Since we use estimated coefficients to parameterize the simulation, we present the simulation results together with the regression results in section 4.1.

3 Data

Our paper combines two main sources of data, firm-level export declarations and books on Champagne producer quality. We start by discussing features of Champagne exporting that appear to conform to the main elements of our model. Then we describe the sources and construction of the data set.

3.1 Why Champagne?

Champagne is an attractive industry for an empirical application of our theoretical predictions. Most importantly, we have good *ex ante* reasons to believe that Champagne producers are vertically differentiated in terms of the quality of their products. Second, the firms that handle exports, and hence are listed on the customs declarations we rely upon, are predominately producers to whom we can assign quality ratings. Third, experts on Champagne have identified a variety of mechanisms that support the Baldwin-Harrigan assumption linking higher quality to higher marginal costs. Finally, Champagne appears to exhibit Armington-style differentiation by place of origin, a key implicit assumption of the model. We discuss each advantage in turn below.

Champagne fits the assumption of firm-level differentiation very well. Geographic distinctions *within* the Champagne region (a single *appellation*) are not emphasized.

“[The E]ssence of Champagne is that it is a blended wine, known in all but a handful of cases by the name of the maker, not the vineyard.” (Johnson and Robinson, 2005)

While Champagne is a single *appellation*, other French wine regions are extensively subdivided, with each *appellation* purported to have distinct taste properties. Since the export product classification stops at the level of regions, it would appear less suited to capture quality variation in regions like Burgundy that put the primary emphasis on detailed geography, rather than firms.

Cost-quality trade-off exist in both grape-growing and Champagne-making. The quality of land has been built into the price of grapes in Champagne through a system called *échelle des crus*, with grapes from vineyards with better reputations commanding higher prices. Thus if we think of $w(s)$ as the factor costs embodied in wine of quality s , we have good reasons to expect $w' > 0$. There is also a productivity trade-off in viticulture since “over-cropping” (more grapes per hectare) is believed to undermine the intensity of the flavors. For any given set of grapes, the making of Champagne also exhibits cost-quality tradeoffs. The longer the time the wine spends on its lees, prior to the disgorgement of the yeast deposit, the more complexity it tends to acquire. Furthermore, the Champagne maker can choose more or less costly liquids to add when the yeast is removed. Depending on this “dosage,” the Champagne may become excessively sweet.

A critical practical consideration is that the major producers of Champagne are also the firms that handle most of the export value of the industry. Customs data lists exports by a firm for each cn8 product. In other firm-level sources of data, the same firms are classified according to a “primary” activity. In other wine regions a large proportion of the firms named on export declarations do not correspond to the producers rated in the wine guides. Some of those “non-producing” exporters are dealers who mainly label and distribute wine made by other firms (as is the case for Bordeaux). Other firms are mainly dealers, but are also vertically integrated backwards into grape growing and even wine making (as is the case for Burgundy).

Table 2 provides the share of exports (and the share of exporters between parentheses) according to primary activity in 2003. For Champagne, the growers and makers add up to 78% of the total, only 35% for white Burgundy, and a very small 18% for white Bordeaux. The picture is even worse for red Bordeaux where most exports are channelled through wholesalers (85%). Champagne is the notable exception to the rule that wholesalers dominate the export business since they account for only 7.2% of Champagne exports in 2003. The problem with exports by dealers is that it is hard to assess the exact wine and therefore the quality exported by the firm. For Bordeaux it is infeasible to obtain exporter quality measures because most of the major dealers are omitted from the guidebooks (presumably because they do not make any wine). For Burgundy, some of the main exporting wholesalers are vertically integrated and we are able to find quality ratings for them in the guides. This is why we also investigated exports of red Burgundy as a robustness check. Note, however, that the large proportion of wholesalers in total export value is likely to add a substantial amount of noise in the measurement of firm-level

quality in our Burgundy regressions.¹¹

Table 2: Who exports wine?

Export shares, % (Shares of exporting firms, %), 2003						
White wines						
Activity	Code	Champ.	Bord.	Burg.	Loire	Alsace
grape-growing	011G	2.4 (30.9)	10.4 (27)	16 (36.6)	36.7 (38)	10.5 (33.4)
wine-making	159G/F	75.2 (13.1)	7.1 (2.2)	18.6 (3.9)	7.3 (4.1)	33.6 (7.8)
wholesale	513J	7.2 (27.8)	73.2 (42)	60.3 (38.2)	43.9 (32.8)	53.5 (34.8)
other	-	15.3 (28.2)	9.4 (28.9)	5.1 (21.2)	12.2 (25.1)	2.3 (24)
Red wines						
Activity	Code	Rhone	Bord.	Burg.	Loire	Beauj.
grape-growing	011G	7.6 (29.2)	6.7 (41.5)	23.9 (39.9)	23.9 (38.3)	2.5 (25.7)
wine-making	159G/F	12.7 (7)	2.4 (2.9)	4.8 (3.3)	6.4 (5.1)	7.3 (4.9)
wholesale	513J	60.5 (38.3)	84.5 (28.6)	62.2 (35.4)	56.2 (31.9)	84.5 (40)
other	-	19.1 (25.5)	6.3 (27.1)	9.2 (21.4)	13.5 (24.7)	5.8 (29.5)

In contrast to regions like Bordeaux and Burgundy where the vintage of the wine is thought to have a decisive influence on quality, the quality of Champagne is considered relatively stable over time. This is because most Champagne producers blend several years of grapes to reproduce a consistent quality over time. Since we observe yearly exports by a firm, but not the precise mix of vintages of the bottles exported, this reduction in inter-temporal quality variance is helpful.

The geographic definition of the Champagne industry makes it particularly appropriate for studying the effect of heterogeneity on the composition of exporters by destination. The relevance of differentiation by place of origin for this study is that the Melitz (2003) model, upon which we base our analysis, assumes that firms face only the option of exporting or not to a given market. Firms cannot relocate production to the consuming market as they can in the Helpman et al. (2004) framework. With footloose production, the implications for quality sorting could be quite different. In particular, the best firms might conduct FDI in the difficult markets, rather than serving them via exports.

Claims of Champagne producer associations and wine critics both support the assumption of Armington-style differentiation by *place* of origin. The Champagne industry has used full-page advertisements and legal actions to reinforce that belief. To qualify as Champagne in the EU, a sparkling wine must be produced within the Champagne geographic *appellation*. The following quote comes from the official Champagne promotion agency.

¹¹Results available at <http://strategy.sauder.ubc.ca/head/sup/>.

“The important thing to remember is that while some processes of Champagne production may be duplicated, the terroir is unique, original, and impossible to replicate.” (www.champagne.us)

Some wine critics agree with the proposition that sparkling wine from Champagne is distinct:

“The Champagne region has certain natural advantages that no amount of money, ambition, or talent can surmount: The combination of chalky soil and fickle northern European weather yields sparkling wines that simply can’t be replicated anywhere else...” (Steinberger, 2005)

We proceed as follows: the body of the paper concentrates on Champagne (cn8 22041011) for the reasons discussed above. However, we also carry out the analysis on red Burgundy (cn8 22042143), notably because it is a type of wine where far more firms are listed in wine guides. While the results for red Burgundy tend to be weaker (as expected) than those we obtain for Champagne, our main results hold up for both wines.

3.2 Trade data

We use the micro-data collected each year based on export declarations submitted to French Customs. It is an almost comprehensive database which reports annual shipments by destination at the 8-digit product level for each French exporting firm. The “almost” is due to EU legislation following the implementation of the single market, which sets different thresholds for compulsory declarations inside and outside the customs union. All exports within the European Union must be declared. Exports outside the EU must be declared unless the total value to a destination country d is smaller than 1000 euros or 1000 kilograms. The average unit value in our sample is slightly higher than 20 euros per kilogram, which can be reasonably taken as the average price of an exported bottle. The declaration threshold is therefore around 50 bottles per destination country. We find very few cases of exports outside the EU that are close to the reporting threshold. Averaged over the 1998–2003 period, the minimum value exported in our sample is 850 euros for Switzerland, and 958 for Canada, the only two non-EU countries under the threshold. The average of the minimum observed values for countries outside the EU, 8400 euros, is more than eight times the declaration threshold.

For each firm, Customs records FOB values and quantities exported to 216 countries. Our extraction from this data spans the six years from 1998 to 2003. We calculate firm-destination-level FOB prices (often referred to as “unit values”) as $p_d(j) = x_d(j)/q_d(j)$.

Customs utilizes 11,578 8-digit combined nomenclature product classifications (abbreviated as “cn8”). The cn8 is the harmonized system 6-digit (hs6) code with a 2-digit suffix that is particular to the European Union. Wine has an hs4 of 2204. Sparkling wine is 220410. For our purposes, it is fortunate that the last two digits of the cn8 distinguish important wine-growing regions in the EU. Thus Champagne, the sparkling wines from the official Champagne region, receive their own cn8 (22041011).

Champagne accounts for 0.45% of French exports. This might not seem large, but is impressive when compared to other goods. The mean good-level contribution to total trade is less than 0.01% and the largest exporting industry at this level of disaggregation (aeroplanes and other aircraft exceeding 15 tons) accounts for only 3.24%. Champagne is clearly among the largest contributors to French exports, and is also a strong outlier in other dimensions. When ranking *cn8* products according to the number of exporting firms, Champagne ranks 21st out of 11,578 products. Its importance is even more striking in terms of the number of destination countries. As Figure 1 shows, this industry exports to a much larger number of countries than the typical French industry. The actual rank is 7th, with an average of 171 countries served in our sample years (the top industry serves an average of 179 countries).¹²




Figure 1: Champagne is an outlier in the distribution of destinations per product

The export declaration data provides us with firm identification numbers, or SIREN, for all 12,314 firms who exported any form of wine (*hs4* = 2204) between 1998 and 2003. Of those, the French national statistical agency (INSEE) provides the names, addresses, and primary activity code for the 10,341 firms in existence as of June 2007. We used the firm-level name and address information to match exporters with wine producers that were rated in two guidebooks.

¹²The good exported to the most destinations is perfume, another industry where quality differentiation is considered important.

3.3 Quality ratings

Wine producer quality ratings come from two different sources: i) a French one: Burtschy, Bernard and Antoine Gerbelle, 2006, *Classement des meilleurs vins de France*, Revue Des Vins De France (Paris), which we refer to as RVF, ii) an internationally recognized one: Parker, Robert, *Wine Buyer's Guide*, 5th Edition, 1999, which we refer to as WBG. For each of the listed producers, the name and location were matched with the exporter's dataset by hand.

In RVF, listed producers receive between 0 and 3 stars. We have 64 Champagne producers listed, and are able to match those with 51 exporters. In WBG, 70 Champagne producers are categorized as “average,” “good,” “excellent,” or “outstanding.” Of those we find 47 Champagne exporters.¹³

Table 3 evaluates how closely those two quality ratings match for Champagne. Kendall's τ index of concordance between ratings (given in the footnote) suggests that while those two ratings are certainly not independent, they are not identical either. We will explain how we exploit those differences later.

Table 3: Champagne quality ratings

Parker's WBG	RVF's <i>Classement</i>					Total
	n/a	Incl.	*	**	***	
n/a	1724	16	6	0	0	1746
Average	3	1	0	0	0	4
Good	7	3	1	2	0	13
Excellent	7	6	4	3	0	20
Outstanding	1	0	3	4	2	10
Total	1742	26	14	9	2	1793

Note: Kendall's τ measure of concordance $-1 \leq \tau \leq 1$ (p -value for test for independence) is 0.58 (p -value of 0.000) for all exporters and 0.43 (p -value of 0.009) for those included in both books.

One difficulty raised by using guidebooks to measure producer quality is that producers deemed to make low quality wine are usually omitted. With a vital caveat, exclusion from the books can be interpreted as a bad signal. We cannot infer that *all* exporters omitted from the guides are low quality because substantial amounts of Champagne are exported by non-producers. Intermediary exporters who do not make Champagne would normally not enter the guides even if they exported exclusively high-quality Champagne. We therefore omit from most of our analysis all firms that ship wine abroad but for which we have no basis to infer quality of the wine exported. The main challenge is to define a reasonably homogenous low quality category. To do this, we use the information contained in Table 4.

¹³For comparison purposes, we conducted the same exercise for red burgundy and found 268 listings in RVF and 159 in WBG, of which 206 and 139, respectively, can be found in the customs dataset.

Table 4: Champagne exporters by primary activity, location, and guidebook inclusion (1998–2003)

Primary Activity	Included in Guide?		
	No		Yes
	Local?		
	No	Yes	Yes
grape-growing	40	392	22
	76	3972	1015
Champagne-making	10	84	33
	50	2914	5846
wholesale	346	94	10
	1624	2229	769
other	678	101	4
	2923	1166	824

Note: In each cell the first row provides the number of firms and the second row gives the number of export observations (firm-destination-years).

Table 4 shows the number of firms and counts of firm-destination-year export observations broken down according to whether the exporter was included in guide, its location, and its primary activity. We classify firms as low quality $s(j) = 1$ if they are (1) unrated by either guide, (2) located within the official Champagne-growing départements (“Local”)¹⁴, and (3) engaged in grape-growing or Champagne-making as their primary activity. These cases are shown in the gray-shaded cells of Table 4. Non-local firms as well as unrated firms with other primary activities will be referred to as having “mixed” quality.

Table 5 shows how we standardized measured quality, $s(j)$, to a range from 1 to 5 for each of the guides. Our firm-level regressions mostly average the standardized RVF and WBG ratings. In the conditional mean regressions, we calculate RVF and WBG country means separately and then average them.

Table 5: Standardized quality mapping

$s(j)$	RVF <i>Classement 2007</i>	WBG (Parker, 1999)
5	***	“Outstanding” (*****)
4	**	“Excellent” (**)**
3	*	“Good” (***)
2	Included in RVF	“Average” (**)
1 (low)		Included in <i>other</i> book OR Local grower/maker
N/A (mixed)		All other exporters

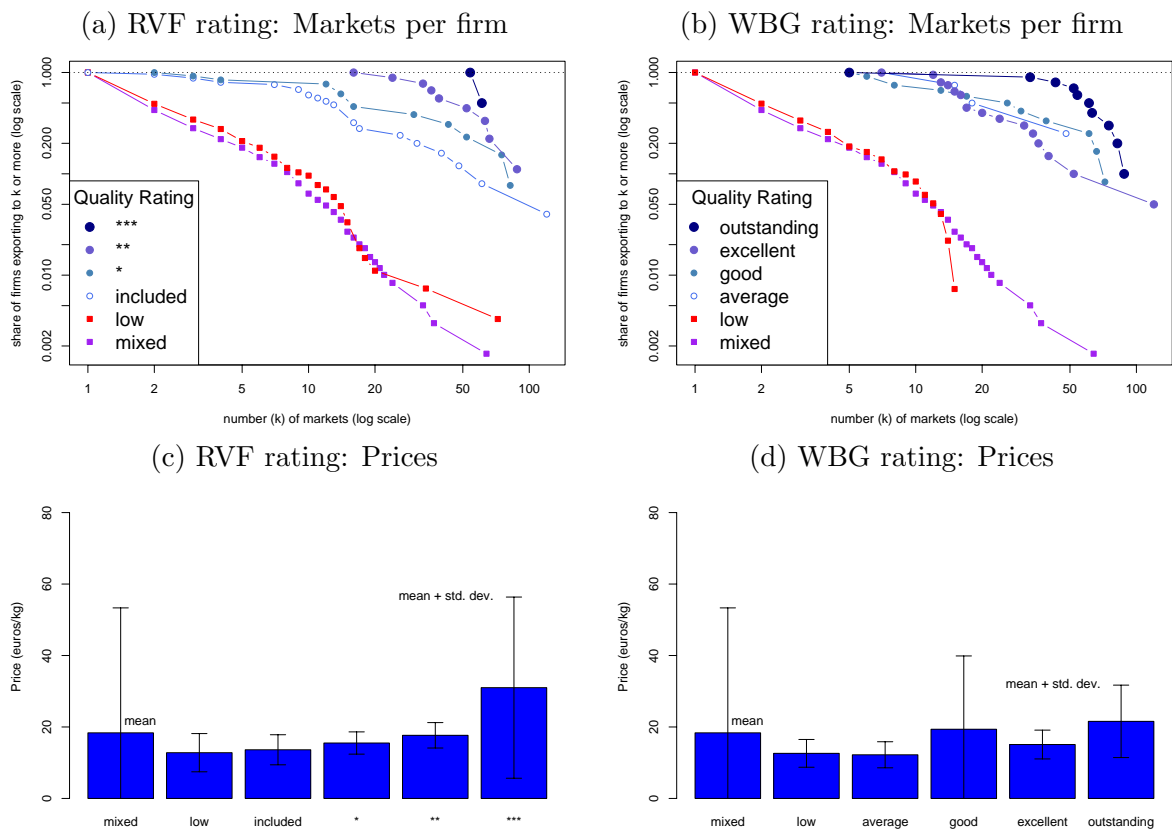
Figure 2 allows us to compare rated firms with the low and mixed quality firms in two model-relevant dimensions: the number of destinations to which they export and the average prices they charge across all destinations. For ratings based on both books, quality ratings from 2 to 5 are associated with larger numbers of export destinations than the Champagnes we classify as low quality ($s = 1$). The mixed quality producers also tend to export to low numbers of markets but there are a few outliers. In terms of price the “mixed” category very much deserves its name: standard deviations are very high. We also see that as quality increases that prices generally rise but there are some non-increasing steps using Parker’s (WBG) ratings.

There are several additional difficulties with using guidebook ratings as quality measures. We list them along with our responses below.

1. *The ratings are hard to interpret*: units of measurement (stars) do not correspond to prices or quantities. Our theory includes the parameter γ to capture the marginal

¹⁴Marne (67% of production), Aube (22%), Aisne, Haute-Marne, and Seine-et-Marne (<http://www.champagne.fr/fr/localisation.aspx>).

Figure 2: Champagne: Markets per firm and Prices (wt. avg.)



utility of quality units. This parametric approach also has the advantage of compactness in the presentation of the results .

2. *The ratings are unreliable*: authors may have idiosyncratic tastes or be influenced by non-taste considerations.¹⁵ In order to minimize this concern, we use two completely independent sets of ratings, for which we have no reason to suspect that author-specific “specificities” would be correlated.
3. *The ratings may influence demand by increasing foreign customer awareness*. For instance, consumers in New Zealand are probably not aware of all varieties of wine produced and available for consumption in France. A guide like Parker’s, because it is in English and widely available, could increase demand merely by increasing awareness (adding varieties to the consumers’ information sets). To eliminate this “advertising” effect, we run a separate set of regressions using only the French guide ratings (RVF) and restricting the sample to non-francophone markets (RVF is not translated).¹⁶

4 Results

We start by presenting firm-level regression results and the next subsection displays the conditional mean (by destination) results using figures and regressions.

4.1 Individual level analysis

Table 6 reports estimates from our firm-level regressions for price (column 1), export probability (column 2), and value exported (columns 3–6). The corresponding equations from the model are (23), (22), and (24). Column (3) uses a fixed effect linear estimator to assess the impact of quality on export value. This specification excludes zeros from the estimation sample. The next column uses the Tobit methodology described above for solving the selection issue. Columns (1)–(4) and (6) average the two quality ratings (WBG and RVF) to obtain $s(j)$. Column (5) attempts to neutralize the “promotional” role of guidebooks by restricting the sample to non-francophone countries and using the French-language guide (RVF) as the sole measure of quality. Finally, we test the Hallak hypothesis that higher incomes increase the demand for quality in column (6).

A first broad statement can be made about the influence of quality. Our estimates reveal that higher quality tends to raise export prices, export probability, and export value as predicted in the model. A second important point is that selection bias shrinks the coefficient on quality in the OLS export value regression shown in column (3). The bias arises because selection into exporting generates a negative correlation between the

¹⁵For example, Parker stopped including Faiveley wines in his guide after a lawsuit brought against him by the wine maker for insinuating that his exported wine was inferior to that served in France.

¹⁶We thank Andrew Bernard for pointing out this concern and for suggesting the solution we have implemented.

Table 6: Firm-level regressions for quality-rated Champagne exporters

	(1)	(2)	(3)	(4)	(5)	(6)
	$\ln p_{dt}(j)$	$x_{dt}(j) > 0$	$\ln x_{dt}(j)$	$\ln x_{dt}(j)$	$\ln x_{dt}(j)$	$\ln x_{dt}(j)$
$\ln s(j)$	0.29 ^a (0.01)	1.77 ^a (0.01)	2.09 ^a (0.06)	7.64 ^a (0.06)	7.95 ^a (0.08)	7.43 ^a (0.06)
$\ln s(j) \times \ln(y_{dt}/y_0)$						0.63 ^a (0.03)
Method	OLS	Probit	OLS	Tobit	Tobit	Tobit
Observations	12426	405189	12426	405189	317516	366749
Within R ² /Pseudo R ²	0.117	0.482	0.269	0.321	0.267	0.324
FE share of variance	0.38		0.30			

Note: Destination-year (dt) fixed effects for all columns. Column (5) restricts the sample to non-francophone countries and $s(j)$ is based on RVF guide only. $y_0 = \$6,800$ is the all-country average GDP per capita (1998–2003). Standard errors in parentheses. Significance levels: ^c $p < 0.1$, ^b $p < 0.05$, ^a $p < 0.01$

quality conditional on being selected and unobserved firm-country demand shocks. Comparing columns (3) and (4) confirms the direction and magnitude of this bias. Using the Tobit estimator multiplies the OLS coefficient on quality by 3.66.

Monte Carlo simulations of our model show that the OLS bias is of the expected order of magnitude. More importantly, they also show that our Tobit method results in an estimate very close to the true impact of quality on exports in the simulated population of firms. This gives us some confidence that the Tobit method successfully corrects for the selection bias described in section 2.6.

The simulation comprises 1000 firms and 10 countries.¹⁷ The first step is to generate a random set of $s_d(j)$, A_d , and $\alpha_d(j)$, with which to create the uncensored vector of $x_d(j)$ based on equation (6). We specify the “true β ,” as 7.64, the estimate from column (4) of Table 6. Since the simulation draws log-normally distributed $\alpha_d(j)$, we expect the regression of uncensored $\ln x_d(j)$ on $\ln s(j)$ to yield a consistent estimate of β . The simulation is repeated with 10,000 different draws on the error term $\ln \alpha_d(j)$ and the results summarized in Table 7. The mean $\hat{\beta}$ is correct out to the level of precision with which we specify the true value.

The censored sample is obtained by imposing the condition that gross profits exceed fixed costs, which holds when $x_d(j) > F_d\sigma$. We choose the parameters of the A_d and F_d distributions such that the share of firm/destination profitable combinations, 3%, replicates the share we observe in our empirical sample.¹⁸ The censoring condition requires an estimate of σ and we use $\sigma = 7$. We then regress $\ln x_d(j)$ on $\ln s(j)$ in the censored sample, which therefore removes 97% of the original set of $x_d(j)$ which have been determined to be unprofitable. This corresponds to the OLS regression shown in column (3) of Table 6. The average $\hat{\beta}$ over the 10,000 simulations is 1.236, although there is consid-

¹⁷Stata code provided at <http://strategy.sauder.ubc.ca/head/sup/>.

¹⁸ Table 6 shows that over the 405,189 possible combinations, only 12,426 are positive.

erable variation (standard deviation of 0.438) across runs. The cause of the downward bias is revealed in the -0.526 average correlation between $\ln s$ and $\ln \alpha$ in the censored sample.

Table 7: Simulation results (assumed true $\beta = 7.64$)

Variable	mean	std. dev.
OLS β before censoring $x_d(j)$	7.639	0.046
Share of profitable firm-destination exports	0.030	0.002
OLS β ($x_d(j) < \sigma F_d$ censored)	1.236	0.438
Correlation($\ln s$, $\ln \alpha$) in censored data	-0.526	0.047
Tobit 1 β (estimate σF_d with $\min x_d(j) > 0$)	7.449	0.578
Tobit 2 β (known σF_d)	7.664	0.591
Magnification: Tobit 1 / OLS	6.966	39.86

To correct for the bias in OLS on censored data, we consider two Tobit regressions, reported in the fifth and sixth rows of Table 7. Tobit 1 is the method used in our econometrics. It estimates the censoring point using the minimum observed trade value: $\widehat{F_d\sigma} = \min_{j \in \mathcal{B}_d} x_d(j)$. Although Tobit 1 is biased downwards ($7.449 < 7.64$), it corrects 97% of the bias found in the OLS.¹⁹ We compare this performance to that of Tobit 2, where we use the censoring value $\hat{x}_d = F_d\sigma$ that was used to generate the simulated data. Tobit 2 obtains an estimate that is very close to the true value but we see that the massive amount of censoring we have incorporated in these simulations leads to considerable imprecision in both sets of Tobit estimates. The final row gives the magnification of the OLS result that the simulation predicts for our Tobit method. The average ratio of the coefficients is almost seven but the magnification varies a huge amount. The magnification ratio in the real data was 3.7, which is lower than expected, but of the right order of magnitude. All in all, the simulations make us confident that our Tobit method does a good job of correcting an otherwise important bias. Since the selection issue arises in any regression of firm-destination-level exports on firm ability measures, we think that the Tobit 1 method may prove useful in other studies.

Returning to our econometric results, we can use the structure of our model to reveal estimated values of the model's structural parameters and thereby obtain a precise quantification of the quality effects. Recall that equation (24) defines the elasticity of quantity with respect to quality as $\beta \equiv (\gamma - \lambda)(\sigma - 1)$. Rearranging, the implied value of γ is $\hat{\lambda} + \hat{\beta}/(\sigma - 1)$. Parameter λ can be obtained as the coefficient on log quality in the price regression, 0.29. Anderson and van Wincoop (2004) report $5 \leq \sigma \leq 10$ as a reasonable range for the CES. Plugging in estimates obtained for the full sample, we infer γ to lie between 1.14 and 2.2. A consumer is willing to trade between 6 and 34 bottles of low quality ($s = 1$) wine for one bottle of the highest quality ($s = 5$). This range is also the

¹⁹The correction share is the ratio of the difference between the means of the Tobit 1 and OLS estimators and the difference between the true value and mean OLS.

one for the ratio of prices between a five-star and a one-star bottle that would leave a consumer's indirect utility unchanged.²⁰

The estimates in column (6) test the Hallak (2006) hypothesis that the preference for quality parameter depends on income: $\gamma_d = \gamma_0 + \gamma_1 \ln(y_d/y_0)$. This formulation normalizes the income per capita of country d by the average world income (y_0) so that γ_0 is the preference parameter for the average country. With this specification of the preference for quality, the export equation becomes

$$\ln x_d(j) = \beta_0 \ln s(j) + \beta_1 \ln s(j) \ln(y_d/y_0) - (\sigma - 1) \ln \omega + \ln A_d + \ln \alpha_d(j), \quad (26)$$

where $\beta_0 \equiv (\gamma_0 - \lambda)(\sigma - 1)$ and $\beta_1 \equiv \gamma_1(\sigma - 1)$. With estimates of $\hat{\lambda} = 0.29$ from the price equation and $\sigma = 7$ from the literature, one can calculate both γ_0 and γ_1 . The interaction term coefficient in column (6) implies $\gamma_1 = \hat{\beta}_1/(\sigma - 1) = 0.63/6 = 0.105$.²¹ The coefficient on $\ln s$ reveals $\gamma_0 = (7.43/6) + 0.29 = 1.53$, which is the preference for quality parameter for a country with an average income per capita ($y_0 = \$6,800$). For the United States in 2003, the preference for quality is $1.53 + 0.105 \times \ln(37658/6800) = 1.71$. Even the poorest importer in our sample (Burundi in 2003) is estimated to have a γ exceeding one: $1.53 + 0.105 \times \ln(85/6800) = 1.07$.

4.2 Country mean regressions

By examining how conditional means of quality, prices, and quantities vary across markets we now test for evidence of quality sorting in Champagne. We start by conducting estimations that follow the prior literature in regressing log means on a set of *gravity* right-hand-side (RHS) variables. Second, we estimate what we see as the preferred relationship between country means and our estimates of destination *attractiveness* (A_d) and *entry thresholds* ($F_d\sigma$). Finally, we estimate the relationships between means and *popularity*—the number of firms (with non-missing s) who export Champagne to destination d .

Table 8 estimates the relationship between the country means (for quality, price, and quantity) and the gravity variables that determine attractiveness (A_d) and, possibly, fixed costs (F_d) as well. We restrict the sample to the countries where at least two French firms export. For the reason discussed in subsection 2.5, country mean regressions are weighted by the number of firms that export to that country.²²

The quality sorting model predicts that any of the gravity variables that raise the number of firms who export to a market (its popularity) should lower average quality. The ratio of the coefficients in columns (1) and (2) should be $-\kappa$. The variables that lower average quality should have the same effect on price, but with the magnitude scaled down by $\hat{\lambda} = .29$. The quantity regressions should have the opposite sign from quality and price, so long as the impacts of the gravity variables on fixed costs are not too large.

The sign pattern of the results shown in table 8 conforms to these priors remarkably well. Market size variables (population, income, wine consumption) all raise popularity

²⁰The CES indirect utility is y_d/P_d . Indirect utility holds constant while s rises if and only if $(p_d(i)\tau_d)/(\alpha_d(i)s(i)^\gamma)$ remains unchanged: Hence $p_d(5)/p_d(1) = 5^\gamma$.

²¹Hallak (2006) reports a median estimate that implies $\gamma_1 = 0.03$ with the same assumption of $\sigma = 7$.

²²Mechanically this implies multiplying LHS and RHS variables by $\sqrt{N_{dt}}$.

Table 8: Mean-gravity regressions for exporter quality, prices, and quantities

	(1)	(2)	(3)	(4)	(5)	(6)
	$\ln N_{dt}$	$\ln \tilde{s}_{dt}$	$\ln \tilde{p}_{dt}$	$\ln \tilde{q}_{dt}$	$\ln \hat{A}_{dt}$	$\ln \widehat{F}_{dt}\sigma$
ln popn. (M_{dt})	0.37 ^a (0.042)	-0.09 ^a (0.011)	0.01 (0.013)	0.50 ^a (0.073)	0.79 ^a (0.099)	-0.21 ^a (0.039)
ln inc. p.c. (y_{dt})	0.69 ^a (0.043)	-0.07 ^a (0.013)	0.05 ^b (0.020)	0.62 ^a (0.080)	1.60 ^a (0.106)	-0.28 ^a (0.037)
ln cons p.c. (b_{dt})	0.07 ^c (0.041)	-0.04 ^a (0.012)	-0.01 (0.018)	0.14 ^c (0.079)	0.25 ^b (0.098)	-0.08 ^c (0.047)
ln prodn ($\searrow P_{dt}$)	-0.05 ^b (0.023)	0.01 (0.005)	0.00 (0.007)	-0.03 (0.032)	-0.11 ^b (0.050)	0.06 ^a (0.022)
ln distance ($\nearrow \tau_d$)	-0.08 (0.072)	0.08 ^a (0.022)	0.12 ^a (0.027)	0.05 (0.089)	-0.15 (0.190)	0.28 ^a (0.068)
French ($\searrow \tau_d$)	1.27 ^a (0.155)	-0.27 ^a (0.048)	-0.05 (0.048)	0.33 ^b (0.165)	2.40 ^a (0.348)	-0.38 ^a (0.137)
Method	OLS	GLS	GLS	GLS	OLS	OLS
Observations	907	775	775	775	775	775
R^2	0.668	0.743	0.334	0.757	0.691	0.242

Note: GLS regressions are performed with weight_{*d*} = N_d . Standard errors in parentheses are robust to arbitrary forms of remaining heteroskedasticity and clustered by country. Significance levels: ^c $p < 0.1$, ^b $p < 0.05$, ^a $p < 0.01$

and *lower* average quality. Distance lowers popularity but raises quality. Speaking French (which is assumed to lower trade costs) raises popularity and lowers average quality. Having high production of wine should reduce the price index (P_d) in a market. This should reduce popularity and therefore raise quality. The estimated signs are as expected. The performance with average price as the dependent variable is disappointing, as population does not enter significantly and income per capita enters positively where the model predicts it should lower average prices (just as it lowers average quality). However, the positive effect of distance effect supports quality sorting. For means of quantity shipped to each market, the quality-sorting model is supported by all variables except distance and local production, which are not statistically different from zero. In terms of the magnitude of coefficients, the average ratio of columns (1) and (2) over the six RHS variables give an estimate of $\hat{\kappa} = 4.57$, while the mean ratio of price to quality coefficients is 0.22, quite close to the $\hat{\lambda} = .29$ we obtained from firm-level regressions.

The last two columns of Table 8 empirically assess how closely our estimates of $\ln A_{dt}$ and $\ln F_{dt}\sigma$ derived from the firm-level regressions can be explained by the gravity variables. Recall that country-time fixed effects estimated in the export value equation

corresponds to $\ln A_{dt}$ in our model, and that the minimum export value per destination-year is our estimator of $F_{dt}\sigma$. As expected, all gravity variables usually associated with higher aggregate bilateral trade volumes (GDP, income per capita, common language, and proximity to France) tend to raise A_{dt} . In keeping with our interpretation of high local production of wine as a variable that reduces the domestic price index, we find it lowers A_d . We did not have strong priors on how destination-specific fixed export costs would relate to the gravity variables. We find that, when significant, each of these gravity variables has the opposite sign from what was estimated in the previous column. In all cases except distance the absolute magnitudes in column (6) are lower. This suggests that the parsimonious $F_d \propto A_d^\phi$ with $-1 < \phi < 0$ is a reasonable approximation. The factor of proportionality between attractiveness and fixed costs can be estimated more precisely using the ratio of column (6) to column (5) for each coefficient. Over the 6 variables, this ratio has an average of $\phi = -.56$. The explanatory power of the gravity determinants for this regression is somewhat lower than for the one on attractiveness, but still quite substantial.²³

Those gravity variables are therefore reasonable proxies for what we really want to capture: Attractiveness and fixed costs of exporting. Several problems arise with the use of these proxies however. First, each of the six RHS variables gives a different result to be compared with the predictions of Table 1. Second, using proxies restricts the analysis to checking the signs of effects, rather than on the precise value predicted by the model. Third, and most important, these proxies are incomplete and use *ad hoc* functional forms. This is a potential source of mis-specification.

We therefore proceed to the two methods which replace a long list of gravity determinants with one or two “indexes” that summarize *all* the relevant country-specific information. We start with the regressions on attractiveness and entry thresholds and then proceed to the regressions on popularity. Equations (14), (16) and (19) all reveal that F_{dt} should enter the regression. In conditional mean quality and price equations, $\ln F_d\sigma$ should enter with the same coefficient as $\ln A_d$, but with the opposite sign. The average quantity equation implies the signs should be the same for the two variables but one cannot impose a coefficient restriction.

The relationships between means and imputed attractiveness and fixed costs are reported in Tables 9 and 10. The first table estimates coefficients freely, while the first three columns of Table 10 constrain the coefficients to be equal but of opposite signs, as implied by the model. Comparing the two tables shows that the data do not object strenuously to these constraints.

Overall, the results support the quality-sorting model. Popularity is positively affected by attractiveness as predicted by equation (12). As predicted also, average quality is negatively related to attractiveness. The predicted coefficient in column (2) is $-1/\beta$. Using our $\hat{\beta} = 7.64$ from the firm-level regressions we would therefore have predicted

²³There is also a purely “statistical” interpretation for the results in columns (5) and (6). Since the dependent variable in (5) is based on the expected value of $\ln x_d(j)$ in each market and the dependent variable of (6) is based on the minimum $x_d(j)$, it is natural to expect these two statistics to be inversely related. The lower fit might arise because the minimum is a noisier statistic than the mean.

Table 9: Mean-attractiveness regressions (unconstrained)

	(1)	(2)	(3)	(4)
	$\ln N_{dt}$	$\ln \tilde{s}_d$	$\ln \tilde{p}_{dt}$	$\ln \tilde{q}_{dt}$
$\ln \hat{A}_{dt}$ (FE estimate of attractiveness)	0.36 ^a (0.008)	-0.06 ^a (0.005)	-0.01 (0.008)	0.27 ^a (0.052)
$\ln \widehat{F}_{dt}\sigma$ (entry threshold)	-0.28 ^a (0.016)	0.04 ^a (0.010)	0.05 ^a (0.018)	-0.12 (0.097)
Method	OLS	GLS	GLS	GLS
Observations	919	857	857	857
R^2	0.916	0.795	0.135	0.663

Note: GLS regressions are performed with $\text{weight}_d = N_d$. Standard errors in parentheses are robust to arbitrary forms of remaining heteroskedasticity and clustered by country. Significance levels: ^c $p < 0.1$, ^b $p < 0.05$, ^a $p < 0.01$

Table 10: Mean-attractiveness regressions (constrained)

	(1)	(2)	(3)	(4)
	$\ln N_{dt}$	$\ln \tilde{s}_d$	$\ln \tilde{p}_{dt}$	$\ln \tilde{q}_{dt}$
$\ln \hat{A}_{dt}$ (FE estimate of attractiveness)	0.34 ^a (0.007)	-0.06 ^a (0.003)	-0.01 ^a (0.005)	0.27 ^a (0.052)
$\ln \widehat{F}_{dt}\sigma$ (entry threshold)	-0.34 ^a (0.007)	0.06 ^a (0.003)	0.01 ^a (0.005)	-0.12 (0.097)
Method	OLS	GLS	GLS	GLS
Observations	919	857	857	857
R^2	0.912	0.792	0.112	0.663

Note: GLS regressions are performed with $\text{weight}_d = N_d$. Standard errors in parentheses are robust to arbitrary forms of remaining heteroskedasticity and clustered by country. Significance levels: ^c $p < 0.1$, ^b $p < 0.05$, ^a $p < 0.01$

an elasticity of -0.13 instead of -0.06 . The discrepancy may arise because the means regressions assume Pareto-distributed s whereas the firm-level regressions use the actual observed distribution of quality. The sign on the price effect is supportive of quality sorting but it is only weakly statistically significant. Furthermore, the implied λ of 0.17 is smaller than the 0.29 estimated using firm-level data. Finally, the quantity relationship is strongly significant for $\ln A_{dt}$. Indeed the elasticity of quantity with respect to attractiveness is too strong to be consistent with the theory's prediction that the price and quantity effects be equal in absolute value. This asymmetry in the magnitudes is not evidence against quality sorting since it also runs counter to the prediction of the productivity sorting model. Rather, we believe the asymmetry casts doubt on the assumption of Pareto-distributed firm heterogeneity.

Since the quality and price relationships with popularity do not involve other variables, we can examine them directly using scatterplots of averages versus the number of exporters. For the quantity-popularity relationship, we have to assume that variation in fixed costs is white noise in order to justify the two-dimensional figure. With Pareto-distributed heterogeneity, the quality sorting and productivity sorting models both predict that all three relationships should be linear in log scale. Furthermore, both models predict equal absolute slopes of opposite signs for the mean price and quantity figures. The quality sorting model predicts the negative average quality-popularity relationship, negative price-popularity relationship, and positive quantity-popularity relationship.

The three scatterplots shown as panels (a)–(c) of figure 3 mainly support the quality sorting predictions. Average quality and popularity exhibit strong negative relationships in panel (a)—once popularity is sufficiently high. The weighted least squares estimate for the $N_d > 4$ sample is -0.21 . This implies a Pareto shape parameter of $\kappa \approx 5$, very close to the average estimate of 4.57 from mean gravity regressions.


Although the quality-popularity relationship is not globally linear, this may be due to small-sample issues for the less popular markets. The mean price panel (b) exhibits considerable noise. The slope is only mildly negative. Some very popular markets like Japan (JPN) have high prices that run counter to the model. Note that we expect the price relationship to be less steep than the quality relationship. Indeed, inspecting (15) and (17), the ratio of the price slope to the quality slope should be equal to λ , which we estimated in the previous section to be equal to 0.29 . The price slope should therefore be around -0.06 , very close to the -0.05 that we obtain.

Panel (c) reveals a strong positive relationship between average quantity and popularity. It seems to be linear in logs as predicted by the model under Pareto distributed quality. However, the slope is too large. Our theoretical predictions summarized in Table 1 imply a positive slope of 0.06 . Instead, we find an effect that is an order of magnitude too large for the Pareto-distributed quality sorting model.²⁴


Table 11 presents regressions that correspond to the three panels of Figure 3, with two differences. First, we pool all the data on A_{dt} and N_{dt} for the six years, rather than taking averages. More importantly we do not assume fixed costs are orthogonal

²⁴Crozet et. al (2009) show that such large effects *can* be obtained in quality sorting models with log-normal heterogeneity.

(a) Quality



(b) Price



(c) Quantity




Table 11: Mean-popularity regressions for exporter quality, prices, and quantities

	(1)	(2)	(3)
	$\ln \tilde{s}_d$	$\ln \tilde{p}_{dt}$	$\ln \tilde{q}_{dt}$
$\ln N_{dt}$ (popularity)	-0.19 ^a (0.013)	-0.05 ^b (0.018)	0.84 ^a (0.144)
$\ln \widehat{F}_{dt}\sigma$ (entry threshold)			0.07 (0.113)
Observations	857	857	857
R^2	0.756	0.105	0.654

Note: All regressions are GLS performed with $\text{weight}_d = N_d$. Standard errors in parentheses are robust to arbitrary forms of remaining heteroskedasticity and clustered by country. Significance levels: ^c $p < 0.1$, ^b $p < 0.05$, ^a $p < 0.01$

to popularity, but rather introduce our estimated fixed costs into the mean quantity regression. Columns (1) and (2) show coefficients on average quality and prices that are quite similar to the ones in the figure. The price slope of -0.05 is remarkably close to the -0.06 predicted by multiplying the quality coefficient, -0.19 , by our estimate of λ from the firm-level regressions (0.29). As in the figure, the quantity coefficient is much higher than expected (0.84 instead of 0.06), but still has a sign consistent with quality sorting. Somewhat surprisingly, our estimate of fixed costs does not enter significantly even though it was predicted by equation (20) to have a unit elasticity in this regression.

Taken together with the previous results, cross-country variation in mean prices appears to be driven by forces outside the basic model. Noise in unit values is to be expected but perhaps greater predictive power would be possible in a model with some pricing-to-market. The Dixit-Stiglitz-Krugman prediction of destination-invariant FOB prices seems hard to reconcile with the data. Alternatively, since many firms produce more than one quality level, average prices may vary across markets due to shifts in the destination-specific *composition* of exports. This “mixing-to-market” (our term) is considered by Manova and Zhang (2009) to be important for explaining the pattern of firm-destination-level unit values of Chinese exporters.

5 Conclusion

Heterogeneous firms theory implies ability sorting: Bad firms tend to serve only the markets where it is easy to be profitable whereas good firms serve those markets as well as the more difficult ones. We have illustrated the importance of quality sorting for trade by examining an industry where could obtain direct measures of quality. We show empirically that firms with higher measured quality are more likely to export, export more, and charge higher prices. We also identify a severe selection bias issue that is likely

to be present in any firm-level regression that tries to assess determinants of firms' export performance. Monte Carlo simulations show that a Tobit method removes almost all the bias and leads to much more reliable estimates of the structural parameters. Depending on the assumed elasticity of substitution, our estimates imply that the value for the consumer of increasing quality (which translates into higher price) is 4 to 7.5 times larger than the costs associated with this higher quality. Quality pays in this industry and this is true with respect to all destinations, even allowing for lower valuations of quality in poor countries.

We also develop predictions for aggregate statistics on French exporters which allow for discriminating criteria between our model of quality sorting and the traditional productivity sorting model. The average of directly measured quality falls with increases in the attractiveness of a market. Average prices also fall and average quantity rises for the markets we estimate to be more attractive, which again supports quality sorting for the Champagne industry.

There are a certain number of points that have been left unanswered in this paper and which will be the focus of further work. First, we would like to know how much of the discriminating criteria we developed here between quality and efficiency sorting is general, and how much is specific to our assumptions. The first suspect for lack of generality is the Pareto distribution. This functional form proves very convenient for working with CES and multiplicative trade costs, but that tractability comes at the cost of fragile predictions. In particular the strict inversion of coefficients between average price and quantity appears to be a Pareto-dependent prediction. The second assumption that should be relaxed is the constant markup rule. In quality-driven competition, it seems natural that higher quality firms would charge higher markups, and therefore their higher prices might not arise solely from higher costs. Departing from Dixit-Stiglitz would also make prices depend on destination market characteristics, which is important given the significant cross-country mean price heterogeneity exhibited in the Champagne FOB prices. Finally, the multiplicative iceberg trade costs prevents an Alchian-Allen effect, which might be at work in the real data, together with the selection effects we have emphasized.

6 References

- Alchian, A. A. and W. R. Allen, 1964, *University Economics*, Belmont, CA: Wadsworth Publishing Company.
- Anderson J. E. and E. Van Wincoop, 2004, "Trade costs," *Journal of Economic Literature* 42(3), 691–751.
- Baldwin, R. E. and J. Harrigan, 2007, "Zeros, Quality and Space: Trade Theory and Trade Evidence," NBER Working Paper No. 13214.
- Baldwin, R. E. and T. Ito, 2008, "Quality competition versus price competition goods: an empirical classification," NBER Working Paper No. 14305.

- Bernard, A., B. Jensen, 1999, "Exceptional exporter performance: cause, effect, or both?" *Journal of International Economics* 47, 1-25
- Bernard, A., B. Jensen, S. Redding and P. K. Schott, 2007, "Firms in International Trade," *Journal of Economic Perspectives* 21(3), 105–130.
- Chaney, T. 2008. "Distorted Gravity: the Intensive and Extensive Margins of International Trade," *American Economic Review* 98(4), 1707-1721.
- Crozet, M., K. Head, and T. Mayer, 2009, "Exporter prices, quantities, and cross-market sorting: Discriminating evidence from France," work-in-progress.
- Eaton, J. and S. Kortum, 2001, "Trade in Capital Goods," *European Economic Review* 45, 1195–1235.
- Eaton J., S. Kortum and F. Kramarz, 2004, "Dissecting Trade: Firms, Industries, and Export Destinations," *American Economic Review, Papers and Proceedings* 93, 150–154.
- Eaton, J., S. Kortum and F. Kramarz, 2008, "An Anatomy of International Trade: Evidence from French Firms," NBER Working Paper 14610.
- Foster, L. J. Haltiwanger and C. Syverson, 2008, "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?" *American Economic Review* 98(1), 394-425.
- Hallak, J-C. 2006. "Product Quality and the Direction of Trade," *Journal of International Economics* 68(1), 238–265.
- Hallak, J-C. and J. Sivadasan, 2008. "Productivity, Quality and Exporting Behavior Under Minimum Quality Requirements," *mimeo*.
- Helpman E., M. J. Melitz and Y. Rubinstein, 2008, "Estimating Trade Flows: Trading Partners and Trading Volumes," *Quarterly Journal of Economics* 123, 441–487.
- Helpman E., M. J. Melitz and S. Yeaple, 2004, "Export Versus FDI with Heterogeneous Firms," *American Economic Review* 94, 300-316.
- Hummels, D. and P. J. Klenow, 2005, "The Variety and Quality of a Nation's Exports," *American Economic Review* 95, 704–723.
- Hummels, D. and A. Skiba, 2004, "Shipping the Good Apples Out? An Empirical Confirmation of the Alchian-Allen Conjecture," *Journal of Political Economy* 112, 1384–1402.
- Iacovone L. and B. Javorcik, 2008, "Shipping Good Tequila Out: Investment, Domestic Unit Values and Entry of Multi-product Plants into Export Markets," *mimeo* University of Oxford.

- Johnson R. C., 2009, “Trade and Prices with Heterogeneous Firms,” mimeo.
- Kugler M, and Verhoogen E., 2008, “The Quality-Complementarity Hypothesis: Theory and Evidence from Colombia,” NBER working paper No. 14418.
- Lawless, M. and K. Whelan, 2007, “A Note on Trade Costs and Distance,” mimeo Central Bank of Ireland.
- Mandel, Benjamin, 2008, “Heterogeneous Firms and Import Quality: Evidence from Transaction-Level Prices” UC Davis mimeo.
- Manova, K. and Z. Zhang, 2009, “Export Prices and Heterogeneous Firm Models,” mimeo Stanford University.
- Mayer T. and G. Ottaviano, 2007, “The happy few: new facts on the internationalisation of European firms”, Bruegel-CEPR EFIM2007 Report, Bruegel Blueprint Series.
- Melitz M.J., 2003, “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity,” *Econometrica* 71(6), 1695–1725.
- Melitz, M.J. and G.I.P. Ottaviano, 2008, “Market Size, Trade, and Productivity,” *Review of Economic Studies* 75, 295–316.
- Schott, P.K., 2004, “Across-Product versus Within-Product Specialization in International Trade,” *Quarterly Journal of Economics* 119(2), 647–678.
- Steinberger, M., 2005, “American Sparkling Wines: Are they ever as good as Champagne?” *Slate*, Posted at <http://www.slate.com/id/2132509/> on Friday, Dec. 30, 2005.
- Verhoogen E., 2008, “Trade, Quality Upgrading and Wage Inequality in the Mexican Manufacturing Sector,” *Quarterly Journal of Economics* 123(2), 489–530.
- Yeaple, S.R., 2009, “Firm heterogeneity and the structure of U.S. multinational activity,” *Journal of International Economics*

A Derivation of conditional expectations

Steps for deriving the conditional expectations shown in subsection 2.4 are provided below.

A.1 Quality

To obtain the numerator of equation 13, we start by integrating over s , conditional on α :

$$\int_{\hat{s}_d(\alpha)}^{\infty} sg(s)ds = \hat{s}_d(\alpha)^{1-\kappa} \underline{s}^{\kappa} \frac{\kappa}{\kappa - 1}. \quad (27)$$

Substituting the expression for $\hat{s}_d(\alpha)$ shown in (9) and integrating over all values of α , the numerator of (13) is

$$\int_0^\infty \int_{\hat{s}_d(\alpha)}^\infty sg(s)h(\alpha)dsd\alpha = \hat{s}_d(1)^{1-\kappa} \underline{s}^\kappa \frac{\kappa}{\kappa-1} \mu_2, \quad (28)$$

where μ_2 is defined as

$$\mu_2 \equiv \int_0^\infty \alpha^{\frac{\kappa-1}{(\sigma-1)(\gamma-\lambda)}} h(\alpha) d\alpha.$$

Dividing (28) by the probability of entry obtained from equation (12), the expected value of quality exported to a given market is

$$E[s \mid \pi_d(j) > 0] = \left(\frac{A_d}{\sigma F_d \omega^{\sigma-1}} \right)^{\frac{-1}{(\sigma-1)(\gamma-\lambda)}} \frac{\kappa}{\kappa-1} (\mu_2/\mu_1), \quad (29)$$

A.2 Price

The expected price conditional on exporting is given by

$$E[p \mid \pi_d(j) > 0] = \frac{\int_0^\infty \int_{\hat{s}_d(\alpha)}^\infty p(s)g(s)h(\alpha)dsd\alpha}{\Pr[\pi_d(j) > 0]}. \quad (30)$$

To obtain the numerator we start by plugging in equation (7) for $p(s)$ and integrating over s , conditional on α :

$$\frac{\sigma\omega}{\sigma-1} \int_{\hat{s}_d(\alpha)}^\infty s^\lambda g(s)ds = \hat{s}_d(\alpha)^{\lambda-\kappa} \underline{s}^\kappa \frac{\sigma\omega\kappa}{(\sigma-1)(\kappa-\lambda)}. \quad (31)$$

For the integral to be finite we need $\kappa > \lambda$. Substituting the expression for $\hat{s}_d(\alpha)$ shown in (9) and integrating over all values of α , the numerator of (30) is

$$\int_0^\infty \int_{\hat{s}_d(\alpha)}^\infty p(s)g(s)h(\alpha)dsd\alpha = \frac{\omega\sigma\kappa\underline{s}^\kappa}{(\sigma-1)(\kappa-\lambda)} \hat{s}_d(1)^{\lambda-\kappa} \mu_3, \quad (32)$$

where μ_3 is defined as

$$\mu_3 \equiv \int_0^\infty \alpha^{\frac{\kappa-\lambda}{(\sigma-1)(\gamma-\lambda)}} h(\alpha) d\alpha.$$

Dividing (32) by the probability of entry obtained from equation (12), the expected value of price exported to a given market is

$$E[p \mid \pi_d(j) > 0] = \left(\frac{A_d}{\sigma F_d} \right)^{\frac{-\lambda}{(\sigma-1)(\gamma-\lambda)}} \omega^{\frac{\gamma}{(\gamma-\lambda)}} \frac{\sigma\kappa}{(\sigma-1)(\kappa-\lambda)} (\mu_3/\mu_1). \quad (33)$$

A.3 Quantity

The expected quantity, conditional on exporting profitably to market d , is given by

$$E[q \mid \pi_d(j) > 0] = \frac{\int_0^\infty \int_{\hat{s}_d(\alpha)}^\infty q(s)g(s)h(\alpha)dsd\alpha}{\Pr[\pi_d(j) > 0]}. \quad (34)$$

Quantity $q_d(s)$ is obtained by dividing (6) by the FOB price equation, (7):

$$q_d(j) = \frac{\sigma - 1}{\sigma} \omega^{-\sigma} s(j)^{(\sigma-1)\gamma - \lambda\sigma} A_d \alpha_d(j). \quad (35)$$

Substituting in equation (18) for q , we start by evaluating $\int_{\hat{s}_d(\alpha)}^\infty q(s)g(s)ds$:

$$\int_{\hat{s}_d(\alpha)}^\infty q(s)g(s)ds = \frac{\kappa \underline{s}^\kappa \omega^{-\sigma} A_d \alpha}{\kappa - (\sigma - 1)\gamma + \lambda\sigma} \frac{\sigma - 1}{\sigma} \hat{s}_d(\alpha)^{(\sigma-1)\gamma - \lambda\sigma - \kappa}. \quad (36)$$

For the integral to be finite, we assume $\kappa > \gamma(\sigma - 1) - \lambda\sigma$. The above expression is more complex than the corresponding equation, (31) obtained for prices. In particular, *both* A_d and \hat{s}_d enter average exports, while only $\hat{s}_d(\alpha)$ enters average price. The reason has to do with the intensive and extensive margins of trade increases in this model. In a Dixit-Stiglitz setup, prices are a constant markup over marginal costs, and in particular do not depend on market size or anything that enters A_d . Therefore, a rise in market attractiveness A_d impact prices only through the extensive margin, the entry of firms into export market d , the $\hat{s}_d(\alpha)$ term in (31). Quantities sold by each firm that exports to d do however depend on A_d . Consequently, (36) depends on the extensive margin $\hat{s}_d(\alpha)$, but also on the intensive one through the independent impact of A_d .

Next, we substitute (9) into (36) and integrate over α to obtain

$$\int_0^\infty \int_{\hat{s}_d(\alpha)}^\infty q(s)g(s)h(\alpha)dsd\alpha = \frac{\kappa \underline{s}^\kappa \omega^{-\sigma} A_d}{\kappa - (\sigma - 1)\gamma + \lambda\sigma} \frac{\sigma - 1}{\sigma} \hat{s}_d(1)^{(\sigma-1)\gamma - \lambda\sigma - \kappa} \mu_4, \quad (37)$$

where μ_4 is defined as

$$\mu_4 \equiv \int_0^\infty \alpha^{\frac{\kappa + \lambda}{(\sigma-1)(\gamma-\lambda)}} h(\alpha) d\alpha.$$

The final step is to divide (37) by (12), the probability of being a profitable exporter to d , yielding

$$E[q \mid \pi_d(j) > 0] = \left(\frac{A_d}{\sigma} \right)^{\frac{\lambda}{(\sigma-1)(\gamma-\lambda)}} F_d^{\frac{(\sigma-1)\gamma - \sigma\lambda}{(\sigma-1)(\gamma-\lambda)}} \omega^{\frac{-\gamma}{(\gamma-\lambda)}} \frac{\kappa(\sigma - 1)}{\kappa - (\sigma - 1)\gamma + \sigma\lambda} (\mu_4/\mu_1). \quad (38)$$