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

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

June 28, 2011

Forthcoming in *The Review of Economic Studies*

Abstract

Empirical investigation of the quality interpretation of the Melitz (2003) model of firm heterogeneity and trade has been limited by lack of direct data on quality. This paper matches firm-level export data with expert assessments of the quality of Champagne producers to estimate the key parameters of that model. Quality monotonically increases firm-level prices, the probability of market entry, and export values. The estimated model—which calibrates the relative importance of firm-level quality and idiosyncratic demand—accurately predicts the average quality exported to each country. Simulations show that the data reject the polar alternatives where outcomes are based entirely on either quality or randomness.

*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 Commission's 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 the French customs administration for making the exports data available at CEPII, Bernice Bray and Isaac Holloway for superb research assistance and Yvan Decreux for additional data. Participants in conferences and seminars at Louvain-la-Neuve, Nottingham, NOITS, ERWIT, LSE, Princeton University, Hong Kong University, Chinese University of Hong Kong, George Washington University, Kyoto University, Stockholm School of Economics, the Graduate Institute in Geneva, and Penn State provided helpful suggestions. We thank in particular Bruno Biais, Andrew Bernard, Jonathan Eaton, James Harrigan, Nathan Nunn, Henry Overman, Esteban Rossi-Hansberg, Andrés Rodríguez-Clare, Scott Taylor, Dan Trefler, Stephen Yeaple, and two anonymous referees.

[†]Paris School of Economics (Paris I) and CEPII

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

[§]Sciences Po, CEPII, and CEPR. Corresponding author: thierry.mayer@sciences-po.fr

1 Introduction

Since firm-level data on trade have become available, researchers have documented overwhelming evidence of dramatic differences in export outcomes. 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. Only the highest performing firms export substantial amounts to large sets of destinations.[1] While such differences in outcomes are well-established, the underlying source of this heterogeneity remains unclear.

In line with the vast literature estimating production functions, empirical work on firm-level heterogeneity has focused on differences in technical efficiency. Foster, Haltiwanger, and Syverson (2008) point out that the bulk of this work uses revenue-based productivity measures that confound the "separate and opposing effects of technical efficiency and demand." Proxies for firm-level productivity such as value-added per worker (Bernard and Jensen, 1999) or sales in the home market (Eaton et al., forthcoming, and Yeaple, 2009) could be driven by primitives other than physical output per unit of input. Indeed the seminal paper linking firm heterogeneity and trade, Melitz (2003), points out that the "productivity" term in the model can be thought of as either a cost shifter or a demand-shifting quality variable. Casual observation suggests that the quality interpretation of productivity could apply in many industries. However, 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 for quality such as unit values. In this respect, we propose in this paper the first empirical attempt to validate the *quality interpretation of the Melitz (2003) model*, combining firm-level data that directly measure quality and trade.

To achieve this, our paper studies the exports of Champagne producers, where firm-destination export flows can be matched to firm quality ratings from a comprehensive guidebook. We argue that the Champagne industry conforms reasonably well to the heterogeneous firm monopolistic competition assumptions of Melitz (2003). 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 country. The results show that there is a payoff to quality in terms of greater presence in export markets. We find that the quality effect remains strong even for the subset of firms for which we can also measure—and control for—physical productivity (bottles per worker). Furthermore, selection is a key feature in the data, as predicted in models with fixed costs of entering each market. The evidence we find of quality sorting follows the predictions of the Melitz (2003) model quite well with

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

two provisos: (i) the principal difference between firms is quality, not physical productivity; (ii) idiosycratic demand plays an important role in determining export outcomes. Since the model was not tailored to fit unique characteristics of the Champagne industry, we expect these results to apply to other industries.

Our paper fits within an evolving literature on the relationship between quality and trade. A first generation of empirical work investigates the attributes of countries that export and import higher quality goods—as inferred from unit values. Schott (2004) and Hummels and Klenow (2005) finds that within goods categories, unit values tend to increase with the exporters' per capita income. Khandelwal (2010) critiques the use of unit values as proxies for quality and instead infers exporter product quality by comparing market shares conditional on price. His results corroborate the earlier finding that higher income countries export higher quality goods. On the destination side, Hallak (2006) finds some evidence that richer countries have relatively greater demand for high unit value source countries. 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.[2]

A second generation of papers uses product-level trade data to test the implications of models of firm-level heterogeneity in quality based upon Melitz (2003). Baldwin and Harrigan (2007) propose a model where increased utility more than compensates for higher production costs when product quality rises. 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. 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 source-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 indirect 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

[2]The Alchian-Allen effect, also known as "shipping the good apples out" arises when freight costs depend on weight, rather than being proportional to value as per the iceberg assumption. An increase in freight costs therefore lowers relative delivered prices and thereby raises the relative attractiveness of high quality goods for distant consumers.

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 (2009) also find a positive relationship between exporting and output prices using data from India, the United States, Chile and Colombia. 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 match direct quality measures obtained from a guidebook (for individual producers) with firm-level export values and quantities. We can then relate the proxies that have been used for quality, unit values and market shares, to the direct measure. This exercise would be impossible for many products where intermediaries, rather than producers, handle exports. However, 94% of Champagne exports are handled by firms for which we could obtain quality ratings.

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 Tobit solution is easy to implement and Monte Carlo simulations indicate that it corrects the selection bias. The firm-level regressions allow us to estimate the structural parameters of the model in terms of producers' marginal cost of quality and the benefit-cost ratio for each quality rating as perceived by consumers. We corroborate the Baldwin and Harrigan (2007) hypothesis that quality raises prices but that it raises valuations by a more than offsetting amount such that higher quality firms sell more.

We then propose several extensions of our model to check whether our most important findings on quality sorting are robust to alternative mechanisms that have been proposed in the literature. A first extension introduces the possibility of sources of firm-level heterogeneity other than quality. We find that our measure of quality can account for 27% of the firm-specific factor explaining export values. For the set of firms for which we observe quality and physical productivity, the former explains 36% of firm-level variation and the inclusion of the latter only raises the R^2 to 38%. Thus, among observables, quality is by far the stronger influence on export success. Another set of extensions deals with deviations from the standard Dixit-Stiglitz-Krugman framework. We first introduce a non-iceberg component to trade costs in our model, leading to pricing-to-market and Alchian-Allen effects. We operationalize this by interacting quality with distance in the export price and value equations. Neither pricing to market (based on distance) nor Alchian-Allen effects show up in a systematic way. We then allow for non-homotheticity in the demand system by letting high-income countries exhibit a higher taste for high-quality Champagne. Although there is some evidence that demand for Champagne increases with income per capita, higher quality raises exports within all income categories.

Finally, we evaluate the fit to the aggregate data of our model, and conduct falsification exercises. Using the structural parameters of the model estimated in the firm-level regressions, we simulate the predicted average quality and price of the Champagne exporters that enter each foreign destination market. We then confront this prediction with actual data and observe a very good fit, much better than alternative models of sorting—one with random entry and one where observed quality is the sole determinant of competitiveness in each market.

The paper proceeds as follows. The next section derives firm-level estimating equations from a Melitz-based model of firm-level heterogeneity in quality. 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, where we also back out the implied values of the key structural parameters. Section 5 proposes several extensions to our baseline model to allow for other sources of firm-level heterogeneity, Alchian-Allen effects, and non-homothetic preferences. Section 6 compares averages of quality and price for each market with the predictions of the model. Our conclusion evaluates the quality version of the Melitz (2003) model in light of the evidence from Champagne and comments on which extensions matter the most.

2 A simple model of heterogeneous quality and exports

The model examined in this paper is based on Melitz (2003). We maintain the Dixit-Stiglitz-Krugman assumptions on demand, market structure, and trade costs in the baseline model but consider generalizations in later sections. The model also draws key ideas from Baldwin and Harrigan (2007) and Eaton, Kortum, and Kramarz (forthcoming).

2.1 Entry, export values, and prices

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, Ω_d, of all varieties, j, available in country d:

$$U_d = \left(\int_{j \in \Omega_d} [a_d(j) b[s(j)] q(j)]^{\frac{\sigma-1}{\sigma}} \, dj \right)^{\frac{\sigma}{\sigma-1}}. \quad (1)$$

In this expression $q(j)$ denotes quantity of variety j consumed and $s(j)$ denotes the quality provided by variety j.[3] The $b()$ function maps quality into quantity equivalents in the utility of the consumer. The $a_d(j)$ are destination d-specific demand parameters capturing country-level deviations in utility relative to the firm-level $b(s(j))$. Such shocks are a feature that Eaton, Kortum and Kramarz (forthcoming) add to the Melitz model.

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. Thus, the $a_d(j)$ provide a structural error term for firm-level regressions, as shown in subsection 2.2.[4]

The sub-utility enters a Cobb-Douglas full utility, implying that consumers of country d spend an exogenous total amount, X_d, on the good. We abstract from the issue of firms producing more than one variety within the narrow trade category (Champagne is an 8-digit product in our application) and model only differentiation across firm-level "brands."[5] Firm j's market share in country d is given by

$$\frac{x_d(j)}{X_d} = \frac{(p_d(j)/(a_d(j)b[s(j)]))^{1-\sigma}}{\int_{i\in\Omega_d}(p_d(i)/(a_d(i)b[s(i)]))^{1-\sigma}di}\mathcal{E}_d(j), \tag{2}$$

where $x_d(j)$ and $p_d(j)$ denote trade-cost inclusive (CIF) export values and prices, respectively, and $\mathcal{E}_d(j)$ is a dummy variable set to 1 when firm j enters market d.

Firms maximize profits in each destination market,

$$\pi_d(j) = (p_d(j) - c[s(j)]\tau_d)q_d(j) - F_d\mathcal{E}_d(j), \tag{3}$$

[3]The mnemonic for s is sensory services. In our application, $s(j)$ is measured as the number of stars that a wine expert assigns to a producer.

[4]Firm-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. (forthcoming) show that such "hierarchy" relationships do not hold strictly for French exports in general.

[5]Thus, we will treat for instance Veuve-Clicquot and Mumm as two competing varieties and disregard (due to data limitations) their *cuvée de prestige* bottles. The CIVC (2010) reports that such high-end Champagne accounts for only 2.3% of total exports in 2009.

where $\tau_d > 1$ is the iceberg trade cost for firms to reach market d and F_d is a destination-specific fixed market entry cost. We interpret s as a quality draw the firm makes upon entering the industry in just the way productivity is drawn in the Melitz (2003) model. As in Baldwin and Harrigan (2007) quality has consequences for the marginal cost of production. In particular, when firms draw a "recipe" for quality level s, this entails a set of ingredients and production methods such that marginal costs are given by $c(s)$. We expect $c'(s) > 0$ for two reasons. First, Kugler and Verhoogen (2009) hypothesize that firms who make high-quality output require high-priced inputs. Second, higher quality may entail greater factor input per physical unit of output. That is, "low yield" or "artisan" methods may be necessary in some industries to achieve high quality. We motivate these two mechanisms in the context of Champagne in subsection 3.1.

With iceberg transport costs, CES monopolistic competition implies a constant mark-up, $\sigma/(\sigma-1)$, yielding the CIF price equation

$$p_d(j) = \frac{\sigma}{\sigma-1} c[s(j)]\tau_d. \tag{4}$$

Substituting (4) into (2), export revenue from destination d is given by

$$x_d(j) = \left(\frac{a_d(j)b[s(j)]}{c[s(j)]}\right)^{\sigma-1} X_d \tau_d^{1-\sigma} P_d^{\sigma-1} \mathcal{E}_d(j), \tag{5}$$

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

$$P_d \equiv \tau_d \left(\int_{i\in\Omega_d} \left(\frac{a_d(j)b[s(i)]}{c[s(i)]}\right)^{\sigma-1} di\right)^{1/(1-\sigma)}.$$

We collect all country-specific determinants of exports into a single factor, A_d, defined as $A_d \equiv X_d \tau_d^{1-\sigma} P_d^{\sigma-1}$. We refer to A_d as the "attractiveness" of a destination market. It depends positively on the size (X_d) and relative accessability ($\tau_d^{1-\sigma} P_d^{\sigma-1}$) of the market. Following the notation of Eaton et al. (forthcoming), we also use $\alpha_d(j)$ to abbreviate $a_d(j)^{\sigma-1}$ and serve as the idiosyncratic demand shifter. Incorporating this notation, the export value equation can be expressed succinctly as

$$x_d(j) = \left(\frac{b[s(j)]}{c[s(j)]}\right)^{\sigma-1} A_d \alpha_d(j) \mathcal{E}_d(j). \tag{6}$$

Since the gross profit to sales ratio in Dixit-Stiglitz is given by $1/\sigma$, the net contribution to firm profits of destination d is given by

$$\pi_d(j) = x_d(j)/\sigma - F_d \mathcal{E}_d(j) \tag{7}$$

Substituting equation (6) into (7), the probability of exporting is

$$\mathbb{P}[\mathcal{E}_d(j) = 1] = \mathbb{P}\left[(b[s(j)]/c[s(j)])^{\sigma-1} A_d \alpha_d(j) > \sigma F_d\right]. \tag{8}$$

2.2 Estimating equations

We can estimate the model using firm-level data for three different dependent variables: the probability of exporting, the price and export value expressed in FOB terms. The reason we use FOB valuation is that it corresponds to the way export value data are collected by customs administrators.

From (4), the FOB price $p_d^{\text{fob}}(j)$ charged by firm j takes the following estimable form:

$$\ln p_d^{\text{fob}}(j) \equiv \ln(p_d(j)/\tau_d) = \ln c[s(j)] + \ln[\sigma/(\sigma-1)]. \tag{9}$$

We estimate (9) non-parametrically as

$$\ln p_d^{\text{fob}}(j) = \sum_s \lambda_s I_s(j) + \text{PFE}_d + \epsilon_d(j), \tag{10}$$

where the $I_s(j)$ are indicators for each level of quality. Each $\hat{\lambda}_s$ estimates $\ln c(s)$. We also estimate a specification that sets $\ln c[s(j)] = \lambda \ln s(j)$. This corresponds to the power function hypothesized by Baldwin and Harrigan (2007) and derived in models by Kugler and Verhoogen (2008), Mandel (2008), and Johnson (2009) where firms choose quality subject to a cost of upgrading. The PFE_d are destination-specific price fixed effects. They are not implied by the model but are included to make the specification robust to differences in markups or destination-specific costs that enter the FOB price such as special labeling, documents, or packaging requirements. The $\epsilon_d(j)$ in equation (10) is a non-structural error term reflecting deviations between a parsimonious model and an undoubtedly more complex data generation process.

Turning to the probability of exporting, we can take logs of (8) to specify that firm j will export to d with probability given by

$$\mathbb{P}[\mathcal{E}_d(j) = 1] = \mathbb{P}\left[(\sigma-1)\ln(b[s(j)]/c[s(j)]) - \ln(\sigma F_d) + \ln A_d + \ln \alpha_d(j) > 0\right]. \tag{11}$$

As with the price equation, we estimate the entry-quality relationship non-parametrically, replacing $(\sigma - 1)\ln(b[s(j)]/c[s(j)])$ with $\sum_s \beta_s I_s(j)$. The parameters can be estimated up to a scaling factor using a binary choice model whose form depends on the assumption made on the distribution of the unobserved heterogeneity term $\alpha_d(j)$. Assuming $\ln \alpha_d(j)$ is distributed normally with variance ψ^2 implies a probit form for the entry equation:

$$\mathbb{P}[\mathcal{E}_d(j) = 1] = \Phi\left[\psi^{-1}\left(\sum_s \beta_s I_s(j) + \mathrm{EFE}_d\right)\right], \tag{12}$$

where Φ represents the standard normal CDF. The logged attractiveness of country d, $\ln A_d$, and its fixed export costs, $\ln(\sigma F_d)$, appearing in equation (11) are captured with country-specific entry fixed effects, denoted EFE_d.

Since $\sigma > 1$, a positive estimate of β_s implies that "quality pays," i.e. consumer's marginal valuation of quality exceeds the marginal cost to producers for a level s of quality. We also use the power function approach. The latter specification assumes $b[s(j)] = s(j)^\gamma$, implying $(\sigma - 1)\ln(b[s(j)]/c[s(j)]) = \beta \ln s(j)$, where $\beta \equiv (\gamma - \lambda)(\sigma - 1)$ is an abbreviation for the elasticity of firm-level exports with respect to quality.

From (6), the log of FOB-valued firm-level exports is

$$\ln x_d^{\mathrm{fob}}(j) \equiv \ln(x_d(j)/\tau_d) = (\sigma - 1)\ln(b[s(j)]/c[s(j)]) + \ln(A_d/\tau_d) + \ln \alpha_d(j) + \ln \mathcal{E}_d(j). \tag{13}$$

Adopting the same non-parametric estimation approach as we use for prices and entry the export value equation is given by

$$\ln x_d^{\mathrm{fob}}(j) = \sum_s \beta_s I_s(j) + \mathrm{XFE}_d + \ln \alpha_d(j) + \ln \mathcal{E}_d(j). \tag{14}$$

The country-specific export value fixed effects, XFE_d, capture $\ln(A_d/\tau_d)$.

Assuming log-normal $\alpha_d(j)$ implies that OLS would be the maximum likelihood estimator for equation (14)—*if* we observed positive exports to all markets, i.e, if $\mathcal{E}_d(j) = 1 \quad \forall d, j$. In fact most Champagne exporters have positive exports to only a small number of destinations. This zero problem is predicted by the model assuming the fixed costs of exporting are non-negligible. The zero problem implies that OLS (with the dependent variable $\ln x_d^{\mathrm{fob}}(j)$ treated as missing for $x_d^{\mathrm{fob}}(j) = 0$) yields inconsistent estimates of the quality effect on exports.

Inspecting equation (11) reveals that among firms with identical quality levels, the probability of passing the cutoff and exporting increases with $\alpha_d(j)$. It follows that firms with high s will become exporters even with relatively low draws of $\alpha_d(j)$, whereas low quality firms need high draws of $\alpha_d(j)$ to be observed as positive exporters. This implied negative correlation in the selected data 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) also assume normally distributed errors and use a Heckman correction to address the sample selection issue. 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 (13). 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 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 is a direct implication of our model. Replacing $x_d(j)$ with $x_d^{\text{fob}}(j)\tau_d$ in equation (7), we see that firm j exports to destination d if and only if $x_d^{\text{fob}}(j) > \sigma F_d/\tau_d$. Thus we specify $\underline{x}_d^{\text{fob}} = \sigma F_d/\tau_d$ as the minimum value of $x_d^{\text{fob}}(j)$ that would be consistent with non-negative profits. Under the assumption of a log-normal distribution for $\alpha_d(j)$, we have Tobit structure in which $\ln x^{\text{fob}}$ is censored at $\ln \underline{x}_d^{\text{fob}}$. Although we do not observe $\underline{x}_d^{\text{fob}}$, Eaton and Kortum (2001) suggest that a maximum likelihood estimate of the censoring point can be obtained from the minimum observed positive value of exports, that is

$$\hat{\underline{x}}_d^{\text{fob}} = \min_{j \in \Omega_d} x_d^{\text{fob}}(j). \tag{15}$$

For $\ln x_d^{\text{fob}}(j) > \ln \hat{\underline{x}}_d^{\text{fob}}$ the likelihood is based on the continuous $\ln x_d^{\text{fob}}(j)$ from equation (13). For $\ln x_d^{\text{fob}}(j) < \ln \hat{\underline{x}}_d^{\text{fob}}$ the likelihood is the probability that $\ln x_d^{\text{fob}}(j) \leq \ln \hat{\underline{x}}_d^{\text{fob}}$.

To assess the magnitude of the OLS bias expected given the frequency of zeros in our data, we conducted Monte Carlo simulations. The simulations also establish the finite-sample properties of using the minimum to estimate the unknown threshold $\underline{x}_d^{\text{fob}}$. In addition, we investigate consequences of deviations from normality. Since we use estimated coefficients to parameterize the Monte Carlo, we present the simulation results together with the regression results in section 4 (with the alternative distributions shown in Appendix B.2).

3 Data

Our paper combines two main sources of data, firm-level export declarations and Juhlin (2008), a guide to Champagne producers. We start by discussing features of Champagne exporting that appear to conform to the main elements of our model. Then we describe the features of our data that affect estimation and interpretation.

3.1 Why Champagne?

Champagne is a leading example of the branded consumer products that are important contributors to exports for high-income countries. More importantly for our purposes, it has many features that make it well-suited for an empirical study of the quality-interpretation of the Melitz model. First, the monopolistic competition assumptions of large numbers of competitors selling products differentiated at the firm level are consistent with information we have on the Champagne industry. Second, Champagne as a whole appears to exhibit Armington-style differentiation by place of origin, an implicit assumption of the model. Third, the firms that handle Champagne exports, and hence are listed on the customs declarations we rely upon, are predominately producers to whom Juhlin (2008) assigns quality ratings. Finally, experts on Champagne have identified mechanisms that support the Baldwin-Harrigan assumption linking higher quality to higher marginal costs. We discuss each of these industry features below.

The monopolistic competition model assumes firms are sufficiently small and numerous to make oligopolistic interactions negligible. While we cannot make any strong statements about the presence or absence of oligopoly behavior in this industry, the data show that production is highly dispersed. Juhlin (2008) notes that of 5119 registered producers, "2454 make their own, unique champagne." For the 517 producers for which Juhlin provides production amounts, we calculate a Herfindahl index of 0.033. This amount of concentration is safely below the 0.2 threshold where oligopoly concerns arise according to Besanko et al. (2007, p. 198).

Champagne fits the assumption of firm-level differentiation *within the region* very well. The name of the maker is emphasized in marketing, rather than the vineyard or vintage. Other French wine regions are extensively subdivided, with each *appelation* purported to have distinct taste properties. Thus, the trade classification for red burgundy aggregates 100 appelations that range enormously in reputation for quality. In contrast Champagne is both a trade classification and a single appelation. Furthermore about 95% of Champagne is non-vintage (CIVC, 2010, p. 4), reflecting the common practice of blending wines from different years to maintain stable quality. This contrasts sharply with regions like Bordeaux and Burgundy that always show vintage on the label and exhibit substantial variation in quality due to weather. Since we do not observe the vintages of the exported bottles, Champagne's lower 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 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.[6]

Champagne producer associations use full-page advertisements and legal actions to reinforce the idea that sparkling wine that is made outside Champagne cannot be confused with "real" Champagne. Wine critics seem to agree with the proposition that sparking 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 anyplace else..." (Mike Steinberger, "American Sparkling Wines: Are they ever as good as champagne?", Slate.com, Dec. 30, 2005)
>
> "But when talking about sparkling wine, let's be honest: There is Champagne and there is everything else. The others are good, but they're not Champagne." (Eric Asimov, *New York Times* wine critic, Dec. 14, 2005)

Putting aside the factual validity of these claims, they do represent a widely held view that Champagne producers are collectively different from producers from other areas. This mitigates concerns over the omission of sparkling wine producers from outside Champagne in the trade and quality data.[7]

Cost-quality trade-offs 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. 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,

[6] A small number of Champagne producers do make sparkling wine outside of France but they do not label it as Champagne.

[7] In the empirics, fixed effects capture cross-country variation in the amount of competition from non-Champagne producers.

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. Data from INSEE classifies firms according to a "primary" activity. In wine regions other than Champagne, a large proportion of the firms named on export declarations do not correspond to the producers rated in the wine guides. Table 1 provides the share of exports according to primary activity in 2005. For Champagne, the growers and makers add up to 80% of the total, only 31% for white Burgundy, 22% for red Burgundy, and just 9% for red Bordeaux. Champagne is the notable exception to the rule that wholesalers dominate the export business since they account for only 13% of Champagne exports in 2005 compared to five or more times that share in the other regions.

Table 1: Exporters classified according to primary activity and size category

		Champagne	Burgundy		Bordeaux
			(white)	(red)	(red)
Activity	**code**	**Share of all Champagne exports**			
grape-growers	011G	2%	14%	20%	7%
wine-makers	159F/G	78%	17%	2%	2%
wholesalers	513J	13%	66%	65%	82%
others	various	7%	4%	12%	8%
Employees		**Share of producer exports**			
0–5		3%	22%	74%	66%
6–9		0%	8%	6%	4%
10–19		4%	44%	0%	1%
20–49		4%	16%	15%	1%
50–99		9%	9%	1%	0%
100–499		48%	2%	4%	29%
500+		32%	0%	0%	0%

The bottom panel of Table 1 suggests a possible reason for why producers, rather than wholesalers, dominate Champagne exports. It provides a breakdown of the shares of exports by growers and makers by size category. Larger firms (100+ employees) account for 80% of Champagne exports by producers, whereas these firms take much smaller shares in the other wine regions. In recent models such as Ahn et al. (2011), there is a trade-off between lower fixed costs of using intermediaries and lower variable costs of direct exporting. This leads to sorting patterns with larger firms choosing to export directly. Differences in production methods may underlie the larger sizes of Champagne makers. The Champagne region has no subdivisions, which allows producers to buy grapes from a much larger region while retaining the Champagne appelation. As a result, they face less strict capacity constraints than Bordeaux or Burgundy wine producers, where narrowly defined geographic areas force greater reliance on own grape production. Thus, Champagne's broad definition allows for firms to produce at greater scale.

3.2 Export data

The micro-data used in this paper are based on export declarations submitted to French Customs. We use the most recent data available at the time of writing, 2005. The customs data is an almost comprehensive record of annual shipments by destination country at the 8-digit product level for each French exporting firm. The "almost" is due to reporting thresholds for compulsory declarations inside and outside the customs union. If a firm's total exports to a non-EU country are below 1000 euros or 1000 kilograms, then it does not need to complete an export declaration. As discussed in the appendix, our Champagne export flows generally lie well above this threshold. Within the EU, firms that exported less than 100,000 euros to the whole EU in the previous year have the option of completing a short declaration that does not include destination country. These declarations do not enter our data set. For reasons considered in the appendix, many firms fill out the full declaration even when they are not required to do so. This discrepancy in statistical selection between EU and non-EU destinations might be important in our case, since our main prediction is about heterogenous firms sorting in different destination markets. In the results section we report estimates that show our results are robust to allowing for a differential impact of quality on EU vs non-EU exporting.

For each firm, product and destination country, Customs records FOB values and quantities. We calculate firm-destination-level FOB prices (often referred to as "unit values") as $p_d^{\text{fob}}(j) = x_d^{\text{fob}}(j)/q_d(j)$. The data provide two measures of quantity, the primary and supplementary. Sometimes the prices using the primary quantity units made no sense. We suspect this may have been caused by reporting quantity in tons when it should have been reported in kilos. We calculated the median price for each firm using the primary quantity measure. Then we recalculated prices using the secondary measure for all cases where this lowered the absolute deviation from the median price. The resulting prices had a first percentile of 8.34 EUR/kg and a 99th percentile of 66.83 EUR/kg. Two extreme outliers remained even after using the supplementary units: 0.33 and 8095.42 EUR/kg. We did not use these in any of our regressions.

Customs utilizes the 8-digit combined nomenclature product classification (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). France exported 1.87 billion euros of Champagne to 176 countries in 2005.

The export declaration data provides us with firm identification numbers, or SIREN, for all of the 1,134 firms who exported Champagne in 2005. The French national statistical agency (INSEE) provides the names, addresses, and primary activity code for most of those. We used the firm-level name and address information to match exporters with wine producers that were rated in Juhlin's Champagne guide. A total of 284 firms can be successfully matched with the rating textbook. All together, the Juhlin-rated firms account for 94% of French exports of Champagne in 2005, and serve 157 countries. Our final database would therefore include $284 \times 157 = 44,588$ observations except for the deletion of the two outliers mentioned above. This leaves us with 44,586 observations, of which 3,205 (7.2%) correspond to positive trade flows.

3.3 Quality ratings

Juhlin (2008) devotes 368 pages to detailed descriptions of "the most important active producers and their wines." He provides two scores for the overall quality of the Champagnes made by 487 different producers. One rating summarizes the quality of current wines while the second includes "all their historical baggage." We use the latter rating since it is likely to correspond better to our 2005 trade data. Producer ratings range from one to five stars, with a five stars reserved for "perfect" producers and one euphemistically assigned to a "producer whose wines have aroused my interest." Basing his producer scores on tastings of 6,500 different Champagnes, Juhlin claims that he has carried out "the most exhaustive evaluation of Champagne ever." This claim finds support in two observations. First, the 482 producers for which Juhlin provides stars and production collectively produce 291 million bottles per year. This represents 90% of 2008 total shipments of Champagne to all destinations, including the domestic market (CIVC, 2010, p. 2). Second, Juhlin (2008) rates about eight times more producers than the broader wine guides that have single chapters devoted to Champagne.

The other books, which formed the basis for the quality ratings we used in the Crozet et al. (2009) working paper, are nevertheless useful as comparisons with Juhlin for the wines they rate in common. The first source we consider is Robert Parker's internationally recognized *Wine Buyer's Guide*, 6th Edition, 2002. Like Juhlin (2008), Parker (2002) contains a rating of producers in which he assigns them up to five stars based on overall quality of the Champagnes they produce. Unfortunately Parker's ratings cover just 71 producers (of which we could match 62 to corresponding producers rated by Juhlin). A less well known source of producer ratings is the French guide, *Classement des meilleurs vins de France*, by Burtschy and Gerbelle (2006), published by the Revue du Vin de France. This book also assigns stars to producers but the maximum rating is three stars. Mere inclusion in the book is considered to be a positive evaluation.

Figure 1: Comparison of different quality ratings

(a) Parker vs Juhlin (b) *Revue du Vin de France* vs Juhlin

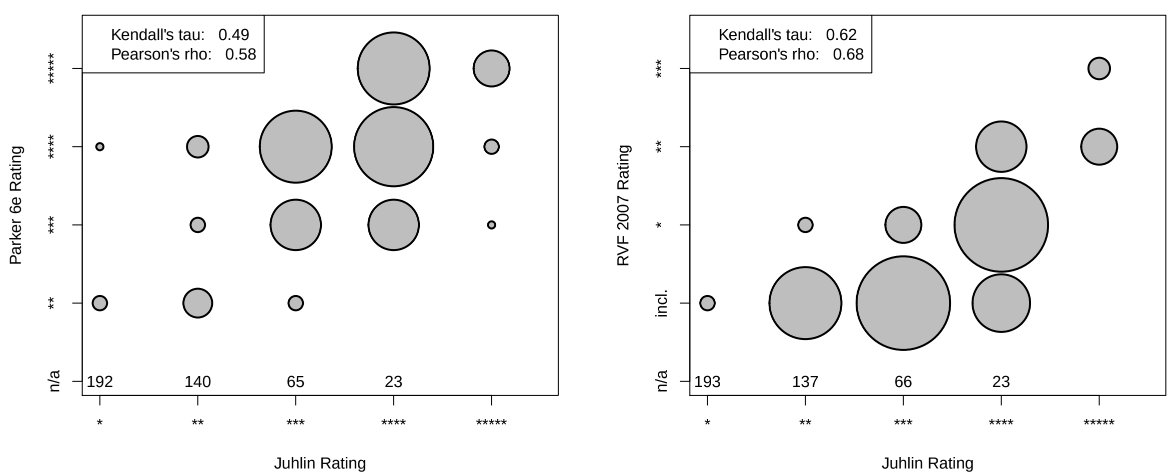

Figure 1 plots the ratings of Juhlin producers against the ratings these producers received from Parker (2002) and the Revue du Vin de France (RVF). Each circle is sized proportionately to the number of producers at that position. Thus, most producers given four stars by Parker are given three or four stars by Juhlin. However, small numbers were given higher or lower ratings. One thing that is evident is that Parker is a more generous marker than the other two raters, with a mode of 4 stars, compared to modes of 1 for Juhlin and "included" for the Revue. Despite these level differences, we see that the ratings are highly correlated. The Pearson correlations are 0.48 and 0.68. Kendall's τ, an index of concordance between ratings that varies between -1 for perfect disagreement and $+1$ for perfect agreement comes in at 0.49 and 0.62. In most cases, the large number of wines that Parker and the RVF omit are given just one star by Juhlin. However, omission does not necessarily imply low quality as there are 23 four-star wines in Juhlin that were omitted in the other two guidebooks. Because Juhlin (2008) is so much more comprehensive than the alternative sources of quality ratings, we rely upon it for the subsequent analysis. However, we conduct two robustness checks related to quality ratings: First, we replicate the main estimations for RVF and Parker ratings. Second, we add to the analysis the firms that export, despite not being rated by Juhlin, with the idea that those would be (with an important caveat detailed below) the lowest quality exporters. We briefly discuss the alternative results after presenting our baseline results but provide full tables and additional discussion in appendices.

Figure 2: Quality and prices

(a) Averages of (weighted) average prices (b) Medians of (weighted) average prices

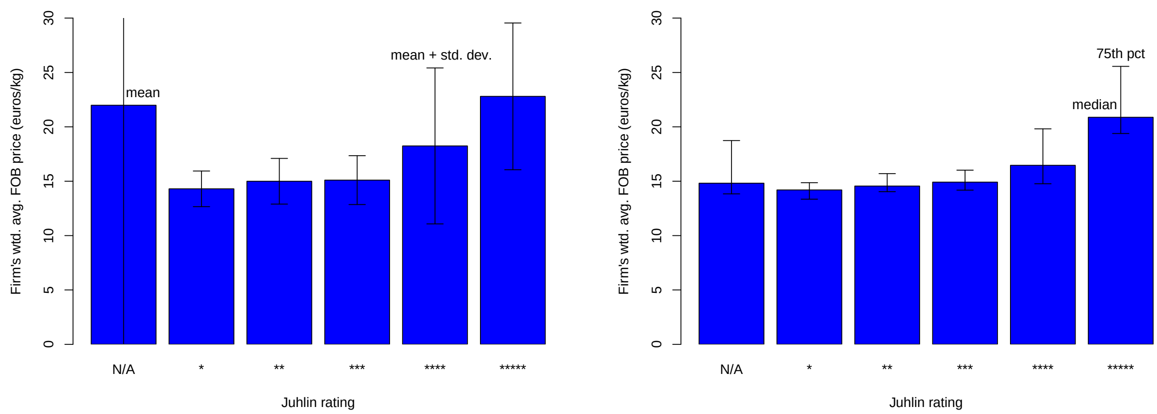

The next way we validate Juhlin's ratings is to examine graphically how they relate to prices (unit values), the proxy used by almost all of the studies of trade and quality mentioned in the introduction. For each firm we divided the total value of its exports (across all destinations) by total quantity to create a weighted average price (in euros per kg). Figure 2 plots the mean and standard deviation for each Juhlin star rating in panel (a) and the median and interquartile range in panel (b). Both the means and medians do not increase rapidly for one to three stars but prices do rise more impressively for four and five star Champagnes. The Champagne exported by non-Juhlin firms (depicted in the bars labeled "N/A") is much more heterogeneous in value, as is to be expected given that many of these exporters are wholesalers of a range of Champagne produced by other firms. The relationship between prices and quality shown in Figure 2 is partly attributable to the mix of destination countries to which each set of firms exports. In our regressions on individual firm prices we purge this compositional influence using country fixed effects.

Figures 3(a) and (b) show how the propensity to export and the number of export destinations relate to the Juhlin quality ratings. In panel (a) the height of the middle

Figure 3: Good Champagne goes further

(a) Exporters vs non-exporters (b) # of destination markets

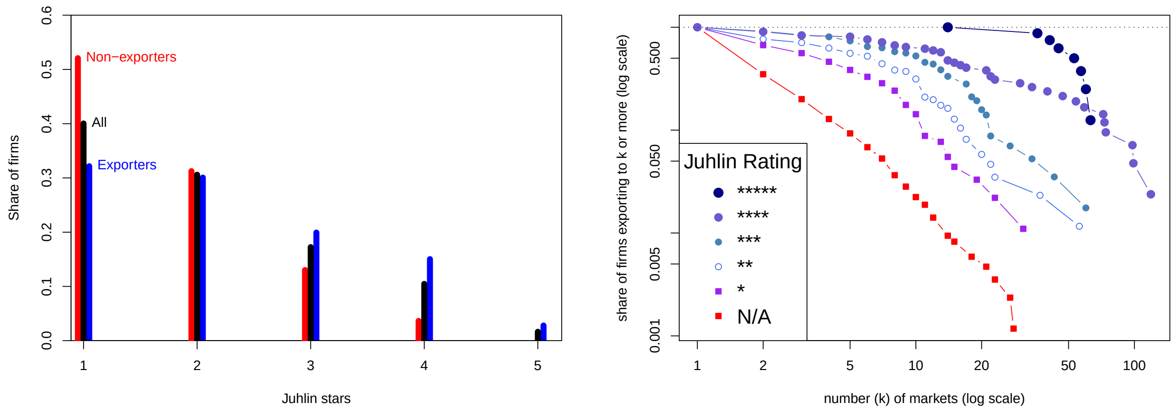

(black) bars corresponds to the shares of the 487 firms that obtain each rating. The (blue) line to the right shows that the 284 exporting firms are under-represented at one star, but over-represented for three to five stars. Meanwhile, about half the non-exporters obtained only a single star and none of them obtained five stars. This superior quality of exporters is consistent with the quality interpretation of the Melitz (2003) model. However, there is no evidence of a quality cutoff above which all firms export and below which no firms export.[8]

Figure 3 (b) looks at the set of all 1134 exporting firms and the number of countries to which each firm exports. The lowest line of red squares corresponds to the 850 exporters that could not be matched to Juhlin listing. The lines to the upper right correspond to increasing numbers of stars. Each marked point on a line shows the fraction of firms that export to k or more countries. Thus each line is the empirical "survival function." Reading up from say $k = 10$ we see that two percent of non-Juhlin firms export to ten or more markets. However, over half the three-star and *all* eight of the five star producers export that widely.

From the descriptive statistics portrayed in Figures 2 and 3 we see that higher quality is associated with higher prices, greater likelihood of exporting, and a wider set of penetrated markets. The next section moves to regression evidence where the estimated parameters can be interpreted within the structure of a Melitz-based model.

[8]Appendix B.1 reports the regression equivalent of this figure, and shows the statistically significant impact of higher quality on the propensity to export (to any destination).

4 Baseline results

Table 2 reports estimates from our firm-level regressions corresponding to the equations (10), (12), and (14) of the model. The first column regresses log FOB prices on quality. Columns (2) and (3) show the marginal effects of quality on export probability using first a linear probability model (LPM) and then Probit. The LPM does not have the ancillary parameters problem associated with large numbers of fixed effects but the Probit has the attractive feature of restricting predicted export probabilities to lie between zero and one. OLS and Tobit estimations of the export value equation are reported in columns (4) and (5). The OLS has the same advantage as the LPM in terms of handling fixed effects but has the added disadvantage of inducing selection bias. The Tobit corrects that selection bias but imposes more structure on the data. Since our regressions map a single quality rating to multiple observations (destinations) for each exporter, we cluster standard errors by firm in all regressions.

The estimates reveal that higher quality tends to raise export prices, export probability, and export value as predicted in the model. Indeed, the quality-outcome relationship is monotonic in all five columns, with the one exception that 5-star producers have lower export values on average than 4-star producers in column (4). The one step down is not statistically significant but all the upward steps in the entry equations and the Tobit are significant at the 5% level. To facilitate interpretation and comparison between the LPM and Probit, we report the parameters in the entry equations as marginal effects. Thus, for example, a 4-star producer is nine percentage points more likely to export to any given country than 3-star producer according to the LPM and six percentage points higher for the Probit.[9] The coefficients in the price and Tobit equations are structural parameters that we interpret later in this section.

Consistent with the model's prediction, selection bias shrinks the coefficient on quality in the OLS export value regression shown in column (4). The bias arises because selection into exporting generates a negative correlation between the quality conditional on being selected and unobserved firm-country demand shocks. Comparing columns (4) and (5) confirms the direction and magnitude of this bias. The Tobit also eliminates the one exception to the monotonic relationship between quality and export outcomes: Five-star producers now export more than their four-star competitors (an F-test reveals that coefficients are statistically different at the 3% level). Using the parametric version of the Tobit estimator multiplies the OLS coefficient on quality by 3.5.[10]

[9]The baseline probability for 1-star producers naturally depends on the destination country.

[10]Our Tobit method relies on the minimum value of exports to each country d. As there might be noise in this estimate of $\sigma F_d/\tau_d$, we also experimented with the second lowest export value to each country as in

Table 2: Firm-level regressions for quality-rated Champagne exporters

	(1)	(2)	(3)	(4)	(5)
Dependent variable:	$\ln p_d^{\text{fob}}(j)$	$\mathcal{E}_d(j)$	$\mathcal{E}_d(j)$	$\ln x_d^{\text{fob}}(j)$	$\ln x_d^{\text{fob}}(j)$
Method	OLS	LPM	Probit	OLS	Tobit
Observations	3205	44586	44586	3205	44586
	Parametric				
ln stars	0.22^a	0.09^a	0.09^a	1.31^a	4.58^a
	(0.04)	(0.01)	(0.01)	(0.19)	(0.54)
ψ, std. dev. of $\ln \alpha_d(j)$					4.30^a
					(0.16)
R^2	0.24	0.27	0.32	0.23	0.62/ 0.15
	Non-Parametric				
2 stars	0.05^a	0.02^a	0.02^b	0.32	1.25^b
	(0.02)	(0.01)	(0.01)	(0.23)	(0.52)
3 stars	0.07^a	0.04^a	0.05^a	0.63^a	2.68^a
	(0.03)	(0.01)	(0.01)	(0.23)	(0.55)
4 stars	0.20^a	0.13^a	0.11^a	1.99^a	5.80^a
	(0.03)	(0.03)	(0.02)	(0.34)	(0.79)
5 stars	0.52^a	0.26^a	0.16^a	1.67^a	7.70^a
	(0.14)	(0.03)	(0.02)	(0.23)	(0.59)
ψ, std. dev. of $\ln \alpha_d(j)$					4.19^a
					(0.16)
R^2	0.32	0.29	0.33	0.26	0.63 / 0.17

Note: Destination (d) fixed effects for all columns. Column (3) reports marginal effects of the probit estimation. R^2 include country dummies. For columns (3) and (5), R^2 are computed as the squared correlation between the predicted and actual values of the dependent variable. Second R^2 in column (5) uses the same sample as column (4). Standard errors clustered at the firm-level in parentheses. Significance levels: ^{c} $p < 0.1$, ^{b} $p < 0.05$, ^{a} $p < 0.01$

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 yields 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 2.2.

Table 3: Simulation results (assumed true $\beta = 4.58$)

Variable	mean	std. dev.
Share of profitable firm-destination exports	0.072	0.003
Correlation(s, α) in censored data	-0.394	0.038
OLS $\hat{\beta}$ before censoring $x_d^{\text{fob}}(j)$	4.580	0.044
OLS $\hat{\beta}$ after censoring $x_d^{\text{fob}}(j) < \sigma F_d/\tau_d$	1.330	0.160
Tobit $\hat{\beta}$ (estimate $\sigma F_d/\tau_d$ with min $x_d^{\text{fob}}(j) > 0$)	4.518	0.213
Contraction: censored OLS $\hat{\beta}$/Tobit $\hat{\beta}$	0.295	0.033

The simulation comprises 1000 firms and 10 countries.[11] The first step is to generate a random set of $s(j)$, A_d/τ_d, and $\alpha_d(j)$, with which to create the uncensored vector of $\ln x_d^{\text{fob}}(j)$ based on equation (13). We specify the "true β," as 4.58, the estimate from the parametric version of column (5) of Table 2. The second step obtains the censored sample by imposing the condition that gross profits exceed fixed costs, which holds when

$$\ln x_d^{\text{fob}}(j) > \ln(\sigma F_d/\tau_d). \tag{16}$$

We choose the parameters of the A_d/τ_d (entering $x_d^{\text{fob}}(j)$ in (16) as can be seen in (13)) and $\sigma F_d/\tau_d$ (the RHS of (16)) distributions such that the share of firm/destination profitable combinations replicates the share we observe in our empirical sample (7.2%).[12] Then we calculate the correlation between s and α after censoring.

The simulation results for 1000 repetitions are summarized in Table 3. The second row shows that censoring induces a strong negative correlation between firm quality and the error term of -0.394. The next row establishes the unsurprising result that OLS estimates $\hat{\beta}$ correctly for the complete sample. After censoring, however, OLS regressions yield an

Eaton and Kortum (2002). Results (shown in the appendix) are very similar.

[11]Stata code provided on the authors' webpage (`http://strategy.sauder.ubc.ca/head/sup/`).

[12] Table 2 shows that over the 44,586 possible combinations, only 3,205 are positive.

average coefficient of 1.33, which is remarkably similar to the coefficient shown in column (4) of Table 2.

The fifth row of Table 3 shows the results for the Tobit method we use. It estimates the censoring point using the minimum observed trade value: $\hat{\underline{x}}_d^{\text{fob}} = \min_{j\in\mathcal{B}_d} x_d^{\text{fob}}(j)$. Although this Tobit exhibits some downward bias ($4.52 < 4.58$), it corrects almost all the bias found in the OLS. The final row shows the contraction in the OLS estimates relative to Tobit. In the simulation, the $\hat{\beta}$ from OLS on censored data is 0.295 of the Tobit estimate, while our estimates from the real data put this contraction at 0.286. Additional simulations, reported in Appendix B.2 investigate departures from the normality assumption for $\ln \alpha$. We find that Tobit estimates remain within 3% of the true value and OLS contraction factors are comparable. All in all, the simulations make us confident that the 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, the Tobit method, which is the maximum likelihood estimator of the model, should prove useful in other applications of the Melitz model.

In discussing the impact of selection on the export value equation, we have emphasized the parametric specification for convenience. However, it is important to note that the non-parametric method of estimation is to be preferred on several grounds. First, the elasticity concept is intended for continuous variables, not the discrete quality ratings we observe. Second, the quality of fit is systematically larger in the lower part of the regression table, especially when explaining the price charged by firms endowed with different quality levels. This comes from the fact that imposing a parametric structure seems to do too much violence to the data. The coefficients on prices, entry probability and exported values vary a lot according to the level of s. Taking prices as an example, increasing quality from 1 star to 2 stars raises prices by just 5% whereas the parametric version imposes a $2^{0.22} - 1 = 16.5\%$ increase. Going to 5 stars is associated with a 68% non-parametric increase, and a 42% parametric one. Thus, the parametric form seems unable to handle the convexity in the relationship between prices and quality.[13] We will therefore focus on the non-parametric estimates for the rest of the paper.

As discussed in subsection 3.2 and the data appendix, there is concern over another potential selection problem. This comes not from the model but the way the data are collected. Briefly, it is possible that some intra-EU exports by small firms are not included because they lack destination information. Table 11 in the appendix reproduces the specifications

[13] This finding is confirmed when adding a new (low) quality category of exporters producing in Champagne but not rated by Juhlin. Details about the creation of this low quality ($s(j) = 0$) category are provided in Appendix B.4 along with the reproduction of the baseline results incorporating these firms.

in Table 2 with the quality indicators interacted with a dummy for intra-EU flows. Most coefficients on the EU interactions are insignificant, suggesting that quality has the same impact on export outcomes inside and outside the customs union.

A second potential concern with our results is reverse causation. Suppose that Juhlin's ratings are influenced by how successful a wine has been in export markets outside France. In that case, exports would cause high quality ratings rather than the line of causation we have presumed. To some extent this critique cannot be ruled out since no blind tasting of the 500 producers has ever been undertaken. It is not obvious which direction the endogeneity bias would run. Some experts appear to take pride in giving high quality ratings to obscure, hard-to-find producers. This would tend to bias the coefficients in the value equation downwards. In any case, we doubt that there are strong firm-level instruments for quality ratings that satisfy the exclusion restriction.[14]

Our response to the endogeneity concern is to investigate the robustness of our results to alternative ratings that we see as being more or less likely to exhibit the reverse causation. It is natural to expect that an internationally recognized reviewer based in the United States would be more focused on the American market and therefore we hypothesize that Parker would be more likely to give high ratings to producers that sell abroad, especially in the US. Correspondingly, the RVF guide is intended for French consumption and therefore its tasters might pay less attention to the presence of Champagnes in foreign markets. Thus, we expect the largest, positive endogeneity bias to be found for Parker ratings and the smallest, negative bias to be found for RVF.

Appendix Tables 12 and 13 re-estimate Table 2 using ratings from Parker and RVF instead of Juhlin. The main impression we draw from these results is how similar they are to each other. For example, the coefficent on 5-star Champagne in the export value Tobit equation (column 5) is 7.37 for Parker, 7.24 for RVF, and 7.70 for Juhlin.[15] The similarity does not, of course, rule out a common endogeneity bias. However, it does suggest that there is a common factor underlying the quality ratings of different tasters and that factor has a robust relationship with prices and export success.

The structure embedded in the estimating equations allows us to infer the key parameters of the model and thereby obtain a precise quantification of the impact of quality. Recall that equation (9) shows that the coefficients in on the quality dummies in the price equations can

[14]Geographic and weather-based variation in quality—which might serve as potential instruments for other types of wine—are not applicable for Champagne given the practices of blending wines from multiple vineyards and years (Johnson and Robinson, 2005, pp. 78–79).

[15]As in Figure 1, we map RVF stars to the 5 star systems used by Juhlin and Parker as follows: $3 \rightarrow 5$, $2 \rightarrow 4$, $1 \rightarrow 3$, included $\rightarrow 2$ and excluded $\rightarrow 1$. Since Parker never assigns a single star, we use that to designate firms that are in Juhlin but not included in Parker.

Figure 4: Structural interpretation of coefficients

(a) prices: the cost of quality (b) values: the benefit-cost ratio

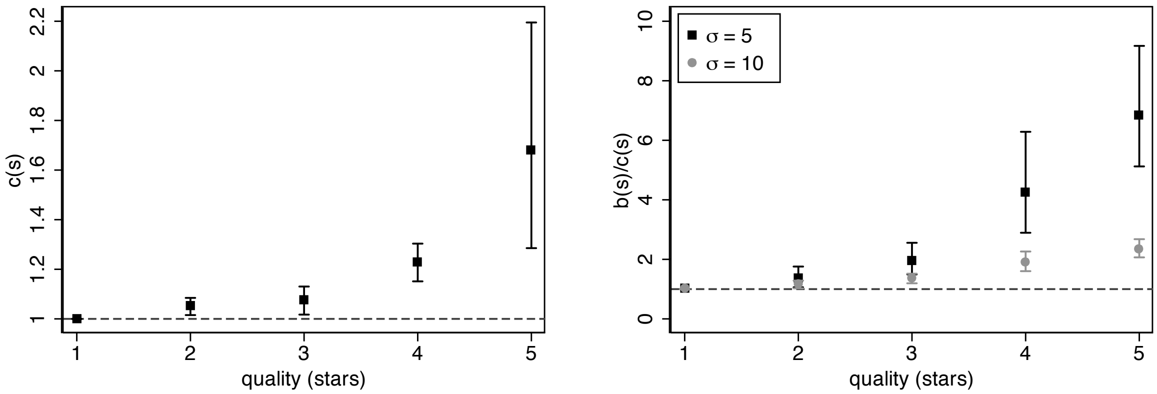

be interpreted as estimates of $\ln c(s)$, with costs for $s = 1$ normalized to zero. Exponentiating, we then see how quality raises the marginal costs of producing Champagne. Figure 4(a) reports the exponentiated coefficients from column 1 of Table 2 accompanied by their 95% confidence intervals, $\exp(\hat{\lambda}_s \pm 1.96\text{s.e.}_s)$. The estimates suggest that 5-star Champagne has 68% higher costs of production than 1-star Champagne.

Figure 4(b) depicts the structural interpretation of the coefficients from the export value equation (column 5 of Table 2). Equation (13) implies that each coefficient estimates $\beta_s = (\sigma-1)\ln b(s)/c(s)$. Rearranging, we obtain $b(s)/c(s) = \exp(\beta_s/(\sigma-1))$. This can be thought of loosely as the benefit to cost ratio for s-star wine compared to one star wine. Even more loosely, we can speak of it as the measure of whether a wine "delivers good value for money." A high σ implies that consumers find different producers to be highly substitutable. Thus a high assumed σ allows for a large impact of quality on exports (high β_s) with relatively low benefit to cost ratios. We show implied $b(s)/c(s)$ for σ assumptions of five (black) and ten (gray), the lower and upper bounds of the range that Anderson and van Wincoop (2004) consider to be consistent with the literature. A b/c ratio of nearly seven for a five-star Champagne producer is implied for $\sigma = 5$ but this falls to two for $\sigma = 10$.

By multiplying the value, $b(5)/c(5)$, and price, $c(5)$, estimates, we can back out the implied benefits consumers receive from quality. Assuming $\sigma = 5$ the $b(5)/c(5) = 6.86$ and $c(5) = 1.68$. Hence, the results imply that a consumer is willing to trade about 12 bottles of the lowest quality ($s = 1$) Champagne for one bottle of the highest quality ($s = 5$). This is also the ratio of prices between a five-star and a one-star bottle that would leave a consumer's indirect utility unchanged.[16] These impacts of quality strike us as economically large but within the plausible range. The underlying coefficients are estimated within a highly parsimonious structure that stays as close as possible to the quality-interpreted Melitz model. This raises the question of whether such high magnitudes would survive generalizations of the model that relax potentially problematic assumptions. Through all the modifications to the baseline model explored in the next section, we find the impact of quality on prices and exports remains remarkably stable.

5 Extensions

In this section we consider three extensions to our baseline model. The first incorporates additional sources of firm heterogeneity, in particular differences in labour productivity. The next two generalize functional form assumptions, allowing for non-iceberg trade costs and then non-homothetic preferences for quality.

5.1 Other sources of firm heterogeneity

The baseline model abstracts from any source of heterogeneity other than quality. However, most empirical work on firm heterogeneity follows the productivity interpretation of the Melitz (2003) model. The assumption of a single factor determining firm-level heterogeneity is grounded in analytical convenience, not realism. In the case of Champagne, anecdotal accounts suggest that marketing efforts by the major brands could play an important role. Unfortunately, no direct measure is available for the marketing efforts of the whole set of Juhlin firms. Only imperfect data are available for their physical productivity. Taking these data limitations into consideration we investigate the role of other types of heterogeneity in two steps. First, we compare the share of the variation in prices and export values that can be explained by quality dummies versus the explanatory power of a full set of firm effects, which capture all firm-level performance determinants. Second, we calculate the share of the variance in firm fixed effects that can be explained by our measures of firm-level quality and physical productivity.

Table 4 decomposes the variance for the main variables of interest, the price and value of exports. It establishes how quality ratings can explain variation in prices and export values compared to other explanations. The explanatory variables are $\ln p_d^{\text{fob}}(j)$ and $\ln x_d^{\text{fob}}(j)$ as in columns (1) and (5) of Table 2. The first row presents the R^2 of a regression including

[16]The CES indirect utility is y_d/P_d. Indirect utility holds constant while s rises if and only if $(p_d(i)\tau_d)/(\alpha_d(i)b[s(i)])$ remains unchanged. For the same draw of α, this is true when $p_d(5)/p_d(1) = b(5)/b(1)$.

Table 4: Explanatory power of quality

Dependent variable:	$\ln p_d^{\text{fob}}(j)$		$\ln x_d^{\text{fob}}(j)$	
Model Criterion:	R^2	Δ AIC$_c$	R^2	Δ AIC$_c$
Country dummies only	0.151	1315.1	0.556	8352.5
Country dummies + quality dummies	0.317	627.5	0.629	5557.5
Country dummies + firm dummies	0.546	0	0.830	0

Note: Δ AIC$_c$ measures information loss relative to the best approximating model.

country dummies only, while the second one adds the Juhlin quality dummies exactly as in the baseline regressions, and the last one replaces quality dummies with firm effects to show how much of the variance can be explained by all sources of firm heterogeneity. The R^2 for Tobit is calculated as the squared correlation between actual trade and predictions using the Tobit regression equation.[17] Since the firm dummies require estimation of 284 parameters rather than the five quality levels, we also should consider the costs of adding extra parameters. The Akaike Information Criterion (AIC) provides an appealing metric since it measures the information loss relative to full reality taking into account the benefits of improved fit with the cost of greater variance in the estimates. We follow the suggestion of Burnham and Anderson (2002) and present Δ_i AIC$_c$ which corrects the AIC for having relatively small number of observations per parameter and express results as the difference between the AIC$_c$of model i and that of the best approximating model.

Table 4 reveals that firm-level effects have substantial explanatory power as one would expect since they collect all systematic differences between firms. The Juhlin quality indicators take the R^2 almost halfway to this upper bound for variables that vary only at the firm level. According to the AIC the best approximating model has firm-level fixed effects. The information loss from using observed quality in lieu of the firm dummies is considerable, even after penalizing the firm effects model for raising the number of parameters from 5 to 284. Nevertheless, starting from a model with only country dummies, the quality dummies reduce the information loss by about one half for price and one third for export values.

In Table 5 we unpack the variation of the firm fixed effects into the contributions of quality and physical productivity. The fixed effects for price are the dependent variable in the left panel and those for export values (Tobit) are used in the right panel. Columns (1) and (4) show the results of regressing the 284 estimated firm effects on the quality dummies. The R^2 indicate that quality explains 36% of the firm-level component of log prices and 27% of the firm level component of log export values. A single dimension, quality, accounts for

[17] The Stata option "ystar()" is used to obtain predicted trade over the range beginning at the minimum observed value.

Table 5: Firm fixed effects, quality and labour productivity

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Fixed effects		ln price (fob)				ln export value		
Observations	284	126	126	126	284	126	126	126
2 stars	0.03^c	0.03	0.04^c		0.83^b	1.37^b	1.20^b	
	(0.01)	(0.02)	(0.02)		(0.36)	(0.56)	(0.59)	
3 stars	0.05^b	0.03	0.04		2.17^a	2.27^a	2.21^a	
	(0.02)	(0.04)	(0.04)		(0.42)	(0.67)	(0.68)	
4 stars	0.20^a	0.22^a	0.22^a		3.81^a	5.17^a	5.13^a	
	(0.04)	(0.04)	(0.04)		(0.66)	(0.99)	(0.99)	
5 stars	0.53^a	0.51^a	0.48^a		7.33^a	6.94^a	7.34^a	
	(0.11)	(0.13)	(0.13)		(0.64)	(0.50)	(0.52)	
ln labour			-0.03^b	-0.06^a			0.54^b	0.26
productivity			(0.02)	(0.02)			(0.26)	(0.28)
Constant	0.09^a	0.11^a	0.46^a	0.78^a	0.85^a	1.81^a	-3.85	1.28
	(0.01)	(0.02)	(0.16)	(0.22)	(0.23)	(0.41)	(2.70)	(3.00)
R^2	0.358	0.446	0.472	0.075	0.268	0.362	0.383	0.005
Δ AIC$_c$		3.95	0.00	62.28		2.00	0.00	51.64

Note: Δ AIC$_c$ measures information loss relative to the best approximating model. Heteroskedasticity-robust standard errors in parentheses. Significance levels: ^{c} $p < 0.1$, ^{b} $p < 0.05$, ^{a} $p < 0.01$

about one third of all firm-level heterogeneity.

These results raise the question of whether a second observed dimension of firm ability might account for an even higher share of the variance. Of much greater concern for the central premise of the paper, observed quality could simply be a surrogate for some other variable. Given that we have motivated this paper as an illustration of the quality interpretation of Melitz (2003), the obvious second dimension to consider is physical productivity. Given limitations of the data, we do not attempt any sophisticated productivity estimation. Rather, we measure labour productivity as the number of Champagne bottles (750ml) produced per year per employee at each company. Of the 284 exporting companies for which we have quality, Juhlin (2008) provides annual production for all but three.[18] Since Juhlin (2008) does not supply employment data, we use the Enquête Annuelle d'Entreprise (EAE), which covers firms larger than 20 employees, and INSEE's SIREN database which covers some small firms but only provides employee ranges.[19] These sources provided employment data for just 126 firms. Although this low coverage rate is a concern, the firms for which we do have both quality and physical productivity measures represent 96% of the total exported value by quality-rated firms (and 90% of total Champagne exports). In addition to the coverage problem and the need to interpolate for the employee range data, we also suspect that some of the employees at a given firm may not be involved in wine production and therefore should not be included in the denominator for productivity. While we have done the best we could with available data, the labour productivity variable probably suffers from considerable measurement error.

Columns (2) and (5) of Table 5 conduct the export price and value regressions of columns (1) and (4) for the sample of firms for which we also measure labour productivity. Columns (3) and (6) add productivity to the set of covariates while holding the sample the same. This hardly affects the coefficients on the quality dummies, suggesting that the two forms of heterogeneity are uncorrelated. This would imply that the main cost of higher quality is that it requires the use of higher priced inputs rather than a greater quantity of labour per bottle produced. The inclusion of physical productivity only nudges the R^2 up slightly in both the price (from 0.45 to 0.47) and value (from 0.36 to 0.38) equations. The coefficient on log physical productivity in the price equation is much closer to zero than the theoretical prediction of -1. In the export value equation, the coefficient should theoretically estimate $\sigma - 1$. The implied σ of 1.54 lies at the low end of the plausible range. The apparent bias in these coefficients away from the model predictions and towards zero is the expected con-

[18] Juhlin (2008) reports only a single production number for each listing and he does not assign a particular year. His production numbers appear to be rounded so we interpret them as some kind of average amount in recent years.

[19] See Appendix for more detail on the data and how we handled range data.

sequence of measurement error in physical productivity. Nevertheless, the results dispel one alternative explanation of our baseline estimates. High quality ratings are not just "awards" that wine reviewers bestow on firms who succeed abroad due to underlying differences in physical productivity.

Columns (4) and (8) of Table 5 remove the quality variables and attempt to explain firm fixed effects only with physical productivity. The R^2 for both price effects and export value effects fall sharply. We again use the Δ_i AIC_c to evaluate the different models. Adding productivity to the model with quality dummies only slightly reduces the information loss. However, eliminating quality and using only productivity would result in very large information loss.[20]

5.2 Non-iceberg transport costs and Alchian-Allen effects

Actual shipping fee structures are based on the volume of wine shipped rather than the value, a fact that casts doubt on the applicability of the iceberg assumption of ad valorem trade costs. Yet there are other trade cost that are likely to be proportional to value such as insurance, loss of reputation due to spoilage, and non-payment risk. We therefore consider a more general specification of trade costs such that CIF price $p_d(j)$ includes both an ad valorem component, τ_d, and a per unit freight charge, f_d. The CIF-FOB price relationship is given by

$$p_d(j) = p_d^{\text{fob}}(j)\tau_d + f_d. \tag{17}$$

Substituting into the export value equation expressed in FOB terms, $x_d^{\text{fob}}(j) = p_d^{\text{fob}}(j) \times x_d(j)/p_d(j)$, and using the CES demand system from (2), we obtain

$$x_d^{\text{fob}}(j) = p_d^{\text{fob}}(j)[p_d^{\text{fob}}(j)\tau_d + f_d]^{-\sigma} b[s(j)]^{\sigma-1} X_d P_d^{\sigma-1} \alpha_d(j)\mathcal{E}_d(j). \tag{18}$$

The Alchian-Allen hypothesis stipulates that the relative demand for high quality firms rises with the non-iceberg part of the transport cost. This is because relative demand for two different quality levels depends on their relative CIF prices. The ratio of the CIF price of high quality wine to low quality wine starts as a number greater than one for $f_d = 0$ and gradually goes to one as $f_d \to \infty$. Hence, relative demand for high quality is expected to increase with f_d. Analysis of this effect under monopolistic competition is complicated because, for $f_d > 0$, profit maximization implies pricing to market, i.e. the FOB price to country d depends on f_d and τ_d. We investigate Alchian-Allen effects in our model by first

[20] Burnham and Lapham (2002, p. 70) deem Δ_i in excess of 10 to show that the level of support for model i is "essentially none."

deriving the elasticity of FOB prices with respect to quality and showing how it depends on f_d. We then do the corresponding calculations for FOB export values, conditional on entry.

The profit-maximizing FOB price is as in Martin (2010, eq. 8),

$$p_d^{\text{fob}}(j) = \frac{\sigma c[s(j)]}{\sigma - 1} + \frac{f_d}{\tau_d(\sigma - 1)}. \tag{19}$$

Using (19), the price elasticity can be shown to be

$$\eta_{ps} = \frac{\eta_{cs}}{1 + \frac{f_d}{\tau_d \sigma c[s(j)]}}, \tag{20}$$

where η_{cs} is the elasticity of costs with respect to quality (λ in the power form of $c(s)$). Inspection of this equation reveals that η_{ps} begins at η_{cs} for $f_d = 0$, and falls to 0 as $f_d \to \infty$. This implies a negative interaction between quality and distance in the price regression, assuming that distance has a larger effect on the per unit trade cost than on the ad valorem cost.

The elasticity for FOB export values (18, maintaining $\mathcal{E} = 1$) is

$$\eta_{xs} = (\sigma - 1)\left[\eta_{bs} - \left(\frac{1}{1/\sigma + \frac{f_d}{\tau_d \sigma c[s(j)]}} - \frac{1}{1 + \frac{f_d}{\tau_d \sigma c[s(j)]}}\right)\frac{\eta_{cs}}{\sigma - 1}\right]. \tag{21}$$

For $f_d = 0$, the term in parentheses simplifies to $\sigma - 1$ and the whole elasticity reduces to be $\eta_{xs} = (\sigma - 1)[\eta_{bs} - \eta_{cs}]$. Imposing the power function specification for $b[s]$ and $c[s]$, the result reverts back to $\beta \equiv (\sigma - 1)[\gamma - \lambda]$, as shown in section 2. As f_d becomes large, the term in parentheses in equation (21) goes to zero, and η_{xs} rises to $(\sigma - 1)\eta_{bs}$. This implies a positive coefficient on the interaction between quality and distance in the export value regression, maintaining the assumption that distance increases f_d/τ_d.

In the empirical implementation, we take the non-parametric price and export value equations of Table 2, and add a complete set of interactions between the quality indicators and dummies for six distance intervals.[21] Figure 5 displays the coefficients for price (left panel) and export value (right panel) for the interactions between quality and the different distance bins. We also add horizontal lines indicating the baseline coefficients from Table 2, i.e. not allowing for Alchian-Allen effects. The symbols are circled in red when the distance-specific estimate is significantly different from the baseline one.

[21] We chose the same distance cuts to determine the dummies as Eaton and Kortum (2002): 375, 750, 1500, 3000, and 6000 miles. Alternative specifications using the fully parametric version (ln quality times ln distance) or semi-parametric (quality dummies times ln distance) are reported in the appendix. The fully parametric test gives a mistaken impression of generality for the AA effect which the semi-parametric form shows to be insignificantly different from zero at most quality levels.

Figure 5: Alchian-Allen effects

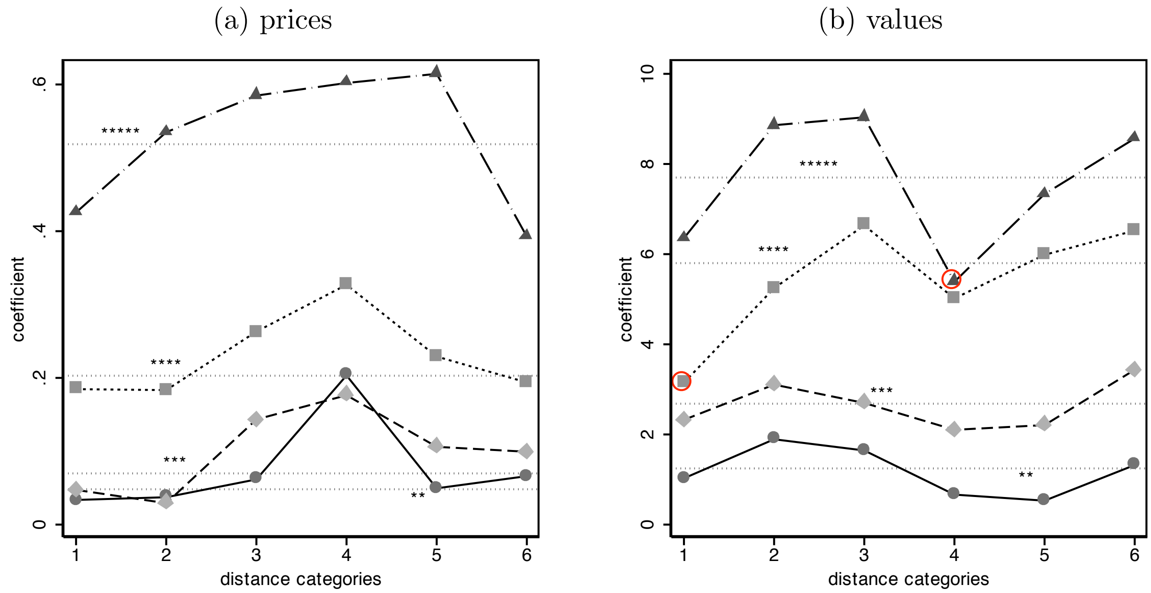

Both panels show that the vast majority of coefficients interacted with distance bins do not differ from the baseline significantly. Four-star Champagne seems to exhibit an increasing advantage over lower quality Champagne with distance, as predicted. However, a significant difference from the constrained average effect for 4-star Champagne is only present for the lowest of the six distance categories. Moreover, the convergence of the four and five star export values as distance increases conflicts with the Alchian-Allen prediction. In sum, Alchian-Allen effects are not entirely absent, but they have little impact on the main finding: higher quality generally increases FOB prices and export values over all distances.

5.3 Non-homothetic preferences

The base model follows Melitz (2003) in assuming homothetic preferences. Relaxing this assumption is desirable in light of recent research emphasizing income effects on the relative demand for quality. For example, Fajgelbaum, Grossman, and Helpman (2009) construct a model of consumer choice in which the share of consumers that choose high quality products is rising in income. Since Champagne is a canonical example of a luxury good, it seems unrealistic to force demand to be proportional to income across all quality levels. The extension is very simple: we replace $b(s(j))$ with $b(s(j), y_d)$ where y_d is per-capita income. Since we do not have strong priors on the functional form for the new $b()$ function we estimate it non-parametrically using six equal-sized income categories. The income thresholds for

Figure 6: Luxury effects (income effects on demand for quality)

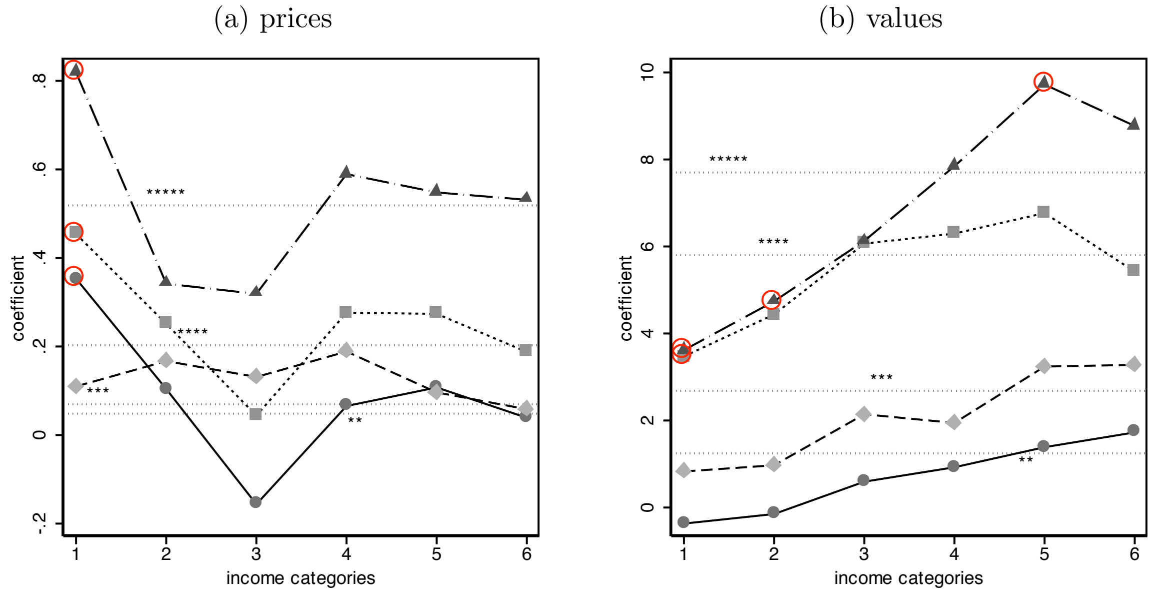

these sextiles are \$645, \$1741, \$3771, \$7943, and \$25914. For additional parallelism with the non-iceberg subsection, we estimate the interaction-saturated model for both log prices and log export values (Tobit).

The results are shown graphically in Figure 6. As in Figure 5, all results are relative to the benchmark of $s = 1$. The price regression results displayed in panel (a) can be seen as investigating whether higher income countries have lower σ and hence higher markups over marginal cost. To the extent that markups rise with income we would expect the profiles for $s = 2..5$ to slope up across the income categories. Instead, we find that the only significant differences from the baseline specification are higher prices for the lowest income category for quality levels $s = 2$, 4, and 5. One possible interpretation is that the low-income markets have so few Champagne producers that the firms that do enter charge higher prices.[22]

Turning to the panel (b), the value of Champagne exports for producers with two or more stars tends to rise with income relative to that for just one star. The gaps are generally widening as would be predicted if higher quality is a luxury. For example, in the top income sextile the gap between five and three stars is $8.8 - 3.3 = 5.5$, which greatly exceeds the $3.6 - 0.8 = 2.8$ gap in the first income sextile. The "fanning out" we see in panel (a) seems to be demand-related given the absence of corresponding price patterns in panel (b). Even

[22]This interpretation is at odds with the Dixit-Stiglitz assumption of a constant markup regardless of the amount of competition. We leave the investigation of variable markups to future work.

for the lowest income categories, there are no significantly negative coefficients associated with increased quality. Thus even the poorest countries do not systematically invert demand rankings.[23]

6 Fit to the aggregate data and falsification exercises

The main prediction of the quality sorting model is that only the best firms can profitably enter the most difficult markets. Empirically, this corresponds to a prediction that the larger the number of exporters to a destination, the lower will be the average quality (and, therefore, average price) of these exporters.[24] While probit results confirm that quality does indeed explain selection of individual firms, they cannot illustrate the aggregate effects of quality sorting. In this section, we use the estimated parameters derived from the firm-level regressions to show that the quality sorting model can replicate the main features of the aggregate (market-level) data. We also contrast the fit of the quality sorting model to data with the predictions of two extreme alternatives. The first falsification exercise predicts average quality in a process where entry is entirely random. The second eliminates all randomness at the firm-market level, and stipulates that selection depends exclusively on measured quality.

We therefore establish the prediction of the quality sorting model in terms of average quality and average price to each destination market. To do so, one needs to establish which firm succeeds in exporting to each country. One method is to derive the quality cut-off for Champagne producers in each market d, and average the qualities of firms above this cut-off. However, the entry threshold for each market d is an endogenous variable depending non-linearly on attributes of the market and both the distributions of quality and of the idiosyncratic demand shocks. In order to derive analytical expressions of the cut-off, one has in particular to assume a particular distribution for quality draws. While the use of Pareto is almost universal in the literature, it might not be either innocuous for predictions or precisely reflecting the true distribution of quality. Fortunately, it is not necessary for us to solve for this cut-off if we take as given the actual outcome in terms of N_d, the number of French Champagne exporters to every country d in the world in 2005.

The quality sorting model specifies that the N_d most competitive firms will be the ones to

[23]To the extent that consumption of Champagne is confined to the top of the income distribution, average incomes may be misleading proxies for luxury effects. In results available upon request, we have reproduced Figure 6, defining sextiles based on the average income of the top decile. Due to the high correlation between the two measures, the patterns displayed in the figures are very similar.

[24]Note that Melitz (2003) uses changes in the equilibrium average performance variable, productivity in his case, to perform the main comparative statics of his model.

export to market d. The ranking in terms of competitiveness in the baseline model depends on the sum of the quality term, $(\sigma - 1)\ln(b[s(j)]/c[s(j)])$, and the idiosyncratic demand shock $\ln \alpha_d(j)$. A firm with quality s has an estimated $(\sigma - 1)\ln(b[s]/c[s])$ given by the $\hat{\beta}_s$ from column (5) of Table 2. The model specifies $\ln \alpha_d(j)$ to be drawn from a normal distribution with mean 0 and standard deviation ψ. The firm-level Tobit estimates $\hat{\psi} = 4.19$ such that the relative importance of randomness in our simulations is calibrated from the firm-level data. Given this sum for each potential entrant j, we sort them by decreasing competitiveness. It is then straightforward to average the quality of the first N_d firms. To generate average prices of the N_d selected firms, we average over the sum of the country d price fixed effect (PFE_d) and the estimate of $\lambda_s = \ln c[s]$ from column (1) in Table 2.

Figure 7 provides results from our simulation. The two top graphs plot average quality against count of exporters, while the lower graphs depict average price. Red and black dots represent data and simulation respectively; the lines are kernel smoothers. The two left graphs use our full sample of 157 destination countries, while the right ones zoom in on the set of 35 countries with 20 or more exporters. Although less than a quarter of the sample, these markets account for 96% of the exports of Juhlin-rated firms (90% of all Champagne exports). For countries that received few entrants, the averages exhibit huge dispersion in average quality and price. This is exactly what the model predicts given the large estimated standard deviation of idiosyncratic demand. For the countries who consume the bulk of exported Champagne, there is a much tighter, negative relationship between average quality and the number of entrants. The relationship between average price and number of entrants is much flatter. This is also in line with the model's prediction since the elasticity of price with respect to quality is just 0.22. In the zoomed region, the kernel smoothers for predicted and observed average quality lie almost on top of each other, demonstrating the close fit of the baseline model to the data.

Is the good fit to the data exhibited in Figure 7 unique to the quality sorting model? A falsification exercise can answer this question by comparing the data to the predictions of two polar alternative entry processes. The first alternative is a very simple departure from the baseline model. It excludes the random firm-destination term in the model, which means that observed quality becomes the sole determinant of which firms enter each market. The black circles at the top of Figure 8 describe what happens in this simulation. Entry follows a strict, hierarchical ordering. In very difficult markets, with few entrants, the average quality is level five, which comes from the fact that there is strict quality sorting in this model, and Juhlin rated just eight firms with five stars. As soon as the number of entrants exceeds 8, firms with quality level 4 become exporters, causing average quality to fall monotically, until a minimum mean quality is attained for the country where 195 out of 284 firms export.

Figure 7: Fit of data to the quality sorting model

(a) Quality: all countries

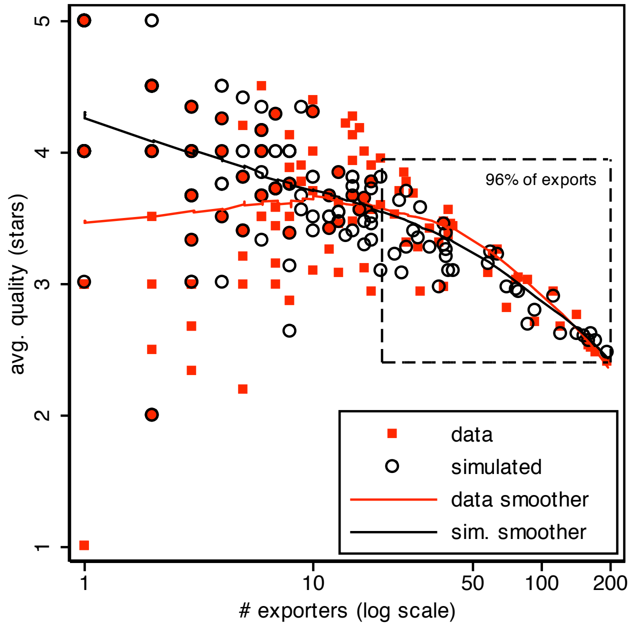

(b) Quality: 20+ entrants

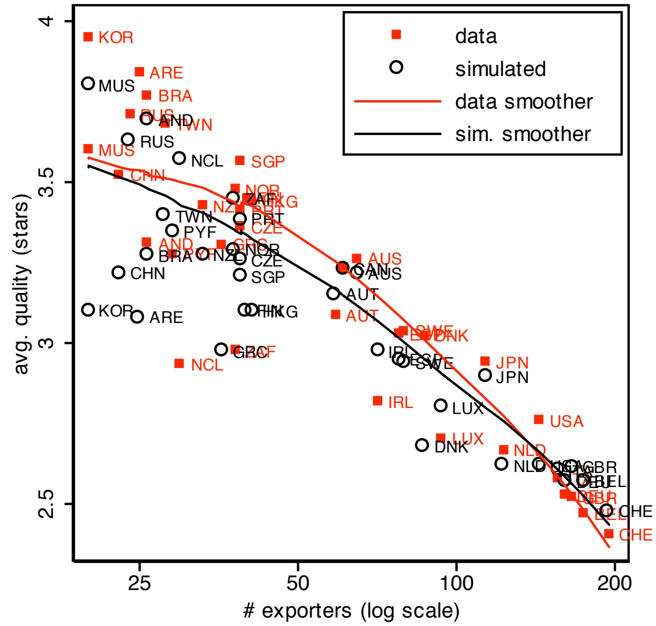

(a) Prices: all countries

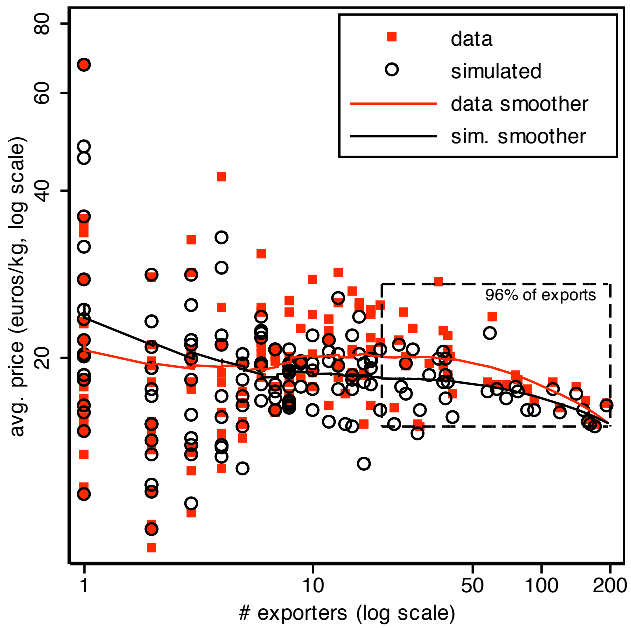

(b) Prices: 20+ entrants

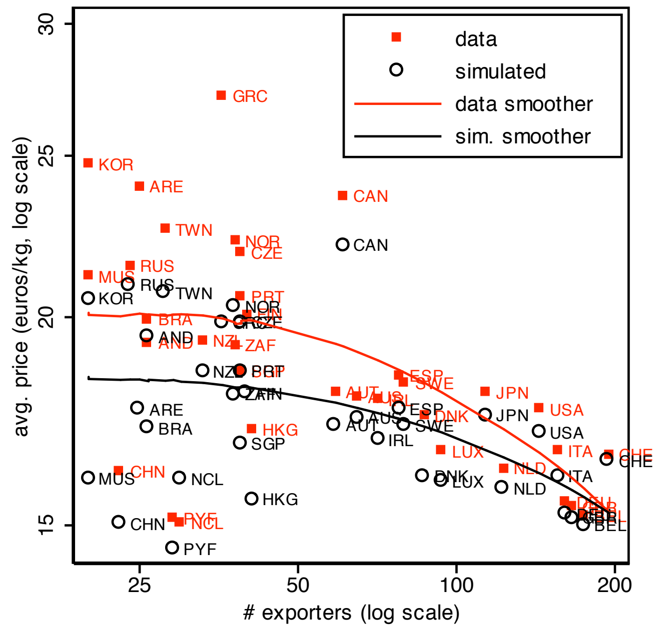

Figure 8: Fit of data to two alternative sorting models

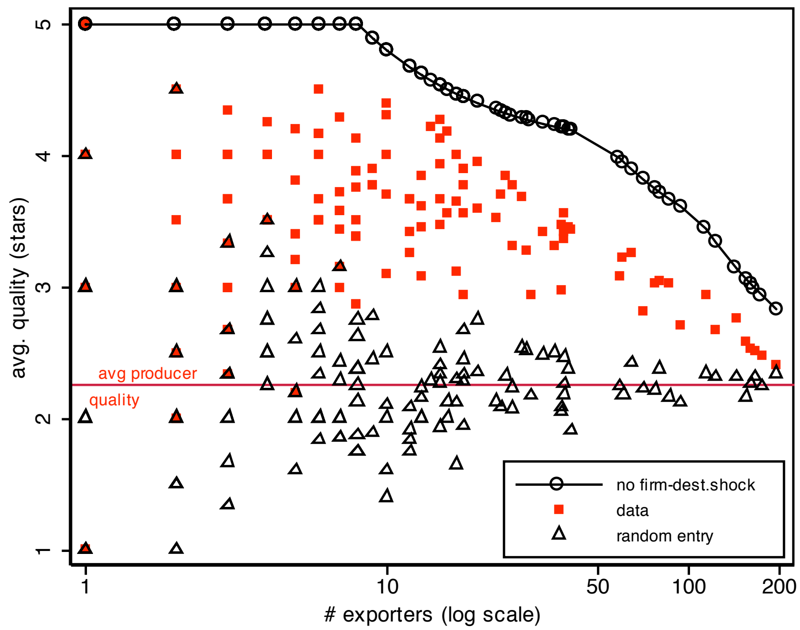

The second simulation takes the opposite extreme and attributes all firm-level heterogeneity to the $\ln \alpha_d(j)$ term, and none to observable quality. This is a random entry model, represented by triangles in Figure 8. For this simulation, predicted average quality is naturally very noisy when the number of draws is small (in countries with small N_d on the left of this figure), and converges to the mean quality of the 284 firms in our sample, around 2.2, when N_d becomes large enough. The squares representing the data make it clear that the alternative entry processes cannot be reconciled with the actual pattern of average quality across destination markets. While the model putting all weight on observable quality exhibits a decreasing relationship that is in the data, it is overestimating average quality in general, and has a particularly bad fit for small destination countries. The random entry model is not even capable of replicating a downward sloping curve between average quality and ease of market entry. Note that the prediction of random entry only approximates reality for very easy countries, where two thirds of the potential exporters enter in reality. This is because average quality of entrants in that case is converging towards the random entry prediction of the population mean.

In sum, Figure 8 reveals that the true entry process appears as a convex combination of the polar alternatives. Figure 7 shows that by calibrating the combined model using parameters from the firm-level regressions, a very close fit can be obtained despite the strong assumptions made on transport costs and preferences in the baseline model.

7 Conclusion

We have illustrated the importance of quality sorting for trade by examining an industry where we could obtain direct measures of quality. Adapting the Melitz (2003) model, we derive three estimable equations for FOB pricing, market entry, and export value. 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. Our estimates imply that the ratio of consumer benefits to producer costs for the highest quality Champagne (5 stars) is two to seven times higher that the ratio for the lowest and most common category (1 star).

Industry information suggests that Champagne provides a reasonably good fit to the Melitz (2003) assumptions of monopolistic competition and firm-level heterogeneity. Two potentially problematic assumptions of the model are iceberg (ad valorem) trade costs and homothetic preferences. Relaxing those assumptions in the extensions section, we find that the fit at the level of individual firm behavior can be improved somewhat by allowing for per-unit freight costs and income-based preferences. However, the results are somewhat mixed and do not present a compelling case for either generalization. This tentative conclusion is reinforced when we simulated the baseline model and found that it does an excellent job at explaining the average quality and price data across markets, suggesting that the extensions are not important for explaining the aggregate patterns in quality and price across markets.

The extension to the Melitz model that clearly matters a great deal—which we include in the baseline—is firm-market demand heterogeneity. With only firm-level heterogeneity, the model predicts a starkly counter-factual entry pattern. Our results therefore align well with a diverse set of findings by Eaton et al. (forthcoming), Bernard et al. (2010), and Foster et al. (2008) pointing to the importance of idiosyncratic differences in demand to explain market outcomes. Producer-level measured quality determines average outcomes very well for all markets that large numbers of Champagne exporters entered. For smaller markets the vagaries of demand play a large role.

8 References

Ahn, J. A. Khandelwal, and S-J. Wei, 2011, The Role of Intermediaries in Facilitating Trade, *Journal of International Economics* 84(1): 73–85.

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.

Bernard, A., B. Jensen, 1999, "Exceptional exporter performance: cause, effect, or both?" *Journal of International Economics* 47.

Bernard, A., B. Jensen, S. Redding and P.K. Schott, 2007, "Firms in International Trade," *Journal of Economic Perspectives* 21(3), 105–130.

Bernard, A., S. Redding and P.K. Schott, 2010, "Multiproduct Firms and Trade Liberalization," manuscript.

Besanko, D., D. Dranove, M. Shanley, and S. Schaefer, 2007, *Economics of Strategy*, Wiley.

Burnham, K.P. and Anderson, D.R., 2002, *Model selection and multimodel inference: a practical information-theoretic approach*, Springer-Verlag.

Burtschy, Bernard and Antoine Gerbelle, 2006, *Classement des meilleurs vins de France*, La Revue du vin de France.

Comité interprofessionnel du Vin de Champagne (CIVC), 2010, "Les expéditions de vins de Champagne en 2009." `www.champagne.fr`

Crozet, M., Head, K. and T. Mayer, 2009, "Quality sorting and trade: Firm-level evidence for French wine" CEPR DP 7295.

Eaton, J. and S. Kortum, 2001, "Trade in Capital Goods," *European Economic Review* 45, 1195–1235.

Eaton, J. and S. Kortum, 2002 "Technology, Geography, and Trade" *Econometrica* 70(5), 1741–177.

Eaton, J., S. Kortum and F. Kramarz, forthcoming, "An Anatomy of International Trade: Evidence from French Firms," *Econometrica.*

Fajgelbaum, P.D., G.M. Grossman, and E. Helpman, 2009, "Income Distribution, Product Quality, and International Trade," NBER Working Paper No. 15329.

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, 2009. "Firms' Exporting Behavior under Quality Constraints," International Policy Center Working Paper No. 88.

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, H. and J. Robinson, 2005, *World Atlas of Wine*, 5th ed., London: Mitchell Beazley.

Johnson R. C., 2009, "Trade and Prices with Heterogeneous Firms," mimeo.

Juhlin, Richard, 2008, *Champagne Guide*, Richard Juhlin Publishing AB.

Khandelwal, A., 2010, "The Long and Short (of) Quality Ladders" *Review of Economic Studies*, forthcoming.

Kugler M., and Verhoogen E., 2008, "The Quality-Complementarity Hypothesis: Theory and Evidence from Colombia," NBER working paper No. 14418.

Mandel, Benjamin, 2008, "Heterogeneous Firms and Import Quality: Evidence from Transaction-Level Prices" UC Davis mimeo.

Manova, K. and Z. Zhang, 2009, "Quality Heterogeneity across Firms and Export Destinations," NBER Working Paper No. 15342.

Martin, J., 2010, "Mark-ups, Quality and Transport Costs," mimeo CREST.

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.

Parker, Robert M., 2002, *Parker's Wine Buyer's Guide, Sixth Edition*, Simon and Schuster, New York.

Schott, P.K., 2004, "Across-Product versus Within-Product Specialization in International Trade," *Quarterly Journal of Economics* 119(2), 647–678.

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 Data Appendix

A.1 Matching Juhlin data to firm-level export declarations

The key data construction task was matching quality data from Juhlin (2008) to the French export data. We refer to each distinct Juhlin listing as a "producer." Our export data are based on what INSEE refers to as an "enterprise" (in French) or "company" (English translation). The INSEE (the French national statistical institute) database called SIREN provides a additional information for many companies. In this paper we have used the "apen," the primary activity of the enterprise and the employment range (available for 62% of the companies that export Champagne in 2005).

From each Juhlin (2008) entry, we extracted the following information: producer name, the number of wines that Juhlin scored for that producer, the range of scores he gave to those wines, total production (number of bottles), establishment date (when available), the type of firm (see below), and, most importantly Juhlin's summary rating of the producer's wines.

Juhlin provides the name and address for each producer but not the SIREN identification number which we need to obtain trade data. The matching procedure we followed was first to match the line one address name or the line two company name from INSEE to the Juhlin name for a producer. If this gave no positive results, we matched street addresses. When neither the name, nor the street address matched exactly but both matched partially, we used the INSEE primary activity and elements from Juhlin's descriptions to corroborate these matches.

The mapping between Juhlin producers and INSEE companies is one-to-one for 433 producer-companies (95% of the companies). There are seven producers listed in Juhlin (2008) that are associated with more than one company. These producers map to 15 distinct SIREN numbers. In each of these cases, one of the exporting companies is a beverage wholesaler (513J). The other has champagnisation (159F) as its main activity (3 cases) or agriculture (2 cases). Two Juhlin-listed producers appear to channel exports between two distinct beverage wholesalers. There is even one producer whose Champagne is exported by three different companies: a wholesaler, a wine maker, and a firm engaged in business administration (741J). For the nine exporting companies that are associated with more than one producer, we summed production amounts and averaged the producer ratings from the corresponding Juhlin listings.

A.2 Customs export data

The annual trade database comprises firms' export declarations listing values and quantities in each product category (8-digit level) for each destination country. We retain the year 2005. The original database contains 2,221,505 observations, for a total of 9,874 products, 103,222 exporters and 234 destinations countries.

This custom data is nearly exhaustive. Firms located in France must declare all export flows to non-EU countries exceeding 1,000 euros or 1000 kilograms per destination. The average unit value in our sample is slightly higher than 18 euros per kilogram, which can be reasonably taken as the average price of an exported bottle. The declaration threshold is therefore around 55 bottles per destination country. We find very few cases of exports outside the EU that are close to the reporting threshold. In 2005, the minimum value exported outside the EU in our sample is 1005 euros for Morocco. The average of the minimum observed values for countries outside the EU, 6244 euros, is more than six times the declaration threshold.

For intra-EU trade, a detailed declaration—which includes the destination country—is still compulsory for large exporters, but not for smaller ones. Firms are allowed to fill a simplified declaration if the value they exported to the EU does not exceeded 100,000 euros over the previous year.[25] Since the simplified declaration does not mention the destination country, the corresponding flows are not included in the data we use. Note that the threshold is not binding and the customs database reports a substantial number of export flows below the threshold: About 39% of firms that export to at least one EU member export less than 100,000 euros to the European Union. There are 123 Juhlin-rated firms in this situation (43% of our exporters), but they account for only 0.14% of the total intra-EU exports of Champagne. We can only speculate on why so many individual small exporters elected to complete a detailed declaration.

A.3 Productivity Data

We measure labour productivity as the count of bottles produced per year per employee at each company. As described in the text, Juhlin (2008) provides annual production for 281 out of our 284 usable quality rated firms.

Employment data comes from two sources: Enquête Annuelle d'Entreprise (EAE), which covers firms larger than 20 employees, and INSEE's SIREN database which covers some small firms but only provides employment ranges. EAE enables to find non missing precise employment figures for 56 Juhlin-rated exporters. SIREN adds information on employment ranges for 70 more firms, putting our total number of firms at 126 for which we have both quality and productivity measured. The employment ranges are 1–2, 3–5, 6–9, 10–19, 20–49, 50–99, and then 100 by 100 until 99,999. We use an interpolation method to predict more precisely the expected number of employees for firms falling into each range. Rather than simply taking the midpoint of the ranges (implicitly assuming uniform distribution of firms inside each range), we assume a pareto distribution (which assumes a greater density for the smaller firms in each range) with minimum value of 1 and shape parameter a, such that $\mathbb{P}(X > x) = x^{-a}$. By summing all the categories greater than x, $\mathbb{P}(X > x)$ is observed as the share of firms with more than x employees, which we denote as $y(x)$. Using the pareto distribution we obtain $\hat{a}(x) = \frac{-\ln y(x)}{\ln(x)}$. We estimate a for every tranche lower bound denoted L (1,3,6,10,20,50,...). The average a is 0.54 with a standard deviation of 0.066. We then denote the upper bound of each tranche as H and generate the estimated employment for a

[25]This threshold was increased to 150,000 euros in 2006.

firm in an (L, H) tranche

$$E[x|L < x < H] = \frac{a/(a-1) \times (1/L^{(a-1)} - 1/H^{(a-1)})}{1/L^a - 1/H^a}.$$

With the estimated value of $\hat{a} = 0.54$, the employment estimation adds an average 15% to the lower bound of the employment range.

B Supplemental results

B.1 Quality effects on being an exporter

Table 6 provides econometric estimates of the effects of quality on the probability of exporting to any destination. The sample is the same as used in Figure 3(a) , that is all Juhlin-rated firms. The first column reports LPM estimates, while the second provides marginal effects of a Probit. This latter column does not report the effects for 5-star champagne, which is a perfect predictor for exporting and therefore cannot be estimated. The ordering of point estimates is as expected from the figure, with both estimation methods: 2-star champagnes have a 10 percentage points larger propensity to export than the reference 1-star champagne, while the difference is around 40 points for 4-star champagne. In terms of statistical significance, while all high-quality champagnes have a larger export probability than the reference group, the difference between 3 and 2 stars is significant at the 11% level in both methods, the significance levels are 5 and 3% for the 4 vs. 3 star difference. Estimated with LPM, the difference between 5 and 4 stars propensity is not statistically significant.

B.2 Monte-Carlo results on log-Tobit method

We have investigated three alternative distributions to the log-normal, with results reported in Table 7. First, we assumed $\ln \alpha$ was logistically distributed (so α is log-logistic). This distribution is known to resemble the normal distribution for $\ln \alpha$. Next we considered two distributions for $\ln \alpha$ that depart substantially from the normal in shape: the Gumbel (implying α is Frechet) and the exponential (α is Pareto). In each case we parameterized the distributions so as to retain the data feature that only 7% of firm-markets have positive exports. The distributions (especially the Frechet and Pareto) have long right tails, so outlier values for α can exert a strong influence on the OLS regressions confined to the 7% of cases that are selected into exporting. For this reason, the simulation trims the bottom and top 1% of the randomly generated values of α for all distributions. The trimming has very little effect on the Tobit results: In the log-normal specification, the untrimmed Tobit

Table 6: Firm-level decision to export (any destination)

	(1)	(2)
Dependent variable:	$\mathcal{E}_d(j)$	$\mathcal{E}_d(j)$
Method	LPM	Probit
Observations	478	470
2 stars	0.11^b	0.10^b
	(0.05)	(0.05)
3 stars	0.22^a	0.21^a
	(0.06)	(0.06)
4 stars	0.38^a	0.42^a
	(0.08)	(0.08)
5 stars	0.52^a	n/a
	(0.17)	
R^2	0.072	0.053

Note: Column (2) reports marginal effects of the probit estimation. For column (2), R^2 is computed as the squared correlation between the predicted and actual values of the dependent variable. Standard errors clustered at the firm-level in parentheses. Significance levels: ^{c} $p < 0.1$, ^{b} $p < 0.05$, ^{a} $p < 0.01$

coefficient is 4.53, against 4.52 in the trimmed version. The OLS coefficient in the censored OLS regression rises from 0.86 to 1.33 on average, remarkably close to the 1.31 we estimated using the real data.

The other distributions reproduce the following key features of the simulation: (1) The correlation between α and s is negative, ranging from -0.36 (Pareto) to -0.41 (log-logistic); (2) The Tobit estimates have slight downward bias, ranging from 2.6% for Pareto to 1.3% for log-normal; (3) The OLS estimates from the censored data have a very large bias towards zero, with the contraction factor ranging 0.29 for log-normal to 0.19 for Pareto.

Table 7: Simulation results for alternative distributions (assumed true $\beta = 4.58$)

	Assumed distribution of $\ln \alpha$			
Variable	normal	logistic	gumbel	expon.
Share of profitable firm-destination exports	0.07	0.07	0.07	0.07
	(0.00)	(0.00)	(0.00)	(0.00)
Correlation(s, α) in censored data	-0.39	-0.41	-0.38	-0.36
	(0.04)	(0.04)	(0.04)	(0.04)
OLS $\hat{\beta}$ before censoring $x_d^{\text{fob}}(j)$	4.58	4.58	4.58	4.58
	(0.04)	(0.04)	(0.04)	(0.04)
OLS $\hat{\beta}$ after censoring $x_d^{\text{fob}}(j) < \sigma F_d/\tau_d$	1.33	1.04	0.93	0.86
	(0.16)	(0.17)	(0.17)	(0.18)
Tobit $\hat{\beta}$ (estimate $\sigma F_d/\tau_d$ with min $x_d^{\text{fob}}(j) > 0$)	4.52	4.51	4.50	4.47
	(0.21)	(0.23)	(0.24)	(0.26)
Contraction: censored OLS $\hat{\beta}$/Tobit $\hat{\beta}$	0.29	0.23	0.21	0.19
	(0.03)	(0.04)	(0.04)	(0.04)

B.3 Tobit with alternative proxy for the entry threshold

In the main text we use the minimum export value to each country as a proxy for $\sigma F_d/\tau_d$, which defines the censoring point. In Table 8, we compare this with taking the second minimum export value to each country. The first column replicates our benchmark results from Table 2, the second column reduces the sample to the set of non-missing observations of column (3) but maintains the minimum export value as a censoring point. The third column uses the second smallest value as an alternative. Although this last specification exhibits lower coefficient estimates, the different never seems economically large or statistically significant.

Table 8: Robustness to alternative censoring value

	(1)	(2)	(3)
Dependent variable:	$\ln x_d^{\text{fob}}(j)$	$\ln x_d^{\text{fob}}(j)$	$\ln x_d^{\text{fob}}(j)$
Method	Tobit (min)	Tobit (min)	Tobit (2nd min)
Observations	44586	38906	38906
	Non-Parametric		
2 stars	1.25^b	1.28^b	1.18^b
	(0.52)	(0.51)	(0.47)
3 stars	2.68^a	2.71^a	2.46^a
	(0.55)	(0.55)	(0.50)
4 stars	5.80^a	5.84^a	5.28^a
	(0.79)	(0.79)	(0.71)
5 stars	7.70^a	7.83^a	7.06^a
	(0.59)	(0.60)	(0.54)
ψ, std. dev. of $\ln \alpha_d(j)$	4.19^a	4.20^a	3.77^a
	(0.16)	(0.17)	(0.14)
R^2	0.629/0.17	0.487/0.16	0.491/0.16

Note: Destination (d) fixed effects for all columns. R^2 include country dummies, and are computed as the squared correlation between the predicted and actual values of the dependent variable. The second R^2 reported uses the sample of positive export flows only. Standard errors clustered at the firm-level in parentheses. Significance levels: ^{c} $p < 0.1$, ^{b} $p < 0.05$, ^{a} $p < 0.01$

B.4 Inferring a low quality category from firms unrated by Juhlin

We test here the robustness of our results to the inclusion of exporters unrated by Juhlin into the analysis. With a vital caveat, exclusion from the book can be interpreted as a bad signal: We cannot infer that *all* exporters omitted from the guide are low quality because substantial amounts of Champagne are exported by non-producers. We classify firms as low quality $s(j) = 0$ if they are (1) unrated by Juhlin, (2) located within the official Champagne-growing départements ("Local")[26], and (3) engaged in grape-growing or Champagne-making as their primary activity. We then re-run the baseline regression table and show results in

[26] Marne (67% of production), Aube (22%), Aisne, Haute-Marne, and Seine-et-Marne (`http://www.champagne.fr/fr/localisation.aspx`).

Table 9. Those low quality exporters increase the sample of non-missing price and export value by 8%. The full sample including zero exports is now 66040, consisting of 418 exporters (284 rated, 134 unrated) times 158 countries served, minus the same two outlier price values mentioned in section 3.2 and two observations with positive volumes but zero value in the customs database. The non-parametric regressions retain $s(j) = 1$ as the reference group. Comparing Table 9 results with the corresponding estimates in Table 2, we see very similar effects for 2 to 5 stars. However, column (5) reveals a very steep drop in exports for 0 stars (coefficient of -2.75 relative to 1-star wine). This forces the parametric estimate to increase substantially (5.74 vs 4.58 in Table 2), pointing once again to the benefits of allowing for non-linearity in the effect of quality.

Table 9: Firm-level regressions for quality-rated and unrated Champagne exporters

	(1)	(2)	(3)	(4)	(5)
Dependent variable:	$\ln p_d^{\text{fob}}(j)$	$\mathcal{E}_d(j)$	$\mathcal{E}_d(j)$	$\ln x_d^{\text{fob}}(j)$	$\ln x_d^{\text{fob}}(j)$
Method	OLS	LPM	Probit	OLS	Tobit
Observations	3459	66040	66040	3459	66040
	Parametric				
ln stars	0.20^a	0.07^a	0.08^a	1.62^a	5.74^a
	(0.04)	(0.01)	(0.01)	(0.18)	(0.47)
ψ, std. dev. of $\ln \alpha_d(j)$					4.51^a
					(0.15)
R^2	0.21	0.23	0.31	0.25	0.68/ 0.14
	Non-Parametric				
0 stars	0.03	-0.02^a	-0.04^a	-0.81^a	-2.75^a
	(0.02)	(0.00)	(0.01)	(0.17)	(0.38)
2 stars	0.05^a	0.02^a	0.02^b	0.32	1.31^b
	(0.02)	(0.01)	(0.01)	(0.23)	(0.53)
3 stars	0.07^b	0.04^a	0.04^a	0.62^a	2.80^a
	(0.03)	(0.01)	(0.01)	(0.23)	(0.57)
4 stars	0.20^a	0.13^a	0.08^a	1.97^a	6.01^a
	(0.03)	(0.03)	(0.01)	(0.34)	(0.81)
5 stars	0.52^a	0.26^a	0.12^a	1.65^a	7.99^a
	(0.14)	(0.03)	(0.01)	(0.23)	(0.61)
ψ, std. dev. of $\ln \alpha_d(j)$					4.41^a
					(0.15)
R^2	0.31	0.25	0.33	0.29	0.69 / 0.16

Note: Destination (d) fixed effects for all columns. Column (3) reports marginal effects of the probit estimation. R^2 include country dummies. For columns (3) and (5), R^2 are computed as the squared correlation between the predicted and actual values of the dependent variable. Second R^2 in column (5) uses the same sample as column (4). Standard errors clustered at the firm-level in parentheses. Significance levels: ^{c} $p < 0.1$, ^{b} $p < 0.05$, ^{a} $p < 0.01$

B.5 More parametric specifications of non-iceberg trade costs

Table 10 reports results of Alchian-Allen regressions with the fully parametric and semi-parametric specification of interaction terms. As in the fully non-parametric version analyzed in the text, the distance interaction terms seem overall not to have an overly dominant influence on our results.

Table 10: Firm-level regressions for quality-rated Champagne exporters - non iceberg costs.

	(1)	(2)	(3)	(4)	(5)
Dependent variable:	ln $p_d^{\text{fob}}(j)$	$\mathcal{E}_d(j)$	$\mathcal{E}_d(j)$	ln $x_d^{\text{fob}}(j)$	ln $x_d^{\text{fob}}(j)$
Method	OLS	LPM	Probit	OLS	Tobit
Observations	3205	44586	44586	3205	44586
			Parametric		
ln stars	0.26^a	0.08^a	0.10^a	1.26^a	5.12^a
	(0.05)	(0.01)	(0.02)	(0.20)	(0.69)
ln stars $\times \ln(\text{dist}_d/\overline{\text{dist}})$	0.03^b	-0.03^a	0.01^a	-0.04	0.62^a
	(0.01)	(0.01)	(0.00)	(0.07)	(0.21)
ψ, std. dev. of $\ln \alpha_d(j)$					4.31^a
					(0.16)
R^2	0.242	0.275	0.321	0.228	0.623/0.15
			Semi-Parametric		
2 stars	0.07^b	0.01^b	0.02	0.28	1.01
	(0.03)	(0.01)	(0.01)	(0.30)	(0.63)
3 stars	0.11^a	0.03^a	0.05^a	0.44^c	2.70^a
	(0.03)	(0.01)	(0.01)	(0.26)	(0.64)
4 stars	0.22^a	0.12^a	0.12^a	1.92^a	6.09^a
	(0.04)	(0.03)	(0.02)	(0.32)	(0.86)
5 stars	0.53^a	0.24^a	0.16^a	1.33^a	7.72^a
	(0.12)	(0.03)	(0.02)	(0.21)	(0.66)
2 stars $\times$ ln distance	0.01	-0.02^a	-0.01	-0.03	-0.26
	(0.01)	(0.01)	(0.00)	(0.11)	(0.22)
3 stars $\times$ ln distance	0.03^b	-0.03^a	-0.00	-0.14	-0.01
	(0.01)	(0.01)	(0.00)	(0.10)	(0.23)
4 stars $\times$ ln distance	0.01	-0.03^a	0.01^b	-0.01	0.52^b
	(0.02)	(0.01)	(0.01)	(0.12)	(0.22)
5 stars $\times$ ln distance	0.00	-0.09^a	-0.01	-0.43^a	-0.07
	(0.04)	(0.01)	(0.00)	(0.11)	(0.20)
ψ, std. dev. of $\ln \alpha_d(j)$					4.19^a
					(0.16)
R^2	0.318	0.293	0.334	0.268	0.634/0.17

Note: Destination (d) fixed effects for all columns. Column (2) reports marginal effects of the probit estimation. R^2 include country dummies. For columns (3) and (5), R^2 are computed as the squared correlation between the predicted and actual values of the dependent variable. Second R^2 in column (5) uses the same sample as column (4). Standard errors clustered at the firm-level in parentheses. Significance levels: ^{c} $p < 0.1$, ^{b} $p < 0.05$, ^{a} $p < 0.01$

B.6 Robustness to EU dummy interactions

Table 11 reports results of our benchmark firm-level set of regressions with European Union interaction terms (In 2005, the EU had 25 members countries, including France of course not part of the interaction terms).

Table 11: Firm-level regressions for quality-rated Champagne exporters - EU interactions

	(1)	(2)	(3)	(4)	(5)
Dependent variable:	$\ln p_d^{\text{fob}}(j)$	$\mathcal{E}_d(j)$	$\mathcal{E}_d(j)$	$\ln x_d^{\text{fob}}(j)$	$\ln x_d^{\text{fob}}(j)$
Method	OLS	LPM	Probit	OLS	Tobit
Observations	3205	44586	44586	3205	44586
			Parametric		
ln stars	0.24^a	0.07^a	0.10^a	1.22^a	4.86^a
	(0.04)	(0.01)	(0.02)	(0.19)	(0.68)
ln stars × EU dummy	-0.03	0.13^a	-0.02	0.17	-0.64
	(0.04)	(0.02)	(0.01)	(0.15)	(0.49)
ψ, std. dev. of $\ln \alpha_d(j)$					4.31^a
					(0.16)
R^2	0.239	0.281	0.318	0.228	0.619/0.15
			Semi-Parametric		
2 stars	0.06^b	0.01	0.02	0.34	0.85
	(0.03)	(0.00)	(0.01)	(0.26)	(0.62)
3 stars	0.10^a	0.03^a	0.05^a	0.54^b	2.45^a
	(0.03)	(0.01)	(0.01)	(0.24)	(0.65)
4 stars	0.20^a	0.11^a	0.11^a	1.93^a	5.84^a
	(0.05)	(0.03)	(0.02)	(0.31)	(0.86)
5 stars	0.51^a	0.20^a	0.15^a	1.30^a	7.18^a
	(0.10)	(0.03)	(0.02)	(0.21)	(0.66)
2 stars × EU dummy	-0.01	0.07^a	0.02	-0.02	0.84
	(0.03)	(0.02)	(0.01)	(0.22)	(0.53)
3 stars × EU dummy	-0.06	0.11^a	0.01	0.14	0.50
	(0.04)	(0.02)	(0.01)	(0.21)	(0.58)
4 stars × EU dummy	0.00	0.16^a	-0.01	0.08	-0.41
	(0.07)	(0.03)	(0.01)	(0.27)	(0.56)
5 stars × EU dummy	0.02	0.41^a	0.06^a	0.85^a	1.75^a
	(0.12)	(0.02)	(0.01)	(0.25)	(0.49)
ψ, std. dev. of $\ln \alpha_d(j)$					4.18^a
					(0.17)
R^2	0.318	0.302	0.334	0.267	0.632/0.17

Note: Destination (d) fixed effects for all columns. Column (2) reports marginal effects of the probit estimation. R^2 include country dummies. For columns (3) and (5), R^2 are computed as the squared correlation between the predicted and actual values of the dependent variable. Second R^2 in column (5) uses the same sample as column (4). Standard errors clustered at the firm-level in parentheses. Significance levels: ^{c} $p < 0.1$, ^{b} $p < 0.05$, ^{a} $p < 0.01$

B.7 Baseline regressions using alternative quality ratings

Table 12: Firm-level regressions using Parker quality ratings

	(1)	(2)	(3)	(4)	(5)
	$\ln p_d^{\text{fob}}(j)$	$x_d^{\text{fob}}(j) > 0$	$x_d^{\text{fob}}(j) > 0$	$\ln x_d^{\text{fob}}(j)$	$\ln x_d^{\text{fob}}(j)$
Method	OLS	LPM	Probit	OLS	Tobit
Observations	3205	44586	44586	3205	44586
			Parametric		
ln stars	0.19^a	0.12^a	0.08^a	1.46^a	4.15^a
	(0.04)	(0.02)	(0.01)	(0.16)	(0.34)
ψ, std. dev. of $\ln \alpha_d(j)$					3.93^a
					(0.15)
R^2	0.267	0.302	0.351	0.328	0.657/ 0.25
			Non-Parametric		
2 stars	0.01	0.10^b	0.08^a	1.70^a	4.33^a
	(0.07)	(0.04)	(0.02)	(0.38)	(1.04)
3 stars	0.15^a	0.15^a	0.11^a	2.07^a	5.46^a
	(0.03)	(0.04)	(0.02)	(0.30)	(0.80)
4 stars	0.18^a	0.11^a	0.09^a	1.76^a	4.70^a
	(0.04)	(0.03)	(0.02)	(0.47)	(0.85)
5 stars	0.40^a	0.26^a	0.15^a	2.51^a	7.37^a
	(0.10)	(0.05)	(0.02)	(0.32)	(0.71)
ψ, std. dev. of $\ln \alpha_d(j)$					3.87^a
					(0.18)
	0.332	0.297	0.340	0.262	0.640/ 0.19

Note: Destination (d) fixed effects for all columns. Column (2) reports marginal effects of the probit estimation. R^2 include country dummies. For columns (3) and (5), R^2 are computed as the squared correlation between the predicted and actual values of the dependent variable. Second R^2 in column (5) uses the same sample as column (4). Standard errors clustered at the firm-level in parentheses. Significance levels: ^{c} $p < 0.1$, ^{b} $p < 0.05$, ^{a} $p < 0.01$

Table 13: Firm-level regressions using RVF quality ratings

	(1)	(2)	(3)	(4)	(5)
	$\ln p_d^{\text{fob}}(j)$	$x_d^{\text{fob}}(j) > 0$	$x_d^{\text{fob}}(j) > 0$	$\ln x_d^{\text{fob}}(j)$	$\ln x_d^{\text{fob}}(j)$
Method	OLS	LPM	Probit	OLS	Tobit
Observations	3205	44586	44586	3205	44586
			Parametric		
ln stars	0.25^a	0.14^a	0.10^a	1.26^a	4.69^a
	(0.06)	(0.02)	(0.01)	(0.20)	(0.42)
ψ, std. dev. of $\ln \alpha_d(j)$					4.17^a
					(0.19)
R^2	0.290	0.293	0.337	0.243	0.636/ 0.18
			Non-Parametric		
included	0.08^a	0.09^a	0.08^a	1.51^a	4.01^a
	(0.03)	(0.03)	(0.02)	(0.44)	(0.90)
1 stars	0.27^a	0.11^a	0.09^a	1.48^a	4.48^a
	(0.04)	(0.04)	(0.02)	(0.38)	(1.03)
2 stars	0.25^a	0.24^a	0.14^a	1.82^a	6.80^a
	(0.06)	(0.05)	(0.02)	(0.43)	(0.81)
3 stars	0.63^b	0.29^a	0.16^a	1.51^a	7.24^a
	(0.30)	(0.04)	(0.02)	(0.26)	(0.63)
ψ, std. dev. of $\ln \alpha_d(j)$					4.14^a
					(0.18)
	0.332	0.297	0.340	0.262	0.640/ 0.19

Note: Destination (d) fixed effects for all columns. Column (2) reports marginal effects of the probit estimation. R^2 include country dummies. For columns (3) and (5), R^2 are computed as the squared correlation between the predicted and actual values of the dependent variable. Second R^2 in column (5) uses the same sample as column (4). Standard errors clustered at the firm-level in parentheses. Significance levels: ^{c} $p < 0.1$, ^{b} $p < 0.05$, ^{a} $p < 0.01$