Journal of International Economics 97 (2015) 100–111

Contents lists available at ScienceDirect

Journal of International Economics

journal homepage: www.elsevier.com/locate/jie

Export prices of U.S. firms☆

James Harrigan [a,b,*], Xiangjun Ma [c], Victor Shlychkov

[a] University of Virginia, United States
[b] NBER, United States
[c] University of International Business and Economics, Beijing, China

ARTICLE INFO

Article history:
Received 4 March 2014
Received in revised form 16 March 2015
Accepted 27 April 2015
Available online 14 May 2015

JEL classification:
F1
F10
F23

Keywords:
Exporters
Firm level data
Pricing
Heterogeneous firms

ABSTRACT

Using confidential firm-level data from the United States in 2002, we show that exporting firms charge prices for narrowly defined goods that differ substantially with the characteristics of firms and export markets. We control for selection into export markets using a three-stage estimator. We have three main results. First, we find that highly productive and skill-intensive firms charge higher prices, while capital-intensive firms charge lower prices. Second, U.S. firms charge substantially higher prices to markets other than Canada and Mexico. Third, the correlation between distance and product-level export prices is largely due to a composition effect.

1. Introduction

Economists now know a lot about the characteristics of exporting plants and firms: they are bigger and more productive than non-exporters, and in many countries they are also more skill-intensive and capital-intensive.[1] For U.S. firms, recent research by Bernard et al. (2005) and Bernard et al. (2007) shows that exporters are also quite likely to import, and big firms trade many different products with many different countries. These facts about heterogeneous exporters have informed a vibrant theory literature, beginning with Melitz (2003) and Bernard et al. (2003).

Economists have also documented systematic heterogeneity in the prices that are charged for the same traded products. Starting with Schott (2004), it has been established that even within narrowly defined product categories, average prices differ systematically with the characteristics of importing and exporting countries. Since many firms may sell even in narrowly defined product categories, explaining these product-level findings requires firm-level data. There have been only a few studies that examine export price variation across markets using firm-level data, including Martin (2012) for France, Bastos and Silva (2010) for Portugal, Gorg et al. (2010) for Hungary, and Manova and Zhang (2012) for China. Our paper is the first to use U.S. firm-level data to look at export pricing, and we establish some new facts:

- Within country–product categories, firms that are more productive and skill-intensive charge higher prices, while more capital-intensive firms charge lower prices.
- Within narrow product categories, exporting firms charge higher prices to countries other than Canada and Mexico.
- The product-level correlation between export prices and destination market characteristics is largely due to a selection effect, where firms that charge higher prices are more likely to select into tougher markets, where "tougher" refers to both the costs of market access and the degree of competition.

☆ We thank David Weinstein, Eric Verhoogen, Antoine Gervais, and seminar audiences in North America, Western Europe, and China for very helpful comments. We also thank the Editor and referees of this journal for their particularly insightful feedback. This paper uses restricted data that was analyzed at the U.S. Census Bureau's Research Data Centers in New York City and North Carolina. Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed. James Harrigan thanks the University of Virginia's Bankard Fund for Political Economy for financial support. Xiangjun Ma thanks the Fundamental Research Funds for the Central Universities in UIBE (7600011127) for support.

* Corresponding author at: Department of Economics, University of Virginia, Charlottesville, VA 22904, United States.
E-mail address: james.harrigan@virginia.edu (J. Harrigan).

[1] The literature documenting these facts is vast. Good summaries of the evidence include Bernard et al. (2005) for the United States and Mayer and Ottaviano (2008) for Europe. Wagner (2012) reviews the most recent evidence.

http://dx.doi.org/10.1016/j.jinteco.2015.04.007
0022-1996/

Understanding firm-level export pricing has implications for both theory and price index measurements in international economics. The facts that we establish are broadly supportive of models where consumers value quality, but quality is expensive to produce. With heterogeneous producers, as in the models of Verhoogen (2008), Kugler and Verhoogen (2012), Johnson (2012), Hallak and Sivadasan (2009) and Baldwin and Harrigan (2011), in equilibrium more successful firms produce higher cost and higher quality goods which command higher prices. What we mean by price in this context is the ordinary definition of money per unit, although consumers who value quality can be thought of as choosing goods on the basis of "quality-adjusted" prices: an expensive, high quality good may have a lower "quality adjusted" price than a cheap, shoddy good.

An implication of these models of quality competition and heterogeneous firms is that the marginal firm has low quality and sells at a low price. When more firms enter, the entering firms charge lower prices, and thus the average unit value in a market will fall. This extensive margin of firm entry can be thought of as happening across markets that differ in their level of competition, with more entry and thus lower average prices in less competitive markets. A simple comparison of average money prices across markets, however, will have misleading implications for welfare, since with quality competition the true price index can be lower when the average money price is higher.

The extensive margin of firm entry may also operate at business cycle frequencies, with less competitive firms entering in booms and exiting during busts, as in the models of Ghironi and Melitz (2005) and Feenstra et al. (2009). In Ghironi and Melitz (2005) and Feenstra et al. (2009) there is no quality competition, firms compete only on price, and the best firms have the lowest prices. Average export prices *rise* in booms, since the marginal entering firms have high costs and prices. With quality competition, as suggested by our results, this implication is reversed: average export prices *fall* in booms, since the marginal entering firms have low costs and prices. While our data analysis looks only at cross-sectional variation in export prices, the support that we find for quality competition thus has implications for models of international business cycles.

2. Analytical framework

In this section we discuss our hypotheses and how we will test them. Subsequent sections discuss data and measurement issues, and report our results.

2.1. Firm-level export prices and destination market characteristics

Baldwin and Harrigan (2011) found that there is a strong and robust relationship between destination market characteristics and export prices at the HS10 product level: product-level U.S. export prices in 2005 increase strongly with distance and decrease with GDP, GDP per capita, and remoteness. Their theoretical explanation for these findings comes from a variation on the Melitz (2003) model. In their model, heterogeneous firms compete on quality as well as price, with the most profitable firms producing high quality, high price goods. Selection implies that only the best firms will enter the toughest markets, which theory suggests will be small, distant, and well-served by other exporters. The empirical findings are then explained as a composition effect: since only the best firms sell in the toughest markets, and these firms charge high prices, average prices at the product level will be increasing in measures of market toughness. An alternative explanation is simple price discrimination, with firms charging systematically higher prices in more distant markets. The importance of these firm-level mechanisms, selection and/or price discrimination, can only be evaluated with firm-level data.

2.1.1. A simple decomposition

As a matter of arithmetic, the average price of a given product exported to destination d is a quantity-weighted average of the prices charged by all the N firms f that export the good,

$$\overline{p}_d = \sum_{f=1}^{N} w_{fd} p_{fd}, \quad w_{fd} = \frac{q_{fd}}{\sum_{f=1}^{N} q_{fd}}, \tag{1}$$

where p_{fd} and q_{fd} are the price and quantity respectively of the good sold by firm f in destination d, N is the number of exporting firms selling the good, and the weight w_{fd} is firm f's quantity market share in market d, that is, the quantity share of firm f among all firms selling in destination d. For each good, we also define a given firm's quantity-weighted average price across all D markets,

$$\overline{p}_f = \sum_{d=1}^{D} \omega_{fd} p_{fd}, \quad \omega_{fd} = \frac{q_{fd}}{\sum_{d=1}^{D} q_{fd}}.$$

The weight ω_{fd} is firm f's quantity share of its total worldwide sales that take place in destination d. The overall average world price for a good is

$$\overline{p} = \frac{\sum_{d=1}^{D}\sum_{f=1}^{N} p_{fd} q_{fd}}{\sum_{d=1}^{D}\sum_{f=1}^{N} q_{fd}} = \sum_{f=1}^{N} \overline{p}_f \overline{w}_f, \quad \overline{w}_f = \frac{\sum_{d=1}^{D} q_{fd}}{\sum_{d=1}^{D}\sum_{f=1}^{N} q_{fd}},$$

where $\overline{w}_f$ is firm f's average quantity market share in the world market, defined as f's total quantity sold divided by total quantity sold by all firms. With these definitions, we establish a decomposition that shows how the destination d average price $\overline{p}_d$ differs from the world average price $\overline{p}$:

Lemma 1. Price decomposition across destinations

$$\overline{p}_d - \overline{p} = \underbrace{\sum_{f=1}^{N}\left(p_{fd} - \overline{p}_f\right)\overline{w}_f}_{\text{price discrimination}} + \underbrace{\sum_{f=1}^{N}(w_{fd} - \overline{w}_f)\overline{p}_f}_{\text{market share}} + \underbrace{\sum_{f=1}^{N}(w_{fd} - \overline{w}_f)\left(p_{fd} - \overline{p}_f\right)}_{\text{interaction}}. \tag{2}$$

Proof. By substitution from the above definitions,

$$\overline{p}_d - \overline{p} = \sum_{f=1}^{N} w_{fd} p_{fd} - \sum_{f=1}^{N} \overline{p}_f \overline{w}_f.$$

Adding and subtracting $\sum_{f=1}^{N} w_{fd}\overline{p}_f$ and $\sum_{f=1}^{N} \overline{w}_f\left(p_{fd} - \overline{p}_f\right)$ from the right hand side of the above, collecting terms, and re-arranging gives Eq. (2).

If a given firm charges the same price in all destinations, then $p_{fd} = \overline{p}_f$, and the first and third summations in Eq. (2) will be zero. As a consequence, the average price across destinations will differ only because of differences in the quantities sold by different firms. More generally, the average export price can also differ because a given firm charges different prices in different destinations, in which case the first and third summations in Eq. (2) will be non-zero.

2.1.2. An econometric model

We now turn to a closer examination of firm-level export pricing behavior across markets. We begin with two descriptive linear equations. Let $\mathbf{X}_d$ denote a vector of destination country characteristics including distance, real GDP, etc. Linear projections of log export prices from the U.S. of product i by firm f to destination d are given by

$$\ln p_{ifd} = \alpha_{1i} + \beta \mathbf{X}_d + \varepsilon_{ifd}, \tag{3}$$

$$\ln p_{ifd} = \alpha_{1if} + \beta \mathbf{X}_d + \varepsilon_{ifd}. \tag{4}$$

The parameter α_{1i} is a product fixed effect, while α_{1if} is a product-firm fixed effect. The error term is ε_{ifd}. The vector β is the parameter of interest, as it answers the question: how do firm-level export prices differ across destinations as a function of the characteristics of those destinations? Eq. (3), which includes only product fixed effects, identifies β through variation both within firms across markets and across firms. In this way it is very similar to the specifications estimated by Baldwin and Harrigan (2011) using product level data. By contrast, Eq. (4), which includes product-firm fixed effects, identifies β using only within-firm variation across markets, and thus allows a direct test of the hypothesis that firms vary their export prices systematically with export market characteristics $\mathbf{X}_d$. In either equation, the error term ε_{ifd} has the interpretation of unmeasured factors that lead firms to optimally charge higher prices in some markets than others. When there are product-firm fixed effects, an intuitive interpretation of ε_{ifd} is as a market-specific demand or supply shock: for a given firm–product, demand is randomly higher, or supply is randomly less costly, in some markets d than in others.

Any model of product market competition suggests that market entry is a key choice for the firm, and that firm characteristics will determine which markets are entered. Theory also suggests that the price charged conditional on entry will be a key determinant of entry, which implies that the interpretation of β in Eqs. (3) and (4) is complicated by a selection bias.[2] In particular, if firms compete on quality so that higher-price firms are the most competitive, β will conflate selection and price discrimination effects. The key statistical issue is that we only observe a firm's pricing decision when it chooses to sell in a market. Consider the reduced form export volume equation

$$\ln y_{ifd} = \mathrm{Max}\left[0, \alpha_{2i} + \delta_i \mathbf{X}_d + u_{ifd}\right], \tag{5}$$

where y_{ifd} is export sales of product i by firm f in market d. Economic theory suggests that the errors u_{ifd} from the export volume (Eq. (5)) will be correlated with the errors ε_{ifd} from the export price (Eqs. (3) and (4)), $E[\varepsilon s|\alpha_{1i}, \alpha_{2i}, \mathbf{X}_d, u_{ifd}] = \rho u_{ifd}$. One reason for such a correlation is market-specific demand shocks that affect both optimal sales and optimal price, in which case ρ would be likely to be positive. Another reason is market specific supply shocks (such as particularly high or low unobserved trade costs) that also affect both optimal sales and price, in which case ρ would likely be negative. Whatever the cause, this correlation is what gives rise to selection bias in the price (Eqs. (3) and (4)). Given a consistent estimate $\hat{u}_{ifd}$ of the errors u_{ifd} from Eq. (5), selection bias can be controlled for by including $\hat{u}_{ifd}$ as a regressor in Eqs. (3) and (4), leading to the estimating equations

$$\ln p_{ifd} = \alpha_{1i} + \beta \mathbf{X}_d + \gamma \hat{u}_{ifd} + \varepsilon_{ifd}, \tag{6}$$

$$\ln p_{ifd} = \alpha_{1if} + \beta \mathbf{X}_d + \gamma \hat{u}_{ifd} + \varepsilon_{ifd}\ . \tag{7}$$

If u_{ifd} in Eq. (5) is assumed to be normally distributed, then Eq. (5) can be estimated by Tobit, with the residuals $\hat{u}_{ifd}$ from the estimated export participation (Eq. (5)) used as a regressor in the export price (Eqs. (6) and (7)), which are estimated by OLS. This is the two-step estimator developed by Wooldridge (1995).

As we show in our discussion in Section 4.1 below, the observed differences in export prices across markets are influenced by two types of selection. The simplest is that firms choose to sell in some markets and not others. A second, more subtle effect is a within-product composition effect, where firms sell a different mix of varieties of the same product in different markets. The selection correction term in Eqs. (6) and (7) is well-suited to controlling for the first type of selection. The within-product composition effect is unobservable in our firm × product level data, and in Section 4.2 we discuss how this effect influences our interpretation.

A notable feature of Wooldridge's estimator in our context is that identification of the price equations does *not* require an exclusion restriction: that is, the model is identified even if the vector of country characteristics $\mathbf{X}_d$ is the same in both the selection (Eq. (5)) and the price (Eqs. (6) and (7)). The intuition is that the export volume $\ln y_{ifd}$ functions as an excluded variable in the price equations. That is, variability in $\ln y_{ifd}$ is an independent source of variation which allows β in the price equations to be identified. Another source of identification is that we allow the effect of $\mathbf{X}_d$ to differ across products i in Eq. (5), which is to say that we estimate Eq. (5) separately for each product while Eqs. (6) and (7) are estimated by pooling across products. This procedure can be interpreted as interacting $\mathbf{X}_d$ with product dummies in a pooled estimation of Eq. (5) but not in Eqs. (6) and (7), which amounts to a set of exclusion restrictions in the second stage. The reason for this strategy is that we are not interested in making inferences about the first stage parameters, so the possible inefficiency of estimating Eq. (5) product-by-product is not a concern. By contrast, inference about β in Eqs. (6) and (7) requires pooling across products.[3]

A drawback of Wooldridge's two-step estimator is that it assumes that the parameter vector δ_i of the index function in Eq. (5) is constant for all values of latent exports — in particular, it assumes that the marginal effect of $\mathbf{X}_d$ on the entry decision is the same as the marginal effect for strictly positive exports. To avoid this assumption, instead of estimating Eq. (5) by Tobit we estimate it using a two-step Heckman estimator, which allows δ_i in the Probit step to differ from δ_i when exports are strictly positive. To be precise, in our first step we estimate the probability of entry using a reduced form Probit,

$$\Pr\left(y_{ifd} > 0\right) = \Phi(\alpha_{2i} + \delta_{1i} \mathbf{X}_d). \tag{8}$$

Eq. (8) is estimated separately for each product i, over all possible firm × destinations. From Eq. (8) we obtain the estimated inverse Mills ratio $\hat{\lambda}_{ifd}$. We then estimate the export volume equation for positive levels of exports by OLS separately for each product i, with $\hat{\lambda}_{ifd}$ as an additional regressor,

$$\ln y_{ifd} = \alpha_{2i} + \delta_{2i} \mathbf{X}_d + \gamma_i \hat{\lambda}_{ifd} + u_{ifd}. \tag{9}$$

The residuals $\hat{u}_{ifd} = \ln y_{ifd} - \hat{\alpha}_{2i} - \hat{\delta}_{2i} \mathbf{X}_d$ from the two-step Heckman procedure are then used as the control for selection in the export price (Eq. (6)). When estimating Eq. (7), which includes product-firm fixed effects, Eq. (9) also includes product-firm fixed effects.

The covariance matrix of $\hat{\beta}$ when estimating Eq. (6) or (7) should include an adjustment for the fact that $\hat{u}_{ifd}$ is a generated regressor, and this adjustment could in principle be done by bootstrapping across the three steps. This proved to be computationally infeasible, and we will discuss the potential implications of this when we report our results below.[4]

[2] This type of selection bias is studied at the country level by Helpman et al. (2008).

[3] Equivalently, there is not enough independent variation in the data to precisely estimate a separate β for each product, though we do report estimates on various subsets of products in our robustness analysis.

[4] The reason that bootstrapping was computationally infeasible is that a single run of our three-step procedure requires on the order of 8000 first-stage Probits. Bootstrapping with 1000 replications over 10 specifications would thus require 80 million first-stage Probits. The Census Bureau facility where this research was conducted did not have the computational resources for us to do this.

2.2. Firm-level export prices and firm characteristics

The analysis above focuses only on how destination market characteristics affect firms' pricing decisions. For the subset of our data that consists of exports of manufactured goods we can go one step further and see how firm characteristics such as productivity, skill intensity, and capital intensity are related to export pricing. We identify the effects of firm characteristics $\mathbf{X}_f$ using a model with destination × product fixed effects α_{1id},

$$\ln y_{ifd} = \text{Max}\left[0, \alpha_{2i} + \delta_{2i}\mathbf{X}_f + \delta_{3i}\mathbf{X}_d + u_{ifd}\right] \tag{10}$$

$$\ln p_{ifd} = \alpha_{1id} + \beta\mathbf{X}_f + \gamma\hat{u}_{ifd} + \varepsilon_{ifd}. \tag{11}$$

For each destination–product, Eq. (11) uses only variation across firms to identify β. Thus, Eq. (11) answers the question: within a group of firms selling product i in destination d, how do firm characteristics covary with the prices that firms charge? The selection Eq. (10) allows selection of firm f into market d to depend on both firm characteristics $\mathbf{X}_f$ and market characteristics $\mathbf{X}_d$, and the selection control $\hat{u}_{ifd}$ in Eq. (11) allows for consistent estimation of the effect of $\mathbf{X}_f$ on pricing.

Because of the destination × product fixed effects in Eq. (11), the specification cannot directly identify the effects of country characteristics $\mathbf{X}_d$ on export prices. To reintroduce these effects, we estimate a model that adds interactions between firm and country characteristics to Eq. (11),

$$\ln p_{ifd} = \alpha_{1id} + \beta_1\mathbf{X}_f + \beta_2\mathbf{X}_f\mathbf{X}_d + \gamma\hat{u}_{ifd} + \varepsilon_{ifd}. \tag{12}$$

Eq. (12) implies that the elasticity of export prices with respect to country characteristics will depend on the characteristics of the firm. The interaction terms can be written out as

$$\beta_2\mathbf{X}_f\mathbf{X}_d = \sum_{l=1}^{D}\sum_{k=1}^{K}\beta_{2kl}x_{ld}x_{kf}.$$

For country characteristic l in destination d, x_{ld}, the derivative as a function of firm characteristics x_{kf} is

$$\frac{\partial \ln p_{ifd}}{\partial x_{ld}} = \sum_{k=1}^{K}\beta_{2kl}x_{kf}. \tag{13}$$

In our empirical results below we will compare estimates of this derivative, which is identified across firms within a destination, with the simple linear effect implied by Eq. (7),

$$\frac{\partial \ln p_{ifd}}{\partial x_{ld}} = \beta_l$$

which is identified across destinations within a firm. In cases where the country characteristics of interest x_{ld} are indicators rather than continuous variables, the derivative (Eq. (13)) is undefined, so the interpretation of the interaction effect is the average difference

$$E\Big(\ln p_{ifd}|x_{ld} = 1, \alpha_{1id}, \ldots\Big) - E\Big(\ln p_{ifd}|x_{ld} = 0, \alpha_{1id}, \ldots\Big) = \sum_{k=1}^{K}\beta_{2kl}x_{kf}. \tag{14}$$

3. Data sources and measurement

We use both firm-level and country-level data, and discuss sources and measurement issues in the next two subsections.

3.1. Firm-level data

We use data on firm-level U.S. exports in 2002. For manufacturing firms, the export data is linked with production data from the 2002 Census of Manufactures. Use of this data was pioneered by Bernard et al. (2005), who provide a detailed discussion of numerous important issues related to construction of the dataset. The data has also been analyzed by Bernard et al. (2007).[5]

We use several classification systems in building our firm-level data. All export data are classified by 10-digit Harmonized System (HS10) product codes. Firms are identified by their Employer Identification Number (EIN) and also by their primary industry, as classified by the 6-digit North American Industrial Classification System (NAICS).

The firm-level export data comes from transaction-level export declarations filed by exporting firms with the U.S. Customs. The transaction-level data contain information about value, HS10-digit product code, quantity, relationship (intra-firm or arm's-length), export destination, date, and transport mode for every shipment. Firm-level data are simply sums of transaction-level data.

Our empirical definition of a product in all of what follows is an HS10-digit code, of which there were almost 9000 in 2002. Our measure of price is unit value (value divided by quantity) for a given exporter–product–country observation. The 10-digit HS system is the most disaggregated product classification system in use in the United States, but it is important to keep in mind that what ordinary people (whether consumers or business managers) think of as a product or variety of a product is yet more disaggregated. For example, consider the following 10-digit HS codes:

8703.10.50.30	Golf carts
8708.30.50.20	Brake drums
8501.10.60.20	Small electric motors, alternating current
9006.53.01.10	35 mm film cameras, with built in flash

It is easy to imagine a given firm exporting many distinct varieties, which have different characteristics and sell for different prices, under one of these product headings. This implies that our observed unit values are averages of these different variety-level prices. Therefore, the unit values of exports sent by a single firm may vary across destinations simply due to differences in variety mix, with no difference in the prices of individual varieties (for example, a golf cart manufacturer may sell its different models for the same price throughout the world, but if it sells relatively more of its high-priced models to Japan than to Canada, then the export unit value will be higher to Japan than to Canada). The trade data includes many different definitions of unit, depending on the product: number, dozens, kilogram, liter, square meter, etc. Comparing unit values across products is inherently meaningless due to this heterogeneity of units, and when constructing our unit values we are careful to make sure that the definition of units is consistent within a given HS10 code. For many products, there are two units reported, where the first unit is (for example) number and the second is kilograms. In these cases, we always use the first unit instead of kilograms, because these natural units are more likely to be comparable within products[6]. As final controls for potential data problems, we drop all observations where quantities are imputed, and also drop the top and bottom 1% of unit values by HS code.

The production data for 2002 comes from the Census of Manufactures, which collects information on the universe of U.S. manufacturing plants. For the purposes of computing firm-level productivity, we also use the Annual Survey of Manufactures data from 1997 to 2002. In each of these years the sample consists of 50,000–60,000 plants.[7]

[5] We are very grateful to J. Bradford Jensen and Peter Schott for their extensive and gracious help with the data.

[6] To see the issue, consider the example of HS 8802.40.00.40, Airplanes weighing at least 15,000 kg. Larger airplanes are more expensive, but might not be more expensive per kilogram, so it is more meaningful to define the unit value of an airplane as "dollars per plane" rather than "dollars per kilogram of plane".

[7] Some 10,000 plants are selected with certainty (including all plants with total employment above 250 workers), and more than 40,000 plants are selected with probability proportional to a composite measure of establishment size. See http://www.census.gov/ for details.

The unified dataset contains annual plant information that includes the total value of shipments, change in inventories, total employment, numbers of production and non-production workers, cost of materials, and 6-digit NAICS industry. As is standard in the literature that uses the Census of Manufactures, we proxy skilled labor and unskilled labor by non-production and production workers respectively. Due to missing data on capital stocks in the Annual Survey of Manufactures, the capital series was constructed using data for capital from the Census of Manufactures, industry depreciation rates from the Bureau of Economic Analysis, and investment series available for all years.

Plant-level revenue total factor productivity (TFP) is computed using the now-standard methods of Olley and Pakes (1996). Firm-level TFP is constructed as a shipment-weighted average of plant level revenue TFP.

The export and manufacturing datasets are linked at the level of the firm. The links between the datasets are made using the Employer Identification Number (EIN) where possible and using "alpha", an identifier of multi-unit firms that have exports to Canada, when the EIN is not available (in particular, for exports to Canada). This identifier is assigned using the business name information from the Census Bureau Business Register, also called Standard Statistical Establishment List (SSEL).

3.2. Country-level data

Our measurements of country characteristics are much the same as those used in Baldwin and Harrigan (2011), and our discussion here is drawn from Baldwin and Harrigan (2011). The objective is to measure features of export markets that affect competition in the market, and that will thus have effects on which firms enter and what prices they charge when they do enter.

3.2.1. Trade costs

While trade costs are likely to be weakly monotonic in distance, there is no reason to expect them to have any particular functional form, so we specify the distance proxy in two ways. The first is simply log distance, which we measure as kilometers from Chicago to the capital city of the importer, which comes from CEPII.[8] Our second trade cost measure breaks distance down into bins, derived from looking for natural breaks in distance among U.S. trading partners. The first bin includes Canada and Mexico. The second bin, 1–4000 km, includes countries in the Caribbean basin and northern South America. The third bin, 4000–7800 km, includes Western Europe and Brazil. The fourth bin, 7800–14,000 km, includes Eastern Europe and most of Asia (Japan, China, India, etc.). The final bin, 14,000 + km, includes Australia and Indonesia.

3.2.2. Market size

Our measure of market size is real GDP, from the Penn World Tables. We also include real GDP per worker as a demand-related control.

3.2.3. Remoteness

The structural gravity literature (including Eaton and Kortum (2002) and Anderson and Van Wincoop (2003)) emphasizes that demand conditions in country d depend on the supply conditions of all countries that potentially sell there. The proper specification of this "remoteness" effect is model-specific, but most theoretically consistent measures of remoteness have a common structure as they all work via the average price of goods sold in a destination market. This average price in turn depends upon the number of varieties produced locally in the destination market, and the number of imported varieties and the bilateral trade costs they face. As the number of varieties coming from each exporting nation is – roughly speaking – related to the origin-nation's size, a reasonable proxy for remoteness involves market-size weighted sums of an inverse power function of trade costs. Following this logic, we adopt the following measure of remoteness in our empirical work,

$$R_d = \left[\sum_{o=1}^{C} Y_o dist_{od}^{-\eta}\right]^{-1}, \tag{15}$$

where Y_o is real GDP in origin country o, and $dist_{od}$ is distance between countries o and d. Harrigan (2003) shows that this remoteness index is an approximation to the model-specific measures of Anderson and Van Wincoop (2003), and Novy (2012) shows that similar expressions hold in the model of Eaton and Kortum (2002) and other bilateral trade models with CES preferences. Empirical implementation of Eq. (15) involves some potentially important choices about how to measure within-country distance $dist_{dd}$ and what value to use for the exponent η. Fortunately, our empirical results are entirely insensitive to any reasonable choice of how to construct Eq. (15), and in what follows we include within country distance as reported in the CEPII data, and set η equal to 1. The reason for this robustness is simply that the cross-section variation in Eq. (15) is overwhelmingly dominated by differences in the GDP-weighted raw distances (consider New Zealand versus Belgium), so that different choices about including own distance and what value to choose for η lead to very highly correlated measures.

4. Empirical results: firm-level export prices

In this section we investigate the relationship between export prices, firm characteristics, and destination market characteristics. We begin by analyzing our full sample of U.S. firms in 2002, which includes exporters of both manufactured and non-manufactured goods. Our second set of results uses data only on manufactured goods exports, and we establish new facts about how export prices vary with firm characteristics.

4.1. Export price decompositions

Our data on export prices in 2002 have three dimensions of variation: HS10 products, firms, and destination markets. Comparing prices across HS10 products is meaningless, so all of our analysis of log price variation is relative to means: either product, firm × product, or destination × product.

We begin with a very simple variance decomposition exercise for log export prices. Once we remove product means, so that price variation is comparable across products, we find that the standard deviation of log export price variation within products is 1.505. The difference between log prices at the 90th and 10th percentiles of the distribution is 3.71, which implies that the prices at the 90th percentile are a factor of 41 higher than prices at the 10th percentile ($e^{3.71} = 41$). Next, we remove firm × product means from log export prices, so that we retain only variation across export markets within firm × products. The resulting standard deviation is 0.707, with a 90–10 ratio in levels of 3.1 (a log 90–10 difference of 1.131), implying much less variation in prices than we find when we remove only product means.[9]

This simple exercise clearly illustrates two features of our data. First, most of the variation in product-level prices is between firms rather than within. Second, there is still a substantial amount of within-firm price variation across destination markets. There are two possible explanations for this within-firm variation. The first is simple price discrimination: a firm sells the same product across 478 markets, but the markup varies. However, a 90–10 price ratio of 3 is inconsistent with a simple price discrimination explanation. To see this, let η_c be the elasticity of demand in market c, implying

[8] http://www.cepii.fr/anglaisgraph/bdd/distances.htm.

[9] Weighting observations by value has almost no effect on these numbers: the within product standard deviation is 1.489, and the within firm–product standard deviation is 0.709.

a markup of $\eta_c/(\eta_c - 1)$. If markups differ by a factor of 3 between two markets A and B, then it must be the case that

$$\eta_A = \frac{3\eta_B}{1 + 2\eta_B}.$$

This relation implies, for example, that highly elastic demand in B such as $\eta_B = 10$ or 20 coexists with an extremely inelastic value of $\eta_A \simeq 1.44$. This much cross-market variation in the demand elasticity for the exact same product seems implausible. It is also inconsistent with the possibility of even very costly arbitrage. Thus we conclude that there must be at least some compositional variation within firm × products across markets.

The analysis above (and almost all of the analysis in this paper) aggregates transactions to the level of firm–product–destination. We thus ignore variation in prices across transactions within a firm–product–destination. As a check, we repeated the above analysis using the original transaction level data, and the overall message is the same: the standard deviation within products is a bit lower at 1.272, and within product–firms it is a bit higher at 0.871.[10]

Our next results come from implementing the product-level decomposition of how export prices differ across destinations which is given by Eq. (2). The decomposition in Eq. (2) holds for each HS10 product, and to make the results comparable across products we divide by $\overline{p}_d - \overline{p}$, so that the three terms on the right hand side in Eq. (2) sum up to one for all products and destinations. We compute the scaled decomposition for the 187,300 product × destination observations in our data for 2002. The results are reported in Table 1, and illustrated vividly in Fig. 1. The figure shows that in the bulk of the cases the market share effect accounts for all or nearly all of the variation in average prices across markets, with the price discrimination and interaction terms tightly clustered around zero. This means that when it comes to explaining cross-country price differences, average price differences across firms are much more important than within-firm differences across markets. This conclusion is consistent with our simple analysis of variance for the log export prices discussed just above.

A further implication of our implementation of Eq. (2) is that the differences in product-level average prices across destination documented by Baldwin and Harrigan (2011) are due primarily to differences in which firms sell to which markets. Since tougher markets have higher product-level prices, it follows that high-price firms have larger market shares in tougher markets. Fig. 1 thus supports the mechanism conjectured by Baldwin and Harrigan (2011).

4.2. Export prices and destination market characteristics

We now look more carefully at what explains export price variation across markets. In this subsection we report the results of estimating Eqs. (6) and (7), which relate log export prices to characteristics of the destination market. Equations are estimated by the three-step selection correction procedure described above, with third-stage standard errors clustered by country. We estimate the equations on various sub-samples of the data:

- all firms/manufacturing firms only
- finished products/parts
- within-firm/arm's length transactions

[10] There are two offsetting effects that probably account for this result. If a firm charges different prices across transactions within a given product–destination then aggregating across transactions will *reduce* measured variability. However, if a firm charges the same price across transactions then aggregating across transactions will *raise* measured variability.

Table 1
Distribution of export unit price decomposition elements.
Sources: U.S. Census Bureau and authors' calculations.

Percentile	Price discrimination	Market share	Interaction
0.05	−0.416	−0.552	−1.680
0.10	−0.053	−0.003	−0.528
0.25	0.000	0.575	−0.003
0.50	0.000	1.000	0.000
0.75	0.057	1.000	0.276
0.90	0.536	1.336	0.854
0.95	1.107	2.072	1.568

Notes: This table reports the results of computing Eq. (2) across all products, scaled by product-means. The table reports quantiles of the empirical distribution of the three terms in Eq. (2).

We also report results using different specifications:

- log linear distance/distance step function
- OLS/controlling for selection
- product fixed effects/product × firm fixed effects

Our estimates in Eqs. (6) and (7) are reported in Tables 2 and 3. Table 2 reports the estimates from the broadest sample (all countries, products, and firms). The first two columns of Table 2 are the simplest: distance is measured as log kilometers, and there is no control for selection. Consistent with the simple decomposition results of the previous section, moving from product to product × firm fixed effects leads to much smaller effects of country characteristics: the distance elasticity falls from 0.267 to 0.199, the real GDP elasticity falls from 0.029 to −0.02, etc. When we control for selection the effects are smaller still: in column 4, the distance elasticity is 0.187, the real GDP effect is −0.013, and the real GDP per worker effect is statistically insignificant. The remoteness effect is statistically significant in all of the first four columns of Table 2, Panel A, but economically small: the sample standard deviation of log remoteness is 0.05, so the estimate of −1.28 in column 4 implies that a one standard deviation increase in remoteness reduces within product × firm export prices by 6 log points, which is less than one-tenth of a standard deviation.

When we allow for a non-linear effect of distance, the story changes somewhat. Focusing on our preferred specification (product × firm fixed effects, selection control) in column 8 of Table 2, we find that the effect of distance is to raise log prices by about 0.45 relative to the excluded category (Mexico and Canada). While smaller than the effect found when we control for neither selection nor firm effects (see column 5, as well as the results of Baldwin and Harrigan (2011)), this is nonetheless a large effect, both in economic terms and relative

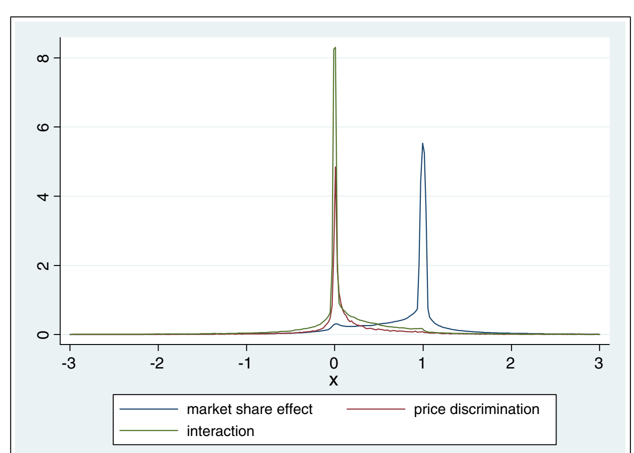

Fig. 1. Distribution of terms in Eq. (2).

Table 2
Impact of destination country characteristics on unit prices of exports.

	Loglinear distance				Distance step function			
	OLS estimation		Selection correction		OLS estimation		Selection correction	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log distance	0.267***	0.199***	0.261***	0.187***				
	(0.016)	(0.021)	(0.016)	(0.022)				
1 < km ≤ 4000					0.330***	0.401***	0.314**	0.376***
					(0.117)	(0.094)	(0.120)	(0.095)
4000 < km ≤ 7800					0.596***	0.503***	0.562***	0.458***
					(0.094)	(0.091)	(0.100)	(0.093)
7,800 < km ≤ 14,000					0.620***	0.511***	0.587***	0.466***
					(0.097)	(0.090)	(0.105)	(0.092)
14,000 < km					0.643***	0.486***	0.611***	0.445***
					(0.085)	(0.088)	(0.091)	(0.090)
Log remoteness	−2.508***	−1.369***	−2.454***	−1.280***	−1.464***	−0.344	−1.430***	−0.317
	(0.347)	(0.190)	(0.343)	(0.188)	(0.531)	(0.264)	(0.535)	(0.265)
Log real GDP	0.029**	−0.020***	0.034***	−0.013*	0.040***	−0.012**	0.047***	−0.003
	(0.011)	(0.008)	(0.011)	(0.007)	(0.014)	(0.006)	(0.014)	(0.006)
Log real GDP/worker	0.079***	−0.019	0.084***	−0.009	0.041	−0.008	0.049	0.003
	(0.026)	(0.016)	(0.026)	(0.016)	(0.047)	(0.018)	(0.046)	(0.017)
Selection control			−0.010**	−0.013***			−0.010***	−0.011***
			(0.004)	(0.002)			(0.002)	(0.001)
Product fixed effects	5019	0	5019	0	4942	0	4942	
Firm × product fixed effects	0	648,067	0	648,067	0	645,732	0	645,732
R^2 (within)	0.031	0.015	0.031	0.017	0.029	0.016	0.030	0.017
Observations	1,047,267	1,047,267	1,047,267	1,047,267	1,046,358	1,046,358	1,046,358	1,046,358

Note: Sample is all U.S. export flows over $250, by product–firm–country. Log unit price is regressed on characteristics of destination countries, controlling for product or firm–product fixed effects. Selection estimator is the three-step generalization of the Wooldridge (1995) estimator. Distance is measured in kilometers from Chicago to destination country's capital city. Standard errors are clustered at the country level. Asterisks denote statistical significance.
*** Significant at the 1% level.
** Significant at the 5% level.
* Significant at the 10% level.

to the variability in export prices.[11] Interestingly, the effect is not increasing in distance, with the estimated effects for the different distance categories all statistically insignificant from each other (although the point estimate for the smallest distance bin is about 8 log points lower than that of the other distance bins). The effects of remoteness, GDP and GDP per capita in this specification are all statistically insignificant and rather small in economic terms.

As discussed above, for computational reasons the standard errors reported in Table 2 (and all subsequent regression tables) are not adjusted due to the fact that the selection control is a generated regressor. The first-order effect of this is that the reported standard errors in the selection terms in columns 3, 4, 7 and 8 are probably underestimated. With this caveat, the selection terms are statistically significant and small, implying about a −1% correlation between the errors in the export price and export volume equations. Following the discussion in Eqs. (3)–(5) above, the negative sign on the selection terms suggests that market-specific supply shocks may be playing a bigger role than demand shocks in influencing which markets firms enter. More important than the sign and size of this correlation is the effect of controlling for selection on the estimated parameters of interest. Generally, controlling for selection leads to somewhat smaller point estimates of the effect of distance on export prices. This can be seen by comparing columns 2 and 4 (the distance elasticity falls by 0.012) and columns 6 and 8 (the distance effects fall by around 0.05). This fall in the distance effect is consistent with selection on quality: if high quality, high price firms are more likely to enter distant markets, then failing to control for selection will bias up the estimated effects of distance on price. However, the changes in the distance effect when controlling for selection are not statistically different from zero, and overall the effect on the parameters of interest is small.

[11] As discussed above, the standard deviation of log export prices across countries, with firm–product means removed, is 0.707.

Table 3 has the same specifications as in Table 2, but the sample is restricted to the manufactured products. The overall message is the same: controlling for product × firm fixed effects leads to smaller estimates of the distance effect, and controlling for selection has a small effect on the parameters of interest. The most interesting difference between Tables 2 and 3 is that in the manufacturing firm samples in Table 3 we find a statistically significant effect of market size (log real GDP) on export prices: larger markets have lower export prices with a small elasticity of −0.02 (column 8) to −0.03 (column 4). This market size effect is consistent with greater price competition in larger markets.

Our conclusions from Tables 2 and 3 can be summarized simply. Controlling for firm effects (through the use of product × firm fixed effects) and selection into exporting leads to smaller estimated effects of country characteristics on export prices than those found in specifications which include only product fixed effects. The distance effect is well approximated by a simple step function, where prices sold to markets other than Canada and Mexico are around 50 log points higher. For manufactured goods, export prices are slightly lower in larger markets.

What might account for the large within product × firm price premium for selling to countries other than Canada and Mexico? The effect seems too large to be accounted for by price discrimination, and in any case there is no particular reason to think that the demand for U.S. exports is more elastic in North America than elsewhere. Our conjecture is that the Canada/Mexico price effect has to do with vertical fragmentation and the associated trade in stages of production. As argued by Yi (2003), low transport costs (such as across a border and within a free trade area) make it possible for firms to adopt offshoring strategies that involve low-value trade transactions which would not be profitable if transport costs and/or tariffs were higher. To the extent that such trade occurs within product categories that also feature higher-value finished exports, it would explain our findings that within product × firm export prices to destinations other than Canada and Mexico are substantially higher. A complementary explanation that

Table 3
Impact of destination country characteristics on unit prices of manufacturing exports.

	Linear distance				Distance step function			
	OLS estimation		Selection correction		OLS estimation		Selection correction	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log distance	0.268***	0.203***	0.273***	0.209***				
	(0.015)	(0.021)	(0.015)	(0.020)				
1 < km ≤ 4000					0.390***	0.430***	0.398***	0.435***
					(0.115)	(0.087)	(0.112)	(0.087)
4000 < km ≤ 7800					0.610***	0.534***	0.630***	0.555***
					(0.097)	(0.086)	(0.091)	(0.084)
78,000 < km ≤ 14,000					0.631***	0.533***	0.652***	0.556***
					(0.094)	(0.084)	(0.086)	(0.082)
14,000 < km					0.646***	0.507***	0.662***	0.523***
					(0.084)	(0.081)	(0.078)	(0.080)
Log remoteness	−2.519***	−1.401***	−2.593***	−1.501***	−1.417**	−0.290	−1.453***	−0.346
	(0.338)	(0.211)	(0.347)	(0.211)	(0.541)	(0.286)	(0.543)	(0.284)
Log real GDP	0.022**	−0.025***	0.016	−0.033***	0.034**	−0.0131**	0.028**	−0.022***
	(0.011)	(0.008)	(0.012)	(0.009)	(0.013)	(0.006)	(0.014)	(0.007)
Log real GDP/worker	0.043*	−0.023	0.036	−0.033*	0.014	−0.012	0.006	−0.022
	(0.024)	(0.017)	(0.024)	(0.017)	(0.045)	(0.018)	(0.047)	(0.018)
Selection control			−0.014***	−0.018***			−0.010**	−0.013***
			(0.005)	(0.003)			(0.005)	(0.002)
Product fixed effects	4070	0	4070	0	4036	0	4036	0
Firm × product fixed effects	0	277,545	0	277,545	0	277,056	0	277,056
R^2 (within)	0.030	0.018	0.031	0.020	0.030	0.019	0.030	0.020
N	553,795	553,795	553,795	553,795	553,723	553,723	553,723	553,723

Note: Sample is U.S. export flows over $250 of manufactured goods, by product–firm–country. Log unit price is regressed on characteristics of destination countries, controlling for product or firm–product fixed effects. The selection estimator is the three-step generalization of the Wooldridge (1995) estimator. Distance measured in kilometers from Chicago to destination country's capital city. Standard errors are clustered at the country level. Asterisks denote statistical significance.
*** Significant at the 1% level.
** Significant at the 5% level.
* Significant at the 10% level.

does not involve stages of production is the Alchian–Allen effect studied by Hummels and Skiba (2004): with per-unit rather than ad valorem transport costs, higher quality goods will be relatively less expensive at greater distances, so that demand within firm–products will shift toward higher priced goods in more distant markets.

4.3. The price–distance effect: robustness

The most striking results of Tables 2 and 3 are the very large estimated distance effects on prices within product × firms (columns 2, 4, 6, and 8 of Tables 2 and 3), and in particular the large negative Canada–Mexico effect implied by the positive and flat estimated step function in distance (columns 6 and 8). In this section we investigate whether this general result is robust to various changes in specification. We consider an alternative estimator of the selection effect, as well as different subsamples.

Before turning to econometric robustness, we illustrate the price–distance effect in a figure. To construct the figure, we estimate the following regression of log export prices on firm × product fixed effects and country dummies,

$$\ln p_{ifd} = \alpha_{if} + \beta_d + \varepsilon_{ifd}.$$

The country dummy β_d has a simple interpretation as the average log price of exports sold in country d, after removing cross-country average prices by firm × product. The estimated $\hat{\beta}_d$s are plotted against distance in Fig. 2 for the 64 countries that each account for at least 0.1% of U.S. exports. The figure is consistent with the regression results in Tables 2 and 3: within firm × products, countries other than Canada and Mexico have much higher export prices, but there is little visual evidence of any other relationship between price and distance.

4.3.1. Robustness: alternative estimator of the selection effect

As discussed in Section 2.1.2 above, our benchmark estimates (reported in columns 4 and 8 of Tables 2 and 3) are computed from a three-step estimator, which is a modest generalization of the two-step estimator due to Wooldridge (1995). In Table A1, we re-estimate our benchmark specifications using Wooldridge's (1995) two-step estimator. The sample sizes are slightly smaller than those in Tables 2 and 3 because some of the first-stage Tobit estimates did not converge, but the estimates of the parameters of interest are not meaningfully different. In what follows, all the specifications with a selection correction use the three-step estimator.

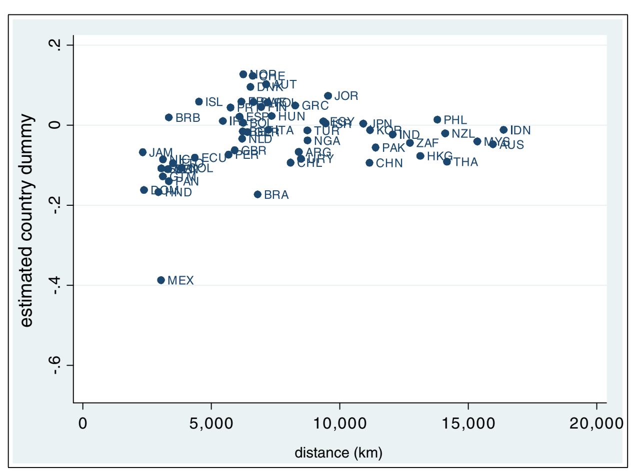

Fig. 2. Average price residuals vs. distance.

Table 4
Impact of destination characteristics on unit prices of exports: subsamples.

Sample	Benchmark	Intra-firm	Arm's-length	Parts	Parts-only
	(1)	(2)	(3)	(4)	(5)
1 < km ≤ 4000	0.376***	0.393***	0.360***	0.142*	0.124**
	(0.095)	(0.064)	(0.109)	(0.072)	(0.061)
4000 < km ≤ 7800	0.458***	0.385***	0.463***	0.254***	0.225***
	(0.093)	(0.066)	(0.107)	(0.067)	(0.055)
78,000 < km ≤ 14,000	0.466***	0.432***	0.458***	0.285***	0.258***
	(0.092)	(0.060)	(0.107)	(0.068)	(0.056)
14,000 < km	0.445***	0.421***	0.443***	0.293***	0.267***
	(0.090)	(0.055)	(0.105)	(0.069)	(0.057)
Log remoteness	−0.317	−0.403	−0.342	−0.729***	−0.716***
	(0.265)	(0.346)	(0.261)	(0.242)	(0.258)
Log real GDP	−0.003	0.016*	−0.005	−0.015*	−0.012
	(0.006)	(0.008)	(0.005)	(0.009)	(0.009)
Log real GDP/worker	0.003	0.039**	−0.006	−0.001	0.006
	(0.017)	(0.019)	(0.018)	(0.016)	(0.015)
Selection control	−0.011***	−0.005***	−0.006***	−0.040***	−0.053***
	(0.001)	(0.001)	(0.001)	(0.002)	(0.003)
Firm × product fixed effects	645,732	86,240	576,808	37,790	29,565
R^2 (within)	0.017	0.016	0.014	0.018	0.021
N	1,046,358	138,780	928,168	67,553	55,181

Note: See notes in Table 2 for details on the specification. Column (1) reproduces column (8) in Table 1. Columns (2) through (5) estimate the same specification as in column (1) using different subsamples of the data. Column (2) includes transactions where the importing and exporting firm are part of the same multinational firm, while column (3) includes only transactions where there is no such relationship. Column (4) includes only HS10 products whose description includes the word "PART", and column (5) includes the subset of products from column (4) whose description includes the phrase "PARTS ONLY". Asterisks denote statistical significance.
*** Significant at the 1% level.
** Significant at the 5% level.
* Significant at the 10% level.

4.3.2. Robustness: subsamples

The estimates reported in Tables 2 and 3 include all available observations. In this subsection we discuss how our inferences about the price–distance effect are affected by looking at different subsamples of products. We focus on our preferred specification, which corresponds to column 8 in Table 2: product × firm fixed effects, distance measured by a step function, and estimation using the three-step sample selection estimator.

We consider two different splits of the data:

1. intra-firm/arm's-length: When trade occurs within the boundary of a firm, the data record this as a "related party" or intra-firm transaction.[12] Otherwise, the transaction is recorded as being at "arm's length", which means that there is no ownership relationship between the U.S. exporting firm and the foreign importing firm.
2. intermediate/final goods: The product classification system does not distinguish between intermediate and final goods. In an attempt to get at this potentially important distinction, we follow Schott (2004) and classify an HS10 code as "parts" if the word "parts" appears in the text describing the code. A smaller number of codes include the phrase "parts only", so we also create a split based on this tighter definition of an intermediate good.

Table 4 reports the results of our analysis of these splits. The first column of the table reproduces column 8 of Table 2, which is our benchmark that includes all products. Columns 2 and 3 are estimated on samples classified by intra-firm/arm's length, and the results show that this split doesn't matter: the results on the subsamples are economically and statistically indistinguishable from the benchmark results (the only exception is that for the intra-firm sample we find a positive and statistically significant effect of GDP per worker).[13]

In columns 4 and 5 we restrict our sample to HS codes that we classify as parts, with striking results: the Canada–Mexico effect is still positive, but the effect is about 20 log points smaller than that found in the full sample reported in column 1. We conjecture that this is because there is less scope for within-product composition effects in the parts categories than in the excluded final goods categories. This interpretation is strengthened by comparing columns 6 and 7: when using the stricter definition of parts used in column 7, we find somewhat smaller point estimates than those in column 6, which includes categories that mix final goods and parts.

In addition to the robustness analysis reported here, we also checked the robustness of our results in other subsamples and specifications. These checks included

1. Estimation on goods classified as differentiated, homogeneous, and reference priced, following Rauch (1999).
2. Additional country characteristics often used in the gravity literature, including indicators for landlocked, English speaking, island, and WTO membership.

All of our conclusions based on our reported results are robust to these unreported results.

4.4. Comparing our results to existing literature

There are four recent papers that also analyze firm-level export pricing across markets: Martin (2012) for France, Bastos and Silva (2010) for Portugal, Gorg et al. (2010) for Hungary, and Manova and Zhang (2012) for China. Each of these four papers works with a specification similar to our Eq. (4), and each finds that firm-level export prices are

12 For the purposes of the trade data, firms are defined as "related parties" if one of them owns, controls, or holds voting power equivalent to 6% of the outstanding voting stock or shares of the other.

13 The sample sizes in columns 2 and 3 sum up to slightly more than the sample size in column 1. The reason is that for a small number of firm–product–destinations, there are some transactions that are intra-firm and others that are at arm's-length. These transactions are aggregated in column 1 but appear separately in columns 2 and 3.

Table 5
Impact of firm characteristics on unit prices of manufacturing exports.

	All	All	Vessel	Air	Intra-firm	Arm's-length
	(1)	(2)	(3)	(4)	(5)	(6)
Log TFP	0.329***	0.320***	0.324***	0.357***	0.185*	0.334***
	(0.051)	(0.051)	(0.080)	(0.096)	(0.102)	(0.049)
Log S/L	0.165***	0.162***	0.119***	0.159***	0.151***	0.174***
	(0.011)	(0.012)	(0.014)	(0.020)	(0.028)	(0.011)
Log K/L	−0.073***	−0.078***	−0.084***	−0.058**	−0.046*	−0.077***
	(0.011)	(0.011)	(0.016)	(0.023)	(0.024)	(0.011)
Log L	−0.008*	−0.013***	0.000	−0.003	0.002	−0.008
	(0.004)	(0.005)	(0.006)	(0.009)	(0.008)	(0.005)
Selection control		−0.032***	0.004**	−0.019***	−0.022***	−0.030***
		(0.002)	(0.002)	(0.004)	(0.005)	(0.002)
Country × product fixed effects	112,049	112,049	54,579	52,644	27,827	104,270
R^2 (within)	0.012	0.014	0.011	0.009	0.010	0.014
N	553,723	553,723	124,515	226,451	97,650	473,349

Note: Sample in columns (1) and (2) is U.S. export flows over $250 of manufactured goods, by product–firm–country. Log unit price is regressed on characteristics of firms, controlling for destination country–product fixed effects. TFP is total factor productivity, S is non-production workers, K is capital stock, and L is total employment (production plus non-production). Selection estimator is the three-step generalization of the Wooldridge (1995) estimator. In the first step, distance is measured using the bins defined in Tables 2 and 3. Standard errors are clustered by firm. Asterisks denote statistical significance. Column (3) includes only transactions shipped by vessel, and column (4) includes only transaction shipped by air. Column (5) includes transactions where the importing and exporting firms are part of the same multinational firm, while column (6) includes only transactions where there is no such relationship. Asterisks denote statistical significance.
*** Significant at the 1% level.
** Significant at the 5% level.
* Significant at the 10% level.

systematically correlated with destination market characteristics. Here we discuss the relevant results of these four papers in some detail.

Manova and Zhang (2012) analyze firm-level export prices from China. They perform a wealth of interesting empirical exercises, including estimating an equation which is essentially identical to our Eq. (4). They find small but statistically significant elasticities of within-firm export prices with respect to export market GDP, distance, and remoteness: for example, their estimated distance elasticity is about 0.01, with a standard error of about 0.002.[14]

Unlike China, France is economically similar to the United States, so it might be reasonable to expect that French and U.S. export prices would behave similarly. Martin (2012) finds no effect of real GDP on French firm-level export prices, but he does find substantial effects of distance: for example, export prices are 11 to 14 log points higher for markets that are at least 3000 km from Paris, when compared to more nearby destinations.[15] The most direct comparison between Martin's results and ours is between his Table 2 and our Table 2. Martin (2012) finds a distance elasticity of between 0.02 and 0.05 with standard errors of around 0.01, while in column 4 of our Table 2 we find an elasticity of 0.19 with a standard error of 0.02. We regard these results as quantitatively similar, although our point estimate is somewhat bigger.

The results of Bastos and Silva (2010) for Portugal are quite consistent with the results of Martin (2012) for France. In the specification closest to our Eq. (4),[16] without a selection correction, Bastos and Silva (2010) find a distance elasticity of around 0.05 with a standard error of 0.013. Like Martin (2012) and we, Bastos and Silva (2010) find no effect of market size on export prices, though they do find a small positive effect of GDP per worker in some specifications.

Gorg et al. (2010) look at firm level export prices for Hungary. Results from their version of our Eq. (4), without a selection correction, are remarkably consistent with the results of Martin (2012) and Bastos and Silva (2010): a distance elasticity of between 0.05 and 0.07 depending on the year, with standard errors of about 0.02.[17] Unlike the other three papers discussed here, Gorg et al. (2010) make an attempt to address the selection issue in later specifications, but they do so in a model without product × firm fixed effects. This makes their results that correct for selection both hard to interpret and not comparable to ours, since their parameters are identified using cross-firm and cross-product variation.

In summary, our results are quite consistent with the results of the four previous papers that have looked at export price variation within product × firms. Data from France, Portugal, and Hungary all give essentially the same answer: within product–firms and across export destinations, the distance elasticity of export unit values is close to 0.05, with a 95% confidence interval of about [0.03, 0.07]. The results from the Chinese data show smaller elasticity, while our results for the United States are somewhat higher. The papers just discussed do not analyze the connection between firm characteristics and export pricing, which is the relationship that we estimate next.

4.5. Export prices and firm characteristics

We now turn to the estimation in Eq. (11), which relates export prices to firm characteristics. Eq. (11) includes country × product fixed effects, so that the estimated effects of firm characteristics are identified purely across firms, within country–products. Because we only have data on the characteristics of manufacturing firms, all the results in this section are for manufacturing firms.

Table 5 reports our estimates. Columns (1) and (2) include all observations, with column (2) controlling for selection. The results are striking: we find that more productive firms charge *higher* prices on average. The TFP elasticity of 0.32 means that firms with a 10% higher total factor productivity charge about 3% higher prices. Equally striking are the large and precisely estimated effects of factor shares on export prices: skill intensity raises export prices with an elasticity of about 0.16, while capital intensity lowers prices with an elasticity of around −0.08. The effect of firm size is very small and negative at −0.01: bigger firms charge very slightly lower prices.[18]

[14] Table 7, columns 5 and 6 (Manova and Zhang, 2012).
[15] Table 3 (Martin, 2012).
[16] Reported in columns 5, 6, 11, and 12 of Table 6 (Bastos and Silva, 2010).
[17] Reported in Table 2 (Gorg et al., 2010). It appears that these standard errors are clustered by importing country, as is appropriate.
[18] Table A2 reports estimation results using the two-step estimator. The differences with Table 5 are inconsequential.

Table 6
Impact of interactions of importing country and firm characteristics on export unit prices of manufacturing firms.

	Main effect	Interactions of firm characteristics (row) with country characteristics (column)						
		1 < km ≤ 4000	4000 < km ≤ 7800	7800 < km ≤ 14,000	14,000 < km	Log remoteness	Log GDP	Log GDP/worker
Log TFP	0.799	−0.002	0.435***	0.280***	0.336***	0.393	−0.036	−0.042
	(0.594)	(0.127)	(0.095)	(0.078)	(0.101)	(0.740)	(0.024)	(0.049)
Log S/L	0.266**	0.060**	0.083***	0.071***	0.101***	−0.013	0.010**	−0.021**
	(0.107)	(0.024)	(0.025)	(0.017)	(0.020)	(0.160)	(0.005)	(0.009)
Log K/L	0.013	−0.055**	−0.104***	−0.039**	−0.085***	−0.222	−0.014***	0.004
	(0.102)	(0.022)	(0.020)	(0.016)	(0.020)	(0.149)	(0.005)	(0.008)
Log L	−0.048	−0.009	−0.002	−0.006	−0.010	−0.019	0.002	0.003
	(0.040)	(0.009)	(0.008)	(0.006)	(0.008)	(0.054)	(0.002)	(0.003)
Selection control	−0.033***							
	(0.002)							
Country × product fixed effects	112,049							
R^2 (within)	0.015							
N	553,723							

Note: Sample is U.S. export flows over $250 of manufactured goods, by product–firm–country. Log unit price is regressed on characteristics of firms, alone and interacted with importing country characteristics, controlling for importing country–product fixed effects. Selection estimator is the three-step generalization of the Wooldridge (1995) estimator. In the first step, distance is measured using the bins defined in Tables 2 and 3. Standard errors are clustered by firm. Asterisks denote statistical significance.
*** Significant at the 1% level.
** Significant at the 5% level.
* Significant at the 10% level.

Columns (3) through (6) estimate Eq. (11) on subsamples.[19] Columns (3) and (4) show that the relationship is invariant to shipping mode, with the estimates neither economically nor statistically different from the benchmark estimates of column (2).[20] The same is true when shipments are disaggregated by intra-firm vs. arm's-length: the point estimate on TFP in column (5) is smaller than that in column (2), but the difference is not statistically significant, and the other parameters in both columns (5) and (6) are similarly no different from those in column (2). We also estimated Eq. (11) on subsamples defined by one- and two-digit HS codes, with no material change in any of our conclusions. One conclusion from the analysis of subsamples is that the negative estimated effect of firm scale on price is fragile, as it disappears in the subsample regressions.

Our conclusions from this section are quite strong: firms that are more productive and more skill-intensive charge substantially higher prices, while more capital-intensive firms charge lower prices. We emphasize how we identify these effects: they are found across firms within narrowly defined destination–product markets. If HS10 products were homogeneous, the law of one price would imply that our results are impossible: the lowest price firm would simply take the entire market. The fact that highly productive, skill-intensive firms charge higher prices is suggestive of quality competition: the higher measured prices in our data are probably hiding important quality variations across firms, with higher quality associated with higher costs and thus higher prices. This interpretation is consistent with the evidence of Gervais (2011), who uses plant-level data from the U.S. Census of Manufactures to show that higher quality firms have both higher productivity and charge higher prices.

4.6. Export prices, firm characteristics, and country characteristics

As a final empirical exercise, we estimate the effects of manufacturing firm characteristics interacted with country characteristics. This interaction model is given by Eq. (12), and we use a specification with a selection control and distance measured using bins. The estimates are reported in Table 6.

The interaction terms reported in Table 6 are the answers to a somewhat subtle question, expressed by Eq. (14): how does the variation of export prices with respect to country characteristics depend on the characteristics of the firm? Most of the interaction effects are not statistically significantly different from zero, so the answer to the question is "not much" or "we don't know". The exceptions are the interaction terms between the firm characteristics TFP, skill intensity and capital intensity and the distance bins. Relative to Canada/Mexico, prices in more distant locations are lower for capital intensive firms and higher for high TFP and skill intensive firms. The biggest effect is clearly TFP: in countries more than 4000 km from the U.S., high TFP firms charge 28 to 43 log points higher prices than they do in nearby countries. This suggests that the aggregate distance effects reported in column (8) of

[19] Unlike in Table 4, we do not report subsamples of parts and parts only here in Table 6. These results were not approved for release by the Census Bureau.
[20] We did not consider the transport subsamples in Table 4 because almost all exports to Canada and Mexico are by surface modes (road and rail).

Table A1
Impact of destination country characteristics on unit prices of exports using the two-step Tobit approach.

	All products		Manuf. products	
	(4)	(8)	(4)	(8)
Log distance	0.263***		0.266***	
	(0.020)		(0.024)	
1 < km ≤ 4000		0.372***		0.429***
		(0.091)		(0.083)
4000 < km ≤ 7800		0.511***		0.555***
		(0.081)		(0.082)
78,000 < km ≤ 14,000		0.523***		0.567***
		(0.083)		(0.084)
14,000 < km		0.491***		0.527***
		(0.076)		(0.073)
Log remoteness	−2.046***	−0.440	−2.117***	−0.410
	(0.198)	(0.303)	(0.251)	(0.313)
Log real GDP	−0.070***	−0.026**	−0.077***	−0.027*
	(0.014)	(0.013)	(0.016)	(0.015)
Log real GDP/worker	−0.083***	−0.028*	−0.093***	−0.024
	(0.020)	(0.014)	(0.022)	(0.015)
Selection control	−0.010***	−0.002	−0.010***	−0.002
	(0.003)	(0.003)	(0.004)	(0.003)
Firm × product fixed effects	638,959	634,735	276,575	275,672
R^2 (within)	0.014	0.014	0.015	0.016
Observations	1,047,653	1,039,392	553,767	548,042

Note: This table reports re-estimation of columns 4 and 8 in Tables 2 and 3. Here the selection effect is controlled for using Wooldridge's (1995) two-step estimator. The sample size is 1% smaller here than that in Tables 2 and 3 because the two-step estimator failed to converge in some cases where the three-step estimator did converge. Asterisks denote statistical significance.
*** Significant at the 1% level.
** Significant at the 5% level.
* Significant at the 10% level.

Table A2
Impact of firm characteristics on export unit prices of manufacturing firms using the two-step Tobit approach.

	OLS estimation	Selection correction
Log TFP	0.357***	0.353***
	(0.053)	(0.053)
Log S/L	0.171***	0.168***
	(0.012)	(0.012)
Log K/L	−0.076***	−0.078***
	(0.012)	(0.012)
Log total employment	−0.009*	−0.012**
	(0.005)	(0.005)
Selection control		−0.003**
		(0.001)
Country × product fixed effects	118,189	118,189
R^2 (within)	0.012	0.012
N	548,042	548,042

Note: Estimated using the two-step Wooldridge (1995) estimator instead of the three-step estimator. The results reported here should be compared to the three-step estimates reported in columns (1) and (2) of Table 5. The sample size is 1% smaller here than that in Table 5 because the two-step estimator failed to converge in some cases where the three-step estimator did converge. Asterisks denote statistical significance.
*** Significant at the 1% level.
** Significant at the 5% level.
* Significant at the 10% level.

Table 3 are driven by the behavior of high TFP, skill-intensive, and less capital-intensive firms.

4.7. Firm-level export prices: summarizing our results

To summarize our findings, we focus on the specifications with the cleanest identification: the columns in Tables 2 and 3 with product × firm fixed effects, and Table 5, which includes product × country fixed effects. The cross-country variation that is used in Tables 2 and 3 shows that firms charge systematically higher prices to destinations other than Canada and Mexico. These results may be partially explained by price discrimination, but our conjecture is that they are driven primarily by within-firm composition effects instead: firms sell fewer semi-finished products to markets other than Canada and Mexico. The analysis of subsamples in Table 4 shows that the overall distance effect is stronger within final goods categories. The interaction effects specification reported in Table 6 suggests that this composition effect is driven by high TFP and skill intensive firms.

There is only weak evidence that country characteristics other than distance are an important determinant of export prices. Remoteness generally has a statistically significant negative effect, though not in our most preferred specification (column 8 in Tables 2 and 3), but the economic magnitude is unimportant. For manufactured goods, there is evidence for a negative effect of market size, but the estimated elasticity is quite small at −0.02 (column 8, Table 3). The effect of per capita income is not statistically significantly different from zero.

The cross-firm variation that is used in Table 5 shows that more productive and skill-intensive firms charge higher prices, while more capital-intensive firms charge lower prices, and these effects are economically sizeable and precisely estimated. This pattern is suggestive of quality competition within export markets, with the most capable and skill-intensive U.S. exporters producing higher quality goods that sell for a premium over goods sold by more capital-intensive and less productive firms. Conversely, more capital intensive firms may produce more standardized products that, loosely speaking, compete on price rather than on quality.

5. Conclusion

This paper is the first to analyze firm-level data on the export pricing decisions of U.S. exporters. We use a three-step estimator to control for firm selection into different export markets. Using restricted firm-level information on exports and firm characteristics, combined with widely available data on country characteristics, we find that

- More productive and skill-intensive firms charge higher unit prices, while more capital-intensive firms charge lower prices.
- In the markets that they choose to serve, firms charge prices that are substantially higher for goods sold outside North America.
- The strong correlations between product-level prices and country characteristics found by Baldwin and Harrigan (2011) are largely due to a selection or composition bias, which is the mechanism that they conjectured but could not test with their data.

Our results on correlations between export prices and country-level characteristics are broadly consistent with earlier studies on export pricing by firms in China, France, Hungary, and Portugal. We are the first to connect firm-level characteristics to export pricing, and our results are supportive of models of monopolistic competition where firms compete on quality rather simply unit cost.

References

Anderson, J., Van Wincoop, E., 2003. Gravity with gravitas: a solution to the border puzzle. Am. Econ. Rev. 93 (1), 170–192.
Baldwin, R., Harrigan, J., 2011. Zeros, quality, and space: trade theory and trade evidence. Am. Econ. J. Microeconomics 3 (2), 60–88.
Bastos, P., Silva, J., 2010. The quality of a firm's exports: where you export to matters. J. Int. Econ. 82 (2), 99–111.
Bernard, A., Eaton, J., Jensen, J., Kortum, S., 2003. Plants and productivity in international trade. Am. Econ. Rev. 93 (4), 1268.
Bernard, A., Jensen, J., Schott, P., 2005. Importers, exporters, and multinationals: a portrait of firms in the US that trade goods. NBER Working Paper.
Bernard, A., Jensen, J., Redding, S., Schott, P., 2007. Firms in international trade. J. Econ. Perspect. 105–130.
Eaton, J., Kortum, S., 2002. Technology, geography, and trade. Econometrica 70 (5), 1741–1779.
Feenstra, R., M. Obstfeld, and K. Russ (2009). In search of the Armington elasticity. *Unpublished manuscript, University of California, Davis.*
Gervais, A. (2011, June). Product quality, firm heterogeneity, and international trade. *Unpublished manuscript, University of Notre Dame.*
Ghironi, F., Melitz, M., 2005. International trade and macroeconomic dynamics with heterogeneous firms. Q. J. Econ. 120 (3), 865–915.
Gorg, H., Halpern, L., Murakozy, B., 2010. Why do within firm–product export prices differ across markets? CEPR Working Paper 7708
Hallak, J. C. and J. Sivadasan (2009). Firms' exporting behavior under quality constraints. *Unpublished paper, University of Michigan.*
Harrigan, J., 2003. Specialization and the volume of trade: do the data obey the laws? Handbook of International Tradep. 85
Helpman, E., Melitz, M., Rubinstein, Y., 2008. Estimating trade flows: trading partners and trading volumes. Q. J. Econ. 123 (2), 441–487.
Hummels, D., Skiba, A., 2004. Shipping the good apples out? An empirical confirmation of the Alchian–Allen conjecture. J. Polit. Econ. 112 (6), 1384–1402.
Johnson, R., 2012. Trade and prices with heterogeneous firms. J. Int. Econ. 86, 53–56.
Kugler, M., Verhoogen, E., 2012. Prices, plant size, and product quality. Rev. Econ. Stud. 79 (1), 307–339.
Manova, K., Zhang, Z., 2012. Export prices across firms and destinations. Q. J. Econ. 127 (1), 379–436.
Martin, J., 2012. Markups, quality, and transport costs. Eur. Econ. Rev. 56, 777–791.
Mayer, T., Ottaviano, G., 2008. The happy few: the internationalisation of European firms. Intereconomics 43 (3), 135–148.
Melitz, M., 2003. The impact of trade on intra-industry reallocations and aggregate industry productivity. Econometrica 71 (6), 1695–1725.
Novy, D., 2013. Gravity redux: measuring international trade costs with panel data. Econ. Inq. 51 (1), 101–121.
Olley, G., Pakes, A., 1996. The dynamics of productivity in the telecommunications equipment industry. Econometrica 64 (6), 1263–1297.
Rauch, J.E., 1999. Networks versus markets in international trade. J. Int. Econ. 48 (1), 7–35.
Schott, P., 2004. Across-product versus within-product specialization in international trade. Q. J. Econ. 119 (2), 647–678.
Verhoogen, E., 2008. Trade, quality upgrading, and wage inequality in the Mexican manufacturing sector. Q. J. Econ. 123 (2), 489–530.
Wagner, J., 2012. International trade and firm performance: a survey of empirical studies since 2006. Rev. World Econ. 148, 235–267.
Wooldridge, J., 1995. Selection corrections for panel data models under conditional mean independence assumptions. J. Econ. 68 (1), 115–132.
Yi, K., 2003. Can vertical specialization explain the growth of world trade? J. Polit. Econ. 111 (1), 52–102.