What do Exporters Know?*

Michael J. Dickstein

New York University and NBER

Eduardo Morales

Princeton University and NBER

May 5, 2018

Abstract

Much of the variation in international trade volume is driven by firms' extensive margin decisions of whether to participate in export markets. We evaluate how the information potential exporters possess influences their decisions. To do so, we estimate a model of export participation in which firms weigh the fixed costs of exporting against the forecasted profits from serving a foreign market. We adopt a moment inequality approach, placing weak assumptions on firms' expectations. The framework allows us to test whether firms differ in the information they have about foreign markets. We find that larger firms possess better knowledge of market conditions in foreign countries, even when those firms have not exported in the past. Quantifying the value of information, we show that, in a typical destination, total exports rise while the number of exporters falls when firms have access to better information to forecast export revenues.

Keywords: export participation, demand under uncertainty, discrete choice methods, moment inequalities

*We are grateful to the editor and five anonymous referees for helpful comments and suggestions. We also thank Tim Bresnahan, Lorenzo Caliendo, Juan Manuel Castro, Jan De Loecker, Dave Donaldson, Jonathan Eaton, Liran Einav, Alon Eizenberg, Guido Imbens, Michal Kolesar, Ariel Pakes, Esteban Rossi-Hansberg, James Tybout and seminar participants at the CEPR-JIE conference on Applied Industrial Organization, Columbia University, Dallas Fed, Dartmouth College, LMU, MIT, McGill University, New York University, the NBER ITI meeting, University of Notre Dame, Northwestern University, Pennsylvania State University, Princeton University, the Stanford/Berkeley IO Fest, Stanford University, UCLA, University of Maryland, University of Minnesota, University of Pennsylvania, University of Virginia, and Yale University for helpful suggestions. The Appendix referenced in the text is available as an Online Appendix. Email: michael.dickstein@nyu.edu, ecmorale@princeton.edu.

1 Introduction

In 2014, approximately 300,000 US firms chose to export to foreign markets (Department of Commerce, 2016). The decision of these firms to sell abroad drives much of the variation in trade volume from the US (Bernard et al., 2010). Thus, to predict how exports may change with lower trade costs, exchange rate movements, or other policy or market fluctuations, researchers need to understand firms' decisions of whether to participate in export markets.

A large literature in international trade focuses on modeling firms' export decisions.[1] Empirical analyses of these decisions, however, face a data obstacle: the decision to export depends on a firm's expectation of the profits it will earn when serving a foreign market, which the researcher rarely observes. Absent direct data on firms' expectations, researchers must impose assumptions on how firms form these expectations. For example, researchers commonly assume firms' expectations are rational and depend on a set of variables observed in the data. The precise specification of agents' information, however, can influence the overall measurement, as Manski (1993, 2004), and Cunha and Heckman (2007) show in the context of evaluating the returns to schooling. In an export setting, the assumptions on expectations may affect both the estimates of the costs firms incur when exporting and predictions of how firms will respond to counterfactual changes in these trade costs.

In this paper, we first document that estimates of the parameters underlying firms' export decisions depend heavily on how researchers specify the firm's expectations. We compare the predictions of a standard model in the international trade literature (Melitz, 2003) under two specifications: the "perfect foresight" case, under which we assume firms perfectly predict observed profits from exporting, and a minimal information case, under which we assume firms use a specific set of observed variables to predict their export profits. For each case, we recover the fixed costs of exporting and the mean profits of firms predicted to export. Finding important differences in the predictions from the two models, we then estimate an empirical model of export participation that places fewer restrictions on firms' expectations.

Under our alternative approach, we do not require the researcher to have full knowledge of an exporter's information set. Instead, the researcher need only specify a subset of the variables that agents use to form their expectations. The researcher must observe this subset, but need not observe any remaining variables that affect the firm's expectations. The set of unobserved variables may vary flexibly across firms, markets, and years.

The trade-off from specifying only a subset of the firm's information is that we can only partially identify the parameters of interest. To do so, we develop a new type of moment inequality, the odds-based inequality, and combine it with inequalities based on revealed preference. Using these inequalities, we show that placing fewer assumptions on expectations

[1] See Melitz (2003), Das et al. (2007), Eaton et al. (2008), Arkolakis (2010), Moxnes (2010), Eaton et al. (2011), Eaton et al. (2014), Arkolakis et al. (2015), Cherkashin et al. (2015) and Ruhl and Willis (2017). The literature has also recently focused on the decisions of importers (Blaum et al., 2017; Antràs et al., 2017).

affects the measurement of the parameters of the exporter's problem. Further, we show our approach generates bounds on these parameters that are tight enough to be informative.

This paper makes four main contributions. First, we demonstrate the sensitivity of the estimated export fixed costs to assumptions the researcher imposes on firms' export profit forecasts. Second, we employ moment inequalities to partially identify the exporter's fixed costs under weak assumptions, applying insights from Pakes (2010) and Pakes et al. (2015). Third, we address the question of "what do exporters know?". We show that, under rational expectations, our moment inequality framework allows us to test whether potential exporters know and use specific variables to predict their export profits. Finally, fourth, we use our model's estimates to quantify the value of information.

To illustrate the sensitivity of export fixed estimates to the researcher's assumptions on exporters' information, we start by estimating a perfect foresight model under which firms perfectly predict the profits they will earn upon entry. Using maximum likelihood, we find fixed costs in the chemicals sector from Chile to Argentina, Japan, and the United States to equal \$868,000, \$2.6 million, and \$1.6 million, respectively. We compare these estimates to an alternative approach, suggested in Manski (1991) and Ahn and Manski (1993), in which we assume that firms' expectations are rational and we specify the variables firms use to form their expectations. Specifically, we assume that firms know only three variables: distance to the export market, aggregate exports from Chile to that market in the prior year, and their own domestic sales in the prior year. We estimate fixed costs of exporting under this approach that are 40-60% smaller than those found under the perfect foresight assumption.

That the fixed cost estimates differ under the two approaches reflects a bias in the estimation. Both require the researcher to specify the content of the agent's information set. If firms actually employ a different set of variables–containing more information or less– to predict their potential export profits, the estimates of the model parameters will generally be biased. Specifically, if the researcher wrongly assumes that firms have perfect foresight, the bias arises for a similar reason to the bias affecting Ordinary Least Squares estimates in linear models when a covariate contains classical measurement error; we show that, in our setting, this bias leads the researcher to overestimate fixed export costs. Thus, we move to employ moment inequalities to partially identify the exporter's fixed costs under weaker assumptions.

Here, we again assume that firms know the distance to the export market, the aggregate exports to that market in the prior year, and their own domestic sales from the prior year. However, unlike the minimal information approach described earlier, the inequalities we define do not restrict firms to use only these three variables when forecasting their potential export profits. We require only that firms know at least these variables. Using our inequalities, we find much lower fixed costs, representing only 10-15% of the perfect foresight values.

Comparing these findings to those in the existing literature is not simple. Our baseline model abstracts from other possible sources of bias, including marketing costs (Arkolakis, 2010), adjustment costs (Ruhl and Willis, 2017), persistent unobserved heterogeneity (Roberts and Tybout, 1997; Das et al., 2007), and buyer-seller relationships (Eaton et al., 2016; Bernard et al., 2017). In extensions to our baseline model, we account for path dependence in export status, as in Das et al. (2007), and allow firms to decide which markets to enter in reaction to unobserved (to the researcher) firm-country specific revenue shocks, as in Eaton et al. (2011).

We next employ our framework to investigate the set of information potential exporters use to forecast export revenues. We run alternative versions of our moment inequality model, holding fixed the model and data but varying the firm's presumed information set. Using the specification tests described in Bugni et al. (2015), we look for evidence against the null that potential exporters use particular variables in their forecasts.

We begin by testing our baseline assumption that exporters know *at least* distance, their own lagged domestic sales, and lagged aggregate exports when making their export decisions. Using data from both the chemicals and food sectors, we cannot reject this null hypothesis. We then test (a) whether firms have perfect foresight about their potential export profits in every country, and (b) whether firms have information on last period's realizations of a destination-time period specific shifter of firms' export revenues that, according to our model, is a sufficient statistic for the effect of all foreign market characteristics (i.e. market size, price index, trade costs and demand shifters) on these revenues. In both sectors, we reject the null that firms have perfect foresight. For the market-specific revenue shifters, we find interesting heterogeneity: we fail to reject that large firms know these shocks, but reject that small firms do. This distinction is not driven by prior export experience. That is, even when we focus only on large firms that did not to export in the previous year, we nonetheless fail to reject the null that these firms use knowledge of past revenue shifters when forecasting their potential export revenue. Large firms therefore have not only a productivity advantage over small firms, but also an informational advantage in foreign markets.

Finally, we use our model's estimates to quantify the value of information. Using our estimated bounds on fixed export costs, we compute counterfactual entry decisions, firm-level profits, and aggregate exports to each destination and in each year under different firm information sets. We find that, as we provide information to potential exporters, these firms choose to export to fewer markets: in the chemicals sector, the expected number of firm-destination pairs with positive export flows *decreases* between 3.5% and 5.7%. Interestingly, although the total number of firm-destination pairs decreases, the overall (aggregated across firms and destinations) export revenue in the sector *increases* between 6.4% and 9.5%. Were all firms able to access information on past export revenue shifters, fewer firms would make mistakes when choosing export markets and, consequently, the average firm's ex post profits in a typical market would increase between 17.5% and 20.6%. In comparison, with information to predict export revenues perfectly, the average firm's ex post profits in a typical market would increase between 46% and 52.9%.

We demonstrate our contributions using the exporter's problem. However, our estimation approach can apply more broadly to discrete choice decisions that depend on agents' forecasts of key payoff-relevant variables. For example, to determine whether to invest in R&D projects, firms must form expectations about the success of the research activity (Aw et al., 2011; Doraszleski and Jaumandreu, 2013; Bilir and Morales, 2016). When a firm develops a new product, it must form expectations of its future demand (Bernard et al., 2010; Bilbiie et al., 2012; Arkolakis et al., 2015). Firms deciding whether to enter health insurance markets must form expectations about the type of health risks that will enroll in their plans (Dickstein et al., 2015). Firms paying fixed costs to import from foreign markets must form expectations about the sourcing potential of these markets (Blaum et al., 2017; Antràs et al., 2017). Finally, in education, the decision to attend college crucially depends on potential students' expectations about earnings with and without a college education (Freeman, 1971; Willis and Rosen, 1979; Manski and Wise, 1983). In these settings, even without direct elicitation of agents' preferences (Manski, 2004), our approach allows the researcher both to test whether certain covariates belong to the agent's information set and to recover bounds on the economic primitives of the agent's problem without imposing strong assumptions on her expectations.

Our estimation approach contributes to a growing empirical literature that employs moment inequalities derived from revealed preference arguments, including Ho (2009), Holmes (2011), Crawford and Yurukoglu (2012), Ho and Pakes (2014), Eizenberg (2014), Morales et al. (2017), Wollman (2017), and Maini and Pammolli (2017). We follow a methodology closest to Morales et al. (2017), but add two features. First, we introduce inequalities in a setting with structural errors that are specific to each observation. The cost of allowing this flexibility is that we must assume a distribution for these structural errors, up to a scale parameter. We also cannot handle large choice sets, such as those considered in Morales et al. (2017). Second, we combine the revealed preference inequalities employed in the prior literature with our new odds-based inequalities, to gain identification power.

We proceed in this paper by first describing our model of firm exports in Section 2. We describe our data in Section 3. In sections 4 and 5, we discuss three alternative estimation approaches and compare the resulting parameter estimates. In sections 6 and 7, we use our moment inequalities both to test alternative information sets and to conduct counterfactuals on the value of information. In Section 8, we discuss extensions of our baseline model. Section 9 concludes. All appendix sections referenced below appear in an Online Appendix.

2 Empirical Model

We model firms' export decisions. All firms located in a country h choose whether to sell in each export market j. We index the firms located in h and active at period t by $i = 1, \ldots, N_t$.[2]

[2] We eliminate the subindex for the country of origin h when possible to simplify notation.

We index the potential destination countries by $j = 1, \ldots, J$.

We model firms' export decisions using a two-period model. In the first period, firms choose the set of countries to which they wish to export. To participate in a market, firms must pay a fixed export cost.[3] When choosing among export destinations, firms may differ in their degree of uncertainty about the profits they will obtain upon exporting. In the second period, conditional on entering a foreign market, all firms acquire the information needed to set their prices optimally and obtain the corresponding export profits.

2.1 Demand, Supply, Market Structure, and Information

Firms face isoelastic demand in every country: $x_{ijt} = \zeta_{ijt}^{1-\eta} p_{ijt}^{-\eta} P_{jt}^{\eta-1} Y_{jt}$. Here, the quantity demanded x_{ijt} depends on: p_{ijt}, the price firm i sets in destination j at t; Y_{jt}, the total expenditure in the sector in which i operates; P_{jt}, a price index that captures the competition firm i faces in market j from other firms selling in the market; and, ζ_{ijt}, a demand shifter.

Firm i produces one unit of output with a constant marginal cost c_{it}.[4] When firm i chooses to sell in a market j, it must pay two export costs: a variable cost, τ_{ijt}, and a fixed cost, f_{ijt}. We adopt the "iceberg" specification of variable export costs and assume that firm i must ship τ_{ijt} units of output to country j for one unit to arrive. The total marginal cost for firm i of selling one unit in country j at period t is thus $\tau_{ijt} c_{it}$. Fixed costs f_{ijt} are paid by firms selling a positive amount in market j at period t, independently of the actual quantity exported.

We denote the firm's potential sales revenue in market j and period t as $r_{ijt} \equiv x_{ijt} p_{ijt}$, and use $\mathcal{J}_{ijt}$ to denote the information firm i possesses about its potential revenue r_{ijt} when deciding whether to participate in market j at t. We assume firm i knows the determinants of fixed costs f_{ijt} for every country j when deciding whether to export. Therefore, if relevant to predict r_{ijt}, these determinants of fixed costs will also enter $\mathcal{J}_{ijt}$.

2.2 Export Revenue

Upon entering a market, a firm observes both η and its marginal cost of selling in this market, and sets its price optimally taking other sellers' prices as given: $p_{ijt} = (\eta/(\eta - 1)) \tau_{ijt} c_{it}$. Thus, the revenue firm i would obtain if it were to sell in market j at period t is:

$$r_{ijt} = \left[\frac{\eta}{\eta - 1} \frac{\tau_{ijt} c_{it}}{\zeta_{ijt} P_{jt}} \right]^{1-\eta} Y_{jt}. \tag{1}$$

We can write an analogous expression for the sales revenue in the domestic market h. As we

[3] In Section 8.1, we consider a fully dynamic export model in which forward-looking firms must also pay a sunk export entry cost as in Das et al. (2007).

[4] The assumption of constant marginal costs is necessary for the export decision to be independent across markets. See Vannoorenberghe (2012), Blum et al. (2013), and Almunia et al. (2018) for models of firms' export decisions with increasing marginal costs.

show in Appendix A.1, taking the ratio of export revenue to domestic revenue for each firm in year t, we can rewrite potential export revenues in a destination market j as

$$r_{ijt} = \alpha_{ijt} r_{iht}, \qquad \text{with} \qquad \alpha_{ijt} \equiv \left(\frac{\zeta_{iht}\tau_{ijt}}{\zeta_{ijt}\tau_{iht}} \frac{P_{ht}}{P_{jt}} \right)^{1-\eta} \frac{Y_{jt}}{Y_{ht}}.$$

Here, α_{ijt} is a firm-destination-year specific shifter of export revenues that accounts for the destination's market size, price index, and the effect of variable trade costs and demand shocks across firms. We can further split this shifter into a component common to firms in a given market and year, and a component that varies across firms:

$$r_{ijt} = \alpha_{jt} r_{iht} + e_{ijt}, \qquad \text{where} \qquad \alpha_{jt} \equiv \mathbb{E}_{jt} \left[\left(\frac{\zeta_{iht}\tau_{ijt}}{\zeta_{ijt}\tau_{iht}} \frac{P_{ht}}{P_{jt}} \right)^{1-\eta} \frac{Y_{jt}}{Y_{ht}} \right], \tag{2}$$

where $\mathbb{E}_{jt}[\cdot]$ denotes the mean across firms in a given country-year pair jt. The term e_{ijt} accounts for firm-market-year specific relative revenue shocks. We assume firms do not know these shocks when deciding whether to export to market j at period t:[5]

$$\mathbb{E}_{jt}[e_{ijt} | \mathcal{J}_{ijt}, r_{iht}, f_{ijt}] = 0. \tag{3}$$

Conversely, we do not restrict the relationship between the information set $\mathcal{J}_{ijt}$ and the component of revenue $\alpha_{jt} r_{iht}$. Thus, for example, more productive firms may be systematically better informed than less productive firms about variables affecting their future domestic sales, r_{iht}, or about the country-year export shifters accounted for by the term α_{jt}. Similarly, we allow firms to have more information about markets that are closer to the domestic market.

2.3 Export Profits

We model the export profits that i would obtain in j if it were to export at t as

$$\pi_{ijt} = \eta^{-1} r_{ijt} - f_{ijt}. \tag{4}$$

We model fixed export costs as

$$f_{ijt} = \beta_0 + \beta_1 dist_j + \nu_{ijt}, \tag{5}$$

where $dist_j$ denotes the distance from country h to country j, and the term ν_{ijt} represents determinants of f_{ijt} that the researcher does not observe. As discussed in Section 2.1, we

[5] Appendix A.2 describes a set of assumptions on the distribution of demand shifters ζ_{ijt} and variable costs τ_{ijt} under which the mean independence condition in equation (3) holds. Furthermore, we extend the model in Section 8.2 to allow for firm-country-year specific export revenue shocks that are known to the firm when it decides whether to export: $r_{ijt} = \alpha_{jt} r_{iht} + e_{ijt} + \omega_{ijt}$ with $\mathbb{E}_{jt}[\omega_{ijt} | \mathcal{J}_{ijt}] = \omega_{ijt}$.

assume that firms know f_{ijt} when deciding whether to export to j at t.[6]

The estimation procedure introduced in Section 4.2 requires ν_{ijt} to be distributed independently of $\mathcal{J}_{ijt}$, and its distribution to be known up to a scale parameter. To match one typical binary choice model, we assume ν_{ijt} follows a normal distribution and is independent of other export determinants:[7]

$$\nu_{ijt}|(\mathcal{J}_{ijt}, dist_j) \sim \mathbb{N}(0, \sigma^2). \tag{6}$$

The assumed independence between ν_{ijt} and $\mathcal{J}_{ijt}$ implies that knowledge of ν_{ijt} is irrelevant to compute the firm's expected export revenue. However, we impose no assumption on the relationship between $\mathcal{J}_{ijt}$ and the observed determinants of fixed export costs, here $dist_j$.

2.4 Decision to Export

A risk-neutral firm i will decide to export to j in year t if and only if $\mathbb{E}[\pi_{ijt}|\mathcal{J}_{ijt}, dist_j, \nu_{ijt}] \geq 0$, where the vector $(\mathcal{J}_{ijt}, dist_j, \nu_{ijt})$ includes any variables firm i uses to predict potential export profits in country j. Combining equations (4) and (5), we can write

$$\mathbb{E}[\pi_{ijt}|\mathcal{J}_{ijt}, dist_j, \nu_{ijt}] = \eta^{-1}\mathbb{E}[r_{ijt}|\mathcal{J}_{ijt}] - \beta_0 - \beta_1 dist_j - \nu_{ijt}. \tag{7}$$

Here, $\mathbb{E}[r_{ijt}|\mathcal{J}_{ijt}, dist_j, \nu_{ijt}] = \mathbb{E}[r_{ijt}|\mathcal{J}_{ijt}]$, following our definition of $\mathcal{J}_{ijt}$ as the set of variables firm i uses to predict r_{ijt}. Given the expression for r_{ijt} in equations (2) and (3), we write:

$$\mathbb{E}[\pi_{ijt}|\mathcal{J}_{ijt}, dist_j, \nu_{ijt}] = \eta^{-1}\mathbb{E}[\alpha_{jt} r_{iht}|\mathcal{J}_{ijt}] - \beta_0 - \beta_1 dist_j - \nu_{ijt}. \tag{8}$$

Let $d_{ijt} = \mathbb{1}\{\mathbb{E}[\pi_{ijt}|\mathcal{J}_{ijt}, dist_j, \nu_{ijt}] \geq 0\}$, where $\mathbb{1}\{\cdot\}$ denotes the indicator function. From equation (8), we can write:

$$d_{ijt} = \mathbb{1}\{\eta^{-1}\mathbb{E}[\alpha_{jt} r_{iht}|\mathcal{J}_{ijt}] - \beta_0 - \beta_1 dist_j - \nu_{ijt} \geq 0\}, \tag{9}$$

and, given equations (6) and (9), we can write the probability that i exports to j at t conditional on $\mathcal{J}_{ijt}$ and $dist_j$:

$$\mathcal{P}(d_{ijt} = 1|\mathcal{J}_{ijt}, dist_j) = \int_\nu \mathbb{1}\{\eta^{-1}\mathbb{E}[\alpha_{jt} r_{iht}|\mathcal{J}_{ijt}] - \beta_0 - \beta_1 dist_j - \nu \geq 0\}(1/\sigma)\phi(\nu/\sigma)d\nu$$

[6] At a computational cost, we can allow fixed export costs to depend on additional variables the firm knows when deciding whether to export to a market, such as shared language (Morales et al., 2017) and the quality of institutions (Antràs et al., 2017). In Appendix B.3, we generalize the specification in equation (5) and present estimates for a model in which we assume $f_{ijt} = \beta_j + \nu_{ijt}$, where β_j varies freely across countries. In Appendix A.11, we discuss an extension in which firms face unexpected shocks to fixed costs.

[7] The assumption that ν_{ijt} is distributed normally is a sufficient but not a necessary condition to derive our moment inequalities. We provide the precise requirements for the distribution of ν_{ijt} when we derive the inequalities in Sections 4.2.1 and 4.2.2.

$$= \Phi\big(\sigma^{-1}\big(\eta^{-1}\mathbb{E}[\alpha_{jt} r_{iht}|\mathcal{J}_{ijt}] - \beta_0 - \beta_1 dist_j\big)\big), \quad (10)$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ are, respectively, the standard normal probability density function and cumulative distribution function.[8] Equation (10) indicates that, after integrating over the unobserved heterogeneity in fixed costs, ν_{ijt}, we can write the probability that firm i exports to country j at period t as a probit model whose index depends on firm i's expectation of the revenue it will earn in j at t upon entry. The key hurdle in estimation, which we discuss in Section 4, is that researchers rarely observe these expectations.

From equation (10), even if the researcher were to observe firms' actual expectations, data on export choices alone would not allow us to identify the scale of the remaining parameter vector $(\sigma, \eta, \beta_0, \beta_1)$. To normalize for scale in export models, researchers typically use additional data to estimate the demand elasticity η (Das et al., 2007). In our estimation, we set $\eta = 5$.[9] For simplicity of notation, we use $\theta \equiv (\theta_0, \theta_1, \theta_2)$ to denote the remaining parameter vector and $\theta^* \equiv (\beta_0, \beta_1, \sigma)$ to denote its true value, as determined by equation (10).

3 Data

Our data come from two separate sources. The first is an extract of the Chilean customs database, which covers the universe of exports of Chilean firms from 1995 to 2005. The second is the Chilean Annual Industrial Survey (*Encuesta Nacional Industrial Anual*, or ENIA), which surveys all manufacturing plants with at least 10 workers. We merge these two datasets using firm identifiers, allowing us to observe both the export and domestic activity of each firm.[10]

The firms in our dataset operate in one of two sectors: the manufacture of chemicals and the food products sector.[11] For each sector, we estimate our model restricting the set of countries to those served by at least five Chilean firms in all years of our data. This restriction leaves 22 countries in the chemicals sector and 34 countries in the food sector.

We observe 266 unique firms across all years in the chemicals sector; on average, 38% of these firms participate in at least one export market in a given year. In Table 1, we report

[8] If knowledge of $dist_j$ helps predict r_{ijt}, then $dist_j \in \mathcal{J}_{ijt}$ and $\mathcal{P}(d_{ijt} = 1|\mathcal{J}_{ijt}, dist_j) = \mathcal{P}(d_{ijt} = 1|\mathcal{J}_{ijt})$.

[9] This value is within the range of values in the literature (Simonovska and Waugh, 2014; Head and Mayer, 2014). Given our model, one can estimate η using data on firms' total sales and variable costs (Das et al., 2007; Antràs et al., 2017). We do not implement this estimation approach given limitations in our measure of variable costs. When presenting our estimates, we indicate which conclusions are sensitive to η.

[10] We aggregate the information from ENIA across plants to obtain firm-level information to match to the customs data. ENIA sometimes identifies firms as exporters when we do not observe exports in the customs data; in these cases, we follow the customs database and treat these firms as non-exporters. We lose a number of small firms in the merging process because, as indicated in the main text, ENIA only covers plants with more than 10 workers. The remaining firms account for roughly 80% of total export flows.

[11] The chemicals sector (sector 24 of the ISIC rev. 3.1) includes firms involved in the manufacture of chemicals and chemical products, including basic chemicals, fertilizers and nitrogen compounds, plastics, synthetic rubber, pesticides, paints, soap and detergents, and manmade fibers. The food sector (sector 151 of the ISIC rev. 3.1) includes the production, processing, and preservation of meat, fish, fruit, vegetables, oils, and fats.

the mean firm-level exports in this sector, which are \$2.18 million in 1996 and grow to \$3.58 million in 2005, with a dip in 2001 and 2002.[12] The median level of exports is much lower, at around \$150,000. In the food sector, we observe 372 unique firms, 30% of which export in a typical year. The mean exporter in this sector sells \$7.7 million, while the median exporter sells approximately \$2.24 million abroad. In the chemicals sector, the average exporter serves 4-5 countries. Firms in the food sector typically export to 6-7 markets on average.

Table 1: Summary Statistics

Year	Share of exporters	Exports per exporter (mean)	Exports per exporter (med)	Domestic sales per firm (mean)	Domestic sales per exporter (mean)	Destinations per exporter (mean)
			Chemical Products			
1996	35.7%	2.18	0.15	13.23	23.10	4.24
1997	36.1%	2.40	0.19	13.29	22.99	4.54
1998	42.5%	2.41	0.17	14.31	22.25	4.35
1999	38.7%	2.60	0.19	14.43	23.95	4.53
2000	37.6%	2.55	0.21	14.41	25.93	4.94
2001	39.8%	2.35	0.12	12.89	21.92	4.68
2002	38.7%	2.37	0.15	13.25	23.73	4.95
2003	38.0%	3.08	0.17	10.41	19.54	5.11
2004	37.6%	3.27	0.15	10.05	18.70	5.17
2005	38.0%	3.58	0.11	12.50	21.65	5.19
			Food			
1996	30.1%	7.47	2.59	9.86	13.68	5.93
1997	33.1%	6.97	2.82	10.56	15.32	6.23
1998	33.3%	7.49	2.86	10.05	14.80	6.34
1999	32.3%	6.71	2.37	9.67	14.88	6.74
2000	30.6%	6.49	2.21	8.44	13.33	5.93
2001	28.0%	6.48	1.74	8.70	14.08	6.09
2002	27.2%	7.82	2.01	7.83	13.59	6.86
2003	29.8%	7.60	1.68	7.15	12.79	6.15
2004	28.5%	9.25	1.68	8.05	13.85	6.69
2005	25.8%	10.72	2.43	9.88	16.27	7.05

Notes: All variables (except "share of exporters") are reported in millions of year 2000 US dollars.

Our data set includes both exporters and non-exporters. Furthermore, we use an unbalanced panel that includes not only those firms that appear in ENIA in every year between 1995 and 2005 but also those that were created or disappeared during this period. Finally, we obtain information on the distance from Chile to each destination market from CEPII.[13]

4 Empirical Approach

In the model we describe in Section 2, r_{ijt}, firm i's potential export revenue to market j at t, is a function of its own marginal costs and demand shifter, and of country j's market size

[12] The revenue values we report are in year 2000 US dollars.

[13] Mayer and Zignago (2011) provide a detailed explanation of the content of this database.

and price index. In Section 2.2, we split these determinants into two terms: $\alpha_{jt}r_{iht}$ and e_{ijt}, where the latter reflects idiosyncratic shifters of firm i's demand and variable trade costs in j. Crucially, while the data described in Section 3 allow us to compute a consistent estimate of $\alpha_{jt}r_{iht}$ for every firm, market and year (see Appendix A.3), e_{ijt} is not observed for all firm-market-year triplets. We thus henceforth refer to the first term as the *observable* determinant of export revenue and label it $r^o_{ijt} = \alpha_{jt}r_{iht}$. We label e_{ijt} as the *unobservable* determinant.[14]

Our model implies that $\mathbb{E}[e_{ijt}|\mathcal{J}_{ijt}] = 0$ and, thus, $\mathbb{E}[r_{ijt}|\mathcal{J}_{ijt}] = \mathbb{E}[r^o_{ijt}|\mathcal{J}_{ijt}]$. The model does not restrict the relationship between $\mathcal{J}_{ijt}$ and r^o_{ijt}. However, identifying the parameter vector θ underlying the fixed export costs f_{ijt} requires additional assumptions (Manski, 1993).

First, we consider a perfect foresight model. With this model, researchers assume an information set $\mathcal{J}^a_{ijt}$ for potential exporters such that $\mathbb{E}[r^o_{ijt}|\mathcal{J}^a_{ijt}] = r^o_{ijt}$. That is, firms are assumed to have ex ante (before deciding whether to enter a foreign market) the same information that the researcher has ex post (when data becomes available). Thus, firms predict r^o_{ijt} perfectly.[15]

Second, we consider a model in which we allow firms to face uncertainty when predicting r^o_{ijt}—for example, they may lack perfect knowledge of the size of the market or the degree of competition they will face. In this model, potential exporters forecast their export revenues in every foreign market using information on three variables: (1) their own domestic sales in the previous year, r_{iht-1}; (2) sectoral aggregate exports to destination j in the previous year, R_{jt-1}; and (3) distance from the home country to j, $dist_j$. That is, we assume that the actual information set $\mathcal{J}_{ijt}$ is identical to a vector of covariates $\mathcal{J}^a_{ijt}$ observed in our data: $\mathcal{J}^a_{ijt} = (r_{iht-1}, R_{jt-1}, dist_j)$. In practice, firms can easily access these three variables in any year. However, this information set is likely to be strictly smaller than the actual information set firms possess when deciding whether to export.[16] Furthermore, specifying $\mathcal{J}_{ijt}$ as in this second model implies that all firms base their entry decision on the same set of covariates. It does not permit firms to differ in the information they use.

Third, we discuss how to identify the model parameters imposing weaker assumptions on the information firms use to predict r^o_{ijt}. We propose a moment inequality estimator that can handle settings in which the econometrician observes only a *subset* of the elements contained in firms' true information sets. That is, we assume that the researcher observes a vector Z_{ijt} such that $Z_{ijt} \subseteq \mathcal{J}_{ijt}$. The researcher need not observe the remaining elements in $\mathcal{J}_{ijt}$. Those unobserved elements of firms' information sets can vary flexibly by firm and by export market.

[14] As an alternative micro-foundation for this structure, one can rule out firm-specific shifters of demand and variable trade costs, and instead assume e_{ijt} reflects error in the researcher's observation of r^o_{ijt}.

[15] Although we denote this case as "perfect foresight", perfectly predicting export revenues only refers to the observable component, r^o_{ijt}. Firms' information sets are still orthogonal to the unobserved component e_{ijt}.

[16] When we indicate that information set $\mathcal{J}^a_{ijt}$ is smaller than information set $\mathcal{J}_{ijt}$, we formally mean that the distribution of $\mathcal{J}^a_{ijt}$ conditional on $\mathcal{J}_{ijt}$ is degenerate.

4.1 Perfect Knowledge of Exporters' Information Sets

Under the assumption that the econometrician's specified information set, $\mathcal{J}^a_{ijt}$, equals the firm's true information set, $\mathcal{J}_{ijt}$, $\mathbb{E}[r^o_{ijt}|\mathcal{J}^a_{ijt}] = \mathbb{E}[r^o_{ijt}|\mathcal{J}_{ijt}]$ and one can estimate θ^* as the value of the unknown parameter θ that maximizes the log-likelihood function

$$\mathcal{L}(\theta|d, \mathcal{J}^a, dist) =$$
$$\sum_{i,j,t} d_{ijt} \ln(\mathcal{P}(d_{jt} = 1|\mathcal{J}^a_{ijt}, dist_j; \theta)) + (1 - d_{ijt}) \ln(\mathcal{P}(d_{jt} = 0|\mathcal{J}^a_{ijt}, dist_j; \theta)), \tag{11}$$

where the vector $(d, \mathcal{J}^a, dist)$ includes all values of the corresponding covariates for every firm, country and year in the sample, and, according to equation (10) and the definition of r^o_{ijt},

$$\mathcal{P}(d_{jt} = 1|\mathcal{J}^a_{ijt}, dist_j; \theta) = \Phi\big(\theta_2^{-1}\big(\eta^{-1}\mathbb{E}[r^o_{ijt}|\mathcal{J}^a_{ijt}] - \theta_0 - \theta_1 dist_j\big)\big). \tag{12}$$

To use equations (11) and (12) to estimate θ^*, one first needs to compute $\mathbb{E}[r^o_{ijt}|\mathcal{J}^a_{ijt}]$. When the researcher assumes perfect foresight, $\mathbb{E}[r^o_{ijt}|\mathcal{J}^a_{ijt}] = r^o_{ijt}$. When the researcher assumes $\mathcal{J}^a_{ijt}$ is equal to a set of observed covariates, one can consistently estimate $\mathbb{E}[r^o_{ijt}|\mathcal{J}^a_{ijt}]$ as the non-parametric projection of r^o_{ijt} on $\mathcal{J}^a_{ijt}$.[17] The key assumption underlying these two procedures is that the researcher correctly specifies the agent's information set.

Bias in estimation will generally arise when the agent's true information set, $\mathcal{J}_{ijt}$, differs from the researcher's specification, $\mathcal{J}^a_{ijt}$, for some firms, countries or years in the sample. To characterize this bias, we begin by defining two types of errors: the agent's expectational error and the researcher's specification error. For the agent, we define $\varepsilon_{ijt} \equiv r^o_{ijt} - \mathbb{E}[r^o_{ijt}|\mathcal{J}_{ijt}]$ as the true expectational error that firm i makes when predicting the observed component of its export revenue. This error reflects the firm's uncertainty about r^o_{ijt}.[18] In contrast, we denote the difference between firms' true expectations and the researcher's proxy as ξ_{ijt}:

$$\xi_{ijt} \equiv \mathbb{E}[r^o_{ijt}|\mathcal{J}^a_{ijt}] - \mathbb{E}[r^o_{ijt}|\mathcal{J}_{ijt}]. \tag{13}$$

Whenever this error term differs from zero, estimates based on equations (11) and (12) will be biased. In Appendix D, we present simulation results that illustrate the direction and magnitude of the bias that arise in three cases: when the researcher assumes perfect foresight, when the researcher's information set is larger than the firm's information set, and when the researcher's information set is smaller than the firm's information set.

To provide intuition on the direction of the bias, we focus here on the perfect foresight case. In this case, we find an upward bias in the estimates of the fixed costs parameters β_0, β_1 and σ. The upward bias arises for a similar reason to the attenuation bias that

[17] See Manski (1991) and Ahn and Manski (1993) for additional details on this two-step estimation approach.
[18] The total expectational error that the firm makes when forecasting export revenue r_{ijt} is $\varepsilon_{ijt} + e_{ijt}$.

affects Ordinary Least Squares estimates in linear models when a covariate contains classical measurement error (see Wooldridge, 2002). Under perfect foresight, the researcher assumes the firm perfectly predicts the observable part of its export revenue, such that $\mathbb{E}[r^o_{ijt}|\mathcal{J}^a_{ijt}] = r^o_{ijt}$. Thus, the measurement error affecting the researcher's specification, $\xi_{ijt} \equiv r^o_{ijt} - \mathbb{E}[r^o_{ijt}|\mathcal{J}_{ijt}]$, is the same as the firm's true expectational error, ε_{ijt}. Rational expectations implies that firms' expectational errors are mean independent of their true expectation and, therefore, correlated with the ex-post realization of the variable being predicted; i.e. rational expectations implies that $\mathbb{E}[\varepsilon_{ijt}|\mathcal{J}_{ijt}] = 0$ and $cov(\varepsilon_{ijt}, r^o_{ijt}) \neq 0$. Thus, if we were in a linear regression setting, wrongly assuming perfect foresight and using r^o_{ijt} as a regressor instead of the unobserved expectation, $\mathbb{E}[r^o_{ijt}|\mathcal{J}_{ijt}]$, would generate a downward bias on the coefficient on r^o_{ijt}.

The probit model in equation (12) differs from this linear setting in two dimensions. First, our normalization by scale $\eta = 5$ sets the coefficient on the covariate measured with error, $\mathbb{E}[r^o_{ijt}|\mathcal{J}_{ijt}]$, to a given value. Thus, the bias generated by the correlation between the expectational error, ε_{ijt}, and the realized export revenue, r^o_{ijt}, will be reflected in an upward bias in the estimates of the remaining parameters β_0, β_1 and σ. Second, the direction of the bias depends not only on the correlation between ε_{ijt} and r^o_{ijt} but also on the functional form of the distribution of unobserved expectations and the expectational error.[19]

4.2 Partial Knowledge of Exporters' Information Sets

In most empirical settings, researchers rarely observe the exact covariates that form the firm's information set. However, they can typically find a vector of covariates in their data that represents a subset of the firm's information set. For example, in each year, exporters will likely know past values of both their domestic sales, r_{iht-1}, and the aggregate exports from their home country to each destination market, R_{jt-1}; the former appears in firms' accounting statements, while the latter appears in publicly available trade data. Similarly, firms can easily obtain information on the distance to each destination, $dist_j$. Thus, while $(r_{iht-1}, R_{it-1}, dist_j)$ might not reflect firms' complete information, they likely know at least this vector.

In this section, we show how to proceed in estimation using a vector of observed covariates Z_{ijt} that is a subset of the information firms use to forecast export revenues, i.e. $Z_{ijt} \subseteq \mathcal{J}_{ijt}$. We show how to test formally whether firms possess this information in Section 6. We form two types of moment inequalities to partially identify θ^*.[20]

[19]If both firms' true expectations and expectational errors are normally distributed, $\mathbb{E}[r^o_{ijt}|\mathcal{J}_{ijt}] \sim \mathbb{N}(0, \sigma^2_e)$ and $\varepsilon_{ijt}|(\mathcal{J}_{ijt}, \nu_{ijt}) \sim \mathbb{N}(0, \sigma^2_\varepsilon)$, one can apply the results in Yatchew and Griliches (1985) and conclude that there is an upward bias in the estimates of β_0, β_1 and σ. This bias increases in the variance of the expectational error, σ^2_ε, relative to the variance of the true expectations, σ^2_e. When either firms' true expectations, $\mathbb{E}[r^o_{ijt}|\mathcal{J}_{ijt}]$, or the expectational error, ε_{ijt}, are not normally distributed, there is no analytic expression for the bias. However, our simulations in Appendix D illustrate that the upward bias in the estimates of all elements of θ^* generally persists under different distributions of $\mathbb{E}[r^o_{ijt}|\mathcal{J}_{ijt}]$ and ε_{ijt}.

[20]As shown in Appendix A.4, given the model described in Section 2, the assumption that the researcher observes a subset of a firm's true information set is not strong enough to point-identify θ^*. Whether the bounds

4.2.1 Odds-based Moment Inequalities

For any $Z_{ijt} \subseteq (\mathcal{J}_{ijt}, dist_j)$, we define the conditional odds-based moment inequalities as

$$\mathcal{M}^{ob}(Z_{ijt};\theta) = \mathbb{E}\left[\begin{array}{c} m_l^{ob}(d_{ijt}, r_{ijt}^o, dist_j;\theta) \\ m_u^{ob}(d_{ijt}, r_{ijt}^o, dist_j;\theta) \end{array}\middle| Z_{ijt}\right] \geq 0, \tag{14a}$$

where the two moment functions are defined as

$$m_l^{ob}(\cdot) = d_{ijt}\frac{1-\Phi\big(\theta_2^{-1}\big(\eta^{-1}r_{ijt}^o - \theta_0 - \theta_1 dist_j\big)\big)}{\Phi\big(\theta_2^{-1}\big(\eta^{-1}r_{ijt}^o - \theta_0 - \theta_1 dist_j\big)\big)} - (1-d_{ijt}), \tag{14b}$$

$$m_u^{ob}(\cdot) = (1-d_{ijt})\frac{\Phi\big(\theta_2^{-1}\big(\eta^{-1}r_{ijt}^o - \theta_0 - \theta_1 dist_j\big)\big)}{1-\Phi\big(\theta_2^{-1}\big(\eta^{-1}r_{ijt}^o - \theta_0 - \theta_1 dist_j\big)\big)} - d_{ijt}. \tag{14c}$$

We denote the set of all possible values of the parameter vector θ as Θ and the subset of those values consistent with the conditional moment inequalities described in equation (14) as Θ_0^{ob}.

Theorem 1 *Let $\theta^* = \big(\beta_0, \beta_1, \sigma\big)$ be the parameter defined by equation* (10)*. Then $\theta^* \in \Theta_0^{ob}$.*

Theorem 1 indicates that the odds-based inequalities are consistent with the true value of the parameter vector, θ^*. We provide here an intuitive explanation of Theorem 1. The formal proof appears in Appendix C.1.

We focus on the intuition behind the moment function in equation (14c); the intuition for equation (14b) is analogous. From the definition of d_{ijt} in equation (9) and the definition of r_{ijt}^o, we can write

$$\mathbb{1}\{\eta^{-1}\mathbb{E}[r_{ijt}^o|\mathcal{J}_{ijt}] - \beta_0 - \beta_1 dist_j - \nu_{ijt} \geq 0\} - d_{ijt} = 0. \tag{15}$$

This equation, using revealed preference, implies the condition that expected export profits are positive, $\eta^{-1}\mathbb{E}[r_{ijt}^o|\mathcal{J}_{ijt}] - \beta_0 - \beta_1 dist_j - \nu_{ijt} \geq 0$, is both necessary and sufficient for observing firm i exporting to country j in year t, $d_{ijt} = 1$. Equation (15) cannot be used directly for identification, as it depends on the unobserved terms ν_{ijt} and $\mathcal{J}_{ijt}$. To account for the term ν_{ijt}, we take the expectation of equation (15) conditional on $(\mathcal{J}_{ijt}, dist_j)$. Given the distributional assumption in equation (6), we use simple algebraic transformations to rewrite the resulting equality as

$$\mathbb{E}\left[(1-d_{ijt})\frac{\Phi(\sigma^{-1}(\eta^{-1}\mathbb{E}[r_{ijt}^o|\mathcal{J}_{ijt}] - \beta_0 - \beta_1 dist_j))}{1-\Phi(\sigma^{-1}(\eta^{-1}\mathbb{E}[r_{ijt}^o|\mathcal{J}_{ijt}] - \beta_0 - \beta_1 dist_j))} - d_{ijt}\middle|\mathcal{J}_{ijt}, dist_j\right] = 0. \tag{16}$$

If we write this equality as a function of the unknown parameter θ, it would only hold at

defined by our inequalities are sharp is left for future research. However, as the results in Section 5 show, in our empirical application, they generate bounds that are tight enough to be informative.

its true value θ^*. This equality, however, still depends on the unknown true information set, $\mathcal{J}_{ijt}$, through the unobserved expectation, $\mathbb{E}[r^o_{ijt}|\mathcal{J}_{ijt}]$. We exploit the property that the moment function in equation (14c) is convex in the unobserved expectation $\mathbb{E}[r^o_{ijt}|\mathcal{J}_{ijt}]$; i.e. $\Phi(\cdot)/(1-\Phi(\cdot))$ is convex. Thus, applying Jensen's inequality, equation (16) becomes an inequality if we introduce the observed proxy, r^o_{ijt}, in place of the unobserved expectation $\mathbb{E}[r^o_{ijt}|\mathcal{J}_{ijt}]$ and take the expectation of the resulting expression conditional on an observed vector $Z_{ijt} \subseteq \mathcal{J}_{ijt}$. Consequently, if the equality in equation (16) holds at the true value of the parameter vector, the inequality defined in equations (14) and (14c) will also hold at $\theta = \theta^*$.[21]

The moment functions in equations (14b) and (14c) are not redundant. For example, consider the identification of the parameter θ_0. Given observed values of d_{ijt}, r^o_{ijt}, and $dist_j$, and given any arbitrary value of the parameters θ_1 and θ_2, the moment function $m^{ob}_l(\cdot)$ in equation (14b) is increasing in θ_0 and, therefore, will identify a lower bound on θ_0. With the same observed values, $m^{ob}_u(\cdot)$ in equation (14c) is decreasing in θ_0 and will thus identify an upper bound on θ_0. The same intuition applies for identifying bounds for θ_1 and θ_2.

4.2.2 Revealed Preference Moment Inequalities

For any $Z_{ijt} \subseteq (\mathcal{J}_{ijt}, dist_j)$, we define a conditional revealed preference moment inequality as

$$\mathcal{M}^r(Z_{ijt};\theta) = \mathbb{E}\left[\begin{array}{c} m^r_l(d_{ijt}, r^o_{ijt}, dist_j;\theta) \\ m^r_u(d_{ijt}, r^o_{ijt}, dist_j;\theta) \end{array}\middle| Z_{ijt}\right] \geq 0, \tag{17a}$$

where the two moment functions are defined as

$$m^r_l(\cdot) = -(1-d_{ijt})\big(\eta^{-1}r^o_{ijt} - \theta_0 - \theta_1 dist_j\big) + d_{ijt}\theta_2 \frac{\phi\big(\theta_2^{-1}(\eta^{-1}r^o_{ijt} - \theta_0 - \theta_1 dist_j)\big)}{\Phi\big(\theta_2^{-1}(\eta^{-1}r^o_{ijt} - \theta_0 - \theta_1 dist_j)\big)}, \tag{17b}$$

$$m^r_u(\cdot) = d_{ijt}\big(\eta^{-1}r^o_{ijt} - \theta_0 - \theta_1 dist_j\big) + (1-d_{ijt})\theta_2 \frac{\phi\big(\theta_2^{-1}(\eta^{-1}r^o_{ijt} - \theta_0 - \theta_1 dist_j)\big)}{1-\Phi\big(\theta_2^{-1}(\eta^{-1}r^o_{ijt} - \theta_0 - \theta_1 dist_j)\big)}. \tag{17c}$$

We denote the values of θ consistent with the moment inequalities in equation (17) as Θ^r_0.

Theorem 2 *Let $\theta^* = \big(\beta_0, \beta_1, \sigma\big)$ be the parameter defined by equation* (10)*. Then $\theta^* \in \Theta^r_0$.*

We provide a formal proof of Theorem 2 in Appendix C.2. Theorem 2 indicates that the revealed preference inequalities are consistent with the true value of the parameter vector, θ^*.

Heuristically, the two moment functions in equations (17b) and (17c) are derived using standard revealed preference arguments. We focus our discussion on the moment function

[21] The assumption that ν_{ijt} follows a normal distribution is sufficient but not necessary to derive the odds-based inequalities. For any distribution of ν_{ijt} with cumulative distribution function $F_\nu(\cdot)$, we need simply that $F_\nu(\cdot)/(1-F_\nu(\cdot))$ and $(1-F_\nu(\cdot))/F_\nu(\cdot)$ are globally convex. This condition will be satisfied if the distribution of ν is log-concave. Both the normal and the logistic distributions are log-concave, as are the uniform, exponential, type I extreme value, and Laplace distributions. Heckman and Honoré (1990), and Bagnoli and Bergstrom (2005) provide more information on the properties of log-concave distributions.

in equation (17c); the intuition behind the derivation of the moment in equation (17b) is analogous. If firm i decides to export to country j in period t, so that $d_{ijt} = 1$, then by revealed preference, it must expect to earn positive returns; i.e. $d_{ijt}\big(\eta^{-1}\mathbb{E}[r_{ijt}|\mathcal{J}_{ijt}] - \beta_0 - \beta_1 dist_j - \nu_{ijt}\big) \geq 0$. Taking the expectation of this inequality conditional on $(d_{ijt}, \mathcal{J}_{ijt}, dist_j)$ and taking into account that $\mathbb{E}[r_{ijt}|\mathcal{J}_{ijt}] = \mathbb{E}[r^o_{ijt}|\mathcal{J}_{ijt}]$, we obtain

$$d_{ijt}\big(\eta^{-1}\mathbb{E}[r^o_{ijt}|\mathcal{J}_{ijt}] - \beta_0 - \beta_1 dist_j\big) + S_{ijt} \geq 0, \tag{18}$$

where $S_{ijt} = \mathbb{E}[-d_{ijt}\nu_{ijt}|d_{ijt}, \mathcal{J}_{ijt}, dist_j]$. The term S_{ijt} is a selection correction that accounts for how ν_{ijt} affects the firm's decision to export, where again ν_{ijt} captures determinants of profits that the researcher does not observe.[22] We cannot directly use the inequality in equation (18) because it depends on the unobserved agents' expectations, $\mathbb{E}[r^o_{ijt}|\mathcal{J}_{ijt}]$, both directly and through the term S_{ijt}. However, the inequality in equation (18) becomes weaker if we introduce the observed covariate, r^o_{ijt}, in the place of the unobserved expectations, $\mathbb{E}[r^o_{ijt}|\mathcal{J}_{ijt}]$, and take the expectation of the resulting expression conditional on Z_{ijt}. As in the case of the odds-based inequalities, we need the moment function in equation (17c) to be globally convex in the unobserved expectation $\mathbb{E}[r^o_{ijt}|\mathcal{J}_{ijt}]$; i.e. $\phi(\cdot)/(1-\Phi(\cdot))$ is convex. Consequently, if the inequality in equation (18) holds at the true value of the parameter vector, the inequality in equations (17) and (17c) will also hold at $\theta = \theta^*$.[23]

The inequalities in equation (17) follow the revealed preference inequalities introduced in Pakes (2010) and Pakes et al. (2015). In our setting, our inequalities feature structural errors ν_{ijt} that may vary across (i, j, t) and that have unbounded support. The cost of allowing this flexibility is that we must assume a distribution for ν_{ijt}, up to a scale parameter.[24,25]

4.2.3 Combining Inequalities for Estimation

We combine the odds-based and revealed preference moment inequalities described in equations (14) and (17) for estimation. The set defined by the odds-based inequalities is a singleton

[22] Appendix C.2 shows that, under the assumptions in Section 2,

$$S_{ijt} = (1 - d_{ijt})\sigma \frac{\phi\big(\sigma^{-1}(\eta^{-1}\mathbb{E}[r^o_{ijt}|\mathcal{J}_{ijt}] - \beta_0 - \beta_1 dist_j)\big)}{1 - \Phi\big(\sigma^{-1}(\eta^{-1}\mathbb{E}[r^o_{ijt}|\mathcal{J}_{ijt}] - \beta_0 - \beta_1 dist_j)\big)}.$$

[23] As in footnote 21, the assumption of normality of ν_{ijt} is sufficient but not necessary. For the inequality equations (17) and (17c) to hold, we need a distribution for ν_{ijt} such that $\mathbb{E}[\nu_{ijt}|\nu_{ijt} < \kappa]$ is globally convex in the constant κ. An analogous condition is needed to derive equation (17b). In addition to the normal distribution, the logistic distribution also satisfies this condition.

[24] In our empirical application, we find σ, the standard deviation of ν_{ijt}, to be greater than zero. Therefore, including the selection correction term S_{ijt} in our inequalities is important: given that $S_{ijt} \geq 0$ whenever $\sigma > 0$, if we had generated revealed preference inequalities without S_{ijt}, we would have obtained weakly smaller identified sets than those found using the inequalities in equation (17).

[25] Pakes and Porter (2015) and Shi et al. (2017) show how to estimate discrete choice models in panel data settings without imposing distributional assumptions on ν_{ijt}. Both models, however, impose a restriction that agents make no errors in their expectations.

only when firms make no expectational errors and the vector of instruments Z_{ijt} is identical to the set of variables firms use to form their expectations. In this very specific case, the revealed preference inequalities do not have any additional identification power beyond that of the odds-based inequalities. However, in all other settings, the revealed preference moments can provide additional identifying power.

The set of inequalities we define in equations (14) and (17) condition on particular values of the instrument vector, Z_{ijt}. Exploiting all the information contained in these conditional moment inequalities can be computationally challenging.[26] In this paper, we base our inference on a fixed number of unconditional moment inequalities implied by the conditional moment inequalities in equations (14) and (17). We describe in Appendix A.5 the unconditional moments we use to compute the estimates discussed in Section 5. We denote the set of values of θ consistent with our unconditional odds-based and revealed-preference inequalities as Θ_0.

Conditioning on a fixed set of moments, while convenient, entails a loss of information. Thus, the identified set defined by our unconditional moment inequalities may be larger than that implied by their conditional counterparts. However, as the empirical results in sections 5, 6 and 7 show, the moment inequalities we employ nonetheless generate economically meaningful bounds on our parameters and on counterfactual choice probabilities, and also allow us to explore hypotheses about the information firms use to forecast export revenue.

4.2.4 Characterizing the Identified Set

Theorems 1 and 2 imply that θ^* will be contained in the set Θ_0 defined by our odds-based and revealed-preference moment inequalities when the instrument vector Z_{ijt} used to define these inequalities satisfies $Z_{ijt} \subseteq \mathcal{J}_{ijt}$ for all i, j, and t. However, these theorems do not fully characterize the set Θ_0. That is, they do not indicate the values of θ other than θ^* that are also included in this set. A full characterization is beyond the scope of this paper, but we conduct a simulation, with full results reported in Appendix E, to explore the content of Θ_0.

In particular, we design a simulation in which the researcher has access to three possible information sets: (1) a small information set, $\mathcal{J}^s_{ijt}$, that contains too few variables relative to the true information set; (2) a medium-sized set, $\mathcal{J}^m_{ijt}$, that coincides with the true information set; and (3) a large information set, $\mathcal{J}^l_{ijt}$, that contains more information than the firm actually possesses. Here, under $\mathcal{J}^l_{ijt}$, every firm can predict perfectly the observable component of its potential export revenues; i.e. $\mathbb{E}[r^o_{ijt}|\mathcal{J}^l_{ijt}] = r^o_{ijt}$. We denote the probability limits of the corresponding maximum likelihood estimators as θ_s, θ_m, and θ_l. For example, the maximum likelihood estimator with probability limit θ_s is computed under the incorrect assumption that the true information set equals $\mathcal{J}^s_{ijt}$. We similarly denote the corresponding identified

[26] Recent theoretical work, including Andrews and Shi (2013), Chernozhukov et al. (2013), Chetverikov (2013), Armstrong (2014), Armstrong (2015), and Armstrong and Chan (2016), provide estimation procedures that exploit all information contained in conditional moment inequality models.

sets defined by our moment inequalities as Θ_s^0, Θ_m^0, and Θ_l^0. For example, the identified set Θ_s^0 is computed under the correct assumption that the true information set includes $\mathcal{J}_{ijt}^s$.

Using our moment inequalities, both the assumptions that exporters know at least the variables in $\mathcal{J}_i^s$ and $\mathcal{J}_i^m$ are compatible with the data generating process. Thus, as discussed in Section 4.2, Θ_s^0 and Θ_m^0 will both contain θ^*. Of the maximum likelihood estimators, only $\mathcal{J}_i^m$ is compatible with the data generating process and, consequently, as discussed in Section 4.1, only θ_m coincides with the true parameter θ^*.

The informational assumptions imposed to compute θ_s and θ_l are compatible with the weaker information assumption imposed to compute Θ_s^0. However, as we show in Appendix E, it need not be the case that Θ_s^0 contains θ_s and θ_l. Their inclusion depends on (a) how different θ_s and θ_l are from the true value θ^* and (b) the span of points in the identified set Θ_s^0 around θ^*.

The distance between θ_s and θ^* depends on the importance of those predictors of export revenues contained in the true information set, $\mathcal{J}_{ijt}^m$, and excluded from the assumed one, $\mathcal{J}_{ijt}^s$; i.e. θ_s and θ^* move further apart as the variance of $\xi_{ijt}^s \equiv \mathbb{E}[r_{ijt}^o|\mathcal{J}_{ijt}^s] - \mathbb{E}[r_{ijt}^o|\mathcal{J}_{ijt}^m]$ increases. The distance between θ_l and θ^* increases in the importance of the variables included in the assumed information set, $\mathcal{J}_{ijt}^l$, and excluded from the true one, $\mathcal{J}_{ijt}^m$. Here, the distance increases in the variance of the firm's true expectational error, $\varepsilon_{ijt} \equiv r_{ijt}^0 - \mathbb{E}[r_{ijt}^o|\mathcal{J}_{ijt}^m]$.

The identified set Θ_s^0 will be larger when $\mathcal{J}_i^s$ excludes important predictors of potential export profits, r_{ijt}^o. Specifically, as the variance of $r_{ijt}^o - \mathbb{E}[r_{ijt}^o|\mathcal{J}_{ijt}^s] = \varepsilon_{ijt} - \xi_{ijt}^l$ increases, the identified set grows larger. Therefore, the same factors that increase the difference between both θ_s and θ_l and the true parameter vector θ^* will also make the identified set Θ_s larger. However, as we show in Appendix E, these factors have a larger effect on the bias of the misspecified maximum likelihood estimators than on the size of the identified set. Consequently, θ_s and θ_l will tend to belong to Θ_s^0 when the two chosen information sets, respectively, are close to the true information set.

5 Results

We estimate the parameters of exporters' participation decisions using the three different empirical approaches discussed in sections 4.1 and 4.2. First, we use maximum likelihood when we assume perfect foresight. Second, we again use maximum likelihood methods, but under the two-step procedure in which we project realized revenues on a set of observable covariates that we assume form a firm's information set. Third, we carry out our moment inequality approach under the assumption that the firm knows the same observed variables as in the two-step approach, but may also use additional variables to forecast revenues.

Before implementing these three procedures, we first need to compute our proxy for the observable component of export revenue, r_{ijt}^o. We describe in Appendix A.3 how to obtain

Table 2: Parameter estimates

Estimator	Chemicals σ	Chemicals β_0	Chemicals β_1	Food σ	Food β_0	Food β_1
Perfect Foresight (MLE)	1,038.6 (393)	745.2 (280)	1,087.8 (394)	1,578.1 (225)	2,025.1 (292)	214.5 (26)
Minimal Information (MLE)	395.5 (83)	298.3 (56)	447.1 (102)	959.9 (146)	1,259.3 (188)	129.4 (16)
Moment Inequality (OB and RP)	[85.1, 115.9]	[62.8, 81.1]	[142.5, 194.2]	[114.9, 160.0]	[167.1, 264.0]	[36.4, 80.2]
Moment Inequality (OB only)	[85.1, 133.3]	[34.8, 129.3]	[101.3, 293.4]	[114.9, 1,000]	[115.1, 1,310]	[6.8, 485]
Moment Inequality (RP only)	[80.0, 133.3]	[44.0, 133.3]	[110.0, 333.3]	[66.7, 487.8]	[64.7, 659.7]	[3.3, 731.4]

Notes: All estimates are reported in thousands of year 2000 USD and their values scale proportionally with η, which is set equal to 5 (see Section 2.4). For the two ML estimators, bootstrap standard errors are computed according to the procedure described in Appendix A.6 and reported in parentheses. For the three moment inequality estimates, extreme points of the 95% confidence set are reported in square brackets. These confidence sets are projections of a confidence set for $(\beta_0, \beta_1, \sigma)$ computed according to the procedure described in Appendix A.7. The "OB and RP" confidence sets exploit both odds-based and revealed-preference moment inequalities. The "OB only" use only odds-based moment inequalities. The "RP only" use only revealed-preference moment inequalities.

this proxy, which requires estimating revenue shifters, α_{jt}, for each market j and period t. We report the estimates $\{\hat{\alpha}_{jt}; \forall j, t\}$ for both the chemicals and food sectors in Appendix B.1.[27]

5.1 Average Fixed Export Costs

In Table 2, we report the estimates and confidence regions for our model parameters. The first coefficient, σ, is the standard deviation of the structural error ν_{ijt}. It controls the heterogeneity across firms and time periods in the fixed costs of exporting to a particular destination j. The remaining coefficients, β_0 and β_1, represent a constant component and the contribution of distance to the level of the fixed costs. We normalize the demand elasticity, η, to equal five.

From Table 2, we see the models that assume researchers have full knowledge of the exporter's information set produce much larger average fixed export costs than does our moment inequality approach. For example, consider the coefficient on the distance variable in models estimated using data from the chemicals sector. Under the moment inequality approach, we find an added cost of \$142,500 to \$194,200 when the export destination is 10,000 kilometers farther in distance. Under the two maximum likelihood procedures, estimates of the added cost equal \$1,087,800 and \$447,100 for the same added distance.

The moment inequality bounds on each of the elements of the parameter vector θ reported in Table 2 arise from projecting a three-dimensional 95% confidence set for the vector

[27] When computing standard errors for the maximum likelihood estimates of θ and computing moment inequality confidence sets for this parameter, we take into account the sampling error affecting our estimates of α_{jt}. See appendices A.6 and A.7 for details.

$(\beta_0, \beta_1, \sigma)$, computed following the procedure in Appendix A.7.[28] The results in Table 2 illustrate the value of using the revealed-preference and odds-based inequalities jointly. Re-running our estimation using each set of inequalities separately, we obtain much larger bounds on the fixed export costs than when we combine both types of inequalities.

We translate the coefficients reported in Table 2 into estimates of the average fixed costs of exporting by country. We report the results in Table 3 for three countries: Argentina, Japan, and the United States. Total exports to these countries account for 29% of total exports of the Chilean chemicals sector and 56% of the food sector in the sample period. In addition, these three countries span a wide range of possible distances to Chile and thus provide an illustration of the impact of distance on average fixed export costs.[29]

Table 3: Average fixed export costs

		Chemicals			Food	
Estimator	Argentina	Japan	United States	Argentina	Japan	United States
Perfect Foresight (MLE)	868.0 (150.1)	2,621.4 (468.2)	1,645.0 (290.7)	2,049.3 (213.8)	2,395.1 (259.4)	2,202.5 (233.8)
Minimal Information (MLE)	348.7 (49.9)	1,069.4 (142.2)	668.1 (90)	1,273.9 (221.8)	1,482.4 (272.9)	1,366.3 (244.2)
Moment Inequality	[79.5, 102.6]	[309.2, 414.3]	[181.3, 240.1]	[175.6, 270.1]	[269.1, 361.0]	[227.3, 308.9]

Notes: All estimates are reported in thousands of year 2000 USD and their values scale proportionally with η, which is set equal to 5 (see Section 2.4). For the two ML estimators, bootstrap standard errors are computed according to the procedure described in Appendix A.6 and reported in parentheses. For the three moment inequality estimates, extreme points of the 95% confidence set are reported in square brackets. These confidence sets are projections of a confidence set for $(\beta_0, \beta_1, \sigma)$ computed according to the procedure described in Appendix A.7.

Table 4: Average fixed export costs relative to perfect foresight estimates

		Chemicals			Food	
Estimator	Argentina	Japan	United States	Argentina	Japan	United States
Minimal Info.	40.2%	40.8%	40.6%	62.2%	61.9%	62.0%
Moment Ineq.	[9.1%, 11.9%]	[11.8%, 15.8%]	[11.1%, 14.6%]	[8.6%, 13.2%]	[11.3%, 15.1%]	[10.3%, 14.1%]

Notes: This table reports the ratio of (a) the minimal information ML point estimates and (b) extremes of the moment inequality confidence set, compared to the perfect foresight ML point estimate. All numbers reported in this table are independent of the value of η chosen as normalizing constant.

[28] Formally, denoting $\hat{\Theta}^{95\%}$ as the 95% confidence set for the vector $(\beta_0, \beta_1, \sigma)$, the confidence set for β_0, for example, contains all values of the unknown parameter θ_0 such that there exists values of θ_1 and θ_2 for which the triplet $(\theta_0, \theta_1, \theta_2)$ is included in $\hat{\Theta}^{95\%}$. Bugni et al. (2016) introduce a new inference procedure that dominates this projection-based inference in terms of power. We report here confidence sets based on the projection of $\hat{\Theta}^{95\%}$ because (a) these one-dimensional confidence sets are nonetheless small enough to illustrate the difference between the maximum likelihood and the moment inequality estimates and (b) they do not require additional computation once we have computed $\hat{\Theta}^{95\%}$. We use $\hat{\Theta}^{95\%}$ directly to compute the results in sections 6 and 7.

[29] In Appendix B.2, we also report quantiles of the distribution of fixed export costs across firms. In Appendix B.3, we relax the assumptions in equation (5) and instead estimate average fixed costs for each country j as a country fixed effect. Moment inequality confidence sets and maximum likelihood confidence intervals are wider in this case, reflecting the larger number of parameters to estimate. The qualitative results are similar.

Under perfect foresight, we estimate the average fixed costs in Argentina, Japan, and the United States in the chemicals sector to equal \$868,000, \$2.62 million, and \$1.64 million, respectively. In the food sector, the average fixed cost estimates in these three countries equal \$2.05 million, \$2.40 million, and \$2.20 million, respectively. As we show in Table 4, when comparing the estimates under perfect foresight to the estimates that assume a minimal information set with only three variables, the latter produces estimates that are about 60% smaller in the chemicals sector and 38% smaller in the food sector. Under our moment inequality estimator, we find 95% confidence sets for the fixed costs of exporting in the chemicals sector between \$79,500 and \$102,600 for Argentina, \$309,200 and \$414,300 for Japan, and \$181,300 and \$240,100 for the United States.[30] In all cases, the estimated bounds we find from the inequalities represent only a fraction of the perfect foresight estimates and the estimates from the two-step approach. Taken together, these results reflect the discussion in Section 4.1 and in Appendix D of the bias that arises if the researcher incorrectly specifies the exporter's information set.[31]

It may seem counterintuitive that the maximum likelihood estimates obtained under the assumption $\mathcal{J}^a_{ijt} = (r_{iht-1}, R_{it-1}, dist_j)$ are not contained in the confidence set computed under the assumption that $(r_{iht-1}, R_{it-1}, dist_j) \subseteq \mathcal{J}_{ijt}$. However, as we discuss in Section 4.2.4 and illustrate in Appendix E in detail, not every information set $\mathcal{J}^a_{ijt}$ consistent with our assumption that $(r_{iht-1}, R_{jt-1}, dist_j) \subseteq \mathcal{J}_{ijt}$ generates a likelihood function whose maximand is contained in the identified set defined by our moment inequalities.

6 Testing Content of Exporters' Information Sets

What do exporters know? We use the moment inequalities introduced in Section 4.2.3 to provide some answers. To do so, we exploit an implication of our empirical model: under rational expectations, any variable in the information set the firm uses to predict export revenues serves as an instrument in our moment inequalities. Thus, we can test whether a set of observed variables, Z_{ijt}, belongs to the firm's information set using the model specification test in Bugni et al. (2015) to test the null hypothesis that there exists a value of the parameter vector that rationalizes the resulting set of moment inequalities.

If we reject that there is a value of the parameter vector at which all our moment inequalities hold, we can conclude either that (a) one of the assumptions embedded in the export model described in Section 2 does not hold in the data or that (b) the set of observed variables Z_{ijt} we specify are not contained in the firm's information set, $\mathcal{J}_{ijt}$. To distinguish between

[30] To find confidence sets for the average fixed costs in j, $\bar{f}_j = \beta_0 + \beta_1 dist_j$, we compute the lower bound on $\bar{f}_j$ as $\min_{\theta \in \hat{\Theta}^{95\%}} \theta_0 + \theta_1 dist_j$ and the upper bound as $\max_{\theta \in \hat{\Theta}^{95\%}} \theta_0 + \theta_1 dist_j$.

[31] Specifically, the upward bias in the minimal information set case is consistent with a simulation in which the distribution of the difference between the true expectation and the one implied by the minimal information set, $\mathbb{E}[r_{ijt}|\mathcal{J}_{ijt}] - \mathbb{E}[r_{ijt}|\mathcal{J}^a_{ijt}]$, is not symmetric. See Table D.3 for details.

these two conclusions, we repeat our test with the same underlying model but different Z_{ijt}.[32]

The p-values for the tests we perform appear in Table 5. In panel A, we test our main specification in which Z_{ijt} contains three covariates: the aggregate exports from Chile to each destination market in the previous year, R_{jt-1}; the distance to each market, $dist_j$; and the firm's own domestic sales in the previous year, r_{iht-1}. We fail to reject, at conventional significance levels, the null hypothesis that potential exporters know at least these three covariates when predicting export revenue.[33] In panel B, we run our moment inequality procedure under the assumption of perfect foresight. Here, we presume the firm knows r^o_{ijt} when it chooses whether to export. We can reject, at conventional significance levels, that firms know this future revenue when deciding whether to export.

In the remaining panels of Table 5, we re-run the same test as in Panel A adding an additional variable to the vector of instruments. In panel C, we add the lagged value of the country-year revenue shifter, α_{jt-1}. From the model in Section 2, this shifter is a sufficient statistic for how destination-specific supply and demand factors affect export revenues. The results in panel C support two broad conclusions: (a) large firms have more information about α_{jt} than small firms; and (b) the information that a firm has about this shifter appears independent of prior export experience to a market and the popularity of the market. Specifically, at the 5% significance level, we cannot reject that α_{jt-1} is in the information set of large firms (defined as firms with above median domestic sales in the previous year) when exporting to either popular or unpopular markets (defined as markets with above or below the median number of Chilean exporters in the previous year). We further rule out that this finding on large firms' information is a result of past export experience: we cannot reject that large firms know α_{jt-1} even if they did not export to j at $t-1$. We perform the same tests for small firms. For most tests, we can reject that small firms have information on α_{jt-1}.

In panel D, we explore whether firms differ in the information they have about a quantity that may be simpler to observe—the number of exporters to a destination in the previous year, N_{jt-1}. Although acquiring information about N_{jt-1} is likely easier than acquiring information about α_{jt-1}, it is not trivial.[34] According to our tests, we cannot reject that large firms, including those without prior export experience, are likely to have information on N_{jt-1} for any destination j. The results are similar for both popular and unpopular destinations. Given

[32]We simultaneously test multiple hypotheses. We describe in Appendix A.8 how we compute individual p-values for each test and how we use the procedure in Holm (1979) to compute family-adjusted p-values.

[33]We will not reject our null hypothesis as long as the expectational error in the firm's revenue forecast, $\varepsilon_{ijt} \equiv r^o_{ijt} - \mathbb{E}[r^o_{ijt}|\mathcal{J}_{ijt}]$, satisfies the condition $\mathbb{E}[\varepsilon_{ijt}|Z_{ijt}] = 0$. This condition will hold when Z_{ijt} is (a) irrelevant to predict r^o_{ijt} or, (b) if relevant, when Z_{ijt} is in the information set $\mathcal{J}_{ijt}$. To make the conclusion from our test clearer, we rule out the "irrelevant" explanation to our findings by running a pre-test on every variable included in any vector Z_{ijt} whose validity as instruments we test. In this pre-test, we check that these variables have predictive power for r^o_{ijt}. The results from this pre-test are included in Appendix B.4.

[34]The annual and monthly reports published by the Chilean Customs Agency (`http://www.aduana.cl/anuarios-compendios-y-reportes-estadisticos/aduana/2016-09-20/165452.html`) include information on total volume of exports by destination country and product, but not on the number of exporters.

Table 5: Testing Content of Information Sets

Set of Firms	Set of Export Destinations	Chemicals Individual p-value	Chemicals Adjusted p-value	Chemicals Reject at 5%	Food Individual p-value	Food Adjusted p-value	Food Reject at 5%
		Panel A: Minimal Information					
All	All	0.111	0.111	No	0.980	0.980	No
		Panel B: Perfect Foresight					
All	All	0.023	0.023	Yes	< 0.001	< 0.001	Yes
		Panel C: Minimal Information & Country Shifter					
Large	Popular	0.144	0.418	No	0.974	1	No
Large	Unpopular	0.114	0.418	No	0.981	1	No
Small	Popular	< 0.001	< 0.001	Yes	< 0.001	< 0.001	Yes
Small	Unpopular	0.024	0.118	No	0.004	0.021	Yes
Large Exporter	All	0.104	0.418	No	0.990	1	No
Large Non-exporter	All	0.140	0.418	No	0.048	0.190	No
Small Exporter	All	< 0.001	< 0.001	Yes	< 0.001	< 0.001	Yes
Small Non-exporter	All	< 0.001	< 0.001	Yes	0.015	0.075	No
		Panel D: Minimal Information & Number of Exporters					
Large	Popular	0.104	0.311	No	0.978	1	No
Large	Unpopular	0.114	0.311	No	0.981	1	No
Small	Popular	< 0.001	< 0.001	Yes	< 0.001	< 0.001	Yes
Small	Unpopular	0.116	0.311	No	0.003	0.015	Yes
Large Exporter	All	0.018	0.080	No	0.988	1	No
Large Non-exporter	All	0.016	0.080	No	0.717	1	No
Small Exporter	All	< 0.001	< 0.001	Yes	< 0.001	< 0.001	Yes
Small Non-exporter	All	< 0.001	< 0.001	Yes	< 0.001	< 0.001	Yes
		Panel E: Minimal Information & Country Group Avg. Shifter					
				(a) Continent Avg. Shifter			
Large	All	0.109	0.828	No	0.986	1	No
Small	All	0.112	0.828	No	0.470	1	No
				(b) Language Group Avg. Shifter			
Large	All	0.116	0.828	No	0.980	1	No
Small	All	0.115	0.828	No	0.003	0.0180	Yes
				(c) Income p.c. Group Avg. Shifter			
Large	All	0.104	0.828	No	0.980	1	No
Small	All	0.152	0.828	No	0.991	1	No
				(d) Border Group Avg. Shifter			
Large	All	0.119	0.828	No	0.981	1	No
Small	All	0.114	0.828	No	0.001	0.008	Yes

Notes: Each panel differs in the content of the information set being tested and is a separate family for the purpose of adjusting p-values. Panel A tests that $(dist_j, r_{iht-1}, R_{jt-1}) \subseteq \mathcal{J}_{ijt}$; panel B tests $\alpha_{jt} r_{iht} \subseteq \mathcal{J}_{ijt}$; panel C tests $(dist_j, r_{iht-1}, R_{jt-1}, \alpha_{jt-1}) \subseteq \mathcal{J}_{ijt}$; panel D tests $(dist_j, r_{iht-1}, R_{jt-1}, N_{jt-1}) \subseteq \mathcal{J}_{ijt}$; panel E(a) tests $(dist_j, r_{iht-1}, R_{jt-1}, \alpha_{c(j)t-1}) \subseteq \mathcal{J}_{ijt}$; panel E(b) tests $(dist_j, r_{iht-1}, R_{jt-1}, \alpha_{l(j)t-1}) \subseteq \mathcal{J}_{ijt}$; panel E(c) tests $(dist_j, r_{iht-1}, R_{jt-1}, \alpha_{i(j)t-1}) \subseteq \mathcal{J}_{ijt}$; panel E(d) tests $(dist_j, r_{iht-1}, R_{jt-1}, \alpha_{b(j)t-1}) \subseteq \mathcal{J}_{ijt}$. The variable $\alpha_{c(j)t-1}$ is the average value of α_{jt-1} across all countries that share continent with country j. The variables $\alpha_{l(j)t-1}$, $\alpha_{i(j)t-1}$, and $\alpha_{b(j)t-1}$ are analogous averages across countries that share language, similar income per capita and border, respectively, with j. *Large* firms are those with above median domestic sales in the previous year. Conversely, firm i is *Small* if its domestic sales fall below the median. *Popular* export destinations are those with above median number of exporters in the previous year. Firm i at period t as an *Exporter* to country j if $d_{ijt-1} = 1$ and as a *Non-exporter* if $d_{ijt-1} = 0$. All reported p-values correspond to the test RC; for details on how to compute these p-values, see Appendix A.8. All numbers reported in this table are independent of the value of η chosen as the normalizing constant.

that large firms were informed about α_{jt-1}, it is not surprising that they also have access to information on N_{jt-1}. For small firms, we find evidence to reject the null hypothesis that these firms have access to information on the past number of exporters by destination.

Finally, in panel E, we relax further the informational requirements. Rather than testing whether firms know country-specific information, we test whether they know variables that help predict exports for large groups of countries. Specifically, we test whether firms know average values of α_{jt-1} across groups of countries that share a continent, language, similar income per capita, or a border. Given the results in panels C and D, we focus on testing heterogeneity between large and small firms. For the chemicals sector, we cannot reject the null hypothesis that all firms, large and small firms alike, know these aggregate shifters. For the food sector, we can still reject that small firms have access to average revenue shifters when these averages are computed across countries that share income per capita or a border.

Overall, we find no evidence that firms learn from other exporters or from their prior export experience. However, firms that are either more productive or sell higher quality products tend to have an informational advantage when forecasting market conditions in foreign countries.[35] While we do not investigate why large firms have better information to predict their revenue in foreign markets, the prior literature offers some insights. Weiss (2008) documents that large firms are more likely to participate in international trade fairs and Álvarez and Crespi (2000) note that firms with larger domestic sales are more likely to gather information on foreign markets via programs sponsored by the Chilean Agency for Export Promotion.[36]

7 Counterfactuals

Finally, we use our model and the estimates in Section 5 to explore how changes in firms' information sets and changes in the fixed costs of exporting affect export decisions.

7.1 Changes in Information Sets

We first compute how export profits for the average firm and aggregate exports vary in a counterfactual in which we endow firms with more information about the determinants of export revenues. Specifically, we consider firms whose initial information set includes only three variables: the firm's own lagged domestic sales, r_{iht-1}; the distance to a destination market, $dist_j$; and Chile's lagged aggregate exports to this market, R_{jt-1}. This is the minimal

[35] Our approach allows us to test *passive learning* about a destination-year aggregate shifter of export revenues; we do not test whether firms learn about a firm-specific demand shock (as in Albornoz et al., 2012) or about the demand shifter in a particular buyer-seller relationship (as in Eaton et al., 2014).

[36] The Chilean National Agency for Export Promotion "manages a system that provides information to firms. It is used by companies interested in obtaining information about international markets, for example: external prices, transport costs, entrance regulations and trade barriers." (Álvarez and Crespi, 2000)

Table 6: Effect of Improving Information

Sector:		Chemicals			Food		
		Percentage Change in:			Percentage Change in:		
Firms	Markets	Number of Exporters	Mean Export Profits	Aggregate Exports	Number of Exporters	Mean Export Profits	Aggregate Exports
Panel A: Impact of Adding Information on Aggregate Revenue Shocks to Minimal Information							
All	All	[-5.7, -3.5]	[17.5, 20.6]	[6.4, 9.5]	[-23.2, -17.2]	[57.4, 71.6]	[-2.1, 4.3]
Large	All	[-9.0, -7.1]	[22.0, 24.9]	[0.6, 2.5]	[-20.1, -17.2]	[42.5, 56.8]	[-6.1, -5.3]
Small	All	[0.0, 0.0]	[0.0, 0.1]	[2.0, 2.7]	[0.2, 0.4]	[0.0, 0.1]	[7.4, 12.8]
All	Large	[-2.4, -1.4]	[5.8, 7.0]	[1.1, 2.0]	[-22.2, -17.0]	[5.8, 7.0]	[-3.9, -0.7]
All	Small	[0.3, 0.3]	[6.3, 9.0]	[21.6, 27.3]	[0.8, 1.8]	[16.0, 24.1]	[23.9, 33.4]
Panel B: Impact of Switching from Minimal Information to Perfect Foresight							
All	All	[-10.2, -6.1]	[46.0, 52.9]	[25.1, 33.5]	[-28.0, -20.0]	[90.6, 111.6]	[6.7, 16.8]
Large	All	[-17.3, -12.7]	[59.2, 67.4]	[13.5, 20.3]	[-26.1, -23.9]	[80.8, 97.2]	[1.5, 3.2]
Small	All	[0.3, 0.5]	[0.1, 0.4]	[21.1, 30.8]	[1.3, 2.8]	[2.7, 6.7]	[112.1, 175.1]
All	Large	[-7.5, -4.5]	[29.6, 32.8]	[14.4, 19.0]	[-29.7, -22.8]	[29.6, 32.8]	[1.0, 6.7]
All	Small	[1.7, 2.3]	[16.2, 23.6]	[68.7, 81.6]	[3.9, 6.2]	[53.2, 76.4]	[77.6, 106.8]

Notes: *Large* firms are those with above median domestic sales in the previous year. Conversely, firm i at period t is defined as *Small* if its domestic sales fall below the median. *Large* export destinations are those with above median aggregate exports from Chile in the previous year. Conversely, *Small* export destinations are those with below median aggregate exports from Chile in the previous year. Panel A studies the impact of switching the information set from $(r_{iht-1}, R_{jt-1}, dist_j)$ to $(r_{iht-1}, R_{jt-1}, dist_j, \alpha_{jt-1})$. Panel B studies the impact of switching the information set from $(r_{iht-1}, R_{jt-1}, dist_j)$ to $\alpha_{jt} r_{iht}$. Extreme points of 95% confidence sets computed according to the procedure described in Appendix A.7 are reported in square brackets.

information set we use to compute the moment inequality bounds in Section 5. As indicated in panel A of Table 5, we cannot rule out that every firm knows these three variables.

For this pool of minimally informed firms, we consider extending their information sets in two ways. First, we provide them with additional information on every destination's lagged aggregate revenue shock, α_{jt-1}. According to Table 5, only large firms have access to this information. This counterfactual thus represents an outcome in which all firms gain the same information known by the most informed firms in the economy. Second, we endow firms with all the information necessary to predict perfectly the observable component of potential export revenues.[37] This perfect foresight case allows us to evaluate the overall importance of the informational frictions firms face when predicting their potential export profits.

As we discuss in Section 2, firms in our model obtain all relevant information once they enter a market and, thus, set their prices optimally upon entry. Consequently, the counterfactual change in information sets that we consider here can only affect aggregate export revenues and average export profits by changing the set of firms that self-select into each destination market. We describe in Appendix A.9 the procedure we follow to compute these counterfactual changes in export participation, and report the results in Tables 6 and 7.

The first row in panel A of Table 6 illustrates that, as we provide information on α_{jt-1} for all markets to all firms, the total number of firm-destination pairs with positive exports

[37] Firms remain uncertain about the unobservable component of revenues e_{ijt} (see equations (2) and (3)).

decreases between 3.5% and 5.7% in the chemicals sector. Interestingly, although the total number of firm-destinations decreases, the overall (aggregated across firms and destinations) export revenue in the sector *increases* between 6.4% and 9.5%. Therefore, across all firms, destinations and years, information frictions operate as barriers to trade.

The value to the average firm of acquiring information on α_{jt-1} is quantitatively important. With this information, firms improve their export revenue forecasts and, thus, fewer firms make mistakes in their entry decisions. As a consequence, the realized ex post profits of the average firm in the average market to which it exports increases between 17.5% and 20.6% in the chemical sector.

Panel A in Table 7 provides a detailed accounting of the basis for these counterfactual changes in exports. In the second row in Table 7, we document that export flows for between 708 and 914 firm-destination pairs present in the baseline case would cease if firms acquired information on α_{jt-1}. This would increase average export profits, as mean ex post export profits in these specific firm-destination pairs are negative; the mean export losses among these observations is between 47,000 and 52,000 USD. In short, as information on destination markets improves for firms in the chemical sector, these firms realize that their expectations of export revenues were too optimistic. Accounting only for reductions due to overly optimistic forecasts, overall exports would fall 209 to 219 million USD. However, as information on α_{jt-1} becomes available to all firms, there are also between 512 and 572 new firm-destination pairs with positive export flows. Average export profits and total export revenues in these new destinations are, respectively, between 62,000 and 77,000 USD and between 456 and 512 million USD. Consequently, although firm-destination pairs are lost on net as information increases, export revenue aggregated over all firms and destinations increases, as does average export profits per firm and market.

The results in tables 6 and 7 also allow us to compare the importance of information to firms of different sizes, and to compare against the perfect foresight benchmark. First, increasing the information firms can access always increases realized average ex post profits, particularly for large firms. This is an implication of better informed firms being less likely to make mistakes. Small firms generally benefit less from improving their information because, for most of these firms, their optimal decision is not to export, both before and after acquiring the extra information. Second, increasing access to information for a subset of firms has ambiguous effects on the total number of firm-destination pairs with positive export flows and on the aggregate exports of these firms. For example, in the food sector, informing large firms of the value of lagged aggregate export shocks α_{jt-1} for all destinations leads to a drop in the number of export markets they enter. Aggregate exports of these large exporters end up falling. Third, the information firms may acquire from lagged variables is fairly limited relative to the perfect foresight benchmark. Specifically, comparing the results in panels A and B in tables 6 and 7, we observe that the predicted effect of acquiring information about

Table 7: Decomposing the Effect of Improving Information for All Firms and Destinations

Firm Groups by Export Status	Chemicals Number of Firm-Dest.-Year	Chemicals Mean Export Profits	Chemicals Aggregate Exports	Food Number of Firm-Dest.-Year	Food Mean Export Profits	Food Aggregate Exports
Panel A: Impact of Adding Information on Aggregate Revenue Shocks to Minimal Information						
Always export	[4,830, 5,146]	[121, 126]	[3,021, 3,506]	[5,147, 6,606]	[245, 270]	[9,319, 11,407]
Switch out	[708, 914]	[-52, -47]	[209, 219]	[2,199, 3,114]	[-119, -108]	[1,230, 1,411]
Switch in	[512, 572]	[62, 77]	[456, 512]	[919, 1,034]	[61, 114]	[1,136, 1,679]
Never export	[37,403, 37,968]	[-205, -157]	[3,148, 3,586]	[74,311, 76,708]	[-312, -229]	[8,598, 10,324]
Panel B: Impact of Switching from Minimal Information to Perfect Foresight						
Always export	[4,096, 4,267]	[158, 162]	[2,978, 3,442]	[4,481, 5,611]	[307, 344]	[9,322, 11,341]
Switch out	[1,387, 1,814]	[-63, -55]	[250, 282]	[2,865, 4,127]	[-129, -113]	[1,227, 1,477]
Switch in	[1,046, 1,196]	[95, 119]	[1,215, 1,331]	[1,375, 1,589]	[121, 183]	[2,338, 2,995]
Never export	[36,781, 37,438]	[-210, -162]	[2,389, 2,769]	[73,754, 76,252]	[-316, -233]	[7,396, 9,007]

Notes: *Mean export profits* and *Total export revenues* indicate the model-predicted potential mean export profits and aggregate export revenues, respectively, for all firms, countries and years belonging to the group indicated in the corresponding row. *Mean export profits* are reported in thousands of year 2000 USD and are conditional on the assumption that $\eta = 5$. *Total export revenues* are reported in millions of year 2000 USD and do not depend on the value of η. Panel A studies the impact of switching the information set of every firm i in every country j and year t from $(r_{iht-1}, R_{jt-1}, dist_j)$ to $(r_{iht-1}, R_{jt-1}, dist_j, \alpha_{jt-1})$. Panel B studies the impact of switching the information set of every firm i in every country j and year t from $(r_{iht-1}, R_{jt-1}, dist_j)$ to $\alpha_{jt} r_{iht}$. Extreme points of 95% confidence sets computed according to the procedure described in Appendix A.7 are reported in square brackets.

$r^o_{ijt} \equiv \alpha_{jt} r_{iht}$ is generally larger than the effect of acquiring information about α_{jt-1} only.

7.2 Changes in Fixed Export Costs

Here, we conduct a counterfactual exercise in which we simulate a reduction in exporters' fixed costs of 40%. With this counterfactual, we aim to capture in a stylized way the effect of export promotion programs, such as those instituted in Canada (Van Biesebroeck et al., 2015), Peru (Volpe Martincus and Carballo, 2008), or Uruguay (Volpe Martincus et al., 2010). It is difficult to quantify the precise savings in fixed costs that these measures imply; our choice of a 40% reduction illustrates one possible level.

Counterfactual export probabilities in our setting are partially identified for two reasons: (a) our parameter of interest θ is partially identified; and, (b) we do not want to impose assumptions on the content of the firm's information set $\mathcal{J}_{ijt}$ beyond those we imposed for estimation. Thus, even given a value of θ, export probabilities are not point identified because we only observe a subset Z_{ijt} of the variables firms use to predict export revenues. Thus, we cannot compute firms' expectations, $\mathbb{E}[r_{ijt}|\mathcal{J}_{ijt}]$, exactly and therefore cannot compute the export probabilities in equation (10) directly. We provide details of our algorithm in Appendix A.10. Here, we show the theorem that allows us to bound the probability of export participation given a value of θ and a set of variables $Z_{ijt} \subseteq \mathcal{J}_{ijt}$.

Theorem 3 *Suppose $Z_{ijt} \subseteq \mathcal{J}_{ijt}$ and, for any $\theta \in \Theta$, define $\mathcal{P}(Z_{ijt}; \theta) \equiv \mathbb{E}[\mathcal{P}_{ijt}(\theta)|Z_{ijt}]$, with*

$\mathcal{P}_{ijt}(\theta)$ *defined as*

$$\mathcal{P}_{ijt}(\theta) \equiv \Phi\big(\theta_2^{-1}\big(\eta^{-1}\mathbb{E}[r_{ijt}^o|\mathcal{J}_{ijt}] - \theta_0 - \theta_1 dist_j\big)\big). \tag{19}$$

Then,

$$\mathcal{P}^l(Z_{ijt};\theta) \leq \mathcal{P}(Z_{ijt};\theta) \leq \mathcal{P}^u(Z_{ijt};\theta), \tag{20}$$

where

$$\mathcal{P}^l(Z_{ijt};\theta) = \frac{1}{1 + B^l(Z_{ijt};\theta)}, \tag{21a}$$

$$\mathcal{P}^u(Z_{ijt};\theta) = \frac{B^u(Z_{ijt};\theta)}{1 + B^u(Z_{ijt};\theta)}, \tag{21b}$$

and

$$B^l(Z_{ijt};\theta) = \mathbb{E}\left[\frac{1 - \Phi\big(\theta_2^{-1}\big(\eta^{-1}r_{ijt}^o - \theta_0 - \theta_1 dist_j\big)\big)}{\Phi\big(\theta_2^{-1}\big(\eta^{-1}r_{ijt}^o - \theta_0 - \theta_1 dist_j\big)\big)}\middle| Z_{ijt}\right], \tag{22a}$$

$$B^u(Z_{ijt};\theta) = \mathbb{E}\left[\frac{\Phi\big(\theta_2^{-1}\big(\eta^{-1}r_{ijt}^o - \theta_0 - \theta_1 dist_j\big)\big)}{1 - \Phi\big(\theta_2^{-1}\big(\eta^{-1}r_{ijt}^o - \theta_0 - \theta_1 dist_j\big)\big)}\middle| Z_{ijt}\right]. \tag{22b}$$

The proof of Theorem 3 appears in Appendix A.10.1. If $\theta = \theta^*$, $\mathcal{P}_{ijt}(\theta)$ in equation (19) equals the true export probability in equation (10) and, thus, $\mathcal{P}^l(Z_{ijt};\theta^*)$ and $\mathcal{P}^u(Z_{ijt};\theta^*)$ bound the conditional probability that firm i exports to j at time t.

We use data from the chemicals sector and compare the counterfactual predictions from our moment inequality approach and from the models that require the researcher to specify the covariates included in firms' information sets.[38] Three elements of our model dictate how a change in fixed export costs translates into a change in the number of firms participating in export markets: (a) the initial level of average fixed costs, (b) the heterogeneity across firms in fixed export costs, and (c) firms' expectations of potential export revenues.

First, from equation (10), the level of fixed export costs, $\beta_0 + \beta_1 dist_j$, affects the number of firms that export. Since we reduce fixed costs by a fixed percentage in the counterfactual, the larger the initial estimate of average fixed costs, the larger the reduction in the *level* of fixed export costs. In our setting, the average fixed export costs we recover are largest under the perfect foresight assumption, and thus our counterfactual change in the number of exporters would be largest under that assumption, holding all else equal.

[38] When computing the effect of the reduction in fixed export costs, we assume the parameters $\{\alpha_{jt}; \forall j \text{ and } t\}$ remain invariant. In Appendix B.5, we provide support for this partial equilibrium assumption in our setting.

Table 8: Impact of 40% Reduction in Fixed Costs in Chemicals

Estimator	1996 Argentina	1996 Japan	1996 United States	2005 Argentina	2005 Japan	2005 United States
	Counterfactual Number of Exporters					
Perfect Foresight	67	38	51	70	37	72
Minimal Info.	68	29	38	71	43	56
Moment Inequality	[68, 72]	[12, 95]	[91, 106]	[68, 72]	[5, 89]	[131, 152]

Notes: For the moment inequality estimates, the minimum and maximum predicted values obtained by projecting the 95% confidence set for θ are reported in squared brackets. We compute counterfactual exporter counts by multiplying the observed number of exporters by the counterfactual changes predicted by each of the three models, rounding to the nearest exporter. For the chemicals sector in 2005, we observe 46, 5 and 24 exporters to Argentina, Japan and United States, respectively. For 1996, we observe 44, 5, and 17. All numbers reported in this table are independent of the value of η chosen as normalizing constant.

Second, the joint distribution of firms' heterogeneity in fixed export costs and expectations, $\{(\nu_{ijt}, \mathbb{E}[r^o_{ijt}|\mathcal{J}_{ijt}]); \forall i\}$ will also affect the participation decision of firms in reaction to a decrease in fixed export costs. Specifically, for a given $100(1-\lambda)\%$ reduction in average fixed export costs, the firms that will start exporting will be those for which $\lambda(\beta_0 + \beta_1 dist_j) < \eta^{-1}\mathbb{E}[r^o_{ijt}|\mathcal{J}_{ijt}] - \nu_{ijt} < \beta_0 + \beta_1 dist_j$. Different features of the distributions of $\mathbb{E}[r^o_{ijt}|\mathcal{J}_{ijt}]$ and ν_{ijt} will thus impact the mass of switchers. As an example, consider the case in which there is no heterogeneity across firms in predicted export revenues—i.e. $\mathbb{E}[r^o_{ijt}|\mathcal{J}_{ijt}] = r^*_{jt}$ for all i—and ν_{ijt} is equal to 0. In this case, the response depends on the level of r^*_{jt}. If r^*_{jt} is less than the baseline fixed costs but greater than the counterfactual ones, all firms will stay out in the baseline and all firms will export in the counterfactual. If $r^*_{jt} < \eta\lambda(\beta_0 + \beta_1 dist_j)$, no firm will export in the baseline or counterfactual.

We report our counterfactual estimates in Table 8 for 1996 and 2005. The counterfactual results differ importantly across our three example markets. For Argentina, the three estimation procedures yield very similar answers. Two features of the market explain this similarity. First, given that Argentina is very close to Chile, changes in the distance coefficient β_1 have little impact on entry into Argentina. Therefore, differences across models in the estimate of β_1 will not translate into large differences in predicted export participation. Second, revenues predicted using the minimal information set approach do not differ much from the predicted revenue under perfect foresight. Thus, with similar predicted revenues entering the export participation decision in equation (9), both the perfect foresight and the minimal information models should generate similar counterfactual predictions. For Japan, the two maximum likelihood estimators yield different predictions, and our moment inequality estimator yields predictions that are wide and thus not very informative. The lack of precision in our predictions in this market relates to the relatively few firms we observe exporting to Japan in the data. Finally, for the United States, the moment inequality approach produces predictions that are informative and larger than both maximum likelihood approaches.

8 Extensions

In this section, we extend the model presented in Section 2 in two directions. First, we relax the assumption that a firm's export decision is static and independent of past export participation. To do so, in Section 8.1 we build on Das et al. (2007) and Morales et al. (2017) to allow for sunk export entry costs and forward-looking exporters. Second, in Section 8.2, we relax the assumption, captured in equations (2) and (3), that all firm-country-year specific export revenue shocks are mean independent of exporters' information sets, $\mathcal{J}_{ijt}$. The extensions we discuss here involve larger dimensional parameter vectors than our benchmark specification and, thus, require more computing time to estimate than our benchmark model (see Ho and Rosen, 2017); we thus restrict our estimation below to the chemicals sector.

8.1 Dynamics

The model introduced in Section 2 is static: the export profits of firm i in country j at period t are independent of the previous export path of i in j. Here we extend this model to allow for dynamics. In this extension, exporting firms must still pay fixed costs f_{ijt} in every period in which they choose to export, but they must also pay sunk costs s_{ijt} if they export to j at t and did not export to j at period $t-1$. Therefore, the potential export profits are

$$\pi_{ijt} = \eta_j^{-1} r_{ijt} - f_{ijt} - (1 - d_{ijt-1}) s_{ijt}. \tag{23}$$

We model sunk export costs as:

$$s_{ijt} = \gamma_0 + \gamma_1 dist_j, \tag{24}$$

and assume that firms know these costs when deciding whether to export to a destination.[39] We further assume that information sets evolve independently of past export decisions:

$$(\mathcal{J}_{ijt+1}, f_{ijt+1}, s_{ijt+1}) | (\mathcal{J}_{ijt}, f_{ijt}, s_{ijt}, d_{ijt}) \sim (\mathcal{J}_{ijt+1}, f_{ijt+1}, s_{ijt+1}) | (\mathcal{J}_{ijt}, f_{ijt}, s_{ijt}). \tag{25}$$

If firms are forward-looking, the export dummy d_{ijt} becomes:

$$\begin{aligned} d_{ijt} = \mathbb{1}\{ & \eta^{-1}\mathbb{E}[r^o_{ijt}|\mathcal{J}_{ijt}] - \beta_0 - \beta_1 dist_j - (1 - d_{ijt-1})(\gamma_0 + \gamma_1 dist_j) - \nu_{ijt} \\ & + \rho\mathbb{E}[V(\mathcal{J}_{ijt+1}, f_{ijt+1}, s_{ijt+1}, d_{ijt})|\mathcal{J}_{ijt}, f_{ijt}, s_{ijt}, d_{ijt} = 1] \\ & - \rho\mathbb{E}[V(\mathcal{J}_{ijt+1}, f_{ijt+1}, s_{ijt+1}, d_{ijt})|\mathcal{J}_{ijt}, f_{ijt}, s_{ijt}, d_{ijt} = 0] \geq 0\}, \end{aligned} \tag{26}$$

[39] Contrary to Das et al. (2007), we do not allow for unobserved heterogeneity in sunk costs. Our moment inequality approach may be generalized, at a loss of identification power, to allow for sunk costs $s_{ijt} = \gamma_0 + \gamma_1 dist_j + \nu^s_{ijt}$, with ν^s_{ijt} independent over time and normally distributed with mean zero and constant variance.

Table 9: Export fixed and sunk costs: firm average

Estimator	Cost	Argentina	Chemicals Japan	United States
Benchmark	Fixed	[79.5, 102.6]	[309.2, 414.3]	[181.3, 240.1]
Dynamics	Fixed	[55.8, 109.3]	[853.3, 1,670.0]	[409.2, 800.8]
	Sunk	[384,2, 734,3]	[5,874.4, 11,224.5]	[2,816.6, 5,382.7]

Notes: All estimates are reported in thousands of year 2000 USD and their values scale proportionally with η, which is set equal to 5 (see Section 2.4). Extreme points of 95% confidence sets computed according to the procedure described in Appendix A.7 are reported in square brackets.

where $V(\cdot)$ denotes the value function, ρ is the discount factor and $(\mathcal{J}_{ijt}, f_{ijt}, s_{ijt}, d_{ijt-1})$ is the state vector on which firm i conditions its entry decision in country j at period t. The parameter to estimate is $\theta_D^* \equiv (\beta_0, \beta_1, \sigma, \gamma_0, \gamma_1)$, and we normalize $\eta = 5$ as in the static case.

The firm's export decision now depends on the firm's expectations of both the observable component of static revenues, r_{ijt}^o, and the difference in the value function depending on whether firm i exports to j in t. We follow the approach from the static case to find a measure of r_{ijt}^o, but finding a measure of the difference in value functions is impossible: $V(\cdot)$ at $t+1$ depends on the observed choice at $t+1$, d_{ijt+1}, which is a function of the observed choice at t, d_{ijt}. Therefore, even if firms were only to account for profits at periods t and $t+1$ when making a decision at t, we can only find a measure of either $\mathbb{E}[V(\cdot)|\cdot, d_{ijt} = 1]$ or $\mathbb{E}[V(\cdot)|\cdot, d_{ijt} = 0]$. To solve this lack of measurement, we adjust the Euler approach in Morales et al. (2017). This approach follows the methodology developed in Hansen and Singleton (1982) and Luttmer (1999) for continuous controls, but adapted for our model with discrete controls.[40]

The moment inequalities we employ to compute a confidence set on θ_D^* are the equivalent of the odds-based and revealed-preference inequalities introduced in Section 4.2, adjusted to account for the forward-looking behavior of firms. In Table 9 we report confidence sets for the fixed and sunk costs of exporting. Sunk entry costs are significantly larger than fixed export costs, consistent with Das et al. (2007). Fixed and sunk export costs are also increasing in distance. Furthermore, the sensitivity of these costs to distance is very similar for both types: relative to the bounds for Argentina, the bounds on fixed and sunk costs for the United States and Japan are approximately eight and fifteen times larger.

Comparing the estimated fixed costs in the static model to those from the dynamic model, we find two key differences. First, the bounds are wider; when we estimate fixed and sunk costs simultaneously, we face difficulties in separately identifying both types of costs. Second,

[40] Appendix F shows how to adapt the Euler approach in Morales et al. (2017) to the model described in Section 2 and in equations (23) to (26). Morales et al. (2017) consider models in which the unobserved component ν_{ijt} is constant across groups of countries for each firm-year specific pair. The Euler approach in Morales et al. (2017) has the advantage that it allows us to partially identify the parameter vector of interest without taking a stand on the information set of each exporter, as in Pakes et al. (2015). We also need not specify the number of periods ahead that each firm takes into account when deciding whether to export.

fixed costs for the United States and Japan are larger in the dynamic model. This difference is due to the parameter β_1, the effect of distance on fixed export costs, which we estimate to be larger when accounting for dynamics.[41,42]

8.2 Expected Firm-Country Export Revenue Shocks

In this section, we generalize the model described in Section 2 and allow the firm to observe determinants of export revenue that the researcher does not observe. Specifically, we assume:

$$r_{ijt} = \alpha_{jt} r_{iht} + \omega_{ijt} + e_{ijt}, \qquad \mathbb{E}[e_{ijt}|\mathcal{J}_{ijt}, r_{iht}, f_{ijt}] = 0, \qquad \mathbb{E}[\omega_{ijt}|\mathcal{J}_{ijt}] = \omega_{ijt}; \tag{27}$$

and define a subset $\mathcal{W}_{ijt}$ of the information set of the firm $\mathcal{J}_{ijt}$ such that the firm's expectations of the observable component of export revenues satisfy $\mathbb{E}[\alpha_{jt} r_{iht}|\mathcal{J}_{ijt}] = \mathbb{E}[\alpha_{jt} r_{iht}|\mathcal{W}_{ijt}]$. The export dummy d_{ijt} therefore becomes

$$d_{ijt} = \mathbb{1}\{\eta^{-1}\mathbb{E}[\alpha_{jt} r_{iht}|\mathcal{W}_{ijt}] - \beta_0 - \beta_1 dist_j - (\nu_{ijt} - \eta^{-1}\omega_{ijt}) \geq 0\}, \tag{28}$$

and the components of revenue known to the firm but not the researcher follow the distribution:

$$\left.\begin{pmatrix} \omega_{ijt} \\ \nu_{ijt} \end{pmatrix}\right| (\mathcal{W}_{ijt}, dist_j) \sim \mathbb{N}\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_\omega^2 & \sigma_{\omega\nu} \\ \sigma_{\omega\nu} & \sigma^2 \end{pmatrix}\right). \tag{29}$$

If we assume both that potential exporters have perfect foresight over the component $\alpha_{jt} r_{iht}$ of their export revenues, $\mathbb{E}[\alpha_{jt} r_{iht}|\mathcal{W}_{ijt}] = \alpha_{jt} r_{iht}$, and that the firm observes all components of export revenue, i.e. $e_{ijt} = 0$ for all i, j, and t, then we can estimate the parameter vector $(\{\alpha_{jt}\}_{j,t}, \beta_0, \beta_1, \sigma_\omega, \sigma_{\omega\nu}, \sigma_\nu)$ using the procedure introduced in Heckman (1979).[43] Appendix G.1 shows how to use moment inequalities to estimate a confidence set for this parameter vector when we both allow e_{ijt} to differ from zero and impose weaker assumptions on the firm's information. We require only that the researcher observes a vector $Z_{ijt} \subseteq \mathcal{W}_{ijt}$.

[41] It may seem counterintuitive that accounting for sunk export costs increases the estimates of fixed export costs. This pattern would not arise if exporters were to decide whether to export at period t by comparing the static profits at t with the sum of fixed and sunk export costs. However, the presence of the value function in equation (26) makes the pattern we observe more likely, as firms in the dynamic model decide whether to export at any given period t taking into account the effect their decision has on subsequent periods' potential export profits. Specifically, when exiting an export destination, exporters take into account that they would have to repay the sunk costs if they were to re-enter in subsequent periods. This implies that, if fixed costs in the dynamic model were to remain at the values estimated in the static model, firms would be less likely to exit than in the static model. Therefore, rationalizing the observed exit behavior in the data requires larger fixed export costs in the dynamic model with forward-looking firms than in the static case.

[42] While our estimates of export sunk costs for Argentina are similar to those in Das et al. (2007), those for the U.S. and Japan are significantly larger. The differences in these estimates could be due to the differences in the datasets or in the specification of sunk costs and information sets.

[43] When estimating the model introduced in Heckman (1979), it is typical to fix one of the components of the variance matrix in equation (29) as a normalization. In our case, we opt to maintain the normalization $\eta = 5$.

Table 10: Export fixed costs: specification with expected unobserved revenue shocks

		Chemicals		
Specification	Estimator	Argentina	Japan	United States
Baseline: $\omega_{ijt} = 0$	Perfect Foresight (MLE)	868.0 (106.5)	2,621.4 (315.9)	1,645.0 (199.3)
	Moment Inequality	[79.5, 102.6]	[309.2, 414.3]	[181.3, 240.1]
Extension: $\omega_{ijt} \neq 0$	Perfect Foresight (MLE)	323.4 (34.8)	983.6 (107.5)	615.9 (66.5)
	Moment Inequality	[75.2, 767.7]	[435.9, 3,449.4]	[284.7, 1,654.2]

Notes: All estimates are reported in thousands of year 2000 USD and their values scale proportionally with η, which is set equal to 5. Estimates reported for the extension assume $\alpha_{jt} = \alpha \; \forall \; j$ and t . Extreme points of 95% confidence sets computed according to the procedure described in Appendix G.1.3 appear in square brackets.

When we allow firms to account for ω_{ijt} in their export entry decision while imposing only that we observe a vector $Z_{ijt} \subseteq \mathcal{W}_{ijt}$, the export revenue coefficients $\{\alpha_{jt}\}_{j,t}$ are only partially identified and must be estimated jointly with the remaining parameters $(\beta_0, \beta_1, \sigma_\omega, \sigma_{\omega\nu}, \sigma_\nu)$. Given that our sample period covers 10 years and 22 countries, we would have to estimate jointly a confidence set for over 200 parameters. While this is theoretically possible, as far as we know, it is infeasible to compute using our chosen inference approach. Therefore, we simplify the problem by assuming $\alpha_{jt} = \alpha$ for every j and t and, thus, estimate the parameter vector $\theta_S \equiv (\alpha, \beta_0, \beta_1, \sigma_\omega, \sigma_{\omega\nu}, \sigma_\nu)$.

We report the results in Table 10. To facilitate comparison, the top portion of this table repeats the results for the baseline case. Allowing the export revenue shock ω_{ijt} to differ from zero yields perfect foresight estimates of the export fixed costs that are smaller than the parallel estimates computed in the baseline case. They are, however, still larger than the baseline moment inequality estimates. For example, the perfect foresight estimates computed in the extended model with $\omega_{ijt} \neq 0$ yields estimates of fixed export costs in Argentina equal to \$323,400, significantly smaller than the baseline perfect foresight estimate of \$868,000. However, they remain larger than the upper bound of the moment inequality baseline estimate, which equals \$102,600. Unfortunately, our moment inequality bounds for the case in which we allow ω_{ijt} to differ from zero are too wide to be informative. A consequence of our bounds being uninformative is that they always contain the corresponding maximum likelihood estimates.[44]

9 Conclusion

We study the extensive margin decision of firms to enter foreign export markets. This participation decision drives much of the variation in trade volume. Thus, to predict how trade

[44] Equations (27) to (29) may be compatible with restrictions on θ_S not accounted for by the inequalities we describe in Appendix G. These additional restrictions may further reduce the size of the confidence set for θ_S. We leave this exploration for future work.

flows will adjust to changes in the economic environment, policymakers first need a measure of the determinants of firms' decisions to engage in exporting. In this paper, we measure these determinants using a moment inequality approach that exploits relatively weak assumptions on the content of exporters' information sets. We show how to use our moment inequalities to recover the fixed costs of exporting, to quantify how firms will react to counterfactual changes in the information they access and in export trade costs, and also to test whether firms use certain key variables to forecast their potential export revenues.

The estimated fixed costs from our inequality model are between ten and thirty percent of the size of the costs found using approaches that require the researcher to specify fully the content of exporters' information sets. When evaluating the effect of endowing potential exporters with better information, we find average export profits of large firms increase significantly while those of small firms remain largely invariant. The overall share of firm-destination pairs with observed positive exports decreases, and this reduction is concentrated among large firms and in large markets. The total volume of exports may increase or decrease.

Finally, we test alternative assumptions on the content of the information sets firms use in their export decision—that is, we test what exporters know. We find important heterogeneity in information sets by firm size: large firms have better information on foreign markets than small firms. While large firms have access to country-by-country information on market-specific demand and trade costs shifters, small firms do not.

References

Ahn, Hyungtaik and Charles F. Manski, "Distribution Theory for the Analysis of Binary Choice under Uncertainty with Nonparametric Estimation of Expectations," *Journal of Econometrics*, 1993, *56*, 291–321. [2, 11]

Albornoz, Facundo, Héctor F. Calvo Pardo, Gregory Corcos, and Emanuel Ornelas, "Sequential Exporting," *Journal of International Economics*, 2012, *88* (1), 1–24. [23]

Almunia, Miguel, Pol Antràs, David Lopez-Rodriguez, and Eduardo Morales, "Venting Out: Exports during a Domestic Slump," *mimeo*, 2018. [5]

Álvarez, Roberto and Gustavo Crespi, "Exporter Performance and Promotion Instruments: Chilean Empirical Evidence," *Estudios de Economía*, 2000, *27* (2), 225–241. [23]

Andrews, Donald W.K. and Gustavo Soares, "Inference for Parameters Defined by Moment Inequalities Using Generalized Moment Selection," *Econometrica*, 2010, *78*, 119–157. [8, 10, 24, 25, 28, 68]

_ and Patrik Guggenberger, "Validity of Subsampling and Plug-in Asymptotic Inference for Parameters Defined by Moment Inequalities," *Econometric Theory*, 2009, *25*, 669–709. [10, 28]

_ and Xiaoxia Shi, "Inference Based on Conditional Moment Inequalities," *Econometrica*, 2013, *81* (2), 609–666. [16]

Antràs, Pol, Teresa Fort, and Felix Tintelnot, "The Margins of Global Sourcing: Theory and Evidence from U.S. Firms," *American Economic Review*, 2017, *107* (9), 2514–2564. [1, 4, 7, 8]

Arkolakis, Costas, "Market Penetration Costs and the New Consumers Margin in International Trade," *Journal of Political Economy*, 2010, *118* (6), 1151–1199. [1, 2]

_, Sharat Ganapati, and Marc-Andreas Muendler, "The Extensive Margin of Exporting Goods: A Firm-level Analysis," *mimeo*, November 2015. [1, 4]

Armstrong, Tim B., "Weighted KS Statistic for Inference on Conditional Moment Inequalities," *Journal of Econometrics*, 2014, *182* (2), 92–116. [16]

_, "Asymptotically Exact Inference in Conditional Moment Inequality Models," *Journal of Econometrics*, 2015, *186* (1), 51–65. [16]

_ and Hock Peng Chan, "Multiscale Adaptive Inference on Conditional Moment Inequalities," *Journal of Econometrics*, 2016, *194* (1), 24–43. [16]

Aw, Bee Yan, Mark J. Roberts, and Daniel Yi Xu, "R&D Investment, Exporting, and Productivity Dynamics," *American Economic Review*, 2011, *101* (4), 1312–1344. [4]

Bagnoli, Mark and Theodore C. Bergstrom, "Log-Concave Probability and Its Applications," *Economic Theory*, 2005, *26* (2), 445–469. [14]

Bernard, Andrew B., Andreas Moxnes, and Karen Helen Ulltveit-Moe, "Two-Sided Heterogeneity and Trade," *Review of Economics and Statistics*, 2017, *forthcoming.* [3]

—, J. Bradford Jensen, Stephen J. Redding, and Peter K. Schott, "The Margins of US Trade," *American Economic Review*, 2010, *99*, 487–493. [1, 4]

Biesebroeck, Johannes Van, Emily Yu, and Shenjie Chen, "The Impact of Trade Promotion Services on Canadian Exporter Performance," *Canadian Journal of Economics*, 2015, *48* (4), 1481–1512. [26]

Bilbiie, Florian, Fabio Ghironi, and Marc Melitz, "Endogenous Entry, Product Variety and Business Cycles," *Journal of Political Economy*, 2012, *120* (2), 304–345. [4]

Bilir, Kamran L. and Eduardo Morales, "Innovation in the Global Firm," *mimeo*, November 2016. [4]

Blaum, Joaquín, Claire Lelarge, and Michael Peters, "The Gains from Input Trade with Heterogeneous Importers," *American Economic Journal: Macroeconomics*, 2017, *forthcoming.* [1, 4]

Blum, Bernard S., Sebastiàn Claro, and Ignatius J. Horstmann, "Occasional and Perennial Exporters," *Journal of International Economics*, 2013, *90* (1), 65–74. [5]

Bugni, Federico A., Ivan A. Canay, and Xiaoxia Shi, "Alternative Specification Tests for Partially Identified Models Defined by Moment Inequalities," *Journal of Econometrics*, 2015, *185* (1), 259–282. [3, 20, 10, 24, 25, 28, 29]

—, —, and —, "Inference for Functions of Partially Identified Parameters in Moment Inequality Models," *Quantitative Economics*, 2016, *8* (1), 1–38. [19]

Cherkashin, Ivan, Svetlana Demidova, Hiau Looi Kee, and Kala Krishna, "Firm Heterogeneity and Costly Trade: A New Estimation Strategy and Policy Experiments," *Journal of International Economics*, 2015, *96* (1), 18–36. [1]

Chernozhukov, Victor, Sokbae Lee, and Adam M. Rosen, "Intersection Bounds: Estimation and Inference," *Econometrica*, 2013, *81* (2), 667–737. [16]

Chetverikov, Denis, "Adaptive Test of Conditional Moment Inequalities," *Econometric Theory*, 2013, *forthcoming.* [16]

Coşar, Kerem A., Paul L. E. Grieco, Shengyu Li, and Felix Tintelnot, "What Drives Home Market Advantage?," *Journal of International Economics*, 2016, *110*, 135–150. [18]

Crawford, Gregory S. and Ali Yurukoglu, "The Welfare Effects of Bundling in Multichannel Television Markets," *American Economic Review*, 2012, *102* (2), 643–685. [4]

Cunha, Flavio and James J Heckman, "Identifying and Estimating the Distributions of Ex Post and Ex Ante Returns to Schooling," *Labour Economics*, 2007, *14*, 870–893. [1]

Das, Sanghamitra, Mark J. Roberts, and James R. Tybout, "Market Entry Costs, Producer Heterogeneity, and Export Dynamics," *Econometrica*, 2007, *75* (3), 837–873. [1, 3, 5, 8, 29, 30, 31, 65, 66]

Department of Commerce, "A Profile of U.S. Importing and Exporting Companies, 2013-2014," *Department of Commerce*, 2016. [1]

Dickstein, Michael J., Mark Duggan, Joe Orsini, and Pietro Tebaldi, "The Impact of Market Size and Composition on Health Insurance Premiums: Evidence from the First Year of the ACA," *American Economic Review Papers and Proceedings*, 2015, *105* (5), 120–125. [4]

Doraszleski, Ulrich and Jordi Jaumandreu, "R&D and Productivity: Estimating Endogenous Productivity," *Review of Economic Studies*, 2013, *80*, 1338–1383. [4]

Eaton, Jonathan, David Jinkins, James R. Tybout, and Daniel Yi Xu, "Two Sided Search in International Markets," *mimeo*, June 2016. [3]

_ , Marcela Eslava, C.J. Krizan, Maurice Kugler, and James R. Tybout, "A Search and Learning Model of Export Dynamics," *mimeo*, February 2014. [1, 23]

_ , _ , Maurice Kugler, and James R. Tybout, "The Margins of Entry into Export Markets: Evidence from Colombia," in Elhanan Helpman, Dalia Marin, and Thierry Verdier, eds., *The Organization of Firms in a Global Economy*, Cambridge: Harvard University Press, 2008. [1]

_ , Samuel Kortum, and Francis Kramraz, "An Anatomy of International Trade: Evidence from French Firms," *Econometrica*, 2011, *79* (5), 1453–1498. [1, 3]

Eizenberg, Alon, "Upstream Innovation and Product Variety in the U.S. Home PC Market," *Review of Economic Studies*, 2014, *81*, 1003–1045. [4]

Freeman, Richard, *The Market for College-Trained Manpower*, Cambridge MA: Harvard University Press, 1971. [4]

Hansen, Lars Peter and Kenneth J. Singleton, "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models," *Econometrica*, 1982, *50* (5), 1269–1286. [30]

Head, Keith and Thierry Mayer, "Gravity Equations: Workhorse,Toolkit, and Cookbook," in Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff, eds., *Vol.4 of Handbook of International Economics*, Elsevier, 2014, pp. 131–195. [8]

Heckman, James J., "Sample Selection Bias as a Specification Error," *Econometrica*, 1979, *47* (1), 153–161. [31]

_ and Bo E. Honoré, "The Empirical Content of the Roy Model," *Econometrica*, 1990, *58* (5), 1121–1149. [14]

Ho, Katherine, "Insurer-Provider Networks in the Medical Care Market," *American Economic Review*, 2009, *99* (1), 393–430. [4]

_ and Adam Rosen, "Partial Identification in Applied Research: Benefits and Challenges," in Bo Honoré, Ariel Pakes, Monika Piazzesi, and Larry Samuelson, eds., *Advances in Economics and Econometrics: Eleventh World Congress (Econometric Society Monographs) (Volume II)*, Cambridge University Press, 2017, pp. 307–359. [29]

— **and Ariel Pakes**, "Hospital Choices, Hospital Prices, and Financial Incentives to Physicians," *American Economic Review*, 2014, *104* (12), 3841–3884. [4]

Holm, Sture, "A Simple Sequentially Rejective Multiple Test Procedure," *Scandinavian Journal of Statistics*, 1979, *6* (2), 65–70. [21, 10, 11, 28]

Holmes, Thomas J., "The Diffusion of Wal-Mart and Economies of Density," *Econometrica*, 2011, *79* (1), 253–302. [4]

Luttmer, Erzo C. J., "What Level of Fixed Costs can Reconcile Consumption and Stock Market Returns?," *Journal of Political Economy*, 1999, *107* (5), 969–997. [30]

Maini, Luca and Fabio Pammolli, "Reference Pricing as a Deterrent to Entry: Evidence from the European Pharmaceutical Market," *mimeo*, 2017. [4]

Manski, Charles F., "Nonparametric Estimation of Expectations in the Analysis of Discrete Choice Under Uncertainty," in William Barnett, James Powell, and George Tauchen, eds., *Nonparametric and Semiparametric Methods in Econometrics and Statistics*, Cambridge: Cambridge University Press, 1991. [2, 11]

— , "Adolescent econometricians: How do youth infer the returns to schooling?," in Charles Clotfelter and Michael Rothschild, eds., *Studies of Supply and Demand in Higher Education*, Chicago: University of Chicago Press, 1993, pp. 43–60. [1, 10]

— , "Measuring Expectations," *Econometrica*, 2004, *5* (9), 1329–1376. [1, 4]

— **and David A. Wise**, *College Choice in America*, Cambridge MA: Harvard University Press, 1983. [4]

Martincus, Christian Volpe and Jerónimo Carballo, "Is Export Promotion Effective in Developing Countries? Firm-Level Evidence on the Intensive and the Extensive Margins of Exports," *Journal of International Economics*, 2008, *76* (1), 89–106. [26]

— **,** — **, and Pablo Garcá**, "Entering New Country and Product Markets: Does Export Promotion Help?," *Review of World Economics*, 2010, *146* (3), 437–467. [26]

Mayer, Thierry and Soledad Zignago, "Notes on CEPII's Distances Measures," *mimeo*, December 2011. [9]

Melitz, Marc, "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 2003, *71*, 1695–1725. [1]

Morales, Eduardo, Gloria Sheu, and Andrés Zahler, "Extended Gravity," *mimeo*, June 2017. [4, 7, 29, 30]

Moxnes, Andreas, "Are Sunk Costs in Exporting Country-Specific?," *Canadian Journal of Economics*, 2010, *43* (2), 467–493. [1]

Pakes, Ariel, "Alternative Models for Moment Inequalities," *Econometrica*, 2010, *78* (6), 1783–1822. [2, 15]

— **and Jack R. Porter**, "Moment Inequalities for Multinomial Choice with Fixed Effects," *mimeo*, December 2015. [15]

— **, — , Katherine Ho, and Joy Ishii**, "Moment Inequalities and their Application," *Econometrica*, 2015, *83* (1), 315–334. [2, 15, 30]

Roberts, Mark J. and James R. Tybout, "The Decision to Export in Colombia: An Empirical Model of Entry with Sunk Costs," *American Economic Review*, 1997, *87* (4), 545–564. [3]

Romano, Joseph P. and Azeem M. Shaikh, "Inference for Identifiable Parameters in Partially Identified Econometric Models," *Journal of Statistical Planning and Inference*, 2008, *138* (9), 2786–2807. [10, 28]

Ruhl, Kim J. and Jonathan L. Willis, "New Exporter Dynamics," *International Economic Review*, 2017, *58* (3), 703–725. [1, 3]

Shi, Xiaoxia, Matthew Shum, and Wei Song, "Estimating Semi-parametric Panel Multinomial Choice Models using Cyclic Monotonicity," *Econometrica*, 2017, *forthcoming.* [15]

Simonovska, Ina and Michael Waugh, "The Elasticity of Trade: Estimates and Evidence," *Journal of International Economics*, 2014, *92* (1), 34–50. [8]

Vannoorenberghe, Gonzague, "Firm-level Volatility and Exports," *Journal of International Economics*, 2012, *86* (1), 57–67. [5]

Weiss, Kenneth D., *Building an Import-Export Business*, Hoboken NJ: John Wiley & Sons, Inc., 2008. [23]

Willis, Robert J. and Sherwin Rosen, "Education and Self-Selection," *Journal of Political Economy*, 1979, *87* (5), 7–36. [4]

Wollman, Thomas G., "Trucks Without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles," *American Economic Review*, 2017, *forthcoming.* [4]

Wooldridge, Jeffrey M., *Econometric Analysis of Cross Section and Panel Data*, Cambridge MA: MIT Press, 2002. [12]

Yatchew, Adonis and Zvi Griliches, "Specification Error in Probit Models," *The Review of Economics and Statistics*, 1985, *67* (1), 134–139. [12, 47]