Demand Volatility and Export Entry☆

Michael Olabisi

University of Michigan

Abstract

Does demand volatility affect exporters' choices of foreign destinations? All export destinations exhibit volatility, with demand from some being more volatile than others. To answer the question, I develop a simple model of trade with heterogeneous firms facing stochastic demand. Firms in the model incur adjustment costs in response to fluctuations in demand. This model predicts that fewer exporters will serve destinations with higher demand volatility, as adjustment costs decrease profits. I test this prediction using data on the universe of Chinese exports from 2000 to 2006 at the firm level. As expected, fewer exporters in the data enter destinations with high demand volatility. Additional firm level regressions show a negative and statistically significant relationship between demand volatility and aggregate trade levels.

Keywords: Demand Volatility; Adjustment Costs; Export Entry

1. Introduction

Demand volatility is an undeniable feature of international trade. The typical exporter serves a narrow set of destinations with non-trivial year-on-year growth shocks that are not reflected in aggregate trade figures. In the average year, 11% of trade destinations represented by country-product pairs experience a shock that reduces demand to zero. More generally, year-on-year growth shocks for trade at that level cover the full spectrum of negative and positive values. The drivers of demand volatility include income shocks, government policy and structural economic factors that shape buyers decisions at the level of destinations - defined as unique country and product pairings like the US imports of bicycles.[1].

Do exporters avoid destinations with high demand volatility? Phrased differently, do exporters choose country and product combinations with low demand volatility over similar destinations high demand volatility? The question is relevant to understanding how volatility determines the margins of international trade, and to the exporter choices that lead to aggregate trade patterns.

To address the question formally, this paper develops a model of trade with adjustment costs.[2] Expected profits fall with adjustment costs, e.g. the costs of hiring in times of peak demand or deactivating equipment during lulls. With lower expected profits, fewer firms will find entry into a given destination profitable. Firms in the model are heterogeneous in terms of productivity, so that the more productive are able to enter destinations across a wider range of demand volatility. The model's predictions for trade's extensive margin are unambiguous: destinations with high demand volatility will have fewer exporters and lower levels of trade.

[1]This paper characterizes each country-product destinations by a level of demand volatility, derived from its history; joining a long tradition of scholarship on demand shocks that are exogenous to firms' choices and traits (Blum et al., 2013; Rob and Vettas, 2003; Staiger and Wolak, 1992; Viner, 1922). In proposing that demand shocks in export destinations may be exogenous to firms and their technical efficiency, the paper follows Foster et al. (2008)

[2]A robust body of work in microeconomics and macroeconomics describes the nature of labor and capital adjustment costs, and how they influence aggregate economic outcomes, e.g. (Bloom et al., 2007; Cooper and Haltiwanger, 2006; Pindyck, 1982; Lucas, 1967). Of these, Lucas (1967) raises the specific concern that adjustments change producer's per-unit costs - an idea that features notably in my model.

I take the model's predictions to a unique combination of firm-level and global trade data: I observe the destination choices of exporters from the universe of Chinese export transactions between 2000 and 2006, and measure demand volatility using aggregate imports between 1995 and 2005 for each of these destinations. I define destinations as the import stream for unique product-country combinations, like the US imports of truck tires, or Kenyan imports of bicycles. The UN COMTRADE database provides this trade information for narrowly defined HS6 product categories at the country level. Using global demand to estimate demand volatility addresses possible concerns that firm-level shocks drive the measured volatility, rather than patterns of markets' demand.[3]

The data show a negative and statistically significant relationship between demand volatility and the number of exporters that serve a destination. First, destinations with high demand volatility were more likely to have zero Chinese exporters. Demand volatility was 105% higher on average in destinations that no Chinese exporter served in all the years covered. After controlling for destination size, country and product factors, the likelihood that no Chinese exporter serves a destination increases by 2.8% with one standard deviation from the mean value of demand volatility. With the liberalization that accompanied China's WTO accession in 2001, the number of exporters nearly tripled between 2000 and 2006. The destinations that exporters choose to serve is informative in this context of trade expansion.

Second, conditional on having at least one Chinese exporter, the number of exporters serving a destination is 5% lower for destinations with demand volatility one standard deviation above the mean. The negative estimated effect of demand volatility on exporter numbers holds up to alternative definitions of demand volatility, exporter numbers and a variety of specifications that address concerns about the causal nature of this relationship. The effect of demand volatility on trade is not limited to the extensive margin – the value of exports to destinations are 2% lower with a one standard deviation increase in demand volatility in my more conservative estimates.

The findings imply that the prospect of adjustment costs deters firms from incurring the up-front costs of exporting to destinations with high demand

[3]Using Chinese firm-level data is informative for an empirical exercise that describes exporter behavior in general. China is the world's largest exporter. Furthermore, the high level of correlation between China's aggregate exports and the rest of the world suggests that its exporters behave like firms of other nationalities.

volatility. I define adjustment costs in this paper to broadly include capital adjustment costs, as well as the costs of firing and hiring employees. From an exporter's perspective, the findings support the argument in Cuñat and Melitz (2012) that countries may derive a comparative advantage in exporting products with highly volatile demand if they have flexible labor regimes – and therefore lower labor adjustment costs. One can make broad statements about exporter behavior and demand volatility based on these findings for two reasons: Chinese exports are correlated with global exports (Amiti and Freund, 2010), and China is currently the world's largest exporter.

For an importing country, demand volatility represents a potential barrier to economic growth. Most producers in developing economies require imported inputs, usually imported capital goods (Connolly, 2003). In the data, I find that prices are slightly higher for Chinese exports of capital goods to destinations with high demand volatility. Controlling for quality differences should make these price differences starker, as the literature generally finds lower prices in the developing economy destinations where high levels of demand volatility are more common, e.g. Manova and Zhang (2012) and Harrigan et al. (2011).[4] Demand volatility also leads to fewer imported varieties after controlling for other determinants of trade. Firms in destinations with high demand volatility therefore lose the potential benefits of new and more imported varieties described in Goldberg et al. (2010).

Based on the foregoing, considering demand volatility can contribute to our understanding of aggregate trade, its extensive margins and the presence of zeros in international trade aggregates. In that sense, the paper complements the large body of work that includes Armenter and Koren (2014), Baldwin and Harrigan (2011) and Helpman et al. (2008). Helpman et al. (2008) shows that zero bilateral trade is more likely with long distances or high marginal costs. This paper extends the idea, suggesting that high adjustment costs due to demand volatility can also increase the occurrence of zero trade. Others describe the determinants of trade in terms of firm-level productivity and geographically-driven trade costs, e.g., Anderson and Van Wincoop (2003) and Melitz (2003). This paper suggests a role for a

[4]This finding in AppendixA.8 agrees with the evidence in Eaton and Kortum (2001) that capital goods have higher relative prices in developing economies (where I find the most volatility). Other papers that find lower prices for exports to developing economies do not focus on specific sub-categories like capital goods, which make up less than 11% of global exports by value.

market feature like demand volatility. In addition to providing evidence that volatility influences exporter choice, the paper contributes to the literature in three ways.

First, by focusing on the decisions of firms to serve specific combinations of products and countries, I provide an approach for explaining firm level trade choices that country-level measures like GDP, exchange rates and geographic distance may not adequately capture. My units of observation are destinations represented by product-country combinations, like the US imports of bicycles: The GDP of the US may not be relevant to a Chinese exporter of bicycles if GDP is a poor predictor of the demand for bicycles in particular. To such an exporter, historical information on imports of bicycles into each potential foreign market is more valuable.[5] Therefore, this paper describes demand from destinations as a random walk with a constant growth trend, following Carroll et al. (2011) and Hall (2004). Exporters forecast profits from the volatility and trend observed in each destination's demand history, and enter markets on that basis. Market-specific volatility for a given product reflects income, policy or other transactional frictions not explained by equilibrium prices.

Second, the paper extends the literature on investment under uncertainty to the context of trade. The initial costs of setting up overseas trading networks are analogous to investments made in the expectation of future returns. Future demand is not known, but rational agents characterize its expected value using historical information. This yields new testable insights on how exporters choose destinations. If adjustment costs are expected to be higher for destinations with high demand volatility, then fewer exporters should serve those destinations. Recent related papers show that policy uncertainty reduces exports, when trade costs are driven by policy (Handley and Limão, 2012, 2013). In that context, exporters are more likely to invest in foreign destinations with stable tariff regimes. Earlier work by Dixit (1989) shows that with uncertain prices, firms require prices above a certain threshold to expand their operations. The same relationship between uncertainty and in-

[5]Most exporters serve few foreign destinations. The median number of products and countries per exporter in the data are 5 and 4 respectively. Country-level measures like GDP mask shocks that may be important to firms that focus on a few product categories. For example, US imports of pure fructose (HS 170250) are more volatile than Rwandan imports of truck tires (HS 401120). The two countries' aggregate measures of GDP and demand volatility are at opposite ends of the spectrum.

vestment holds for exchange rate uncertainty (Das et al., 2007; Frankel and Rose, 2002; Glick and Rose, 2002). The novelty in this paper is its focus on demand, as the volatility of trade costs, prices and exchange rates only explain small shares of the year-on-year variation in trade.

Finally, I provide a simple and intuitive measure of volatility: i.e. the sum of squared deviations from a trend for a series. For this paper, the trend used to derive the measure is linear, though the definition is flexible enought to admit other trend specifications. Related papers define volatility as the standard deviation of year-on-year growth rates, but measuring growth rates is problematic when observations include zero.[6] The index of volatility introduced by this paper avoids such issues of measurement.

The rest of the paper is organized as follows: Section 2 presents a stylized model in the tradition of Melitz (2003) and Chaney (2008) to motivate the empirics. Section 3 follows, with the data, formal definitions for key variables, empirical specifications and results. Section 4 discusses the implications and concludes.

2. Model

2.1. Adjustment Costs with Stochastic Demand

In this model, exporters decide on foreign destinations using information on historical demand shocks. Exporters can form unbiased expectations of demand for each year in a forward-looking planning horizon, given the trajectory of past demand. Previous related papers model demand uncertainty in a framework that requires exporters to learn about demand e.g. Akhmetova and Mitaritonna (2012) and Nguyen (2011). In that context, the conditional distribution of possible demand outcomes is taken as unknown. In this paper, the conditional distribution of demand outcomes depends only on historical demand realizations.

In other words, firms form expectations of volatility and future demand for each destination from its historical demand. For example, a destination that imports exactly $1m for all years between 1995 and 2005 will lead all

[6]One can correct the conventional measure of growth volatility by using a mid-point growth measure, which bounds growth between -2 and 2, but it does not help that those extreme values of growth may be outliers that skew the measure of volatility. Papers that define volatility as the standard deviation of growth include di Giovanni and Levchenko (2009) and Koren and Tenreyro (2007).

exporters to expect little no growth, and negligible demand volatility. (Past volatility is taken as the predictor of future volatility). This aggregate import volatility described in the paper applies to all exporters, as historical demand is common knowledge to all firms.[7]

I illustrate the section's main idea using refineries. Updating the capacity of a refinery in production is costly. For a refinery that exports, each production run incurs upfront customization costs, as gasoline and diesel blends differ by country. Therefore, before a refiner enters a foreign destination, it must consider the usual per-unit marginal costs, the upfront customization costs and the costs of capacity adjustments expected in its planning horizon for that destination. The hypothetical demand trajectories in Figure 1 illustrate the relevance of capacity adjustments. Judging from the plots of historical demand, the two destinations in the graph have the same expected size, but the scales of deviations from the expected trajectory of demand differ. As the producer must consider the relative costs of scaling production to match demand in each period, the destination on the right panel becomes more preferable if the historical pattern of volatility persists.

Exporters only serve destinations with non-negative expected profits – after accounting for adjustment costs. With adjustment costs, some marginally profitable destinations with high demand volatility become unprofitable.[8] For the refinery example that motivates this section, a rational exporter may find the destination on the right panel of the figure profitable, but avoid the destination on the left. Large production scale adjustments reduce realized profits, as documented in the literature on adjustment costs (Cooper and Haltiwanger, 2006; Lucas, 1967).

Destinations with high demand volatility and expected adjustment costs will attract fewer exporters – as long as expected demand volatility reflects

[7]One could conceptually create firm-specific measures of demand volatility that use weighted composites of the demand history for each firm's target destinations. Larger shocks at the firm-level may make such a firm-specific measure higher than the currently proposed destination-specific measure – the differences may be more notable for new entrants or small firms. However, the relationship between these two conceptions of demand volatility depends on the 'diversification effect' - where demand shocks offset one another for multi-destination exporters, which account for the largest share of exports.

[8]The model proposed here reverts to a conventional model of trade if one collapses the planning horizon to one period, and presumes perfect information about future demand. Adjustment costs disappear with these additional assumptions. In that sense, this paper extends the conventional model of trade to include multi-period sales and information.

Figure 1: Demand with Stochastic Shocks

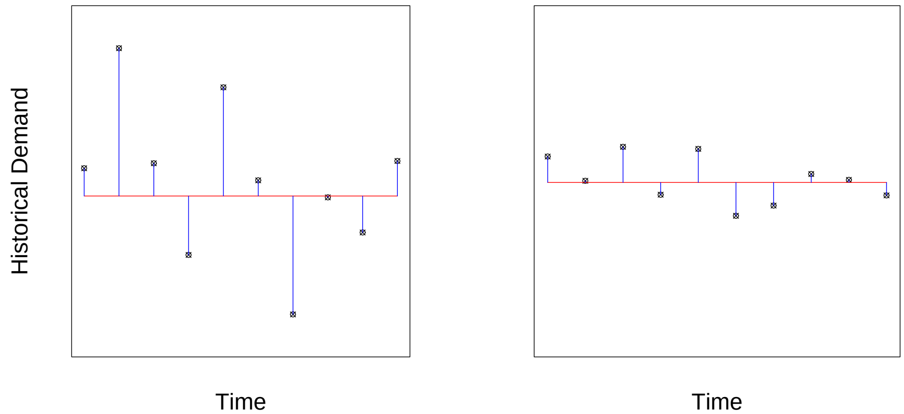

The graph is purely illustrative - its two hypothetical destinations have equal average size, but different levels of demand volatility. The horizontal line represents projected demand and vertical lines represent deviations from the demand trajectory.

the observed demand history. Demand volatility measures the variability in demand over time, and as expected, high demand volatility corresponds to high adjustment costs for firms serving a destination. The next steps relate demand volatility to profits.

Formally, profits for a producer i considering exports of product j to country k:

$$\Pi_{ijk} = \sum_{t=1}^{H} \{p_{ijk} * q_{ijkt} - \hat{c}_{ijk}(1 + \text{adjustment costs}_{ijkt})q_{ijkt}\} - S_{jk} \quad (1)$$

$$p_{ijk} = \text{unit price}$$

$$q_{ijkt} = \text{quantity sold}$$

$$S_{jk} = \text{sunk costs of production}$$

$$\hat{c}_{ijk} = \frac{\tau_{jk}}{\phi_{ij}} = \text{standard unit costs}$$

p and q represent the prices and quantities for firm i in the planning horizon that covers periods $t \in [1, H]$. $\hat{c}_{ijk}$, the standard unit production cost captures τ_{jk}, the combined per-unit costs of inputs like labor and materials, which are specific to product j, and trade factors like shipping and tariffs for destination k. It also accounts for the firm's productivity ϕ_{ij}. Firms with higher productivity ϕ will therefore have lower unit costs and higher profits per unit sold. (For parsimony, the model ignores temporal discounting and simply sums profits from period 1 to H; a reasonable approximation if the planning horizon is short and the discount rate is small).

Adjustment costs may be overtime wage costs, equipment capacity modification or hiring and firing costs. I adopt the convex quadratic form proposed in Cooper and Haltiwanger (2006):[9]

$$\text{adjustment costs}_{ijkt} = \gamma_j \left[(q_{ijkt} - q^*_{ijkt})/q^*_{ijkt}\right]^2 \tag{2}$$

The γ_j term is a product-specific scaling parameter. It enables comparisons in the cross-section of destinations. For example, the cost implications of a 20% growth shock to demand are different for an auto manufacturer, compared to a maker of tee-shirts. Each sector faces different relative costs of updating production capacity. γ_j captures these differences in scale adjustment cost differences. q^*_{ijkt} is the planned production scale or capacity for firm i in period t.[10]

In the adjustment costs function, $q - q^*$ represent deviations from the production scale, or the height of the vertical lines in Figure 1. Unit production costs should depend on the proximity of actual production to the expected production scale. Several recent papers show that production scale adjustments alter marginal costs (Blum et al., 2013; Soderbery, 2013; Ahn and McQuoid, 2012). (The foregoing implicitly assumes that the cost of changing production scale from q^*_{jkt-1} to q^*_{jkt} is zero, because such changes follow the planned trajectory).[11]

[9]Costs are symmetric around q^* in equation (2); this makes the model tractable, although cost symmetry may only be a crude approximation to the data. See AppendixB.2 for a brief consideration of how asymmetric adjustment costs may affect the model.

[10]The additive form specified in equation (1) for adjustment costs ensures that unit costs will not be zero for the hypothetical destination with zero demand volatility. This structure also allows one to measure adjustment costs' effects separately from other components of unit costs, which improves on related papers that also consider changing marginal costs with market-specific shocks (Vannoorenberghe, 2012; Ahn and McQuoid, 2012; Liu, 2012).

[11]Alternatively, one could change the definition in (2):
$\text{adjustment costs}_{ijkt} = \gamma^v_j \left[\frac{(q_{ijkt} - q^*_{ijkt})}{q^*_{ijkt}}\right]^2 + \gamma^p_j \left[\frac{(q^*_{ijkt} - q^*_{ijkt-1})}{q^*_{ijkt-1}}\right]^2$
The inclusion of planned investment costs, scaled by γ^p suggests biased estimates if γ^p

Aggregate demand is stochastic, but the growth process for demand is known, given the demand history. Therefore, exporters can characterize aggregate demand Q in each destination and estimate expected profits from equations (1) and (2):

$$Q_{jkt} = Q^*_{jkt}(1 + \nu_{jkt}) \tag{3}$$

Q^*_{jkt}, period t's expected aggregate demand is estimated from the trajectory:

$$Q^*_{jkt} = Q^*_{jk0}(1 + t\hat{g}_{jk}) \tag{4}$$

The Q_{jk0} baseline and the expected growth trend $\hat{g}_{jk}$ in (4) come from historical data.[12]

Demand volatility σ^2_{jk} represents the second moment of the distribution of the growth innovations ν_{jkt} in equation (3), given that $\nu_{jkt} \sim N(0, \sigma^2_{jk})$. Assuming a normal distribution for ν helps to obtain a tractable form for expected profits shortly.[13] (I adopt the linear growth form for simplicity; a multiplicative growth model in (4) gives $Q^*_{jkt} = Q^*_{jk0}(1+\hat{g}_{jk})^t$, which approximates linear growth for small values of $\hat{g}_{jk}$).

From historical data one gets $\hat{g}_{jk}$, which characterizes Q^*_{jkt}, as well as σ^2_{jk}, which fully describes the expected shocks to demand, even if specific realizations of Q_{jkt} are not known. In the model, exporters estimate σ^2 for each product-country destination once, and do not update their estimates of demand volatility. This simplifying assumption helps to justify tests in the cross-section of destinations in Section 3. I show that this assumption is reasonable, based on tests in Section 3.2.3.

is not equal to zero. Taking γ^p as zero seems reasonable for two reasons: (1) the costs of adjusting scale upwards must be less than profits from increased scale if one is to observe more firms in larger markets in equilibrium; (2) early studies that do not consider expected demand find no statistically significant estimates for labor or capital adjustment costs (Hall, 2004).

[12] Firms can forecast q_{ijkt}, given the history of Q_{jkt}. Even with perfect information about future demand, e.g., a firm like Boeing making airplanes to order, the demand stream with greater deviations from a stable growth trajectory will incur higher adjustment costs, regular planned investments of labor and capital in each period are less costly to implement than large swings. Having information on future demand may reduce, but not eliminate those adjustment costs.

[13] Growth innovations in the data resemble a normal distribution near the mean. The $(1+\nu)$ term correspond to growth shocks in the Euler equations proposed by Carroll et al. (2011) and Hall (2004).

The next subsection derives exporters' expected demand q^*_{ijk} from the expected aggregate demand Q^*_{jk}. Firms' profits and the decision to export to jk depend on q_{ijk} and q^*_{ijk}.

2.2. Export Entry

The primary variable of interest is the number of exporters that find a destination profitable, and therefore export to that destination. In the empirics, this translates to gross export entry in the long run for trade destinations. This section explores the relationship between this variable and demand volatility.

Exporter i producing its unique variety of product j for market k can expect to sell q_{ijkt} in period t. Following conventional models of trade with CES demand preferences:

$$q_{ijkt} = \frac{p_{ijk}^{-\varepsilon}}{P_{jk}^{1-\varepsilon}} Q_{jkt} \tag{5}$$

p_{ijk} is the firm's expected price, P is the Dixit-Stiglitz aggregate price index for product j and ε is the elasticity of substitution between varieties of the product. Q_{jkt} is the aggregate demand for product j in country k. The steps that follow assume no exporter is large enough to affect the P index. As previously mentioned, Q_{jkt} is stochastic.

For each destination, $q^*_{ijk} = E(q_{ijk})$, the optimal production scale for an exporter is the expected demand.[14]

From equations (2), (3) and (5):

$$\begin{aligned}\text{adjustment costs}_{ijk} &= \gamma_j \left[\frac{\frac{p_{ijk}^{-\varepsilon}}{P_{jk}^{1-\varepsilon}}(Q_{jkt} - Q^*_{jkt})}{\frac{p_{ijk}^{-\varepsilon}}{P_{jk}^{1-\varepsilon}} Q^*_{jkt}} \right]^2 \\ &= \gamma_j (\nu_{jkt})^2 \end{aligned} \tag{6}$$

The expected profits over the planning horizon from equation (1), (with risk-

[14] The Envelope Theorem justifies this cost minimization, given the assumption that adjustment costs are symmetric in equation (2). See more on this in AppendixB.1

neutral exporters and known sunk costs S):

$$E(\Pi_{ijk}) = E\{[p_{ijk} - \frac{\tau_{jk}}{\phi_{ij}}(1 + \text{adjustment costs}_{ijkt})]q_{ijkt}\} - S_{jk}$$

Substituting equation (6) and discarding t subscripts yields:

$$\begin{aligned}
E(\Pi_{ijk}) &= (p_{ijk} - \frac{\tau_{jk}}{\phi_{ij}})E(q_{ijk}) - \frac{\tau_{jk}}{\phi_{ij}}\gamma_j E(q_{ijk}\nu_{jk}^2) - S_{jk} \\
&= (p_{ijk} - \frac{\tau_{jk}}{\phi_{ij}})q_{ijk}^* - \frac{\tau_{jk}}{\phi_{ij}}\gamma_j\{E\left[(q_{ijk} - q_{ijk}^*)(\nu_{jk})^2\right] + q_{ijk}^* E\left[(\nu_{ijk})^2\right]\} - S_{jk} \\
&= (p_{ijk} - \frac{\tau_{jk}}{\phi_{ij}})q_{ijk}^* - \frac{\tau_{jk}}{\phi_{ij}}\gamma_j q_{ijk}^*\{E\left[\frac{(q_{ijk} - q_{ijk}^*)}{q_{ijk}^*}(\nu_{jk})^2\right] + E\left[(\nu_{ijk})^2\right]\} - S_{jk} \\
&= (p_{ijk} - \frac{\tau_{jk}}{\phi_{ij}})q_{ijk}^* - \frac{\tau_{jk}}{\phi_{ij}}\gamma_j q_{ijk}^*[E(\nu_{jk}^3) + E(\nu_{ijk}^2)] - S_{jk}
\end{aligned}$$

The $E(\nu_{jk}^2)$ term is σ_{jk}^2, as defined in the notes to equation (3). The $E(\nu_{jk}^3)$ term is zero, being the third moment of a normal distribution.[15] This gives the expected profit:

$$E(\Pi_{ijk}) = [p_{ijk} - \frac{\tau_{jk}}{\phi_{ij}}(1 + \gamma_j\sigma_{jk}^2)]q_{ijk}^* - S_{jk} \tag{7}$$

As in equation (1), standard unit costs are $\frac{\tau_{jk}}{\phi_{ij}}$: the product of variable costs τ (material, capital, labor inputs, shipping and tariffs), and the firm level productivity index $\frac{1}{\phi}$. Adjustment costs reduce expected profits.[16]

Firm-level prices determine expected profits, so the next steps focus on deriving prices p_{ijk}. (Note that I abstract away from period-to-period price changes – exporter's expected price is actually a proxy for its productivity and its resulting share of Q_{jk} in the planning horizon). I also assume rational

[15]One can get (7) from the profit function for adjustment costs that are any real-valued function of the growth innovation ν. For a normal distribution with mean zero, σ^2 the second moment of ν can fully describe the terms of such a function i.e. higher order moments of ν.

[16]While the form in equation (7) is tractable enough to yield closed form solutions for firm level prices, the model's predictions will hold even if adjustment costs have a component that is fixed, or that does not scale linearly with expected production q_{ijk}^*.

risk-neutral firms that maximize expected profits:[17]

$$\frac{\mathrm{d}E(\Pi_{ijk})}{\mathrm{d}p_{ijk}} = \frac{\mathrm{d}E(\Pi_{ijk})}{\mathrm{d}q^*_{ijk}}\frac{\mathrm{d}q^*_{ijk}}{\mathrm{d}p_{ijk}} = 0 \implies \frac{\mathrm{d}E(\Pi_{ijk})}{\mathrm{d}q^*_{ijk}} = 0 \tag{8}$$

$$0 = p_{ijk} - \frac{\tau_{jk}(1+\gamma_j\sigma^2_{jk})}{\phi_{ij}} - q_{ijk}\left(\frac{1}{\varepsilon}\frac{p_{ijk}}{q^*_{ijk}}\right)$$

$$p_{ijk} = \frac{\varepsilon}{\varepsilon-1}\frac{\tau_{jk}(1+\gamma_j\sigma^2_{jk})}{\phi_{ij}} \tag{9}$$

Prices in (9) take the same form as in conventional models of trade with heterogeneous firms with one difference; unit costs include $(1+\gamma_j\sigma^2_{jk})$ to reflect adjustment costs. This is the expected price for firm i in destination jk. Firms with high productivity ϕ_{jk} will have lower prices, assuming no quality differences.

Expected export profits based on price p_{ijk}, from substituting p back into equation (7):

$$\begin{aligned}
E(\Pi_{ijk}) &= \frac{1}{\varepsilon-1}\frac{\tau_{jk}(1+\gamma_j\sigma^2_{jk})}{\phi_{ij}}q^*_{ijk} - S_{jk} \\
&= \frac{1}{\varepsilon-1}\frac{\tau_{jk}(1+\gamma_j\sigma^2_{jk})}{\phi_{ij}}\left[\frac{\varepsilon}{\varepsilon-1}\frac{\tau_{jk}(1+\gamma_j\sigma^2_{jk})}{\phi_{ij}}\right]^{-\varepsilon}\frac{Q^*_{jk}}{P^{1-\varepsilon}_{jk}} - S_{jk} \\
E(\Pi_{ijk}) &= \frac{Q^*_{jk}}{\varepsilon}\left[\frac{\varepsilon}{\varepsilon-1}\frac{\tau_{jk}(1+\gamma_j\sigma^2_{jk})}{\phi_{ij}}\frac{1}{P_{jk}}\right]^{1-\varepsilon} - S_{jk} \qquad (10)
\end{aligned}$$

Only firms above a certain productivity threshold will be profitable in destination jk. Applying the zero-profit condition to equation (10) identifies those firms. One gets the productivity threshold ϕ^*_{jk} by setting the LHS to zero in (10):

$$\phi^*_{jk} = \frac{\varepsilon}{\varepsilon-1}\frac{\tau_{jk}(1+\gamma_j\sigma^2_{jk})}{P_{jk}}\left[\frac{\varepsilon S_{jk}}{Q^*_{jk}}\right]^{\frac{1}{\varepsilon-1}} \tag{11}$$

Of the N_j firms producing j, only a fraction N_{jk} will export to destination

[17] From equation (5), $\frac{\mathrm{d}p_{ijk}}{\mathrm{d}q_{ijk}} = \frac{P^{1-\varepsilon}}{-\varepsilon p^{-1-\varepsilon}_{ijk}}\frac{1}{Q^*_{jk}} = \frac{-1}{\varepsilon}\frac{p_{ijk}}{q^*_{ijk}}$

jk. That fraction could be as low as zero if none meets the ϕ^*_{jk} threshold. Deriving the fraction N_{jk} is straightforward if one can describe the productivity of all producers of j with the distribution $G(.)$. I model N_j as an exogenous variable:[18]

$$N_{jk} = N_j(1 - G(\phi^*_{jk})) \tag{12}$$

I take $G(.)$ as the Pareto distribution.[19].

$$N_{jk} = N_j[1 - (1 - (\phi^*_{jk})^{-\theta_j})] \quad = N_j(\phi^*_{jk})^{-\theta_j} \tag{13}$$

θ_j is the Pareto shape parameter for product j.

Equation (13) shows an unambiguous relationship between σ^2 and N_{jk}, (which suggests a focus on the extensive margin of trade). ϕ^*_{jk} is a function of $\tau_{jk}(1+\gamma_j\sigma^2_{jk})$, therefore N_{jk} is a function of σ^2_{jk}. In contrast, AppendixB.3 models the relationship between demand volatility and trade volumes, which is not as pointed as the relationship in (13).[20]

---

[18] In assuming an exogenous mass of exporters, I follow others − notably, Chaney (2008) and Eaton et al. (2004). Here N_j is the number of firms making product j, e.g., the number of firms that make bicycles, regardless of export status or productivity. N_{jk} represents firms whose productivity exceeds the threshold for jk, given the assumed productivity distribution. Some producers of j will not export at all, if the lowest threshold ϕ^* of all possible markets is higher than firm productivity ϕ_{ij}.

[19] This choice follows Chaney (2008) and is consistent with the firm size distributions described in Hsieh and Ossa (2011) and Axtell (2001) Any of the general class of power law distributions should yield similar predictions, given reasonable assumptions about how the distribution is truncated.

The Pareto distribution function is $Pr(X < x) = 1 - \left(\frac{x_m}{x}\right)^{\theta}$ for $x \geq x_m$. The two parameters that characterize the distribution are x_m, the minimum productivity for a firm that produces j and θ, the shape parameter. For simplicity, I define the range of productivities on a scale [1,∞), this sets x_m equal to one, so $G(x) = Pr(X < x) = 1 - (x)^{-\theta}$.

[20] The dominance of the extensive margin is consistent with other papers that model the responses of heterogeneous firms to trade costs, e.g. Crozet and Koenig (2010) and Helpman et al. (2008). The adjustment costs associated with demand volatility increase exporters' per unit costs, just as trade costs do.

Substituting the threshold defined in equation (11) into (13):

$$N_{jk} = N_j \left\{ \frac{\varepsilon}{\varepsilon - 1} \frac{\tau_{jk}(1 + \gamma_j \sigma_{jk}^2)}{P_{jk}} \left(\frac{\varepsilon S_{jk}}{Q_{jk}^*} \right)^{\frac{1}{\varepsilon - 1}} \right\}^{-\theta_j}$$

Focusing on N_{jk} and σ^2.

$$ln(N_{jk}) = ln(N_j) - \theta_j \left[ln(1 + \gamma_j \sigma_{jk}^2) + ln\left(\frac{\varepsilon}{\varepsilon - 1} \frac{\tau_{jk}}{P_{jk}} \right) + \frac{1}{\varepsilon - 1} ln\left(\frac{\varepsilon S_{jk}}{Q_{jk}^*} \right) \right]$$

$$\frac{\mathrm{d}ln(N_{jk})}{\mathrm{d}\sigma_{jk}^2} = \frac{-\theta_j \gamma_j}{1 + \gamma_j \sigma_{jk}^2} \tag{14}$$

Plotting $ln(N_{jk})$ against σ^2 should give a line with a negative slope. The elements of the RHS term in equation (14) are all non-negative by definition: the Pareto shape parameter, θ, the adjustment cost scaling parameter γ and the demand volatility σ^2. Formally:

$$\frac{\mathrm{d}ln(N_{jk})}{\mathrm{d}\sigma_{jk}^2} < 0 \tag{15}$$

Restating equation (15):
Prediction: *Higher levels of destination demand volatility σ_{jk}^2 reduce the numbers of exporters in equilibrium.*

To restate the hypothesis, the adjustment costs associated with demand volatility reduce profitability, such that the mass of firms that find a destination profitable decreases with increases in demand volatility. If demand volatility is zero, the model reverts to the conventional model of trade. One way to take this prediction to the data is a linear regression of N_{jk} on σ^2; the sign of the coefficient on demand volatility should be negative.

Lemma: *Holding other factors equal, the minimum productivity of firms in destinations with high demand volatility is higher.*

From equation (11), it is clear that the productivity threshold ϕ_{jk}^* for entering a destination increases with demand volatility, therefore one expects the minimum level of other proxies for productivity like exporters' share of a product's exports to increase with demand volatility, all other things being equal:

$$\phi_{jk}^* = \frac{\varepsilon}{\varepsilon - 1} \frac{\tau_{jk}(1 + \gamma_j \sigma_{jk}^2)}{P_{jk}} \left[\frac{\varepsilon S_{jk}}{Q_{jk}^*} \right]^{\frac{1}{\varepsilon - 1}}$$

$$\frac{\mathrm{d}\phi_{jk}^*}{\mathrm{d}\sigma_{jk}^2} > 0 \tag{16}$$

Corollary: *The rate of decline in exporter numbers with demand volatility is not constant, but varies nonlinearly with gamma γ.*

The cost implications of demand volatility are non-linear, so are its impacts on an exporters' profit. Destination choice by extension is nonlinear with respect to demand volatility. The non-linearity will depend on γ_j. For very small γ, such that low labor-firing costs for example make adjustments relatively costless, the slope of $log(N_{jk})$ with respect to σ_{jk}^2 will always be approximately linear and proportional to θ_j, the productivity distribution's shape parameter. For large γ_j, the slope is increasingly non-linear, and decreases in absolute terms with σ^2. (If the γ term was a function of σ^2 because adjustment costs are not linear proportions of ν^2 as proposed in equation (2), the slope could increase or decrease with σ^2).

Testing the relationship between N_{jk} and the square of the demand volatility term directly addresses this corollary in Section 3.2.4.

3. Empirics

This section examines the relationship between exporter numbers and demand volatility. First, I describe key variables and data sources. Regression estimates follow the definitions: a baseline specification and variations that further test the model's predictions. I address the most important alternative explanations with robustness checks at the end of the section.

3.1. Data and Definitions

The key variables come from two trade datasets: firm level Chinese export data describes exporters' destination choices, and UN COMTRADE data on global imports by product and country describes the history of each destination in terms of size and demand volatility.

The firm level export data captures exporter numbers in each destination, derived from the universe of Chinese export transactions between years 2000 and 2006. From this dataset, which identifies firms, the year of each transaction, the product and the country to which it was shipped. Products are defined at the HS8 level, which I tally up to 4,903 HS6 categories – the narrowest global standard for defining traded products.

The UN COMTRADE data include global imports of each narrowly defined HS6 product category for all countries between 1995 and 2010. Annual imports up to 2005 were collapsed to product-country-year observations – imports of bicycles (HS871200) into Kenya in the year 2000 from all countries would be one observation, for example. (Restricting historical data to 2005 and earlier is motivated by the fact that firm level export data stop at 2006). I estimate demand volatilities using this global import demand history for each destination.[21].

The combined datasets represent the destination choices of more than 243,000 Chinese exporters, mostly in the period of export expansion that followed China's entry into the WTO in 2001. Exporters in the data cover more than 390,000 of the roughly one million possible product-country combinations that define destinations. (These are imports into one of more than 200 countries in any of the more than 5000 possible narrowly defined HS6 product categories). AppendixA.1 describes these data sources further and outlines how I merge the two. Information on GDP, distance and other predictors of trade come from the CEPII gravity dataset (Head et al., 2010).

The next sub-sections describe the key variables in the empirical specification: the dependent variable, which is a multi-year count of unique exporters and the key independent variables, derived from the historical demand data described in the preceding paragraphs.

3.1.1. Dependent Variable: Multi-Year Unique Exporter Counts

The paper uses two definitions of exporter numbers, N_{jk}. The primary measure represents the number of Chinese firms that exported good j to a

[21]The COMTRADE database of global trade was compiled and cleaned up by Gaulier and Zignago (2010) and released to the public through the Centres d'Études Prospectives et d'Information Internationales (CEPII). I will refer to this database as COMTRADE from here. See www.cepii.fr/anglaisgraph/bdd/baci.htm.

destination jk between 2000 and 2006, counting each exporter only once in the entire period. This definition reflects the equilibrium number of exporters entering a potentially profitable destination. (The measure extends over multiple years because theory provides no clear guidance on how long it takes to reach equilibrium).[22] Furthermore, exports drop to zero and rebound in the following year for many destinations. Defining exporter numbers on an annual basis may misrepresent markets that simply run on a multi-year demand cycle. I use logged values of this unique exporter count for the regressions.[23] For parsimony and consistency with the definition in the next paragraph, I call this measure gross entry.

The second measure I use is *gross entry* from 2001 to 2006. This is the log of the difference between [1] N_{jk} the number of unique exporters over all years and [2] the number of unique exporters observed in 2000, the first year in the data, which happens to be the year before China joined the WTO. China's WTO entry in 2001 marked a new regime of low trade costs and easier export access, evidenced by an increase in aggregate exporter numbers from around 62,000 in 2000 to more than 170,000 in 2006. By its construction, this measure of gross entry in the new trade regime controls for undocumented destination attributes that may influence export entry. For example, destinations with lower entry costs are more likely to have exporters in 2000, and to have more exporters.

This alternative measure could be more relevant, as export restrictions that may have distorted observed exporter counts were removed in the course of WTO accession. For example, before WTO accession, trading license requirements barred Chinese firms below a certain size from exporting directly (Ahn et al., 2011). In sum, the number of new exporters after 2000 was large, with ample variation across destinations, which helps the analysis.

[22] Later sections of the paper include tests with annual counts, which provide results comparable to typical estimates using annual data. Defining exporter counts as gross entry is a better fit to the model because the model remains silent about exits from destination after entry. Table A.13 relaxes this connection between the model and the empirics, using annual exporter counts as the dependent variable, ignoring how that measure includes gross entry and exit.

[23] For convenience, I call N_{jk} gross entry in the tables that follow. The data show that exporters in 2000 represent only a quarter of the full set of observed unique exporters. Fortunately, China's accession to the WTO in 2001 suggests an alternative measure of gross entry in the era of liberalization, which I discuss in the next paragraph.

3.1.2. Key Variable: Demand Volatility

I measure demand volatility as the sum of the squared deviations of demand from a linear trend over the years 1995 to 2005. The trend is calculated using the reported dollar value for each destination's annual demand – total imports from all countries in a given HS6 category. As Chinese exports represent less than 13% of the global total in this period, this measure mitigates concerns about reverse causality that may have resulted from defining volatility with only Chinese export data. (The value-based demand volatility measure ensures comparability across destinations, given how quantity measures differ by product. I also include estimates that use demand volatility calculated from quantities in the robustness checks).

The growth trend estimation uses a linear regression, with each destination scaled by total demand over all years to yield a scale-free measure for cross-sectional comparisons. To get σ^2_{jk}, the volatility term, I estimate the trend and intercept for each destination jk:[24] For this exercise, the definition of destinations in the empirics is consistent with the model. Each Q_{jk} in the model corresponds to a specific product j and a country k, just as destinations in the data are defined as the combination of a narrow HS6 product and a country.

Formally, I run the following regression:[25]

$$\frac{Q_{jkt}}{\sum_t Q_{jkt}} = \zeta_{jk}t + \alpha_{jk} + \epsilon_{jkt} \tag{17}$$

Using the residuals, I derive estimates for σ^2:

$$\hat{\sigma}^2_{jk} = \sum_t (\epsilon_{jkt})^2 \tag{18}$$

The incidental assumption in this setup is that exporters form expecta-

[24]To avoid the bias that may result from imposing the same non-linear form on all demand trajectories, I opt for a linear regression of historical trade levels on time. This measure directly interprets the model's definition of demand volatility. Some measurement error is expected, given that I use only 11 years of demand history. However, using longer demand histories comes with the risk of including irrelevant information – exporters may completely discount information on demand from the too-distant past.

[25]I repeat the baseline regressions using a quantity-based volatility measure and report the results in the robustness checks section 3.3.2.

tions of future volatility based on observed aggregate volatility. As described in the previous section, each exporter's expected revenue from a destination is proportional to the aggregate demand from that destination, therefore volatility at the product-country level is a reasonable measure for evaluating the profitability of each destination at the firm level.

After the main tests in the next subsection, I also use the standard deviation of year-on-year growth as an alternative definition. This alternate definition is also broadly consistent with the model, and with other papers in the literature.

Nonetheless, the measure of demand volatility has the advantage of addressing the two main challenges to measuring volatility for time series: (1) making the measure of demand volatility independent of the size of each series and (2) separating baseline growth from volatility. I control for size by scaling all series by the total value over all periods, and control for growth by introducing the linear trend that best fits the data. Controlling for size ensures that the volatility measure is comparable across series with different initial levels. For example, using the standard deviation of historical values to compare the volatility of US aggregate imports with Rwandan imports would lead to the flawed conclusion that US imports are more volatile, simply because the absolute values are larger. (One could try to fix this by using the coefficient of variation - i.e. dividing by the mean, but that still leaves concerns about how the measure accounts for growth trends).

Furthermore, defining demand volatility as deviations around a trend avoids mis-measurement when the data include instances of zero demand. The common measure of volatility as the standard deviation suffers from the problem of measuring growth from or to zero. If one uses the mid-point growth measure of Davis and Haltiwanger (1992), growth at these instances of zero will fall at the extreme values of -2 and 2. While those values are usable, they may represent outliers that bias the volatility measure, especially if most growth observations are clustered near zero. Measuring demand volatility as deviations from a trend simply avoids this concern by used scaled levels, rather than transforming those levels into a growth index before calculating the volatility of a series.

3.1.3. Key Variable: Destination Size

In estimating the effects of demand volatility, I use aggregate historical demand as an indicator of product-country destination size. Conventional estimates of international trade measure size as GDP. This would be appropriate for a model of trade where the countries define market boundaries and firms' narrow product specializations were not relevant to competition. Using aggregate historical demand for each destination as a measure of its size fits the structure provided by the previous section, one that emphasizes competition within narrow product categories and exporters' choices of markets on that basis.

I define the terms as the logged sum of aggregate demand in each destination between 1995 and 2005. (I use the first 11 years of aggregate data available for the same reasons that I use those years to measure demand volatility). In principle, this logged sum represents the projected future demand for a destination. One only needs to assume that destination growth g_{jk} is consistent to use this history as a proxy for the projected aggregate demand over the exporter's planning horizon, (Q^* in the model). Formally, the projected size of a destination with average historical growth rate g_{jk} is $log(\sum_t Q_{jkt}) \simeq log(Q_{jk0}) + log[\sum_t (1 + g_{jk})^t]$. As long as past growth rates are a reasonable proxy for expected growth rates, this measure allows me to control for destination size in terms of its present value and growth.

This measure offers a finer level of control for testing export destination choice than a country-level measure like GDP. The regressions in this section will show that it explains more of the variation in exporter numbers than conventional variables like GDP and distance. (This is in part because; historical demand is explained by GDP and distance, so that the inclusion of current GDP in an estimation exercise that includes historical demand provides little additional information). I run versions of the regressions that follow this section without this market size variable and obtain results that are similar in sign, but with larger coefficients. This is unsurprising, given the correlation between market size and volatility (see Figure A.5). To avoid the implied omitted variable bias, the next section reports only regressions that include this destination size variable.

3.2. Results

3.2.1. Demand Volatility and Exporter Counts

Table 1 summarizes the key variables.

About 40 unique exporters served the average destination between 2000 and 2006; with 35 of these being the firms that entered the destination after the year 2000. This number is highly skewed; both variables have a median value of 5. The variation in exporter counts is large; products like

Table 1: Summary of Key Variables

Variable	Mean	Std. Dev.	Min.	Max.	N
Gross Entry	39.95	190.77	1	15643	397547
Gross Entry post-2000	35.38	168.38	0	13497	397547
Log(Gross Entry)	1.96	1.66	0	9.66	397547
Log(Entry post-2000)	1.91	1.63	0	9.51	387916
Destination Size	8.65	2.63	0	20.6	397547
Demand Volatility	0.053	0.086	0	0.91	380372

Chinese exporters sent goods to 397,547 out of a possible 992,302 destinations between 2000 and 2006. Only 380,372 had the two or more non-zero observations required to compute demand volatility. 9,170 had no new exporters after 2000. Destination size is the log of total historical demand in the COMTRADE data.
Number of countries (205); products (4,902)

buttons naturally had many producers, while airplanes had few. Country-specific variations also existed; large ones like the US had more exporters. However, countries and products alone leave much of the variation in the data unexplained. Unreported regressions of exporter numbers at the level of product-country destinations on product and country fixed effects alone yield R^2 values of 0.20 and 0.01 respectively.

Given that more than half of the possible destinations had zero Chinese exporters, the first order of inquiry is whether demand volatility is higher on average for those destinations that had no Chinese exporters. Figure A.5 in the appendix suggests that the destinations served by Chinese exporters tend to be larger and have lower demand volatility, but the visual comparisons of distributions in that figure is only suggestive at best. In the summary statistics, destinations with zero Chinese exporters had an average demand volatility of 0.1098, while those matched to the firm-level export data reported an average of 0.053, the former category's average demand volatility is 105% higher.

Table 2 presents the conditional expectation. Destinations with higher demand volatility are on average more likely to report zero imports from Chinese exporters. The specification adopted is a linear probability model with a dependent variable that is 1 if no Chinese firm exported to the destination between 2000 and 2006; it is 0 otherwise. Product fixed effects address the fact that some items are more likely to be exported than others for time-invariant reasons outside the model, and country fixed effects or variables like GDP and distance address the fact that country-level factors also determine the prevalence of zeros in trade.

Table 2: Demand Volatility and Incidence of Zero Chinese Exporters
(Dependent Variable: **1**[Number of Exporters in Destinations = 0])

VARIABLES	(1)	(2)	(3)	(4)
Demand Volatility	0.391***	0.277***	0.335***	0.259***
	(0.007)	(0.007)	(0.006)	(0.006)
Destination Size		-0.045***		-0.040***
		(0.001)		(0.001)
Log(GDP)	-0.086***	-0.054***		
	(0.001)	(0.001)		
Log(GDP per capita)	0.018***	0.022***		
	(0.001)	(0.001)		
Log(Distance)	0.103***	0.098***		
	(0.002)	(0.002)		
Constant	-0.390***	-0.561***		
	(0.023)	(0.023)		
Observations	573,997	573,997	708,802	708,802
R-squared	0.465	0.480	0.515	0.525
Country FE			Y	Y
Product FE	Y	Y	Y	Y

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

The dependent variable is a dummy equal to 1 for destinations with no Chinese exporter. The units of observation are destinations: unique HS6-product and country combinations.

Increasing demand volatility by one standard deviation corresponds to a 2.8% decrease in the likelihood that a destination is served by any Chinese exporter, after controlling for destination size and country features. The standard deviation of demand volatility is 0.11 for this set of destinations. I control for market size in columns 2 and 4 to address the concern that larger destinations will generally have more exporters, as market size is correlated with demand volatility. Table 2 suggests Poisson estimates for the effect of demand volatility on exporter numbers, which are in Table 9.[26]

Figure 2 shows that more Chinese exporters serve destinations with low demand volatility, on average. The plot sets the logged number of exporters that served destinations between 2000 and 2006 against demand volatility. The predicted averages in the plot control for market size, country features and product-fixed effects. Each average is calculated separately for 50 equal-frequency bins of demand volatility.[27]

Figure 2: Average Chinese Exporter Counts and Demand Volatility

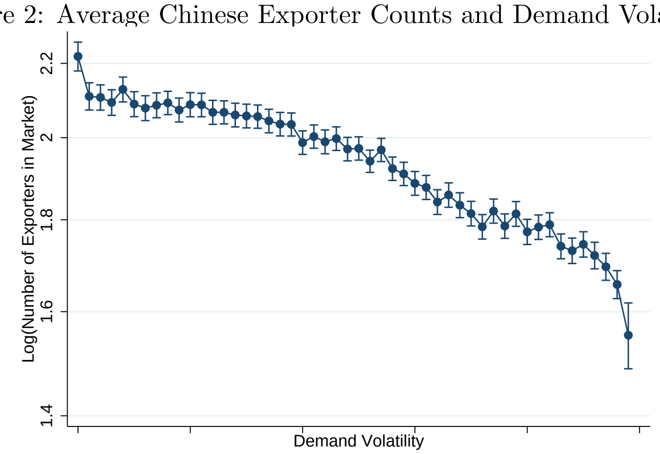

Predicted unique exporter counts over 2000 - 2006. Estimated means and 95% confidence intervals after controlling for market size, country factors and HS6 product fixed effects. Observations are destinations grouped into 50 Quantiles, least to largest by demand volatility. Data Sources: China GAC Export Data (2000-2006), COMTRADE

[26] I calculated demand volatility for 708,802 destinations. The balance of destinations were not usable because the demand volatility variable is not defined for destinations with only one or two years of imports.

[27] To interpret the graph, note that the average destination's logged exporter numbers is 1.96 (with a standard deviation of 1.66). The curvature of the graph is addressed later in this section.

I estimate the following baseline specification:

$$\begin{aligned} log(N_{jk}) &= \beta_0 \sigma_{jk}^2 + \beta_1 X_{jk} + \alpha_j + \alpha_k + \epsilon_{jk} && (19) \\ X_{jk} &= \text{a vector of gravity model variables e.g., GDP, distance} \\ \alpha_j &= \text{product fixed effects} \\ \alpha_k &= \text{country fixed effects} \end{aligned}$$

To identify the effects of demand volatility on export destination choice, this reduced form specification plays on differences between destinations and the fact that certain products or countries tend to have higher volatility. It is necessary to control for product-specific factors; the number of potential entrants and the sunk costs of entry vary significantly along this dimension. For example, between narrowly defined HS6 product categories, the number of exporting firms ranges from 1 for nuclear reactor fuel cartridges (HS 840130) to more than 40,000 for miscellaneous plastic articles (HS 392690). Estimates of the Pareto Distribution parameter also ranged from less than 5 to greater than 15, with varying degrees of fit for these product categories. Applying product fixed effects in the cross-section helps to address these differences.

Differences in export entry by country are expected, given factors like GDP, distance, language and currency. Furthermore, differences in volatility by country exist, as shown in Figure A.6. (Papers like di Giovanni and Levchenko (2012a); Koren and Tenreyro (2007) provide similar evidence). To ensure differences in exporter numbers due to these factors are not conflated with demand volatility at the product-country level, I introduce either country fixed effects or direct measures of these variables (for the year 2006).

The specification with product fixed effects can simply be described as a comparison for a product like bicycles, using countries like Portugal and Greece that have similar GDP, GDP per Capita and distance from China, if bicycle imports into these countries differ in terms of demand volatility. The way the data is set up makes it possible to identify which country has the higher level of demand volatility for bicycles, knowing that the similar comparisons for other products are not guaranteed to be identical. Table 3 presents the results.

Fewer exporters enter destinations with high demand volatility, after controlling for the common predictors of exporter numbers. The observed number of exporters is about 5.1% lower on average for destinations one standard deviation above the mean. The estimated effect in Table 7 is 11% when demand volatility is measured with quantities , not dollar values. (The lower estimated effect is not surprising – measuring volatility with dollar values will capture demand shocks as well as the mitigating price-changes that come with them). Columns 1 and 2 show the number of unique exporters observed between 2000 and 2006 as the dependent variable; the other columns use gross entry after 2000, the first year in the data.

The results in Table 3 translate to about 2 fewer exporters in the average destination with each increase in demand volatility by one standard deviation. (The response is calculated as $\{40 * [1 - exp(-0.612 * 0.086)]\}$). The next paragraphs describe the findings further. They also show how I identify the effects of demand volatility. Columns 3 and 4 use gross entry as the dependent variable, i.e. the number of new exporters in a destination after 2000. The columns have fewer observations because the log transform excludes destinations with no new exporters after 2000.[28] As previously discussed, gross entry in columns 3 and 4 measures how many exporters found a destination potentially profitable in the less restrictive trade regime that started in 2001, the year of China's WTO accession.

The last two columns of Table 3 agree in sign and significance with the first two, although the estimated coefficient on demand volatility decreases. (Note that the gross entry post 2000 variable is by definition always less than gross entry, and several destinations report zero entry after 2000). Destination size takes away most of the statistical significance associated with gravity variables like GDP and distance. The variable, which I measure as the logged sum of imports between 1995 and 2005, represents both observed and projected demand growth, like Q^*_{jk} in the model. By its definition, it also addresses concerns that historical average growth rates affect exporter numbers.[29] GDP may not adequately represent the size of specific destinations,

[28] Columns 1 and 3 also have fewer observations due to missing GDP, distance or other control variables from the CEPII gravity dataset.

[29] A possible challenge to the definition of this variable is that total absorption each destination includes imports and the domestic production. That poses no real problem for this paper, the fact that imports and domestic production are generally close substitutes within the narrow product categories suggests that imports can be used as a proxy for aggregate demand. Regressing destination demand on annual lagged demand gives an R^2 of 0.93, indicating that imports patterns are persistent or autocorrelated, even in the absence of data on domestic production.

Table 3: Exporter Counts and Demand Volatility:
(Dependent Variable: Log Number of Exporters in Destinations)

	(1)	(2)	(3)	(4)
VARIABLES	Log(Gross Export Entry)		Log(Gross Entry Post-2000)	
Demand Volatility	-0.612***	-0.613***	-0.584***	-0.582***
	(0.035)	(0.032)	(0.035)	(0.032)
Destination Size	0.421***	0.421***	0.412***	0.412***
	(0.003)	(0.003)	(0.003)	(0.003)
Log(GDP)	0.001		0.001	
	(0.001)		(0.001)	
Log(GDP per capita)	0.001		0.001	
	(0.002)		(0.002)	
Log(Distance)	0.002		0.001	
	(0.006)		(0.006)	
Constant	-1.645***		-1.619***	
	(0.079)		(0.079)	
Observations	272,926	371,531	266,915	363,381
R-squared	0.547	0.545	0.546	0.544
Country FE		Y		Y
Product FE	Y	Y	Y	Y

Robust standard errors in parentheses. Errors clustered by HS6 products.
*** p<0.01, ** p<0.05, * p<0.1

The units of observation are destinations: unique HS6-product and country combinations, e.g., Ethiopian imports of women's cotton overcoats (HS 610220). Gross entry is the log of unique firms with recorded exports to a destination between 2000 and 2006. The change in exporter count captures the difference between all firms that served a destination and firms that served the destination in 2000. This difference measures new exporters that appeared with China's trade liberalization from 2001 onwards. Control variables used but not shown in the table include geographic remoteness and dummies for shared borders, common languages and WTO membership. Missing observations in Columns 1 and 3 are because GDP data is not always available. However, estimates on the largest common sample are almost identical to columns 2 and 4.

even if it represents market size well in country-level gravity models. I do not show colonial relationships, WTO membership and other gravity variables in the table to conserve space. The gravity model variables all come from the year 2006.

The demand volatility measure varies in the cross-section across destinations, but not over time. (This also matches the model and baseline empirical specification, which described each destination destination with one demand volatility parameter and one count of exporters). To justify the implicit assumption that exporter estimates are not updated over time, I provide two tests. First, in Section 3.2.3 I repeat the main regression using demand volatilities defined over various periods and show that the predicted outcomes remain largely unchanged. Second, Section 3.3.1 shows that changing the weights allocated to the deviation from trend in calculating demand volatility does not substantively alter the findings of the paper. These tests suggest that demand volatility is itself a stable feature of destinations, or that exporters do not substantially update their perceptions of destinations.

In identifying the effect of demand volatility, note that the demand data is a global aggregate, which mitigates concerns about reverse causation, as described in the variable definition. (I go further in AppendixA.5.2 and use global aggregates that exclude Chinese exports to show that the findings are robust). As mentioned earlier, product fixed effects capture differences in the γ adjustment parameter, the mass and distribution of exporters N_j and θ_j, as well as the setup costs and fixed costs associated with specific products. Country fixed effects also control for factors that include exchange rates, exchange rate volatility, country size, trade costs and policies like tariffs and trade agreements. Testing in the cross-section helps to avoid concerns about other time-varying factors, as long as the variables I use are stable over the period under review. Column 1 allows the GDP, GDP per capita and other gravity variables to explain country-specific determinants of trade costs. The gravity model variables are either constant, like distance, or highly auto-correlated.[30] Columns 2 and 4 apply both country and product fixed effects simultaneously.[31] The country fixed effects absorb the gravity model

[30]Many small economies were missing GDP and GDP per capita, hence the differences in the number of observations between the even and odd-numbered columns

[31]For computational efficiency, I follow the algorithm proposed by Guimarães and Portugal (2009) for multiple high-dimensional fixed effects. This algorithm iteratively estimates the coefficients, unlike conventional OLS estimation that directly calculates the matrix

variables, as expected.

The foregoing shows that Chinese exporters entered destinations with lower demand volatility in greater numbers. This is after accounting for product and country characteristics that recognize the potential costs and profits from exporting. The results hold whether the dependent variable is a count of unique exporters or gross entry – the increase in unique exporter counts between 2001 and 2006. As I do not discuss how diversification or choosing multiple markets may mitigate firm-level volatility, the estimated effects are on the conservative side.[32]

In describing whether trade levels fall with increasing demand volatility, one must consider how much of the predicted effect on trade volumes lies on the extensive margin – the number of exporters as shown above, or the intensive margin - exports per exporter. The model predicts that more of demand volatility's effects are observed in exporter numbers – the extensive margin. The higher expected prices associated with demand volatility imply lower expected demand and profits, given non-zero sunk costs. In a world with heterogeneous firms, those with lower productivity will generally self-select out of destinations with high demand volatility.

Table 4 presents the results. Higher demand volatility is associated with fewer exporters, (exporter numbers shows the largest coefficients and the highest level of explained variation in the table). Total exports from China summed across all years is lower for destinations with high demand volatility (columns 1 and 2); with fewer exporters as predicted, exports per exporter increase (columns 5 and 6). The coefficients in columns 3 and 5 sum to column 1, as the regressions are linear in logs. Destination size explains much of the variation in exporter counts in this table, just as in Table 3. Section AppendixB.3 develops the model to show a relationship between equilibrium trade levels and demand volatility; the relationships established in the model

inverse. The coefficients represent a vector for the selected fixed effects and independent variables that yield the least squared error, within a 1e-6 tolerance. The effects are not fully interacted, as that would eliminate all degrees of freedom in the data.

[32]Most exporters serve more than one destination – and destinations are not perfectly correlated – entering two destinations simultaneously generally yields a firm-level portfolio volatility that is lower than the demand volatility of either destination. (The upper bound of the portfolio volatility being the higher of the two markets). In sum, the reported effects of demand volatility on exporter numbers are muted. When $\beta_0 = [log(N_{jk}) - (\beta_1 X_{jk} + \alpha_j + \alpha_k)]/\sigma_{jk}$ in equation (19); if the true volatility perceived by exporters $\sigma^*_{jk} \leq \sigma_{jk}$, then the true coefficient $|\beta^*_0| \geq |\beta_0|$

for both variables are found in this table.

In sum, trade levels are lower for destinations with high demand volatility, and most of the effect comes from the extensive margin, represented by columns 3 and 4. Columns 1 and 2 of the table are consistent with the prediction in equation (B.7). The imperfect matching of countries between the trade and CEPII gravity data set means that GDP and GDP per capita are missing for many observations. (AppendixA.1 describes the matching). Unreported regressions on the even-numbered columns give nearly identical coefficients on a sample restricted to those with no missing variables in the odd-numbered columns.

To assuage concerns that counts of unique exporters over multiple years may not be comparable to conventional estimates of gravity models with annual export measures, AppendixA.4 presents regressions that include estimates with annual exporter counts, annual control variables and country-year fixed effects. These show that the reported estimates are not due to periodic shocks, or exchange rate volatility – other potential drivers of trade in the literature. The first two columns of Table 4 represent firm level gravity model regressions, (as do the first two of Table A.13). The two tables show that fewer exporters enter destinations with high demand volatility.

Finding a consistent pattern of lower exporter counts with demand volatility suggests that prices will be higher in those destinations. However, reliable tests of demand volatility's effect on prices are difficult with no data on quality, given the well-documented link between prices and quality (Hallak and Schott, 2011; Hallak and Sivadasan, 2009). Regressions of price on demand volatility yield statistically insignificant coefficients for the largest product categories. This was after including various sets of controls that included gravity model variables, product-year fixed effects, country year fixed effects and firm fixed effects. (See AppendixA.8.)

Demand volatility compares favorably with conventional variables like GDP and geographic distance in predicting trade. I run separate regressions (reported in Table A.16) of exporter counts on demand volatility, GDP, geographic distance and destination size, in the absence of additional controls. The omitted-variable regressions involve only the dependent variable, the selected variable and product fixed effects. The respective adjusted R^2 values are 0.29, 0.22, 0.22 and 0.54. Only destination size explains more exporter count variation on this crude test.

Table 4: Exports and Exporter Counts vs. Demand Volatility
(Dependent Variable: Log Export Measure in Destinations)

VARIABLES	(1) Log(Exports)	(2)	(3) Log(Exporters)	(4)	(5) Log(Exports per Exporter)	(6)
Demand Volatility	-0.219***	-0.189***	-0.612***	-0.613***	0.392***	0.424***
	(0.075)	(0.067)	(0.035)	(0.032)	(0.057)	(0.050)
Destination Size	0.834***	0.836***	0.421***	0.421***	0.413***	0.415***
	(0.005)	(0.005)	(0.003)	(0.003)	(0.003)	(0.003)
Log(GDP)	0.001		0.001		-0.001	
	(0.003)		(0.001)		(0.002)	
Log(GDP per capita)	0.003		0.001		0.002	
	(0.004)		(0.002)		(0.003)	
Log(Distance)	0.008		0.002		0.006	
	(0.011)		(0.006)		(0.008)	
Constant	4.176***		-1.645***		5.821***	
	(0.159)		(0.079)		(0.113)	
Observations	272,926	371,531	272,926	371,531	272,926	371,531
R-squared	0.490	0.488	0.547	0.545	0.390	0.387
Country FE		Y		Y		Y
Product FE	Y	Y	Y	Y	Y	Y

Robust standard errors in parentheses. Errors clustered by HS6 products.
*** p<0.01, ** p<0.05, * p<0.1

The units of observation are destinations: unique HS6-product and country combinations, e.g., Ethiopian imports of women's cotton overcoats (HS 610220). Exports in columns 1-2 are for all Chinese exporters. Exporters in columns 3-4 is the log of unique firms with recorded exports to a destination between 2000 and 2006. It is the same number used for columns 1-2 of Table 3. Destination size is the log of total demand from a destination between 1995 and 2005. Control variables used but not shown in the table include geographic remoteness and dummies for shared borders, common languages and WTO membership.

3.2.2. Demand Volatility and Exporter Productivity Thresholds

Figure 3 supports the model's claim that demand volatility filters out producers with low productivity. The high costs of adjustment mean that in general, fewer firms will have the profit margins required to succeed in destinations with high volatility. (As the data allows no direct measures of productivity, I use producers' market shares within product categories as a proxy).

Figure 3: Minimum Export Thresholds and Demand Volatility

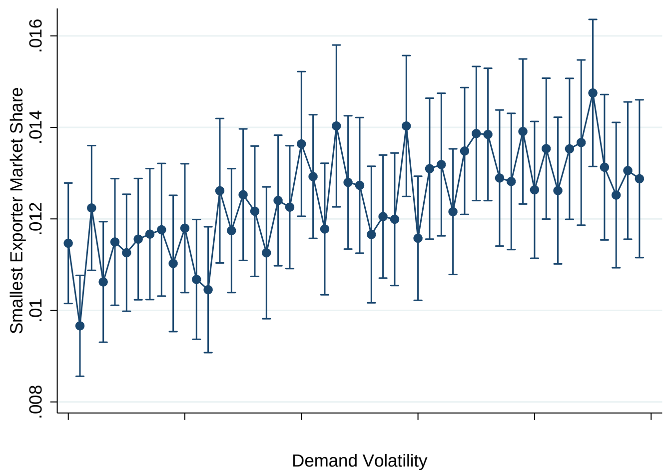

The minimum market share of the exporters in each destination jk. I use firms' share of exports within a narrow HS6 product category as a proxy for productivity, so that destinations served by firms with the least market share get the lowest value of this proxy for ϕ^*_{jk}. The y-axis averages these minimum market share measure for all destinations in each quantile. Destinations Grouped Into 50 Quantiles, Least to Largest by Demand Volatility. Data Sources: China GAC Export Data (2000-2006), COMTRADE

To facilitate comparisons in the cross-section, I compute each exporter's share of Chinese exports in 2006 within its HS6 category, and identify the minimum share held by a firm serving each destination:
$Share_{ij} = \frac{q_{ij}}{\sum_i q_{ij}} = \frac{\sum_k q_{ijk}}{\sum_i \sum_k q_{ijk}}$. For example, among firms making product j, say bicycles (HSHS871200), each Chinese exporter of the good is assigned a value $Share_{ij} \in (0, 1]$. That value represents its share by dollar value of all Chinese exports of bicycles. That share is firm-specific regardless of country destination.

I represent the productivity threshold ϕ^*_{jk} by the smallest market share recorded by any firm in destination jk. That is, for each destination, ϕ^*_{jk} is represented by $min(Share_{ij})$. In the bicycle example from the previous paragraph, if all Chinese exporters of bicycles serve the Tanzanian market for imported bicycles, then the ϕ^*_{jk} measure is the minimum possible firm share for bicycles. Similarly, destinations served by only the largest exporter in the product category will report a higher threshold than the destination served by both the largest and smallest exporter, (if more than one firm exports the product from China).

I use simple regression methods to check whether this measure of productivity thresholds increases with demand volatility. The regressions mimic equation (19), replacing the demand volatility measure with a dummy for each of 50 equal-frequency bins for destinations, ranked by demand volatility. Product fixed effects control for differences in the distributions of market shares by product. AppendixB.4 includes additional plots that use other proxies for firm-level productivity. That section also includes results of regression exercises that support the finding of higher productivity thresholds with increasing demand volatility.

The plot shows an increasing trend in the predicted size-rank of exporters with increasing demand volatility; the standard errors suggest that the trend is statistically significant. In other words, the destinations with the highest demand volatility have on average, exporters that are among the largest producers for the related product. This is after controlling for destination size, and variables like GDP and Distance. This supports the mechanism proposed in the model that expectations of adjustment costs drive more low-productivity firms away from the most volatile destinations.

3.2.3. Temporal Variations in Demand Volatility

The previous test steps assume that expected demand volatility does not change over time from the exporter's perspective. As this assumption is important to how I simplify the estimation, I test it using the specification in (20):

$$log(N_{jk}) = \beta_0 \sigma_{jk}^{t_1-t_2} + \beta_1 X_{jk} + \alpha_j + \alpha_k + \epsilon_{jk} \tag{20}$$

$$\sigma_{jk}^{t_1-t_2} = \text{Demand volatility for years } t_1 \text{ to } t_2$$

$$\text{the other terms mirror the definitions in equation (19)}$$

The variants of $\sigma_{jk}^{t_1-t_2}$ use (1995,2000), (2000,2005) and (2003,2008) as

$(t_1,\ t_2)$ pairs. These shorter demand histories yield less precise estimates. The measure for (1995,2000) gives more weight to historical demand before exporters in the data made their destination choices. The two additional 6-year history samples draw on years that put more weight on in-sample and after-the-sample data, i.e. 2000-2005 and 2003-2008. (I stop at 2008 because the severe drop in trade across many products for 2009 is exceptional).

Table 5 shows that the assumption of stable expected demand volatility is not far-fetched. The coefficient of the 1995-2000 measure is only one standard deviation away from the 2003-2008 measure. The standard errors and explained variations are similar across all three measures of demand volatility, whether the measure is weighted toward the past, or toward the future. While the size of demand volatility's coefficient is higher for the mid-data sample, the sign and significance remain unchanged.

Table 5: Exporter Numbers with Past and Future Demand Volatility
(Dependent Variable: Log Number of Exporters in Destinations)

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	Log(Gross Export Entry)			Log(Gross Entry Post-2000)		
Dem. Volat.(95-00)	-0.182***			-0.136***		
	(0.029)			(0.029)		
Dem. Volat.(00-05)		-0.493***			-0.483***	
		(0.030)			(0.031)	
Dem. Volat.(03-08)			-0.220***			-0.235***
			(0.033)			(0.033)
Destination Size	0.435***	0.424***	0.430***	0.426***	0.415***	0.420***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Observations	356,899	369,240	368,012	349,331	361,372	360,395
R-squared	0.545	0.544	0.544	0.545	0.543	0.543
Country-Year FE	Y	Y	Y	Y	Y	Y
Product FE	Y	Y	Y	Y	Y	Y

Robust standard errors in parentheses. Errors clustered by HS6 products.
*** p<0.01, ** p<0.05, * p<0.1

The units of observation are destinations: unique HS6-product and country combinations, e.g., Ethiopian imports of women's cotton overcoats (HS 610220). Exporter counts represent the log of the number of unique firms with recorded exports to a destination between 2000 and 2006. Gross entry post-2000 captures the difference between all firms that served a product market and firms that served the destination in 2000.

Compared to the predicted 5.1% change in exporter numbers from Table 3, the predicted declines in exporter numbers using only 1995-2000 to measure demand volatility drops to 1.6% if only years 1995-2000 were used to form expectations. The respective predictions for the other sample periods are 4.1 and 1.9%. Correlations between the measures of demand volatility are 0.37 (95-00, 00-05), 0.44 (95-00,03-08) and 0.59 (00-05,03-08).

In sum, the paper's main qualitative findings are robust to the period used for measuring demand volatility. Had the model and estimates explicitly considered exporters updating their estimates of destinations' demand volatility, the findings are expected to remain consistent. (Tables 3 and 5 broadly agree in terms of the sign and significance of the estimated coefficients of demand volatility). Section 3.3.1 provides further evidence – using the demand history from 1995-2005, but applying different weights to each year.

3.2.4. Non-Linear Responses to Demand Volatility

The non-linear form of σ^2 in equation (14) suggests a test of the dependent variable on higher orders of demand volatility, as well as its interactions with market size, among other variables.

Table 6 includes the square of demand volatility as a regressor, as well as interactions with destination size. (If the γ term is constant, the model implies a positive coefficient on the squared volatility term. However, negative coefficients are possible if the adjustment costs term is not a simple linear function of deviations from expectation).

The results confirm expectations of non-linearity. The coefficients of the squared demand volatility term are statistically significant in all specifications. However, unlike what a naive interpretation of the model would suggest, Table 6 indicates a steeper decline in exporter counts at the higher levels of demand volatility. This suggests a more complex structure to adjustment costs than I outline in the model. Concave relationships between demand volatility and the log of exporter numbers are possible with adjustment costs that depend on aggregate demand, or adjustment costs that are non-linear functions of squared deviations from projected trend. (If demand volatility enters exporters' considerations of profit strictly as described in equation (14) and γ is a constant, the curve should be convex). The more remarkable finding is that the higher order terms are non-zero and statistically significant.

Interactions with destination size also show consistency with the main

Table 6: Exporter Counts and Interacted Terms of Demand Volatility
(Dependent Variable: Log Number of Exporters in Destinations)

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	Log(Gross Export Entry)			Log(Gross Entry Post-2000)		
Demand Volatility	-0.613***	-0.301***	3.294***	-0.582***	-0.231***	3.321***
	(0.032)	(0.073)	(0.142)	(0.032)	(0.074)	(0.144)
(Demand Volatility)^2		-0.687***	-0.529***		-0.777***	-0.637***
		(0.131)	(0.135)		(0.132)	(0.135)
Destination Size			-0.503***			-0.495***
*Demand Volatility			(0.016)			(0.017)
Destination Size	0.421***	0.424***	0.442***	0.412***	0.415***	0.433***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Observations	371,531	371,531	371,531	363,381	363,381	363,381
R-squared	0.545	0.545	0.548	0.544	0.544	0.547
Country FE	Y	Y	Y	Y	Y	Y
Product FE	Y	Y	Y	Y	Y	Y

Robust standard errors in parentheses. Errors clustered by HS6 products.
*** p<0.01, ** p<0.05, * p<0.1

The units of observation are destinations: unique HS6-product and country combinations, e.g., Ethiopian imports of women's cotton overcoats (HS 610220). Gross entry is the log of unique firms with recorded exports to a product market between 2000 and 2006. Destination size is the log of total demand from a destination between 1995 and 2005. The coefficients suggest that γ in equation (14) is a non-linear functions of squared deviations from projected trend. If γ was a constant, as described in the model, the coefficient of the squared σ^2 term should be positive.

predictions. While the coefficient of the linear demand volatility term switches signs, the corresponding change associated with the market size interaction is larger. The squared volatility measure remains statistically significant and negative. Columns 4 to 6 repeat the same pattern for gross entry from 2001.

This section's results corroborate the predictions of the model for the effects of demand volatility on the number of exporters. The estimated effects of demand volatility on exporter choice are non-linear, and more exporters enter destinations with low demand volatility, holding other factors equal. The next steps check the robustness of the main findings.

3.3. Robustness Checks

This section addresses potential concerns about how exporter counts and demand volatility are measured. AppendixA.4 to AppendixA.6 include additional tests.

3.3.1. Demand Volatility Weighted by Recency

In estimating volatility, firms may ascribe greater weight to recent information (Bloom et al., 2007). Therefore, recent shocks may carry a disproportionate share of exporter's demand volatility estimates, (or less in times of high uncertainty). Figure 4 tests the idea by plotting the coefficient and R^2 values obtained for definitions of demand volatility with different weight indices η. The weights w_t, indexed from 1 to 10 put more emphasis on recent information with higher values of η; setting η to 1 reverts to the default scheme of equal weights. By design, the weighting scheme does not affect destinations with uniform deviations from the demand trajectory in all periods. (η specifies the relative size of the first and last terms of an arithmetic series that sums to H. H is the most recent period).

$$\sigma^2_{jk} = \sum_{t=1}^{H} w_t(\epsilon)^2, \qquad \epsilon \text{ is the residual defined in equation (17):} \tag{21}$$

$$\text{for } t \in [H, 1], \text{ where } \eta \in [1, 10]$$

$$w_t = \frac{2}{\eta+1} + \frac{2(t-1)(\eta-1)}{(H-1)(\eta+1)}, \quad \text{such that } \sum w_t = H = 11 \tag{22}$$

To create the figure, I repeat the baseline regression in Table 3 for each weighted variant of σ^2 and collate the estimated coefficients and R^2s.

Figure 4: Predictions for Weighted Demand Volatility

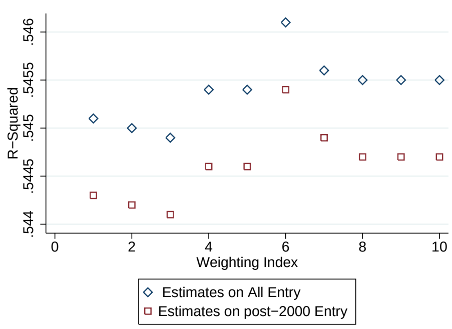

(a) R^2 values

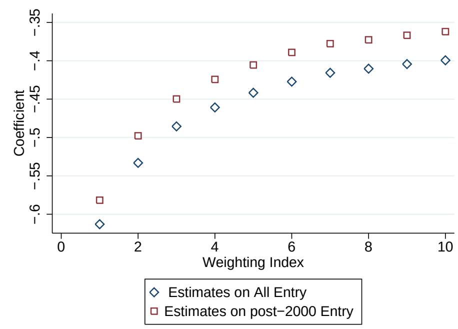

(b) Coefficient β_0 on demand volatility

Estimates after controlling for Market Size, Country and HS6 Product fixed effects. The demand volatility estimated with weight index 1 assigns equivalence to all periods, (as do all tables in the paper). The demand volatility estimated with weight 10 assigns deviations in 2005 10 times the weight of deviations in 1995. The weights are a linear density function. Data Sources: China GAC Export Data (2000-2006), UN COMTRADE

The left panel of Figure 4 indicates that increasing the weight of recent information in the estimation of demand volatility does not increase the explained variation in the number of exporters serving a destination. (The R^2 remain between 0.545 and 0.546). The right panel remains consistent. The estimated coefficient of demand volatility on exporter counts remains statistically significant, but decreases slightly in scale from the default estimate of 0.06 to 0.04 for index 10, the volatility estimate that ascribes 10 times the weight of the first year (1995) to the most recent year of demand history (2005).

These results suggest that the particular period used to estimate demand volatility does not affect the paper's main qualitative predictions.

3.3.2. Alternative Demand Volatility Measures

Table 7 replicates Table 3, measuring demand volatility as the sum of squared deviations from trend for quantities, not values in each destination. There are advantages and disadvantages to this approach: using quantities helps the econometrician avoid conflating exchange rates and price volatility with demand volatility, and quantities are a more faithful representation of demand volatility from the model in Section 2. However, quantity data is not always as reliable or available as data on trade values. The customs authorities that collect trade data have a stronger incentive to collect accurate information on values, and units of measurement for quantities are not always reported consistently by the original trade data sources (Gaulier and Zignago, 2010).

The two measures of demand volatility are nevertheless similar, with the quantity measure having a mean and standard deviation of 0.085 and 0.130, and the value-measure having 0.053 and 0.086. The quantity-based volatility measure is larger, as expected, given the inverse correlation between prices and quantities. That inverse correlation ensures that the volatility of prices and exchange rates, if any, reduces the observed volatility of export values relative to the volatility of quantities.

The coefficients in Table 7 are consistent with the numbers in Table 3. The estimated effect sizes are also larger: increasing this measure of demand volatility by one standard deviation corresponds to a 10.7% decline in exporter numbers, (calculated as 1-exp(-0.88*0.130)). The quantity-based definition yields estimates with greater economic and statistical significance than the value-based definition of demand volatility. This supports previous statements in the paper that the estimated effects in Table 3 tend to be conservative.

Table 8 replicates Table 3, measuring demand volatility as the standard deviation of annual growth rates of destinations. Section 2 summarizes the intuition behind this new measure of demand volatility. For a destination with a constant growth rate, measured demand volatility and adjustment costs would be zero, as the constant growth rate would match exporters' projections. The standard deviation of period-to-period growth is a common proxy for volatility, from finance and microeconomics (Guiso and Parigi, 1999; Sharpe, 1966) to macroeconomics and international trade (di Giovanni and Levchenko, 2009; Koren and Tenreyro, 2007). It may not fit the model in this paper exactly, but it is fairly simple to derive and explain. di Giovanni and Levchenko (2012b) use this measure to construct export-weighted average measures of risk for countries, in a study that links trade volatility to aggregate macroeconomic volatility. Koren and Tenreyro (2007) apply the same definition to GDP volatility in explaining differences between rich and poor countries. I include tests using this measure for consistency with related works on trade and volatility.

Table 8 shows that the estimated relationship between exporter numbers and demand volatility is robust to using this popular alternative measure.

Table 7: Exporter Counts and Quantity-Based Demand Volatility
(Dependent Variable: Log Number of Exporters in Destinations)

	(1)	(2)	(3)	(4)
VARIABLES	Log(Gross Export Entry)		Log(Gross Entry Post-2000)	
Demand Volatility	-0.882***	-0.889***	-0.866***	-0.871***
	(0.022)	(0.021)	(0.022)	(0.021)
Destination Size	0.391***	0.392***	0.383***	0.384***
	(0.003)	(0.003)	(0.003)	(0.003)
Log(GDP)	0.002		0.001	
	(0.001)		(0.001)	
Log(GDP per capita)	0.000		0.000	
	(0.002)		(0.002)	
Log(Distance)	0.002		0.001	
	(0.006)		(0.006)	
Constant	-0.666***		-0.663***	
	(0.077)		(0.077)	
Observations	280,060	381,196	273,567	372,373
R-squared	0.539	0.538	0.538	0.536
Country-Year FE		Y		Y
Product FE	Y	Y	Y	Y

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

The units of observation are destinations: unique HS6-product and country combinations, e.g., Ethiopian imports of women's cotton overcoats (HS 610220). Exporter counts represent the log of the number of unique firms with recorded exports to a destination between 2000 and 2006. The gross entry post-2000 captures the difference between all firms that served a destination and firms that served the destination in 2000. Demand volatility is constructed as in Table 3, but with quantities not dollar values of demand from 1995-2005 in each destination. Control variables used but not shown in the table include geographic remoteness and dummies for shared borders, common languages and WTO membership.

For each destination, I measure demand volatility as the standard deviation of demand growth between 1995 and 2005. The mean and standard deviation of this new measure are 0.78 and 0.55 respectively. (Differences between these numbers and Table 1 are indicative of how each measure is constructed).

The coefficients in Table 8 are consistent with the numbers in Table 3, both in sign and significance. The predicted effects sizes are also similar: a 4.8% decline in exporter numbers is expected for the average destination for a standard deviation increase in this measure of demand volatility, compared with 5.1% in the first table. In sum, this new definition yields results that broadly agree with the former definition of demand volatility.

I replicate Tables 4 to 9 using this measure and obtain the same pattern: slightly smaller coefficients with the same sign and significance as the original measure.

AppendixA.5 presents results that show the main findings are robust to the selection of markets used to define demand volatility. That section replicates Table 3 using two sets of demand histories to define demand volatility: product-country volatility for Chinese goods, and destination volatility for the rest of the world. The first definition of volatility corresponds to a model where Chinese exporters' expectations of demand volatility is linked to only the history of goods imported from China for any given destination. The definition of demand volatility that excludes imports from China follows an alternative scenario in which expectations neither correspond to firms' own historic sales to a destination, nor to past supply shocks from China.

The two sets of results in AppendixA.5 yield the same finding as the previous tables in the paper, we see that destinations with high demand volatility are served by fewer exporters. The coefficients for the main variables in the tables retain their sign and statistical significance, though there are small differences in the R^2 values, as with the number of observations. This is not unexpected, as the number of destinations with years of usable history changes.

3.3.3. Poisson Regressions to Include Zeros

Tables 3 to 8 ignored destinations with zero exporters – the estimations used logarithms of a count variable. Table 9 addresses concerns of possible bias from ignoring these destinations with a Poisson regression. The inclusion of observations with zero exporters should accentuate the claims made in the previous section because demand volatility is higher on average for these destinations, as shown in Figure A.5.

Table 8: Exporter Numbers and Growth-Based Demand Volatility
(Dependent Variable: Log Number of Exporters in Destinations)

VARIABLES	(1) Log(Gross Export Entry)	(2)	(3) Log(Gross Entry Post-2000)	(4)
Demand Volatility	-0.095***	-0.152***	-0.082***	-0.134***
	(0.011)	(0.012)	(0.010)	(0.012)
Destination Size	0.283***	0.240***	0.278***	0.236***
	(0.004)	(0.003)	(0.004)	(0.003)
Log(GDP)	0.224***		0.224***	
	(0.003)		(0.003)	
Log(Distance)	-0.540***		-0.514***	
	(0.005)		(0.005)	
Constant	6.410***		6.030***	
	(0.092)		(0.091)	
Observations	311,337	362,490	304,545	354,337
R-squared	0.710	0.748	0.708	0.745
Product FE	Y	Y	Y	Y
Country FE		Y		Y

Robust standard errors in parentheses. Errors clustered by HS6 products.
*** p<0.01, ** p<0.05, * p<0.1

The units of observation are destinations: unique HS6-product and country combinations, e.g., Ethiopian imports of women's cotton overcoats (HS 610220). Exporter counts represent the log of the number of unique firms with recorded exports to a destination between 2000 and 2006. The gross entry post-2000 captures the difference between all firms that served a destination and firms that served the destination in 2000. Demand volatility is constructed as the standard deviation of destinations' annual growth rates from 1995-2005. Control variables used but not shown in the table include geographic remoteness and dummies for shared borders, common languages and WTO membership.

Table 9: Exporter Counts and Demand Volatility: Poisson Estimates
(Dependent Variable: Number of Exporters in Destinations)

	(1)	(2)
VARIABLES	Gross Export Entry	Gross Entry Post-2000
Demand Volatility	-2.518***	-2.408***
	(0.092)	(0.093)
Destination Size	0.288***	0.276***
	(0.006)	(0.006)
Observations	688,005	686,997

Robust standard errors in parentheses. Errors clustered by HS6 products.
*** p<0.01, ** p<0.05, * p<0.1

The units of observation are destinations: unique HS6-product and country combinations, e.g., Ethiopian imports of women's cotton overcoats (HS 610220). All 708,000 destinations with measurable demand volatility were included, many of which had zero Chinese exporters. Product and country fixed effects limited the usable observations to about 690,000. Gross entry is the log of unique firms with recorded exports to a destination between 2000 and 2006. The gross entry post-2000 variable captures the unique number of entrants from the year China joined the WTO.

The estimates using Poisson regressions with fixed effects are consistent with those from Tables 3, and larger, as expected. Interpreting the coefficients in Table 9 suggests that a 19% decline in the number of exporters should be associated with a standard deviation increase in demand volatility from the mean, holding other factors constant. (The response is calculated as $40 * \{1 - exp[(-2.5) * 0.086]\}$, using negative terms to represent the negative predicted effect). The estimates are higher than the comparable numbers in Table 4. The destinations with no Chinese exporters had higher levels of demand volatility and the estimates are consistent with expectations in terms of size, sign and significance.

4. Discussions and Conclusion

This paper examines how the margins of trade are shaped by observed volatility – the patterns of exogenous fluctuations in aggregate demand from potential export destinations. The nature of these exogenous fluctuations is well documented (Carroll et al., 2011; Hall, 2004). Demand volatility has country and product-specific patterns, such that imports for some raw materials like steel are particularly volatile for most countries, and countries with low GDP per capita generally have high import volatility, as documented by papers that link national output volatility to GDP per capita.[33] Nevertheless, neither product categories nor countries as isolated variables explain most of the variation in the demand volatility of destinations – unique combinations of products and countries.

I provide evidence that higher demand volatility in a destination leads to fewer Chinese exporters and lower trade volumes. (The findings do not depend on whether demand volatility is measured using data that excludes Chinese exporters as in AppendixA.5, that is weighted to overrepresent recent demand data as in section 3.3.1 or is defined conventionally as the standard deviation of year-on-year growth as in section 3.3.2). Further, I show that a notable share of the variation in trade is explained by observed volatility. The R^2 figures explaining the variable's relevance at the beginning of the paper

[33]Papers in this vein hypothesize that aggregate volatility is due to the correlation between national policy shocks and sectoral shocks, the capacity of economic institutions, as well as specialization in sectors that are unusually volatile (di Giovanni and Levchenko, 2012b; di Giovanni and Levchenko, 2009; Krishna and Levchenko, 2009; Koren and Tenreyro, 2007).

come from Table A.16 in the appendix, which shows that demand volatility's explanatory power compares reasonably with those of established variables like GDP and distance.

In showing a statistically significant first-order relationship between the extensive margin of trade and demand volatility, this paper suggests a new factor that explains observed trade volumes, prices and quantities. In explaining trade, demand volatility complements existing explanations that rely on exporter productivity, geography and economic size (Anderson and Van Wincoop, 2003; Melitz, 2003). As I show in Section AppendixA.7, import volatility is higher in developing economies, consistent with the findings of di Giovanni and Levchenko (2012b) and Koren and Tenreyro (2007). The paper's main finding suggests that these economies may be hurt by the absence of imported varieties.

As imports can be a source of productive efficiency in developing economies (Goldberg et al., 2010; Connolly, 2003), the foregoing implies that low demand volatility can be an advantage for producers. Firms in locations with low demand volatility can enjoy more variety in inputs and possibly lower prices. This matters because imported inputs are a common feature of production processes – for developing and developed economies alike. One could extend this idea to assessing the impacts of trade diversification policies designed to reduce aggregate economic volatility, as discussed in Cadot et al. (2011).

Rational economic agents respond to volatility, and firms engaged in international trade are no exception. This paper extends the conventional model of trade to include adjustment costs in the presence of stochastic demand. (Adjustment costs reflect the fact that demand shocks affect unit costs: firms must pay overtime wage premiums and equipment deactivation costs when demand either peaks or crashes). These costs motivate exporters' choices, with clear implications for the number of firms that export, and the level of trade between countries.

Firm level data on the universe of Chinese export transactions between 2000 and 2006 support this hypothesis: Fewer exporters serve destinations with high demand volatility. The results are robust to how exporter numbers are counted – as the unique number of exporters in all years, as the unique number of new exporters after 2000 when China adopted liberalization policies in preparation for WTO membership, or as annual counts of exporters in each of the years. The results are also robust to how demand volatility is defined, and to several empirical specifications. Furthermore, the data confirm the model's predictions about how the relationship between exporter counts and demand volatility is non-linear. [34]

The estimated effects of demand volatility are statistically and substantially significant. Increasing demand volatility by one standard deviation for the average destination in my conservative specification predicts a 5% drop in the number of exporters that serve the destination. The explained variation in exporter counts and predicted effects of this new variable are comparable to those obtained from conventional predictors of trade like GDP and distance.

These findings suggest further work to evaluate how volatility affects the process of economic development, given how instrumental trade has been to growth in the last half-century. The paper also motivates an inquiry into how differences in labor and capital adjustment costs explain the heterogeneity in producers' responses to demand volatility.

[34]The estimates of exporter responses to demand volatility use Chinese firm-level export data. This choice was necessitated in part by the absence of a dataset that reports exporter numbers for all countries. Furthermore, China is the world's largest exporter, and its exports are highly correlated with global exports. (The correlation is 0.995 for China's aggregate exports and the rest of the world's between 2000 and 2006). Therefore, the empirical exercises in the chapter should be informative about the global pattern of exporter behavior, which is what the model describes.

References

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

Ahn, J., McQuoid, A. F., 2012. Capacity Constrained Exporters: Micro Evidence and Macro Implications. Working Paper.

Akhmetova, Z., Mitaritonna, C., 2012. A Model of Firm Experimentation under Demand Uncertainty with an Application to Multi-Destination Exporters.

Amiti, M., Freund, C., 2010. The Anatomy of China's Export Growth. In: China's Growing Role in World Trade. University of Chicago Press, pp. 35–56.

Anderson, J. E., Van Wincoop, E., 2003. Gravity with Gravitas: A Solution to the Border Puzzle. American Economic Review 93 (1), 170–192.

Armenter, R., Koren, M., 2014. A Balls-and-Bins Model of Trade. American Economic Review 104 (7), 2127–2151.

Axtell, R. L., 2001. Zipf Distribution of US Firm Sizes. Science 293 (5536), 1818–1820.

Baldwin, R., Harrigan, J., 2011. Zeros, Quality, and Space: Trade Theory and Trade Evidence. American Economic Journal-Microeconomics 3 (2), 60–88.

Bas, M., 2009. Trade, Foreign Inputs and Firms Decisions: Theory and Evidence. Working Paper.

Bloom, N., Bond, S., Van Reenen, J., 2007. Uncertainty and Investment Dynamics. Review of Economic Studies 74 (2), 391–415.

Blum, B. S., Claro, S., Horstmann, I. J., 2013. Occasional and Perennial Exporters. Journal of International Economics 90 (1), 65–74.

Cadot, O., Carrère, C., Strauss-Kahn, V., 2011. Export Diversification: What's Behind the Hump? Review of Economics and Statistics 93, 590–605.

Carroll, C. D., Slacalek, J., Sommer, M., 2011. International Evidence on Sticky Consumption Growth. Review of Economics and Statistics 93 (4), 1135–1145.

Chaney, T., 2008. Distorted Gravity: the Intensive and Extensive Margins of International Trade. American Economic Review 98 (4), 1707–1721.

Chevassus-Lozza, E., Gaigné, C., Le Mener, L., 2013. Does Input Trade Liberalization Boost Downstream Firms' Exports? Theory and Firm-Level Evidence. Journal of International Economics 90 (2), 391–402.

Connolly, M., 2003. The Dual Nature of Trade: Measuring its Impact on Imitation and Growth. Journal of Development Economics 72 (1), 31–55.

Cooper, R. W., Haltiwanger, J. C., 2006. On the Nature of Capital Adjustment Costs. Review of Economic Studies 73 (3), 611–633.

Crozet, M., Koenig, P., 2010. Structural Gravity Equations with Intensive and Extensive Margins. Canadian Journal of Economics 43 (1), 41–62.

Cuñat, A., Melitz, M. J., 2012. Volatility, Labor Market Flexibility, and The Pattern Of Comparative Advantage. Journal of the European Economic Association 10 (2), 225–254.

Das, S., Roberts, M. J., Tybout, J. R., 2007. Market Entry Costs, Producer Heterogeneity, and Export Dynamics. Econometrica 75 (3), 837–873.

Davis, S. J., Haltiwanger, J., 1992. Gross Job Creation, Gross Job Destruction, and Employment Reallocation. Quarterly Journal of Economics 107 (3), 819–863.

di Giovanni, J., Levchenko, A. A., 2009. Trade Openness and Volatility. Review of Economics and Statistics 91 (3), 558–585.

di Giovanni, J., Levchenko, A. A., 2012a. Country Size, International Trade, and Aggregate Fluctuations in Granular Economies. Journal of Political Economy 120 (6), 1083–1132.

di Giovanni, J., Levchenko, A. A., 2012b. The Risk Content of Exports: A Portfolio View of International Trade. In: NBER International Seminar on Macroeconomics. Vol. 8. University of Chicago Press, pp. 97–151.

Dixit, A., 1989. Entry and Exit Decisions Under Uncertainty. Journal of Political Economy, 620–638.

Eaton, J., Kortum, S., 2001. Trade in Capital Goods. European Economic Review 45 (7), 1195 – 1235, international Seminar On Macroeconomics.

Eaton, J., Kortum, S., Kramarz, F., 2004. Dissecting Trade: Firms, Industries, and Export Destinations. American Economic Review 94 (2), 150–154.

Foster, L., Haltiwanger, J., Syverson, C., 2008. Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability? American Economic Review 98 (1), 394–425.

Frankel, J., Rose, A., 2002. An Estimate of the Effect of Common Currencies on Trade and Income. Quarterly Journal of Economics, 437–466.

Gaulier, G., Zignago, S., 2010. BACI: International Trade Database at the Product-Level. The 1994-2007 Version. CEPII Working Papers (Working Papers 2010-23).

Glick, R., Rose, A. K., 2002. Does a Currency Union Affect Trade? The Time-Series Evidence. European Economic Review 46 (6), 1125–1151.

Goldberg, P. K., Khandelwal, A. K., Pavcnik, N., Topalova, P., 2010. Imported Intermediate Inputs and Domestic Product Growth: Evidence from India. Quarterly Journal of Economics 125 (4), 1727–1767.

Guimarães, P., Portugal, P., 2009. A Simple Feasible Alternative Procedure to Estimate Models with High-Dimensional Fixed Effects. SSRN 1329573.

Guiso, L., Parigi, G., 1999. Investment and Demand Uncertainty. Quarterly Journal of Economics 114 (1), 185–227.

Hall, R. E., 2004. Measuring Factor Adjustment Costs. Quarterly Journal of Economics, 899–927.

Hallak, J. C., Schott, P. K., 2011. Estimating Cross-Country Differences in Product Quality. Quarterly Journal of Economics 126, 417–474.

Hallak, J. C., Sivadasan, J., 2009. Firms' Exporting Behavior Under Quality Constraints. NBER Working Paper 14928.

Handley, K., Limão, N., 2012. Trade and Investment Under Policy Uncertainty: Theory and Firm Evidence. NBER Working Paper 17790.

Handley, K., Limão, N., 2013. Policy Uncertainty, Trade and Welfare: Theory and Evidence for China and the US. NBER Working Paper 19376.

Harrigan, J., Ma, X., Shlychkov, V., 2011. Export Prices of US Firms. NBER Working Paper.

Head, K., Mayer, T., Ries, J., 2010. The Erosion of Colonial Trade Linkages after Independence. Journal of International Economics 81 (1), 1–14.

Helpman, E., Melitz, M., Rubinstein, Y., 2008. Estimating Trade Flows: Trading Partners and Trading Volumes. Quarterly Journal of Economics 123 (2), 441–487.

Helpman, E., Melitz, M. J., Yeaple, S. R., 2004. Export versus FDI with Heterogeneous Firms. American Economic Review 94 (1), 300–316.

Hsieh, C.-T., Ossa, R., 2011. A Global View of Productivity Growth in China. NBER Working Paper 16778.

Hummels, D., Klenow, P., 2005. The Quality and Variety of a Nations Trade. American Economics Review 95 (3), 704–723.

Jones, C. I., 2011. Intermediate Goods and Weak Links in the Theory of Economic Development. American Economic Journal: Macroeconomics, 1–28.

Koren, M., Tenreyro, S., 2007. Volatility and Development. Quarterly Journal of Economics 122 (1), 243–287.

Krishna, P., Levchenko, A. A., 2009. Comparative Advantage, Complexity and Volatility. NBER Working Paper 14965 (w14965).

Lee, J.-W., 1995. Capital Goods Imports and Long-run Growth. Journal of Development Economics 48 (1), 91–110.

Liu, Y., 2012. Capital Adjustment Costs: Implications for Domestic and Export Sales Dynamics. Working Paper.

Lucas, R. E., 1967. Adjustment Costs and the Theory of Supply. Journal of Political Economy, 321–334.

Manova, K., Yu, Z., 2012. Firms and Credit Constraints along the Global Value Chain: Processing Trade in China. NBER Working Paper 18561.

Manova, K., Zhang, Z., 2012. Export Prices Across Firms and Destinations. Quarterly Journal of Economics 127 (1), 379–436.

Melitz, M. J., 2003. The Impact Of Trade On Intra-Industry Reallocations And Aggregate Industry Productivity. Econometrica 71 (6), 1695–1725.

Milgrom, P., Segal, I., 2002. Envelope Theorems for Arbitrary Choice Sets. Econometrica 70 (2), 583–601.

Nguyen, D. X., 2011. Demand Uncertainty: Exporting Delays and Exporting Failures. Journal of International Economics 86 (2), 336–344.

Pindyck, R. S., 1982. Adjustment Costs, Uncertainty, and The Behavior Of The Firm. American Economic Review, 415–427.

Rob, R., Vettas, N., 2003. Foreign Direct Investment and Exports with Growing Demand. Review of Economic Studies 70 (3), 629–648.

Sharpe, W. F., 1966. Mutual Fund Performance. Journal of Business, 119–138.

Soderbery, A., 2013. Market Size, Structure, and Access: Trade with Capacity Constraints. Working Paper.

Staiger, R. W., Wolak, F. A., 1992. The Effect of Domestic Antidumping Law in the Presence of Foreign Monopoly. Journal of International Economics 32 (3), 265–287.

Vannoorenberghe, G., 2012. Firm-level Volatility and Exports. Journal of International Economics 86 (1), 57–67.

Viner, J., 1922. The Prevalence of Dumping In International Trade. Journal of Political Economy 30 (5), 655–680.

AppendixA. Data, Variable Definitions and Additional Tests

AppendixA.1. Exporter-Level Data

Firm level export choices are taken from the universe of Chinese export transactions between 2000 and 2006, collapsed to annual values for firms in each product-country. This dataset identifies the year of trade, firms by unique IDs, the countries to which they export, and the products sent to each destination. The full data set exceeds 24 million observations. The raw data report the f.o.b. value of exports in nominal U.S. dollars in an unbalanced panel over 7 years of more than 240,000 firms serving 390,000 destinations comprising 200 economies and about 4,100 HS6 product categories. This rich dataset provides no identifiers for buyers in overseas markets, unfortunately. Thus, each destination conceptually stands for one representative consumer.

To link this data to the destinations identified in the COMTRADE data, I match the two sources on countries and product categories. The product categories are originally reported as eight-digit HS categories in the firm-level data, of which the leading 6 digits correspond to standardized categories.[35] To ensure that the product category definitions remain consistent over time, I convert all years to the 1992 HS standard using the concordances provided by the UN at `http://unstats.un.org/unsd/trade/conversions/HSCorrelationandConversiontables.htm`. The destination data that I describe next are reported using the 1992 HS standard. I complete the matching between the two data sources by mapping the country-codes in the export data to the standardized ISO categories used by COMTRADE.

For Chinese exporters, 2000 to 2006 was a period of remarkable growth and entry into foreign destinations. It spans China's entry into the WTO at the end of 2001, which lowered trade costs, reduced internally mandated barriers to trade and created export opportunities for Chinese firms. The nominal dollar value of Chinese goods exports nearly quadrupled to \$968bn in 2006 from \$250bn 2000. A large share of this growth was at the extensive margin —the number of exporters went from 62,600 to more than 170,000. (Table A.10 decomposes this growth in exporter numbers into its intensive

[35]The 6-digit harmonized system (HS6) is a global standard used for reporting trade between most countries; revisions to its roughly 5,000 product categories occurred in 1996, 2002 and 2007. Each country may have more detailed HS8 or HS10 categories that further refine the HS6 product categories

and extensive margins).[36] Among firms that remained exporters, entry into new destinations was pervasive. Four out of five exporters chose or added a new product-country destinations in the average year.[37]

Table A.10: Exporter Dynamics in China: 2000 - 2006

Year	Exporter Count	Incumbents	Entrants	Leavers
	$A = B + C$	$B = A_{t-1} - D$	C	D
2000	62,603			
2001	68,347	52,201	16,146	10,402
2002	78,567	57,263	21,304	11,084
2003	95,627	68,506	27,121	10,061
2004	120,363	82,858	37,505	12,769
2005	143,583	103,724	39,859	16,639
2006	170,642	124,419	46,223	19,164

	Exporter-ProdMkt Count	Incumbents	Entrants	Exits
	$A = B + C$	$B = A_{t-1} - D$	C	D
2000	1,782,803			
2001	2,011,808	696,379	1,315,429	1,086,424
2002	2,464,544	828,853	1,635,691	1,182,955
2003	3,076,358	1,059,347	2,017,011	1,405,197
2004	3,827,074	1,307,810	2,519,264	1,768,548
2005	4,846,699	1,593,626	3,253,073	2,233,448
2006	5,895,393	1,907,010	3,988,383	2,939,689

Table A.10 also helps to justify the paper's focus on gross entry as a measure of exporter numbers in Tables 3 to 6. The first panel shows exporter numbers nearly tripled from 2000 to 2006. The last three columns in each panel break down the annual changes into entry (column C), exits (column D) and holdovers (column B). The increase in the number of exporters fits

[36] See Ahn et al. (2011); Manova and Yu (2012) for fuller descriptions of how WTO accession reduced trade costs for Chinese exporters.

[37] The average exporter in 2000 served 28 destinations - 13 countries and 7 HS6 products, while the corresponding number for 2006 was 34 (16 countries and 8 products). Destinations had a skewed distribution of Chinese exporter participation —40 firms on average served each destination. The median exporter count was 5. The reported central moments exclude destinations with zero Chinese exporters.

the pattern expected for trade liberalization with China's WTO accession.

Entry is also the dominant dynamic at the finer level of firm destination combinations (in the second panel). Here, turnover rates are higher. The fact that many exporters exit a destination in one year only to return in a later year suggests the need to redefine entry over periods longer than a year.

AppendixA.2. Product-Country Destination Data

Global trade in nominal US dollar terms grew at an average annual rate of 7% to reach $12tn in 2006, covering more than 220 countries and 5,000 HS6 products. Approximately 990,000 unique destinations registered imports in the COMTRADE database, though many had zero demand in several years. The COMTRADE database indicates a moderate expansion of destinations from 695,000 to 738,000 between 2000 and 2006. (In the same period, destinations served by Chinese exporters increased by 64% – from 179,000 to 292,000).

The original data set reports more than 63 million observations of trade at the HS6 product level for importing and exporting country-pairs in years from 1995 to 2005. (The full dataset goes to 2012, but only the first 11 years are usable as history because the firm level data stops in 2006). I collapsed this data to importing country-HS6 combinations for each of the years, noting that HS6 categories remain consistent over time. The UN COMTRADE data is reporterd in the 1992 version of the HS6 system for all years (Gaulier and Zignago, 2010). This collapsed form represents the history of imported demand into each destination in the analysis. I expand the data to a balanced panel of 11 years for all 990,000 destinations. This identifies instances of zero-demand. (Missing data is coded as zero). Market sizes and demand volatility come from this expanded data.

About 390,000 unique destinations were served at least once by Chinese exporters between 2000 and 2006. About 200,000 of these had no Chinese exporters in 2000 and only show up later in the exporter-level data. It is important to evaluate the differences between the destinations with Chinese exporters and the other destinations, of the 990,000 possibilities. Some of the variables of interest in this comparison of destinations that drew Chinese exporters include size and demand volatility.

Figure A.5 graphs the distribution of sizes and demand volatility simultaneously for all destinations in the data, comparing those served by Chinese exporters with the remainder.

Figure A.5: Destinations Served (and Not Served) by Chinese Exporters

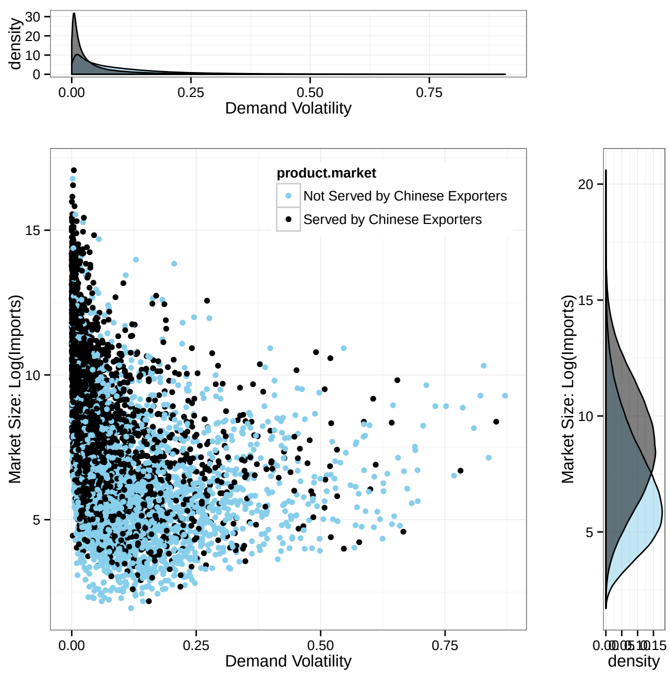

The scatter plot only shows a random 1% sample of the more than 900,000 destinations. The density plots for market size and demand volatility use the full data set. Demand Volatility represents deviations around the import trend as defined in Section 3.1.2, and market size is the log of the total USD value of imports into the destination between 1995 and 2005.

Data Sources: China GAC Export Data (2000-2006), UN COMTRADE

The density plots in Figure A.5 indicate that the regressions in the main body of the paper provide broad coverage of destinations in terms of size and demand volatility. The ranges of demand volatility and destination sizes covered by the two categories appear similar in the graph. Nevertheless, from the distribution at the top of the graph, one sees that a larger share of Chinese exporters – more than the global average – were concentrated in destinations with low demand volatility. The average demand volatility for destinations served by Chinese exporters is .053, while markets not served by Chinese exporters have an average of .110. Table 2 makes the same point with linear regressions. The density on the right panel shows that the destinations Chinese exporters serve are also larger on average.

To facilitate replication, here are country-level variables used but not reported in Tables 3, 4 and 8 to test the relationship between demand volatility and the number of Chinese exporters that enter a destination. They are from Head et al. (2010).

Table A.11: Additional Regression Variables

VARIABLES	N	Mean	Std. Dev.	Min.	Max.
GDP	289801	10.06095	2.438002	4.258556	16.39586
GDP per Capita	288048	8.018416	1.572457	4.634483	11.11192
Log(Distance)	346217	8.989019	.5693675	6.925665	9.857974
Log(Remoteness)	346217	-9.105768	.4982548	-10.53251	-8.298696
Contiguity	346217	.0830606	.2759742	0	1
Language	346217	.0277225	.1641769	0	1
Legal Origin	346217	.1771115	.3817636	0	1
GATT/WTO	346217	.7430744	.4369386	0	1

GDP: GDP of the destination economy
GDP per Capita: of the destination economy
Distance: from China to the destination economy
Remoteness: Geographic remoteness, i.e. Destination's GDP-weighted distance from all countries
Contiguity: Dummy indicating whether country has shared borders with China
Language: Dummy indicating whether country shares ethnic or official languages with China
Legal Origin: Dummy indicating whether country shares legal origin with China
GATT/WTO: Whether destination economy was a member of WTO

AppendixA.3. Product and Country Variation in Demand Volatility

Country-specific factors explain as much of the variation in demand volatility as product-specific factors. This is consistent with the patterns observed for output volatility in Koren and Tenreyro (2007).

Table A.12 presents linear regressions of demand volatility on destination size. The first panel in the table is limited to destinations with at least one Chinese exporter, while the second panel extends the regressions to include all destinations with aggregate demand volatility data. Although destination size explains a notable share of the variation in demand volatility in the first panel, it is interesting to note that country-specific factors also explain a larger share than product specific factors. This comes from a comparison of columns 1 and 2. When combined with destination size, the two sets of fixed effects explain comparable incremental shares of the variation in the dependent variable, if one compares the R^2 values in columns 4 and 5.

The second panel of Table A.12 follows the same pattern as the first panel; smaller destinations tend to have higher demand volatility. (This called for the inclusion of market size as a control in the regressions that I report in the main body of the paper). Furthermore, product-fixed effects generally explain less of the variation in demand volatility if one does not control for destination size. As expected, destination size is driven in part at the country-level. Larger economies import more of most product categories.

AppendixA.4. Annual Exporter Counts and Demand Volatility

To exploit the annual variations in export flows and exporter counts, Table A.13 provides results of regressions structured after conventional gravity model estimations. The regressions include estimates with country-year fixed effects in the even-numbered columns.

The results indicate that demand volatility decreases export volumes and exporter counts, just as in Table 4. Unlike Table 4, the dependent variables here measure annual export volumes and exporter numbers like most gravity model estimations.[38] In this table, exports per exporter decrease with increased volatility, another point of difference with the cumulative version

[38]One could try to reconcile this to the model with claims that each year represents its own equilibrium; a claim that requires justification, but that may change how demand volatility should be defined.

Table A.12: Demand Volatility: Product and Country Fixed Effects
(Dependent Variable: Demand Volatility)

For Destinations Served by Chinese Exporters						
VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
Destination Size			-0.02***	-0.02***	-0.01***	-0.01***
			(0.00)	(0.00)	(0.00)	(0.00)
Constant	0.08***	0.08***	0.21***	0.23***	0.19***	
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	
Observations	259,963	259,963	259,963	259,963	259,963	259,963
R-squared	0.10	0.15	0.15	0.25	0.20	0.31
Product FE	Y			Y		Y
Country FE		Y			Y	Y
For All Destinations						
VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
Destination Size			-0.02***	-0.02***	-0.01***	-0.01***
			(0.00)	(0.00)	(0.00)	(0.00)
Constant	0.08***	0.08***	0.21***	0.22***	0.19***	
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	
Observations	708,802	708,802	708,802	708,802	708,802	708,802
R-squared	0.09	0.14	0.15	0.24	0.20	0.29
Product FE	Y			Y		Y
Country FE		Y			Y	Y

Robust standard errors in parentheses. Errors clustered by HS6 products.
*** p<0.01, ** p<0.05, * p<0.1

Notes: The units of observation are destinations: unique HS6-product and country combinations, e.g., Ethiopian imports of women's cotton overcoats (HS 610220). Demand volatility is the sum of the squared deviations of demand from a linear trend over the years 1995 to 2005. Destination size is the sum of aggregate demand in each destination from 1995 to 2005.

in the main body of the paper. However, this table does not include full-fixed effects for China, only its GDP is included as an additional control for changes over time.

Table A.13: Annual Trade Estimates with Demand Volatility
(Dependent Variable: Log Export Indicator in Destinations by Year)

VARIABLES	(1) Log(Exports)	(2)	(3) Log(Exporter Counts)	(4)	(5) Log(Exports per Exporter)	(6)
Demand Volatility	-2.619***	-2.398***	-1.201***	-0.921***	-1.418***	-1.477***
	(0.084)	(0.083)	(0.040)	(0.041)	(0.056)	(0.054)
Log(GDP)	0.687***		0.343***		0.345***	
	(0.006)		(0.003)		(0.003)	
Log(GDP per Capita)	-0.095***		0.016***		-0.111***	
	(0.006)		(0.003)		(0.004)	
Log(GDP China)	2.983***		1.848***		1.134***	
	(0.051)		(0.027)		(0.034)	
Log(Distance)	-0.573***		-0.372***		-0.200***	
	(0.011)		(0.005)		(0.007)	
Constant	19.190***		8.133***		11.056***	
	(0.275)		(0.127)		(0.186)	
Observations	1,485,716	1,590,273	1,485,716	1,590,273	1,485,716	1,590,273
R-squared	0.435	0.442	0.570	0.575	0.329	0.340
Country FE		Y		Y		Y
Product FE	Y	Y	Y	Y	Y	Y

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

The units of observation are destination - years: unique year, HS6-product and country combinations, e.g., Ethiopian imports of women's cotton overcoats (HS 610220) in 2006. The dependent variable is log (the number of firms with recorded exports to a destination) in each of the years between 2000 and 2006. Destination size is the value of all imports into each destination, for the respective year.

To address concerns about time-varying factors like exchange rate volatility, country-specific shocks or trade deals, the even-numbered columns include country-year fixed effects. These improve the estimated coefficients, usable observations and explained variation. A caveat is necessary: The demand volatility term used in this table does not change over time, nor do the country-level measures like GDP change within products for each country. In other words, the estimated variables are not well matched. (Variants of this table that include the destination size term, or a measure of logged annual imports from all destinations also predict smaller exporter numbers with higher demand volatility, even if they do not consistently predict lower volumes).

AppendixA.5. Other Measures of Demand Volatility

To avoid confronting criticism that that the measure of demand volatility does not reflect the expectations of Chinese exporters, or is biased, this section repeats the estimates in Table 3 using two variations on the measure of demand volatility.

AppendixA.5.1. Chinese Aggregate Demand Volatility

First, I use only imports from China to estimate the volatility measure for each destination. This goes with the idea that Chinese goods are poor substitutes for other countries' products, *á la* the Armington national differentiation narrative, so that in each destination, there is a market for Chinese goods, and Chinese exporters expect demand fluctuations to reflect the history of demand for only Chinese goods.

Formally, this alternative measure of demand volatility:

$$\sigma^2_{jk} = \sum_t (\epsilon_{jkt})^2$$

where ϵ_{jkt} is derived from:

$$\frac{Q_{jkt}}{\sum_t Q_{jkt}} = \zeta_{jk} t + \alpha_{jk} + \epsilon_{jkt}$$

Q_{jkt} is the sum of imports into destination jk from only Chinese sources in year t. I use the Chinese component of the UN COMTRADE data for this purpose.

Table A.14 shows that destinations with high demand volatility have fewer exporters, just as in Table 3. Destination size, GDP, and Distance keep the same signs as the comparable results from the body of the chapter, with more of the coefficients being statistically significant in this specification.

There are other differences, all of which are reasonable for this specification. More of the variation is explained by the right-hand side variables. (R^2 values in Table A.14 are higher by about 0.24 on average). There are fewer

Table A.14: Exporter Counts and Demand Volatility (China Only):
(Dependent Variable: Log Number of Exporters in Destinations)

	(1)	(2)	(3)	(4)
VARIABLES	Log(Gross Export Entry)		Log(Gross Entry Post-2000)	
Demand Volatility	-1.140***	-1.113***	-1.081***	-1.062***
	(0.022)	(0.020)	(0.022)	(0.020)
Destination Size	0.366***	0.333***	0.352***	0.319***
	(0.002)	(0.002)	(0.002)	(0.002)
Log(GDP)	0.200***		0.205***	
	(0.002)		(0.002)	
Log(GDP per capita)	0.002		-0.001	
	(0.002)		(0.002)	
Log(Distance)	-0.321***		-0.308***	
	(0.004)		(0.004)	
Constant	3.515***		3.289***	
	(0.063)		(0.064)	
Observations	244,139	279,780	240,832	275,951
R-squared	0.784	0.814	0.777	0.805
Country-Year FE	N	Y	N	Y
Product FE	Y	Y	Y	Y

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

The units of observation are destinations: unique HS6-product and country combinations, e.g., Ethiopian imports of women's cotton overcoats (HS 610220). Gross entry is the log of unique firms with recorded exports to a destination between 2000 and 2006. The change in exporter count captures the difference between all firms that served a destination and firms that served the destination in 2000. This difference measures new exporters that appeared with China's trade liberalization from 2001 onwards. Control variables used but not shown in the table include geographic remoteness and dummies for shared borders, common languages and WTO membership. Missing observations in Columns 1 and 3 are because GDP data is not always available. However, estimates on the largest common sample are almost identical to columns 2 and 4.

usable observations, as there are destinations with multiple years of available demand history globally, but for which demand history for Chinese exporters alone is not sufficient to construct a volatility measure.

One possible weakness of Table A.14 is that the results may be biased for destinations served by only a few Chinese exporters. For those destinations, the measure of volatility largely reflects shocks to the Chinese exporters; firm-level shocks not related to the destination's attributes could make the measured volatility higher than its implicit value.

Nevertheless, the estimates support the main finding that exporters on average, enter destinations with low demand volatility in greater numbers.

AppendixA.5.2. Aggregate Demand Volatility: Rest of the World

Second, I use imports from the rest of the world, excluding China, to measure demand volatility for each destination. Compared with Section AppendixA.5.1, this takes a different perspective in assuming that Chinese goods are substitutes for items in the same narrow product category from other countries. By that reasoning, aggregate imports for each product-country combination form a reasonable basis for setting expectations of demand fluctuations for Chinese exporters.

The definition for demand volatility for this table is analogous to that for Table 3, except that the Q_{jk} used to construct the measure represents each destination's imports from all countries, excluding China. The exclusion implies that expectations of volatility exclude exporters' own sales and are tied to shocks exogenous to the exporting firm.

Table A.15 presents results that are largely similar to Tables A.14 and 3. As in the previous table, the main finding stands: exporters serve destinations with low demand volatility in higher numbers. Destination size is also a statistically significant predictor of exporter numbers between 2000 and 2006, as in the two previous tables.

The differences between Tables A.15 and 3 follow the pattern seen in Table A.14 for the most part. A greater share of the variation in the dependent variable is explained by the right-hand side variables, and more of these have statistical significant coefficients, while retaining the same sign as in Table 3.

In sum, the results are robust to definitions of demand volatility that exclude the history of imports from China, or demand from the rest of the world. High demand volatility is consistently linked to lower numbers of Chinese exporters in foreign destinations.

Table A.15: Exporter Counts and Demand Volatility (Rest of the World):
(Dependent Variable: Log Number of Exporters in Destinations)

VARIABLES	(1) Log(Gross Export Entry)	(2)	(3) Log(Gross Entry Post-2000)	(4)
Demand Volatility	-0.967***	-0.900***	-0.930***	-0.869***
	(0.033)	(0.029)	(0.033)	(0.030)
Destination Size	0.251***	0.216***	0.245***	0.211***
	(0.003)	(0.003)	(0.003)	(0.003)
Log(GDP)	0.262***		0.260***	
	(0.003)		(0.003)	
Log(GDP per capita)	-0.029***		-0.031***	
	(0.003)		(0.003)	
Log(Distance)	-0.562***		-0.538***	
	(0.005)		(0.005)	
Constant	6.351***		6.028***	
	(0.087)		(0.086)	
Observations	331,153	385,885	324,042	377,196
R-squared	0.695	0.736	0.693	0.733
Country-Year FE	N	Y	N	Y
Product FE	Y	Y	Y	Y

Robust standard errors in parentheses
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

The units of observation are destinations: unique HS6-product and country combinations, e.g., Ethiopian imports of women's cotton overcoats (HS 610220). Gross entry is the log of unique firms with recorded exports to a destination between 2000 and 2006. The change in exporter count captures the difference between all firms that served a destination and firms that served the destination in 2000. This difference measures new exporters that appeared with China's trade liberalization from 2001 onwards. Control variables used but not shown in the table include geographic remoteness and dummies for shared borders, common languages and WTO membership. Missing observations in Columns 1 and 3 are because GDP data is not always available. However, estimates on the largest common sample are almost identical to columns 2 and 4.

AppendixA.6. Comparing Demand Volatility with Other Predictors

Table A.16 shows simple OLS regressions of exporter counts on demand volatility, GDP, Distance, and Destination Size. This regression with omitted variables provides only correlations to facilitate comparisons. The correlations suggest that demand volatility is informative for analyses of trade and exporter counts.

Table A.16: Comparing Demand Volatility and Conventional Predictors of Trade (Dependent Variable: Log Number of Exporters)

VARIABLES	(1)	(2)	(3)	(4)	(5)
Demand Volatility	-5.156***				-0.607***
	(0.054)				(0.034)
Log(GDP)		-0.001			0.001
		(0.001)			(0.001)
Log(Distance)			0.002		0.001
			(0.005)		(0.004)
Destination Size				0.431***	0.421***
				(0.003)	(0.003)
Constant	2.286***	1.938***	1.917***	-1.822***	-1.717***
	(0.003)	(0.011)	(0.041)	(0.026)	(0.047)
Observations	371,531	289,801	346,217	371,531	276,459
R-squared	0.287	0.217	0.215	0.544	0.546
Product FE	Y	Y	Y	Y	Y

Robust standard errors in parentheses. Errors clustered by HS6 products.
*** p<0.01, ** p<0.05, * p<0.1

The units of observation are destinations: unique HS6-product and country combinations, e.g., Ethiopian imports of women's cotton overcoats (HS 610220). The dependent variable is the log of the count of unique firms with recorded exports to a product market between 2000 and 2006. Destination size is the value of all imports into each destination in 2000-2006. The other variables follow conventional definitions.

The R^2 values reported in the table alone indicate that demand volatility compares favorably with conventional predictors in explaining the variation in exporter counts. Comparing each variable's column and column 5, which includes all predictors, provides further evidence. The sign and statistical significance of demand volatility remains consistent between columns (1) and (5). Of the other variables, only destination size explains as much of the variation in exporter counts.

AppendixA.7. GDP per Capita and the Volatility of Imports

How does demand volatility as observed by exporters relate to economic development?

Figure A.6 suggests that in the context of trade, import demand volatility declines with increasing GDP per capita. The trend in the graph is downwards, indicating a negative correlation between the volatility of imports and GDP per capita. Not reported in this paper is a graph that shows a strikingly high correlation between the volatility of imports and exports. The small island nations on the right of the graph like Bermuda (BMU), Vanuatu (VUT) and Palau (PLW) also follow the trend, even if it appears shifted. The pattern is comparable to plots of output volatility against GDP per capita in related papers (di Giovanni and Levchenko, 2012b; Koren and Tenreyro, 2007). The vertical axis shows the standard deviation of each country's year-on-year aggregate import growth, while the horizontal axis represents the GDP per capita in 2006.

Economic output is generally more volatile in low-income economies (Koren and Tenreyro, 2007; di Giovanni and Levchenko, 2009; Krishna and Levchenko, 2009). Some pointers on the direction of causation come from di Giovanni and Levchenko (2009), which shows that openness to trade generally increases aggregate output volatility; trade is usually a higher share of output for smaller economies. Krishna and Levchenko (2009) and Koren and Tenreyro (2007) both provide evidence that low-GDP economies tend to specialize in volatile sectors, with the former attributing this tendency to the quality of institutions.

From the perspective of firms importing inputs for industrial processes, the reduction in exporter counts associated with demand volatility matters. Where fewer exporters are available to supply the inputs required for industrial production, the resulting higher prices for those inputs discourage potential growth in the related productive sectors. Several papers show that imports are critical to economic growth and the expansion of exports in developing economies (Goldberg et al., 2010; Bas, 2009; Chevassus-Lozza et al., 2013).

Figure A.6: The Volatility of Imports and GDP per Capita

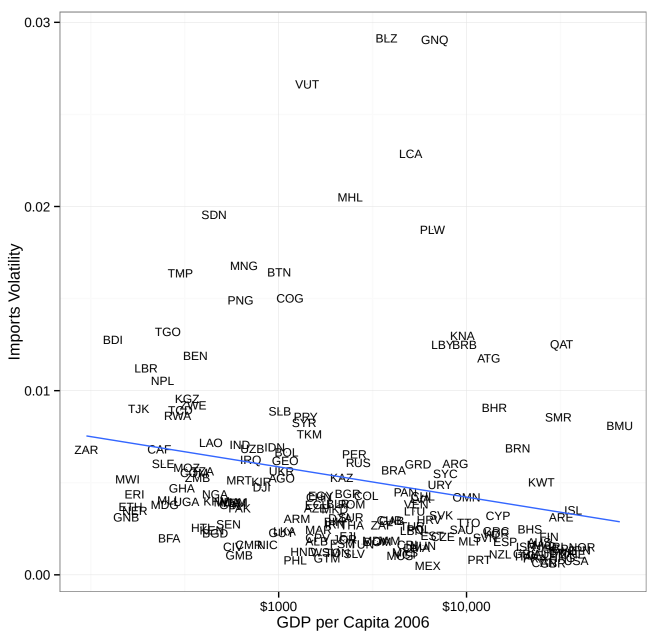

The unit of observation is a country, represented by its 3-letter ISO code. The vertical axis shows the volatility of imports over 1995-2005 and the horizontal axis shows 2006 GDP per capita. Here, the volatility of imports is the sum of the deviations of imports from trend at the country-level, as in equation (18).
Data Sources: UN COMTRADE, CEPII (Head et al., 2010)

AppendixA.8. Demand Volatility and Prices

Table A.17 tests for a relationship between prices and demand volatility separately for each of the UN broad economic categories or BEC goods classifications. The classifications are important for this paper, given how the contributions of each of these imported categories to economic development differ (Jones, 2011). Imports of capital goods have been cited often as a source of productivity growth in developing economies (Eaton and Kortum, 2001; Lee, 1995).

Table A.17: Demand Volatility and Prices, by Broad Economics Categories
(Dependent Variable: Log Unit Prices)

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	Capital Goods		Consumer Goods		Intermediates	
Demand Volatility	0.099**	0.079*	-0.080***	-0.077**	-0.014	0.023
	(0.040)	(0.042)	(0.030)	(0.032)	(0.022)	(0.022)
Prod. Mkt. Size	-0.021***	-0.012***	-0.022***	0.000	-0.013***	-0.005***
	(0.002)	(0.003)	(0.002)	(0.002)	(0.001)	(0.001)
Log (GDP)	0.009***	0.176*	0.027***	-0.022	0.019***	-0.009
	(0.002)	(0.100)	(0.001)	(0.087)	(0.001)	(0.054)
Log(GDP per cap)	0.022***	-0.128	0.027***	0.096	0.023***	0.043
	(0.002)	(0.096)	(0.002)	(0.080)	(0.001)	(0.052)
Log(Distance)	0.009**	-0.604	0.019***	-3.412	0.015***	0.106
	(0.004)	(41.277)	(0.002)	(215.224)	(0.002)	(35.871)
Constant	2.450***		0.033		0.101***	
	(0.054)		(0.039)		(0.027)	
Observations	2,123,105	2,123,105	10,752,290	10,752,290	8,214,698	8,214,698
R-squared	0.964	0.964	0.920	0.920	0.919	0.919
Firm-Product-Year FE	Y	Y	Y	Y	Y	Y
Country FE		Y		Y		Y

Robust standard errors in parentheses. Errors clustered by product-year.
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

The units of observation are unique combinations of firms, HS8 products, destination countries and years, e.g., exports of women's cotton overcoats (HS 61022000) to Ethiopia in 2006 by firm #311996528A. The dependent variable is the log of value divided by quantity for each observation. I used the finer HS8 rather than HS6 product categorization was used because the units for measuring quantity are consistent within HS8 but not HS6 categories. Repeating the regressions for products with only one type of quantity unit within HS6 categories gives substantially similar estimates. The classification into capital, consumer and intermediate goods follows the UN's correspondence between HS6 products and broad economic categories (BEC).

The results indicate that prices are higher for capital goods in destinations with high demand volatility, if only slightly so. This result holds even in the absence of explicit controls for product quality, which could matter for a category with quality differentiation potential like capital goods. Prices are lower on average with high demand volatility for consumer goods, which also tend to have quality differentiation potential but represent a larger share of exports. The estimated effect of volatility on prices is not as clear for intermediate goods and other products that do not fit into any of these three categories (like automobiles). While the findings for these other categories do not conform to the model, they could be explained by related works that imply destinations with low GDP per capita receive lower import prices (Manova and Zhang, 2012; Hummels and Klenow, 2005).

The coefficients of other variables in Table A.17 are consistent with related work on prices in international trade; prices increase with GDP per capita for the specifications that are statistically significant, increase with distance and show mixed effects with GDP.

The firm-product-year fixed effects address several potential concerns with the estimate, including possible changes in HS8 categories from one year to the next (Amiti and Freund, 2010), quality differentiation by firm (Manova and Zhang, 2012) and product differences. That prices vary by product is rudimentary, but when some HS8 products are reclassified to other categories from one year to the next, it is important to include fully-interacted HS8-year effects as a control, even if the reclassifications affect only a small share of product categories. This pair interacted with firms to create fixed effects that reflect consistent differences due to firm productivity, and investments in quality. The country fixed effects address country-specific factors that may consistently affect prices like exchange rates, but are not captured by GDP, distance or GDP per capita.

AppendixB. Model Features

AppendixB.1. The Envelope Theorem Allows Optimal $q^* = E(q)$

Exporters maximize expected profits in the model by working with two parameters: [1] they set prices p, which is equivalent to setting quantities q in monopolistic competition and [2] they set production scale q^*, given that deviations of q from q^* are costly, as defined in equation (2).

The Envelope Theorem justifies my approach of treating this optimization as a one-parameter choice, with the second parameter, in this case q^* fixed at the optimal level. Formally, if profits are a function of both the production scale q^* and prices (which predicts actual quantities sold), the set of optimal profits with respect to prices should be at values of q^* that maximize profits.

Formally, one may define profits as the objective, prices as the parameter that determines profits and the production scale q^* as the maximizer. In optimizing, i.e. setting the derivative equal to zero, the derivative of the profit objective with respect to the production scale equals the partial derivative of profits with respect to prices or quantities, holding the maximizer fixed at its optimal level. Expected profits $E(\Pi)$ is a function of both p and $q*$:

$$\max_p E(\Pi) = \max E(\Pi) \implies \frac{\delta E(\Pi)}{\delta p} = 0|_{\frac{\delta E(\Pi)}{\delta q^*}=0} \tag{B.1}$$

Prices should map one-to-one to quantities, given equation (5), thus one can maximize the preceding equation with respect to q.

In the main body of the paper, I assume production scale will always be set to E(q), and justify the claim with the Envelope Theorem. Here I provide the formal derivation. The optimization exercise fixes quantities and prices for trade in monopolistic competition, with the optimal exporters' production scale q^*:

$$\begin{aligned} \frac{\mathrm{d}E(\Pi)}{\mathrm{d}q^*} &= 0 \\ &= \frac{\mathrm{d}E(\Pi)}{\mathrm{d}E(\text{adjustment costs})} \frac{\mathrm{d}E(\text{adjustment costs})}{\mathrm{d}q^*} \end{aligned} \tag{B.2}$$

From equation (1), $\frac{\mathrm{d}E(\Pi)}{\mathrm{d}E(\text{adjustment costs})} \neq 0$, therefore:

$$0 = \frac{\mathrm{d}E(\text{adjustment costs})}{\mathrm{d}q^*} \tag{B.3}$$

$$\begin{aligned} \frac{\mathrm{d}E(\Pi)}{\mathrm{d}q^*} &= \frac{\mathrm{d}[\gamma \left(\frac{E(q)-q^*}{q^*}\right)^2]}{\mathrm{d}q^*} \\ &\implies q^* = E(q) \end{aligned} \tag{B.4}$$

This leaves us with a one-parameter optimization, as long as the production scale is fixed at expected quantities. (One can generalize equation (B.4) to a set of production scales, one for each period of the exporter's planning horizon).

The Envelope Theorem has also been applied to the analysis of incentive constraints in contract theory and non-convex production problems (Milgrom and Segal, 2002).

AppendixB.2. Export Profits with Asymmetric Adjustment Costs

This section shows that asymmetric adjustment costs do not necessarily change how the model captures exporters' destination choices, although asymmetric adjustment costs may change the zero-profit condition.

The asymmetry in adjustment costs is represented below by a product-specific parameter κ_j. For convenience, I allow adjustment costs to retain the quadratic form in equation (2), but the costs differ for negative and positive shocks if $\kappa \neq 0.5$:

$$\begin{aligned}\text{adjustment costs}_{ijkt} &= \gamma_j \left|\kappa_j - I(q_{ijkt} > q^*_{ijkt})\right| \left[(q_{ijkt} - q^*_{ijkt})/q^*_{ijkt}\right]^2 \\ &= \gamma_j \left|\kappa_j - I(\nu_{ijkt} > 0)\right| \left[(q_{ijkt} - q^*_{ijkt})/q^*_{ijkt}\right]^2\end{aligned}$$

with the product-specific asymmetry factor $\kappa_j \in [0,1]$. The $I()$ term indicates whether a given destination-specific demand shock is positive. $\kappa_j < 0.5$ implies that adjustment costs are disproportionately higher for positive shocks, so that all negative demand shocks are associated with zero adjustment costs at $\kappa_j = 0$. In the same way, $\kappa_j > 0.5$ implies that positive shocks cost relatively less.

The adjustment costs after substituting in equation (6) will be:

$$\text{adjustment costs}_{ijk} = \gamma_j(1-\kappa)(\nu_{jkt})^2|_{(\nu_{ijkt}>0)} + \gamma_j(\nu_{jkt})^2(\kappa)|_{(\nu_{ijkt}\leq 0)}$$

Which gives the following expected profits, the equivalent of equation (7).

$$E(\Pi_{ijk}) = (p_{ijk} - \frac{\tau_{jk}}{\phi_{ij}})E(q_{ijk}) - \frac{\tau_{jk}}{\phi_{ij}}\gamma_j E\left[q_{ijk}(1-\kappa_j)\nu^2_{jk}|_{\nu_{ijkt}>0} + q_{ijk}\kappa_j\nu^2_{jk}|_{\nu_{ijkt}\leq 0}\right] - S_{jk}$$

With $\nu \sim N(0, \sigma^2)$, positive and negative shocks are symmetric around the destination's demand growth trend, which simplifies the expression:

$$E(\Pi_{ijk}) = (p_{ijk} - \frac{\tau_{jk}}{\phi_{ij}})E(q_{ijk}) - \frac{\tau_{jk}}{\phi_{ij}}\gamma_j q^*_{ijk} E\left[(1-\kappa_j)\nu^3_{jk}|_{\nu_{ijkt}>0} + \kappa_j\nu^3_{jk}|_{\nu_{ijkt}\leq 0} + \nu^2_{ijkt}\right] - S_{jk}$$

Going with the assumption that shocks are symmetric, even if costs are not:

$$E(\Pi_{ijk}) = (p_{ijk} - \frac{\tau_{jk}}{\phi_{ij}})E(q_{ijk}) - \frac{\tau_{jk}}{\phi_{ij}}\gamma_j q^*_{ijk} E\left[|(1-2\kappa_j)||\nu^3_{jk}| + \nu^2_{ijkt}\right] - S_{jk}$$

In sum, the expected profits with symmetric shocks and asymmetric adjustment costs should differ, depending on the value of κ_j and demand volatility for a given destination. Nevertheless, the case of symmetric costs remains a useful general guide.

AppendixB.3. Trade Volumes with Demand Volatility

In equilibrium, trade volumes are the integral of firm level sales over the distribution of productivities that meet the threshold ϕ^*_{jk}:

$$X_{ijk} = p_{ijk} q^*_{ijk} = p^{1-\varepsilon}_{ijk} \frac{Q^*_{jk}}{P^{1-\varepsilon}_{jk}}$$

Summing across all firm varieties for destination jk, where $G(.) = 1 - \phi_{ij}^{-\theta_j}$.

$$\begin{aligned}
X_{jk} &= \frac{Q^*_{jk}}{P^{1-\varepsilon}_{jk}} \int_{\phi^*_{jk}}^{\infty} p^{1-\varepsilon}_{ijk} \mathrm{d}G(\phi_{ij}) \qquad \text{(B.5)}\\
&= \frac{Q^*_{jk}}{P^{1-\varepsilon}_{jk}} \int_{\phi^*_{jk}}^{\infty} \left[\frac{\varepsilon}{\varepsilon-1} \frac{\tau_{jk}(1+\gamma_j\sigma^2_{jk})}{\phi_{ij}}\right]^{1-\varepsilon} (\theta_j \phi_{ij}^{-\theta_j-1}) \mathrm{d}\phi_{ij}\\
&= \frac{\theta_j Q^*_{jk}}{P^{1-\varepsilon}_{jk}} \left[\frac{\varepsilon\tau_{jk}(1+\gamma_j\sigma^2_{jk})}{\varepsilon-1}\right]^{1-\varepsilon} \int_{\phi^*_{jk}}^{\infty} (\phi_{ij}^{-\theta_j-1+(\varepsilon-1)}) \mathrm{d}\phi_{ij}\\
&= \frac{\theta_j Q^*_{jk}}{P^{1-\varepsilon}_{jk}} \left[\frac{\varepsilon\tau_{jk}(1+\gamma_j\sigma^2_{jk})}{\varepsilon-1}\right]^{1-\varepsilon} \frac{-\phi_{jk}^{*(\varepsilon-\theta_j-1)}}{\varepsilon-\theta_j-1}
\end{aligned}$$

substituting ϕ^*_{jk} from equation (11):

$$X_{jk} = \frac{-\theta_j \varepsilon S_{jk} \left(\frac{Q^*_{jk}}{\varepsilon S_{jk}}\right)^{\frac{\theta_j}{\varepsilon-1}}}{P_{jk}^{-\theta_j}(\varepsilon-\theta_j-1)} \left[\frac{\varepsilon\tau_{jk}(1+\gamma_j\sigma^2_{jk})}{\varepsilon-1}\right]^{-\theta_j} \qquad \text{(B.6)}$$

Taking logs, (while noting that $\varepsilon - 1 < \theta$, for sales to be finite):

$$ln(X_{jk}) = -\theta_j ln\left[\frac{\varepsilon\tau_{jk}(1+\gamma_j\sigma_{jk}^2)}{\varepsilon - 1}\right] + \frac{\theta_j}{\varepsilon - 1} ln\left(\frac{Q_{jk}^*}{\varepsilon S_{jk}}\right) + ln\left[\frac{-\theta_j\varepsilon S_{jk}}{P_{jk}^{-\theta_j}(\varepsilon - \theta_j - 1)}\right] \tag{B.7}$$

Equation (B.7) implies that trade should decrease with increasing demand volatility. The slope of X with respect to σ^2 in levels and logs should be negative. However, the ε, Q and P, τ and S terms may lead to estimated effects that are smaller than the extensive margin. This pattern is consistent with the assertion in Chaney (2008) that if trade levels change due to changes in costs, the extensive margin dominates.

The intensive margin on the other hand, depends on the productivity distribution. In practice, it should depend on the combination of the productivity slope parameter θ_j and the elasticity of demand ε. For example, exports per exporter $\frac{X_{jk}}{N_{jk}}$ may rise if high demand volatility leads to higher prices that reduce demand, but the slope of the productivity distribution is high enough that the changing export productivity threshold leaves few exporters to meet demand, leading to higher exports per exporter. In principle, it is independent of demand volatility.

Formally;

$$ln(X_{jk}) - ln(N_{jk}) = ln\left[\frac{-\theta_j\varepsilon S_{jk}}{(\varepsilon - \theta_j - 1)N_j}\right]$$

So that:

$$\frac{\mathrm{d}ln\left(\frac{X_{jk}}{N_{jk}}\right)}{\mathrm{d}\sigma_{jk}^2} = 0 \tag{B.8}$$

The relationship in (B.8) is a distinctive feature of the Pareto distribution of productivity (or the class of power laws in general). For these productivity distributions, any change in the exports per exporter that would have resulted from the changes in prices due to demand volatility is perfectly offset by the change in the number of exporters. This requires the usual assumption of large numbers of atomistic exporters. In the firm level data, the average product is associated with 1300 exporters; the median has 400. In sum, deviations from a Pareto productivity distribution and a small pool of potential exporters may skew the findings away from the prediction in (B.8). However, it is still expected that the extensive margin dominates the intensive margin, regardless of the exact nature of the productivity or the size of N_j. Chaney (2008) derived a similar relationship between trade levels and trade costs.

Total exported value to destination jk is the integral over the distribution of productivities of firm-level exports, as given in equation (B.5). The elasticity:

$$\begin{aligned}
\frac{\mathrm{d}lnX_{jk}}{\mathrm{d}ln\sigma_{jk}^2} &= \frac{\mathrm{d}X_{jk}}{\mathrm{d}\sigma_{jk}^2}\frac{\sigma_{jk}^2}{X_{jk}} \\
&= \underbrace{\frac{\int_{\phi_{jk}^*}^{\infty}\frac{\mathrm{d}p_{ijk}^{1-\varepsilon}}{\mathrm{d}\sigma_{jk}^2}\sigma_{jk}^2\mathrm{d}G(\phi_{ij})}{\int_{\phi_{jk}^*}^{\infty}p_{ijk}^{1-\varepsilon}\mathrm{d}G(\phi_{ij})}}_{\text{intensive margin}} - \underbrace{\frac{p_{ijk}^{1-\varepsilon}(\phi_{jk}^*)(1+\gamma_j\sigma_{jk}^2)\frac{\mathrm{d}\phi_{jk}^*}{\mathrm{d}(1+\gamma_j\sigma_{jk}^2)}\mathrm{d}G(\phi_{ij}^*)}{\int_{\phi_{jk}^*}^{\infty}p_{ijk}^{1-\varepsilon}\mathrm{d}G(\phi_{ij})}}_{\text{extensive margin}} \qquad \text{(B.9)} \\
&= \frac{\gamma_j(1-\varepsilon)}{(1+\gamma_j\sigma_{jk}^2)} - \frac{\gamma_j(\theta_j-\varepsilon+1)}{(1+\gamma_j\sigma_{jk}^2)} \qquad \text{(B.10)}
\end{aligned}$$

Because,

$$\frac{\int_{\phi_{jk}^*}^{\infty}\frac{\mathrm{d}p_{ijk}^{1-\varepsilon}}{\mathrm{d}\sigma_{jk}^2}\sigma_{jk}^2\mathrm{d}G(\phi_{ij})}{\int_{\phi_{jk}^*}^{\infty}p_{ijk}^{1-\varepsilon}\mathrm{d}G(\phi_{ij})} = \frac{\gamma_j(1-\varepsilon)}{1+\gamma_j\sigma_{jk}^2}$$

and,

$$\frac{p_{ijk}^{1-\varepsilon}(\phi_{jk}^*)(1+\gamma_j\sigma_{jk}^2)\frac{\mathrm{d}\phi_{jk}^*}{\mathrm{d}(1+\gamma_j\sigma_{jk}^2)}\mathrm{d}G(\phi_{ij}^*)}{\int_{\phi_{jk}^*}^{\infty}p_{ijk}^{1-\varepsilon}\mathrm{d}G(\phi_{ij})} = \frac{\gamma_j(\theta_j-\varepsilon+1)}{(1+\gamma_j\sigma_{jk}^2)}$$

given that

$$\frac{\mathrm{d}\phi_{jk}^*}{\mathrm{d}(1+\sigma_{jk}^2)} = \frac{\phi_{jk}^*}{(1+\gamma_j\sigma_{jk}^2)}$$

and,

$$\frac{\mathrm{d}p_{ijk}^{1-\varepsilon}}{\mathrm{d}\sigma_{jk}^2} = \frac{\gamma_j(1-\varepsilon)}{(1+\gamma_j\sigma_{jk}^2)}p_{ijk}^{1-\varepsilon}$$

Equation (B.10) shows why the extensive margin is prominent.

AppendixB.4. Demand Volatility and Exporter Size Thresholds

One implication of the model is that destinations with higher demand volatility also have higher exporter productivity thresholds.

From equation (11)

$$\phi^*_{jk} = \frac{\varepsilon}{\varepsilon - 1} \frac{\tau_{jk}(1 + \gamma_j \sigma^2_{jk})}{P_{jk}} \left[\frac{\varepsilon S_{jk}}{Q^*_{jk}} \right]^{\frac{1}{\varepsilon - 1}}$$

As demand volatility σ^2 increases, so does ϕ^*. This has two implications: first, that more volatile destinations have fewer exporters, holding other factors constant, and that those exporters on average will be the most productive within their product categories. This complements a similar argument for productivity thresholds and sunk costs in related papers (Helpman et al., 2004; Melitz, 2003).

Figure 3 in the main body of the paper illustrates the idea, derived from this empirical specification:

$$\phi^*_{jk} = \beta^\phi_0 \hat{\sigma}^2_{jk} + \beta^\phi_1 X_{jk} + \alpha^\phi_j + \alpha^\phi_k + \epsilon^\phi_{jk} \tag{B.11}$$

where ϕ_{jk} is the lowest index of productivity for firms in destination jk, the index could be market share or the number of destinations in 2006

X_{jk} = variable(s) representing market size

α = product or country fixed effects

$\hat{\sigma}^2_{jk}$ = demand volatility or demand volatility quantile dummy

Figure B.7 presents a pattern that is similar to the trend observed in Table 3 in the main body of the paper. In the first panel of the figure, we see that the minimum number of destinations served by firms selling goods to a location jk is higher for destinations with high demand volatility. Taking the number of destinations served by a firm as a proxy for productivity follows the finding in Eaton et al. (2004) that the most productive firms serve more export destinations. Note that the number of destinations in this first panel is limited to about 200 countries, as only destinations for the specific product j are considered.

The second panel of Figure B.7 presents a stronger pattern for the relationship between the minimum number of destinations served by firms and demand volatility. The measure of productivity in this panel of the graph uses all destinations served by firms, rather than the number of destinations within a given product. This alternative measure recognizes that many exporters are multi-product firms, so that their productivity within a narrow product category may not fully represent the ability to compete. This is particularly relevant for multi-product exporters. For this measure of productivity, the same pattern of higher productivity thresholds for destinations with high demand volatility is observed.

Figure B.7: Export Thresholds and Demand Volatility

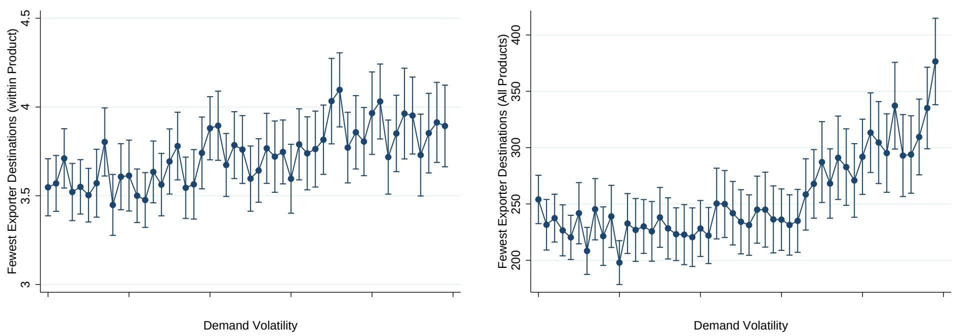

The minimum number of destinations served by exporters in each destination jk. The number of destinations served by an exporter is a proxy for its productivity, so that a destination served by only the most productive firm gets the highest value of this measure. The y-axis averages these threshold counts for all destinations in each quantile. Destinations Grouped Into 50 Quantiles, Least to Largest by Demand Volatility. Data Sources: China GAC Export Data (2000-2006), COMTRADE

The increase in thresholds associated with demand volatility in the right panel of Figure B.7 appears mostly in the upper half of the 50 volatility quantiles used for the plot. The first half of the panel may not indicate a strong relationship because, as the model indicates, the effect of demand volatility on profits is non-linear. Nevertheless, the first order implications of both panels in the graph consistently agree with the paper's main finding.

Table B.18 presents the results of the exercises described in equation (B.11). As expected, columns 2 and 3 of the table show that the threshold for firm productivity, measured by the number of destinations firms serve, is higher in destinations with high demand volatility. I control for destination specific drivers of the exporter threshold with destination size, as used for other estimates in the paper. I also include country and product fixed effects to represent product or country-related factors broadly determine the exporters that can access a foreign destination. While demand volatility does not yield statistically significant predictions if the firm share of output for a product is the assumed measure of productivity, in column 2 the fewest destinations served by firms in a destination increases by a number that is statistically significant at the 90% level. A difference in demand volatility of 0.086, i.e. one standard deviation, should correspond to a 0.05 increase from the mean of 3.8 in the threshold of markets served within each narrow product category. The table also predicts that larger markets are expected to have lower minimum export thresholds, in line with the literature.[39].

Column 3 of Table B.18 completes this exercise, showing that the minimum thresholds for destinations with high demand volatility are higher. Given the mean value of 262 for destination counts and standard deviation of 0.086 for demand volatility, a difference in demand volatility of one standard deviation should correspond to an increase of 30 in the threshold of destinations served.

The increase in ϕ^* associated with demand volatility in (11) implies that the average rank should be higher for destinations with high demand volatility. The higher threshold implies a marginal increase in the average productivity, and size of observed exporters.

[39]Repeating the regression with firm shares multiplied by 1000 to scale up the variable, given its mean of 0.0017, yields 0.493 as the coefficient on σ^2_{jk} - positive but not statistically significant.

Table B.18: Demand Volatility and Export Thresholds
(Dependent Variable: Minimum Market Share and Destination Counts for Firms in jk)

VARIABLES	(1)	(2)	(3)
Demand Volatility	0.000	0.616**	347.170***
	(0.002)	(0.270)	(41.265)
Market Size	0.000	-0.414***	-55.225***
	(0.000)	(0.018)	(2.421)
Observations	285,675	285,661	285,661
R-squared	0.377	0.235	0.086
Country-Year FE	Y	Y	Y
Product FE	Y	Y	Y

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

The units of observation are destinations: unique HS6-product and country combinations. The dependent variable in column 1 is the minimum share of exports within a product category provided by any firm in the destination represented by an observation. Firm share of exports within a product category is a proxy for productivity in this column as in Figure 3. (The average market share is 0.0017, so the zero coefficient is unsurprising). In column 2, it is the minimum number of destinations within the product category served by any exporter with sales to a given destination. In column 3 it is the minimum number of destinations served by any exporter with sales to a given destination. The count of export destinations for the firms are not limited to any products for this column.