

DISCUSSION PAPER SERIES

No. 11013

**SHIPMENT FREQUENCY OF EXPORTERS
AND DEMAND UNCERTAINTY: AN
INVENTORY MANAGEMENT APPROACH**

Gábor Békés, Lionel Fontagné, Balázs Muraközy
and Vincent Vicard

***INTERNATIONAL TRADE AND
REGIONAL ECONOMICS***



SHIPMENT FREQUENCY OF EXPORTERS AND DEMAND UNCERTAINTY: AN INVENTORY MANAGEMENT APPROACH

Gábor Békés, Lionel Fontagné, Balázs Muraközy and Vincent Vicard

Discussion Paper No. 11013
December 2015
Submitted 15 December 2015

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: (44 20) 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **INTERNATIONAL TRADE AND REGIONAL ECONOMICS**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Gábor Békés, Lionel Fontagné, Balázs Muraközy and Vincent Vicard

SHIPMENT FREQUENCY OF EXPORTERS AND DEMAND UNCERTAINTY: AN INVENTORY MANAGEMENT APPROACH[†]

Abstract

This paper studies how exporting firms optimize their inventory management in order to adapt to the uncertainty stemming from demand volatility. We motivate our analysis with a stochastic inventory management framework. We use monthly micro data on French exporters and find that greater uncertainty is associated with lower sales volume. We decompose this effect to its two margins, the number of shipments and average shipment size to find that the number of shipments decreases more quickly as uncertainty increases which is in line with firms adjusting their inventory policy as well as their exported volume as a result of increased uncertainty. Also, uncertainty was found to matter more at distant markets where the uncertainty between firm actions and the arrival of the products is the largest.

JEL Classification: D40, F14 and R41

Keywords: firms, frequency of trade, gravity, inventory model and transport costs

Gábor Békés bekes.gabor@krtk.mta.hu
CERS-HAS Institute of Economics and CEPR

Lionel Fontagné lionel.fontagne@univ-paris1.fr
Paris School of Economics, University Paris 1 and CEPII

Balázs Muraközy murakozy.balazs@krtk.mta.hu
CERS-HAS Institute of Economics

Vincent Vicard vincent.vicard@banque-france.fr
Banque de France

[†] We thank Zsuzsa Holler for excellent research assistance. For comments and suggestions, we are grateful to Carlo Altomonte, Colin Cameron, Jonathan Eaton, Mikos Koren, Thierry Mayer, Adam Szeidl and seminar participants at UC Davis, CEU (Budapest), LMU (Munich), ETH (Zurich), PSE & Sciences Po (Paris), Banque de France (Paris), EEA 2012, ETSG 2013. The authors gratefully acknowledge financial assistance from the European Firms in a Global Economy: Internal policies for external competitiveness (EFIGE), a collaborative project funded by the European Commission's Seventh Framework Programme (contract number 225551), MTA Momentum Programme and MTA Bolyai Grant. Békés thanks the hospitality of CEPII. Fontagné thanks the hospitality of CESifo and acknowledges financial support from the Banque de France. This paper represents the views of the authors and should not be interpreted as reflecting those of the Banque de France.

1. Introduction

Demand volatility matters for business optimizing inventories and supply chains. According to an authoritative industry publication, the Gartner 2014 Supply Chain Study, 45% of supply chain managers mentioned "Forecast accuracy and demand variability", as a major obstacle. Similarly, the 2014 Study of the Chief Supply Chain Officer of SCM World, a leading consulting firm recorded that 70-80% of respondents in a wide variety of industries such as consumer packaged goods, high-tech, chemicals, as well as distribution and retail, are "concerned about demand volatility"¹. The importance of looking at demand uncertainty from a business as usual standpoint is underlined by volatility stemming from characteristics of a modern and globalized economy, including "customer choices, product customization, rapid technological improvements, global competition and upstream supply fluctuations"² as argued by Rajesh Gangadharan of Celuro, a consultancy. These trends not only make demand more volatile, but make its forecast significantly harder, and hence increase uncertainty about demand. Given the centrality of demand volatility in inventory and supply chain management, firms invest significant resources to improving demand prediction and accommodating inherent volatility.

Demand volatility and the uncertainty related to it may be even more pronounced in international commerce as time-to-ship products exacerbate the impact of demand uncertainty on firms' decisions (Coleman, 2009; Steinwender, 2015). Many authors have investigated the effect of uncertainty resulting from the occasional macro-economic shocks, such as currency crises, and political or institutional changes. Demand volatility, which is an integral part of business as usual, has received less attention. The fact that demand volatility seems to be a key concern of logistic experts suggests that it may affect the structure of logistics decisions and, in turn, the cost structure and amount of international trade. All this implies that the role of uncertainty should be studied not only at the time of policy shocks, but also in an equilibrium setting, where firms face an unpredictable fluctuation in demand for their products. The centrality of logistics in determining cost structures also suggests that inventory management models may be key in understanding these effects.

¹www.scdigest.com/ASSETS/FIRSTTHOUGHTS/14-06-26.php?cid=8223 and
www.e2open.com/assets/pdf/papers-and-reports

²www.sdexec.com/article/10289792/supply-chain-strategies-to-manage-volatile-demand

Measuring uncertainty at national level is hard. However, detailed trade data provide a way to test inventory models describing firms reaction to demand volatility because two central margins of inventory policies, shipment frequency and are also observable. In the presence of uncertainty regarding distant buyers, firms adjust their inventory policies by using frequency of delivery as a margin of adjustment. Consequently, from a trade perspective, shipment frequency is another margin of exports worth studying: the question is not whether you export (extensive firm margin), or how much you export (intensive margin), but how often you export, conditional on your foreign sales.³ Such a “transaction margin” has already been observed in the trade literature, although it received limited attention.⁴

In the absence of uncertainty, the optimal policy of exporters would be straightforward. An analogy with the well-known Baumol-Tobin framework provides a good description for determining the optimal combination of the number of shipments and value per shipment for different export markets in a deterministic setting, given the transport technology: the number of shipments and the quantity per shipment increase in proportion to the square root of demand intensity. This question has already been studied in the trade literature. Kropf and Saure (2012) derive a Melitz-type model encapsulating the fixed costs of shipping; they show that higher fixed costs reduce the frequency of shipments and increase the value per shipment.⁵ In Hornok and Koren (2015), consumers have a preference for frequent shipments as timely consumption offers higher utility. In their model per-shipment costs reduce shipment frequency, increase shipment size and the product price.⁶

In this paper we depart from this deterministic case and ask how idiosyncratic demand uncertainty abroad affect the size and frequency of trade flows. This issue

³Firms can also adjust by shipping different sets of products and/or to a different sets of destinations. Iacovone and Javorcik (2010) examine how uncertainty affects trade patterns considering product level dynamics within firms for Mexico. The margin of adjustment here falls on products and uncertainty leads to product churning and limited value for new flows. In these cases, the experience discussed by Araujo and Ornelas (2007) and Albornoz et al. (2012) help explain exporters’ behavior. We will here focus on product-destination decisions: we assume that firms have already chosen their portfolio of exported products to each destination.

⁴When analyzing Colombian transaction-level data, Eaton et al. (2008) show that the distribution of the number of transactions is highly skewed. Ariu (2011) also decomposes trade using the number of transactions using monthly trade data for Belgium and finds the transaction margin to be important at both the firm-level and country level decompositions.

⁵Using Swiss data they estimate per shipment fixed costs at USD 6,500 in 2007.

⁶Also, on explicit modeling of trade technology, see Behrens and Picard (2011) or Kleinert and Spies (2011).

is complex because uncertainty increases logistics costs and makes holding relatively larger inventories optimal in order to reduce back-order costs. Larger logistics costs lead to smaller trade volume. Reduced export volume together with larger optimal inventories both affect trade *frequency* negatively. For shipment size, however, the effect of increased uncertainty is less clear cut: the negative effect of reduced exports may be subdued by the increase in the relative size of shipments required to maintain a relatively larger inventory level. We address these issues based on an extension of Zipkin (2000) – an established text from the area of logistics and inventory optimization – and apply it to a CES demand framework.⁷

Inventory models – standard in the theory of logistics management – have been shown to be useful frameworks when explaining the impact of large demand shocks in the presence of transaction and inventory costs (Alessandria et al., 2011). The models in Alessandria et al. (2010) and Alessandria et al. (2011) were designed to explain time series evidence after large trade shocks. They consider a dynamic version of the decision to be made by the importer as to importing or not. Instead, our simple and tractable approach will reflect the inventory decisions of a firm exporting to many markets in the presence of uncertainty⁸.

In order to study these decisions made by exporters, we use the highly disaggregated nature of monthly export data for individual exporters and consider the frequency of shipments as a new margin of trade. We define this frequency as the number of months within a year in which an international shipment is recorded for a given firm-product-destination. We use data from the French Customs at the level of the individual exporting firm. Importantly for our exercise, the Customs provides monthly firm export data by destination and product category.

We firstly provide evidence on how the frequency of shipments and the value per shipment adjust to different levels of foreign demand. We observe that the frequency of shipments is used by firms to smooth the impact of business conditions on their

⁷We deviate from the literature (Hornok and Koren, 2015; Hummels and Schaur, 2013) by focusing on the supply side and firm level maximization. More precisely, we concentrate on logistics decisions and hence, the cost function of transportation, rather than organizational decisions.

⁸Inspired by the Great Recession, Novy and Taylor (2013) also investigate the role of macro uncertainty on trade volumes. They relate a real option model of stochastic inventory management to the trade reaction model of Bloom (2009) emphasizing the role of imported intermediate inputs. Using monthly US import and industrial production data, they suggest a link between uncertainty and macro-economic cyclicalities.

different markets. As predicted by an analogy with Baumol-Tobin, the adjustment to market size is roughly channeled half through the number of shipment and half through their size.

We then address our focal question concerning the impact of uncertainty on sales, and the number and size of shipments. We rationalize our analysis within a stochastic inventory framework focusing on uncertainty of the demand. Firms will pay per shipment cost to reach their foreign clients, pay a storage cost at destination, and serve clients as they appear. We build this technology into an international trade framework with CES preferences. This framework predicts that: (i) higher uncertainty reduces the export value, the number of shipments and has an ambiguous impact on the average value of shipments; (ii) holding export value fixed, higher uncertainty reduces the number of shipments and increases the average value per shipment; (iii) the effect of uncertainty is magnified by shipment time: uncertainty only matters if transportation take a relatively long time. These predictions are confirmed on our cross-section of detailed firm level export data.

Several additional issues that may be relevant to firms' sales and logistic decisions are excluded from the scope of this paper. When time matters, firms can optimize transportation by choosing between modes of air and maritime cargo (Harrigan, 2010; Hummels and Schaur, 2013).⁹ Evans and Harrigan (2005) argue that an additional adjustment path is location choice: products that need to be served in a timely fashion will be produced closer to destination markets, thus affecting specialization patterns.¹⁰ We consider here that location choices are given, and address firms' strategies conditional on these choices.¹¹

The remaining of the paper is structured as follows. Section 2 provides descriptive statistics on shipment frequency for a cross-section of individual exporters at the product and destination level. Section 3 presents the basic setup and insights from a theory of optimal shipment policy in the presence of uncertainty, while details of the

⁹Indeed, as Hummels and Schaur (2010) demonstrated, uncertainty of demand will affect transport behavior, in the presence of higher demand uncertainty, a greater share of shipments will be by air transport. We will address this issue by restricting our estimations to maritime transport in the robustness test section.

¹⁰Uncertainty is indeed impacting many other dimensions of individual firms decisions like investment in the presence of irreversibility (Bloom and Van Reenen, 2007), in line with the traditional real option argument (Dixit and Pindyck, 1994). We focus here on trade models.

¹¹We use export data for one exporting country, and consider as a robustness maritime routes only.

theoretical model are presented in the Appendix. Testable predictions are confronted to the data in Section 4, where we describe estimation methodology, core results and robustness checks. The last section concludes.

2. Stylized facts

In our dataset from the French Customs for 2007, the unit of observation is monthly export value by a firm i of product j to a destination k ; products are treated at the 6-digit HS level¹². Such data enable us to calculate a proxy for the number of shipments: it will be approximated by the frequency of shipments defined as the number of months within a year in which an international shipment is recorded for a given firm-product-destination. Naturally it is easily possible that a firm sends more than one shipment in one month, in which case our measured frequency will underestimate the true number of shipments. We will return to this issue in our robustness checks.

As we are interested in firm optimization, having firm identifiers as well as product-destination information is crucial. Furthermore, the fact that we can link exports onto different markets and products of the same firms will allow assessing the marginal association with demand, partialling out unobserved firm characteristics. Hence, we can study how the very same firm adjusts to different market conditions.

Figure 1 provides a first look at the distribution of trade frequency¹³. It shows that a large number of firms ship their products only a few months in a year, providing us a large sample and sufficient variation to estimate the determinants of frequency. The relatively low share of observations with 11 or 12 months, where underestimating the true number of shipments is the most likely, and the smooth shape of the function for low frequencies suggest that our frequency variable measures the true number of shipments quite closely especially for frequencies below 10. As a result, in most applications, we will truncate the frequency distribution at 9 and check the importance of this censoring in the robustness section.


We observe a fair number of single shipments: almost 45% of firms ship only once a year to EU markets (left panel of Figure 1) but more than 60% do so to extra-EU markets (resp. right panel). Less than 10% of exporters ship their products every month

¹²We excluded Ships and Aircraft because these items are not exported through usual transport technology but through self-propulsion.

¹³We restrict the sample to incumbent exporters, i.e. firms which exported the same product to the same country in 2006, to exclude censored information from firms which entered during 2007

to their EU markets, as opposed to a smaller fraction towards extra-EU markets.

Figure 1: Frequency of shipments, number of months, 2007, all and extra-EU



Notes: firm-destination-product (HS6) level. Source: French Customs, authors' calculation.

Table 1 provides statistics on the heterogeneity of export frequency. Interestingly, the upper quartile is corresponding to five shipments only: choosing when to export is really a choice to be made by firms. The corresponding strategy has many dimensions worth looking at. Different destinations will be served differently. In EU destinations, the frequency is higher: the median is 3 shipments compared to one for extra-EU trade relationships. This higher frequency can be driven by proximity (authorizing less costly shipments), by market size (large markets can be served more frequently), by type of products exported.¹⁴ Considering market size, we observe that destinations with above the median GDP receive more frequent shipments than destinations below the median. Finally, the largest companies ship their products more often: the 75th percentile of the frequency distribution for large firms is 7 compared to 3 for smaller firms. This observation underlines the importance of controlling for firm characteristics.

¹⁴Composition of exporters would impact the results the other way round (smaller exporters ship less and less frequently).

Table 1: Descriptive statistics of the number of shipments (2007)

	Mean	p25	Median	p75
Total	3.7	1	2	5
Within EU27	4.6	1	3	8
Extra EU27	2.3	1	1	2
Large countries (GDP)	3.8	1	2	6
Small countries (GDP)	1.8	1	1	2
Large firms	4.3	1	2	7
Small Firms	2.8	1	1	3

Let us now interpret frequency and ‘shipment size’ (yearly export value over frequency) as new margins of trade. Using a simple analogy with the Baumol-Tobin model, when demand is deterministic and exogenous, optimal frequency and optimal shipment size should increase in proportion to the square root of the demand. A simple gravity framework can help decomposing the different margins of exports at the firm-product-country level and illustrate how frequency is used by firms to smooth the impact of different business conditions on their different markets. In this simple gravity framework, we consider only market size (GDP of the destination country) and variable trade costs – as proxied by the time to ship from Le Havre in France to the destination market – as explanatory variables.¹⁵

Results are reported in Table 2. As expected, the total value sold by the firm is increasing in GDP and decreasing in time to ship. The number of shipments accounts for the bulk of the impact of trade costs, while shipment frequency and size capture larger demand on a more equal footing: roughly 40 percent of the increase in exports to larger markets channels through the number of shipments. This is in line with the theoretical prediction inspired from Baumol-Tobin, though below the 50 percent predicted by such simple theoretical framework.

The stylized facts presented in this Section brush a simple world where the median firm ships exports on an infrequent basis (twice a year on a given destination for a given product), and the more so on remote destinations (once a year for extra-EU destinations). In a deterministic setting, firms adjust to larger market size on two margins in order to minimize logistics costs by increasing both the number of shipments and shipment size. But this reasoning hardly makes justice to the complexity of decisions

¹⁵Annual GDP data are from the World Bank. Using alternatively great cycle distance in kilometers or total imports – excluding imports from France – as a measure of demand varying across destination-and-product provides qualitatively similar results.

to be made by exporters. To better characterize this optimization problem, we firstly need to model the decision to be made by the exporter in presence of uncertainty. This is done in the next section.

Table 2: Frequency as a trade margin

Dep.Var.	(1) Tot. Value	(2) Nbr shipment	(3) Avg. Value
Log GDP	0.161*** (0.008)	0.069*** (0.003)	0.092*** (0.008)
Log time to ship	-0.045** (0.018)	-0.054*** (0.003)	0.009 -0.018
Firm*product FE	Yes	Yes	Yes
Observations	548,813	548,813	548,813
Number of id	307,622	307,622	307,622
R-squared	0.041	0.043	0.023

Extra-EU exports sample. Shipment frequency truncated for 1-9 months. Robust standard errors in parentheses, clustered by destination-product level. ***, **, * denotes significance at 10%, 5%, 1% level, respectively.

3. Inventory model

Consider the following management problem of an exporting firm: how many Barbie dolls should I ship from my Chinese assembly line to the UK for next Christmas? This question illustrates the role of uncertainty on demand (children might prefer electronic devices this year) aggravated by distance (Guangzhou to Southampton is up to a 30 days sailing route) and inventory costs. Posting new orders in case of underestimation of demand will be very costly (air delivery), while ordering more than demanded means storing unsold dolls. Hence, uncertainty faced by the firm is positively affected by the uncertainty of demand and the time to ship (aggravated potentially by perishability given season specific demand).

This example is indeed an extreme case of a more fundamental optimization problem faced by firms making decisions on the size of likely sales and on the modes and details of how to best serve their clients. In this section, we present the building blocks of an inventory management model tailored for exporting. We discuss key insights, while the model is fully presented in the Appendix. This framework will help us in formulating predictions to be tested in the empirical part of the paper.

In what follows, we will first present the model structure, discuss demand and uncertainty. Then we describe the functioning of the inventory model and offer relevant insights and predictions.

Inventory model for exporters – In this paper, we propose a framework which is suitable for analysing the effect of demand uncertainty in international trade. Firms consider external demand at each of their (product-destination) markets and optimize their shipment process based on available cost information. Looking at exporters has the advantage of identifying variation in demand from rather detailed international trade data. We will see that volatility of demand differs substantially across destination and product markets. Firms face uncertainty regarding their sales in all their markets – an important feature of business.

Note that in this model we investigate a direct exporter, and assume that the firm pays all logistic costs and sells to foreign clients directly from its warehouse in the foreign country. We do not consider this a very strong restriction, however. While many firms work with several buyers and customers, a great deal of firms are supplying overseas firms under an agreement, or even making specific products ordered by foreign partners. Yet, this does not change the overall validity of this setup, because the two parties face the same structure of costs and sources of uncertainty.¹⁶

Model structure – Modeling inventory management of firms is widely studied¹⁷ and there are several frameworks to choose from when studying optimal inventory behavior. In this paper, we build upon a stochastic inventory management model with a continuous revision order-quantity/re-order point, called (r, q) policy. Stochastic demand is important in order to encompass uncertainty, and the inventory management literature has also shown that the “order-quantity/re-order point”, called (r, q) policy is optimal under many conditions and it is frequently used in practice. In an (r, q) policy, firms monitor the fluctuation of sales and reorder goods of worth q should inventory fall below

¹⁶For instance, consider a case when a French company exports a consumption good to the US on orders of a large retail chain. In this case, it may be the retailer who pays the costs and takes the risks but it will make decisions based on the same information set and facing the same demand and cost conditions. As a result, the effect of demand and cost variables on frequency of shipments shall be qualitatively similar to the case where exporters conditions minding final consumers. Of course, transaction prices may be different in the two settings; logistic costs will then be excluded from sales price. Also, the impact of firm characteristics such as financial strength (capacity for lending) may depend upon which partner we consider but this is not central to our case.

¹⁷See for instance the Hadley Within review from 1984, the textbook by Zipkin we rely on, or Frederick S. Hillier, Gerald J. Lieberman (2004) Chapter 18 of Introduction to Operations Research

a certain limit, r . Continuous revision implies that firms consider re-ordering at any point in time. Given the significant shipping times characteristic in international trade using a continuous time approximation does not constrain the generality of the model. Also, we assume that lead-time (the shipping time) is deterministic.

Trade cost – Following the international trade literature, we assume that trade costs are composed of four parts. First, sending a shipment requires k per-shipment cost. Second, holding one unit of the product for one unit of time requires spending h on inventory costs. Third, when firms run out of stocks, back-ordering one item costs b . Fourth there is an iceberg type trade costs. Note that iceberg trade costs are paid independently of the inventory policy decisions.

Demand and uncertainty – A crucial part of the theory is the way we model demand uncertainty. Let us first assume that demanded quantity is exogenous, i.e. prices and quantities are given, and demand during a unit of time is normally distributed, with $N(\lambda, \Psi^2\lambda)$. Importantly, relevant demand will be a product of per time unit average demand, λ and time required for the shipment to arrive, L , while Ψ reflects market-specific demand uncertainty proportional to its mean. The key measure of uncertainty for the firm is the variability of demand between its actions and the arrival of the shipment overseas, a product of L and the markets specific volatility of demand. We will be interpreting L broadly, taking up the time difference between order and delivery. Based on anecdotal evidence, in overseas transport, L typically ranges between 2-12 months as it would not only include actual shipment but administration and with just in time production, assembly as well.¹⁸

Defining the optimal quantity – The aim of this model is to express the optimal size of quantities to be shipped in function of expected demand and uncertainty. In section 6.5 of the Appendix, we will show that the optimal base-stock policy has no closed form solution. However, it will be derived that instead, we can offer an implicit functional form, and calculate functional form of bounds. Indeed, the optimal quantity may be approximated as:

$$\ln q^* = \alpha + \beta_{dem} \ln \lambda + \beta_{unc} \Psi \sqrt{L} \tag{1}$$

¹⁸It may include: notification of partner, discussion of details of delivery, implementation, production – depending on the good this may include assembly, order of parts, production of certain components, retooling, resetting machines – packaging, administration of delivery, delivery from factory gate to harbor, containerizing, shipping, delivery to distribution, check and verification of delivery.

The optimal quantity is an increasing function of demand intensity and decreasing function of the product of shipment time and uncertainty (λ and $\Psi\sqrt{L}$). Importantly, betas are functions of structural parameters (k, h, b) of the model.

Concerning the level of cost, one can show that the optimal cost of general (r, q) policies can be approximated as:

$$C_{(r,q)}^* = \tau\lambda + \left(C_1 + C_2\sqrt{\Psi^2 L} \right) \sqrt{\lambda} \quad (2)$$

Costs are positively related to $\sqrt{\lambda}$ and are also increasing in $(\Psi\sqrt{\lambda L})$. Also, C_1 and C_2 are positive parameters, which are functions of k , h and b . For international trade transactions, we added the iceberg-type cost (τ) to this function.

Insights and predictions – Our model generates four insights. First, there are economies of scale in logistics, owing to the per-shipment cost component: as demand grows, firms send more units in each shipment, generating a declining marginal cost function. In other words, firms can serve large markets with a lower marginal cost than smaller markets. Second, it is optimal for firms to adjust to larger demand on two margins: they send both larger and more frequent shipments. Third, logistics costs increase in demand uncertainty, because firms are more likely to pay back-order costs when demand is more volatile. As a result, the marginal cost curve in a more uncertain market is above that of a more stable market - as illustrated in Figure 2. Moreover, economies of scale are also more important on more uncertain markets. Fourth, as a response to uncertainty, it is optimal for firms to adjust the parameters of their inventory policy, r and q in such a way, that they send less frequent but larger shipments.

Let us highlight the intuition behind the uncertainty-shipment linkage result. In our setting, there is an inventory level (s^*) which would be optimal if the per shipment costs was zero. The expected inventory cost is an increasing function of the distance of the inventory level from this optimal level - i.e. one can define a U-shaped expected cost function of inventory levels with its minimum at s^* . A key insight is that this function becomes less convex when uncertainty is higher because inventory levels will change more anyway during the time the shipment arrives, hence the starting level of inventories matter less (similarly to call options for stocks with different volatility levels).

Let us consider a firm's decision whether to send an extra shipment. Its marginal cost is k independently from the level of demand uncertainty. Its marginal benefit is

that the inventory level will be held more tightly in the vicinity of s^* . But staying close to s^* matters less when volatility is large (and the expected cost function is less convex), so the marginal benefit of sending the extra shipment decreases in volatility. As a result, firms send less frequent shipments under more demand uncertainty.

Optimization across markets – Firms will take into account demand conditions, including uncertainty, when deciding about shipment schemes to each of their markets. Let us now consider optimal decisions of firms facing a downward-sloping demand function. To do this, we embed the downward sloping marginal cost function into a CES demand framework¹⁹. For one particular market, the demand for firm i may be given by the CES demand function in that export market:

$$\lambda_i = \frac{I}{P^{1-\sigma}} p_i^{-\sigma}, \quad (3)$$

where I denotes income, and P the price level in the export market.

To illustrate, consider in Figure 2 the marginal cost functions (MC) approximated earlier, together with a marginal cost function for a firm only facing iceberg-type trade costs.²⁰ We may make several observations. Obviously, marginal costs are higher when logistics costs are taken into account. The marginal cost function is also decreasing in quantity owing to increasing returns in inventory management. Finally, the marginal cost is larger and more convex in the stochastic case thanks to the increasing returns in handling demand volatility.


Let us now compare two markets of different size, as depicted in Figure 3. Firms obviously sell more on the larger market, and economies of scale in logistics makes their output elasticity with respect to market size to be larger than one. This implies that firms will send both more frequent and larger shipments when demand is higher. Note that a farther away market will have higher iceberg trade costs and a smaller effective market size.

Insights and predictions across markets – As we have seen an increase in uncertainty pushes up the marginal cost curve, leading to a fall in exports. As for the components of export quantity - its frequency and its size - we have a strong prediction for frequency

¹⁹Note that modelling dynamic pricing of the firm following realized demand shocks goes beyond the scope of this paper.

²⁰Note that due to uncertainty on demand, the MC curve is above what would be a deterministic setting.

Figure 2: Optimal decisions with different market size




only. In fact, the number of shipments is affected by higher uncertainty through two channels. Smaller quantity implies that firms will send a lower number of shipments and higher uncertainty in itself leads to less frequent shipments even when holding quantity fixed. Hence, these two forces point to the same direction: we expect a negative effect of increased uncertainty on the number of shipments.

This is not case for shipment size. On the one hand, the reduced export quantity predicts a negative effect. Higher uncertainty, on the other hand, leads to larger shipments when holding quantity fixed. Hence, the net effect of the two channels can go either way depending on the cost structure and the elasticity of demand. We can predict, however, that increased uncertainty is associated with larger shipments when we control for quantity.

Finally, as we have discussed, the amount of uncertainty relevant for the firm is the product of uncertainty and shipment time. This predicts an interaction between uncertainty and distance: the effects of uncertainty discussed in the previous paragraphs should be larger in farther away markets.

Figure 3: Optimal decisions with different market size



4. Estimation methodology and results

This section will first present the estimation method and present our approach to measuring uncertainty. Then, we present the core results, followed by robustness tests regarding measurement of uncertainty, sample selection and econometric specification.

Our theoretical framework yields a number of predictions on quantity sold as well as number of shipments. Unfortunately, in our database many firms fail to report quantity (while all report sales values). Furthermore, sometimes firms do only report quantity for some products/destinations (for instance, coverage to shipments to Central-America is rather low.) The probability of non-reporting may well be correlated with our variables of interest: this variable is more likely to be missing for less frequently exporting firms or in smaller markets. In order to handle this potential endogenous selection issue, we will use sales revenue as our main dependent variable rather than quantity.

To mitigate this disconnect, we will control for conditions of demand and competition with both firm-product and country fixed effects. In this cross-sectional data, this should capture a great deal of variation of prices. Importantly, as shown in the Appendix (Table 9), core results do not change importantly when we use quantity as a dependent variable on the smaller sample.

4.1. Estimation method

Our theoretical predictions relate the uncertainty that firms face on specific markets to (i) their shipment frequency (F_{ijk}), and (ii) their average shipment value (q_{ijk}) and (iii) their total (annual) shipment value ($\lambda_{ijk} \equiv F_{ijk}q_{ijk}$). We take logs in each variable and estimate the model at firm-product-destination (i, j, k) level in 2007. The estimating equation is:

$$\ln y_{ijk} = \alpha + \beta_{unc}\Psi_{jk} + \theta_{ij} + \eta_k + \epsilon_{ijk} \quad (4)$$

Where y_{ijk} represents any of the three dependent variables. Given our focus on estimating uncertainty, we will use a set of fixed effects to capture unobserved heterogeneity. In particular, θ_{ij} are firm-product fixed effects controlling for composition effects due to the self-selection of most productive firms into difficult markets (Eaton et al., 2004).²¹ Additionally, unobserved cost characteristics may be related to both firm features such as discount rate or product attributes such as weight. We also include destination specific fixed effects (η_k) that shall pick distance and market size (represented in the model as I_k) as well as costs associated with destination market interest rates, or doing business types of costs. Regarding the identification in this double fixed effect model, we compare the sales of a given firm selling the same (6-digit) product at two markets with different level of demand volatility. For both markets, we look at how sales may deviate from the average sales of French firms to that country²². Standard errors are clustered at the destination-product level²³.

One concern regarding our estimation strategy is that the number of months with positive shipments is a noisy proxy of the number of shipments in a given year. While

²¹Another selection issue arises from the fact that we do not observe all firms/products on all markets and that this selection is not random. The potential negative correlation arising from selection into export markets would however bias our results towards zero. As a robustness test, we follow Crozet et al. (2012) and estimate Equation 4 for $f_{q_{ijk}}$ considering explicitly censoring points depending on markets using the minimum positive value of exports observed on each market. This methodology is however not directly applicable to the breakdown in the shipment and average value margins. As expected, this methodology yields a larger coefficient on uncertainty. Results are available on request.

²²For example, we compare firm A's sales of shirts in Argentina less the average sales by a French firm in Argentina to firm A's sales of shirts in Brazil less the average French firm sales in Brazil. We relate this difference in relative sales to the difference between shirt market uncertainty in Brazil and Argentina.

²³According to Cameron and Miller (2011), it is maybe useful to use two-way clustering at the destination and the product level or the destination and firm level. We tested models accordingly, and it did not change results qualitatively. Results are available on request.

it is a reliable approximation for low frequency exports, it may be biased for high frequency exporters (see Figure 1). To handle this problem, we opted for the simplest and most conservative treatment and excluded the upper tail of the distribution focusing on firm-product-destination observations for 1-9 shipments. In terms of generality, this is not a very serious problem, as this requires dropping only 7.6 percent of observations. However, it implies losing the larger exporters and the ability to identify the role of very frequent shipments. In robustness section 4.4, we show that our results are robust to a direct treatment of censoring using a tobit estimator.

We also exclude within-EU exports. The main reason for doing this is the fact that data on intra- and extra-EU trade are measured differently: different reporting thresholds apply to the collection of export data for EU and non-EU countries which may lead to systematic measurement error in measured frequencies.²⁴ Logistics effects of uncertainty may also be less important within the EU, because transport is much closer to just-in-time there. Finally, excluding EU countries also guarantees that we will disregard neighboring countries to France.²⁵

4.2. *Measuring uncertainty*

From an empirical point of view, an important issue is the *measurement* of uncertainty faced by exporters. In our model of inventory management, stochastic demand makes firms uncertain about their final demand on a market. The relevant dimension of uncertainty we aim at capturing is, therefore, related to sales dynamics over time at the firm level. This measured uncertainty may be related to both variations in overall annual demand on the market and the process of reallocation of market shares across firms.

Our uncertainty variable Ψ_{jk} is measured at the product-destination level. It captures the volatility of sales firms have experienced in the past on a given market: it is the average of firm-level sales variation. To measure it, we use 8 years of data previous to our point of estimation, and look at the variance of annual sales changes of all French firms exporting a given product j to a given destination k . First we calculate the variance of change in (log) sales from one year to another for each firm in the sample, and

²⁴All extra EU export shipments over 1,000 Euros are to be declared to the French Customs whereas for exports to other EU Member states the declaration is compulsory if the yearly cumulated value of exports to the other 26 EU Member states taken together is larger than 150,000 Euros.

²⁵With the exception of Switzerland.

then, take the mean across firms. Formally, consider N French firms (indexed by i) selling on a specific market (j, k) over the period 1999-2006:

$$\Psi_{jk} = \ln \left((1/N) \sum_{i=1}^N \sqrt{\left[1/(n-1) \sum_{t=1}^n \dot{X}_{ijk,t}^2 - \left(1/(n-1) \sum_{t=1}^n \dot{X}_{ijk,t} \right)^2 \right]} \right) \quad (5)$$

where $\dot{X}_{ijk,t}$ is the log change in the value of exports of product j by firm i to destination k over the one-year period $(t-1)$ to t . We use information on changes in firm sales observed at least three times on the same (destination/product) market over 1999-2006. In all regression analysis, we take log of this average.

This log uncertainty measure has a mean of 0.15 and varies substantially across countries (for instance 0.1 in Australia, 0.21 in Brazil, 0.26 in South Africa). Importantly for our exercise, it varies across product markets within countries. It has a standard deviation of 0.37 in our sample (the top 5% reaching 0.67, and lowest 5% -0.48).

We are aware that several studies considered uncertainty, stemming from productivity shocks (Bloom et al., 2012), price volatility (Hummels and Schaur, 2010) or instability of political-institutional variables (Handley and Limao, 2012). Our approach is dedicated to the choices of individual exporters confronted with demand uncertainty is somewhat different, as we are concerned with the volatility of demand itself, and its use will complement existing approaches. Robustness estimations with alternative measures of uncertainty are provided in section 4.4.

4.3. Core results

We now turn to testing the predictions of the inventory model with results shown in Table 3. A firm facing one standard deviation (or 0.37) higher log uncertainty for a given product/destination market tends to sell 1.67% less on average at that market (column (1)). This is a significant economic magnitude: in terms of trade volumes, one standard deviation higher level of uncertainty at a market is equivalent to approximately 37% larger trade costs.

Columns (2) and (3) decompose this elasticity into two margins as suggested by the model. Here, the negative relationship is mainly explained by fewer shipments per year in that destination: shipment frequency is less by 0.34% (column(2)), while shipment

Table 3: Core results

Dep.Var.	(1) Tot. Value	(2) Nbr shipment	(3) Avg. Value	(4) Nbr shipment	(5) Avg. Value
Log uncertainty	-0.042*** (0.009)	-0.034*** (0.004)	-0.008 (0.007)	-0.021*** (0.003)	0.021*** (0.003)
Log tot. value				0.316*** (0.001)	0.684*** (0.001)
Firm*product FE	Yes	Yes	Yes	Yes	Yes
Destination FE	Yes	Yes	Yes	Yes	Yes
Observations	548,813	548,813	548,813	548,813	548,813
Number of id	307,622	307,622	307,622	307,622	307,622
R-squared	0.058	0.057	0.047	0.494	0.811

All regressions include firm-product fixed effects as well as destination fixed effects. Extra-EU exports sample. Shipment frequency truncated for 1-9 months. Robust standard errors in parentheses, clustered by destination-product level. ***, **, * denotes significance at 10%, 5%, 1% level, respectively.

size does not change significantly (column (3)). These results are much in line with our predictions: an increase in uncertainty leads to larger marginal cost, which results in a fall in export value and shipment frequency.

Remember that uncertainty affects average shipment size via two channels. Directly, higher uncertainty increases average package size, but indirectly, it lowers overall sales (λ) and this reduces shipment size, too. The outcome of this depends on model parameters in general. In our case the total effect is close to zero (column(3)).

In columns (4) and (5) we focus on the direct effect only, by running the same regression as before, but this time, controlling for the total value exported by the firm to the given market. This allows one to see whether firms choose a different combination of margins under different levels of uncertainty even when they deliver the same total quantity. Our estimated equation becomes:

$$\ln y_{ijk} = \alpha + \beta_1 \Psi_{jk} + \beta_2 \ln \lambda_{ijk} + \eta_k + \theta_{ij} + \epsilon_{ijk}, \quad (6)$$

which is estimated both for F_{ijk} and q_{ijk} . The results confirm that controlling for the annual total sales, the product-destination markets with higher uncertainty are associated with fewer but larger transactions.

For given value exported, a 10% higher uncertainty is associated with 0.21% larger shipment size (and hence, 0.21% less transactions). This is the estimated value of

the direct impact of uncertainty on optimal logistics choice - i.e. disregarding the indirect channel. This exercise is also an illustration of how firms with larger sales have on average both more frequent shipments (31.6%) and larger sized shipment packages (68.4%), hence the direct channel accounts for about 2/3 of the impact.

Next, we test whether demand volatility matters more for far away markets, as predicted by the stochastic inventory framework where effective uncertainty is the product of time to ship and demand volatility. We test this by interacting the volatility variable with time to ship.

We use data computed by Berman et al. (2012) on the time to ship between any two countries, assuming a speed of 20 knots for maritime transports and 60 knots for shipments by road. Further, we use travel time by maritime (and road for neighboring and landlocked countries) transportation as a proxy for the time gap between sending and delivering goods to final consumers.²⁶ As travel time is hard to measure when air freight is included, we therefore restrict our sample to maritime shipments which represent 52% of our sample for this exercise.²⁷

To test this prediction, we interact in Table 4 our time to ship and volatility variables. Equation (4) becomes:

$$\ln F_{ijk} = \alpha + \beta_1 \Psi_{jk} + \beta_2 \ln L_k * \Psi_{jk} + \beta_3 \ln Value_{ijk} + \eta_k + \theta_{ij} + \epsilon_{ijk}. \quad (7)$$

We estimate two versions of this model. First, we consider time to ship linearly in the interaction, followed by a model introducing a dummy for far away transactions. The parametric model presented in column (1) provides empirical evidence that shipment time magnifies the impact of uncertainty: the number of shipment decreases more with uncertainty when the destination market takes a long time to reach. Alternatively, the non-parametric specification presented in column (2) confirms that the negative impact of uncertainty on the number of shipment falls particularly on destination market whose time to ship is larger than the median.

²⁶Using great circle distance provides quantitatively similar results.

²⁷Data on the mode of transport at the frontier are from Comext, which details the mode of transport of extra-EU trade by destination and HS6 digit level and differentiate between sea, rail, road, air, postal consignment, fixed transport installations, inland water transports or own propulsion. We use the information on the main mode of transport by market (product \times destination) to identify shipments by sea from other modes of transport.

Table 4: The role of travel time

Dep. Var	(1) Nbr ship.	(2) Nbr ship.
Log uncertainty	0.019 (0.014)	-0.006 (0.005)
Time to ship \times log uncertainty	-0.012*** (0.005)	
High time to ship dum. \times log uncertainty		-0.023*** (0.006)
log tot value	0.326*** (0.001)	0.326*** (0.001)
Firm*product FE	Yes	Yes
Destination FE	Yes	Yes
Observations	289,781	289,781
Number of id	165,294	165,294
R-squared	0.499	0.499

Maritime shipment only. All regressions include firm-product fixed effects as well as destination fixed effects. Extra-EU exports sample. Shipment frequency truncated for 1-9 months. Robust standard errors in parentheses, clustered by destination-product level. ***, **, * denotes significance at 10%, 5%, 1% level, respectively.

4.4. Robustness

Our benchmark measure of uncertainty is the average variation in sales faced by French firms in their (j, k) markets. There are several alternative approaches. In this section, we analyze the robustness of our findings using alternative measures of uncertainty, estimation methods and sample restrictions.

Measurement of uncertainty:

In terms of uncertainty measure, we are suggesting three alternatives.

First, firms may look at demand uncertainty from the vantage point of market, rather than firm-level demand fluctuations based on past experience. To capture this, we first consider aggregate uncertainty over time, computed as the relative standard deviation of French quarterly sales (j, k) for 32 quarters (1999-2006). We added zeros to quarters when annual sales that year were non zero and applied a simple seasonal adjustment by calculating quarter dummies as deviations from a trend.

Second, we note that both our benchmark and the aggregate uncertainty variable may be endogenous to the (i, j, k) shipment. To avoid this, we construct a second aggregate uncertainty variable using the volatility, at the (j, k) level, of imports from

countries other than France. Specifically, our second aggregate measure of uncertainty is the standard deviation of growth rates (in log) of imports by (j, k) markets from all countries except France over 2005/2006. Product level import data are taken from BACI. Such measure of aggregate uncertainty in a market is likely to be less sensitive to idiosyncratic shocks to large French firms on narrow product/destination markets. A similar strategy is applied in Autor et al. (2013).

Third, we exploit the history of firms on specific markets and consider a firm's experience in a given market as a measure of uncertainty about its own demand on specific markets. To the extent that local experience helps a firm to learn about its market, it can reduce uncertainty. Our measure of firm's experience in a given market (j, k) is simply the number of years since entry on the export market (1994 being the first available year). Of course this variable captures firm age and overall export experience. However, given our firm-product fixed effect specification, this shall be partialled out. Note that this variable has the opposite expected sign as all other, as a greater number represents more certainty while for other variables, it implies greater uncertainty.

Results are presented in Table 5 for the number of shipments, comparing the effect to the benchmark case (column 1), one by one. Results presented before are confirmed as all uncertainty variables behave similarly to our benchmark.

Estimation and censoring: Here we test the sensitivity of our results to the use of alternative estimators and treatment of the censoring of our dependent variable. For all of this, we focus on the number of shipments only, and results are presented in Table 6. Estimates for the unconstrained model, for average shipment size and total sales may be found in Appendix B.

Column (1) reproduces column (4) from Table 3 to ease comparability. Consider first the issue of censoring. The fact that we have monthly data to proxy frequency has led us to focus on markets with no more than 9 shipments, and hence we now also use the full sample with all transactions. We present results of our preferred specification for alternative estimators and systematically compare with estimations done with the full sample (instead of keeping only the 1-9 shipments). Column (2) shows the baseline results for the whole sample, i.e. also including observations with shipments 10-12. Note that the number of observations is not so different: this restriction involved dropping 4.8% of observations. The estimated coefficient is slightly larger in absolute terms (-0.027 vs -0.021), and this remains true for other specifications. In column (3), as an alternative way to handle censoring, we present results using a panel data Tobit model,

Table 5: Robustness: alternative uncertainty measures

Dep. Variable	(1)	(2)	(3)	(4)
		log nbr shipments		
Log uncertainty	-0.022*** (0.003)			
Log aggregate uncertainty over time		-0.070*** (0.003)		
Log aggregate uncertainty across countries			-0.015*** (0.004)	
Experience by (j, k) market				0.017*** (0.000)
Log tot. value	0.318*** (0.001)	0.317*** (0.001)	0.312*** (0.001)	0.319*** (0.001)
Firm*product FE	Yes	Yes	Yes	Yes
Destination FE	Yes	Yes	Yes	Yes
Observations	504,078	504,078	504,078	504,078
Number of id	287,510	287,510	287,510	287,510
R-squared	0.498	0.500	0.504	0.498

The first column is benchmark index used so far. Log aggregate uncertainty over time is a j, k aggregate volatility over 32 quarters. Log aggregate uncertainty across countries is variation of annual change for all countries but France. Experience is the number of shipment years since 1994. Extra-EU exports and incumbent firms only. Shipment frequency truncated for 1-9 months. Robust standard errors in parentheses, clustered by destination-product level. ***, **, * denotes significance at 10%, 5%, 1% level, respectively.

in which all observations with more than 9 months are treated as censored.²⁸ Once again, the key negative relationship between shipment frequency and uncertainty is confirmed.

Our second concern is potential inconsistency resulting from heteroscedasticity in data. To handle this, we use Poisson pseudo-maximum likelihood estimator proposed by Santos Silva and Tenreyro (2006). This methodology is consistent with average value of shipment estimation and the number of shipments proxied by the number of non-zero monthly exports, at the firm-destination-product level. Poisson PML results with destination fixed effects – presented in column (4) of Table 6 – confirm key results.

Table 6: Robustness: estimators and censoring

	(1)	(2)	(3)	(4)
	log nbr shipments			
Sample	OLS Fixed Effects baseline	All transactions	Tobit RE All transactions	Poisson PML baseline
Log uncertainty	-0.021*** (0.003)	-0.027*** (0.003)	-0.069*** (0.002)	-0.028*** (0.003)
Log tot. value	0.316*** (0.001)	0.335*** (0.001)	0.300*** (0.000)	0.338*** (0.001)
Firm*product FE	Yes	Yes	-	Yes
Destination FE	Yes	Yes	Yes	Yes
Observations	548,813	575,999	575,999	312,865
Number of id	307,622	310,961	310,961	71,674
R-squared	0.494	0.577		

All regressions include firm-product fixed effects as well as destination fixed effects. Extra-EU exports sample. Robust standard errors in parentheses, clustered by destination-product level. ***, **, * denotes significance at 10%, 5%, 1% level, respectively.

Sample restrictions: So far we have focused on a full sample of observations outside of the EU, with the only sample restriction due to censoring. Here, we propose two additional sample reductions for robustness check.

First, a potentially important issue when looking at frequency in a particular year is the impact of entries. New entrants on a market do not necessarily have a full year

²⁸Changing the censoring limit to 8 or 10 months does not change the results importantly. Results are available on request.

of presence (e.g. they can enter in June 2007) which is flawing any attempt of measuring export frequency or even export performance (Berthou and Vicard, 2013). In order to address this issue, in Table 7 we limit the sample to incumbent firms (firms that exported in any month of 2007 a given product to a given destination and were present in 2006 for this product-destination pair). Results remain qualitatively unchanged for total value and shipment frequency. As for the quantity, we now get a small but negative coefficient suggesting that for incumbents, the indirect (sales volume optimization) effect of higher uncertainty slightly outweighs the direct (inventory management reorganization) effect.

Second, Table 8 provides results of all our specifications on the sample of markets mainly served through maritime transportation used in Table 4. This robustness check reflects to the concern that the effect of demand uncertainty may be quite different for maritime and other types of transportation. Here as well we find a significantly negative average shipment size coefficient.

Table 7: Excluding new entrant firms

VARIABLES	(1)	(2)	(3)	(4)	(5)
	Total value	Avg.Value	Nbr shipments	Nbr shipments	Nbr shipments
Log uncertainty	-0.071*** (0.011)	-0.020** (0.008)	-0.052*** (0.005)	-0.028*** (0.004)	0.026 (0.016)
Time x Log uncertainty					-0.017*** (0.005)
Log tot. value				0.339*** (0.001)	0.343*** (0.001)
Observations	313,062	313,062	313,062	313,062	173,599
R-squared	0.063	0.056	0.055	0.522	0.520
Number of id	149,105	149,105	149,105	149,105	82,377

Excludes firm-product-market triads when the firm did not export in 2006. All regressions include firm-product fixed effects as well as destination fixed effects. Robust standard errors in parentheses, clustered by destination-product level. ***, **, * denotes significance at 10%, 5%, 1% level, respectively.

Considering all these robustness checks, we can see that point estimates of the uncertainty elasticity of shipment frequency to vary between -0.021 (our baseline estimate) and -0.069 suggesting that presenting results on the non-EU, truncated sample with no additional restrictions and with fixed effects is a rather conservative approach.

Table 8: Only maritime transport

	(1)	(2)	(3)	(4)
VARIABLES	Total value	Avg. Value	Nbr shipments	Nbr shipments
Log uncertainty	-0.058*** (0.010)	-0.022*** (0.008)	-0.036*** (0.005)	-0.017*** (0.003)
Log tot. value				0.326*** (0.001)
Observations	289,781	289,781	289,781	289,781
R-squared	0.064	0.059	0.051	0.499
Number of id	165,294	165,294	165,294	165,294

Includes firms and destination-product markets when there is evidence on maritime transport use. All regressions include firm-product fixed effects as well as destination fixed effects. Robust standard errors in parentheses, clustered by destination-product level. ***, **, * denotes significance at 10%, 5%, 1% level, respectively.

5. Conclusion

Commerce is always confronting firms with uncertainty requiring planning and operational optimization to use optimal logistics. Selling on foreign markets exacerbates uncertainty as it takes time – often months – to re-optimize inventories following fluctuations in demand. This paper focuses on the exporting firms problem aiming to understand how can firms adapt to volatility. Markets with high uncertainty tend to create more transaction costs for firms. This can be mitigated by flexibly adjusting inventory management policies: altering the reorder point and shipment frequency. Microeconomic decisions, made by firms on the number of shipments, may therefore dampen the impact of uncertainty on aggregate trade flows. The purpose of this paper was to propose a simple modeling framework of such decisions and to bring it to the data.

To this end, we took a stochastic inventory management model, directly derived the logistics cost function from this model and included it to a CES framework. In this framework the marginal cost function is larger and steeper for more uncertain markets, hence expected sales is a direct (decreasing) function of demand uncertainty. The model helped us understand that shipment size will be affected by two opposing channels, an indirect one (via lower sales in more uncertain markets) and a direct one (under higher uncertainty, it is optimal to shop the same quantity in larger shipments) effect. The indirect channel is expected to be negative – higher uncertainty is translated into higher expected optimization costs and hence, lower expected sales. The direct

channel is expected to be positive firms will hope to have larger batches at site to serve unexpected demand. At the same time, shipment frequency is lowered by both channels in the case of higher uncertainty. Importantly, as the cost of not being able to serve a customer is a function of delivery time, uncertainty matters only for shipments to distant markets.

Key model predictions were confirmed by evidence using French monthly trade data for 2007. In particular, shipment frequency was proxied by the number of monthly non-zero shipments. Our empirical analysis confirms that firms behavior varies by market conditions in terms of uncertainty. For markets with 10% higher demand uncertainty, firms on average reduce their annual sales by 0.4% and cut the number of their shipments by 0.3%. Finally, we predict and observe that decreasing time to ship increases more the number of shipments and total exports to more distant and more uncertain markets.

Modeling the effect of uncertainty with logistics models shows an additional and very intuitive channel through which uncertainty affects trade: under larger uncertainty firms have to hold larger inventories, which increase their marginal cost in a non-linear way. We have also shown, that the effective uncertainty faced by firms is the product of time to ship and demand volatility. This framework suggests a number of macro and policy consequences. First, this framework provides an additional factor when considering trade costs: time-to-ship can increase logistics costs to a much larger degree than its direct impact of shipping costs. This suggests an additional advantage of reducing time-to-ship either by technological progress or policy interventions (doing business etc): with shorter shipping times firms will also face lower effective uncertainty, which further reduces their marginal cost and increases trade volume. Second, these advances may have a heterogeneous effect for different countries: as the added cost of inventory management is most likely to be largest for distant and small countries with volatile demand, such countries will benefit the most from reducing trading times when logistics costs are considered. For instance, some of such countries may gain better chances to integrate into global value chains. Third, our framework shows that besides direct shipping costs per-shipment cost and inventory costs may also be important in determining trade. Hence policies and technological progress which affects these (reducing the cost of smaller shipments, providing low-cost and high quality logistics centers) should also have trade-enhancing effects.

6. Appendix A: Theoretical framework

We now provide a more detailed description of our theoretical framework. First we present key features of the inventory management model we chose and explain how it will be applied for exports. This is followed by a detailed presentation of the model setup, inventory management behavior and firm optimization. Before introducing the (r, q) inventory policies, we start with the special case of base-stock policy in subsection A:basestock. In subsections and 6.6, we show the logic of optimization in the general (r, q) policy setup building on the previous results about base-stock policies. Once optimal quantity and shipment size is determined, we extend this by allowing for endogenous demand - shaped by local market conditions and shipment reactions.

6.1. Key features

The theoretical model builds on Chapter 6 of Zipkin (2000), an established textbook on inventory management. We gather important features, simplify the presentation, refocus the model on the problem at hand and extend it to a CES demand system. Streamlining and extending the model leads to a simple framework used for the first time – to the best of our knowledge – in the framework of international trade.

As discussed in the main text, we will use a “world-driven”, stochastic demand model with a deterministic leadtime. In a “world-driven” stochastic demand model, demand occurs independently once every unit of time and the state of the world will partly determine demand. The state of the world includes anything relevant from macro-economic conditions to industry shocks. This initial approach assumes that demand fluctuation is external to the given firm’s optimization framework: its actions will not affect the distribution of possible demand in the next period and beyond.

We also assume that firms follow an (r, q) inventory policy: it will reorder goods of worth q should inventory fall below a certain limit, r .

Out of several options, we picked this setup for several reasons. First, stochastic demand is crucial for our purposes, as we are mainly interested in the effect of demand volatility on exporting. Uncertainty has a structure that resonates exporters’ problem: demand is a product of some basic feature (such as market structure, institutions, technology) and market size, as featured in world-driven systems. Second, the model captures the fact that it takes time to export, and hence adds a time dimension to

uncertainty.²⁹ Third, while we'll use the normal approximation of the MCDC process, we know that this structure is flexible enough to include several scenarios, including a model where importers also optimize their inventories.

Note finally that this model assumes independent inventory management process for multi-product firms and hence, allows determining rules for single products only. Under – a fairly realistic – assumption of product-specific production, shipment and distribution, this may be an acceptable simplification.

6.2. *Exporters facing uncertainty*

In the next step, we embed such steady state inventory behavior into a framework with CES demand on foreign markets. Firms will consider external demand at each of their (product-destination) markets and optimize their shipment process based on available cost information. This is much in line with our empirical framework, which has the advantage of identifying variation in demand from rather detailed international trade data.

We assume that demand during a unit of time is normally distributed, with $N(\lambda, \Psi^2\lambda)$. Variance is written as a product of market size and Ψ , which represents the variability of demand for the product-country combination relative to its mean. This is a useful abstraction as it allows considering it as personal demand aggregated for countries. Relevant demand is a product of per time unit average demand, λ and time required for the shipment to arrive, L .³⁰

We assume that trade costs are composed of four parts. First, sending a shipment requires k per-shipment cost. Second, holding one unit of the product for one unit of time requires spending h on inventory costs. Third, when firms run out of stocks, backordering one item costs b . Fourth, there are iceberg type trade costs, to be paid independently of the inventory policy decisions.

To follow the inventory behavior, we will define some stock variables. First, $I(t)$ will be the inventory on hand at time t . Second, $B(t)$ will denote outstanding backorders.

²⁹The model does not include stochastic leadtime, i.e. random variation in transportation time. Using a stochastic leadtime model would complicate things a lot, while having modest benefit, as we cannot observe individual leadtimes, only proxy time spent in transportation by the distance between France and the destination countries

³⁰This can be interpreted as an approximation of a Poisson or, more generally, a Markov-chain-driven counting (MCDC) process. Zipkin (2000) shows (p206) that the Normal approximation in our context works rather well

Third, $IO(t)$ will represent inventory on order, the number of units which were ordered, but have not arrived.

The total average cost in the long run depends on the average order frequency, inventory stock and backorder stock per unit of time. These will be denoted by \bar{F} , \bar{I} and \bar{B} , respectively. Under these assumptions, average cost per period (C) will be:

$$C = k\bar{F} + h\bar{I} + b\bar{B} \quad (8)$$

According to the (r, q) inventory policy, the firm orders q units of the product whenever its inventory level fall below r ³¹.

6.3. Base-stock policy: setup

In order to approach the general (r, q) inventory policies, we start with the special case of base-stock policy. The definition of the base-stock policy is that the firm can order infinitesimally small quantities, $q = \epsilon$ at any time t . To analyze total cost, we assume that the stock level in time $t = 0$ is exactly $s = r + \epsilon > 0$ hence $I(0) = B(0) = s$. Consider any $t > L$. The number of outstanding orders only includes orders between $t - L$ and t . The logic of base policy also means, that each of these orders corresponds to a demand, hence the stock of outstanding orders equals the leadtime demand $D(t - L, t)$ during this period:

$$IO(t) = D(t - L, t) \quad (9)$$

Consider now the net inventory position, $I(t) - B(t)$. In the base-stock policy, by definition, the sum of the net inventory position and that of outstanding orders is always s .

$$I(t) - B(t) = s - IO(t) \quad (10)$$

As this is just a linear transformation of $IO(t)$, it also has a normal limiting distribution, with mean s and variance of $\Psi^2\lambda L$.

The normal distribution of the net inventory position allows us calculating the prob-

³¹Throughout this Appendix, we will use the following notation: λ : demand rate; L : Lead time; Ψ : long-run variance to mean ratio; k : per-shipment cost, b : backorder cost; h : holding cost; $\omega = b/(b+h)$; c : variable production cost; τ : variable transportation cost; r : reorder point; q : batch size; F : frequency, $f = \lambda/q$; C : Average cost per period; \bar{F} : Average order frequency, \bar{I} : Average inventory stock; \bar{B} : Average backorder stock; s : $r + 1$, base-stock level; $I(t)$: inventory level at time t ; $B(t)$: number of outstanding backorders at time t ; $IO(t)$: Inventory on order; $D(t)$: cummulative demand through time t .

ability of a positive inventory level and the expected backorder and inventory positions. The probability that the inventory level is positive:

$$P(I(t) \geq 0) = \int_s^\infty \frac{1}{\Psi\sqrt{\lambda L}} \phi\left(\frac{x-\lambda}{\Psi\sqrt{\lambda L}}\right) dx = 1 - \Phi\left(\frac{s-\lambda}{\Psi\sqrt{\lambda L}}\right), \quad (11)$$

where ϕ is the standard normal probability density function, and Φ is the c.d.f. of the standard normal distribution.

Given this probability, one can calculate the expected inventory level and backorder cost:

$$\bar{B}(s) = \int_s^\infty (x-s) \frac{1}{\Psi\sqrt{\lambda L}} \phi\left(\frac{x-\lambda}{\Psi\sqrt{\lambda L}}\right) dx = (s-\lambda) \Phi\left(\frac{s-\lambda}{\Psi\sqrt{\lambda L}}\right) + \Psi\sqrt{\lambda L} \phi\left(\frac{s-\lambda}{\Psi\sqrt{\lambda L}}\right) - (s-\lambda) \quad (12)$$

$$\bar{I}(s) = s - \lambda + \bar{B} = (\lambda - s) \Phi\left(\frac{\lambda - s}{\Psi\sqrt{\lambda L}}\right) + \Psi\sqrt{\lambda L} \phi\left(\frac{\lambda - s}{\Psi\sqrt{\lambda L}}\right) - (\lambda - s) \quad (13)$$

One may simplify the notation by introducing the standard loss function, $\Phi^1\left(\frac{s-\lambda}{\Psi\sqrt{\lambda L}}\right) = \int_s^\infty \left(x - \frac{s-\lambda}{\Psi\sqrt{\lambda L}}\right) \phi(x) dx$ and get:

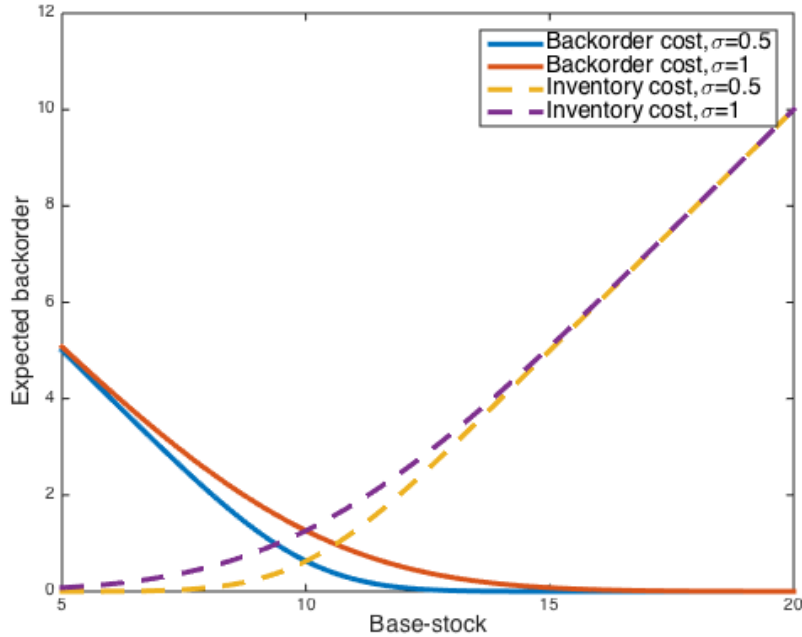
$$\bar{B}(s) = \Phi^1\left(\frac{s-\lambda}{\Psi\sqrt{\lambda L}}\right) \Psi\sqrt{\lambda L} \quad (14)$$

$$\bar{I}(s) = \Phi^1\left(-\frac{s-\lambda}{\Psi\sqrt{\lambda L}}\right) \Psi\sqrt{\lambda L} \quad (15)$$

Figure 4 shows how these functions behave under different conditions. The expected stock of backorders increases while that of inventories decreases nonlinearly as s increases. Also, larger variance is associated with larger expected backorders as well as larger inventory costs for a given s . Finally, both functions are convex.

Another important property is that both performance measures are affected by L and Ψ through their effect on variance of demand, which, in turn, is a multiplicative function of the two variables. The effective uncertainty a firm faces is the product of leadtime and volatility of demand. Hence volatile demand plays an important role only in far-away countries.

Figure 4: Expected backorder and Inventory costs



Notes: In the figure σ is a shorthand for $\Psi\sqrt{\lambda}$. The parameters used for the figure are $\lambda = 10$, $h = 1$, $b = 1.5$, $k = 1.5$.

6.4. General (r, q) policies

The (r, q) policy can be interpreted as a generalized case of the base-stock policy, where q may take any value. The main insight when considering these policies is that their performance measures are the unweighted averages all the base-stock policies with $s \in [r, r + q]$.

To gain some intuition for this statement, imagine what happens in very short time intervals. We are interested in the behavior of the inventory position, $I(t) - B(t) + IO(t)$. As long as $I(t) - B(t) + IO(t) \geq r$, the new demand is satisfied from the inventory, while it jumps instantly to $r + q$ as soon as it falls below r . Hence, we can imagine the state space as a circle with a perimeter of q , going from r to $r + q$. The inventory position always moves clockwise on this cycle according to the normal distribution of the demand. As the magnitude of this movement does not depend on our position on the circle, the inventory position can take any value between r and $r + q$ with equal probability so its limiting distribution is uniform on this interval. Hence, this is a

continuous-time Markov-chain³².

It is also intuitive to show, that for a fixed t , the random variables $IP(t)$ and $D(t, t+L)$ are independent, because $IP(t)$ only depends on the demand process realized before t .

Third, we can generalize equation (10) for this case. For all $t \geq 0$, the difference of inventory and backorder stocks (net inventory) is the sum of the realized demand in the previous L time and the inventory position L periods ago:

$$I(t) - B(t) = I(t - L) - B(t - L) + IO(t - L) - D(t - L, t] \quad (16)$$

This formula provides a description of the distribution of $I(t) - B(t)$, because $I(t - L) - B(t - L) + IO(t - L)$ is distributed uniformly and $D(t - L, t]$ has a normal distribution, and the two are independent from each other³³.

Thanks to the uniform distribution of $I(t - L) - B(t - L) + IO(t - L)$, the distribution of each performance measure can be expressed as the unweighted integral of the performance measure of base-stock policies with $s \in [r, r + q]$. In particular,

$$\bar{B}(r, q) = \int_r^{r+q} \bar{B}(s) ds = \int_r^{r+q} \Phi^1 \left(\frac{s - \lambda}{\Psi \sqrt{\lambda L}} \right) \Psi \sqrt{\lambda L} ds \quad (17)$$

Similarly,

$$\bar{I}(r, q) = \int_r^{r+q} \bar{I}(s) ds = \int_r^{r+q} \Phi^1 \left(\frac{\lambda - s}{\Psi \sqrt{\lambda L}} \right) \Psi \sqrt{\lambda L} ds \quad (18)$$

Finally, frequency, the expected number of shipments per time unit is the fraction of the mean demand and shipment size, which is the expression used in the core of the text:

$$\bar{F} = \frac{\lambda}{q}$$

6.5. Optimal base-stock policy

In this subsection we first show the optimal inventory policy for the base-stock model followed by showing the bounds for the optimal solution in the general (r, q) case

³²A more formal proof is provided by Zipkin (2000), p. 193.

³³A more formal derivation is given by Zipkin (2000), p. 195.

where no closed form solution exists.

Under the base-stock model, the firm minimizes:

$$C(s) = h\bar{I}(s) + b\bar{B}(s) = \left[h\Phi^1\left(-\frac{s-\lambda}{\Psi\sqrt{\lambda L}}\right) + b\Phi^1\left(\frac{s-\lambda}{\Psi\sqrt{\lambda L}}\right) \right] \Psi\sqrt{\lambda L} \quad (19)$$

Two such functions with different uncertainty levels are illustrated in Figure 6. Importantly, the partial derivative of the cost function with respect to variance of demand is a bell-shaped curve:

$$\frac{\partial C}{\partial(\Psi\sqrt{\lambda L})} = (h+b)\phi\left(\frac{s-\lambda}{\Psi\sqrt{\lambda L}}\right) \quad (20)$$

Logistic costs can be minimized simply by differentiating the cost function with respect to s . The solution for the optimal base stock level can be expressed with using the inverse normal c.d.f, $\Phi^{-1}()$:

$$s^* = \lambda + \Phi^{-1}(\omega) \Psi\sqrt{\lambda L} \quad (21)$$

Let us turn now to the characterization of the optimal policy, $q(r, q)^*$, in the general (r, q) case. The optimal q is bounded by two functions³⁴. First, the optimal q is larger than the optimal q when uncertainty is zero and positive backorders are allowed ($\sqrt{\frac{2k\lambda}{\omega h}}$). Second, it has an upper bound which is the sum of the optimal q in the deterministic case and the optimal cost in the base-stock policy divided by $\frac{bh}{b+h}$. In particular:

$$\sqrt{\frac{2k\lambda}{\omega h}} \leq q(r, q)^* \leq \sqrt{\frac{2k\lambda}{\omega h}} + \frac{(b+h)}{\omega h} \phi(1 - \Phi^{-1}(\omega)) \Psi\sqrt{\lambda L} \quad (22)$$

Taking logs and using the formula $\ln 1 + x \leq x$ yields:

$$\frac{1}{2} \ln \frac{2k}{\omega h} + \frac{1}{2} \ln \lambda \leq \ln q(r, q)^* \leq \frac{1}{2} \ln \frac{2k}{\omega h} + \frac{1}{2} \ln \lambda + \frac{(b+h)^{\frac{1}{2}}}{2\omega h k} \phi(1 - \Phi^{-1}(\omega)) \Psi\sqrt{L} \quad (23)$$

The optimal cost of general (r, q) policies is also bounded by similar functions. bounded by two functions, depending on the optimal cost in the deterministic case ($C_D^* = \sqrt{2kh\omega\lambda}$) and the optimal cost of the base-stock policy³⁵:

³⁴See Zipkin (2000) (p. 218)

³⁵This theorem is proved in Zipkin (2000) pp. 218-221.

$$C_s^* = (b + h) \phi([1 - \Phi]^{-1}(\omega)) \Psi \sqrt{\lambda L} \quad (24)$$

In particular, the optimal cost of the (r, q) policy, $C_{(r,q)}^*$, is bounded by:

$$\sqrt{2kh\omega\lambda + C_s^{*2}} \leq C_{(r,q)}^* \leq \sqrt{2kh\omega\lambda} + C_s^* \quad (25)$$

Again, as both bounds increase in C_s^* which, in turn, increases both in λ and $\Psi\sqrt{L}$, we will assume that the cost function of the firm behaves this way.

6.6. Optimization with an (r, q) model

In this subsection, we will show the logic of optimization in the general (r, q) policy setup building on the previous results about base-stock policies. In the general case, $k > 0$ so q is not infinitesimal. While there is no closed-form solution for these general policies, one can characterize them effectively by relying on the fact that their performance measures can be calculated as an unweighted average of all base-stock policies between r and q . Let us denote the cost of an (r, q) policy with $C(r, q)$. This function takes the following form:

$$C(r, q) = \frac{k\lambda + \int_r^{r+q} C(s) ds}{q}, \quad (26)$$

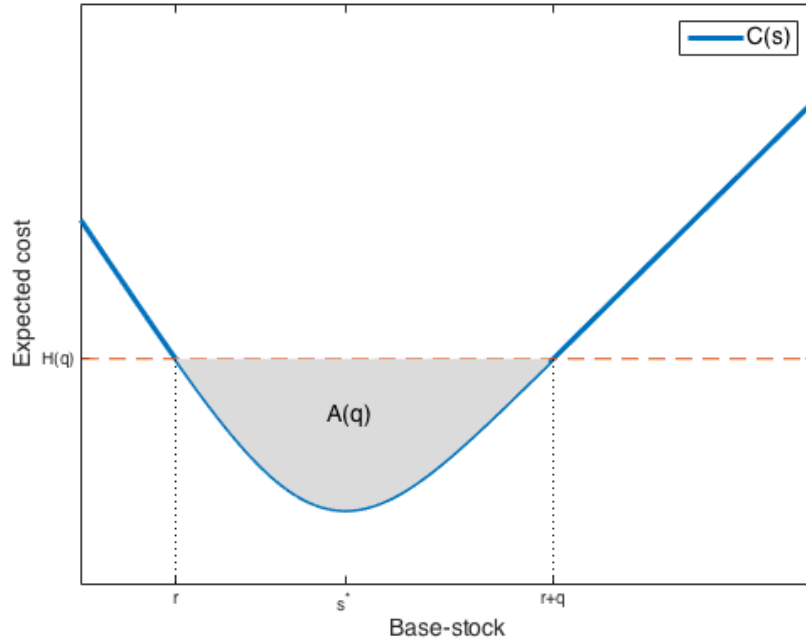
hence, the cost is the sum of the per-shipment costs and the average cost of corresponding base-stock policies.

The optimal q of such a policy can be illustrated in Figure 5. The $C(s)$ function shows the cost function of the base-stock policies for different values of s . For any optimal q , the cost of the base-stock policies corresponding to r and $q + r$ should be the same, which is represented by $H(q)$. Let us define $A(q)$ as the shaded area between the U-shaped curve and the horizontal line defined by $H(q)$. $A(q)$ is given by:

$$A(q) = qC(r + q) - \int_r^{r+q} C(s) ds. \quad (27)$$

The graph helps understand how we find the optimal value of shipment size, q . Let us assume, that we send just one less shipment. The marginal benefits of this is k for each unit of demand λ . The marginal cost of sending one shipment less is that the q units in it will be distributed across the other shipments and the logistic cost on each of these units will be approximately $H(q)$. Integrating over these units (from r to $r + q$)

Figure 5: Optimality condition for a general (r, q) policy



Notes: The graph depicts an inventory cost functions for various base-stock levels, s as well as the optimal reorder point, r and the optimal shipment size, q . The parameters used for the figure are $\lambda = 10$, $h = 1$, $b = 1.5$, $k = 1.5$, $\Psi\sqrt{\lambda} = 0.5$. The shaded area, $A(q)$ illustrates the marginal cost of sending one less shipment.


the cost difference, $H(q) - C(s)$, is equal to $A(q)$. As a result, in optimum:

$$A(q) = k\lambda \quad (28)$$

Understanding the process, we can now illustrate the impact of uncertainty, i.e. how the optimal shipment size, q is affected by an increase in the variance of demand, $\Psi\sqrt{\lambda L}$. Let us assume that volatility increases by a small amount, ε . The original optimal quantity will be denoted by q^1 . Two functions with different volatilities and the optimal choices are shown on the same Figure 6, the dotted line curve corresponding to higher variance of demand. Given that $C(s)$ is always convex, the cost function corresponding to larger volatility will be 'less convex'³⁶. As a result, using q^1 when

³⁶The partial derivative of $C(s)$ w.r.t the variance is $(h+b)\phi\left(\frac{\lambda-s}{\Psi\sqrt{\lambda L}}\right)$. Hence the difference between the $C(s)$ and $C'(s)$ can be approximated by the bell-curve shaped $\varepsilon(h+b)\phi\left(\frac{\lambda-s}{\Psi\sqrt{\lambda L}}\right)$ where ϕ denotes the standard normal density function.

Figure 6: Inventory cost functions



Notes: The graph depicts the optimal (r, q) policy for two uncertainty levels. When uncertainty is lower, $\Psi\sqrt{\lambda} = 0.5$ the optimal policy is represented by q^1 and r^1 . When uncertainty is high, $\Psi\sqrt{\lambda} = 1$, the optimal choices are q^2 and r^2 . The parameters used for the figure are $\lambda = 10$, $h = 1$, $b = 1.5$, $k = 1.5$.

volatility is increased by ε cannot be optimal, because $A(q)$ will be smaller for this less convex function – the new optimum will be q^2 in our figure. Hence, the optimal q should be increasing in the variance; larger $\Psi\sqrt{\lambda L}$ is associated with larger and less frequent shipments.

As noted earlier, there is no closed form solution to the optimal shipment size. However, we can get a bit more precise picture about the behavior of the optimal quantity $q(r, q)^*$ by using results about its bounds (equation 23).

The bounds show that optimal q is a function of technology parameters (k, ω, h, b) as well as two of our key variables, λ and $\Psi\sqrt{L}$. Hence we will use an implicit functional form:

$$q_i = f(\lambda, \Psi\sqrt{L}, \bullet) \quad (29)$$

Based on the functional form of the bounds, we can approximate this as:

$$\ln q^* = \alpha + \beta_{dem} \ln \lambda + \beta_{unc} \Psi \sqrt{L} \quad (30)$$

Based on the bounds above, one can derive some predictions for the parameters in this equation. First, it is reasonable to assume that $0 < \beta_{dem} < 1$. Second, the model also predicts that $\beta_{unc} > 0$, hence higher uncertainty increases shipment size.

Furthermore, given equation (2), we can write

$$\ln f^* = -\alpha + (1 - \beta_{dem}) \ln \lambda - \beta_{unc} \Psi^2 L \quad (31)$$

Finally, let us look at the cost function. Concerning the level of cost, one can show that the optimal cost of general (r, q) policies is bounded by two functions, as described earlier. Both bounds increase in the cost of the optimal base-stock policy, C_s^* which, in turn, increases both in λ and $\Psi\sqrt{L}$: we will assume that the cost function of the firm behaves this way. These bounds are increasing proportionally with $\sqrt{\lambda}$ and are also increasing in $(\Psi\sqrt{\lambda L})$. Hence, we can assume that this is the case with the optimal cost of the general (r, q) policy. Based on these bounds one can write up an approximation of the logistics cost function which is increasing in demand intensity and uncertainty. In the next section, in particular, we will use the following approximation:

$$C_{(r,q)}^* = \tau\lambda + \left(C_1 + C_2\sqrt{\Psi^2 L} \right) \sqrt{\lambda} \quad (32)$$

Note that we add the iceberg-type cost to this function. Also, C_1 and C_2 are positive parameters, which are functions of k , h and b . As this function is only an approximation from the two bounds, one cannot write up a closed form for these parameters. Still, this form provides a tractable function for our simple derivations in the next section.

6.7. Extension to an endogenous demand system

Until now, we have assumed that demanded quantity is exogenous. It is, however, an endogenous decision affected by cost structure and exogenous demand parameters. Let us consider one particular market and assume that the demand for firm i is given by the CES demand function in that export market:

$$\lambda_i = \frac{I}{P^{1-\sigma}} p_i^{-\sigma}. \quad (33)$$

Where I denotes income, and P the price level in the export market. Combining

the CES demand function with the cost function previously derived yields the following profit function, $\Pi_i = (p_i - c_i)\lambda_i - C_{(r,q)}^*$:

$$\Pi_i = \frac{I^{\frac{1}{\sigma}}}{P^{\frac{1-\sigma}{\sigma}}} \lambda_i^{\frac{\sigma-1}{\sigma}} - (c + \tau)\lambda_i - C_1 \lambda_i^{\frac{1}{2}} - C_2 \sqrt{\Psi^2 L} \lambda_i^{\frac{1}{2}} \quad (34)$$

Where is c_i the variable cost of production. This function is maximized when:

$$\frac{d\Pi_i}{d\lambda_i} = \frac{\sigma-1}{\sigma} \frac{I^{\frac{1}{\sigma}}}{P^{\frac{1-\sigma}{\sigma}}} \lambda_i^{-\frac{1}{\sigma}} - (c + \tau) - \frac{1}{2} \left(C_1 + C_2 \sqrt{\Psi^2 L} \right) \lambda_i^{-\frac{1}{2}} = 0. \quad (35)$$

This equation once again does not have a closed-form solution, and thus we will only consider it in a implicit functional form:

$$\lambda_i = g(I, \tau, \sqrt{\Psi^2 L}, \bullet) \quad (36)$$

Importantly, one can calculate how optimal quantity is affected by the exogenous variables by implicit function differentiation. In the following paragraphs, we are interested in the effects of market size, I , trade cost, τ and demand volatility.

To fix ideas, note that the marginal cost function is the following:

$$MC_i = c_i + \tau + \frac{1}{2} (C_1 + C_2 \sqrt{\psi^2 L}) \frac{1}{\sqrt{\lambda}} \quad (37)$$

In Figure 3 in the main text, we immediately observe that market size, I is impacting the optimal decision (here there are two markets with different I): taking into account inventory costs leads to lower quantity on both markets than a firm which only pays iceberg transportation costs. Also, the firm facing inventory costs increases its sales to a larger degree as market size increases, thanks to increasing returns to logistics. It is also shown by implicitly differentiating the first order condition:

$$\frac{d\lambda_i}{dI} = \frac{\lambda_i}{I - \frac{1}{4} \frac{\sigma^2}{\sigma-1} \left(C_1 + C_2 \sqrt{\Psi^2 L} \right) \left(\frac{I}{P} \right)^{\frac{\sigma-1}{\sigma}} \lambda_i^{\frac{2-\sigma}{2\sigma}}} \quad (38)$$

Consider now the effect of τ on λ . With implicit differentiation:

$$\frac{d\lambda_i}{d\tau} = - \frac{\lambda_i^{\frac{1+\sigma}{\sigma}}}{\frac{\sigma-1}{\sigma^2} \frac{I^{\frac{1}{\sigma}}}{P^{\frac{1-\sigma}{\sigma}}} - \frac{1}{4} \left(C_1 + C_2 \sqrt{\Psi^2 L} \right) \lambda_i^{\frac{2-\sigma}{2\sigma}}} \quad (39)$$

Again, when inventories are costless ($h = 0 \Rightarrow C_1 = 0$) and if there were no uncertainty, we would get back the usual result for the CES case with iceberg trade costs.³⁷ In presence of uncertainty, the predictions of the model with respect to market size and τ are unchanged, but the nonlinearities are further magnified by uncertainty. Expressing the effect of uncertainty by implicit differentiation yields:

$$\frac{d\lambda}{d\sqrt{\psi^2 L}} = \frac{\frac{1}{2}C_2\lambda}{\frac{1-\sigma}{\sigma^2} \frac{I^{\frac{1}{\sigma}}}{P^{\frac{1-\sigma}{\sigma}}} \lambda^{\frac{\sigma-2}{2\sigma}} + \frac{1}{4}(C_1 + C_2\sqrt{\psi^2 L})} < 0 \quad (40)$$

Hence, uncertainty reduces quantity sold, and the effect depends on the elasticity of demand and the slope of the marginal cost curve. When analyzing the effect of uncertainty on shipment size and the number of shipments, we should combine this observation with the previous result that – with fixed quantity – uncertainty leads to larger and less frequent shipments.

6.8. Wrapping up

In the previous subsections, we have derived the following expressions for our three endogenous variables, exported quantity, shipment size and number of shipments. Here we will wrap up our empirical predictions regarding the effect of uncertainty on these variables.

$$\lambda_i = g(I, \tau, \sqrt{\Psi^2 L}, \bullet) \quad (41)$$

$$\ln f^* = -\alpha + (1 - \beta_{dem}) \ln \lambda(I, \tau, \sqrt{\Psi^2 L}, \bullet) - \beta_{unc} \Psi^2 L \quad (42)$$

$$\ln q^* = \alpha + \beta_{dem} \ln \lambda(I, \tau, \sqrt{\Psi^2 L}, \bullet) + \beta_{unc} \Psi \sqrt{L} \quad (43)$$

Regarding the first equation, $\frac{d\lambda}{d\sqrt{\psi^2 L}} < 0$, as shown before. As a result, we expect a negative effect of uncertainty on total exports.

³⁷When inventories are costly, however, the smaller effective demand on more distant markets drives up the marginal costs because firms cannot benefit from increasing returns to scale in logistics technology. As a consequence, the fall in export is larger in our model than in models assuming iceberg transportation costs, and this difference is increasing in $C_1 = \sqrt{2hk\omega}$.

Regarding frequency, we have:

$$\frac{df^*}{d\sqrt{\psi^2 L}} = (1 - \beta_{dem}) \frac{d\lambda}{d\sqrt{\psi^2 L}} - \beta_{unc} < 0 \quad (44)$$

The first term in this equation represents the indirect effect of uncertainty: when uncertainty increases, total exports decrease, hence frequency decreases. $-\beta_{unc}$ shows the direct effect of increased uncertainty on frequency: even when quantity is fixed, larger uncertainty leads to more frequent shipments to minimize logistics costs. As both terms are negative, we expect the total effect of uncertainty to be negative on frequency. Also, when controlling for total quantity, we expect a negative direct effect of uncertainty on frequency.

Regarding shipment size, we get:

$$\frac{dq^*}{d\sqrt{\psi^2 L}} = \beta_{dem} \frac{d\lambda}{d\sqrt{\psi^2 L}} + \beta_{unc} \quad (45)$$

Now, the indirect effect is negative while the direct effect is positive. Hence, the sign of total effect of uncertainty of shipment size is uncertain, depends on the magnitude of β_{unc} relative to $\beta_{dem} \frac{d\lambda}{d\sqrt{\psi^2 L}}$. The direct effect, β_{unc} , however, is expected to be positive.

Overall, equations (43), (42), (41) provide the predictions for our empirical analysis.

7. Appendix C: Additional tables

Table 9: Core results - with quantity as dependent variable

Dep.Var.	(1) Tot. Quantity	(2) Nbr shipment	(3) Avg. Quantity	(4) Nbr shipment	(5) Avg. Quantity
log uncertainty	-0.050*** (0.010)	-0.028*** (0.004)	-0.022*** (0.008)	-0.016*** (0.003)	0.016*** (0.003)
Log total quantity				0.231*** (0.001)	0.769*** (0.001)
Firm*product FE	Yes	Yes	Yes	Yes	Yes
Destination FE	Yes	Yes	Yes	Yes	Yes
Observations	417,098	417,098	417,098	417,098	417,098
R-squared	0.039	0.053	0.032	0.398	0.868
Number of id	226,545	226,545	226,545	226,545	226,545

Reduced sample, where quantity is available. All regressions include firm-product fixed effects as well as destination fixed effects. Extra-EU exports sample. Shipment frequency truncated for 1-9 months. Robust standard errors in parentheses, clustered by destination-product level. ***, **, * denotes significance at 10%, 5%, 1% level, respectively.

Table 10: Pseudo Poisson Maximum Likelihood

	(1)	(2)	(3)	(4)	(5)
VARIABLES	Total value	Nbr shipments	Avg.Value	Nbr shipments	Nbr shipments
Log uncertainty	-0.051 (0.053)	-0.049*** (0.004)	0.042 (0.038)	-0.028*** (0.003)	0.032** (0.016)
Time x log uncertainty					-0.019*** (0.005)
Log tot. value				0.338*** (0.001)	0.343*** (0.002)
Observations	312,865	312,865	312,865	312,865	161,866
R-squared					
Number of id	71,674	71,674	71,674	71,674	37,379

All regressions include firm-product fixed effects as well as destination fixed effects. Robust standard errors in parentheses, clustered by destination-product level. ***, **, * denotes significance at 10%, 5%, 1% level, respectively.

Table 11: Uncensored OLS

	(1)	(2)	(3)	(4)	(5)
VARIABLES	Tot. Value	Avg.Value	Nbr. Shipment	Nbr. Shipment	Nbr. Shipment
Log uncertainty	-0.112*** (0.010)	-0.048*** (0.007)	-0.064*** (0.005)	-0.027*** (0.003)	0.011 (0.013)
time x log uncert					-0.012*** (0.004)
Log tot. Value				0.335*** (0.001)	0.343*** (0.001)
Observations	575,999	575,999	575,999	575,999	303,945
R-squared	0.094	0.068	0.095	0.577	0.576
Number of id	310,961	310,961	310,961	310,961	167,036

Includes above 9/year firm-product-destination observations. All regressions include firm-product fixed effects as well as destination fixed effects. Robust standard errors in parentheses, clustered by destination-product level. ***, **, * denotes significance at 10%, 5%, 1% level, respectively.

8. References

- Albornoz, F., Calvo Pardo, H. F., Corcos, G. and Ornelas, E. (2012), ‘Sequential exporting’, *Journal of International Economics* **88**, 17–31.
- Alessandria, G., Kaboski, J. P. and Midrigan, V. (2010), ‘Inventories, lumpy trade, and large devaluations’, *The American Economic Review* **100**(5), pp. 2304–2339.
- Alessandria, G., Kaboski, J. P. and Midrigan, V. (2011), ‘Us trade and inventory dynamics’, *American Economic Review* **101**(3), 303–307.
- Araujo, L. and Ornelas, E. (2007), Trust-based trade, Technical report, CEP Discussion Papers dp0820, Centre for Economic Performance, LSE.
- Ariu, A. (2011), The margins of trade: Services vs goods, mimeo, University of Leuven.
- Autor, D. H., Dorn, D. and Hanson, G. H. (2013), ‘The China Syndrome: Local Labor Market Effects of Import Competition in the United States’, *American Economic Review* **103**(6), 2121–68.
- Behrens, K. and Picard, P. M. (2011), ‘Transportation, freight rates, and economic geography’, *Journal of International Economics* **85**(2), 280–291.
- Berman, N., de Sousa, J., Martin, P. and Mayer, T. (2012), ‘Time to Ship during Financial Crises’, pp. 225–260.
- Berthou, A. and Vicard, V. (2013), Firms’ Export Dynamics: Experience versus Size, Working paper 445, Banque de France.
- Bloom, N. (2009), ‘The impact of uncertainty shocks’, *Econometrica* **77**(3), 623–685.
- Bloom, N., Floetotto, M., Jaimovich, N., Saporta-Eksten, I. and Terry, S. J. (2012), Really uncertain business cycles, NBER Working Papers 18245, National Bureau of Economic Research, Inc.
- Bloom, N. and Van Reenen, J. (2007), ‘Measuring and explaining management practices across firms and countries’, *Quarterly Journal of Economics* **122**(4), 1351–1408.
- Cameron, A. C. and Miller, D. L. (2011), *Robust Inference with Clustered Data*, CRC Press, chapter Handbook of Empirical Economics and Finance, pp. 1–28.

- Coleman, A. (2009), ‘Storage, slow transport, and the law of one price: Theory with evidence from nineteenth century u.s. corn markets.’, *Review of Economics and Statistics* **91** (2), 332–350.
- Crozet, M., Head, K. and Mayer, T. (2012), ‘Quality sorting and trade: Firm-level evidence for french wine’, *Review of Economic Studies* **79**(2), 609–644.
- Dixit, A. K. and Pindyck, R. S. (1994), *Investment under Uncertainty*.
- Eaton, J., Eslava, M., Kugler, M. and Tybout, J. (2008), Export dynamics in colombia: Firm-level evidence, Borradores de Economia 522, Banco de la Republica de Colombia.
- Eaton, J., Kortum, S. and Kramarz, F. (2004), ‘Dissecting trade: Firms, industries and export destinations’, *American Economic Review* **94**(2), 150–154.
- Evans, C. L. and Harrigan, J. (2005), ‘Distance, time, and specialization: Lean retailing in general equilibrium’, *American Economic Review* **95**(1), 292–313.
- Handley, K. and Limao, N. (2012), Trade and investment under policy uncertainty: Theory and firm evidence, NBER Working Papers 17790, National Bureau of Economic Research, Inc.
- Harrigan, J. (2010), ‘Airplanes and comparative advantage’, *Journal of International Economics* **82**(2), 181–194.
- Hornok, C. and Koren, M. (2015), ‘Administrative barriers to trade’, *Journal of International Economics* **96** (S1)(14), S110–S122.
- Hummels, D. L. and Schaur, G. (2010), ‘Hedging price volatility using fast transport’, *Journal of International Economics* **82**(1), 15–25.
- Hummels, D. and Schaur, G. (2013), Time as a trade barrier, NBER Working Papers 17758, forthcoming American Economic Review, 2013, National Bureau of Economic Research, Inc.
- Iacovone, L. and Javorcik, B. S. (2010), ‘Multi-product exporters: Product churning, uncertainty and export discoveries*’, *The Economic Journal* **120**(544), 481–499.
- Kleinert, J. and Spies, J. (2011), Endogenous transport costs in international trade, IAW Discussion Papers 74, Institut für Angewandte Wirtschaftsforschung.

Kropf, A. and Saure, P. U. (2012), Fixed costs per shipment, Working Papers 2012-13, Swiss National Bank.

Novy, D. and Taylor, A. (2013), Trade and uncertainty, Working Papers mimeo.

Santos Silva, J. M. C. and Tenreyro, S. (2006), ‘The log of gravity’, *The Review of Economics and Statistics* **88**(4), 641–658.

Steinwender, C. (2015), Information frictions and the law of one price: When the states and the kingdom became united, Technical report, Mimeo, Princeton University and Harvard Business School.

Zipkin, P. H. (2000), *Foundations of inventory management*, McGraw-Hill.