Red Nacional de Investigadores en Economía

Air or Sea? Quality-Driven Sorting in International Trade

Esteban Jaimovich (University of Turin/Collegio Carlo Alberto)

Vincenzo Merella (Prague University of Economics and Business/University of Cagliari)

DOCUMENTO DE TRABAJO N° 388

Febrero de 2026

Citar como:
Jaimovich, Esteban, Vincenzo Merella (2026). Air or Sea? Quality-Driven Sorting in International Trade. Documento de trabajo RedNIE N°388.

Air or Sea?
Quality-Driven Sorting in International Trade*

Esteban Jaimovich† Vincenzo Merella‡

January 2026

Abstract

We study how product quality shapes the choice between air and sea freight. In a model with non-homothetic demand and income uncertainty, exporters face a timing wedge: sea shipments must commit quantities before uncertainty resolves, whereas air shipments can be better timed but at higher cost. The model implies a sharp modal sorting for vertically differentiated products: high-quality varieties fly while lower-quality varieties sail. It also predicts that the average air–sea quality gap shrinks as sea routes lengthen relative to air routes. Using U.S. customs data at the exporter–district–product level, and proxying quality with unit values, we confirm both patterns. In addition, leveraging within origin–destination time variation in relative freight costs in a 2SLS design, we show that higher sea costs (relative to air) raise the average unit value of air shipments, consistent with quality sorting. Calibrated counterfactuals show that higher demand volatility reshapes selection in a distance-dependent way, tightening it on short routes while expanding air-based participation on long routes, whereas COVID-type air-freight disruptions and tariff-type marginal-cost shocks disproportionately hit profits for high-quality, long-distance exporters, and push marginal exporters to switch from air to sea.

JEL Classification: F1; F14

Keywords: International trade; Transportation mode; Unit values

*The authors gratefully acknowledge the financial support of the Czech Science Foundation under grant No. GA CR 19-16764S for the project "Revisiting the relationship between import quality and importer income and its effects on international trade patterns."

†University of Turin (ESOMAS Department) and Collegio Carlo Alberto. E-mail: esteban.jaimovich@unito.it

‡Prague University of Economics and Business and University of Cagliari. Email: merella@unica.it.

1 Introduction

Two central factors shape the choice of transportation mode in international trade: freight costs and delivery timeliness. This trade-off is most evident when comparing air and sea freight.[1] Air transport offers much faster delivery but at a substantially higher cost than shipping by sea. For some industries, product characteristics effectively dictate mode choice. Perishable or time-sensitive goods require rapid delivery, while goods that are heavy relative to their value (such as crude oil, grain, or coal) are generally uneconomical to fly (Alessandria, Kaboski and Midrigan, 2010). Yet, a striking feature of modern trade flows is that many products are shipped by both air and sea, even within the same origin–destination pair. This "mode mixing" is pervasive in the data (Hummels and Schaur, 2010) and suggests that transportation mode is not merely a technological constraint, but an endogenous margin of adjustment.

This paper investigates how the quality composition of trade interacts with the air–sea trade-off. We begin by documenting two stylized facts for U.S. imports of vertically differentiated goods. First, air-shipped varieties exhibit higher unit values than sea-shipped varieties within the same product–exporter–district cell. Interpreting unit values as proxies for product quality, this pattern points to quality sorting across modes: higher-quality varieties fly while lower-quality varieties sail. Second, the average air–sea unit-value gap is smaller on routes where sea travel is long relative to air travel. Taken together, these facts suggest that relative sea-to-air distance governs the extent of heterogeneity in quality sorting across air and sea shipments.

We propose a simple mechanism that ties these facts to demand risk stemming from income volatility. The key wedge is timing. Sea shipments require committing quantities well before final demand is realized, whereas air shipments can be better timed to reduce exposure to demand uncertainty in the destination market, albeit at a higher per-unit shipping cost. We embed this timing wedge into a model of vertically differentiated products with non-homothetic demand.[2] When destination income is uncertain and shipment delays matter, such non-homotheticities ties demand volatility to quality: demand for premium varieties is more sensitive to (and thus more volatile with) income realizations. In the presence of time wedges, higher-quality exporters face greater risk from early quantity commitment under sea freight and hence have stronger incentives to pay for the flexibility of air shipping. The model thus yields a

[1] For the case of the U.S., which is the importer we focus in this paper, the vast majority of imports are shipped either by air or sea, except for those originating from Canada and Mexico.

[2] An extensive literature has documented nonhomotheticities along the quality dimension, e.g., Schott (2004), Hallak (2006), Verhoogen (2008), Bastos and Silva (2010), Manova and Zhang (2012), Dingel (2017), Jaimovich, Madzharova, and Merella (2023).

sharp quality threshold for modal choice: high-quality varieties are shipped by air, low-quality varieties by sea. It also predicts that the air–sea quality gap shrinks on routes where sea transit times are long relative to air. The reason is that the commitment cost of sea freight rises under increased uncertainty, shifting the sorting cutoff and inducing more varieties to fly.

A central implication of endogenizing quality jointly with shipping mode is that changes in relative transport costs do not simply reallocate volumes across modes; they also reshape the average quality shipped by each mode. Guided by this result, we estimate a 2SLS specification at the product–origin–district–year level in which the average unit value of air shipments responds to air-freight costs and the import share shipped by sea, instrumenting sea share with sea-freight costs conditional on air-freight costs and absorbing a rich set of fixed effects. The instrument exploits within-route changes in relative freight costs that shift modal shares. Consistent with the model's re-sorting mechanism, we find that increases in the (instrumented) sea share are associated with higher average unit values among air shipments, indicating systematic changes in the quality composition shipped by air.

Beyond within-exporter reallocation across modes, the model also delivers implications for the extensive margin of exporting. When serving a destination requires incurring a fixed (or sunk) cost, export participation becomes jointly shaped by quality, relative distance, and the option value of transport-mode flexibility under uncertainty. Higher income volatility tilts the mode trade-off toward flexibility: it magnifies the cost of early quantity commitment under sea freight, thereby increasing the value of being able to delay shipment decisions to better accommodate demand swings. This generates a distance-dependent entry response: at short relative distances, volatility discourages marginal low-quality exporters who would ship by sea, while at long relative distances—where sea is uneconomical—volatility can increase participation by allowing marginal exporters to profitably ship by air. This extensive-margin channel differs from the standard logic in heterogeneous-firm trade models, where export participation is often summarized by one-dimensional productivity cutoffs (Melitz, 2003; Chaney, 2008; Das, Roberts, and Tybout, 2007). Instead, in our setting, participation is not separable from transport technology: timing risk and mode flexibility make the extensive margin depend both on route characteristics and product quality.

These mechanisms are quantitatively meaningful in our counterfactuals. Using a baseline calibration disciplined by U.S. data, we quantify how three distinct shocks—a mean-preserving increase in income volatility, a COVID-type rise in the relative cost of air freight, and a tariff-type increase in exporters' marginal production costs—reconfigure entry and mode choice across qualities and distances. Rising income volatility shifts the entry condition in opposite directions across routes: on short relative sea-to-air distances it raises the minimum quality required to profitably export, while on long relative distances it lowers the minimum quality required to export by air, thereby expanding participation at the lower end of the quality distribution. In contrast, shocks to relative air-to-sea transport costs and to marginal costs disproportionately affect high-quality, long-distance exporters that rely most heavily on fast transport, as well as mid-quality producers at intermediate distances that are either pushed out of exporting or induced to switch into slower (and riskier) sea shipping. These heterogeneous incidence patterns emerge because quality, geography, and nonhomothetic demand interact in our setting, and they are difficult to detect in reduced-form approaches that abstract from the joint determination of quality and transport mode.

Our analysis connects to the literature that studies delivery timeliness and demand uncertainty as determinants of firms' choices between air and sea freight (Hummels and Schaur, 2010, 2013; Martincus, Carballo, and Graziano, 2015; Hornok and Koren, 2015). The timing trade-off across transport modes is central in Hummels and Schaur (2010), who model air freight as a real option under demand uncertainty: firms can commit a baseline quantity by sea and adjust via air once uncertainty resolves, implying that air use should rise with demand volatility—a prediction they confirm using U.S. HS-10 import data.[3] Our work builds on this timing-under-uncertainty insight but embeds it in a setting where product quality is endogenously chosen and uncertainty originates in income volatility, so that demand risk is systematically higher for more income-elastic (higher-quality) varieties. This joint perspective generates endogenous sorting of quality across modes and yields testable implications both for the cross-section linking unit values and modal choice, as well as for the re-sorting dynamics when relative air–sea costs shift over time.

Beyond Hummels and Schaur's option-value approach, a number of papers study how the trade-off between speed and freight cost shapes specialization. Harrigan (2010) embeds transit-time differences in a Ricardian model and shows that more distant exporters have a comparative advantage in lightweight goods, which can more economically use air transport; a corollary is that longer-distance imports tend to exhibit higher unit values.[4] Hummels and Schaur (2013) quantify consumers' valuation of timeliness by estimating per-day ad valorem equivalents for time in transit, showing that timeliness is a key determinant of modal choice between air and

[3]Two recent working papers (Brancaccio, Kalouptsidi and Papageorgiou, 2024; and Blaum, Esposito and Heise, 2025) document how unexpected delays and demand volatility can impose significant revenue losses on exporters relying on maritime shipping.

[4]These results echo the findings by, e.g., Hummels and Skiba (2004), Baldwin and Harrigan (2011), and Harrigan, Ma and Shlychkov (2015), among others.

sea. Relatedly, Evans and Harrigan (2005) argue that time-sensitive goods are more likely to be imported from producers located geographically nearer, while Djankov, Freund and Pham (2010) document that such goods tend to be exported by countries with better infrastructure or institutional frameworks to avoid long delays. While these papers deliver sharp and testable specialization predictions, they generally abstract from the joint determination of quality and mode choice that is central in our framework.

Finally, our paper bridges the literatures on modal selection and quality-dependent nonhomothetic demand. The importance of income-dependent demand for quality is well documented in international trade (e.g., Bastos and Silva, 2010; Feenstra and Romalis, 2014; Lugovskyy and Skiba, 2016; Chen and Juvenal, 2022). In particular, higher-quality goods are associated with greater income elasticity of demand (e.g., Manova and Zhang, 2012; Fajgelbaum, Grossman, and Helpman, 2011; Jaimovich and Merella, 2012, 2015). By allowing quality to be jointly chosen with transport mode under uncertainty, we uncover a mechanism through which income volatility shapes trade composition across both goods and shipping modes, yielding cross-sectional and dynamic predictions that are absent in prior work.

The remainder of the paper is organized as follows. Section 2 describes the U.S. customs data, the construction of air and sea distances, and documents the main stylized facts on mode mixing and quality sorting. Section 3 lays out the model environment and demand structure. Section 4 characterizes optimal shipment choices and export entry across routes. Section 5 derives the model's implications for quality sorting and relative distance, and presents the panel evidence linking changes in relative freight costs to changes in the quality composition of air shipments. Section 6 reports counterfactual exercises. Section 7 concludes.

2 Dataset and Stylized Facts

This section documents a number of stylized facts on the choice of transport mode for U.S. imports, focusing on how product quality and relative sea-to-air distances shape this decision. The analysis is based on U.S. customs import data in year 2024, disaggregated at the HS 10-digit level. The dataset records, among other variables, the value, weight, and quantity of imports by country of origin and U.S. district of entry, as well as the mode of transport (air or vessel). In total, imports into the U.S. are classified across 42 districts of entry.[5]

[5]The U.S. customs data records both the district of entry and the district of unladen for imports. To increase the reliability that the district of entry is the relevant area of final destination of imports, in our empirical analysis we keep only the observations for which the district of entry and unladen coincide.

We augment the U.S. customs import dataset with information on sea and air distances between the country of origin and the U.S. district of entry of imports. Since the U.S. Census Bureau reports only the country of origin—not the precise port of departure—we compute distances as averages between each U.S. district of entry and the main ports/airports in the exporting country.

Several restrictions were imposed on the raw data in order to focus on the most informative observations. We first retained only those districts where both air and sea shipments are economically meaningful, keeping districts whose total import value exceeds USD 25 million for each mode of transport. (To put this threshold in perspective, the median value of seaborne imports across districts is USD 2.65 billion, while the median for airborne imports is USD 3.08 billion.) Districts outside the U.S. mainland—Honolulu Harbor, Anchorage, and San Juan—were also dropped. These two steps reduce the number of districts from 42 to 26.

On the exporter side, we excluded four different sets of countries: *i*) landlocked countries, *ii*) Mexico and Canada (the only U.S. land-border neighbors), *iii*) countries with fewer than one million inhabitants, and *iv*) countries whose territory exceeds 2.5 million square kilometers.[6] The population restriction serves to eliminate very small countries where import records are often noisy, while the area restriction limits potential measurement error in distance calculations, which is particularly severe for very large countries when the exact port of departure is not observed. After applying these four criteria, we are left with 83 countries of origin.

Finally, because the analysis is concerned with the possibility that transport choice varies along the quality dimension, we restrict attention to manufacturing goods. These account for 10,622 of the 11,949 HS10 product categories in the original dataset.

Mode of Transport Choice and Alchian-Allen Effect

One well-documented stylized fact in the trade literature is that many products exported from a given country are shipped to foreign markets using a combination of air and sea transport (see, e.g., Hummels and Schaur, 2010, 2013; Martincus, Carballo, and Graziano, 2015; Hornok and Koren, 2015). In particular, using HS10 product-level data, Hummels and Schaur (2010) show that roughly 35% of product–exporter pairs shipped to the United States between 1990 and 2004 actively employed both modes of transport.

[6] There are eight countries, besides the U.S., above this threshold: Russia, Canada, China, Brazil, Australia, India, Argentina, and Kazakhstan. As mentioned above, Canada is already excluded owing to sharing a land border with the United States. In addition, Kazakhstan is also excluded as a result of being landlocked. Our results are robust to keeping countries with an area smaller than 7,000,000 sq km (which would amount to keeping India and Argentina in the sample).

Figure 1. Mode of Transport at the Product-Exporter-District Level

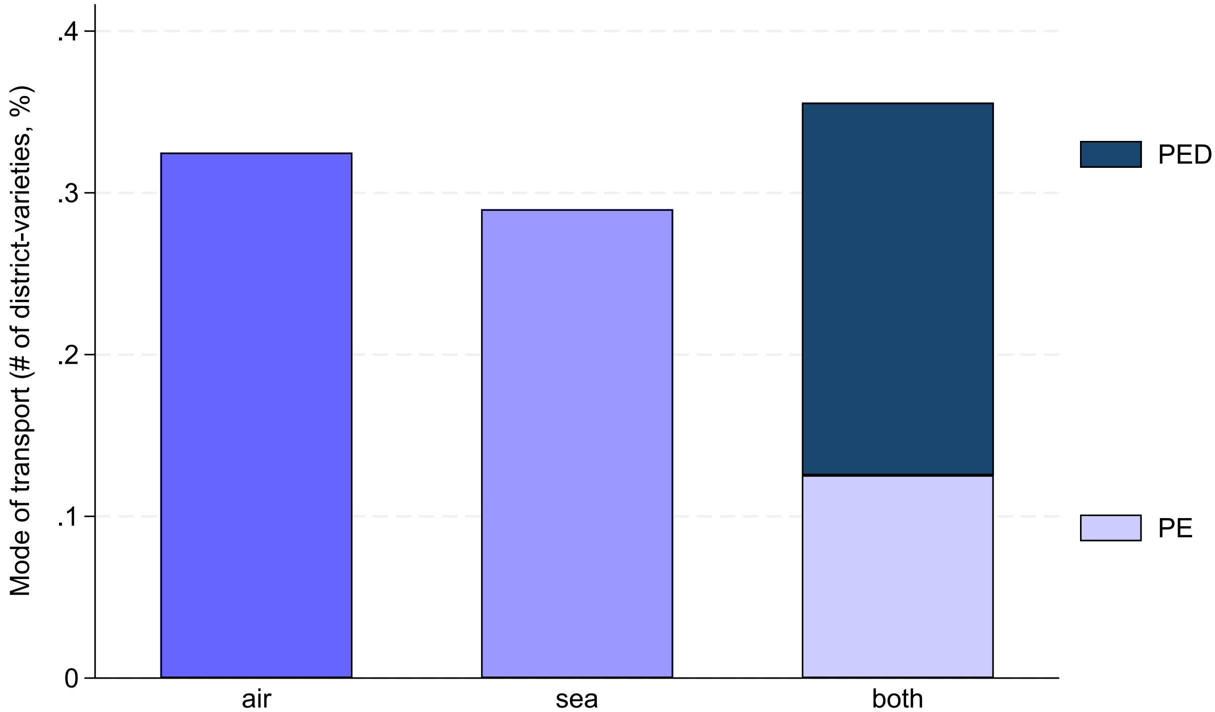

Note. The left bar represents the share of exporter-product pairs shipped exclusively by air to all U.S. districts of entry; the center bar shows the share shipped exclusively by sea; and the right bar corresponds to the share involving mode mixing. In this case, we further distinguish between mode mixing across districts (PE) and within the same district of entry (PED).

In Figure 1, we build on the findings of Hummels and Schaur (2010) by further disaggregating trade flows according to the specific U.S. district of entry. More precisely, we slice the data at the product–exporter–district (PED) level. The bar graph reports the share of products per exporter shipped exclusively by air, exclusively by sea, or through a combination of modes. Among those employing both air and sea transport, we distinguish two cases: (i) mode mixing observed only *across* different entry districts (the lighter blue portion at the bottom of the bar, labeled PE); and (ii) mode mixing at the product-exporter level *within* the same district (the darker blue portion at the top, labeled PED). The results reveal that more than 37% of product–exporter pairs rely on both transport modes. Strikingly, in about two-thirds of these cases (that is, roughly 25% of all product–exporter flows) air and sea transport are used simultaneously at the product-exporter level for shipments to the very same U.S. district.

A key benefit of analyzing trade flows at the product–exporter–district level is that it allows us to link transport mode choice to differences in relative distances between sea and air routes

Table 1. Mode of Transport and Relative Sea vs. Air Distance

	(1)	(2)	(3)
relative air-to-sea distance	-0.297^{***}	-0.209^{***}	-0.224^{***}
	(0.055)	(0.025)	(0.027)
observations	512, 947	468, 013	468, 013
exporter-product FE	yes	yes	yes
district FE	no	yes	no
district-product FE	no	no	yes

Note. The dependent variable is the share of U.S. imports (in weight) transported by sea for a specific product-exporter-district combination. Relative sea-to-air distance is the sea-route distance divided by the air-route distance between the exporter and the U.S. entry district of the exports. Robust standard errors clustered at exporter-district level. Significance levels: $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

from the country of origin to the district of entry. As Hummels and Schaur (2013) emphasize, the East–West coast geography of the United States generates substantial variation in sea-versus-air distances across districts on different coasts. Table 1 exploits this variation to study how relative distances shape transport decisions by regressing the share of imports arriving by sea on the relative sea-to-air distance between the exporter and the U.S. district of entry. Column (1) shows that exporters are more likely to ship by air when the ratio air-to-sea distance increases. Column (2) confirms that this relationship remains robust after controlling for district fixed effects, indicating that it is not driven by persistent district-level patterns in transport mode. Column (3) further includes district–product fixed effects, allowing for product-specific modal patterns to vary by U.S. districts, with results remaining essentially unaltered.

The Alchian-Allen effect predicts that exporters will send higher-quality varieties of products to more distant markets. This prediction has been widely supported in the data that exploits cross-country variation in bilateral distances –e.g., Hummels and Skiba (2004), Baldwin and Harrigan (2011). In Table 2, we show that the Alchian-Allen predictions hold as well when exploiting variation in distances from a country of origin to different U.S. districts and by separating trade flows according to the mode of transport.

The dependent variable in all regressions reported in Table 2 is the Free on Board (FOB) unit value of imports at the HS10 product level, with trade flows separated by mode of transport. Columns (1)–(3) focus on imports transported by sea, regressing (log) unit values of seaborne imports on the (log) sea distance between the exporter and the U.S. district of entry. Columns (4)–(6) repeat the analysis for imports transported by air. The baseline results in columns (1)

Table 2. Alchian-Allen Effect by Mode of Transport

	(1)	(2)	(3)	(4)	(5)	(6)
(log) sea distance	0.107^{***}	0.128^{***}	0.181^{***}			
	(0.027)	(0.019)	(0.025)			
(log) air distance				0.419^{***}	0.384^{***}	0.462^{***}
				(0.093)	(0.058)	(0.071)
observations	192, 201	192, 201	159, 879	252, 263	252, 263	224, 190
exporter-product FE	yes	yes	yes	yes	yes	yes
district FE	no	yes	no	no	yes	no
district-product FE	no	no	yes	no	no	yes

Note. The dependent variable in columns (1)-(3) is the (log) unit value for imports transported by sea, whereas in columns (4)-(6) is the (log) unit value of imports transported by air. All unit values are based on FOB import values. Robust standard errors clustered at exporter-district level. Significance levels: $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

and (4) show that, for a given exporter and product, shipments to more distant U.S. districts tend to have higher unit values, conditional on the chosen mode of transport (sea or air).

To test robustness, columns (2) and (5) add district fixed effects, which capture factors such as regional income differences that could affect average import quality under nonhomothetic preferences. Finally, columns (3) and (6) introduce a full set of district-by-product fixed effects, allowing for example for heterogeneity in product-specific quality preferences across districts.

Quality, Distance, and Mode of Transport

Disaggregating exporter-by-product imports by district of entry and mode of transport provides additional insight into how transportation mode choices vary along the quality dimension. To that end, column (1) of Table 3 regresses the (log) unit values of imports —regardless of transport mode—on a dummy variable *Air Trans*, which equals 1 when air transport is used. The specification includes a full set of product-exporter-district fixed effects. The results show that, within a given product-exporter-district combination, airborne imports are approximately twice as expensive as seaborne imports. Interpreting unit values as proxies for product quality, this suggests that air freight is disproportionately used for higher-quality varieties.

We further exploit variation in relative sea versus air distances across U.S. entry districts. Because these relative distances differ substantially across exporters and certain U.S. districts, they allow us to test whether transport mode choice along the quality dimension responds at the

Table 3. Unit Values by Mode of Transport

	(1)	(2)	(3)	(4)	(5)
air transp	1.017^{***}	1.027^{***}			
	(0.016)	(0.017)			
air transp x (log) distance diff		-0.245^{***}	-0.208^{***}	-0.224^{***}	-0.192^{***}
		(0.076)	(0.066)	(0.044)	(0.058)
observations	131,718	131,718	128,504	128,496	75,548
district-exporter-product FE	yes	yes	yes	yes	yes
product-transport FE	no	no	yes	yes	no
district-transport FE	no	no	no	yes	no
exporter-transport FE	no	no	no	yes	no
district-product-transport FE	no	no	no	no	yes
exporter-product-transport FE	no	no	no	no	yes

Note. The dependent variable in all columns is the (log) unit value by exporter-product-district-transport cells. 'air trans' is a dummy variable equal to 1 when the product is transported by air, and to 0 when is done by sea. '(log) distance diff' is the (log) sea-to-air bilateral distance between the exporter and the U.S. district of entry. Robust standard errors clustered at exporter-district-transport level. Significance levels: $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

extensive margin to changes in sea–air distance ratios. Columns (2)–(5) in Table 3 address this question by including an interaction term, *Air Trans* × *(log) Distance Diff*, where *Distance Diff* is the ratio of sea distance to air distance between the exporter and the U.S. district of entry. Column (2) augments the regression in column (1) with this interaction term. The coefficient associated with the interaction term is negative and statistically significant. This implies that the difference in the average price of airborne and seaborne imports tends to narrow as relative sea–air distance increases. Columns (3)–(5) demonstrate that this result remains robust across specifications with additional sets of fixed effects.[7]

Considering unit values as a proxy for quality, the results in Table 3 showcase two inter-related empirical observations. When both means of transport are simultaneously used at the product-exporter-district level:

[7] For example, in column (3), product-transport fixed effects would control for the possibility that certain products tend to use more intensively one specific mode of transport. In column (4), district-transport fixed effects will account for different intensities of the mode of transport across U.S. districts. Similarly, exporter-transport fixed effects will account for the variation of the mode of transport by country of origin of imports. Finally, the sets of fixed effects in column (5) allow the intensity of mode of transport to vary heterogeneously by district of entry, exporter, and product.

1. Product varieties sent by air tend to be of higher quality than those shipped by sea.
2. The average quality gap of the varieties sent by air relative to those shipped by sea narrows as the relative bilateral distance sea vs air increases.

In the next section, we present a model that will intend to rationalize the above two observations as a result of shifts in the extensive margin along which vertically differentiated producers choose their mode of transport optimally. The model will lead to modal sorting along the quality dimension: firms producing higher-quality varieties will use air freight as a mode of transport. In addition, this selection will also respond to differences in relative bilateral distances between air and sea routes. In particular, consistent with the evidence in Table 1, higher relative distances sea vs air will lead to lower use of sea relative to air freight. This result will materialize as a shift on the extensive margin for the mode of transport selection, leading in turn to a reduction of the average gap in quality between airborne and seaborne imports.

3 Setup of the Model

This section introduces the main elements of our framework and next derives the import demand schedules. This is instrumental to the analysis conducted in the next section, where we investigate the exporters' choice in terms of transport mode.

3.1 Goods Space, World Geography, and Exporters' Technology

We study a framework in which a differentiated consumption good is available in a continuum of quality levels, denoted by $\lambda \in [0, \bar{\lambda}]$. The good is non-perishable and can be stored indefinitely at zero cost.

Our focus is on a large destination market that imports distinct varieties of the good from geographically dispersed producers.[8] More precisely, there is a large number of countries, each one hosting producers of distinct varieties of the consumption good who are potential exporters to the destination market. We order these countries by increasing sea distance to the destination market, denoting this distance by $x > 0$. For ease of notation, we also use x to index exporting countries throughout the analysis.

In each country x, there is a continuum of producers that differ in the range of quality varieties they are able to supply. Specifically, consider a producer in country x who can achieve

[8] In line with the empirical evidence discussed in Section 2, we interpret this market as the United States.

a maximum quality level $\lambda \in [0, \bar{\lambda}]$. We assume that such a producer can manufacture any quality within the interval $[0, \lambda]$ at a constant marginal cost $c > 0$. This cost structure ensures that producers with a broader quality scope—that is, with higher attainable λ—enjoy larger potential profits, since they can target higher-quality segments without incurring a higher marginal cost. In this sense, the marginal cost specification can be interpreted as a reduced form that captures the endogenous choice of quality by producers.[9] In addition, since each producer will optimally choose the highest level of achievable quality λ, we can thus denote the producer of variety λ in country x by the pair (λ, x).

Exporters face both entry and transportation costs. To access the destination market, a producer must incur a lump-sum entry cost $\Upsilon > 0$ at $t = 0$. That is, exporters must commit to paying Υ upfront in order to establish an export route.

Conditional on the entry decision, the exporter must eventually choose a mode of transportation between air (A) and sea (S). Air shipping is assumed to be instantaneous (i.e. delivery to the destination market is immediate). By contrast, sea shipping requires one unit of time per unit of (sea) distance. Transportation costs differ across modes as well: in the case of air, goods are subject to an iceberg cost $\tau_x > 1$, whereas sea shipping involves no iceberg loss. We adopt this assumption for analytical convenience, as it allows us to disentangle the effect of higher transportation costs (borne by air cargoes) from the effect of demand uncertainty arising from the longer time lags of sea shipping.[10]

Alongside foreign producers of the differentiated good, there exists a competitive fringe of local suppliers capable of producing a "copycat" variety that is indistinguishable from the imported one. Local copycats can supply this variety at a constant marginal cost $\bar{c} > \tau_x c$.

3.2 Endowments and Preferences

The destination market is populated by a unit mass of individuals. Individuals are born at time $t = 0$, live for a unit time interval, and consume only at the last instant of their lives, $t = 1$. We assume that every individual is endowed with an identical amount of resources at all times, and denote the endowment at time t by $Y_t > 1$. The endowment Y_t follows a random process, with the probability and cumulative distributions denoted by $f(\cdot)$ and $F(\cdot)$, respectively, and such that income volatility between any two dates is proportional to the time elapsed between

[9]For a discussion on this type of reduced form modeling, see, e.g., Baldwin and Harrigan (2011).

[10]In this respect, it should be noted that while the *absolute* sea distance and the *relative* sea-to-air distance are effectively equivalent, the measure x will impact the exporter's choice only through the delayed products' delivery that sea shipping entails relative to air shipping.

them. Longer transit lags will be then associated with greater uncertainty about the income level that will prevail at the time of consumption.

Each individual in the destination market has identical preferences represented by a two-tier PIGLOG expenditure function:

$$\log Y = \bar{\alpha} \sum_{\lambda \in \Lambda} \log P_{\lambda} + u \prod_{\lambda \in \Lambda} (P_{\lambda})^{\lambda}, \tag{1}$$

where

$$P_{\lambda} \equiv \exp\left(\int_{\mathcal{X}} \ln p_{\lambda,x} dx \right) \tag{2}$$

is a Cobb-Douglas price aggregator comprising all varieties of quality λ, and where $p_{\lambda,x}$ is the price of the variety of quality λ supplied by country x.[11]

Applying Shepherd's lemma to the log-expenditure function yields the Marshallian demand functions (see the Appendix for the derivation of the demand function):

$$q_{\lambda,x} = \frac{(\bar{\alpha} + \lambda \log Y) Y}{p_{\lambda,x}}.$$

To ease notation, we define $\Omega_{\lambda}(Y) \equiv (\bar{\alpha} + \lambda \log Y) Y$, and let

$$q_{\lambda,x} = \frac{\Omega_{\lambda}(Y)}{p_{\lambda,x}}. \tag{3}$$

We interpret the term $\Omega_{\lambda}(Y)$ as a *demand shifter* revealing the nonhomothetic nature of preferences along the quality dimension of consumption goods. As it will become clearer later on, it will be through the convexity of $\Omega_{\lambda}(Y)$ with respect to Y that demand uncertainty will influence the choice between (cheap and slow) maritime vs. (fast and expensive) airborne transportation heteoegeneously at different levels of quality λ.

4 Trade Routes, Mode of Transport, and Entry

Foreign producers must make two sequential decisions. The first is whether or not to incur the entry cost Υ. The second, conditional on entry, is which mode of transport to use for exporting goods. This second decision implicitly involves the exact departure time $t_x \in [0,1]$ at which goods are placed at the relevant port: either their local seaport ($t_x = t_x^S$) or airport ($t_x = t_x^A$).

[11]See the Appendix for a derivation of this specific preference specification. To be perfectly rigorous, the indirect utility, expenditure and prices of all goods should be indexed by $t = 1$, i.e., the time when trade occurs. Whenever possible, we henceforth choose to avoid this detail to ease notation.

Since consumption takes place at $t = 1$, an exporter from country x that ships by air will optimally set $t_x^A = 1$, as air transport is instantaneous. Conversely, an exporter choosing sea transportation will optimally set $t_x^S = 1 - x$; this ensures that goods arrive at the destination market no later than $t = 1$, while minimizing at the same time the exposure to demand volatility during transit.[12]

4.1 Market-Clearing Prices and Expected Profits

Exporters set the quantity to ship to the destination market, $q_{\lambda,x}^T$, where $T = A, S$ denotes the mode of transport. Given that quantity, and the realized Y_1 at $t = 1$, market prices are determined to clear markets at that moment. The presence of a local competitive fringe who may supply identical alternative varieties to those imported at marginal cost $\bar{c} > \tau_x c$ implies that, given $q_{\lambda,x}^T$, market-clearing prices will be such that:

$$p_{\lambda,x} = \begin{cases} \bar{c} & \text{if } q_{\lambda,x}^T \leq \dfrac{\Omega_\lambda (Y_1)}{\bar{c}}, \\ \dfrac{\Omega_\lambda (Y_1)}{q_{\lambda,x}^T} & \text{if } q_{\lambda,x}^T > \dfrac{\Omega_\lambda (Y_1)}{\bar{c}}. \end{cases} \tag{4}$$

Air Transport Optimal Quantities and Profits

When choosing air transport, the exporter will optimally set the quantity $q_{\lambda,x}^A$ to send to the destination market at $t_x^A = 1$. This means that the exporter will face no uncertainty about the demand conditions at the destination market. Since (ex-post) profits are equal to $\pi_{\lambda,x}^A = (p_{\lambda,x} - c)\, q_{\lambda,x}^A$, with $p_{\lambda,x}$ given by (4), in the case of air transport the exporter (λ, x) will optimally set the quantity:

$$q_{\lambda,x}^A = \frac{\Omega_\lambda (Y_1)}{\bar{c}}, \tag{5}$$

consistent with a market-clearing price $p_{\lambda,x} = \bar{c}$, and yielding a profit $\pi_{\lambda,x}^A = (1 - \tau_x c/\bar{c})\, \Omega_\lambda (Y_1)$.

Although the exact amount of production $q_{\lambda,x}^A$ will be set at $t_x^A = 1$, the transport mode decision will be taken actually at the latest moment when sea transport is still feasible; that is, at $t = t_x^S = 1 - x$. As a consequence, at the moment $t = 1 - x$, an exporter will assess air shipping in terms of its expected profits, given the information available at that time (i.e., given Y_{1-x}). More precisely, when assessing the expected profits of shipping by air, the exporter knows that they will set a quantity (5) to achieve a price $p_{\lambda,x} = \bar{c}$, yet the exact endowment

[12]Seaborne shipments departing at $t < t_x^S$ are suboptimal (as they would entail higher demand uncertainty than those departing at $t = t_x^S$), while shipments departing at $t > t_x^S$ would fail to deliver the goods on time.

value Y_1 remains still unknown. The expected profit if using air transport, computed at $t = 1-x$ is then given by:

$$E\left(\pi^A_{\lambda,x}\middle| Y_{1-x}\right) = E\left(\Omega_\lambda\left(Y_1\right)\middle| Y_{1-x}\right)\left(1 - \frac{\tau_x c}{\bar{c}}\right). \tag{6}$$

Sea Transport Optimal Quantities and Profits

When exporters choose sea transportation, they must set the quantity $q^S_{\lambda,x}$ to ship in advance, before the actual value of $\Omega_\lambda\left(Y_1\right)$ is realized. More specifically, an exporter must set the quantity to deliver in the destination market at $t^S_x = 1-x$, owing to the time lag x involved in sea cargoes, while the market-clearing price will only be determined at $t = 1$.

We denote by $\hat{Y}(q^S_{\lambda,x})$ the value of Y that would set the market-clearing price $p_{\lambda,x} = \bar{c}$ for a quantity $q^S_{\lambda,x}$ transported by sea. That is, bearing in mind (4), $\hat{Y}(q^S_{\lambda,x})$ is implicitly determined by the equality:

$$\Omega_\lambda\left(\hat{Y}(q^S_{\lambda,x})\right) = \bar{c}q^S_{\lambda,x}, \tag{7}$$

which in turn yields $\hat{Y}(q^S_{\lambda,i}) = \Omega^{-1}_\lambda\left(\bar{c}q^S_{\lambda,i}\right)$. Then, the expected profit if using sea transportation, given the quantity shipped from x ($q^S_{\lambda,x}$) and the value of the endowment at $t = 1-x$ (Y_{1-x}), is given by:

$$\begin{aligned} E\left(\pi^S_{\lambda,x}(q^S_{\lambda,x})\middle| Y_{1-x}\right) = \textstyle\int_{Y_{\min}}^{\hat{Y}(q^S_{\lambda,x})} \Omega_\lambda\left(Y_1\right)\, f\left(Y\middle| Y_{1-x}\right)\, dY_1 + \\ q^S_{\lambda,x}\left[\bar{c}\left(1 - F\left(\hat{Y}(q^S_{\lambda,x})\middle| Y_{1-x}\right)\right) - c\right]. \end{aligned} \tag{8}$$

The optimal value of $q^S_{\lambda,x}$ given Y_{1-x} stems from maximizing expected profit (8). Namely,

$$\left[\Omega_\lambda\left(\hat{Y}(\cdot)\right) - \bar{c}q^S_{\lambda,x}\right]\, f\left(\hat{Y}(\cdot)\right)\frac{d\hat{Y}(\cdot)}{dq^S_{\lambda,x}} + \left[\bar{c}\left(1 - F\left(\hat{Y}(q^S_{\lambda,x})\middle| Y_{1-x}\right)\right) - c\right] = 0, \tag{9}$$

which bearing in mind (7) yields the following equality:

$$F\left(\hat{Y}(q^S_{\lambda,x})\middle| Y_{1-x}\right) = 1 - \frac{c}{\bar{c}}. \tag{10}$$

Lemma 1 *Given the optimal quantity to be shipped by sea cargoes from country x, $q^S_{\lambda,x}$, the cutoff value $\hat{Y}_x$ is independent of λ for any exporter from x.*

Lemma 1 states that, when exporters set optimally $q^S_{\lambda,x}$, the threshold $\hat{Y}_x$ above which the market-clearing price will turn out to be equal to $\bar{c}$ is exactly the same for all exporters from country x, no matter the quality level λ of the variety produced. Based on this result, we can plug (10) into (8), bearing in mind (7), to obtain the expected profit by an exporter from x producing a variety of quality λ if opting for sea shipping:

$$E\left(\pi^{S}_{\lambda,x}\middle| Y_{1-x}\right) = E\left(\Omega_{\lambda}\left(Y_{1}\right)\middle| Y_{1-x}; Y_{1} \leq \hat{Y}_{x}\right) \cdot \left(1 - \frac{c}{\bar{c}}\right). \tag{11}$$

Equation (11) reveals that the expected profit obtained by an exporter from x when choosing sea shipping is proportional to the conditional expectation of $\Omega_{\lambda}\left(Y\right)$ given the value of Y_{1-x}, but evaluated over a distribution of Y_1 that is truncated at $\hat{Y}_x$. By contrast, equation (6) indicates that expected profits under air shipping are proportional to the conditional expectation of $\Omega_{\lambda}\left(Y_1\right)$ given Y_{1-x} over the entire distribution of Y_1.

The source of discrepancy (6) and (11) lies in the different timing of quantity decisions. With air shipping, exporters can set quantities precisely at $t = 1$, once income uncertainty has been resolved. This guarantees the market-clearing price equals exactly $\bar{c}$, as exporters can always adjust quantities optimally to realized demand. By contrast, sea shipping requires exporters to commit quantities before $t = 1$ (at $t = 1 - x$), based only on partial information (i.e., based on Y_{1-x}). As a consequence, they cannot adjust production to accommodate large, unexpected income shocks, which effectively truncates the relevant distribution and lowers expected profits.

4.2 Choice of Transport Mode

The exporter chooses the mode of transport at time $t^{S}_{x} = 1 - x$ by comparing the expected profits (6) and (11). Notice that since $\Omega_{\lambda}\left(Y_1\right)$ is strictly increasing in Y_1, it follows that $E\left(\Omega_{\lambda}\left(Y_{1}\right)\middle| Y_{1-x}; Y_{1} \leq \hat{Y}_{x}\right) < E\left(\Omega_{\lambda}\left(Y_{1}\right)\middle| Y_{1-x}\right)$, since the former conditional expectation is applied on a distribution of Y_1 truncated above at $\hat{Y}_x$. This inequality reflects the timing advantage of air cargoes: its greater speed enables more flexible production planning, allowing exporters to better align output with demand fluctuations driven by income uncertainty. Exporters must therefore weigh these flexibility gains against the higher transportation costs of air freight when optimally choosing their mode of transport. In particular, an exporter from country x will opt for air transportation when the following condition holds:

$$\frac{E\left(\Omega_{\lambda}\left(Y_{1}\right)\middle| Y_{1-x}\right)}{E\left(\Omega_{\lambda}\left(Y_{1}\right)\middle| Y_{1-x}; Y_{1} \leq \hat{Y}_{x}\right)} \geq \frac{\bar{c} - c}{\bar{c} - \tau_{x} c}. \tag{12}$$

The inequality in (12) contrasts the production-flexibility benefits of air over sea transport (captured on the left-hand side) with the higher costs of air cargo in the form of the iceberg cost τ_x (on its right-hand side). The next proposition shows that there exists an exporter-specific quality threshold above which air dominates sea as the preferred mode.

Proposition 1 (Air-transportation condition) *There exists an exporter-specific quality threshold $\tilde{\lambda}_x$ such that, if choosing to export to the destination market, producer (λ, x) will use airborne cargo when $\lambda \geq \tilde{\lambda}_x$, and maritime freight otherwise.*

The quality threshold $\tilde{\lambda}_x$ can be further characterized by introducing an explicit random process for the endowment. Henceforth, we let Y_t be governed by a diffusion process, such that $Y_1 | Y_{1-x} \sim LogNormal\left(m_x, s_x^2\right)$, with $m_x \equiv \ln Y_{1-x} - s_x^2/2$ and $s_x^2 \equiv \sigma^2 x$.[13] Under this assumption, it follows that:

$$\tilde{\lambda}_x = \frac{\bar{\alpha}}{\Gamma_x\left(s_x, \tau_x\right) - \ln Y_{1-x}}, \tag{13}$$

where $\Gamma_x\left(s_x, \tau_x\right) \equiv s_x\left(\phi_L\left(s_x\right) + \Psi\left(\tau_x\right)\Delta\phi\left(s_x\right)\right) / \left(\Psi\left(\tau_x\right)\Delta\Phi\left(s_x\right) + \Phi_L\left(s_x\right) - 1\right) - s_x^2/2$, with $\Psi\left(\tau_x\right) \equiv \left(\bar{c} - c\right) / \left(\bar{c} - \tau_x c\right) > 1$. The terms with $\phi\left(\cdot\right)$ and $\Phi\left(\cdot\right)$ denote the density and cumulative distribution function of the standardized normal process, ϕ_L the density at the lower-bound of the distribution ($Y_1 = 1$), $\Delta\phi$ and $\Delta\Phi$ the differences between the upper-bound ($Y_1 = \hat{Y}$) and lower-bound pdf and cdf.[14]

We collect the most relevant comparative statics on the transport-mode threshold $\tilde{\lambda}_x$ in the following corollary.

Corollary 1 (Comparative statics on $\tilde{\lambda}_x$) *The threshold $\tilde{\lambda}_x$ is decreasing and s_x (hence, decreasing in both σ and x) and increasing in τ_x.*

Corollary 1 points out that factors heightening demand uncertainty (s_x)—either via higher volatility σ or sea distance x—raise the option value of speed, so the flexibility premium of air grows and $\tilde{\lambda}_x$ falls. Conversely, a higher (relative) cost of air transport (τ_x) lifts the threshold. Therefore, along a given route, the air share rises with product quality and volatility, while, across routes, it rises with (sea) distance.

4.3 Export Entry Choice

The producer (λ, x) must choose whether or not to incur the entry cost Υ at $t = 0$. Entry occurs if the expected profits from exporting, evaluated at $t = 0$, exceed this cost. Importantly,

[13] Note that this process means that $E\left[Y_1 | Y_{1-x}\right] = Y_1$.

[14] Lemma 2 in the appendix describes the geometric Brownian motion generating the stochastic process for Y_t and derives the closed-form expression for the left-hand side of (12)—see equation (19) therein. The threshold (13) is then obtained by jointly considering (19) and (23) evaluated at $\tilde{\lambda}_x$, and imposing that the resulting expression equals Ψ.

at the time of the entry decision the shipping mode is not yet determined, since it depends on the realization of Y_{1-x}, a random variable with distribution is given by $f_0(Y_{1-x}) \equiv f(Y_{1-x}|Y_0)$.

From the perspective of $t = 0$, the threshold condition (12) implies that transport mode itself is a random variable, as its exact value varies with Y_{1-x} according to $\partial\tilde{\lambda}_x/\partial Y_{1-x} > 0$. As a consequence, for a given λ, sufficiently low realizations of Y_{1-x} imply $\lambda \geq \tilde{\lambda}_x$, in which case the producer will use air transport; for sufficiently high realizations of Y_{1-x}, the opposite holds. We may thus define a threshold $\tilde{Y}_{\lambda,x}$, such that for any $Y_{1-x} < \tilde{Y}_{\lambda,x}$, (resp. $Y_{1-x} \geq \tilde{Y}_{\lambda,x}$) the exporter of the variety with quality λ from country x ships by air (resp. by sea).

Denoting by $E_0(\pi_{\lambda,x})$ the expected profit of producer (λ, x) from exporting computed at $t = 0$, we can write:

$$E_0(\pi_{\lambda,x}) = \int_{Y_{1-x}<\tilde{Y}_{\lambda,x}} E_{1-x}\left(\pi^A_{\lambda,x}\right) f_0(Y_{1-x})\, dY_{1-x} + \int_{Y_{1-x}\geq\tilde{Y}_{\lambda,x}} E_{1-x}\left(\pi^S_{\lambda,x}\right) f_0(Y_{1-x})\, dY_{1-x}.$$

Since $\partial E_{1-x}\left(\pi^{\mathcal{T}}_{\lambda,x}\right)/\partial\lambda > 0$ for $\mathcal{T} = \{A, S\}$, it also follows that $\partial E_0(\pi_{\lambda,x})/\partial\lambda > 0$. This allows us to characterize the entry threshold for exporters from country x.

Proposition 2 (Export entry condition) *There exists a quality threshold $\hat{\lambda}_x$ implicitly defined by the zero-profit condition:*

$$E_0\left(\pi_{\hat{\lambda}_x,x}\right) \equiv \int_{Y_{1-x}<\tilde{Y}_{\hat{\lambda}_x,x}} E_{1-x}\left(\pi^A_{\hat{\lambda}_x,x}\right) f_0(Y_{1-x})\, dY_{1-x} + \int_{Y_{1-x}\geq\tilde{Y}_{\hat{\lambda}_x,x}} E_{1-x}\left(\pi^S_{\hat{\lambda}_x,x}\right) f_0(Y_{1-x})\, dY_{1-x} = \Upsilon, \tag{14}$$

such that only producers from country x with $\lambda > \hat{\lambda}_x$ will enter the export market.

The result in Proposition 2 incorporates the classic "shipping the good apples out" result into our framework: only sufficiently high-quality producers choose to export.

5 Quality Sorting, Distance, and Costs

This section brings together the model's implications for quality sorting with our empirical evidence on distance and transport costs. We first use the analytical solution for the entry and mode thresholds to characterize how exporters sort themselves in terms of modal choice, and to show how relative sea-to-air distance shapes both export participation and the quality gap between air- and sea-shipped varieties. We then map these predictions into the data, relating within-route variation in modal shares and unit values to relative distances and freight costs.

5.1 Variations in Relative Distance

The results in Proposition 1 establish a clear quality-sorting pattern when both modes of transport are actively used in equilibrium: lower-quality varieties travel by sea and higher-quality varieties by air. More precisely, if the parametric configuration satisfies $0 < \hat{\lambda}_x < \tilde{\lambda}_x < \bar{\lambda}$, then all varieties originating from country x with quality $\hat{\lambda}_x < \lambda < \tilde{\lambda}_x$ will be shipped by sea, whereas those with λ above $\tilde{\lambda}_x$ will be transported by air.

These equilibrium sorting patterns imply that the average quality of air-shipped varieties exceeds that of sea-shipped varieties. This prediction is consistent with the stylized facts reported in column (1) of Table 3, which show that, when comparing U.S. imports of finely disaggregated products originating from a given country and entering a given U.S. district, the average unit values of airborne imports are about twice as high as those arriving by sea. Beyond documenting quality sorting by transport mode, our model also characterizes how the average quality gap between airborne and seaborne imports evolves across different levels of relative sea-to-air distance. We formalize these patterns in the following proposition.

Proposition 3 *Consider the case where $0 < \hat{\lambda}_x < \tilde{\lambda}_x < \bar{\lambda}$, so that both modes of transport are employed for imports from x. Let $\mu_x^S \equiv \left(\tilde{\lambda}_x + \hat{\lambda}_x\right)/2$ and $\mu_x^A \equiv \left(\bar{\lambda} + \tilde{\lambda}_x\right)/2$ denote, respectively, the average quality of seaborne and airborne imports for a country located at relative sea-to-air distance x. Then the following results hold: (i) $\mu_x^A > \mu_x^S$; (ii) the quality gap, $\mu_x^A - \mu_x^S$, decreases with relative sea-to-air distance x.*

Figure 2 offers a graphical representation of entry and transport choice discerned by product quality and relative sea-to-air distance to the destination market. The values of parameters and exogenous variables used to generate the figure are discussed in Section 6, where we engage in a counterfactual analysis. The horizontal axis sorts qualities; the vertical axis the relative sea-to-air distances. For each quality-distance pair, the downward-sloping blue line portrays the transport-mode threshold $\tilde{\lambda}_x$; the upward-sloping red line the entry threshold $\hat{\lambda}_x$. The figure is drawn from a $t = 0$ perspective. As a result, the $\tilde{\lambda}_x$-curve represents the locus of the expected mode thresholds at the initial date, whose values will eventually depend on the endowment realization at $t = 1 - x$. The $\hat{\lambda}_x$-curve divides the entry actual choices, since these are taken at $t = 0$. The grayed-out area on the left of the entry threshold curve depicts varieties that are not exported. Conversely, the varieties located on the right of the curve will be shipped to the destination market. Those above the transport-mode threshold curve fly, those below sail.

The result in Proposition 3 helps interpret the remaining results in Table 3. In particular, the fact that the quality gap $\mu_x^A - \mu_x^S$ narrows as x increases is consistent with the negative

Figure 2. Entry and Mode of Transport in the Quality–Relative-Distance Space

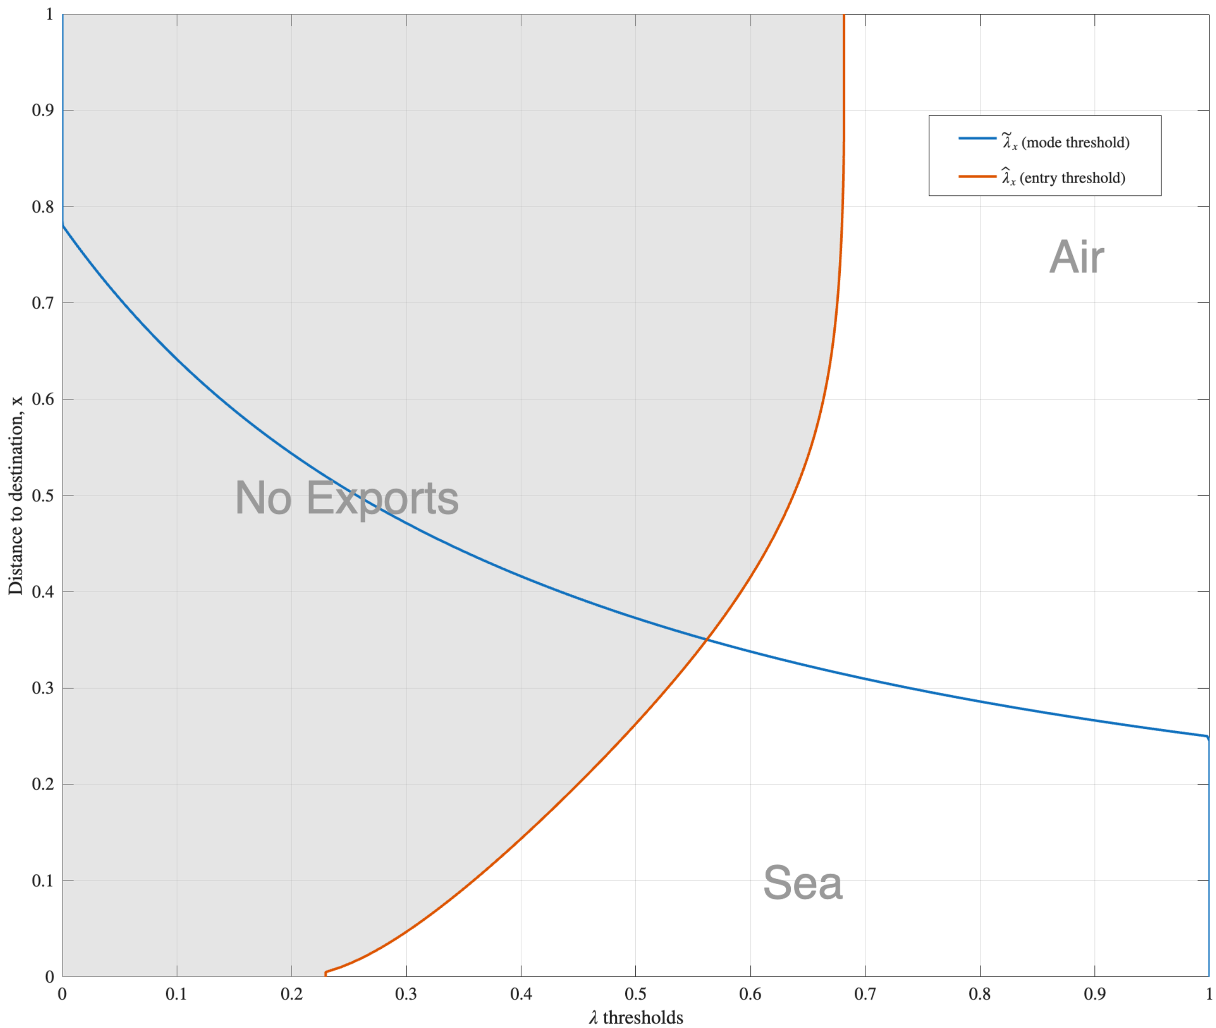

Note. The downward-sloping blue line "$\tilde{\lambda}_x$(mode threshold)" represents the expected value at time $t = 0$ of the transport mode choice threshold, which will materialize at $t = 1 - x$. The upward-sloping red line "$\hat{\lambda}_x$(entry threshold)" represents the entry choice threshold, which emerges at $t = 0$. The area on the left of the $\hat{\lambda}_x$-curve portrays varieties that are not exported. The varieties on the right of the curve reach the destination market. The ones above the $\tilde{\lambda}_x$-curve are expected to travel by air, those below the curve by sea.

coefficient estimated for the interaction term in that table. In our model, a larger relative sea-to-air distance x makes sea transport increasingly uncertain relative to the (instantaneous) air alternative, generating two effects. First, on the entry margin, a higher x reduces the attractiveness of exporting, raising the entry cutoff $\hat{\lambda}_x$. Second, on the sorting margin, it shifts the threshold $\tilde{\lambda}_x$ that separates varieties across modes. The entry effect truncates the lower end of the quality distribution served by sea, while the sorting effect reallocates varieties across modes. Together, these forces narrow the quality gap as x increases.

5.2 Variations in Relative Transport Costs

The choice of transport mode reflects a trade-off between the income risk induced by longer delivery times under sea freight and the higher monetary cost of air freight. In the model, changes in relative transport costs shift the sorting threshold, modifying the share of imports shipped by each mode and, in turn, the quality differential between air- and sea-shipped goods.

From an empirical viewpoint, a useful feature of relative transport costs is that they may vary over time for a given origin–destination pair. Our model predicts that, when both modes of transport are used by exporter x, then the average quality of their airborne exports will respond to changes in relative transport costs (τ_x), as a result of reallocations of import shares across air and sea for a given exporter.[15] We assess this mechanism by exploiting variation in air and sea transport costs at the product–origin–destination level. In our model, shifts in average quality by mode arise precisely when import shares traveling by sea or air respond to changes in relative transport costs. Accordingly, our empirical strategy instruments changes in those shares over time at the product–origin–destination level with changes in relative air-to-sea freight costs at the same level of disaggregation.

To that end, we use U.S. customs data by country of origin and district of entry for 2010–2024.[16] Exploiting the panel structure, we estimate the following two-stage least squares (2SLS) specification on HS-10 products that enter the United States by air (and for which the product–origin–district is observed using both air and sea at some point in the sample):

$$\ln(uv_{pxdt}^{Air}) = \alpha \cdot \ln(airfreight_{pxdt}) + \beta \cdot \ln\left(sea\_share_{pxdt}\right) + \xi_{pxd} + \varsigma_{pt} + \epsilon_{pxdt}, \tag{15}$$

The dependent variable (uv_{pxdt}^{Air}) is the average unit value (measured based on FOB prices) of product p sourced from country x arriving in the U.S. district d in year t when shipped by air. The main regressors are: i) $\ln(airfreight_{pxdt})$, the (log) ratio of total air freight cost to the total amount of air-shipped imports for (p, x, d) in year t; ii) $\ln(sea\_share_{pxdt})$, the log share of imports for (p, x, d) shipped by sea in year t. The coefficient on $airfreight_{pxdt}$ captures the direct effect of air transport costs on the quality of exported products in the absence of re-sorting across modes, so we expect $\alpha > 0$. The coefficient on sea_share_{pxdt} captures the sorting mechanism: as the sea share rises, the average quality of airborne imports increases, implying $\beta > 0$. The specification includes product–origin–district fixed effects ξ_{pxd} and product-year

[15] Recall that exporters who exhibit mode-mixing satisfy $0 < \hat{\lambda}_x < \tilde{\lambda}_x < \bar{\lambda}$. Thus, for them, the impact on the entry threshold ($\hat{\lambda}$) will only affect the average quality of their maritime exports, but will not have any effect on the average quality of their airborne exports (which will only be affected by shifts of $\tilde{\lambda}_x$).

[16] The panel includes data mirroring the structure of the 2024 cross-sectional dataset used in Section 2.

Table 4. Relative Transport Cost and Average Quality

	$\ln\left(uv^{Air}\right)$ (1)	$\ln\left(uv^{Air}\right)$ (2)	$\ln\left(uv^{Air}\right)$ (3)	$\ln\left(\frac{uv^{Air}}{uv^{Sea}}\right)$ (4)
	Panel A. Two-Stage Least Squares			
(log) share seaborne imports	0.559^{***}	0.517^{***}	0.512^{***}	1.860^{***}
	(0.032)	(0.037)	(0.037)	(0.045)
(log) air freight cost	0.174^{***}	0.164^{***}	0.164^{***}	0.107^{***}
	(0.007)	(0.007)	(0.007)	(0.006)
observations	584,240	309,656	307,487	584,240
district-exporter-product FE	yes	yes	yes	yes
product-year FE	yes	no	no	yes
exporter-product-year FE	no	yes	yes	no
district-product-year FE	no	yes	yes	no
district-exporter-year FE	no	no	yes	no
	Panel B. First Stage for (log) share seaborne imports			
(log) sea freight cost	-0.105^{***}	-0.122^{***}	-0.122^{***}	-0.105^{***}
	(0.004)	(0.005)	(0.005)	(0.004)
(log) air freight cost	0.067^{***}	0.081^{***}	0.080^{***}	0.067^{***}
	(0.004)	(0.006)	(0.006)	(0.004)

Note. The dependent variable in columns (1)-(3) is the (log) unit value of airborne imports and in column (4) is the (log) difference in unit values between airborne and seaborne imports. 'share seaborne imports' is the share of imports shipped via sea routes. 'air freight cost' is the ratio of total air freight cost to total air-shipped imports (measured in weight). Panel A reports the two-stage least square estimates, instrumenting '(log) share seaborne imports' with '(log) sea freight cost', defined as the (log) ratio of total sea freight cost to the total sea-shipped imports (measured in weight). Panel B reports the corresponding first stage estimates. Robust standard errors clustered at the exporter-district level in parenthesis. Significance levels: $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

fixed effects ς_{pt}. With δ_{pxd}, identification comes from within-route time variation in modal shares and air freight costs, whereas ς_{pt} absorbs common product-specific price movements over time.

Because sea_share_{pxdt} is endogenous, we isolate the sorting channel driven by relative transport costs by instrumenting $\ln(sea\_share_{pxdt})$ with $\ln(seafreight_{pxdt})$, defined as the log ratio of total sea freight cost to the total amount of of sea-shipped imports for (p, x, d) in year t. Since the regression (15) conditions on $\ln(airfreight_{pxdt})$, the exclusion restriction requires that time variation in sea freight costs affects the average unit value of airborne imports only through its impact on the quality cutoff that allocates shipments between sea and air (i.e., through sorting), and not through any direct effect on air-shipped unit values.

The results of the two-stage least square estimation of (15) are displayed on column (1) of Table 4. Both coefficients are positive and statistically significant, aligning with the model's predictions. Our main coefficient of interest in this specification is the one associated to $\ln(sea\_share_{pxdt})$. The estimate implies that a 10% increase in the sea share—arising from an increase in the relative cost of air versus sea freight—is associated with roughly a 6% increase in air-shipped unit values. This pattern is consistent with quality sorting: relative cost movements shift the cutoff between modes, so a larger fraction of lower-quality varieties travels by sea and the average quality (and unit value) of air shipments increases.[17]

Columns (2) and (3) sequentially introduce additional fixed effects as robustness checks. Specification (2) includes product–origin–year and product–district–year fixed effects, which absorb country-specific product shocks (e.g., industry-level productivity changes within a country) and district-specific demand or preference shocks, respectively. These fixed effects mitigate concerns about instrument exogeneity, as they account for product-level supply- and demand-side factors that may jointly influence freight costs and export prices. In specification (3), we further add origin–district–year fixed effects, capturing shocks that affect bilateral trade links between the exporter country and the U.S. district of entry. In both specifications, the core coefficients retain their sign and precision, in line with the proposed mechanism.

Finally, column (4) re-estimates the baseline specification replacing the dependent variable with the log difference in unit values between air- and sea-shipped imports. The results show that increases in the sea share are accompanied by a widening air–sea unit-value gap, further corroborating the sorting mechanism operating through changes in relative transport costs.

6 Counterfactual Analysis

We now use the model to quantify how changes in demand uncertainty, transport cost, and production cost reshape export participation and mode choice along the quality (λ) and distance (x) dimensions. In practical terms, we examine the expected profit (before entry cost) at time $t = 0$ as measured by the left-hand side of (14) for a pair of representative quality levels and distances, and visually represent the entry and transport-mode decisions for all qualities and

[17]Panel B of Table 4 indicates that the sea-share variation used for identification in Panel A is generated by changes in relative transport costs: conditional on air-freight costs, higher sea-freight costs reduce the sea share. This pattern confirms that the instrument shifts modal shares through relative cost movements.

distances building on Figure 2.

For these numerical exercises, we adopt a simple parameterization that anchors the model to U.S. data while preserving transparency. We normalize the marginal cost of production for the fringe competition to $\bar{c} = 1$ and bound the quality index on $\left[0, \bar{\lambda}\right]$ with $\bar{\lambda} = 1$. These standard choices fix units and have no impact on comparative statics.

The marginal production cost for foreign exporters is set to $c = 0.6$, so that at the competitive fringe price $\bar{c}$ the implied price–cost margin is $1 - c/\bar{c} = 40\%$. U.S. studies typically find average markups in the 20–60% range with a fat upper tail (De Loecker, Eeckhout, and Unger 2020), so a 40% margin is conservative but well within the observed U.S. distribution. Interpreting $\bar{c}$ as the cost of a local "copycat", this configuration delivers a plausible mark-up on imported varieties.

The air-shipping premium is set to $\tau_x = \tau = 1.48$, interpreted as an ad-valorem wedge relative to sea. Using U.S. import data, Hummels and Schaur (2013) estimate that a one-day reduction in transit time is worth roughly 0.6–2.1% ad valorem; over typical 20–30-day ocean lags, this implies a 12–63% willingness to pay for faster modes. At the same time, freight-rate data indicate that air cargo is typically priced at an order of magnitude higher per kilogram than ocean shipping (IATA, 2021; World Bank, 2009); because these per-kg charges represent a small share of product value for the high-value, time-sensitive goods that select into air, the implied ad-valorem wedge is much smaller than the raw rate ratio. Choosing τ within the 1.15–1.6 range is therefore empirically plausible.[18]

We set the sunk export-market entry cost to $\Upsilon = 0.02$, which corresponds to 1.8% of the calibrated model's initial income per capita. This magnitude is economically modest relative to observed trade-cost components. For the United States, the World Bank Doing Business *Trading Across Borders* module reports documentary and border compliance costs on the order of 0.4% of income per capita. However, entering and serving a foreign market typically entails additional fixed costs—such as logistics, distribution, and market development expenses—that are often of larger magnitude (e.g., Das, Roberts, and Tybout, 2007). Our calibration thus implies a sunk market-access cost that exceeds the compliance burden without being disproportionately large, yielding a non-trivial but not dominant extensive margin.

The demand shifter Y_t follows a geometric Brownian motion without drift ($\mu = 0$) and annualized volatility $\sigma = 0.10$. Under GBM, σ is exactly the standard deviation of the log growth rate per unit time. For the U.S., the standard deviation of four-quarter real GDP growth

[18] Per-kg freight-rate ratios do not map one-for-one into ad-valorem wedges because the latter depend on value-to-weight ratios.

fell to roughly 1.5 percentage points after the mid-1980s "Great Moderation". Taking $\sigma = 0.10$ therefore constitutes a deliberate "stress test" for income risk; values in the $\sigma \in [0.03, 0.08]$ range would be more conservative for U.S. macro volatility (Stock and Watson 2003).

We have verified that the qualitative behavior of the entry and mode thresholds $\hat{\lambda}_x$ and $\tilde{\lambda}_x$ is robust to parameter variations within ranges that remain compatible with this evidence: $\tau \in [1.15, 1.60]$, $\Upsilon \in [0.01, 0.10]$, $c \in [0.6, 0.8]$, and $\sigma \in [0.03, 0.15]$. Within these intervals, the comparative-statics patterns emphasized below are unchanged: low-quality varieties serving nearby markets sail; high-quality varieties, especially on long routes, fly; and increases in distance or volatility lower the mode threshold and raise the entry cutoff.

6.1 Variations in Income Volatility

We begin by examining how a mean-preserving increase in income volatility affects exporters' entry and mode choice. In the model, this corresponds to a rise in the volatility parameter σ of the geometric Brownian motion governing Y_t, holding its drift and initial level fixed. By Corollary 1, higher volatility makes the risky sea option less attractive: the mode threshold $\tilde{\lambda}_x$ falls, so that, conditional on exporting, a larger set of varieties finds it optimal to rely on air transport. Intuitively, more volatile income raises the value of flexibility provided by fast delivery.

This exercise reveals nontrivial implications for export variety and route re-sorting, which prove especially intriguing in light of the recent work re-examining the Great Moderation in the context of exchange-rate volatility (Stavrakeva and Tang, 2024), cross-country inflation dynamics (Stracca, 2025), and supply chain disruption and energy shocks (De Santis and Tornese, 2025). These contributions cast doubt on the idea that macro volatility has kept trending down, and suggest that the post-COVID environment is characterized by more frequent and complex macro shocks relative to the Great Moderation benchmark. Our counterfactual captures the impact of a shift from a relatively stable macro environment to one in which destination markets are hit by larger and more frequent shocks.

Figure 3 illustrates these comparative statics in the quality–distance plane. Relative to the baseline in Figure 2, the increase in σ shifts the mode threshold $\tilde{\lambda}_x$ downward across all (relative sea-to-air) distances: for any given x, the quality level at which firms switch from sea to air falls, reflecting the greater option value of air shipping under more volatile demand. The adjustment in the entry threshold $\hat{\lambda}_x$ is more nuanced. Its curvature becomes more pronounced, yielding a flatter segment for routes with small sea-to-air distance gaps and a steeper segment for routes with wider gaps, where volatility and transit-time exposure interact to sharpen

Figure 3. Increase in Income Volatility

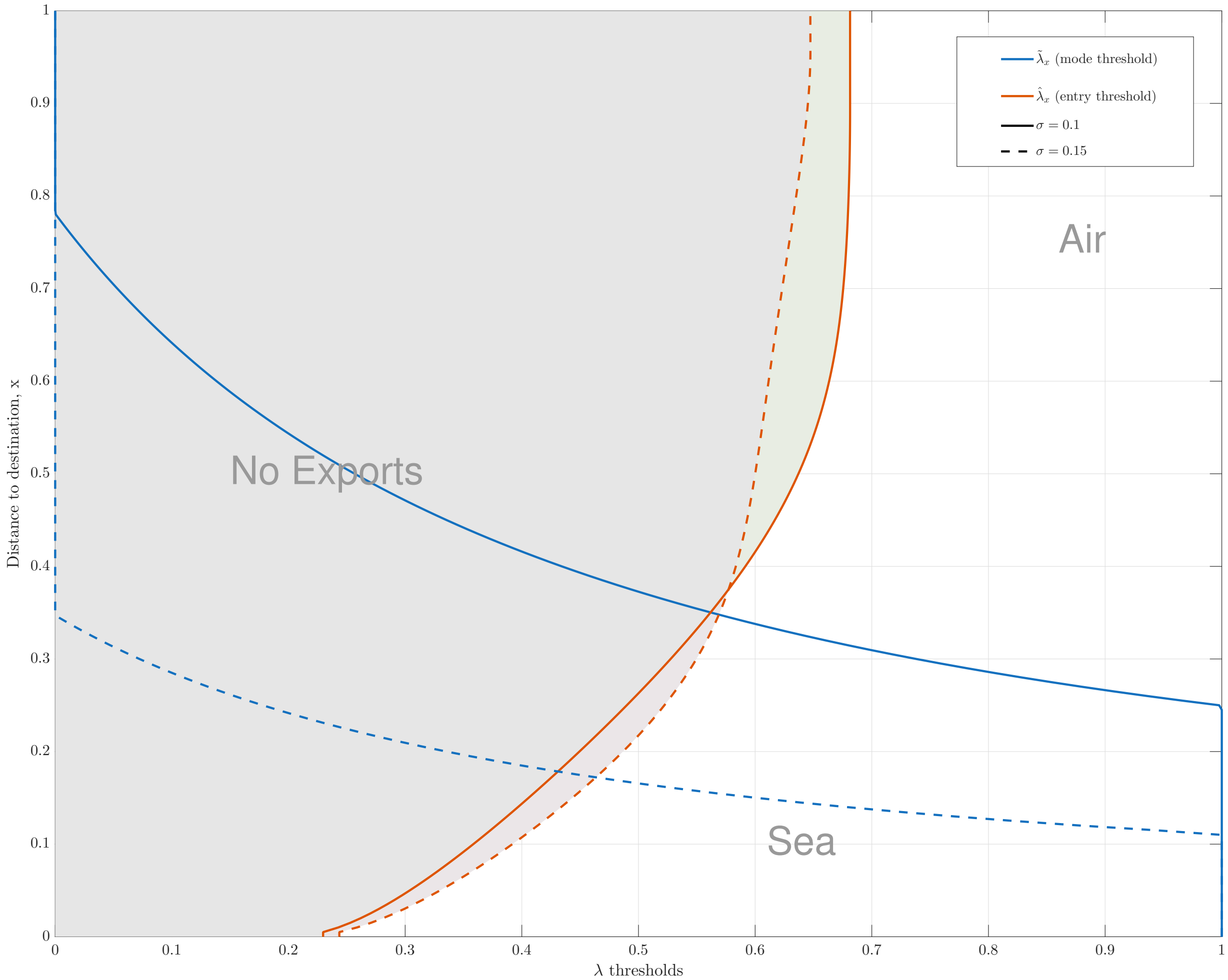

Note. The downward-sloping blue lines "$\tilde{\lambda}_x$(mode threshold)" represent the expected value at time $t = 0$ of the transport mode choice threshold, which will materialize at $t = 1 - x$. The upward-sloping red lines "$\hat{\lambda}_x$(entry threshold)" represent the entry choice threshold, which emerges at $t = 0$. The area on the left of the $\hat{\lambda}_x$-curve portrays varieties that are not exported. The varieties on the right of the curve reach the destination market. The ones above the $\tilde{\lambda}_x$-curve are expected to travel by air, those below the curve by sea. The solid lines represent the benchmark scenario, the dashed lines the one emerging from the comparative statics exercise. Relative to the baseline in Figure 2, the increase in endowment volatility shifts the mode threshold downward, while the entry threshold rotates counter-clockwise, expanding the region of high-quality, long-distance varieties served by air, reallocating mid-to-high quality exporters either into air shipping or out of the export market and cutting out a wider range of low-to-mid quality, short-to-intermediate-distance varieties.

the difference between shipping modes. The combined effect is tighter selection among lower-quality varieties—especially at short-to-intermediate distances—and an expansion of the region in which higher-quality varieties enter and ship by air.

To gauge the magnitude of these effects, we counterfactually impose a 50% increase in volatility, from $\sigma = 0.10$ to $\sigma = 0.15$, in a mean-preserving spread of the demand process, and quantify how this affects expected profits computed at $t = 0$ (i.e., at the export market entry moment) for representative varieties at different combinations of quality and relative sea-to-air distance. For low-quality varieties ($\lambda = \bar{\lambda}/3$) shipped over short distances ($x = 1/3$), expected profits fall by 9.31%. At the same quality but over a longer route ($x = 2/3$), the drop reaches 10.07%, about 1.08 times larger than for low-quality, short-distance exporters. At the opposite end of the quality spectrum, high-quality varieties ($\lambda = 2\bar{\lambda}/3$) shipped from nearby exporters see expected profits rise by 3.81%, a swing of roughly 13.1 percentage points relative to their low-quality counterparts. At the wider route gap, the gain for the same high-quality class reaches 6.43%, roughly 16.5 percentage points better than low-quality exporters and about 1.69 times the increase experienced by high-quality, short-distance exporters.

These numbers reveal a strong complementarity between quality, relative distance, and volatility in shaping the incidence of macro-risk variations. For low-quality exporters, higher volatility is unambiguously bad: they rely primarily on sea shipping, cannot profitably absorb the cost of switching to air, and are thus exposed to adverse income realizations. For high-quality exporters, volatility turns out to be beneficial. The intuition for this rests on the demand of high-quality varieties being more income-elastic. As a consequence, the upside from good states is magnified; and because they optimally choose air in a wider range of scenarios when σ rises, they are better able to limit their downside exposure inherent in sea shipping.

This exercise extends the findings of the existing literature connecting greater volatility to a transport mode switching in favor of airborne freight by identifying which varieties are involved in the switch along the quality dimension. It also unveils heterogeneous effects in export participation along the quality dimension. Furthermore, it complements recent contributions on U.S. ports such as Brancaccio, Kalouptsidi, and Papageorgiou (2024), who show that increases in aggregate demand volatility sharply raise the welfare gains from infrastructure investment, precisely because congestion costs are convex in demand and shocks are transmitted through the transport network. In our model, volatility plays an analogous role on the exporter side: it amplifies the value of flexibility, tilts sorting further toward air for high-quality, long-distance exporters, and intensifies the selection of who participates in trade in the first place.

6.2 Variations in the Relative Cost of Air Transport

We next consider an increase in the relative cost of air transportation, holding the sea cost fixed. In the model, this corresponds to a rise in τ, which, by Corollary 1, shifts up the mode threshold $\tilde{\lambda}_x$ for all distances: a higher quality level is required to justify paying for speed.

This experiment may be interpreted in light of the COVID-19 pandemic, when air traffic was the first transport mode to be severely disrupted. Travel bans and the collapse of passenger flights sharply reduced belly-hold cargo capacity, with global air-cargo capacity falling by around 30–35% in 2020 relative to 2019 and air-freight rates spiking on many routes. At the same time, container shipping also came under pressure, but with a lag: blank sailings and port congestion eventually pushed container freight indices to historical highs from late 2020 into 2021 (USITC, 2021; Pulido, 2023). Our counterfactual thus refers to the early phase of the crisis, when the effective increase in the relative price of air versus sea was particularly pronounced.

Figure 4 depicts how a marginal increase in τ reshapes the quality–distance tradeoff. Relative to the baseline in Figure 2, the air–sea cutoff $\tilde{\lambda}_x$ shifts upward: at any given distance, only higher-quality varieties continue to fly. The entry threshold $\hat{\lambda}_x$ also moves rightward, especially at long distances, reflecting the fact that higher air costs compress expected profits and make entry less attractive precisely on routes where the value of speed was previously greatest. The region in the figure corresponding to "high quality and long distance" shrinks markedly, while mid-to-high quality exports at intermediate distances that do not exit altogether are rerouted from air to sea.

The profit effects of this variation are highly heterogeneous. For a 1% increase in the relative air-to-sea cost, our simulations imply that, for low-quality varieties ($\lambda = \bar{\lambda}/3$) shipped from a short sea distance ($x = 1/3$), expected profits fall by 1.48%. For the same variety class exported from a longer sea distance ($x = 2/3$) the drop is 2.86%, 1.93 times larger than for short-distance exporters. Profits for high-quality producers ($\lambda = 2\bar{\lambda}/3$) at $x = 1/3$ fall by 2.72%, 1.84 times their low-quality counterparts. The decline reaches 7.29% for high-quality varieties sourced from an exporter located at distance $x = 2/3$ (i.e., 2.55 times the loss for low-quality, long-distance exporters and 2.68 times that for high-quality, short-distance exporters).

These elasticities illustrate a strong complementarity between quality and distance in exposure to air-cost shocks. High-quality varieties are more income-elastic and thus more reliant on the flexibility afforded by fast delivery; distant routes, in turn, magnify the income risk associated with shipping by sea. An increase in τ thus disproportionately hits those exporters for whom the option value of faster air routes is most valuable—high-quality varieties produced in countries with longer sea distances.

Mid-quality producers at intermediate distances are affected in two ways. First, those whose pre-shock profits were close to the entry threshold see expected payoffs fall below Υ and exit. Second, others remain profitable but switch from air to sea: they cross the new, higher $\tilde{\lambda}_x$

Figure 4. Increase in the Air Transport Relative Cost

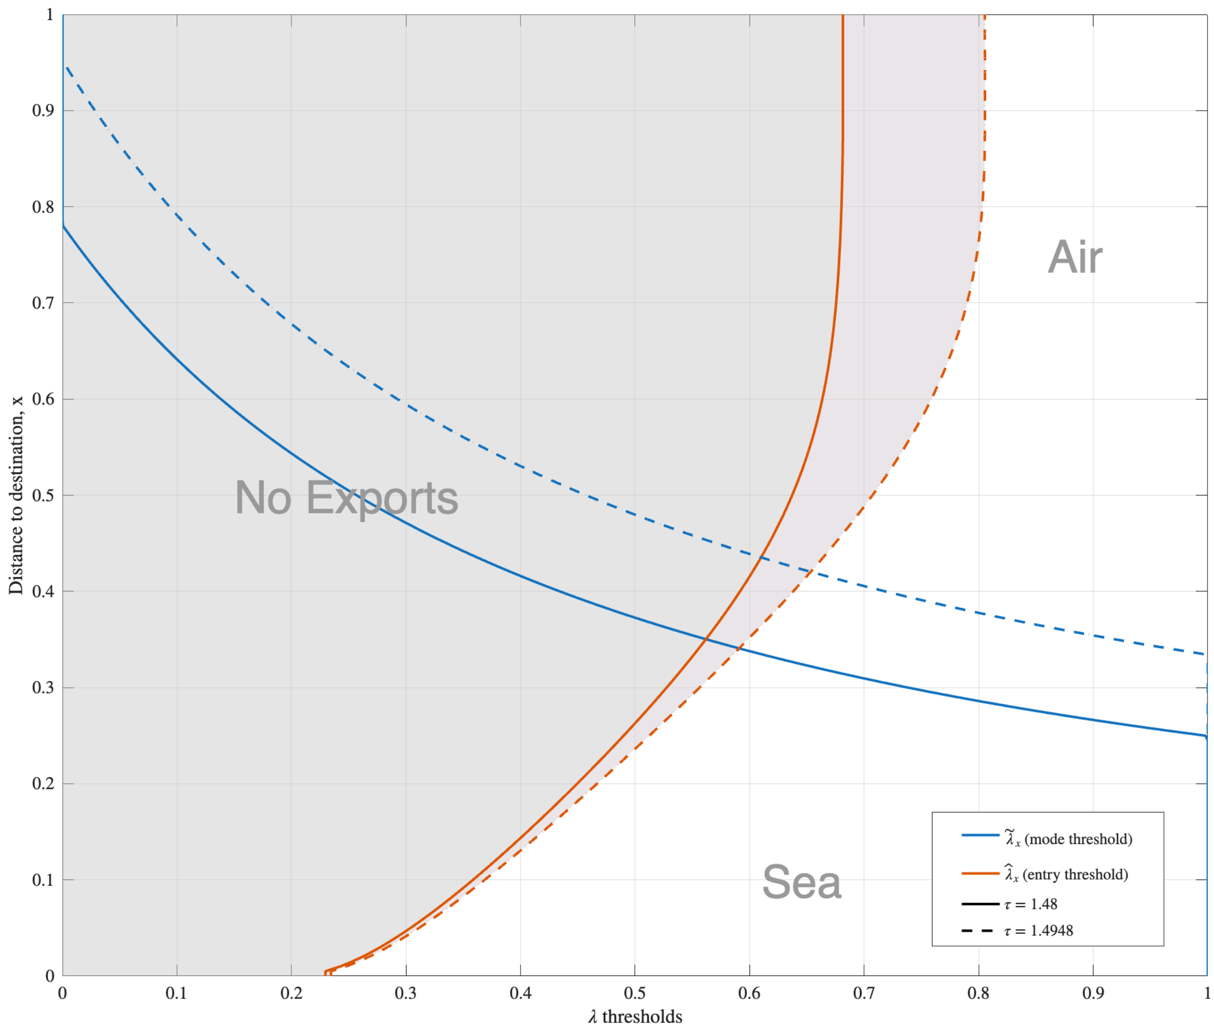

Note. The downward-sloping blue lines "$\tilde{\lambda}_x$(mode threshold)" represent the expected value at time $t = 0$ of the transport mode choice threshold, which will materialize at $t = 1 - x$. The upward-sloping red lines "$\hat{\lambda}_x$(entry threshold)" represent the entry choice threshold, which emerges at $t = 0$. The area on the left of the $\hat{\lambda}_x$-curve portrays varieties that are not exported. The varieties on the right of the curve reach the destination market. The ones above the $\tilde{\lambda}_x$-curve are expected to travel by air, those below the curve by sea. The solid lines represent the benchmark scenario, the dashed lines the one emerging from the comparative statics exercise. Relative to the baseline in Figure 2, the increase in the air-to-sea cost shifts the mode threshold upward and the entry threshold rightward, shrinking the region of high-quality, long-distance varieties served by air and reallocating mid-to-high quality exporters either into sea shipping or out of the export market altogether.

cutoff and no longer find it worthwhile to pay for speed. For these firms, the shock bites twice: profits are lower ex ante, and ex post they operate in a riskier environment, since quantities must be committed earlier and are exposed to the truncation of the income distribution under sea shipping.

6.3 Increase in Marginal Production Costs

Finally, we consider an increase in the marginal production cost c, holding $\bar{c}$ fixed. In practice, this can be interpreted as a rise in trade costs that acts like a proportional cost wedge on foreign varieties—for instance, an increase in U.S. import tariffs on manufactured goods. Since 2018, the U.S. has implemented several waves of tariff hikes, initially in the context of the U.S.–China trade war and subsequently broadened and partially reshaped under later administrations. Empirically, tariffs were raised on thousands of Chinese products to around 25% on average in 2018–2019, and more recent measures have extended higher import duties to sectors such as steel, aluminum, electric vehicles, and semiconductors (Asdourian and Wessel, 2025). These episodes motivate our second counterfactual as a generic "tariff-type" cost shock.

In the model, a higher c has two main effects, which are graphically represented in Figure 5. First, it erodes the markup relative to the domestic fringe, depressing expected profits and shifting the entry threshold $\hat{\lambda}_x$ to the right; this represents a standard discouragement effect of tariffs on imports. Second, $\tilde{\lambda}_x$, the modal-sorting threshold, moves upwards: as production becomes more expensive, the value of being able to better fine-tune quantities (allowed by air shipping) decreases, so for a given distance firms are less willing to pay the air premium. Unlike the air-cost shock—which primarily reallocates varieties from air to sea—an increase in c primarily operates through the entry margin, with a more subtle rebalancing of mode choice.

The profit effects of a 1% increase in c are again heterogeneous across the quality–distance space, greater in magnitude in general but less volatile across varieties than in the air-cost experiment. The expected profits fall by 4.11% for low-quality varieties shipped from a short distance. The drop is 5.09% for varieties of the same quality but sources from a long distance, only 1.24 times larger than for low-quality, short-distance exporters. Producers of high-quality products exported from a short distance face a profit fall by 5.02%, about 1.22 times the low-quality, short-distance benchmark. From a long distance, the decline for this class of varieties reaches 7.71%, roughly 1.52 times the loss for low-quality, long-distance exporters and 1.54 times that for high-quality, short-distance exporters.

Relative to the modal-cost shock, the amplification across quality and distance is more muted: all exporters see a sizable proportional drop in profits, with somewhat larger losses for high-quality and/or distant varieties, but the ratios are closer to one. This reflects the nature of the shock. A higher c compresses margins on every unit sold, irrespective of mode, so the main channel is a roughly proportional squeeze on profits, reinforced—but not dominated—by the stronger variability of the demand shifter at higher λ and by the longer exposure to income risk on distant routes.

Figure 5. Increase in the Exporter's Marginal Cost of Production

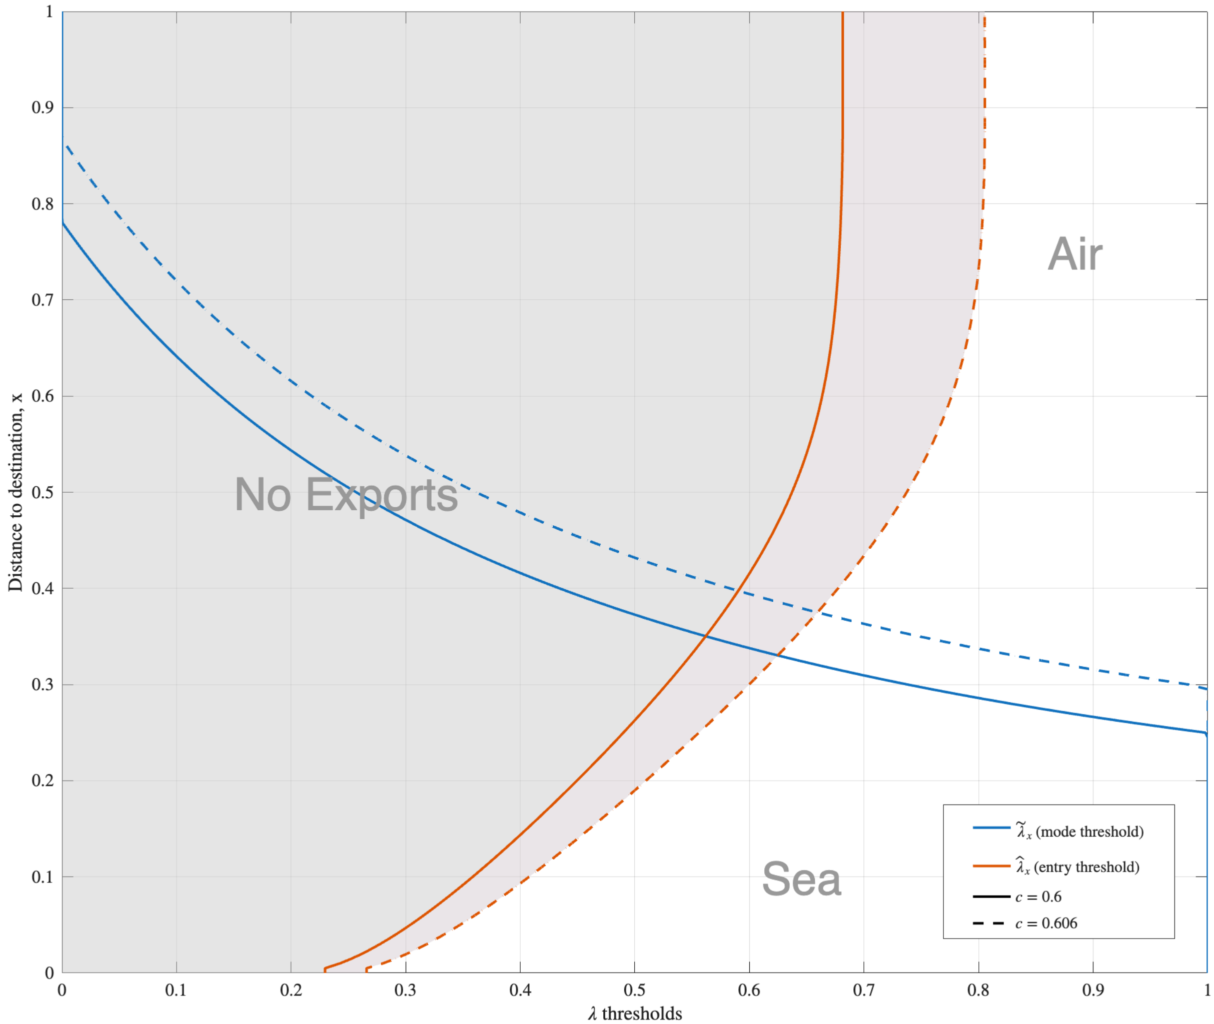

Note. The downward-sloping blue lines "$\tilde{\lambda}_x$(mode threshold)" represent the expected value at time $t = 0$ of the transport mode choice threshold, which will materialize at $t = 1 - x$. The upward-sloping red lines "$\hat{\lambda}_x$(entry threshold)" represent the entry choice threshold, which emerges at $t = 0$. The area on the left of the $\hat{\lambda}_x$-curve portrays varieties that are not exported. The varieties on the right of the curve reach the destination market. The ones above the $\tilde{\lambda}_x$-curve are expected to travel by air, those below the curve by sea. The solid lines represent the benchmark scenario, the dashed lines the one emerging from the comparative statics exercise. Relative to the baseline in Figure 2, the increase in the marginal cost of production shifts the entry threshold markedly to the right and moves the mode threshold upward, tightening export participation overall and inducing a more selective use of air transport that is concentrated on higher-quality varieties, especially at longer distances.

7 Conclusion

This paper has studied how the trade-off between freight costs and timeliness of delivery shapes firms' modal choices in international trade of quality-differentiated products. The interplay between nonhomothetic demand and income uncertainty tilts the trade-off in favor of speed for higher-quality varieties. As a consequence, exporters sort by mode of transport along the quality dimension: high-quality varieties fly and lower-quality varieties sail. In addition, as income uncertainty increases with transit delays, the quality gap between air- and sea-shipped goods shrinks as sea routes lengthen relative to air. Our findings thus link the classic speed-versus-cost trade-off to endogenous quality choice under demand risk.

Furthermore, when serving a market incurs an entry cost, this mechanism operates on the extensive margin as well: participation becomes a quality–distance selection problem rather than a separable productivity cutoff. In our framework, entry and mode choice move together, since the option value of timeliness affects both the profitability of exporting and modal choice. This implies that policies or disruptions that affect the relative cost or reliability of modes can change not only how firms ship, but also which exporters remain active in a market once entry is costly.

We substantiate these mechanisms through counterfactuals. A COVID-type shock that raises the relative cost of air freight primarily reallocates varieties from air to sea and hits high-quality, long-distance exporters hardest, both by compressing their profits and by exposing them to greater demand risk. A tariff-type increase in marginal production costs reduces profits more evenly but still disproportionately affects high-quality and distant varieties, chiefly by pushing marginal exporters out of the market and modestly shifting the mode threshold. In both cases, shocks to trade and logistics costs do not merely dampen total trade volumes; they also reshape the composition of exporters and their modal choice, with particularly pronounced consequences for premium varieties. A mean-preserving increase in destination-demand volatility further shifts high-quality, long-distance exporters toward air while tightening selection among marginal exporters, making the distributional incidence of macro risk strongly quality- and geography-dependent.

We have also shown that shifts in relative modal prices, in geographical frictions, or in the volatility of destination demand reallocate varieties across air and sea. Our empirical and numerical exercises suggest that policy-induced changes in modal costs or reliability (e.g., carbon pricing for air cargo, port congestion, or insurance premia tied to maritime delays) may generate non-trivial composition effects on traded quality and on the distribution of profits across exporters, not just on aggregate trade volumes. From this viewpoint, it indicates a natural complementarity with recent work on port infrastructure and congestion, which shows that volatility can sharply increase the welfare returns to relaxing capacity constraints in key nodes of the transport network. A natural next step is to embed these shocks in a richer welfare and incidence analysis, allowing for stochastic transit delays and general-equilibrium feedbacks on prices and entry. Such extensions would help assess how recent supply-chain disruptions or green-shipping mandates shift the quality thresholds that govern modal sorting, and who ultimately gains or loses from the resulting reorganization of trade flows.

References

[1] Alessandria, G., Kaboski, J.P. and Midrigan, V., 2010. Inventories, lumpy trade, and large devaluations. *American Economic Review*, 100(5), pp.2304-39.

[2] Asdourian, E., & Wessel, D. (2025). What are tariffs, and why are they rising? Brookings Institution, available at brookings.edu/articles/what-are-tariffs-and-why-are-they-rising.

[3] Baldwin, R. and Harrigan, J., 2011. Zeros, quality, and space: Trade theory and trade evidence. *American Economic Journal: Microeconomics*, 3(2), pp.60-88.

[4] Bastos, P. and Silva, J., 2010. The quality of a firm's exports: Where you export to matters. *Journal of International Economics*, 82(2), pp.99-111.

[5] Bernard, A.B. and Jensen, J.B., 2004a. Exporting and productivity in the USA. *Oxford Review of Economic Policy*, 20(3), 343–357.

[6] Bernard, A.B. and Jensen, J.B., 2004b. Why some firms export. *Review of Economics and Statistics*, 86 (2), 561–569.

[7] Blaum, J., Esposito, F. and Heise, S., 2025 Input sourcing under supply chain risk: evidence from us manufacturing firms. *FRB of New York Staff Report 1141.*

[8] Brancaccio, G., Kalouptsidi, M. and Papageorgiou, T., 2024. Investment in Infrastructure and Trade: The Case of Ports. NBER Working Papers 32503.

[9] Chen, N. and Juvenal, L., 2022. Markups, Quality, and Trade Costs. *Journal of International Economics*, 137, 103627.

[10] Das, S., Roberts, M.J. and Tybout, J.R., 2007. Market entry costs, producer heterogeneity, and export dynamics. *Econometrica*, 75(3), pp.837–873.

[11] De Loecker, J., Eeckhout, J. and Unger G., 2020. The rise of market power and the macroeconomic implications." *Quarterly Journal of Economics*, 135(2), 561–644.

[12] De Santis, R.A. and Tornese, T., 2025. Macroeconomic regime change and the size of supply chain disruption and energy supply shocks. *European Economic Review*, 105077.

[13] Dingel, J.I., 2017. The determinants of quality specialization. *Review of Economic Studies*, 84(4), pp.1551-1582.

[14] Djankov, S., Freund, C. and Pham, C. S., 2010. Trading on Time. *Review of Economic and Statistics*, 92(1), pp.166-173.

[15] Evans, C. and Harrigan, J., 2005. Distance, Time, and Specialization: Lean Retailing in General Equilibrium. *American Economic Review*, 95(1), pp. 292-313.

[16] Fajgelbaum, P., Grossman, G.M. and Helpman, E., 2011. Income distribution, product quality, and international trade. *Journal of Political Economy*, 119(4), pp.721-765.

[17] Feenstra, R.C. and Romalis, J., 2014. International prices and endogenous quality. *Quarterly Journal of Economics*, 129(2), pp.477-527.

[18] Hallak, J.C., 2006. Product quality and the direction of trade. *Journal of international Economics*, 68(1), pp.238-265.

[19] Harrigan, J., 2010. Airplanes and comparative advantage. *Journal of International Economics*, 82(2), pp.181-194.

[20] Harrigan, J., Ma, X. and Shlychkov, V., 2015. Export prices of US firms. *Journal of International Economics*, 97(1), pp.100-111.

[21] Hornok, C. and Koren, M., 2015. Per-shipment costs and the lumpiness of international trade. *Review of Economics and Statistics*, 97(2), pp.525-530.

[22] Hummels, D., 2007. Transportation costs and international trade in the second era of globalization. *Journal of Economic Perspectives*, 21(3), pp.131-154.

[23] Hummels, D.L. and Schaur, G., 2010. Hedging price volatility using fast transport. *Journal of International Economics*, 82(1), pp.15-25.

[24] Hummels, D.L. and Schaur, G., 2013. Time as a trade barrier. *American Economic Review*, 103(7), pp.2935-59.

[25] Hummels, D.L. and Skiba, A., 2004. Shipping the good apples out? An empirical confirmation of the Alchian-Allen conjecture. *Journal of Political Economy*, 112(6), pp.1384-1402.

[26] International Air Transport Association (IATA), 2021. *Competitive advantage of air cargo: Not only speed but also price*. IATA Economics, Economic Report, May 28, 2021.

[27] Jaimovich, E., Madzharova, B. and Merella, V., 2023. Inside the white box: Unpacking the determinants of quality and vertical specialization. *European Economic Review*, 152, 104374.

[28] Jaimovich, E. and Merella, V., 2012. Quality ladders in a Ricardian model of trade with nonhomothetic preferences. *Journal of the European Economic Association*, 10(4), pp.908-937.

[29] Jaimovich, E. and Merella, V., 2015. Love for quality, comparative advantage, and trade. *Journal of International Economics*, 97(2), pp.376-391.

[30] Lugovskyy, V. and Skiba, A., 2016. Positive and negative effects of distance on export prices. *Journal of Economic Behavior & Organization*, 127, pp.155-181.

[31] Manova, K. and Zhang, Z., 2012. Export prices across firms and destinations. *Quarterly Journal of Economics*, 127(1), pp.379-436.

[32] Martincus, C. V., Carballo, J. and Graziano, A., 2015. Customs. *Journal of international economics*, 96 (1), pp. 119-137.

[33] Mohammed, S. and Williamson, J., 2004. Freight rates and productivity gains in British tramp shipping 1869–1950. *Explorations in Economic History*, 41 (2), pp. 172-203

[34] Pulido, J., 2023. Pandemic-induced increases in container freight rates: Assessing their domestic effects in a globalised world. BIS Working Papers 1132, Bank for International Settlements.

[35] Schott, P.K., 2004. Across-product versus within-product specialization in international trade. *Quarterly Journal of Economics*, 119(2), pp.647-678.

[36] Stock, J.H. and Watson, M.W., 2003. Has the business cycle changed? Evidence and explanations. In *Monetary Policy and Uncertainty: Adapting to a Changing Economy,* Jackson Hole Symposium, Federal Reserve Bank of Kansas City, 9–56.

[37] Stracca, L., 2025. The Great Moderation at 40: learning from the cross section. ECB Working Paper Series No. 3124.

[38] Stavrakeva, V. and Tang, J., 2025. Explaining the great moderation exchange rate volatility puzzle. *IMF Economic Review*, 73(1), pp.196-238.

[39] USITC, 2021. The impact of the COVID-19 pandemic on freight transportation services and U.S. merchandise imports. In *Shifts in U.S. Merchandise Trade*, United States International Trade Commission, Washington, DC.

[40] Verhoogen, E.A., 2008. Trade, quality upgrading, and wage inequality in the Mexican manufacturing sector. *Quarterly Journal of Economics*, 123(2), pp.489-530.

[41] World Bank, 2009. *Air Freight: A Market Study with Implications for Landlocked Countries*. Transport Papers TP-26, Washington, DC: World Bank.

A Appendix

A.1 Auxiliary Results and Proofs

Derivation of (1). Consider the indirect utility function

$$u = F\left[\left(\frac{Y}{a(\mathbf{p})}\right)^{\frac{1}{b(\mathbf{p})}}\right],$$

with F a well-behaved increasing function, and $a(\mathbf{p})$ and $b(\mathbf{p})$ price aggregators satisfying:

$$a(\mathbf{p}) = \exp\left(\bar{\alpha}\sum_{\lambda}\log P_{\lambda}\right), \quad b(\mathbf{p}) = \exp\left(\sum_{\lambda}\lambda\log P_{\lambda}\right),$$

where we have imposed constant expenditure shares ($\alpha_{\lambda} = \bar{\alpha}$) across goods of quality λ and we have set to zero all terms associated to the second order terms in the AIDS price aggregators. For each level of quality λ, let P_{λ} be a nested price aggregator given in (2). Assume that $F(\cdot) = \log(\cdot)$. Then we have:

$$u \cdot b(\mathbf{p}) = \log Y - \log a(\mathbf{p}),$$

from where $\exp\left(\int_{\mathcal{X}}\alpha_{\lambda,x}\ln p_{\lambda,x}dx\right)$ obtains after rearranging and replacing $a(\mathbf{p})$ and $b(\mathbf{p})$ with the above expressions.

Note that Deaton and Muellbauer (1980) show that the price index P associated to the nominal endowment $D \equiv P \cdot Y$ obeys:

$$\log P = \bar{\alpha}\sum_{\lambda\in\Lambda}\log P_{\lambda} = \log a(\mathbf{p}).$$

By writing (1) in terms of the real endowment, we are implicitly defining Y as the *numeraire* of the economy and letting $P = 1$, which in turn implies that $u \cdot b(\mathbf{p}) = \log Y$. ■

Proof of Proposition 1. Given the definition $\Omega_\lambda(Y) \equiv (\alpha + \lambda \log Y) Y$, then

$$E\left(\Omega_\lambda(Y) \middle| Y_{1-x}\right) = \alpha Y_{1-x} + \lambda E\left(\log Y^Y \middle| Y_{1-x}\right).$$

and

$$E\left(\Omega_\lambda(Y) \middle| Y_{1-x} \cap Y \leq \hat{Y}\right) = \alpha E\left(Y \middle| Y_{1-x} \cap Y \leq \hat{Y}_x\right) + \lambda E\left(\log Y^Y \middle| Y_{1-x} \cap Y \leq \hat{Y}_x\right).$$

Notice now that we may write:

$$E\left(Y \middle| Y_{1-x} \cap Y \leq \hat{Y}\right) = \delta(\hat{Y}) \cdot E\left(Y \middle| Y_{1-x}\right) \tag{16}$$

and

$$E\left(Y \log Y \middle| Y_{1-x} \cap Y \leq \hat{Y}\right) = \gamma(\hat{Y}) \cdot E\left(\log Y^Y \middle| Y_{1-x}\right), \tag{17}$$

where $\delta(\hat{Y}), \gamma(\hat{Y}) < 1$ and $\delta'(\hat{Y}), \gamma'(\hat{Y}) > 0$ for $\hat{Y} < Y_{\max}$, and $\lim_{\hat{Y} \to Y_{\max}} \delta(\hat{Y}) = \lim_{\hat{Y} \to Y_{\max}} \gamma(\hat{Y}) = 1$. In addition, since the function $\log Y^Y$ is strictly convex in Y, it follows that $\gamma(\hat{Y}) < \delta(\hat{Y})$ when $\hat{Y} < Y_{\max}$. Therefore, plugging (16) and (17) into (12) as:

$$\Psi_x(\lambda) \equiv \frac{\alpha Y_{1-x} + \lambda E\left(\log Y^Y \middle| Y_{1-x}\right)}{\alpha \delta(\hat{Y}) Y_{1-x} + \lambda \gamma(\hat{Y}) E\left(\log Y^Y \middle| Y_{1-x}\right)}.$$

Lastly, differentiating$\Psi_x(\lambda)$ with respect to λ yields that

$$\Psi_x'(\lambda) = \frac{\alpha\left(\delta(\hat{Y}) - \gamma(\hat{Y})\right) Y_{1-x} E\left(\log Y^Y \middle| Y_{1-x}\right)}{\left[\alpha \delta(\hat{Y}) Y_{1-x} + \lambda \gamma(\hat{Y}) E\left(\log Y^Y \middle| Y_{1-x}\right)\right]^2} > 0, \tag{18}$$

since the sign of (18) is determined by $\delta(\hat{Y}) - \gamma(\hat{Y})$. ■

Lemma 2 (Preliminary results) *Under the lognormal distribution of endowment* $Y | Y_{1-x} \sim LogNormal\left(m_x, s_x^2\right)$, *with* $m_x \equiv \ln Y_{1-x} - s_x^2/2$ *and* $s_x^2 \equiv \sigma^2 x$, *the left-hand side of (12) becomes:*

$$S = \frac{B\left(1 - \Phi_L\right) + \lambda s_x \phi_L}{B \Delta \Phi - \lambda s_x \Delta \phi}, \tag{19}$$

and its derivative with respect to a generic parameter or exogenous variable is given by:

$$S' = \frac{V_U' V_L - V_L' V_U}{\left(V_L - V_U\right)^2} = \frac{\left(1 - \Phi_L\right)\left(1 - \Phi_U\right)}{\left(V_L - V_U\right)^2} \mathcal{C}, \tag{20}$$

with:

$$\mathcal{C} \equiv \frac{V_U' V_L - V_L' V_U}{(1-\Phi_L)(1-\Phi_U)}, \tag{21}$$

and, for $j = \{L, U\}$*:*

$$V_j \equiv (1-\Phi_j)(B + \lambda s_x r_j), \qquad V_j' \equiv (1-\Phi_j)(B' + K_j r_j), \tag{22}$$

$$B \equiv \bar{\alpha} + \lambda \left(m_x + s_x^2\right) > 0, \qquad K_j \equiv (\lambda s_x)' - a_j' A_j, \qquad A_j = B + \lambda s_x a_j, \tag{23}$$

$$r_j = \phi_j / (1-\Phi_j), \qquad \Delta\Phi \equiv \Phi_U - \Phi_L, \qquad \Delta\phi \equiv \phi_U - \phi_L, \tag{24}$$

where we have used $z_j \equiv z(a_j)$ *for* $z = \{\phi, \Phi, V, r\}$ *and:*

$$a_L = \frac{b_L - m_x - s_x^2}{s_x}, \qquad a_U = \frac{b_U - m_x - s_x^2}{s_x}, \qquad b_L = \ln 1 = 0, \qquad b_U = \ln \hat{Y}. \tag{25}$$

Proof of Lemma 2. Suppose that the endowment follows the diffusion process $dY_t = \mu Y_t dt + \sigma Y_t dW_t$, where μ and σ are the instantaneous trend and standard deviation (in fractional terms) of the process. In log form, letting $X_t \equiv \ln Y_t$, the process reads $dX_t = \left(\mu - \frac{1}{2}\sigma^2\right) dt + \sigma dW_t$. The solution between $t = 1 - x$ and $t = 1$ is $X_1 | X_{1-x} \sim \mathcal{N}\left(X_{1-x} + \left(\mu - \frac{1}{2}\sigma^2\right) x, \sigma^2 x\right)$, hence $Y_1 | Y_{1-x} \sim LogNormal\left(\ln Y_{1-x} + \left(\mu - \frac{1}{2}\sigma^2\right) x, \sigma^2 x\right)$ or, equivalently, $Y_1 = Y_{1-x} e^{\left(\mu - \frac{1}{2}\sigma^2\right)x + \sigma\sqrt{x}Z}$, with $Z \sim \mathcal{N}(0,1)$.

Henceforth, set $\mu = 0$ to have a driftless geometric Brownian motion, let $m_x \equiv \ln Y_{1-x} - \frac{1}{2}s_x^2$, and $s_x^2 \equiv \sigma^2 x$ and, abstracting from the time subscript when $t = 1$, work with $y \equiv Y | Y_{1-x} \sim LogNormal(m_x, s_x^2)$. Assume that $Y_t > 1$ for all $t \in [0,1]$. Then, $\ln Y_t > 0$. Therefore, $E(\ln y) = m_x > 0$ must hold.

For a truncation at the lower bound 1 and upper bound $\hat{Y} > 1$, use (25) and denote Φ and ϕ the standard normal CDF and PDF, respectively. The standard truncated lognormal identities give

$$E(y; 1 < Y) = e^{m_x + \frac{1}{2}s_x^2} (1 - \Phi(a_L)),$$

$$E(y \ln y; 1 < Y) = e^{m_x + \frac{1}{2}s_x^2} \left[\left(m_x + s_x^2\right)(1 - \Phi(a_L)) + s_x \phi(a_L)\right],$$

$$E\left(y; 1 < Y \leq \hat{Y}\right) = e^{m_x + \frac{1}{2}s_x^2} \Delta\Phi,$$

$$E\left(y \ln y; 1 < Y \leq \hat{Y}\right) = e^{m_x + \frac{1}{2}s_x^2} \left[\left(m_x + s_x^2\right)\Delta\Phi - s_x \Delta\phi\right],$$

with $\Delta\Phi$ and $\Delta\phi$ as in (24).

Denote S the the left-hand side of (12), which can be rewritten as

$$S \equiv \frac{\bar{\alpha} E(y; 1 < Y) + \lambda E(y \ln y; 1 < Y)}{\bar{\alpha} E\left(y; 1 < Y < \hat{Y}\right) + \lambda E\left(y \ln y; 1 < Y < \hat{Y}\right)},$$

where $\hat{Y} = e^{m_x + s_x \Phi^{-1}(1-c/\bar{c})}$.[19] Applying the truncated-lognormal moment identities, it follows that

$$S = \frac{\left[\bar{\alpha} + \lambda\left(m_x + s_x^2\right)\right]\left(1 - \Phi\left(a_L\right)\right) + \lambda s_x \phi\left(a_L\right)}{\left[\bar{\alpha} + \lambda\left(m_x + s_x^2\right)\right]\Delta\Phi - \lambda s_x \Delta\phi},$$

from which, using the first definition in (23), (19) immediately follows.

Using the first definition in (22) for $j \in \{L, U\}$, (19) becomes

$$S = \frac{V_L}{V_L - V_U}. \tag{26}$$

Differentiate (26), and use (21) and (22) to get (20).

Note that, from (25) and the definition of $\hat{Y}$, it follows that $a_U = \Phi^{-1}\left(1 - c/\bar{c}\right) - s_x$. Furthermore, exploiting the first definition in (23) and the first and third definition in (24), we have $A_L = \bar{\alpha} + \lambda\left(m_x + s_x^2\right) + \lambda\left(-m_x - s_x^2\right) = \bar{\alpha}$. Finally, r_j represents the inverse Mills ratio for the upper tail of the standard normal distribution; hence, $r_U > r_L > 0$.

We may rewrite condition (12) as $S \geq \Psi$. Note that S is strictly increasing in λ, since operationalizing (21) for λ leads to $\mathcal{C}\left(\lambda\right) = \bar{\alpha} s_x \left(r_U - r_L\right) > 0$, and (20) implies $S'\left(\lambda\right) \propto \mathcal{C}\left(\lambda\right)$. The boundary values of S are then identified by $S^{\min} = S\left(\lambda = 0\right) = \left(1 - \Phi_L\right)/\Delta\Phi > 1$ and $S^{\max} = S\left(\lambda = \bar{\lambda}\right) > S^{\min}$. Therefore, if $\Psi \geq S^{\max}$, then the condition is *never* satisfied and maritime transportation is the only economically viable option for the producer. Conversely, $\Psi < S^{\min}$ implies that the condition is *always* satisfied and maritime transportation is not a desirable option. In the intermediate cases, when $S^{\min} \leq K < S^{\max}$, we can identify a closed-form expression for threshold $\tilde{\lambda}_x$ that solves $S\left(\tilde{\lambda}_x\right) = K$, such that, for any $\lambda < \tilde{\lambda}_x$, the producer will opt for maritime transportation and, for any $\lambda \geq \tilde{\lambda}_x$, for air transportation, as stated in Proposition 1. ■

Proof of Corollary 1. Recall Lemma 2, and note that, since the "sea choice" only makes sense when the denominator of S in (19) is positive, we must have[20]

$$B\Delta\Phi - \lambda s_x \Delta\phi > 0. \tag{27}$$

We subdivide the proof in sections to ease readability.

Effect of a variation in volatility (s_x) on S. Firstly, since $a_U = \Phi^{-1}\left(1 - c/\bar{c}\right) - s_x$, then $a_U' = -1$. Furthermore, from the first definition in (23), we have that $B'\left(s_x\right) = \lambda s_x$. Using

[19]The cutoff value $\hat{Y}$ is obtained from the equilibrium condition (10), which under lognormality reads $\Phi\left(\left(\ln\hat{Y} - m_x\right)/s_x\right) = 1 - c/\bar{c}$, leading to $\ln\hat{Y} = m_x + s_x\Phi^{-1}\left(1 - c/\bar{c}\right)$.

[20]Note that condition (27) must hold for every $\lambda \in \left[0, \bar{\lambda}\right]$.

(22) at a_L and a_U into (21), and noting that $(\lambda s_x)' = \lambda$, we can then write

$$\begin{aligned}\mathcal{C}(s_x) &= B(K_U r_U - K_L r_L) - \lambda^2 s_x^2 (r_U - r_L) + \lambda s_x r_U r_L (K_U - K_L) \\ &= \left(\lambda B - \lambda^2 s_x^2 + A_U B\right)(r_U - r_L) + \lambda s_x r_U r_L (K_U - K_L) + B(A_U + a_L' \bar{\alpha}) r_L,\end{aligned} \tag{28}$$

where $r_j \equiv r(a_j)$, $K_j \equiv K(a_j)$, $\Phi_j \equiv \Phi(a_j)$ for $j \in \{L, U\}$, and therefore, from (20),

$$S'(s_x) = \frac{(1-\Phi_L)(1-\Phi_U)}{(V_L - V_U)^2} \mathcal{C}(s_x). \tag{29}$$

Recall $r(a)$ is strictly increasing in a (inverse Mills ratio for the upper tail), hence (**i**) $r_U - r_L > 0$. Moreover, $m_x > 0$ implies $B > \lambda s_x^2$, so (**ii**) $\lambda B - \lambda^2 s_x^2 > 0$. From (25), we can also write $a_U = a_L + b_U/s_x > a_L$ because $b_U > b_L$, hence $A_U > A_L$. Together with the feasibility condition (27), this also implies the useful bound

$$\frac{B}{\lambda s_x} > \frac{\Delta\phi}{\Delta\Phi} = -E(Z | a_L < Z < a_U) > -a_U,$$

and consequently, (**iii**) $A_U = B + \lambda s_x a_U > 0$. Furthermore, $a_U' = a_L' - 1 - m_x/s_x^2 < a_L'$. Therefore, (**iv**) $K_U > K_L$ since $-a_U' > -a_L'$ and $A_U > A_L$. Finally,

$$\begin{aligned} A_U = -a_U' A_U &= -\left(a_L' - 1 - m_x/s_x^2\right) A_L - a_U' \lambda s_x (a_U - a_L) \\ &= -a_L'\bar{\alpha} + \left(1 + m_x/s_x^2\right) A_L + \lambda s_x (a_U - a_L) > -a_L' \bar{\alpha}, \end{aligned}$$

from which it follows that (**v**) $A_U + a_L'\bar{\alpha} > 0$.
As a result, the first term of $\mathcal{C}(s_x)$ in (29) is strictly positive because (**i**) $r_U - r_L > 0$, (**ii**) $B > \lambda s_x^2$, and (**iii**) $A_U > 0$; the second because (**iv**) $K_U > K_L$; the third because (**v**) $A_U + a_L'\bar{\alpha} > 0$. Since $\mathcal{C}(s_x) > 0$, $(1-\Phi_U)(1-\Phi_L) > 0$, and $V_L > V_U$ by construction, $S'(s_x) > 0$ must also hold.

Effect of variations in marginal costs $(c, \bar{c})$ on S. The marginal costs impact S solely via $\hat{Y}$, which in turn affects a_U. Since $a_U = \Phi^{-1}(1 - c/\bar{c}) - s_x$, we have $a_U'(c) = -1/(\bar{c}\phi(a_U + s_x)) < 0$ and $a_U'(\bar{c}) = c/(\bar{c}^2 \phi(a_U + s_x)) > 0$. Using (22) at a_L and a_U into (21), and noting that $(\lambda s_x)' = 0$, we can then write

$$\mathcal{C} = -a_U' r_U A_U (B + \lambda s_x r_L).$$

Since $r_L, r_U, A_U > 0$, the sign of $-a_U'$ determines the sign of $\mathcal{C}$. Therefore, $S'(c) > 0$ and $S'(\bar{c}) < 0$.

Impact of variations in the fundamentals on $\tilde{\lambda}_x$. Differentiating $S\left(\tilde{\lambda}_x\right) = \Psi$ with respect to a generic parameter or exogenous variable z yields $S'\left(\tilde{\lambda}_x\right)\tilde{\lambda}'_x(z) + S'(z) = \Psi'(z)$, hence

$$\tilde{\lambda}'_x(z) = -\frac{S'(z) - \Psi'(z)}{S'\left(\tilde{\lambda}_x\right)}. \tag{30}$$

With $z = s_x$, $\Psi'(z) = 0$. Since $S'\left(\tilde{\lambda}_x\right) > 0$, it follows that $\tilde{\lambda}'_x(z) \propto -S'(z)$. Hence, $\tilde{\lambda}'_x(s_x) < 0$ because $S'(s_x) > 0$.

With $z = c$, $\Psi'(c) = (\tau_x - 1)\bar{c}/(\bar{c} - \tau_x c)^2 > 0$. Thus, $\tilde{\lambda}'_x(c) < 0$, because $S'(c), \Psi'(c) > 0$.

With $z = \bar{c}$, $\Psi'(\bar{c}) = -(\tau_x - 1)c/(\bar{c} - \tau_x c)^2 < 0$. As a result, $\tilde{\lambda}'_x(c) > 0$, because $S'(\bar{c}), \Psi'(\bar{c}) < 0$.

Finally, With $z = \tau_x$, $\Psi'(\tau_x) = c/(\bar{c} - \tau_x c)^2 > 0$. Therefore, $\tilde{\lambda}'_x(c) > 0$, because $S'(\tau_x) = 0$ and $\Psi'(\tau_x) > 0$. ■

Derivation of $\partial\tilde{\lambda}_x/\partial Y_{1-x}$. The endowment level Y_{1-x} impacts S via m_x, which in turn affects B and a_L. From the definition of m_x, (23), and (25), it follows that $m'_x = 1/Y_{1-x} > 0$, $a'_L = -m'_x/s_x < 0$ and $B' = \lambda m'_x \geq 0$. Using (22) at a_L and a_U into (21), and noting that $(\lambda s_x)' = 0$, we can then write

$$\begin{aligned}\mathcal{C}(m'_x) &= \lambda m'_x(B + \lambda s_x r_L) - (\lambda + A_L r_L/s_x)\, m'_x(B + \lambda s_x r_U)\\ &= -m'_x\left[\lambda^2 s_x(r_U - r_L) + A_L r_L(B + \lambda s_x r_U)/s_x\right] < 0,\end{aligned}$$

where the inequality follows from $m'_x, A_L > 0$ and $r_U > r_L$. As a result, $S'(Y_{1-x}) < 0$. Using (30) and noting that $\Psi'(Y_{1-x}) = 0$ and $S'\left(\tilde{\lambda}_x\right) > 0$, it follows that $\tilde{\lambda}'_x(Y_{1-x}) \propto -S'(Y_{1-x})$. Hence, $\tilde{\lambda}'_x(Y_{1-x}) > 0$. ■

Proof of Proposition 3.

Part (i). Suppose $\mu^A_x \leq \mu^S_x$. Then $\bar{\lambda} + \tilde{\lambda}_x \leq \tilde{\lambda}_x + \hat{\lambda}_x$, which in turn requires $\bar{\lambda}_x \leq \hat{\lambda}_x$, leading to a contradiction.

Part (ii). Using the definition of $\Delta\mu_x$, μ^A_x, and μ^S_x, we can write:

$$\frac{d\Delta\mu_x}{dx} = \frac{d\mu^A_x}{dx} - \frac{d\mu^S_x}{dx} = \frac{1}{2}\left(\frac{d\bar{\lambda}}{dx} - \frac{d\hat{\lambda}_x}{dx}\right).$$

Since $\bar{\lambda}$ is exogenous, $d\bar{\lambda}/dx = 0$. Moreover, $\hat{\lambda}_x$ is pinned down by the entry condition (14), so by the implicit function theorem, $d\hat{\lambda}_x/dx = -\left(\partial E_0\left(\pi_{\hat{\lambda}_x,x}\right)/\partial x\right)/\left(\partial E_0\left(\pi_{\hat{\lambda}_x,x}\right)/\partial\hat{\lambda}_x\right)$. The denominator is positive due to $\partial E_{1-x}\left(\pi^{\mathcal{T}}_{\lambda,x}\right)/\partial\lambda > 0$ for $\mathcal{T} = \{A, S\}$, and thus $\partial E_0\left(\pi_{\lambda,x}\right)/\partial\lambda >$

0. The numerator is negative because $E_0\left(\pi_{\lambda,x}\right)$ is the ex-ante value of optimally choosing air versus sea for the producer of quality λ using the signal Y_{1-x}: as x increases, the mode decision is taken earlier (at $1-x$), with a coarser information set, which lowers the value of the optimal choice. Hence $\partial E_0\left(\pi_{\lambda,x}\right)/\partial x \leq 0$. Therefore, $d\hat{\lambda}_x/dx > 0$, and it follows that $d\Delta\mu_x/dx = -\left(d\hat{\lambda}_x/dx\right)/2 \leq 0$, which proves the claim. ■