Time as a Trade Barrier

January 2012

David L. Hummels, Purdue University and NBER

Georg Schaur, The University of Tennessee

Abstract: A large and growing share of international trade is carried on airplanes. Air cargo is many times more expensive than maritime transport but arrives in destination markets much faster. We model firms' choice between exporting goods using fast but expensive air cargo and slow but cheap ocean cargo. This choice depends on the price elasticity of demand and the value that consumers attach to fast delivery and is revealed in the relative market shares of firms who air and ocean ship. We use US imports data that provide rich variation in the premium paid for air shipping and in time lags for ocean transit to identify these parameters and extract consumer's valuation of time. By exploiting variation across US entry coasts we are able to control for selection and for unobserved shocks to product quality and variety that affect market shares. We estimate that each day in transit is equivalent to an ad-valorem tariff of 0.6 to 2.3 percent and that the most time-sensitive trade flows are those involving parts and components trade. These results suggest a link between sharp declines in the price of air shipping and rapid growth in trade as well as growth in world-wide fragmentation of production. Our estimates are also useful for assessing the economic impact of policies that raise or lower time to trade such as security screening of cargo, port infrastructure investment, or streamlined customs procedures.

1 Introduction

Moving traded goods over long distances takes time. Ocean-borne cargo leaving European ports takes an average of 20 days to reach US ports and 30 days to reach Japan. Air borne cargo requires only a day or less to most destinations, but it is also much more expensive. In 2005, goods imported into the US faced per kilogram charges for air freight that were, on average, 6.5 times higher than ocean freight charges.

Despite the expense, a large fraction of world trade travels by air. Excluding Canada and Mexico, 36 percent of US imports by value and 58 percent of US exports by value were airborne in 2000. Numbers are similar for many other countries. In 2004, air cargo as a share of export value was 29 percent for the UK, 42 percent for Ireland, and 51 percent for Singapore; 22 percent of Argentine and 32 percent of Brazilian imports were airborne.[1] Further, the reliance on airplanes is rising steadily over time. From 1965-2004, worldwide use of air cargo grew 2.6 times faster than use of ocean cargo.[2]

In sum, airplanes are fast, expensive, and increasingly important to trade. In this paper we examine two hypotheses suggested by these facts: lengthy shipping times impose costs that impede trade and firms engaged in trade exhibit significant willingness-to-pay to avoid these costs.

What are these time costs? Lengthy shipping times impose inventory-holding and depreciation costs on shippers. Inventory-holding costs include both the capital cost of the goods while in transit, as well as the need to hold larger buffer-stock inventories at the final destination to accommodate variation in arrival time. Depreciation captures any reason that a newly produced good might be preferable to an older good. This could include literal spoilage (fresh produce or cut flowers), or rapid technological obsolescence for goods such as consumer electronics.

Evans and Harrigan (2005), Aizenman (2004), and Hummels and Schaur (2010) emphasize the interaction between timeliness and demand uncertainty.[3] Long lags between ordering and delivery require firms to commit to quantities supplied well before uncertain demand is resolved. Forecast errors then result in lost profitability as firms over- or under-supply the

---

[1] Cristea et al (2011) provide systematic data on trade by transport mode for many countries in 2004.

[2] Hummels (2007) calculates that worldwide use of airborne cargo (measured in kg-km) grew 11.7 percent per year from 1965-2004 compared to 4.4 percent per year for ocean cargo.

[3] Evans and Harrigan (2005) provide a model of and empirical evidence on lean retailing in the apparel industry and emphasize location choice – the higher the restocking rate for particular clothing items, the greater is the reliance on proximate import sources. Hummels and Schaur (2010) provide evidence on uncertainty revealed by past price volatility and emphasize modal choice, showing that goods with high volatility are more likely to be air shipped.

market. Similarly, firms that adapt their products to reflect changing consumer tastes will find it difficult to respond to late arriving signals about those tastes when resupply is a slow boat ride away. This results in lost sales as firms face a mismatch between the product characteristics on offer and those desired by consumers.

These costs will be magnified in the presence of product fragmentation. When countries specialize in stages of production and trade intermediate inputs the inventory-holding and depreciation costs for early-stage value-added accrue throughout the duration of the production chain. Similarly, demand uncertainty for the final product can ripple throughout upstream stages. Perhaps most importantly, the absence of key components due to late arrival or quality defects can idle an entire assembly plant, making the ability to ship rapidly worth potentially many times the value of the components being transported.[4]

In the same vein one can construct many examples, both in an international trade context or more broadly, in which timeliness is important for consumers and producers. The challenge is finding data that would allow the econometrician to go beyond anecdote and toward careful measurement. To identify the magnitude and the source of time costs it is necessary to examine some tradeoff in the data, some choice in which firms exhibit their willingness to pay to avoid time delays. In this paper we examine the modal choice decisions of firms engaged in trade and the trade-off between fast and expensive air transport versus slow and inexpensive ocean shipping.

In the model consumers have preferences over goods that are differentiated along both horizontal and quality dimensions, and slow delivery reduces consumers' perception of product quality. This creates an incentive for producers to improve perceived quality of the delivered product by shipping it via airplane. Freight carriers charge shipping costs that are proportional to quantity, not value, shipped, and the unit cost of air shipping exceeds that of ocean shipping. Unit shipping costs imply that the air freight premium, measured in ad-valorem terms, is decreasing in product prices. That is, high price firms incur a smaller increase in delivered prices when they upgrade quality using airplanes, and are more likely to air ship goods, while low price firms are more likely to employ ocean shipping. This is consistent with patterns in the data: for 75 percent of trade by value and 90 percent of trade in manufacturing categories we see similar goods (same HS-6 group) coming from the same exporter shipped using both air and ocean modes.

---

[4] Harrigan and Venables (2006) argue that this is an important force driving economic agglomeration, but firms need not cluster geographically if long distances can be rapidly bridged with airplanes.

We show that consumers' valuation of time is then revealed in the relative revenues of the two types of firms.   Since air shipping is expensive, consumers shift purchases away from the firm that air ships in proportion to the price elasticity of demand.   Conditional on prices, consumers shift purchases toward the firm that air ships in proportion to their valuation of time. This revenue shifting will be strongest when demand is more price elastic and when the time delays are greatest.   A consumer buying goods from a nearby exporter may be unwilling to pay the air premium to save a few days in transit, but that same consumer will pay the air premium if the exporter is many weeks of ocean travel away.   By combining our estimates of these two effects we can extract the price-equivalent of the consumers' valuation of each day of delay.

The model also allows us to examine an alternative explanation for the use of air shipment and its rapid rise over time.   Hummels (2007) shows that the cost of air shipping a kilogram of cargo dropped an order of magnitude between 1955 and 2005, and that the value of trade grew faster than the weight of trade as bulk commodities represent a falling share of traded goods. Because the premium paid for air shipping is lower for high value/weight products, a fall in the weight of trade pushes goods toward air shipping. Harrigan (2010) formalizes this insight in a model of comparative advantage to predict that distant countries will specialize in lightweight goods that are air shipped. He shows that that the longer the distance to the destination market, the higher the unit values of the products delivered and that Canada and Mexico have a relatively low US market share for products that are air shipped buy the rest of the world.

While Harrigan (2010) assumes there exists some preference for timely delivery the question is how large this preference must be to generate observed patterns of trade.   For sufficiently high value products, the ad-valorem air freight premium becomes vanishingly small, and even a very small value of time will shift goods onto planes.   In short, the use of airplanes is not by itself evidence that consumers place a large value on timeliness.   It is necessary to demonstrate that consumers are willing to pay a premium in ad-valorem terms, and conditional on that premium, to shift purchases more rapidly toward airplanes the longer is the ocean voyage.

To estimate this model we use data on US imports 1991-2005 that allow us to construct, for air and ocean modes, measures of revenues, prices, shipping costs, and numbers of shipments that are specific to each exporter x HS 6 digit product x US entry coast x year.   We combine this with a detailed ocean shipping schedule for all ocean vessels worldwide that provides us with shipping times for each exporter x US entry coast.   We then relate relative (air/ocean) revenues to relative

prices, relative shipping costs and time delays.   We exploit variation in the price/speed trade off across exporters, products, entry points and time in order to identify consumers' willingness to pay for time savings.

Importantly, the rich structure of the data allows us to address a host of econometric concerns that have been raised in the trade literature. Helpman, Melitz and Rubinstein (2008) and Chaney (2008) point out that trade costs induce selection out of markets and that regressions of export sales on trade costs incorporate both this selection effect and an extensive margin (number of firms trading) response to the costs.   We control for selection with a two-step estimator that uses the exporter's sales to the world (less the US) for each product x year to predict the probability that it will sell that product to the US.   While this does not affect our estimates in the second stage, the success of the first stage suggests that this strategy could be useful for future studies that need selection variables operating at highly disaggregated levels.

We face an additional and somewhat different extensive margin problem than the literature, which contemplates firms selecting out of markets as trade costs rise.   We have a modal extensive margin in which firms stay in the market but potentially switch between modes in response to changes in relative trade costs.   To control for this modal extensive margin we normalize revenues by the number of shipments made.   The normalized dependent variables are akin to average sales per firm, and the transformed data exhibits the muted trade cost response predicted by Helpman, Melitz, Rubinstein (2008) and Chaney (2008) theories.   This suggests that our strategy may also be useful for future studies that wish to control for firm entry and exit occurring within disaggregated product categories.

A recent literature emphasizes the importance of quality differentiation in trade. Empirically, quality is measured either as price variation[5] or as a residual of quantity demanded controlling for prices.[6]   Unlike this literature we have an explicit measure of one aspect of quality, timely delivery, for which we directly estimate consumers' valuation.   In addition, we employ various fixed effect estimators to provide strong controls for unobserved quality and variety that affect relative revenues.

In the most robust treatment we hold fixed unobservables that are specific to an exporter x product x time and exploit variation across US entry points in order to identify the effects.   For

---

[5] Examples include Schott (2004), Hallak (2006), Choi et al (2009), and Baldwin and Harrigan (2011).
[6] Examples include Hummels-Klenow (2005), Hallak-Schott(2011), Khandelwal (2010).

example, we can examine air relative to ocean revenues for German firms selling machine tools to customers on the US East and West coasts. The need to traverse the Panama Canal means that German ocean cargo arrives on the East coast 12 days sooner than the West coast. We can then hold fixed the relative quality of air and ocean shipped German machine tools in a given year and exploit this quirk of geography to generate variation across US coasts in the relative share of air shipping as a function of relative time delays, and relative freight prices.

The ability to exploit variation in modal revenues across exporter-product-coast-time observations allows us to control for unobserved quality variation in a manner that is considerably more general than what is found in the literature on estimating import demand elasticities or in the literature on quality and trade. This technique also allows us to control for a final reason that air cargo is employed in trade. For some countries, air shipping may be especially effective as a way to bridge firms and customers located in continental interiors. Air cargo avoids overland shipping and the cost, congestion, and occasionally corruption found in ocean ports. By exploiting variation across US coasts for a given exporter-product-time, we hold fixed the characteristics of exporters – their geography, income, infrastructure – that may affect usage of air shipments.

We find that air revenues are high relative to ocean revenues when the air freight premium is low, and when shipment lags are long. In the pooled specifications we estimate that each day in transit is worth from 0.6 to 2.1 percent of the value of the good. We also estimate the model separately for each End-Use category and find considerable heterogeneity across products in time sensitivity. The most striking result from the disaggregated product regressions is that parts and components have a time sensitivity that is 60 percent higher than other goods.

The econometric technique employed here directly identifies the value of time saving from transport modal choice, but the estimates are informative about many policies and sources of technological change that speed goods to market. For example, imposing strict port security procedures such as those being currently contemplated in the US could significantly slow the flow of goods into the domestic market. Streamlining elaborate customs procedures or investing in more efficient port infrastructure may allow goods to reach their destinations more quickly and boost trade. Djankov, Freund and Pham (2010) investigate this possibility using product-specific estimates of per day time costs taken from an earlier draft of this paper. They find that countries with long customs delays see reduced trade volumes, and the largest reductions in trade occur in the most time sensitive products. In a related paper, Hummels (2007b) directly calculates the

5

tariff equivalent of customs and port delays taken from World Bank "Doing Business" data. He combines estimates of per day time costs taken from an earlier draft of this paper with data on days lost to customs delays and port clearance and finds that customs delays prove a far larger barrier to trade than do applied tariffs.[7]

Our estimates also have implications for changing patterns of trade and the international organization of production. In the post-war era, world trade has grown much faster than output with typical explanations attributing this growth to declining tariffs and improved technology (information and transportation).[8] To the extent that time is a barrier to trade, declines in air shipping prices may help explain both aggregate trade growth and a shift toward trade in especially time sensitive goods or forms of production organization.

As an example, an important recent feature of trade is especially rapid growth in the fragmentation of production. Hummels, Ishii and Yi (2000) document that vertical specialization (the use of imported inputs in exported goods) has increased 30%, and been responsible for roughly half of overall trade growth from 1970-1990.[9] The explosive growth in China's exports – a tenfold increase from 1992-2006 – is also directly linked to fragmentation as roughly half of these exports are categorized as processing trade.[10] Our estimates show that parts and components are among the most time sensitive products. This suggests that the rapid declines in air transport costs, and the corresponding reduction in the cost of time-saving, may be responsible for the growth of time and coordination-intensive forms of integration.

The paper proceeds as follows. Section 2 models the firm's choice of shipping mode and generates predictions for relative export revenues. Section 3 describes the data and specification issues in estimation. Section 4 provides results. Section 5 concludes.

2 Theory

In our data we see exporter-by-product trade flows into the U.S. disaggregated by transportation mode (air and ocean vessel). In many instances, data for a single trade flow indicates that both air and ocean modes were used in the same time period. In other instances, an

---

[7] For example, Latin American exporters face applied average tariffs of 4 percent, and time costs of customs delays of 8 percent ad-valorem, while Sub-Saharan exporters face applied tariffs of 4 percent and time costs of customs delays equal to 16 percent.
[8] Two important papers are Yi (2003) and Baier and Bergstrand (2001).
[9] See Johnson and Noguera (2011) for related but more detailed calculations for a recent cross-section.
[10] Feenstra and Wei (2009)

(exporter-by- product) trade flow enters the US either by ocean or by air in a single time period, but the mode chosen may vary across periods, across exporters, or across varieties within a narrowly defined product classification. Our goal in this section is to provide a simple theoretical structure that yields these outcomes in order to organize our analysis of modal use and its implications for the value of time savings.

We focus on US import demands within a narrow product category k (in the data, an HS6 or HS10 digit good). Since all variables below are product specific, and all destinations are in the US, we suppress production and destination superscript for notational ease, reintroducing it where appropriate in the empirical section.

Import demand is CES across varieties, summed across export locations j and across firms z within each location j.

$$U = \left(\sum_j \sum_i \lambda_j^z \left(q_j^z \right)^\theta \right)^{1/\theta} \qquad\qquad \theta = (\sigma - 1)/\sigma$$

where σ is the elasticity of substitution between goods and $\lambda_j^z = v_j^z \exp(-\tau \cdot days_j^z)$ is a price-equivalent demand shifter that depends on a firm z, location j-specific quality, v_j^z, and a term $\exp(-\tau \cdot days_j^z)$ that captures the consumer disutility of slow delivery.

This formulation of demand is similar to the literature on quality in trade, including Hallak (2006), Hummels-Klenow (2005), and Hallak-Schott (2011), with the exception that these papers treat all elements of quality as unobservable. In contrast, we measure timeliness as an important measurable component of quality. Time in transit, $days(z)_j$, depends on exporter location because of differences in distance to the import market and infrastructure quality, but also depends on the endogenous choice of firm z to pay a premium for timely delivery.

With real expenditures on product k given by E, demands for firm z from exporter j, selling at a delivered price $p_j^z *$ are

$$q_j^z = E \left(\frac{p_j^z *}{v_j^z \exp(-\tau \cdot days_j^z)} \right)^{-\sigma} . \qquad\qquad (1)$$

Other things equal, a consumer gets more utility from a good that arrives sooner rather than later, which is expressed by increasing demand for that good. A 1% price reduction raises demand by σ %, and a 1 day reduction in delivery times raises demand by $\sigma\tau$. That is, the time valuation

parameter τ translates days of delay into a price (or tariff) equivalent form, and the elasticity of substitution σ translates this into the quantity of lost sales.

Turning to the production side of the model, the firm z marginal cost of delivering a product from export location j to the market via mode m=air,ocean is $z + g_j^m$, where z is the marginal cost of production (potentially correlated with unobserved quality v_j^z) and g_j^m is a location-j, mode-m specific per unit shipping charge. That is to say, g_j^m is proportional to the quantity, not the value shipped as with iceberg costs (see Hummels-Skiba 2004 for evidence on this point). We assume (and our data bear out) that air shipping is more expensive than ocean shipping, $g_j^A > g_j^O$.

The firm pays fixed costs FC at the beginning of the period and commits to a mode of transportation. The firm charges prices that are a markup over marginal costs, $p_j^z* = (z + g_j^m)/\theta$. Multiplying by the quantity demanded from (1) and subtracting fixed and variable costs yields

$$\pi(z)_j^m = \frac{(z + g_j^m)}{\sigma - 1} E \left(\frac{(z + g_j^m)/\theta}{v_j^z \exp(-\tau \cdot days_j^m)} \right)^{-\sigma} - FC \qquad (2)$$

To determine the optimal transport mode the firm compares the profitability of air versus ocean shipping. The firm chooses air if $\pi(z)_j^a > \pi(z)_j^o$. Taking logs of (2), assuming that airborne cargoes can reach their destination in one day, and simplifying implies

$$(1-\sigma)\left[\ln(z + g_j^a) - \ln(z + g_j^o) \right] + \sigma\tau\left[days_j^o - 1 \right] > 0 \qquad (3)$$

Equation (3) shows that a firm trades the greater expense of air shipping against the improved "quality" of a product that arrives $days_j^o - 1$ days earlier. Long ocean shipping times are more likely to induce a switch to air shipping when consumers attach greater value to timeliness, and when goods are closer substitutes. The latter effect operates because we have defined τ in price equivalent terms in order to measure the effect of timeliness on quantities shipped. Higher elasticities of substitution translate into larger quantity effects.

The additive form of shipping costs also implies that modal choice depends on marginal costs of production. Since $g_j^a > g_j^o$, using air shipping always results in a higher delivered price,

$\ln(z + g_j^a) - \ln(z + g_j^o) > 0$. But the magnitude of this difference -- the impact that the air shipping premium has on delivered prices -- is decreasing in marginal costs of production,

$$\partial \ln \left[\frac{z + g_j^a}{z + g_j^o} \right] / \partial z = \frac{g_j^o - g_j^a}{\left(z + g_j^a \right)\left(z + g_j^o \right)} < 0.$$

To see the intuition, suppose a pair of shoes can be shipped by air for $11 or by ocean for $1. For $9 shoes, the air freight premium is equivalent to a 100 percent tariff. For $99 shoes, the air freight premium is equivalent to only a 10 percent tariff.

To complete the firm's problem, it must be the case that for the optimal mode, profits from exporting exceed fixed costs, or

$$(1 - \sigma) \ln(z + g_j^m) + \sigma \ln v_j - \sigma \tau \cdot days_j^m + \kappa(\theta) > \ln FC \tag{4}$$

This defines a selection equation indicating whether or not a particular location successfully exports a product to the importer.

This relationship is shown in Figure 1, which graphs the log profitability (before fixed costs) of each mode against the marginal cost of production. At the crossing point, firms are indifferent between modes. For costs below $\tilde{z}$ only ocean shipping is chosen, above $\tilde{z}$ only air shipping is chosen. An increase in τ, or a decrease in g_j^a / g_j^o shifts the π^a (air profits) curve up and shifts $\tilde{z}$ to the left. An increase in FC reduces the range of marginal costs at which a firm can successfully export.

From this, we can derive two cases that correspond to modal-use patterns in the data. For a single firm, shipping a product to a particular destination at a point in time, it will generally be optimal to choose either air or ocean shipping. As we show in the theory appendix, it is straightforward to derive a probit model from equation (3) relating the probability of air shipment to relative shipping prices and days in transit. If all variables are observed we can extract consumers' valuation of time saving from that model. However, this case poses two significant challenges for estimation: we do not observe shipping costs for the transport mode not chosen and we cannot control for unmeasured product quality variation.

Consider a second case. When we look at national trade data we are aggregating over multiple firms. Suppose we have two firms from exporter j with different marginal costs. If

$z_j^1 < \tilde{z}_j < z_j^2$, the first firm ocean ships and the second firm air ships.   The firm that ocean ships

(now denoted with an "o") generates export revenues inclusive of shipping charges

$$r_j^o* = E\left(\frac{1}{v_j^o \exp(-\tau \cdot days_j^o)}\right)^{-\sigma} \left(\frac{\sigma}{\sigma-1}(z^o + g_j^o)\right)^{1-\sigma} \tag{5}$$

and similarly for the firm that air ships.   Writing revenue in relative terms and taking logs

$$\ln\frac{r(z_j^a)*}{r(z_j^o)*} = \sigma\tau(days_j^o - 1) + (1-\sigma)\left[\ln\left(z_j^a + g_j^a\right) - \ln\left(z_j^o + g_j^o\right)\right] + \sigma\ln\left(\frac{v_j^a}{v_j^o}\right) \tag{6}$$

To take the theory to the data we need to transform equation (6) so that all variables are

observable.   The first step is to rewrite the difference in marginal costs (unobservable) as a

function of (observable) export prices and shipping charges.   Delivered prices inclusive of

shipping charges (denoted *) are $p_j^m* = (z_j^m + g_j^m)/\theta$.   We can then take the difference in marginal

costs, from equation (6) and transform it into the difference in delivered prices so long as the

markup on the origin price is independent of delivery mode.[11]

$$ln\left(z_j^a + g_j^a\right) - ln\left(z_j^o + g_j^o\right) = \ln\frac{1}{\theta} + \ln(z_j^a + g_j^a) - \ln\frac{1}{\theta} - \ln(z_j^o + g_j^O) = \ln p_j^a* - \ln p_j^o* \tag{7}$$

While shipping costs are imposed on a per unit basis they can be rewritten in ad-valorem

terms by dividing through by origin prices, p_j^m

$$\ln p_j^m* = \ln\left(p_j^m\left(1 + \frac{g_j^m}{p_j^m}\right)\right) = \ln p_j^m + \ln f_j^m \quad , \qquad f_j^m = \left(1 + \frac{g_j^m}{p_j^m}\right) > 1 \tag{8}$$

where f_j^m is the ad-valorem equivalent of unit shipping charges facing the firm. Of course, f_j^m is

not an exogenous technological parameter, but instead depends on the origin price that the firm

---

[11] Origin prices are independent of delivery mode if the CES markup on the marginal cost inclusive of delivery is the same for both modes.   Amazon.com for example does not alter the list price of a book depending on whether a buyer chooses next day or delayed shipping; it simply adds the freight premium on to the published price.   Theoretically this assumption is justified if the supply of goods into the shipping channel is perfectly competitive so that factory gate prices equal marginal costs.   If we are instead in a monopolistic competition setting there is a nonlinear interaction between markups and entry mode.   This nonlinear interaction can be approximated by a Taylor expansion that yields the same form used here, which is to say that the feedback effects onto markups are of second order importance.

charges. Using data on ad-valorem charges while omitting data on prices and quality would be problematic, a point we address in detail in Section 3. Using (7) and (8) in (6), we can express the difference in marginal costs as the difference in origin prices plus differences in ad-valorem shipping charges.

$$\ln(z_j^a + g_j^a) - \ln(z_j^o + g_j^o) = \ln p_j^a - \ln p_j^o + \ln f_j^a - \ln f_j^o \tag{9}$$

We do not want to induce bias in our estimates by including shipping charges on both sides of the estimating equation. Noting that $r*/f = r$, we rewrite revenues inclusive of shipping charges in (6) as revenues exclusive of these charges by subtracting $\ln(f_j^a / f_j^o)$ from both sides.

$$\ln \frac{r(z_j^a)}{r(z_j^o)} = \sigma\tau(days_j^o - 1) + (1-\sigma)\ln\left[\frac{p_j^a}{p_j^o}\right] - \sigma\ln\left[\frac{f_j^a}{f_j^o}\right] + \sigma\ln\left(\frac{v_j^a}{v_j^o}\right) \tag{10}$$

Equation (10) captures a trade-off similar to that in equation (3), only expressed in revenue rather than probability terms. Consumers view goods from the two firms as imperfect substitutes, and alter their relative purchases as a function of relative price and relative quality. We identify this in the data as a tradeoff – ocean shipped goods have lower costs but are perceived by consumers to be of lower quality because they arrive days or weeks later than an air shipped good. The cost difference induces larger movements in revenues when σ is large (the goods are close substitutes). Time delays induce larger movemenets in revenues when σ is large and consumers have a higher valuation for timeliness, τ. Combining estimates of σ and $\sigma\tau$ we can extract consumers willingness to pay for timely delivery. Finally, we account for the possibility that consumers may also perceive a quality difference between the two types of firms that is unrelated to timeliness. This appears as the last term in equation (10). We discuss this in depth in Section 3.

Equation (10) generalizes to the case of many firms. Let N_j^m denote the number of firms of type z_j^m, and write aggregate revenues $R(z_j^m)$ as an aggregation over all firms that export using mode m. In relative terms, aggregate revenues are

$$\frac{R(z_j^a)}{R(z_j^o)} = \frac{N_j^a r(z_j^a)}{N_j^o r(z_j^o)} = \sigma\tau(days_j^o - 1) + (1-\sigma)\ln\left[\frac{p_j^a}{p_j^o}\right] - \sigma\left[\frac{f_j^a}{f_j^o}\right] + \sigma\ln\left(\frac{v_j^a}{v_j^o}\right) + \ln\left(\frac{N_j^a}{N_j^o}\right)$$ (11)

The distinction between revenues per firm and revenues aggregated over N_j^m firms displays the *potential* importance of a modal extensive margin – defined not as the number of firms exporting but the number of firms within an industry that export using a given mode of transportation.

How this modal extensive margin adjusts in the data is not immediately clear. For example, if a given exporter has a small number of firms exporting a specific product to the US, and the marginal cost of those firms is not close to $\tilde{z}$, then the modal extensive margin will not adjust to small cost shocks.

However, if an exporter has many firms close to the $\tilde{z}$ cutoff, then an endogenous adjustment of the modal extensive margin must be addressed. We show in appendix I that equation (11) is a second order approximation of a model in which heterogenous firms draw marginal costs z from a distribution as in Melitz (2003). In that case the mass (number) of firms in each mode adjusts continuously in response to changes in time delays and shipping costs, and the included variables for cost and quality are weighted averages over the firms in each mode. We discuss the treatment of the modal extensive margin at length in the Section 3.

3 Data and Specifications

We estimate equations (10) and (11) and exploit variation in the relative revenues for air and ocean shipped goods to identify σ and τ. In this section we describe the data employed and the exact sources of variation we will use, and then discuss how we will address selection, endogeneity, and the possibility of an endogenous extensive margin.

3.1 Data

We employ highly disaggregated data from the U.S. Imports of Merchandise database, which reports US imports at monthly frequencies from 1991-2005. We have quantities (in kg), the total value of the shipments (in US$), shipping charges (US$), and number of distinct shipment records reported separately for each exporter x HS10 product category x US customs district (the point where the imports enter the US) x transportation mode (m=air, ocean) x time period. These data allow us to calculate mode-specific revenues, origin prices, shipping costs in ad-valorem

terms, and number of shipments.[12]

We begin with roughly 45 million trade observations that arrive in the US by air or ocean. We drop inland customs districts[13] along with Puerto Rico, Hawaii, and the Virgin Islands, which accounts for 7 percent of imports by value. Most imports from Canada and Mexico arrive overland and so do not appear in the air or ocean shipments data. We drop the remaining imports from Canada and Mexico that arrive by air or ocean in U.S. coastal districts (4 percent of import value), as we lack reliable transport cost and time data for these shipments and because these shipments have a dominant outside option (overland transport).

When taking equations (10) and (11) to the data, an observation is an HS6 digit good k (roughly 5000 distinct products), exported from country j, arriving at US coast c (c=west, east), via mode m (m=air,ocean) in year t. That is, we aggregate over all HS10 goods within an HS6, aggregate over all entry ports within the US east or west coasts, and aggregate over all months within a year. At this level of aggregation we have 2.1 million jkct observations.

Conceptually, aggregating over products in this way is equivalent to treating an HS6 product code as an industry and HS10 products as individual varieties within each industry.[14] Aggregating over months within a year and over customs districts within a coast may represent an aggregation of different exporting firms shipping within an HS6 code, or it may represent an aggregation of shipments for a given firm as it sells at different points in time or to customers located in different places within the US. We return to this point in Section 3.4.

A potential difficulty with aggregating over HS10 codes is that we may combine products that are fundamentally dissimilar in their shipment characteristics and shipping costs. This can be seen most clearly by inspecting the distribution of relative prices and relative freight prices, and we see very large differences in these variables in some cases. Accordingly, we trim our sample by dropping observations with either relative prices or relative freight costs below the 1st percentile and above the 99th percentile. We further trim our sample by eliminating HS codes in which the air share of revenues (calculated over all exporters and time periods) is less than 1 percent or greater than 99 percent.[15]

---

[12] See section 3.4 for more details on the construction and use of the number of shipments variable.

[13] Shipments into inland districts have very high air shares. While these districts do record some ocean shipments these presumably include overland transport for which we lack data on both costs and transit times.

[14] In a robustness check we maintain maximum disaggregation and use a probit model to predict the probability that air shipping is employed as a function of relative freight prices and ocean shipment days.

[15] Some HS codes have air shares very close to zero or very close to 1. This suggests that one mode is used almost

Table 1 reports data on the use of air shipment in our sample. Over all observations, air revenues represent 28 percent of import value[16], with higher shares for Europe (39 percent) and Asia (27 percent) than for other regions. This primarily represents differences in the product composition of trade across regions, as 52 percent of capital goods and 31 percent of consumer goods are air shipped, with smaller numbers elsewhere. The automotive category has the lowest air share (2 percent) because finished cars are rarely air shipped, but has higher air shares if we focus more narrowly on parts and components within automotive. Looking over all product codes that contain some parts and components trade, the air share is 41 percent.

A modest degree of aggregation allows us to compare revenues, prices, and shipping costs for very similar products coming from the same exporter that nevertheless use different shipping modes. Table 1 shows that in the sample as a whole we observe "mode mixing observations" – both air and ocean shipping employed – for trade equal to 75 percent of total import value. The mixing observations are much more common in Asia and Europe than in other regions, again reflecting product composition. Mode mixing is less common for food (50 percent) and industrial supplies (36 percent), but in other categories ranges from 85 to 92 percent of trade. Of note, the air share of trade for the mode mixing observations is similar to the air share of trade over all obervations. This indicates that trade omitted from our mode mixing observations is roughly balanced between observations using only air and using only ocean shipments.

Figure 2 shows the time series on the use of air shipping in the sample. Air revenues as a share of imports rise steadily until 2000, after which they fall. This pattern is found when using all observations, or only mode mixing observations, and it is found within every regional and product group listed in Table 1. That is to say, the large changes in air usage in our sample are not due to compositional change in what is traded but reflect within group changes. The pattern is also consistent with movements in cargo prices in this period, as the cost of air shipping fell until 2000, then rose sharply. These facts suggest that this is an ideal period for identifying modal substitution in the data and the extent to which higher air shipping prices trade off against more rapid delivery times.

Table 1 also reports on the premia paid to air ship goods. For each *jkct* observation we calculate air freight costs relative to ocean freight costs, both on a per weight and an ad-valorem

exclusively, and the outlying observations may be unusual situations or data errors.
[16] This is considerably smaller than the 33 percent share of air shipments in non-North American imports. The difference comes from dropping inland shipments from our estimation sample.

basis.   We calculate the air premium per kg as a ratio, g^A_{jckt} / g^O_{jckt}, and report the median value over all observations within the group.   For All Imports, air freight costs per kilogram are at the median 6.46 times higher than ocean freight costs per kilogram.   We calculate the ad-valorem air premium as a difference, $f^A_{jckt} - f^O_{jckt}$, and again report the median value over all observations within the group.   For All Imports, the median ad-valorem air premium is 5 percent. That is, ocean shipping costs are equivalent to a 3 percent tariff and air shipping costs are equivalent to an 8 percent tariff, so the use of air cargo raises delivered prices for the median good by 5 percentage points.   Figure 3 shows that there is significant variation over observations in the extent of these premia and at the high end firms are paying substantially more to air ship similar goods.   At the 90[th] percentile air freight costs per kg are 27 times higher than ocean freight, and the ad-valorem air premia reaches a hefty 34 percent.

The remaining variable needed is ocean shipping time to the US.   We employ a master shipping schedule of all vessel movements worldwide derived from the Port2Port Evaluation Tool. For most large exporters it is possible to construct a direct routing between the dominant ocean port in that exporter and a port or ports on the US east or west coast.   If there are multiple port-port combinations within a coast we take the average time to that coast.   For some smaller exporters there are no direct routings to one or both US coasts.   In these cases we construct all possible indirect routings (e.g. transiting through Hamburg, through Rotterdam, etc.) and choose the time minimizing indirect routing to each coast.

In some of our specifications we exploit cross-exporter variation, while in others we exploit within exporter variation across entry coasts.   We display transit times in Figure 4.   The horizontal axis measures the total transit time to the US, averaging over coasts, while the vertical axis measures the difference between transit times to the east coast and west coast for a given exporter.   Total transit time varies enormously across countries, from as little as a few days to as many as 48 days for some African exporters.   A key point here is that, due to quirks of geography, the shipment time difference to the US coasts varies considerably across countries.   For Latin America countries there is a minimal difference (0-4 days) in travel time to east and west coast, European shipments arrive on the east coast 10-14 days before the west coast, and some Asian shipments arrive on the east coast up to 14 days after the west coast.

3.2   Specification

We can now rewrite equation (11) in terms of observable and unobservable components,

providing subscripts to reflect the exporter j, product k, time t and coast c variation that we will exploit in the data.

$$\ln \frac{p_{jktc}^{A} q_{jktc}^{A}}{p_{jktc}^{O} q_{jktc}^{O}} = (1-\sigma)\ln\left(\frac{p_{jktc}^{A}}{p_{jktc}^{O}}\right) - \sigma\ln\left(\frac{f_{jktc}^{A}}{f_{jktc}^{O}}\right) + \sigma\tau\cdot\left(days_{jc}^{O}-1\right) + \varepsilon_{jktc},$$

$$\text{where } \varepsilon_{jktc} = \sigma\ln\left(\frac{v_{jktc}^{A}}{v_{jktc}^{O}}\right) + \ln\left(\frac{N_{jktc}^{A}}{N_{jktc}^{O}}\right) + \mu_{jktc}$$

(12)

At its most general, we will exploit variation across all dimensions (exporter j-product k –coast c-time t) of the data. In other specifications we experiment with different combinations of fixed effects to control for unobservable components in the errors, and thereby exploit only certain kinds of within-variation. We will also experiment with pooling over different product groups at different levels of aggregation. In our baseline regressions we pool over all HS6 industries, which implies that the key elasticities (σ, τ) are identical across all products. In some cases we pool over all observations within a given end-use category and estimate parameters specific to that end-use category.

3.3 Selection

Recalling equation (4) and Figure 1, we only observe exports from a given exporter-product if profits from exporting net of fixed costs are positive for some firms in the industry. When exploiting the variation detailed in (12), firms could be selected out of the sample because they have high marginal costs of production, face high shipping costs or fixed costs of exporting, or because they are selling a time sensitive good and their exports take a long time to travel to the US.

Concern about selection bias and zero trade flows at the aggregate level has motivated researchers to estimate a 2 step selection model. For example, Helpman, Melitz and Rubinstein (2008) use variables intended to capture differences across exporters (or exporter-importer pairs) in the level of exporting fixed costs to identify the probability of selection, and then include the inverse Mill's ratio as a regressor in the second stage. This works in their context because they are focused on aggregate trade flows. We are focused on highly disaggregated trade flows and in many specifications will incorporate exporter fixed effects. If we use country specific information in the first stage to predict the probability of observing a particular trade flow, the

16

inverse Mill's ratio varies across exporters but not within exporters over products or over time and so is absorbed by a country level fixed effect.

We are interested in selection operating at a much lower level, and want to ask whether a particular exporter j ships an HS6 good k to the US at time t. Rather than applying fixed cost proxies in the first stage, we use the volume of j's exports of k at time t to markets other than the US to indicate the latent profitability of jkt exports to the US.[17] For example, suppose Germany has a comparative advantage in machine tools.   Then Germany will export a high volume of machine tools to the rest of the world and it will be more likely that machine tool exports to the US will be sufficiently profitable to exceed fixed costs of trade.   In terms of Figure 1, instead of exploiting country level variation in the fixed cost schedule as in Helpman et al (2008), we are exploiting jkt variation in the position of the profit schedules.

We also include (the log of) ocean transit times in the selection equation as we are independently interested in how time affects the probability of a shipment to the US occurring (as well as its modal switching effects conditional on exporting). We do not observe the price or freight rate information if we don't observe trade and so these variables are not included in the selection equation.

3.4   Modal Extensive Margin

The distinction between revenues per firm (10) and revenues aggregated over N_j^m firms (11) displays the *potential* importance of the modal extensive margin -- defined as the number of firms within an HS6 product that export using a given mode of transportation.   Accounting for the modal extensive margin is an important contribution of this paper.   We use two approaches: controlling for the extensive margin with fixed effects; and directly measuring the extensive margin with data on the number of shipments.

How the *modal* extensive extensive margin adjusts in the data is not immediately clear, and it is not a margin that has been contemplated in the literature. Helpman, Melitz and Rubinstein (2008), for example, are focused on a cross-sectional question of whether at least one firm from any industry successfully exports to a given destination country at a point in time.   This is relevant

---

[17]Two recent papers, Autor, et al (2011) and Hummels et al (2011) use a similar strategy to identify latent profitability of exports in an instrumental variables context.   They use an exporter's sales to the rest of the world in a particular product code to instrument for its sales in a specific import market.

because with a continuum of firms spanning all of manufacturing activity, it seems highly likely that some firms are close to the point where small changes in costs induce selection in and out of the market.

In contrast, we employ data that are highly disaggregated (by exporter, HS6 product, time and entry point into the US) so there may be relatively few firms involved in any *jkct* trade flow.[18] If none of those firms is close to $\tilde{z}$ in Figure 1, then we will not see switching between modes in response to small cost shocks. In this case, a judicious use of fixed effects can absorb the modal extensive margin. In the next section we will describe a number of fixed effects strategies designed to absorb unobserved variation in product variety and quality.

Our second strategy supposes that firms within a given *jkc* trade flow switch modes over time in response to cost shocks so that fixed effects estimators will be insufficient to absorb the modal extensive margin. Here we use data on the number of shipments to control for the number of firms participating in the market.

When a firm exports into the US they electronically file a Shipper's Export Declaration Form, and the data on that form constitute one record. The public use imports data remove firm identifiers and aggregate over all the records with the same characteristics (i.e. same exporter, HS10 product, US customs district, month, and transportation mode), but include a count of records as a variable in the data. At the most disaggregated level of the imports data, most monthly observations consist of a single shipment, though some have multiple records. As we aggregate the data over products, months, and customs districts we are then counting the number of distinct shipments that occurred within each mode.[19]

Having multiple shipments for a *jkct* observation could reflect multiple shipments by the same firm (during different months within the year or to customers in different customs districts within the US), or it could reflect distinct shipments by multiple firms. Using the latter interpretation, the shipment count variable becomes a useful proxy for the number of firms participating in the market. We can then account for variation in the number of shipments (or

---

[18] There are not many data sets that allow a systematic evaluation of how many firms export in a very specific (destination x product) trade flow, but those that do suggest very few firms are involved. As an example, Hummels et al (2011) look at Danish firm-level exports at the HS6 x destination level, and report that in 90 percent of observations 3 or fewer firms are involved in the trade flow.

[19] Strictly speaking, the record count includes both air and ocean shipments within a given exporter-hs10-customs district-time observation. However, in 91.3% of these observations (by count) we can uniquely distinguish the mode used. For the remaining observations where we can't distinguish the number of shipments by mode, we assign a share of the shipments to each mode according the shipment value.

firms) across the two modes at the same jkct level of variation found in the rest of our data.

Of course, including the number of firms in each mode as a regressor would pose endogeneity concerns. If the firms switch between modes of transportation in response to demand fluctuations, shocks to the relative revenues may be correlated with the relative number of shipments. Happily, the structure of the model solves this problem. Starting from equation (11), we divide relative revenues by relative shipments to get back to equation (10). Provided that the number of shipments is a useful proxy for the number of firms, we now have an expression for the average revenues per firm that has eliminated the modal extensive margin problem. Rewriting estimating equation (12) with this adjustment, we have

$$
\ln\left(\frac{p_{jktc}^A q_{jktc}^A / N_{jktc}^A}{p_{jktc}^O q_{jktc}^O / N_{jktc}^O}\right) = (1-\sigma)\ln\left(\frac{p_{jktc}^A}{p_{jktc}^O}\right) - \sigma\ln\left(\frac{f_{jktc}^A}{f_{jktc}^O}\right) + \sigma\tau\cdot\left(days_{jc}^O - 1\right) + \varepsilon_{jktc},
$$

$$
\text{where } \varepsilon_{jktc} = \sigma\ln\left(\frac{v_{jktc}^A}{v_{jktc}^O}\right) + \mu_{jktc}
$$

(13)

3.5 Prices and Unmeasured Quality

The standard concern with including prices in a demand equation is that there are components of the error terms that are correlated with quantities demanded and with prices. This could reflect quality variation, or if the consumer values variety at the firm level, the number of distinct varieties for sale. Recalling estimating equation (12), the error term contains unobservable components

$$
\varepsilon_{jktc} = \sigma\ln\left(\frac{v_{jktc}^A}{v_{jktc}^O}\right) + \left(\frac{N_{jktc}^A}{N_{jktc}^O}\right) + \mu_{jktc}
$$

The terms v_{jktc}^m and N_{jktc}^m reflect demand shifters that are jktc and mode m specific and potentially correlated with regressors of interest, while the remaining term μ_{jktc} is uncorrelated with regressors. It is not feasible to construct instruments for prices that are jktcm varying, and so we use the rich panel structure of the data to account for the unobserved components of the demand equation.

19

In what follows we refer to "quality" but this should be read as any demand shifter that is potentially correlated with prices. For example, in specification (12) that uses relative revenues as a dependent variable we are treating the N_{jktc}^m extensive margin terms as if they were quality, so one can substitute the phrase "variety and quality" everywhere "quality" appears in the discussion. In specification (13) using relative revenues per shipment as a dependent variable we will eliminate the N_{jktc}^m terms by dividing both sides of the equation by a proxy for N_{jktc}^m.

The appropriate fixed effect estimators to use depend on the structure of the error term ε_{jktc}. We step through several cases and the associated estimator we employ. The simplest case is that quality varies by countries, commodities, US entry coasts and time, but for a given $jktc$ quality does not vary across transportation modes. In this case, expressing the equation in shares eliminates the relative qualities from the expression, or $v_{jktc}^a / v_{jktc}^o = 1$, $\varepsilon_{jktc} = \mu_{jktc}$. No fixed effects are needed, and OLS provides a consistent estimator.

Next suppose that quality varies across modes in an exporter-specific manner (i.e. the ratio of air/ocean quality is consistently high for German firms and low for Brazilian firms), but assume that the ratio of air/ocean qualities is time invariant and the same for each product and coast. In this case quality for mode m can be decomposed as $v_{jktc}^m = \upsilon_j^m v_{jktc}$. Expressing in shares eliminates the exporter-time-specific term, leaving an error of

$$\varepsilon_{jktc} = \sigma \ln\left(\upsilon_j^a / \upsilon_j^o\right) + \mu_{jktc}$$

Inclusion of exporter fixed effects (implemented via mean differencing) into the differenced equation eliminates the remaining problematic correlation. In this case we are exploiting within-exporter variation across products, entry coasts, and time.

Similarly, it may be that quality differs across modes and changes systematically across countries and across HS6 commodities such that we can write $v_{jktc}^m = \upsilon_j^m \gamma_k^m v_{jktc}$ and the combined disturbance is

$$\varepsilon_{jktc} = \sigma \ln\left(\upsilon_j^a / \upsilon_j^o\right) + \sigma \ln\left(\gamma_k^a / \gamma_k^o\right) + \mu_{jktc}$$

The nuisance terms here are eliminated by a combination of exporter fixed effects and HS6 commodity fixed effects. In this case the ratio of air/ocean quality can be different for Germany than Brazil, and different for machine tools versus microchips. Taking this one step further, we can

allow unobserved quality to be exporter-HS6 product specific or $v_{jktc}^m = \upsilon_{jk}^m v_{jktc}$. We absorb this unobserved quality variation with fixed effects that are specific to an exporter-HS6 product. In this case we are exploiting variation across time and entry coasts with a given exporter-HS6.

Finally, suppose that even this is not sufficient, and that the unobserved quality shifters differ across modes, countries and time in a way that cannot be log additively decomposed as above. This could be related to functional form. It may also be that an exporter experiences quality change over time that is product specific and where the degree of quality change is systematically related to modal choice. For example, Germany rapidly innovates in machine tool quality and new innovations are more likely to be airborne than older and more standardized products. To deal with this case we write quality differences as $v_{jktc}^m = \upsilon_{jkt}^m v_c$ and exploit the presence of coastal variation in the data, so that

$$\varepsilon_{jktc} = \sigma \ln\left(\upsilon_{jkt}^a / \upsilon_{jkt}^o\right) + \mu_{jktc}$$

Here we employ fixed effects that have exporter x HS6 product x time variation, and identify relevant parameters by exploiting cross-coast variation in all relevant variables. To see how this would work, our firms in the German machine tool industry have customers on the US East and West coasts. When selling to West coast customers ocean cargo must traverse the Panama Canal and requires 14 days longer than for shipments to the East Coast. This yields variation across coasts in the relative share of air shipping, relative time delays, and relative freight prices.

It should be noted that the ability to exploit variation in modal shares across $jktc$ observations allows us to control for unobserved quality variation in a manner that is considerably more general than what is found in the literature on estimating import demand elasticities or in the literature on quality and trade. The literature on quality and trade focuses on quality differences that are explained by exporter variables such as endowments, or importer variables such as per capita income, but makes no attempt to explain cross-product variation. The literature on import demand elasticities either ignores quality, or assumes quality-related shocks to demand and supply are uncorrelated.[20]

As an ancillary benefit, the aggressive use of exporter, exporter-product, and exporter-product-time fixed effects controls for many variables that may affect the likelihood of or

---

[20] See Hillberry-Hummels (2011) for a recent review of this literature. The state of the art papers, Feenstra (1994), and Broda-Weinstein (2006) assume demand and supply shocks are uncorrelated.

revenues from air shipping. This could include the exporter's level of development, delays associated with customs clearance, the quality of their infrastructure (in absolute terms or infrastructure for air relative to ocean shipping), or quirks of geography (being land-locked or having significant inland production).

3.6 Unit Values as Prices and the Endogeneity of Ad-Valorem Freight Rates

In many applications trade economists employ "prices" from trade data that are really unit values, or value/quantity. In our data, the quantity measure is kilograms, and so the "prices" we observe are values per kilogram. That may not correspond to a unit of quantity that is sensible from a utility perspective, which can create problems when using "prices" to estimate demand equations.

To see this, define prices p and shipping costs g as above, in terms of a quantity unit q that enters the utility function and is consistent across firms and shipping modes. We construct unit values as the ratio of total value and total kilograms shipped. $\hat{p} = pq / wq = p / w$, where $w = (\text{Total Weight in kg}) / q$ is a measure of product bulk. Compared to semiconductors, iron ore has very low unit values measured in dollars per kg, but this simply reflects differences in product bulk. Similarly, shipping prices per kilogram are $\hat{g} = g / w$.

If shippers set freight rates in terms of costs per kg, $\hat{g}$, the price of shipping relative to the utility measure of quantity, q, can vary across goods because product bulk (w) varies. In other words, we can rewrite the optimization problem keeping in mind the translation between q and kg when we write marginal costs per unit of quantity, $z + g = z + \hat{g}w$.

To see how this matters, consider the relationship between origin and destination unit values, captured in equation (8). When employing unit values and kg units of quantity, we can rewrite that as

$$\hat{p}\left(1 + \frac{\hat{g}}{\hat{p}}\right) = \frac{p}{w}\left(1 + \frac{g}{p}\right) = \frac{p}{w} \cdot f$$

We now have an expression where unit production costs and unit transportation costs are expressed in kg quantity units. When we use this in equation (12), we have

$$\ln \frac{p^{A}_{jktc} q^{A}_{jktc}}{p^{O}_{jktc} q^{O}_{jktc}} = (1-\sigma) \ln \left(\frac{p^{A}_{jktc} / w^{A}_{jktc}}{p^{O}_{jktc} / w^{O}_{jktc}} \right) - \sigma \ln \left(\frac{f^{A}_{jktc}}{f^{O}_{jktc}} \right) + \sigma \tau \cdot \left(days^{O}_{j} - 1 \right) + \varepsilon_{jktc},$$

The unit value "price" terms now reflect differences in the bulk factor, but the ad-valorem freight terms do not because bulk appears in both the numerator and denominator and therefore cancels. If the bulk factor varies across firms within a product category, then high bulk firms will choose to ocean ship and low bulk firms will air ship. This implies unit values differences will overstate price differences, or $\hat{p}^{a} / \hat{p}^{o} > p^{a} / p^{o}$. In this case, the coefficient on unit values will be biased toward zero even if we have perfectly controlled for unobserved quality variation.[21] For this reason, we will use the coefficients on relative freight prices and not unit value differences to identify σ.

A potential concern is that freight prices themselves are endogenous. This could arise because ad-valorem freight rates, *f*, are constructed by dividing per unit shipping charges, *g*, by prices. Or it could be that the unit shipping charges themselves are responsive to the quantities shipped. Evidence for both kinds of endogeneity are analyzed at length in Hummels et al (2009).

Hummels et al (2009) show that the major cause of endogeneity in ad-valorem freight prices is differences in the prices of products shipped. This is not a concern in the present context for two reasons. As we have just discussed, unit values can vary substantially for reasons unrelated to prices per (utility-relevant) quantities. If shippers charge per kg shipped, then high bulk translates into higher shipping costs both in per unit and ad-valorem terms. When we measure freight differences in air and ocean shipping, part of this reflects differences in bulk that increase the freight price per (utility-relevant) units of quantity. This moves the relative prices of air and ocean in a way that identifies the elasticity of substitution. In other words, it provides precisely the variation that is useful in identifying the relevant parameters.

Suppose however, that the unit value differences between air and ocean shipments do reflect true price variation and not merely differences in bulk. Were we to omit unit values from the regression, or were the regression to omit important quality variation, this could potentially bias the coefficients on freight charges. However, as discussed at length in the previous section,

---

[21] In the limiting case, suppose that prices per (utility-relevant) quantity unit were the same for two firms.   Bulk and value per kg differ across firms in a way that shifts some goods to boats and some to planes. Since the unit value difference, in and of itself, is irrelevant to the consumer there should be no response of modal choice to unit value differences.

the regressions explicitly include unit values and employ an aggressive differencing strategy to remove quality variation as a source of differences in relative demands.

A secondary concern is that unit shipping charges, g, are themselves endogenous to quantities shipped. For example, exporters that trade higher quantities of goods will invest in better transportation infrastructure, and there will be more entry by transport providers on densely traded routes. As Hummels et al (2009) show, these scale effects are characteristics of trade routes (i.e. particular exporter-importer combinations) and are much stronger when considering thinly traded developing country routes, not densely traded routes involving the US. Since we employ only US imports data, use mode differencing and include exporter, exporter-product, and exporter-product-time fixed effects, these scale differences are differenced out or swept into constants.

However, for the sake of completeness we experiment with instrumenting strategies for freight rates. Finding instruments that vary across jkct observations is difficult. However, if the regessors are sequentially exogenous such that their lags are not systematically related with the contemporaneous disturbance, then lagged variables are valid instruments. The instrumental variable assumption is that past freight rates are correlated with contemporaneous freight rates, but they do not significantly explain today's relative revenues.

4   Results

To summarize the discussion to this point, we have shown that consumers who value time savings will trade off the higher cost of air shipping against the higher implicit quality of a good that arrives several days earlier. The precise value consumers attach to time savings can be extracted by estimating the parameters in that tradeoff using equations (12) and (13) along with various fixed effect estimators. The coefficient on relative freight prices identifies consumer sensitivity to price changes, σ, and the coefficient on days in transit identifies the quality-reducing effect of shipment delays, measured in terms of reduced quantities sold, $\sigma\tau$. Combining the two yields the price or tariff-equivalent of the time delay, τ.

4.1 Baseline Specification

We begin with estimates that do not condition on selection and pool over all products. Recall that differences across commodities in the level of relative revenues and relative freight

prices will be swept out in specifications using commodity fixed effects. Pooling in this way maximizes the available observations, simplifies presentation, and yields estimated parameters that are observation weighted average of the commodity level response. We provide commodity specific parameter estimates below.

Table 3 reports the results for the relative revenue equation (12) with different sets of fixed effects. Standard errors are clustered on exporters. Across all five columns, we see two clear patterns: increased ocean shipment times induce substitution toward air shipping, and high relative freight prices for air shipping induce substitution toward ocean shipping. Table 3 also shows that the coefficient on unit values ranges from small and negative to small and positive. This is consistent with our discussion in Section 3.6 noting that unit values are not prices and that differences in product bulk make variation in freight rates, not unit values, a more reliable measure of the price elasticity of demand.

Consumers valuation of timeliness τ is constructed as the ratio of two coefficients: ocean days divided by relative freight prices, and the standard errors are constructed using the Delta Method. We estimate that time sensitivity τ ranges from 0.003 for the OLS estimates to .021 for the fixed effects that exploit differences across coasts for a given exporter-product-time period. At the high end this implies that one additional day in transit is equivalent to a 2.1 percent tariff.

There are significant differences in magnitudes across the specifications and so it is worth understanding where those differences come from. Across the specifications the number of observations changes depending on the requirements of the fixed effects. The most pronounced change comes in the coast-differenced specification because it requires that we observe for a given exporter-HS6 product both air and ocean shipments to both coasts. However, if we apply specifications 1-4 to the sample for column 5 we get very little change in the table.

Rather, the coefficient pattern reveals the importance of controlling for unobserved heterogeneity. As increasingly stringent fixed effects absorb progressively more variation, the estimated price elasticity of demand falls and the impact of the transit time on the relative revenues increases. The most pronounced change in transit days comes once we introduce exporter fixed effects (singly, or in combination with other FE). Across countries there are likely unobserved differences in the relative quality of airport and ocean port infrastructure. There may also be unobserved differences across countries in inland shipment costs and inland transit time that primarily affect ocean shipments. That is, a country with many exporting firms in its interior will

25

endure costs and long inland transit times as goods are trucked to the coast for ocean shipping but these can be avoided by using interior airports. This variation is removed from our data, leaving identification of the days in transit variable to come from differences across countries in shipping time to the US east versus west coast. If these unobserved country characeristics change over time they will still plague specifications in columns 2-4, but they will be eliminated in column 5.

There is a pronounced change in the freight coefficient (the price elasticity of demand) when we include commodity fixed effects (additively, or interacted with other FE). Some commodities are more likely to be air shipped than others due to physical characteristics such as perishability or weight and size and this unobserved information is absorbed by the commodity effect. Rather than identifying this coefficient across variation in dissimilar goods (the small air freight premium and high air shares for electronics compared to the large air freight premium and low air share for bulky furniture), the commodity FE columns identify the coefficient from freight cost variation across different source countries and time periods for a given HS6 product.

A final reason we may see differences across the columns is heterogeneity in product quality and variety across observations. In our discussion of specification issues in Section 3 we indicated many possible dimensions of quality heterogeneity that are controlled for across the different specifications in Table 1. What we see here is consistent with the view that the OLS estimates overstate the response of relative quantities to relative freight prices differences because of that unobserved heterogeneity. To understand the direction of the bias, suppose that air shipped goods are higher quality than ocean shipped goods, and that higher quality goods have lower ad-valorem freight rates (following the discussion of per unit freight charges in Section 3.6). In the absence of fixed effects that control for quality this generates a negative correlation between quality and relative freight rates, and the omitted variable bias is towards finding a larger negative effect. More stringent fixed effects eliminate the bias.

4.2 Accounting for Selection and the Extensive Margin

The revenue specifications in Table 3 estimate equation (12) assuming that the modal extensive margin, the relative number of firms employing ocean and air transport for a given jkct observation, is uncorrelated with the regressors after including various fixed effects. This is a reasonable approach if the modal extensive margin exhibits little within-jkct variation, but it is problematic if firms substitute between modes in response to cost shocks. We address this case

by estimating equation (13) using revenues per shipment as a dependent variable. If the number of shipments is a good proxy for the number of firms operating in each mode, then our dependent variable measures average revenues per firm.

Table 4 provides estimates of equation (13) with various fixed effects, and shows the same sign pattern as Table 3: high relative air freight prices reduce relative air revenues, and longer transit times raise relative air revenues. Our estimates of τ range from 0.004 to 0.006 (one day is equivalent to a 0.6 percent tariff). Notably, all coefficient estimates are smaller than in Table 3. This suggests that high air freight prices and long transit times lower both the number of shipments and revenues per shipments, and the smaller estimates in Table 4 are due to eliminating the number of shipments channel. We also see more consistency across the columns in Table 4, in contrast to Table 3. This suggests that the number of shipments is an important source of unobserved heterogeneity removed by the Table 3 fixed effects. In Table 4 they are differenced out of the dependent variable and so the fixed effects have less impact on the estimates.

In understanding the economics behind Tables 3 and 4, the key question is what the number of shipments are actually capturing. One view of the data is that we are capturing an active modal extensive margin. As we lower air freight prices or increase shipping times we see higher air revenues, and some of this response takes the form of firms switching from ocean to air shipping. When we control for this channel we identify a per firm revenue response and so the estimated elasticities in Table 4 correspond more closely to the parameters from the model. This interpretation is consistent with the Helpman, Melitz, and Rubinstein (2008) argument that ignoring the effect of trade costs on the extensive margin will tend to overstate their impact at the firm level.

An alternative view is that changes in the number of shipments do not reflect firms switching between modes, but instead reflect changes in the number of shipments made by a fixed set of firms. Consider how a single exporting firm might respond to a cost shock that boosts demand for its products. It might make shipments to customers in several different customs districts instead of one, or ship every month rather than every other month. This shows up in the data as a rise in the number of shipments. In this case, calculating revenues per shipment as in Table 4 eliminates an important channel through which a single firm could see enhanced revenue. We have no way of distinguishing which of these views is correct, and think the truth lies somewhere in the middle, that is, with per day time costs somewhere between 0.6 and 2.2 percent

27

ad-valorem.

There is another extensive margin potentially at work in these data, the possibility that high costs and long shipping times could cause a country to have zero exports to the US in a particular product. To address this Table 5 reports the result of a 2 stage Heckman selection estimator. As detailed in Section 3.4, in the first stage we predict the probability that exporter j has positive sales of product k to the US at time t using two variables: j's exports of k to the rest of the world at time t, and (log) ocean days for exporter j to the closest US coast. We then include the inverse Mills ratio in the second stage of the specifications used in Tables 3 and 4.[22]

Table 5 reports estimates of the first and second stages of the Heckman estimator.[23] The first stage is of independent interest for future studies that might desire a selection estimator that operates at the exporter-product-time level. The value of country j's exports of product k to the rest of the world excluding the US is an excellent predictor of the probability of observing those same exports to the US. Long transit times are negatively correlated with the likelihood of exporting to the US, with an elasticity of -0.135 and a marginal effect (at the means) of -0.024. The first stage is not a fully specified model of the exporting decision but taking the marginal effect at face value we can calculate the impact of a reduction in shipping times on the probability of seeing trade. The mean shipment time in our data is 23 days, about the average trip length from East Asia to the US. Decreasing this by 3 days (the average trip length from Europe to the US) increases the probability of any one product being exported by 0.13 percent.

Turning to the second stage we see that the inverse Mills ratio is strongly correlated with relative revenues and relative revenues per shipment. However, the coefficients of interest are very similar to those found in Table 3 and 4. Taken together this suggests that the selection correction does affect relative revenues, but is not correlated with the variables of interest once we have included other controls in the estimation.

4.3. Additional Robustness Checks

Table 6 reports a set of additional robustness checks. For brevity we report only the

---

[22] We do not include the coast-differenced specification. Our selection variables generate an inverse Mills ratio with jkt variation but it does not vary across coasts for a given exporter-product-time. When we difference all variables across coasts, the Mills ratio is eliminated. Put another way, once we control for exporter-product-time effects in the coast differencing estimation we have no variation left to predict selection into the sample.

[23] Note that the coefficient estimates in the selection specification are identical across the specifications. That is because for each specification we estimate exactly the same selection model over the same sample.

coast-differenced specification using relative revenues and relative revenues per shipment (similar to column 5 from Tables 3 and 4), and do not include the Heckman correction. Results are similar with other specifications.

In our main specifications we estimate a linear effect in transit time, which treats an increase from 6 to 7 days the same as an increase from 26 to 27 days. However, at sufficiently long horizons consumers may be indifferent to marginal changes in delivery time. In columns 1 and 2 we experiment with a quadratic in transit time, and find that delays have diminishing impact at longer time horizons. At the sample mean of 23 days (the average travel time for Asia), our estimated effects match those from Tables 3 and 4: ad-valorem time costs of 2.3 percent per day (for revenues), and 0.7 percent per day (for revenues per shipment). Central and South America have shorter travel times and much higher time costs at the margin. At 35 days of travel time (the average for Africa, the most temporally distant region) the effect just reaches zero.[24]

In our main sample we trimmed outlying observations for relative prices and relative freight rates, and dropped products in which the average air share for an HS code was less than 1% or greater than 99%. In columns 3 and 4 of Table 6 we include all these dropped observations. We find somewhat larger estimates on transit days, smaller estimates on freight prices and larger estimated values for time sensitivity.

In Section 3.6 we discussed how our fixed effects specifications account for possible sources of endogeneity in freight rates. For completeness we also experiment with using lagged values as instruments. In columns 5 and 6 we instrument the current period freight rate with its first lag, and to examine the impact of dynamics, we also include the lagged dependent variable and instrument for today's freight rate with the second lag of the freight rate.[25] The conclusion is the same for both dependent variables. Compared to the baseline estimates reported in Tables 3 and 4, we find somewhat higher elasticities for the freight rate variables, somewhat smaller coefficients on days in transit, but the fundamental message is unchanged.

We also experimented with with re-estimating our main specifications on sub-samples of the data (e.g. using only observations from Europe and Asia) or allowing slope coefficients to vary across sub-samples (by exporter income, by year, by season within each year). While intercepts

---

[24] Note that these estimates do not rely on the full range of transit time (from 3 to 48 days) in the data, but instead fit the quadratic on the variation in coast-differenced transit time shown in Figure 4.

[25] It is worth emphasizing that the mean difference panel IV estimator in Column 3 and 4 subtracts the mean over the two coasts, not the mean over time. A mean difference estimator that subtracts the mean over time is not consistent when we include the lagged dependent variable.

varied over these sample cuts, indicating differences in the average reliance on air shipment, we found no significant changes in slope coefficients relative to results reported in Tables 3 and 4. There were, however, significant differences across product categories, a point we take up in the next section.

As a final robustness check we experimented with a probit estimation based on a sub-sample of the data in which only a single transport mode was chosen. The estimation is based directly on equation (3), and is conceptually similar to the relative revenue specification, except that here we estimate the probability that air shipping is chosen as a function of freight prices and transit time. Details of the derivation and specification are reported in the appendix. We find coefficients with the same sign pattern as those in Tables 3-6 and using point estimates, time effects of similar magnitude. However, the estimated effect of shipment time is not statistically significant once we cluster standard errors. We attribute this loss in precision to three factors: losing information about the quantity of sales in the dependent variable; the inability to incorporate rich controls for quality heterogeneity in the demand equation; and the need to estimate rather than observe shipping costs for the transport mode not chosen.

4.4. Estimating the Value of Time by Commodity

The specifications above allow for heterogeneity in the intercepts, but impose homogenous slope coefficients across broad product groups. In other words, we assume that all product categories have the same modal use response to changes in freight prices and to time delays. This has the advantage of maximizing available observations and sources of variation but at the cost of losing potentially interesting information about how time values differ across commodities.

To examine heterogeneity in the coefficient estimates we grouped products by End-Use Category and re-estimated equations (12) and (13) separately for each, using Exporter x HS6 fixed effects. We report results for 1-digit End Use groupings in Table 7. Focusing on relative revenues, equation (12), we find that results are qualitatively similar to Table 3 across all groups. However, we find substantially higher time values for Automotive goods (.043, that is, one day is equal to a 4.3 percent ad-valorem tariff) and for foods and beverages (.031) than for Industrial Supplies, Capital Goods and Consumers Goods (all around .02). In the case of Automotive goods this is driven by a combination of a higher coefficient on transit days and a lower coefficient on freight costs.

30

When we examine relative revenues per shipment to control for the modal extensive margin, equation (13), we see similar sign patterns, but much more dispersion in the estimates. Here the high time value categories are Automotive (.013) and Capital Goods (.009), with much lower estimates for Consumer Goods and Industrial Supplies (.004), and an (insignificant) negative estimate on both transit days and time value for Foods and Beverages. Above we highlighted two alternative stories (firms changing modes, or firms changing the number of shipments to reach a greater number of customers) for the more modest coefficients found when using equation (13). What seems likely in the case of Foods and Beverages, where storability is particularly important, is that firms respond to long shipment times by making more frequent shipments on airplanes. Once we control for this channel there is no remaining response in terms of revenues per shipment.

The one-digit End Use categories are still fairly broad and we next group products at the most disaggregated End-Use Category and re-estimate equations (12) and (13) separately for each, using Exporter x HS6 fixed effects. Figures 5 and 6 show the distribution of time values, with the histogram of all estimated coefficients shaded in grey and the histogram featuring only statistically significant estimates shaded in black.[26] For Figure 5, the mean over the individual group estimates shows an average time sensitivity of about 0.02, which is very similar to Table 3, column 4.[27] However, there is significant heterogeneity in the coefficient estimates. Most of the mass of this distribution is positive, and we see some time values ranging as high as .072 or one day being worth 7.2 percent ad-valorem. For Figure 6, we again see an average effect similar to Table 4, column 4, and again see considerable dispersion.

As we disaggregate we face a tradeoff – greater flexibility in allowing the model to fit different coefficients for different product categories versus the possibility of greater imprecision due to the reduced number of observations from which to identify those coefficients. The question is then whether the coefficient heterogeneity in Figures 5 and 6 reflect true variation in response parameters or noise.

A possible indication that these estimates reflect true variation can be found in Djankov, Freund, and Pham (2010). They used time cost estimates taken from an earlier draft of this paper and showed that countries with long customs delays experienced relatively sharp reductions in

---

[26] The histograms omit insignificant point estimates lying 2 standard deviations from the mean.
[27] This is to be expected, as the pooled estimates in Table 3 are a consistent estimate of the average impact over products (see Zellner, 1969; Pesaran and Smith, 1995).

exports for goods that exhibited the highest time sensitivity.

To get at this issue more systematically, we focus on two characteristics of products that seem especially relevant for timeliness: perishability, and whether the product is a manufactured intermediate input.   To capture these characteristics we drill down to the HS-10 digit level and identify product descriptions that contain the word "fresh" (for perishability) or "parts" or "components" (for intermediate inputs).   We then calculate, for each exporter-HS6-time observation the value share of HS-10 products containing those words, and include this variable both independently and interacted with transit time.[28]

Results are reported in Table 8.   Focusing on the revenues per shipment we see two interesting findings.   First note that a higher "fresh" share increases the use of air shipment, but does not significantly interact with transit days.   A likely explanation is that products like "fresh fish" are so time sensitive that any delay longer than a few days ruins the product.   As a consequence the effect shows up entirely in a higher use of air shipment for all exporters, regardless of ocean transit time to the US.

There is a very different pattern with parts and components.   An increase in the parts and components share of trade for a given exporter-HS6 product results in a sharp increase in the time sensitivity of that trade. Comparing a product with zero component share to one that is 100 percent components raises time sensitivity by 60 percent.

4   Conclusion

Airplanes are fast, expensive, and carry a large and rising share of world trade.   In this paper we model substitution between the use or air and ocean cargo in trade and show how to extract consumers' willingness to pay for time savings from that choice. Our estimates control for selection intro trade, for unobserved variation in quality, for endogeneity of freight charges, and for extensive margin changes in the sets of firms participating in trade by mode.   We estimate that each day in transit is worth 0.6 to 2 percent of the value of the good, and that long transit delays significantly lower the probability that a country will successfully export a good.   Our estimates vary over goods, with especially high time sensitivity exhibited in end use categories motor vehicles and parts, and capital goods, and in HS classifications with high shares of parts and

---

[28] Recall that our observations are at the level of HS-6digit products and that we include exporter - hs6 fixed effects. By using the movements in the shares of the HS-10 digit products we induce changes over time for a given exporter-hs6 that can be used to identify differences in the coefficient.

components. Comparing a product with zero component share to one with a 100 percent component share raises time sensitivity by 60 percent.

This last result connects two important changes in patterns of international specialization and trade.   In the last several decades the cost of air cargo has dropped an order of magnitude, and the use of air cargo has risen 2.6 times faster than ocean cargo.   At the same time there has been a sharp rise in intermediate input trade as firms fragment production across multiple locations. While many products are time sensitive due to inventory holding costs, perishability, rapid technological obsolence, and uncertain demand, these problems are magnified in the presence of fragmentation.   It seems reasonable to conclude that the sharp reduction in the cost of linking far flung production sites through fast moving airplanes has been an important factor in growing fragmentation worldwide.

Finally, our results are relevant to the increased emphasis on trade facilitation – identifying regulatory or other nontariff barriers to trade – in trade negotiations and among aid and development groups such as USAID and the World Bank.   Many efforts to facilitate trade, such as streamlining customs procedures or improving port infrastructure, generate benefits measured in days saved.   With our estimates of the value of each day saved one can then calculate the monetary benefits of these initiatives and how they compare to the cost incurred.

References

Aizenman, J. (2004) "Endogenous pricing to market and financing cost" *Journal of Monetary Economics* 51(4), 691–712

Autor, D., Dorn, D. and Hanson, G. (2011) "The China Syndrome: Local Labor Market Effects of Import Competition in the United States" mimeo.

Baier, Scott and Bergstrand, J (2001): Tariffs, Transport Costs, and Income Similarity," *Journal of International Economics,* Vol. 53, No. 1, Feb 2001, 1-27

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

Broda, C. and Weinstein D. (2006) "Globalization and the Gains from Variety," *Quarterly Journal of Economics,* Volume 121, No. 2.

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

Choi, YC, Hummels, D. and Xiang, C. (2009) "Explaining Import Quality: the Role of the Income Distribution" *Journal of International Economics* 77, pp. 265-276.

Cristea, A., Hummels, D., Puzello, L., and Avetisyan, M. (2011) "Trade and the Greenhouse Gas Emissions from International Freight Transport"   NBER WP 17117.

Djankov, S., Freund, C., and Pham, C. (2010). "Trading on Time," *The Review of Economics and*

Statistics, 92(1), pp. 166-173.

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

Feenstra, R. (1994), "New Product Varieties and the Measurement of International Prices" *American Economic Review*, 84(1) March 157-177.

Feenstra, R. and Wei, S. (2009), "Introduction to 'China's Growing Role in World Trade'" NBER 14716.

Hallak, J. C. (2006) "Product Quality and the Direction of Trade", *Journal of International Economics,* 68 (1) pp. 238-265.

Hallak, J.C. and Schott, P. (2011) "Estimating Cross Country Differences in Product Quality" *Quarterly Journal of Economics*, forthcoming.

Harrigan, J (2010), "Airplanes and Comparative Advantage" *Journal of International Economics* 82, pp. 181-194.

Harrigan, J and Venables, A (2006) "Timeliness and Agglomeration" *Journal of Urban Economics* 59, March, pp 300-316.

Heckman J. (1976), " The common structure of statistical models of truncation, sample selection, and limited dependent variables and a simple estimator for such models, " *Annals of Economic and Social Measurement*, 5: 475-492.

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

Hummels, D. (2007a) "Transportation Costs and International Trade In the Second Era of Globalization", *Journal of Economic Perspectives* 21 pp. 131-154

Hummels, D. (2007b), "Calculating Tariff Equivalent for Time in Trade" Nathan Associates report for US Agency for International Development.

Hummels, D, Ishii, J., and Yi, Kei-Mu, "The Nature and Growth of Vertical Specialization in World Trade", *Journal of International Economics*, 54, pp 75-96.

Hummels, D., Jorgensen, R., Munch, J., and Xiang, C. (2011), "The Wage Effects of Offshoring: Evidence from Matched Worker-Firm Data" NBER 17496

Hummels, D. and Klenow, P. (2005), "The Variety and Quality of a Nation's Exports" *American Economic Review* 95 pp. 704-723.

Hummels, D., Lugovskyy, V., and Skiba, A. (2009) "The Trade Reducing Effects of Market Power in International Shipping"   *Journal of Development Economics* 89 pp. 84-97.

Hummels, D. and Schaur, G. (2010) "Hedging Price Volatility using Fast Transport", *Journal of International Economics* 82 pp. 15-25.

Johnson, R. and Noguera, G. (2011) "Accounting for Intermediates: Production Sharing and Trade in Value Added" *Journal of International Economics*, forthcoming.

Khandelwal, A (2010), "The Long and Short of Quality Ladders" *Review of Economic Studies* 77(4) pp. 1450-1476.

Melitz, M. (2003) "The impact of trade on intra-industry reallocations and aggregate industry productivity" *Econometrica*. 1(6); pg. 1695

Pesaran M. and Smith R. (1995), "Estimating long-run Relationships from Dynamic Heterogenous Panels," *Journal of Econometrics,* 68, pp. 79-113.

Schott,P. (2004) "Across-Product versus Within-Product Specialization in International Trade. *Quarterly Journal of Economics* 119(2):647-678.

Yi, KM. (2003)"Can Vertical Specialization Explain the Growth in World Trade" *Journal of Political Economy* 111.

Zellner A. (1969) "On the Aggregation Problem: A New Approach to a Troublesome Problem," in K.A. Fox et al., eds., Economic Models, Estination and Risk Programming: Essays in Honor of Gerhard Tintner, Springer-Verlag, Berlin, 365-378.

Appendix 1.   The Probit specification

Suppose that there is a single firm in j producing good k (or that all firms producing k in location j are symmetric). Then we can rewrite the inequality in (3) as a probability model for modal choice. First we use equations (7), (8), and (9) to rewrite the difference in marginal costs (unobservable) as a function of (obervable) export prices and shipping charges. If a firm charges the same origin price regardless of shipping mode, the difference in marginal costs reduces to the difference in ad-valorem shipping costs.   Using this in (3) implies that air will be chosen if

$$(1-\sigma)\left(\ln f_j^A - \ln f_j^O\right) - \sigma\tau\left(days_j^O - 1\right) > 0 \tag{14}$$

With normally distributed random shocks to modal choice we have a standard probit model.

$$P\left(m_{jt} = a \mid x\right) = \Phi\left[\delta + (1-\sigma)\left(\ln(f_j^a) - \ln(f_j^o)\right) - \sigma\tau\left(days_j^o - 1\right)\right] \tag{15}$$

Sample construction:

We begin with the sample of 45 million observations discussed in Section 3.1, dropping imports from Canada, Mexico, and those imports entering in inland customs districts or outside the Mainland US.   We aggregate by coasts, but we do not aggregate over months within a year or over HS10 codes within an HS6 code. In the data we only observe the freight rate for the mode that was actually chosen.   We predict the the unobserved freight rates by fitting a cost equation based on observations where we do see freight data and then fitting them to the remaining observations using observable characteristcs.   We employ data on unit freight costs (per kg) and product prices (per kg) that are specific to the exporter j, HS10 product k, customs distrct d, time t and mode m, and shipment distances that are specific to jdm.   We estimate the following equation separately for each mode-HS4 industry so that the intercepts and coefficients are mode-HS4 specific

$$ln\left(g_{jckt}^m\right) = \beta_1^{m,hs4} + \beta_2^{m,hs4}ln\left(p_{jckt}^m\right) + \beta_3^{m,hs4}ln\left(distance_{cj}^m\right) + u_{jckt} \tag{16}$$

Using the fitted values, we then predict the unobserved freight rate using observable characteristics (mode-HS4, unit prices, and shipment distance).   Much of this data is not specific to the firm, and is equivalent to identifying a kind of cost schedule facing firms with particular shipment characteristics at a point in time.   The price data are unique to the firm, and following

the model, we assume that a given firm charges the same price for both shipment modes. For reasons detailed above, differences in the air-ocean freight cost differential across firms also results from product price differences, and we include the product price in the regressions. We use customs district data to be more precise about shipment distances but to keep the number of observations manageable in the probit equations, we aggregate the fitted freight rates up to the coast level. Finally, we omit HS10 products with fewer than 30 observations and drop outlier cases (highest and lowest 1% of freight differences, and cases where fitted air shipping charges are less than fitted ocean shipping charges.

Additional Specification Issues:

As in the relative revenue specification we use a Heckman two-step in which the first stage predicts the probability of positive exports (to the US in product k at time t) using country j's exports of world k to the world as a whole, and shipping time from j to coast c.

Quality differences enter as demand shifters in equation (12). The probits are somewhat different; since price and quality for a single firm is the same regardless of mode both terms cancel in the equation. However, when we go to the data we exploit variation across firms, and their prices and quality will vary. Consider a model where quality is expensive to produce, like Baldwin and Harrigan (2010). High productivity firms produce high quality goods at higher marginal cost. Using the logic of our model with per unit shipping charges, this translates to a higher likelihood of using air shipping. When we look across firms unobserved quality will be correlated with prices, the gap between air and ocean shipping charges, and the likelihood of using air shipping. To account for this possibility we incorporate product prices directly in the probit equation to absorb quality variation.

We can then rewrite (15) to reflect the variation across exporter j, HS10 products k, US entry coasts c, and time t (at monthly frequencies) and incorporate shipment prices. For simplicity, we follow Table 1 and estimate a single pooled regression.

$$P\left(m_{jkct} = a \mid x\right) = \Phi\left[-1.157\left(\ln(f_{jkct}^{a}) - \ln(f_{jkct}^{o})\right) + 0.008\left(days_{jc}^{o} - 1\right) + 0.708\ln(p_{jkct})\right]$$
$$\quad (.41)^{***} \qquad\qquad (.008) \qquad\qquad (.036)^{***}$$

High relative air shipping costs and low product prices reduce the probability of using air shipping, and both effects are highly significant. An additional day in transit has a positive effect on the

probability of using air shipping with magnitudes similar to those found in the main model in the paper. However, once we cluster the standard errors it is not signficant. There are a few possible reasons for these weaker results in contrast to those from them revenue equation. First, it may be that the model variables have a weak effect on the probability that any particular firm chooses air shipping but a strong effect on quantities of trade conditional on a mode being chosen. Second, and consistent with the changing coefficients across columns 1-5 of Table 1, there may be substantial heterogeneity in quality or variety that is correlated with use of air shipment and with the regressors. We successfully eliminate this in the revenue equations but not the probability equations. Third, it may be that predicting rather than observing freight rates introduces noise in the regressor and attenuation bias.

Appendix 2:   Relative Revenues in a Model with Firm Heterogeneity

Let z be the constant marginal cost of productions and assume that firms draw their productivity from some probability distribution as in Melitz (2003). Based on equation (5), the sales of a firm located in country j are

$$r(z)_j = p(z)_j q(z)_j = E \left(\frac{1}{v_j^m \exp(-\tau \cdot days_{rj}^m)} \right)^{-\sigma} \left(\frac{\sigma}{\sigma-1}(z + g_j^m) \right)^{1-\sigma} \tag{17}$$

Assume that the cumulative distribution $W(z)$ describes the distribution of z across firms in country j with support H, L such that $H > L > 0$. Let $L < z_j^o < z_j^e < H$ be a cutoff such that only firms with $z < z_j^e$ export, and only firms with $z < z_j^o$ ocean ship. These cutoffs are endogenously determined by the profits at firm realizes employing alternative modes of transportation as illustrated in Figure 1, where across countries $\tilde{z}$ represents z_j^o and $\bar{z}$ represents z_j^e. With a mass of M_j active firms in the economy, the aggregate export revenue generated by the ocean shippers is then the integral of (17) over all $z < z_j^o$ applying the distribution $W(z)$ and the mass of firms M_j. This integration is difficult as it is nonlinear in the marginal cost. Therefore we apply a second order Taylor approximation to obtain

$$(z + g_j^m)^{1-\sigma} \approx (\mu_j^m + g_j^m)^{1-\sigma} + \frac{(1-\sigma)(\mu_j^m + g_j^m)^{1-\sigma}(z - \mu_j^m)}{\mu_j^m + g_j^m}$$

where μ_j^m is the expected marginal cost of production conditional on transport mode m.

Applying the linearization to (17), the aggregate export revenues generated by the exporters in country j that ocean ship are

$$\int_L^{z_j^o} p_j(z)q_j(z)M_j w(z)dz = \left(\frac{\sigma}{\sigma-1}\right)^{1-\sigma} \frac{E}{\left(v_j^o \exp(-\tau \cdot d_j)\right)^{-\sigma}} M_j \tag{18}$$

$$\times \left((\mu_j^o + g_j^o)^{1-\sigma}\right) \int_L^{z_j^o} \left[1 + \frac{(1-\sigma)(z - \mu_j^o)}{\mu_j^o + g_j^o}\right] w(z)dz$$

Multiply and divide the right hand side by the probability of being and ocean shipper, $W(z < z_j^o)$.

Then, because $\int_L^{z_j^o} \left[(z - \mu_j^o)\right] \frac{w(z)}{W(z < z_j^o)} dc = 0$, the aggregate export revenues of the ocean

shippers simplify to

$$R_j^o = \int_L^{z_j^o} p_j(z)q_j(z)w(z)dz = \left(\frac{\sigma}{\sigma-1}\right)^{1-\sigma} E\left(v_j^o \exp(-\tau \cdot d_j)\right)^\sigma (\mu_j^o + g_j^o)^{1-\sigma} N_j^o, \tag{19}$$

where $N_j^o = M_j W(z < z_j^o)$ is the mass of ocean shippers that serve the export market. Similarly for the revenues generated by the air shippers we obtain

$$R_j^a = \left(\frac{\sigma}{\sigma-1}\right)^{1-\sigma} E\left(v_j^a \exp(-\tau)\right)^\sigma (\mu_j^a + g_j^a)^{1-\sigma} N_j^a \tag{20}$$

where $N_j^a = M_j W(z_j^o < z < z_j^e)$ is the mass of air shippers. Equations (19) and (20) explain the industry level observation in the data that mix air and ocean export revenues in a given time period. They show that the relative export revenues derived in equation (11) for firms that are homogenous within each mode of transportation are a second order approximation of the relative revenues that we obtain if the productivity is heterogenous across firms even within each mode of transportation.

Applying the same linearization strategy as for the aggregate revenues, we obtain the aggregate quantities

$$Q_j^o = E\left[\frac{\left(\frac{\sigma}{\sigma-1}(\mu_j^o + g_j^o)\right)}{v_j^o \exp(-\tau d_j)}\right]^{-\sigma} N_j^o, \quad Q_j^a = E\left[\frac{\left(\frac{\sigma}{\sigma-1}(\mu_j^a + g_j^a)\right)}{v_j^a \exp(-\tau)}\right]^{-\sigma} N_j^a \tag{21}$$

Dividing the aggregate revenues by the aggregate quantities we then obtain

$$\overline{p^o_j} = \frac{R^o_j}{Q^o_j} = \frac{\sigma}{\sigma - 1}\left(\mu^o_j - g^o_j\right), \quad \overline{p^a_j} = \frac{R^a_j}{Q^a_j} = \frac{\sigma}{\sigma - 1}\left(\mu^a_j - g^a_j\right) \tag{22}$$

To derive the estimation specification divide the aggregate air revenues (20) by the aggregate

ocean revenues (19) and apply $\ln(\mu^a_j + g^a_j) - \ln(\mu^o_j + g^o_j) = \ln \overline{p^a_j} - \ln \overline{p^o_j} + \ln f^a_j - \ln f^o_j$,

where $f^m_j = 1 + \dfrac{g^m_j}{\overline{p}^m_j}$, similar as in equation (9).

Figure 1

Profits

fixed
costs

$\ln \pi^a$

$\ln \pi^o$

z

$\tilde{z}$ $\overline{z}$
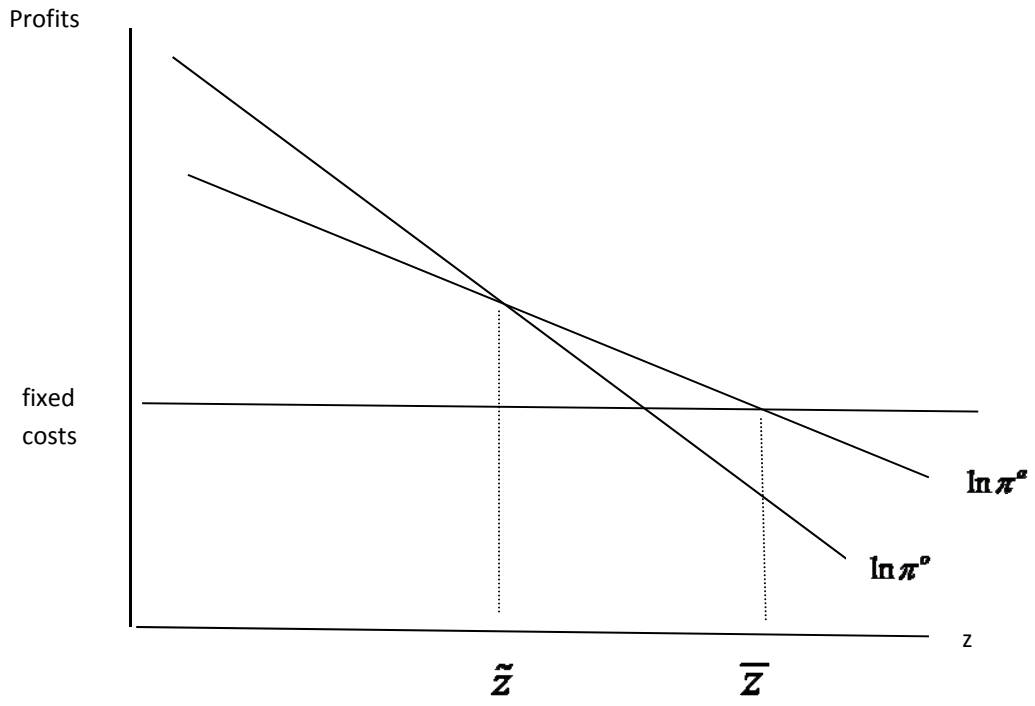
Figure 2: Trends in Air-Value Shares
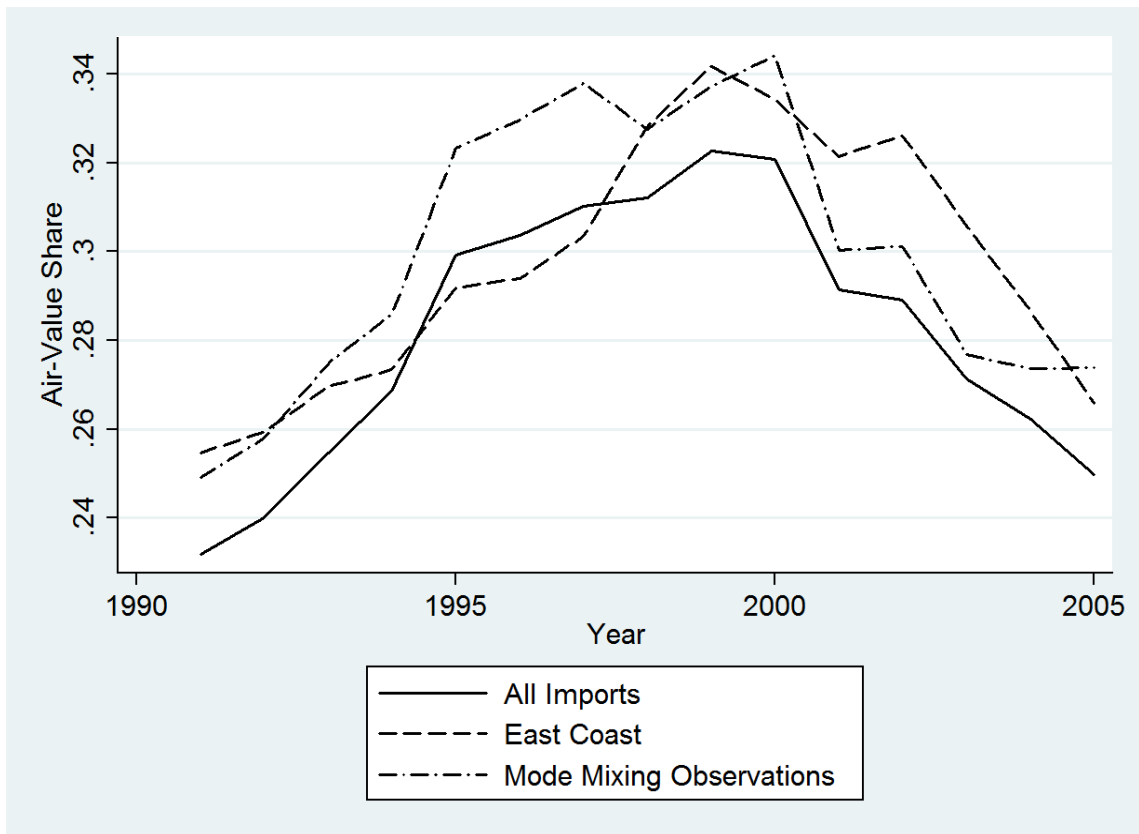
Figure 3: Distribution of Air Freight Premia
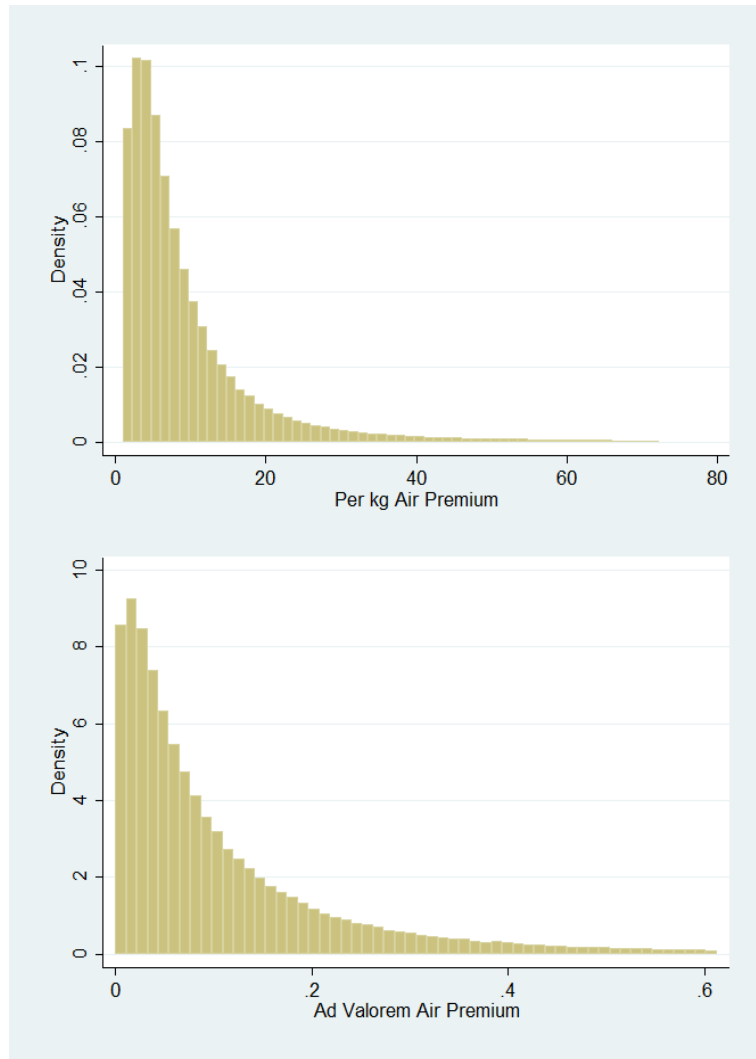
Unit of observation: HS6×Coast×Year. Air Premium Value $=f^a - f^o$ =(1+air charge/air value)-(1+vessel charge/vessel value). Air Premium Weight $=g^a/g^o$ =(air charge/air weight)/(vessel charge/vessel weight). Both distribtuions drop the 99th percentile of the air premia. The bottom figure drops negative air permia.

Figure 4: Transit Times
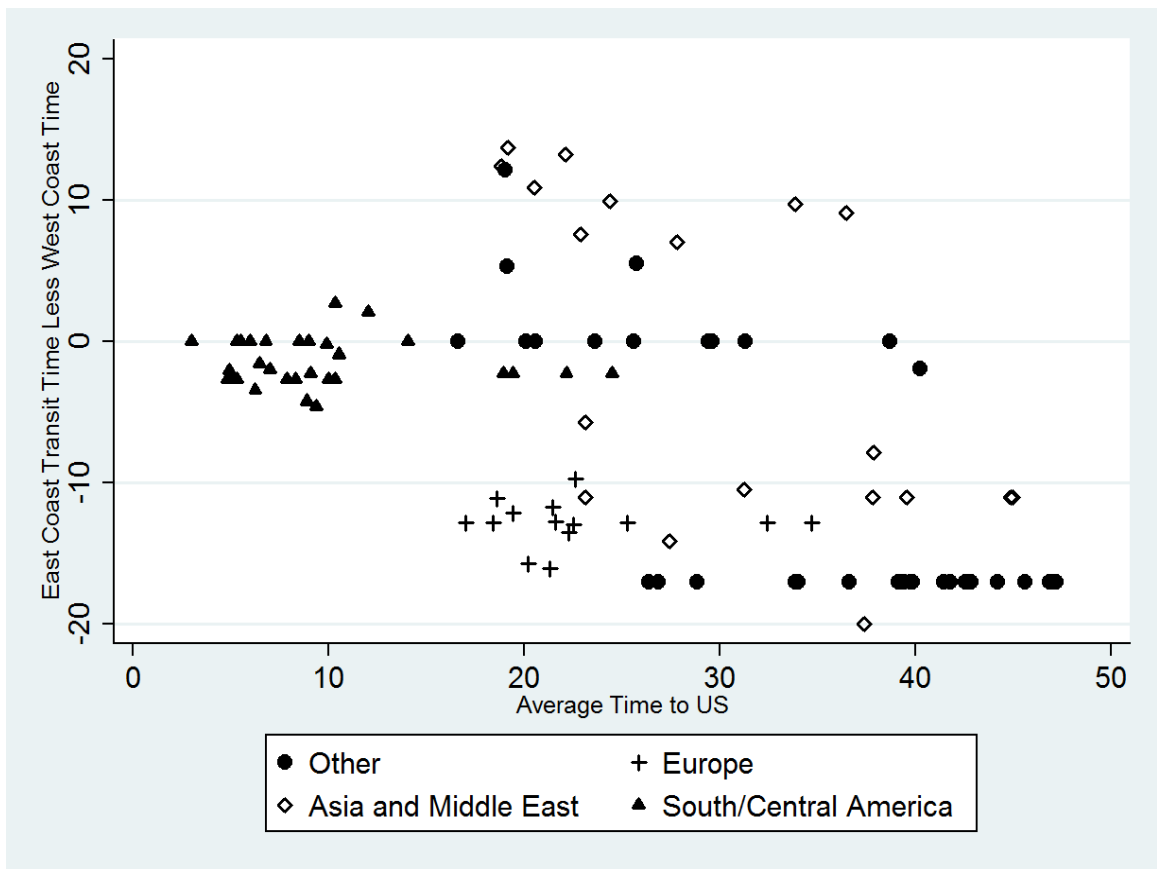
Figure 5: Distribution of Tau Estimates by 5 digit End-Use Category
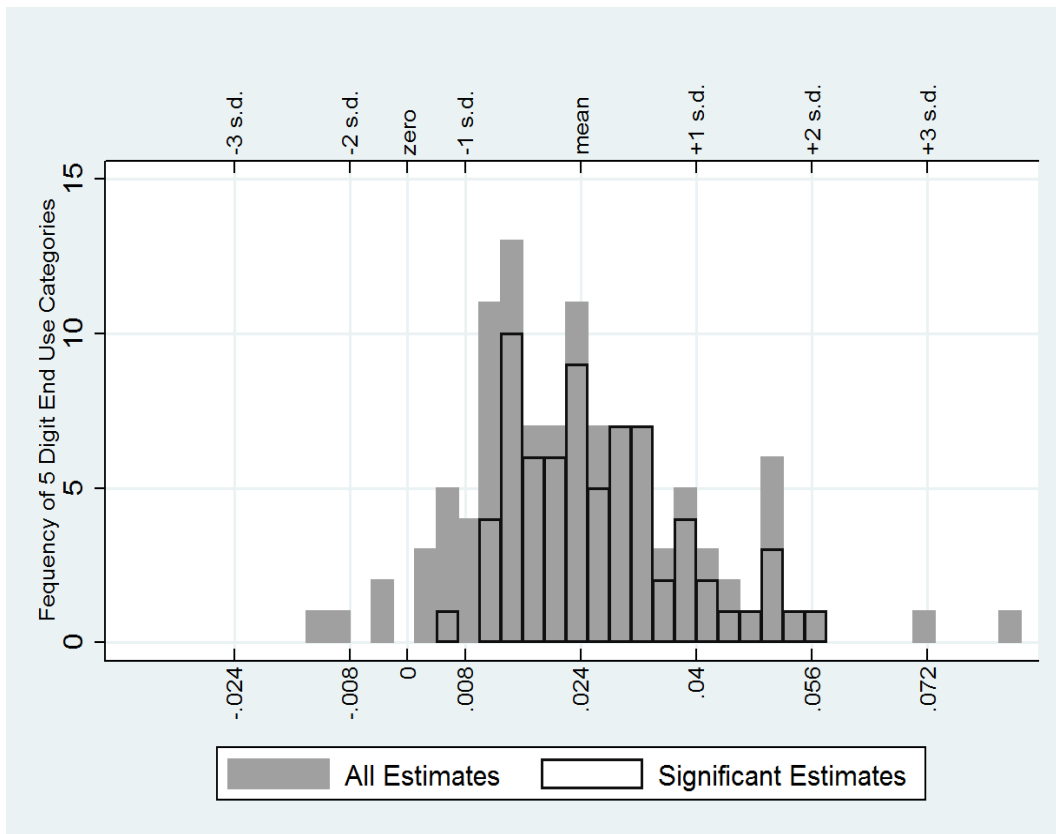
Note: Time costs are estimated for 110 5-digit end-use categories; 70 of these are significantly different from zero at the 10 percent level. Model: equation (12). Fixed Effects: HS6×Exporter.

Figure 6: Distribution of Tau Estimates by 5 digit End-Use Category

Note: Time costs are estimated for 110 5-digit end-use categories; 38 of these are significantly different from zero at the 10 percent level. Model: equation (13). Fixed Effects: HS6×Exporter.

Table 1: Variation in Air and Trade Shares across Commodity Groups and Regions

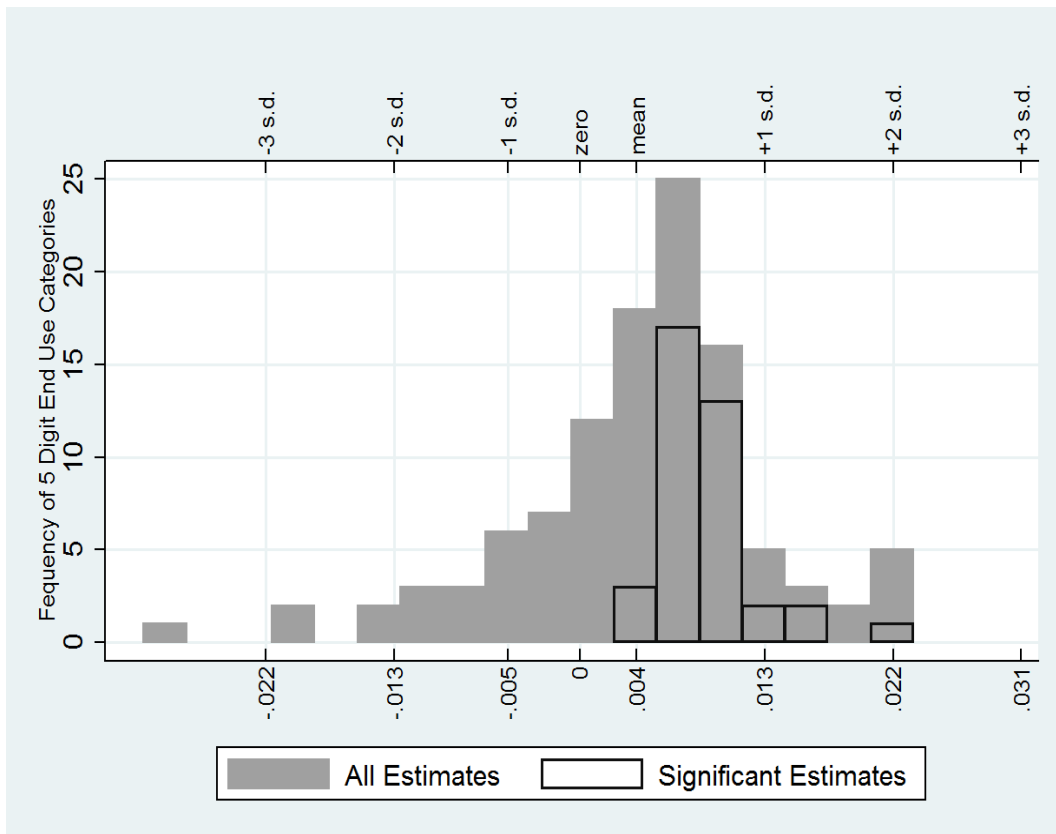| Group	All Observations		Mode Mixing Observations		Median Air Premium	
	Group Share of Imports	Air Share	Share of Group Imports	Air Share	Value	Weight
All Imports:						
	1.00	.28	.75	.30	.05	6.46
Region:						
Central America	.03	.13	.60	.14	.03	3.25
South America	.06	.12	.33	.21	.06	4.91
Europe	.28	.39	.73	.38	.03	5.88
Asia	.59	.27	.85	.28	.08	7.73
Australia/Oceania	.01	.18	.48	.23	.04	6.00
Africa	.04	.10	.15	.11	.08	5.42
End Use Categories:						
Food	.04	.04	.50	.06	.14	7.11
Industrial Supplies	.23	.09	.36	.14	.07	8.90
Capital Goods	.28	.52	.89	.52	.02	6.77
Automotive	.12	.02	.92	.02	.06	7.58
Consumer Goods	.30	.31	.86	.24	.06	5.20
Other	.02	.80	.95	.80	-.01	7.70
Product Group:						
Components	.12	.41	.94	.39	.04	6.45
Fresh	.01	.23	.49	.23	.17	5.55

Air Premium Value $= f^a - f^o =$ (1+air charge/air value)-(1+vessel charge/vessel value). Air Premium Weight $= g^a/g^o =$ (air charge/air weight)/(vessel charge/vessel weight). A mode mixing observation is a HS6$\times$Exporter$\times$Year$\times$Coast observation that shows positive air and ocean values.

Table 2: Summary Statistics

Group	Log Revenue		Log Revenue per Shipment		Log Price		Log Freight Cost		Transit Time	
	Mean	Sdv.	Mean	Sdv.	Mean	Sdv.	Mean	Sdv.	Mean	Sdv.
All	-0.97	2.08	-0.73	1.07	1.00	1.09	0.08	0.12	19.06	8.30
Foods	-1.76	1.95	-0.75	0.95	0.83	1.05	0.16	0.17	17.00	8.51
Industrial Supplies	-1.31	1.99	-0.9	1.10	1.10	1.19	0.11	0.14	17.92	7.31
Capital Goods	-0.59	2.09	-0.69	1.11	1.23	1.05	0.05	0.08	18.48	7.06
Automotive Vehicles/Parts	-1.50	1.86	-0.45	0.97	0.90	1.00	0.09	0.12	19.19	7.09
Consumer Goods	-0.91	2.08	-0.64	1.00	0.76	0.96	0.08	0.11	20.29	9.35

All variables are in relative terms (air to ocean).

Table 3: Revenue Specification

	(1)	(2)	(3)	(4)	(5)
Log Rel. Price	-.078	-.074	.027	.009	.067
	(.027)***	(.020)***	(.011)**	(.009)	(.014)***
Log Rel. Freight Cost	-6.460	-5.823	-3.346	-2.673	-3.301
	(.355)***	(.299)***	(.136)***	(.113)***	(.196)***
Transit Days	.018	.045	.049	.060	.069
	(.008)**	(.010)***	(.014)***	(.017)***	(.018)***
Tau	.003	.008	.015	.022	.021
	(.001)**	(.002)***	(.004)***	(.007)***	(.006)***
Fixed Effects	None	Exporter	Exporter +HS6	Exporter ×HS6	Coast Differenced
Obs.	528977	528976	528721	513424	244530
R^2	.121	.157	.356	.571	.159

Estimation of equation (12). Dependent Variable: log(air revenue/ocean revenue). Standard errors are robust and clustered by exporter. Regressions include a constant.

Table 4: Revenue per Shipment Specification

	(1)	(2)	(3)	(4)	(5)
Log Rel. Price	.037	.039	.038	.046	.070
	(.005)***	(.004)***	(.006)***	(.006)***	(.007)***
Log Rel. Freight Cost	-1.861	-1.900	-1.584	-1.509	-1.554
	(.094)***	(.081)***	(.077)***	(.075)***	(.095)***
Transit Days	.008	.011	.008	.009	.010
	(.002)***	(.002)***	(.002)***	(.002)***	(.002)***
Tau	.004	.006	.005	.006	.006
	(.0008)***	(.001)***	(.001)***	(.001)***	(.001)***
Fixed Effects	None	Exporter	Exporter +HS6	Exporter ×HS6	Coast Differenced
Obs.	528977	528976	528721	513424	244530
R^2	.049	.057	.144	.351	.041

Estimation of equation (13). Dependent Variable: log(air revenue per shipment/ocean revenue per shipment). Standard errors are robust and clustered by exporter. Regressions include a constant.

Table 5: Controlling for Selection

Dependent Variable:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Revenue				Revenue per Shipment			
Main Specification								
Log Rel. Price	-.088	-.091	.017	.007	.035	.036	.037	.044
	(.003)***	(.003)***	(.003)***	(.003)**	(.001)***	(.002)***	(.002)***	(.002)***
Log Rel. Freight Cost	-6.506	-5.886	-3.390	-2.684	-1.862	-1.912	-1.586	-1.510
	(.026)***	(.026)***	(.026)***	(.027)***	(.014)***	(.014)***	(.015)***	(.017)***
Transit Days	.018	.046	.051	.061	.008	.011	.009	.009
	(.0003)***	(.0005)***	(.0004)***	(.0004)***	(.0002)***	(.0003)***	(.0003)***	(.0002)***
Mills Ratio	.139	.247	.538	.305	.023	.065	.037	-.005
	(.006)***	(.007)***	(.009)***	(.019)***	(.003)***	(.004)***	(.005)***	(.012)
Selection Specification								
Log Average Days to U.S.	-.134	-.134	-.135	-.125	-.134	-.134	-.135	-.125
	(.004)***	(.004)***	(.004)***	(.004)***	(.004)***	(.004)***	(.004)***	(.004)***
Log Export Value to ROW	.425	.425	.428	.442	.425	.425	.428	.442
	(.0007)***	(.0007)***	(.004)***	(.004)***	(.0007)***	(.0007)***	(.004)***	(.004)***
Tau	.003	.008	.015	.023	.004	.006	.005	.006
	(.00005)***	(.00009)***	(.0002)***	(.0003)***	(.0001)***	(.0001)***	(.0002)***	(.0002)***
Fixed Effects:	None	Exporter	Exporter +HS6	Exporter ×HS6	None	Exporter	Exporter +HS6	Exporter ×HS6
Obs. Censored	736383	736383	721948	704747	736383	736383	721948	704747
Not Censored	517950	517950	517700	502631	517950	517950	517700	502631

Estimation of equations (12) and (13) with Heckman selection correction. The estimates in columns 1, 2, 5, 6 were obtained from Stata's 2-Step Heckman procedure which does not allow for clustering or robust standard errors. For columns 3, 4, 6 and 8 we implemented the 2 stage procedure by first estimating the selection equation and then including the inverse Mill's ration in the dummy variable regression implemented using Stata's areg command. Both the selection specification and main specification include a constant.

	(1)	(2)	(3)	(4)	(5)	(6)
	\multicolumn{6}{c}{Table 6: Robustness Checks}					
Dep. Variable in Log:	Revenue	Revenue per Ship.	Revenue	Revenue per Ship.	Revenue	Revenue per Ship.
Robustness Exercise	Squared Transit Days		Full Sample		Freight Rate Endogeneity	
Log Rel. Price	.066 (.013)***	.070 (.007)***	.040 (.013)***	.064 (.007)***	-.042 (.034)	.037 (.020)*
Log Rel. Freight Cost	-3.147 (.233)***	-1.532 (.094)***	-2.488 (.186)***	-1.242 (.091)***	-5.796 (.828)***	-2.763 (.480)***
Transit Days	.225 (.038)***	.032 (.005)***	.077 (.020)***	.012 (.002)***	.028 (.0006)***	.007 (.0003)***
Transit Days Squared	-.003 (.0006)***	-.0005 (.00008)***				
Lag Dep. Var.					.648 (.004)***	.383 (.004)***
Lag Freight Cost					1.446 (.189)***	.776 (.115)***
Tau	.026 (.006)***	.007 (.0009)***	.031 (.010)***	.010 (.002)***		
Fixed Effects	Coast Differenced		Coast Differenced		Coast Differenced, IV	
Obs.	244530	244530	321744	321744	110754	110754
R^2	.21	.045	.153	.045		

Standard errors in columns 1-4 are robust and clustered by exporter. The panel IV in columns 5 and 6 was implemented using Stata's xtivreg command which does not accomodate robust or clustered standard errors. The first stage R^2 for columns 5 and 6 are 0.43 and 0.42. In both IV regressions the instrument for the contemporaneous log relative freight rate is the second lag of the log relative freight rate. In both first stage regressions the instrument significantly predicts the endogenous variable and the F test rejects the null hypothesis that the instruments don't have a joint impact on the endogenous variable. Regressions include a constant.

Table 7: Time Costs by End-Use Group

End Use:	(1) Foods Beverages	(2) Indust. Supplies	(3) Capital Goods	(4) Auto-motive	(5) Consumer Goods	(6) Foods Beverages	(7) Indust. Supplies	(8) Capital Goods	(9) Auto-motive	(10) Consumer Goods
	Log Revenues					Log Revenues per Shipment				
Log Rel. Price	-.026 (.034)	-.087 (.011)***	.072 (.016)***	.021 (.030)	.016 (.013)	.047 (.028)*	-.042 (.007)***	.117 (.010)***	.077 (.014)***	.046 (.010)***
Log Rel. Freight Cost	-1.522 (.208)***	-2.539 (.125)***	-3.132 (.319)***	-1.643 (.241)***	-2.969 (.126)***	-.784 (.103)***	-1.647 (.079)***	-1.463 (.188)***	-.897 (.094)***	-1.619 (.094)***
Transit Days	.048 (.010)***	.062 (.013)***	.063 (.014)***	.071 (.015)***	.058 (.023)**	-.004 (.003)	.007 (.003)**	.013 (.004)***	.012 (.002)***	.007 (.003)**
Tau	.031 (.008)***	.024 (.006)***	.020 (.006)***	.043 (.011)***	.019 (.008)**	-.005 (.004)	.004 (.002)*	.009 (.002)***	.013 (.003)***	.004 (.002)**
Fixed Effects:										
Exporter × HS6										
Obs.	12065	143860	138079	18527	211155	12065	143860	138079	18527	211155
R^2	.61	.552	.568	.498	.562	.452	.365	.348	.354	.324

Estimation of equations (12) and (13) by product subsamples. Standard errors are robust and clustered by exporter. Regressions include a constant.

Table 8: Time Costs by Product Characteristics

	(1)	(2)	(3)	(4)
Dependent Variable:	Log Revenue		Log Revenue per Shipment	
Log Rel. Price	.009	.009	.046	.046
	(.009)	(.009)	(.006)***	(.006)***
Log Rel. Freight Cost	-2.676	-2.676	-1.510	-1.510
	(.114)***	(.114)***	(.075)***	(.075)***
Transit Days	.060	.060	.008	.009
	(.017)***	(.017)***	(.002)***	(.002)***
Component Share	.122		-.054	
	(.094)		(.036)	
Component Share× Days	-.002		.004	
	(.005)		(.002)***	
Fresh Share		.697		.306
		(.391)*		(.137)**
Fresh Share × Days		-.037		-.010
		(.018)**		(.007)
Tau (Share=0)	.023	.023	.005	.006
	(.007)***	(.007)***	(.001)***	(.001)***
Tau (Share=1)	.022	.009	.008	-.0006
	(.006)***	(.005)	(.001)***	(.004)
Fixed Effects:	Exporter × HS6			
Obs.	512012	512012	512012	512012
R^2	.571	.571	.352	.352

Estimation of equations (12) and (13) by product characteristics. Standard errors are robust and clustered by exporter. Regressions include a constant.