A1. Related Literature

Our paper relates to literatures on demand uncertainty, on the use of air cargo, on time delays related to port infrastructure and customs procedures, and to the importance of quality in trade.

Evans and Harrigan (2005), Aizenman (2004), and Hummels and Schaur (2010) emphasize the interaction between timeliness and demand uncertainty.   Aizenman (2004), and Hummels and Schaur (2010) focus on models in which there is uncertainty in the quantity demanded in foreign markets. Firms who ocean ship must commit to quantities supplied to the market before uncertain demand is resolved. This creates the possibility of forecast errors that result in lost profitability as firms over- or under-supply the market. Firms who air ship can wait until demand shocks are realized, but must pay a premium to do so. Aizenman (2004) emphasizes the implications of this model for pricing to market, since exchange rate shocks to demand may come after the factory gate prices of imports are set. Hummels and Schaur (2010) emphasize the implications for modal choice, equilibrium effects on price volatility, and the value of airplanes as a real option to smooth demand shocks. Evans and Harrigan (2005) emphasize the need for firms to adapt their products to reflect changing consumer tastes. They focus on location decisions and empirical evidence on lean retailing in the apparel industry. They show

that the higher is the restocking rate for a particular clothing item the greater is the reliance on proximate import sources.

Previous papers have examined the rise in air cargo 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.

A few papers have examined whether time delays imposed during customs procedures or while passing through ocean ports affect quantities of 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.

Previous papers have examined the importance of quality differentiation in trade. Empirically, quality is measured either as price variation (examples include Schott (2004), Hallak (2006), Choi et al (2009), and Baldwin and Harrigan (2011) among others) or as a residual of quantity demanded controlling for prices (examples include Hummels-Klenow (2005), Hallak-Schott(2011), Khandelwal (2010)) 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 our most robust strategy, we control for quality that is specific to an exporter-product-time trade flow and exploit variation across entry points into the US. 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, as in Feenstra (1994) and Broda-Weinstein (2006).

A.2 Theory

A.2.1   Modal Choice and Selection

Equations (3) and (4) in the text describe the problem facing a single firm, including the choice of mode and a selection into exporting equation. These equations can be represented graphically in Figure A1, 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.

A.2.2 Deriving Equation (6)

Starting from equation (5) in the text, which gives revenue for firms using ocean cargo,

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

we write the ratio of air relative to ocean revenues and take logs. All variables that are not mode-specific drop from the epxression.

$$\ln \frac{r(z_j^a)*}{r(z_j^o)*} = \sigma\tau(days_j^o - 1) + (1-\sigma)\left[ln\left(z_j^a + g_j^a\right) - ln\left(z_j^o + g_j^o\right) \right] + \sigma \ln\left(\frac{v_j^a}{v_j^o} \right) \qquad (A1)$$

Next, rewrite the difference in delivered 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 (A1) and transform it into the difference in delivered prices so long as the markup on the origin price is independent of delivery mode.[1]

---

[1] 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

$$ln\left(z_j^a + g_j^a\right) - ln\left(z_j^o + g_j^o\right) = \ln\frac{1}{\theta} + \ln(z_j^a + g_j^a) - \ln\frac{1}{\theta} - \ln(z_j^o + g_j^o) = \ln p_j^a * - \ln p_j^o * \qquad \text{(A2)}$$

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

$$\ln p_j^m* = \ln\left(p_j^m\left(1 + \frac{g_j^m}{p_j^m}\right)\right) = \ln p_j^m + \ln f_j^m \quad , \quad f_j^m = \left(1 + \frac{g_j^m}{p_j^m}\right) > 1 \qquad \text{(A3)}$$

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 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 (A2) and (A3) in (A1), we can express the difference in marginal costs as the difference in origin prices plus differences in ad-valorem shipping charges.

$$\ln(z_j^a + g_j^a) - \ln(z_j^o + g_j^o) = \ln p_j^a - \ln p_j^o + \ln f_j^a - \ln f_j^o \qquad \text{(A4)}$$

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 (A1) as revenues exclusive of these charges by subtracting $\ln(f_j^a / f_j^o)$ from both sides.

A.2.3: Relative Revenues in a Model with Firm Heterogeneity

Equation (7) describes relative revenues when there are N_j^m identical firms for each mode. In this section we describe how to derive this equation in a model with heterogeneous firms. Let z be the constant marginal cost of production and assume that firms draw their productivity from some probability distribution as in Melitz (2003). From equation (5) in the text, the sales of a firm located in country j are

$$r(z)_j = p(z)_j q(z)_j = E\left(\frac{1}{v_j^m \exp(-\tau \cdot days_{rj}^m)}\right)^{-\sigma}\left(\frac{\sigma}{\sigma-1}(z + g_j^m)\right)^{1-\sigma} \qquad \text{(A5)}$$

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

---

expansion that yields the same form used here, which is to say that the feedback effects onto markups are of second order importance.

the profits at firm realizes employing alternative modes of transportation as illustrated in Figure A1, 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

$$\int_L^{z_j^o} p_j(z)q_j(z)M_j w(z)dz = \left(\frac{\sigma}{\sigma-1}\right)^{1-\sigma} \frac{E}{\left(v_j^o \exp(-\tau \cdot d_j)\right)^{-\sigma}} M_j \qquad (A6)$$

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

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

$$R_j^o = \int_L^{z_j^o} p_j(z)q_j(z)w(z)dz = \left(\frac{\sigma}{\sigma-1}\right)^{1-\sigma} E\left(v_j^o \exp(-\tau \cdot d_j)\right)^{\sigma} (\mu_j^o + g_j^o)^{1-\sigma} N_j^o, \qquad (A7)$$

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

$$R_j^a = \left(\frac{\sigma}{\sigma-1}\right)^{1-\sigma} E\left(v_j^a \exp(-\tau)\right)^{\sigma} (\mu_j^a + g_j^a)^{1-\sigma} N_j^a \qquad (A8)$$

where $N_j^a = M_j W(z_j^o < z < z_j^e)$ is the mass of air shippers. Equations (A7) and (A8) 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 applied in equation (9) 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.

A.3.1 Sample construction and Data Description

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.

We begin with roughly 45 million trade observations that arrive in the US by air or ocean. We drop inland customs districts, which 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. We also drop Puerto Rico, Hawaii, and the Virgin Islands, which together with inland districts account 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 (8) and (9) 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. 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.

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 1[st] percentile and above the 99[th] 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. Some HS codes have air

shares very close to zero or very close to 1. This suggests that one mode is used almost exclusively, and the outlying observations may be unusual situations or data errors.

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. 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.

Our data on shipping times are derived from a master 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 the text we refer to several data displays available in the online appendix.

Table A1 provides summary statistics for our sample on variables used in the estimation.

Figure A2 displays a histogram for air freight premia, measured on a per kilogram and on an ad-valorem basis. An observation corresponds to an exporter j-HS6 product k- coast c-time t value. As we note in the text, while the median observation has an ad-valorem air freight premia of 5 percent, many observations in the data have much higher values. At the 90[th] percentile, air freight premia are 34 percent on an ad-valorem basis. On a per kilogram basis, at the 90the percentile air freight costs are 27 times higher than ocean freight.

A.4. Results and Robustness checks

A.4.1 Selection

In the text we discuss using a Heckman selection equation. High costs and long shipping times could cause a country to have zero exports to the US in a particular product. 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 2 and 3. (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.)

Table A2 reports estimates of the first and second stages of the Heckman estimator. 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. 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 2 and 3. 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.

A.4.2 Heterogeneity in coefficients across products

The one-digit End Use categories reported in text Table 5 are still fairly broad and we next group products at the most disaggregated End-Use Category and re-estimate equations (8) and (9) separately for each, using Exporter x HS6 fixed effects. Figures A3 and A4 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. (The histograms omit insignificant point estimates lying 2 standard deviations from the mean.)   For Figure A3, the mean over the individual group estimates shows an average time sensitivity of about 0.02, which is very similar to Table 2, column 4. 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 A4, we again see an average effect similar to Table 3, column 4, and again see considerable dispersion.

A.4.3 Sourcing from Local Markets

In the text we describe a robustness check designed to identify whether highly time sensitive goods are selected out of our sample.   For a given product and time period we calculate the North American (i.e. Mexico plus Canada) share of imports. We include this in our base specifications, both in levels and interacted with days. The text describes the main findings and implications. The full results of this specification are found in appendix Tables A3 and A4.

A.4.4 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 (A2), (A3), and (A4) 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

$$(1-\sigma)\left(\ln f_j^A - \ln f_j^O\right) - \sigma\tau\left(days_j^O - 1\right) > 0 \tag{A9}$$

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

$$P\left(m_{jt} = a \mid x\right) = \Phi\left[\delta + (1-\sigma)\left(\ln(f_j^a) - \ln(f_j^o)\right) - \sigma\tau\left(days_j^o - 1\right)\right] \tag{A10}$$

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 freight rates for both modes 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-year so that the intercepts and coefficients are mode-HS4-year specific

$$ln\left(g_{jckt}^m\right) = \beta_1^{m,hs4,t} + \beta_2^{m,hs4,t} ln\left(p_{jckt}^m\right) + \beta_3^{m,hs4,t} ln\left(distance_{cj}^m\right) + u_{jckt} \tag{A11}$$

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:

Quality differences enter as demand shifters in equation (9). 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 (A10) 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.

.

$$P\left(m_{jkct} = a \mid x\right) = \Phi\left[-0.543\left(\ln(f_{jkct}^a) - \ln(f_{jkct}^o)\right) + 0.007\left(days_{jc}^o - 1\right) + 0.726\ln(p_{jkct})\right]$$
$$\quad\quad\quad\quad (.304)* \quad\quad\quad\quad\quad\quad (.008) \quad\quad\quad\quad (.036)***$$

High relative air shipping costs and low product prices reduce the probability of using air shipping, and both effects are 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 additional noise into the regression.

A.5 Appendix References

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

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.

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.

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.

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.

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

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

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

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