PEDL Research Papers

This research was partly or entirely supported by funding from the research initiative Private Enterprise Development in Low-Income Countries (PEDL), a Foreign, Commonwealth & Development Office (FCDO) funded programme run by the Centre for Economic Policy Research (CEPR).

PEDL Twitter

The value of face-to-face: Search and contracting problems in Nigerian trade*

Meredith Startz†

March 2021

Abstract

Distance between buyers and sellers can create search and contracting problems: how to find out what goods are available and ensure they are actually delivered. I estimate the magnitude of these frictions faced by Nigerian importers by exploiting the fact that travelling to do business in person is a common and observable strategy for coping with them. Using original data that includes the process of trade, I build and estimate a model that embeds search and moral hazard into a trade framework, in which buyers may choose to travel to find the latest goods and conduct spot transactions. Welfare from imported consumer goods increases by 14% in the absence of both frictions, with one-third of gains coming from search. Counterfactual scenarios focused on airline regulation and financial services suggest that attention to such topics not typically considered market integration policies could have large welfare effects via trade in developing countries.

*Thanks to David Atkin, Dean Karlan, Sam Kortum, Mark Rosenzweig, and Chris Udry for their guidance and support, as well as Treb Allen, Simon Board, Aaron Flaaen, Sharat Ganapati, Bob Gibbons, Penny Goldberg, Matthew Grant, Shelby Grossman, Dan Keniston, Shelly Lundberg, Giovanni Maggi, Chiara Margaria, Costas Meghir, Ameet Morjaria, Tommaso Porzio, Ana Reynoso, Nick Ryan, Gabriella Santangelo, Pete Schott, Dick Startz, Glynis Startz, Jeff Weaver, and Chris Woodruff, for helpful conversations and comments. This project was made possible by research assistance from Hakeem Bishi, Dylan Groves, Yemi Ogundare, Nicole Wilson, Anifat Ibrahim, Mayor Ajala, Lanre Aluko, and the Lagos Trader Project enumerator team. I gratefully acknowledge financial support from the Centre for Economic Policy Research, the Tokyo Foundation, the Kauffman Foundation, and the NSF Graduate Research Fellowship program.

†Dartmouth College. Email: meredith.l.startz@dartmouth.edu

1 Introduction

Search and contracting problems can arise when buyers and sellers are in different locations. Consider a shoe wholesaler in Lagos, Nigeria. If he sources shoes in Lagos, he can see what is available and exchange money for goods on the spot. But in China, an immense variety of shoes are produced at low cost. How can the Nigerian wholesaler find out what styles of shoes are currently available in Guangzhou? How can he be sure that a supplier will actually deliver the shoes after taking his money? He will face these problems each time he buys, because shoe styles change and sellers will still be opportunistic.

One solution is to travel to China to do business in person. Travel eliminates the distance between a buyer and seller, effectively transforming a remote transaction into the type that would be conducted at home. However, travel is not obligatory. A Lagos wholesaler can search for products using phones, internet, and word-of-mouth, and trade partners can be motivated to behave well through repeated interaction. The insight of this paper is that heterogeneity in the observable choice to travel or not can be used to quantify the value of solving distance-related information problems, and to identify the underlying mechanisms relating them to firm and product characteristics.

Distance between potential transaction partners creates both a search problem (the difficulty of finding out what goods are available) and a contracting problem (the difficulty of knowing whether goods or money have actually been sent), which I refer to together as "information costs".[1] This may help explain a persistent empirical puzzle: the estimated trade costs implied by flows of goods are consistently much higher than the observable components, such as transportation costs and tariffs (Anderson and van Wincoop (2004)). Information costs may help account for this gap. Without a clear understanding of the underlying sources of trade costs, it is difficult to design appropriate market integration policies. This is particularly important in developing countries, where trade costs are high and information costs may be especially relevant due to weak contract enforcement institutions, limited access to information technology, and small firm size. However, neither the empirical size of information costs in trade nor the underlying mechanisms are well-understood. What we do know is mostly confined to homogenous goods traded under perfect competition and per-

[1]Note that although I make use of data on international trade, the situation is analagous to transactions over distance within a country.

fect contract enforcement, a special case in which a search friction can be inferred from spatial price dispersion (Jensen (2007); Aker (2010); Allen (2014); Steinwender (2018)).

I show that search and contracting frictions in differentiated goods trade are large, and can have a substantial impact on welfare. My approach is to quantify unobserved information problems by focusing on an easily observable strategy that importers use to cope with those problems: traveling to buy in person. To do so, I use panel data I collected from Nigerian importers, which combines information on the type and value of goods traded with variables describing the actual process of firm-to-firm trade (e.g. travel and payment terms) on a transaction level. The data cover 620 importers of differentiated consumer goods, such as clothing and electronics, who were randomly sampled from a census of over 50,000 shops in commercial districts of Lagos. The data captures imports over two years, totaling almost four thousand purchases from over thirty source countries and over a thousand foreign suppliers.

I begin by documenting four empirical patterns that motivate the structure of a model relating travel and importing. First, travel is common but not universal, and is more likely when importing from countries that are cheap to reach and in sectors in which new products appear frequently. Second, travel expenditures are large—equivalent to amounts spent on transportation and regulatory trade costs combined. Third, travel by importers is persistent over time, and does not decline significantly in experience with particular countries or suppliers, suggesting motives beyond one-time matching or learning. Finally, transactions that involve travel look different from those that don't. For purchases conducted in-person, traders pay lower unit costs, buy less frequently in larger batches, and are more likely to buy new product varieties and to switch suppliers.

I build a model that accounts for these patterns by embedding optimal stocking, search, and a repeated game with moral hazard into a monopolistically competitive trade framework. Differentiated goods are produced in a foreign source market. Traders source these goods from suppliers, and resell them to consumers at home who value having up-to-date varieties. Traders make forward-looking choices about how frequently to restock, and when doing so, whether or not to pay a fixed cost to travel to the source country. Traveling allows importers to search more effectively for new vintages and avoid a contract enforcement problem by conducting a spot transaction. Ordering remotely has a lower fixed cost but yields less up-to-date prod-

ucts as a result of the search friction and incurs higher unit costs as a result of the contracting friction. I model the contracting problem as a repeated game of moral hazard with an endogenous period of time between stages. Consistent with my data, which show that post-payment and non-delivery are both rare, I focus on equilibria that feature full pre-payment and honest behavior along the equilibrium path. The solution is analagous to an efficiency wage: the trader pays a price premium to satisfy the supplier's incentive compatibility constraint, so that the supplier prefers to behave honestly and continue to earn the associated profits than to cheat once and never do business with the trader again.

Traders choose whether to travel or order and how frequently to restock based on underlying heterogeneity in how quickly the products they sell evolve over time, how popular they are, and idiosyncratic components of the trader's perceived travel cost. A key departure from most trade models is that the frequency of trade matters for welfare; it determines both the vintage of goods available and the unit cost paid by traders (because smaller, more frequent purchases reduce the temptation for suppliers to renege), and therefore mediates the effects of search and contracting problems on consumers. Barriers remain even when an importer has extensive experience with a particular source country or supplier. Compared to a world without information frictions, consumers face higher prices, less product variety, and (sometimes) less up-to-date goods.

Selection into travel reflects the value of solving both search and contracting problems. In order to separately identify the role of each channel, I make use of additional observables with distinct relationships to search and contracting: the probability of switching suppliers, the probability of finding new vintages, the frequency of purchases, and variable profits. The size of the contracting friction is pinned down by the probability of buying from the same supplier again, because this determines the effective discount rate on future interactions in the repeated game between buyers and sellers. In contrast, the search friction is related to the probability of finding new product styles and the difference in profits when traveling versus ordering. If finding new products is rare, but supplier switching is common, we infer all else equal that contracting motives are stronger relative to search in driving the decision to travel.

I estimate the model using the Nigerian data to uncover the size of search and contracting frictions, and allow the parameters to vary freely across 12 different source country-sector pairs. On average, importing without traveling yields goods that are

1.7 months behind the frontier available in the source country (the search friction) and requires paying a 15.5% price premium to induce good behavior from suppliers (the contracting friction). Removing both frictions increases welfare from the traded consumer goods sector by 14%—roughly two-thirds of the gains that would come from eliminating physical and regulatory trade costs in this sector. This estimate is large, but plausible considering that it represents the complete elimination of these barrier, and applies to traded goods in a sector that accounts for roughly 17% of consumer spending. The welfare gains from eliminating the search problem alone would be 4.5%, and the gains from contracting alone would be 7%—the whole is greater than the sum of the parts due to an interaction through the discrete choice to travel. These aggregate gains reflect sensible underlying variation across countries and sectors—for instance, consumer welfare losses due to information frictions are larger in the observed equilibrium when importing apparel from China relative to hardware from China (because apparel products change more often), and relative to apparel from Benin (because travel to Benin costs less than travel to China).

Information frictions also distort the set of importers and imported products. Importer profits are higher in a frictionless world, but average importer size is smaller because information frictions push the smallest firms out of the market. Information frictions reduce the set of varieties available to consumers, and the varieties that are lost are those that change rapidly or for which total demand is small. Counterintuitively, the effect of information problems on the vintage of goods available is ambiguous—the search problem always makes available goods more out-of-date, but this is in some cases offset by the fact that the contracting game induces traders to buy more frequently than is efficient in order to keep the unit cost premium down.

By fully specifying the relationship between information frictions and underlying product and firm characteristics, the model is useful for understanding when information costs are likely to pose a substantial barrier to trade in other contexts. When firms are large or the costs of travel or other coping strategies are low, information problems will have little effect on welfare. In the first of three counterfactual scenarios, I increase Nigerian consumer spending to match that in the United States, and show that when average revenue per firm is higher, the same fundamental search and contracting frictions yield much smaller welfare costs. Intuitively, fixed cost strategies for addressing information problems in trade are inconsequential relative to revenue for Walmart; not so for small traders in Lagos. The model also suggests that welfare

effects of information problems should be smaller for slow-changing types of goods, and in markets in which trade partners are able to establish longstanding relationships or reputational forces are strong, driving the solution to the repeated contracting game toward what would be achieved under perfect enforcement.

This evidence on information frictions directs attention to a range of trade facilitation policies and services beyond the classic infrastructure investment and tariff reduction. It suggests, for instance, that people's ability to move freely within and between countries is relevant to the efficient movement of goods. Many developing countries suffer from heavy visa restrictions and costly and poorly managed travel options. In a second counterfactual scenario, I show that if China and Nigeria liberalized air travel between the two countries, the fraction of traders who travel would increase, and Nigerian consumers would gain substantially. In a third counterfactual scenario, I investigate why existing financial services have not gone further in ameliorating the contracting problem. I model the option for traders to purchase an escrow service similar to that offered by Alibaba.com at the time of writing, and find that demand is low at actual prices and plausible expectations about the effectiveness of dispute resolution. This result is not limited to the specifics of the particular service – use of agents or financial services more broadly simply transfers the contracting problem onto another party; they will be used only when they offer access to some technology or relationship that allows the contracting problem to be solved at lower cost than the other options available (including travel).

The idea that information problems contribute to trade costs is an old one. Estimates of gravity models often include proxies for information barriers, such as shared language or the quality of legal institutions (summarized in Anderson and van Wincoop (2004)). Empirical evidence on the source of information costs comes from studies on incomplete information about the prices of homogenous goods, often in developing countries, which show that the introduction of better information technology (e.g. cell phones) reduces geographic price dispersion (e.g. fish in India in Jensen (2007), grain in Niger in Aker (2010), agricultural products in the Philippines in Allen (2014), cotton in the mid-19th century UK in Steinwender (2018)).

I go beyond this work by microfounding and quantifying both imperfect contract enforcement and search over products with differentiated characteristics. Rauch and Trindade (2002) find that trade flows are larger between countries with larger ethnic Chinese populations, particularly for differentiated goods. They argue ethnic

networks make matching easier (search) and deter opportunistic behavior through community sanctions (contracting), but do not model the mechanisms or empirically distinguish the two channels. A small related literature shows relationships between observable communication and aggregate trade flows (particularly for differentiated goods), including air travel in Cristea (2011) and Poole (2010), and telephone calls in Portes and Rey (2005). I observe communication and outcomes at the transaction level, allowing me to build a fully-specified model of firm behavior with underlying heterogeneity.

I contribute to a literature on the role of imperfect contract enforcement in firm relationships. While a large theoretical literature (summarized in MacLeod (2007)) suggests a variety of mechanisms, microempirical evidence is limited, and has mostly focused on learning about types and the evolution of transaction terms over time within relationships (Antras and Foley (2015); Macchiavello and Morjaria (2015); Macchiavello (2010); Banerjee and Duflo (2000); McMillan and Woodruff (1999)). I provide evidence that firms can opt out of this relational contracting mode, choosing to use costly monitoring strategies instead (as in the model of Taylor and Wiggins (1997)), and that contracting frictions can have persistent effects even in the context of long relationships. I also link microempirical evidence on the contracting mechanism to a model of trade that allows for an analysis of welfare and interactions with other frictions in equilibrium.

Finally, I provide new evidence on the role of intermediaries and middlemen. A theoretical literature speculates about what middlemen do, and most of the proposed mechanisms have to do with solving matching or contracting problems (Biglaiser (1993); Antras and Costinot (2011); Bardhan, Mookherjee and Tsumagari (2013); Krishna and Sheveleva (2016)). The limited empirical evidence available supports the idea that intermediaries' role involves economies of scale in paying fixed costs (Ahn, Khandelwal and Wei (2011); Crozet, Lalanne and Poncet (2013)), and work by Atkin and Donaldson (2015) shows that they are important in determining the actual prices received by consumers. I do not endogenize market structure or pass-through by traders in this paper, but provide empirical evidence on the production function of intermediaries.

2 Context and data

I study traders operating in Lagos, Nigeria, who import consumer goods from around the world. Nigeria is the 7th most populous country in the world and the largest economy in Africa. Lagos is the commercial capital of Nigeria, and the main port of entry (Nigerian Ports Authority, 2016). The manufactured consumer goods I focus on in this paper—such as apparel, personal electronics, furniture, toiletries and cosmetics, light hardware, and home appliances—account for approximately 17% of Nigerian consumer expenditure.[2] The wholesale and retail trading sector is the second largest contributor to Nigerian GDP (Leke et al. (2014)), after agriculture, which suggests that understanding of market failures in the trading sector is economically important not only due to its potential effects on consumers, but also on aggregate productivity.

2.1 Data: Lagos Trader Survey

The empirical evidence in this paper is based on original survey data collected from traders located in Lagos, Nigeria as part of the Lagos Trader Survey (LTS). Additional details on the data and data collection procedures are provided in Appendix A.

Traders were identified through a census of 52,830 shops in commercial and wholesaling areas of Lagos, conducted between October 2014 and April 2015.[3] Research assistants covered all markets and plazas on foot, enumerating individual shops by location and the type of products being sold. After excluding offices, services, warehouses, and vacant or inaccessible shops, a simple random sample was taken from the remaining 24,159 shops. Interviews were conducted with traders in the sampled shops between April and August 2015. The interview covered trader and firm characteristics, business history in all source countries, relationship history with specific suppliers in those countries, and all international purchases made in 2013 and 2014.

The LTS data combines the type of information found in detailed official customs data with a novel set of variables describing the process of firm-to-firm trade. At the most basic level, the data contains information that would be found in any trade statistics: the value and type of goods being imported from a given source country. It also contains more granular information, such as the dates of individual transactions,

[2]Euromonitor Passport database, consumer spending data for 2010 - 2015.

[3]The census did not include most residential or manufacturing areas. It also excludes traditional markets, which are mainly comprised of small retail vendors selling food and household items.

buyer and seller identifiers, the price at which goods are sold in Nigeria, and physical and regulatory trade costs including shipping and clearing the port in Lagos.[4] The survey also includes a unique set of variables that describe *how* transactions are conducted, such as whether the buyer traveled to the source country, whether the product was exactly the same as something that had been purchased previously (as opposed to a new product, style, model, etc), method of payment, fraction of payment made pre-shipment, mode of communication with suppliers, use of hired agents, and whether there were any delays or defects compared to expected product delivery.

The survey was designed to create a retrospective transaction-level panel, in which relationships can be traced out over time. In addition to information about firm characteristics at the time of the interview, the survey also asks about the firm's history, including when the business was started, start and end dates for all countries it has ever imported from, and whether the trader has ever imported or traveled abroad in the past. It also asks about the history with any suppliers they have purchased from in the past two years, including the year when the relationship was originally established. Transaction specific information (e.g. product type, quantity purchased) can be tied to shipment level information (e.g. shipping costs) and to supplier- and country-specific histories. The data covers 1,179 traders, of whom 620 imported in 2013-14. Among importers, there are 2,481 international shipments and 3,907 import purchases of individual products captured in the data. Traders import from 32 different source countries, and report transacting with 1,073 (not necessarily unique) suppliers.[5]

2.2 Context: Importing consumer goods in Lagos, Nigera

Traders are owner-operators of wholesale and retail trading firms. All analysis focuses on the 53% of traders who imported at least once in 2013 or 2014. Table 1 shows summary statistics describing trader demographics and the characteristics of

[4]The costs of clearing the port include tariffs, fees, and any bribes or other costs—in order to elicit more honest responses, I did not ask respondents to specify how much of the total is attributable to each, since "tipping" officials to facilitate processing and misreporting the contents of containers to reduce taxes are both common practices.

[5]In order to obtain information at this level of detail, several strategies were employed that introduce some limitations in the data. First, individual suppliers are uniquely identified within but not across respondents. Second, a small fraction of respondents who report a very large number of import shipments were asked about a randomly selected subset or a "typical" shipment. These observations are reweighted in all analysis. Details are provided in Appendix A

Table 1: Summary statistics about importing traders

	Mean	Median
Trader		
Male	74%	
Age	38.7	37
Schooling beyond secondary	54%	
Uses smartphone	88%	
Business		
Wholesaler	84%	
Business age	10.1	9
Number of shops	1.3	1
Number of workers (paid & unpaid)	0.98	1
Annual revenue from direct imports (USD)	$107,581	$27,113
Annual shop rent (USD)	$2,482	$1,819
Annual direct international shipments	2.87	2
Traveled for business	61%	

Notes: N=620 traders.

their importing businesses. The typical trader is male and middle-aged, with almost all having completed secondary school. Traders are at ease using modern information technology in their businesses: 88% have smartphones, and the most common methods of communication with overseas suppliers are phone calls and text messages.

Traders generally do not transform goods, instead sourcing ready-made products abroad and reselling them in Nigeria. Branding or customization of products is rare. Traders have average annual revenue of $107,581 from sales of directly imported stock, but in other respects, they look more typical of small owner-operated enterprises in developing countries—the median trader has one shop and one worker in addition to the owner (including paid employees, apprentices, and unpaid family workers). Most specialize in one category of products, such as women's shoes, home appliances, or cosmetics. The majority of sales are wholesale, although a significant minority of revenue is from retail sales.

Of those who imported directly in 2013-14, 61% traveled for business at least once during that period. The imports represented in the LTS are sourced from more than thirty different countries. Figure 1 shows the most common source countries. By far the most common source is China, but the United Arab Emirates (specifically, Dubai), Turkey, Hong Kong, Benin, India, the United States, and the United Kingdom are

Figure 1: Most common import source countries

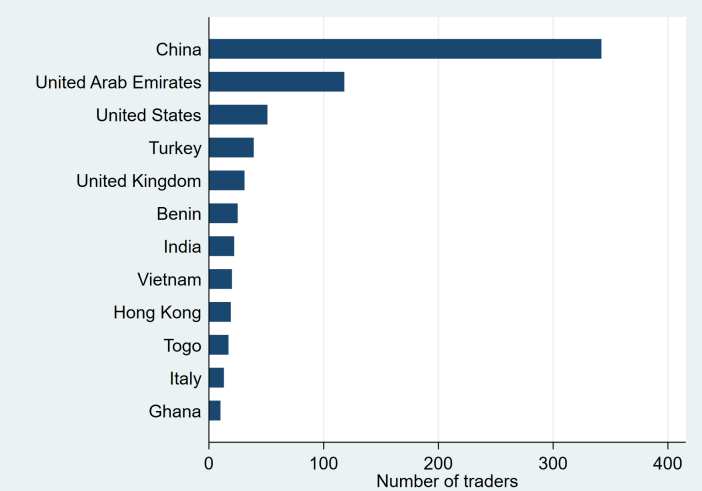

also common.[6]

For the purposes of analysis, I group products into six categories: apparel (which includes shoes, bags, and textiles), electronics, beauty, hardware, home goods, and miscellanous other products. Table 2 gives examples of individual products in each category. The largest number of traders deal in apparel, followed by electronics. China is the only source country with large import flows across all six categories – other countries provide goods mostly in one or two categories. Traders deal in goods they perceive as differentiated and evolving over time. Column 3 shows the fraction of traders who report that the style or model of products available in their line of business change at a quarterly or higher frequency, and column 4 shows their self-assessed importance of having the most recent product or styles in stock, on a scale from 1 (least important) to 10. The patterns across sectors are consistent with general intuition – apparel changes the most frequently, followed by personal electronics, and hardware changes the least frequently.

Traders' own explanations for why they travel highlight search and contracting problems consistently and explicitly. Survey responses to this question almost universally

[6]A number of source countries appear to serve as entrepôts—for instance, clothing from Benin and electronics from Dubai are unlikely to have been originally manufactured in those locations. Indeed, 70% of traders' suppliers are wholesalers rather than manufacturers. I do not make any distinctions based on the type of supplier, but it is worth noting that these transactions may be only one step in a supply chain featuring similar information frictions at other stages, both up- and downstream. This idea is explored further in Grant and Startz (2020).

Table 2: Summary statistics by product sector

Sector	Example products	Number of traders	New styles available at least quarterly	Important to have recent styles (1 - 10 most)
Apparel	Polo shirt, sandals, French lace	325	83%	8.64
Electronics	iPad, memory card, phone charger	77	77%	8.28
Beauty	Lipstick, hair relaxer, earring	62	75%	7.53
Hardware	Brake pad, electric cable, halogen bulb	58	34%	6.72
Homewares	Gas cooker, home theater, dining chair	59	59%	7.62
Miscellanous	Diaper, umbrella, tennis ball	38	53%	7.55

Notes: N=619 traders.

focused on these issues, with comments such as "Designs change so one needs to know what the market is selling" and "It is very risky to order products from abroad because of [...] dubious suppliers". Some comments on the motivation to travel also emphasize the necessity to achieve a certain scale to make it worthwhile: "I could not gather enough money to buy [a] large quantity of goods and it is better to order for small quantities [...] so that I can make profit".

3 Empirical patterns in trade and travel

In this section, I describe four facts about the empirical relationship between importing and travel, which suggest quantitatively important information frictions and guide the structure of the model.

Fact 1 - Travel is common, but varies by source country and product

Traveling to buy goods is a common but far from universal strategy. Taking a trader-source country pair as an observation, 34% of traders always traveled when importing from a particular country in 2013-14, 53% always ordered remotely, and the remaining 13% did some of both.

Panel A of Figure 2 shows, however, that there is substantial heterogeneity across source countries. The probability of traveling, conditional on importing from a given source country, is decreasing in the cost of travel. Traders almost always travel when importing from easy to reach neighbors in West Africa, but are much less likely to do

so when buying from countries that require expensive plane tickets and visas to visit, such as China or the United States.

Figure 2: Likelihood of traveling by source country and sector

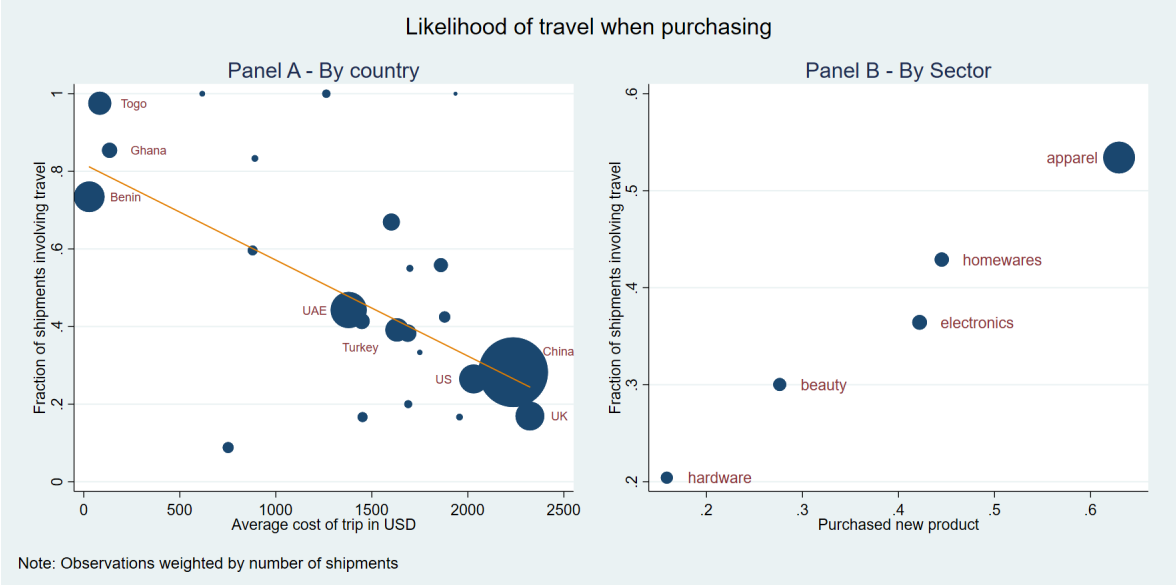

The probability of traveling also varies across products categories, and is positively correlated with the probability of finding new products when restocking. Panel B of Figure 2 shows that travel and purchasing new products (defined as a product that is not exactly the same as something purchased before, in terms of style, model, etc) are by far the highest for apparel and the lowest for hardware.

Fact 2 - Travel costs are large, and comparable in size to physical and regulatory trade costs combined

Figure 3 shows the average costs per shipment of goods imported from China[7], broken out by whether the trader traveled or ordered. The average ordered shipment contains $18,240 worth of goods, and involves $1,468 in transportation and regulatory costs.

The average shipment involving travel was larger, at $26,912 worth of goods and $1,936 in transport and regulatory costs.[8] Travel expenditures averaged $2,238 per

[7]China is taken as an example, as the largest single source country. Similar patterns hold on different scales for other source countries, and overall as shown in Table 4.

[8]The ad valorem rates of transport and regulatory costs are not significantly different between shipments involving travel and those that don't, suggesting that smuggling or carrying goods home are not major motivations for travel. Traders carried any goods home in their luggage on 28% of trips, and the average value of goods carried conditional on doing so was less than one-sixth of the total value.

trip. This number includes the cost of plane tickets, visas, and other reported travel expenses, such as hotels, translation services, and so on. It does not, of course, include indirect costs of travel, such as the opportunity cost of time (many traders shut down their shops and forego sales while traveling) and any (dis)utility of travel. However, explicit travel costs alone are still revealing. First, travel is sufficiently costly relative to the value of goods purchased to be of quantitative importance. Second, travel costs are on the same order of magnitude as the transportation and regulatory costs.

Figure 3: Travel, transport, and tariff expenditures relative to goods value

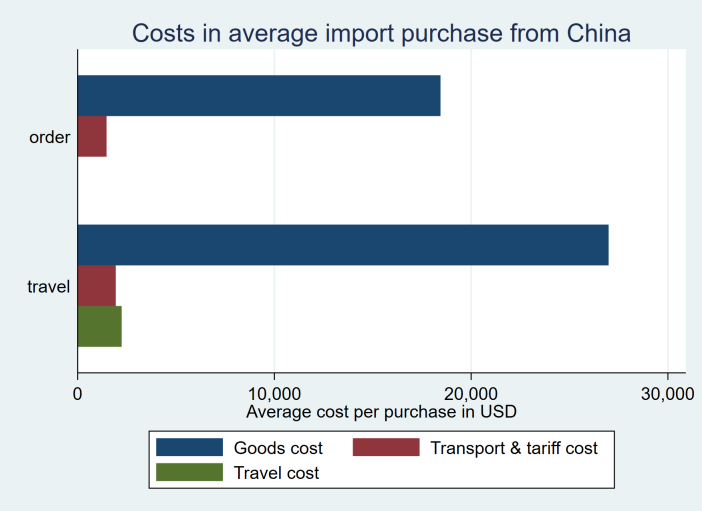

Fact 3 - Travel is persistent over time, even when traders have experience with a particular country or supplier

Given the focus of the existing literature on dynamic patterns in information in trade (e.g. learning about partners' types in Antras and Foley (2015) or network formation in Chaney (2014)), a reasonable starting hypothesis would be that travel facilitates learning or initial matching with suppliers, and becomes less valuable over time. A trader might travel mainly to learn the ropes in a new country, or to evaluate the quality or trustworthiness of a supplier early in the relationship. This would suggest that the probability of travel should decline with experience. Empirically, that is not the case. Table 3 shows regressions relating whether or not an importer traveled when making a purchase to the importing firm's age, years of experience buying from a particular country, and the length of the relationship with a particular supplier. No

specification provides evidence that travel is decreasing in experience with a particular supplier. Column (2) is consistent with a slight decrease in the probability of travel as a trader gains experience in a particular country but this effect is small and explains little of the total variation in travel. Overall, the data does not support learning or information problems that are mitigated with experience as the main mechanisms driving travel.

Table 3: Probability of traveling when importing

Dependent variable:	(1) Traveled	(2) Traveled	(3) Traveled
Business age	0.007		
	(0.005)		
Years buying from country	0.008	-0.014**	
	(0.007)	(0.007)	
Years buying from supplier	-0.004	0.005	0.001
	(0.006)	(0.004)	(0.005)
Observations	3048	3019	3241
Sector x country FEs	yes		
Trader and country FEs		yes	
Trader x country FEs			yes
Mean business age	10.65	10.64	10.79
Mean years buying from country	5.46	5.45	5.46
Mean years buying from supplier	3.84	3.84	4.04

Notes: Observations are at the purchase level. All columns are linear probability models. Standard errors clustered at the trader level are shown in parentheses. Observations for which the profit ratio (i.e. revenue divided by purchase cost) is in the top 5% or bottom 5% are excluded. Results are robust to various treatments of outliers, available on request.

* significant at 10% ** significant at 5% *** significant at 1%.

Fact 4 – Traveling is associated with lower costs, new styles, switching suppliers, and larger and less frequent shipments

Purchases made while traveling look systematically different from those ordered remotely. Table 4 shows the relationship between traveling and five key transaction features: 1) log unit cost of goods purchased (paid to suppliers), 2) buying a new product, 3) buying from a new supplier, 4) frequency of shipments, and 5) total cost of goods per shipment. Columns (1) - (5) regress these outcomes on an indicator

variable for whether or not the purchase involved travel, with source country-sector fixed effects. Goods purchased while traveling have 32% lower unit costs. Traders are 7 percentage points more likely to buy a new product when they travel, and 8 percentage points more likely to buy from a new supplier for the first time. Shipments involving travel are larger and less frequent, with traders waiting 2 months longer between shipments, and buying an additional $7,047 worth of goods relative to shipments made by ordering.

These patterns reflect both selection into travel and any causal effect of travel on transaction outcomes. Although I do not interpret the coefficients as measuring a treatment effect of travel, it is reasonable to ask whether average differences in transactions involving travel versus ordering simply reflect unobserved differences between traders themselves that correlate with both travel and outcomes (e.g. taste for travel and ability). Columns (6) - (10) show the same outcomes in regressions with trader-country fixed effects[9], which yield results qualitatively similar to the first specification.[10]

Taken together, the empirical patterns presented in this section motivate the structure of a model relating travel and importing via search and contracting problems. Traveling when importing is common, and systematically related to the costs of reaching a particular country and the type of product being sourced. Expenditures on travel are comparable in size to transportation and regulatory costs combined, suggesting a motivation of quantitative importance. Traveling and ordering are both feasible methods of importing, and the choice between the two appears to be driven by persistent underlying heterogeneity, rather than learning or matching needs related to experience. Finally, transactions that involve travel and those that don't differ in ways potentially relevant to consumer welfare. These patterns can all be accounted for by the existence of substantial search and contracting frictions. However, to gauge the size of these frictions and their effects on welfare, it is necessary to disentangle

[9] All traders deal in products in a single sector, so this is equivalent to trader-country-sector fixed effects.

[10] These coefficients are identified off the subset of traders who both travel to and order from the same country during the study period, and so should be closer to zero based on any model of selection into travel—those who switch strategies over time in response to small changes costs or benefits will be those who had smaller marginal gains from travel in the first place. Or, these traders may pursue different strategies when buying different products. Either explanation is consistent with the model I show in the following section, in which selection into travel is driven by heterogeneity in trader and product characteristics.

Table 4: Relationship between travel and transaction features

Dependent variable:	**Panel A – Across traders**				
	(1) Log unit cost	(2) New style	(3) New supplier	(4) Months between purchases	(5) Total goods cost ($US)
Traveled	-0.32**	0.07*	0.08***	1.98***	7,046.65**
	(0.113)	(0.043)	(0.030)	(0.412)	(2,889.024)
Observations	2,736	3,531	3,349	1,047	1,346
Mean of outcome	1.90	0.51	0.20	5.77	21,225.37
Sector x country FEs	yes	yes	yes	yes	yes
Trader x country FEs					
Dependent variable:	**Panel B – Within trader**				
	(6) Log unit cost	(7) New style	(8) New supplier	(9) Months between purchases	(10) Total goods cost ($US)
Traveled	-0.14**	0.01	0.12***	0.59	9,966.93***
	(0.064)	(0.046)	(0.042)	(0.57)	(2,173.04)
Observations	2,644	3,428	3,256	1,042	1,134
Mean of outcome	1.90	0.51	0.20	5.37	21,225.37
Sector x country FEs					
Trader x country FEs	yes	yes	yes	yes	yes

Notes: Outcomes in columns 1 - 3 and 6 - 8 are measured at the purchase level, and observations for which the profit margin of the transaction is in the top 5% or bottom 5% are excluded, . Outcomes in columns 5, 6, 9 and 10 are measured at the shipment level, and observations for which the profit ratio is in the top or bottom 2% or for which the total revenue is in the top or bottom 2% are excluded. Results are robust to various specifications of outliers, available on request. Columns 2, 3, 7, and 8 are linear probability models.

* significant at 10% ** significant at 5% *** significant at 1%.

the underlying heterogeneity that drives selection into travel from the causal effects of search and contracting problems. I turn to this question in the following section by building a model that maps heterogeneity in the type of product being imported to variation in observable choices in equilibrium.

4 A Model of Importing With Search and Contracting Frictions

I develop a model of importing in the presence of search and contracting frictions. Traders based in a "home" location face demand from consumers who value variety and low prices, but also care about the vintage of goods, preferring to buy up-to-date items. Traders play an intermediary role, finding ready-made goods around the world, and reselling them at home. Traders restock periodically, and face search and contract enforcement problems every time they do so. Two sourcing methods are available: traveling to make purchases, or ordering remotely. The model describes how forward-looking traders will choose the optimal sourcing method and frequency, endogenously selecting into travel based on heterogenous product characteristics and travel costs.

4.1 Environment

There are three types of agents: firms in upstream markets that manufacture goods (suppliers), firms downstream that act as intermediaries (traders), and final consumers in a home market. Goods must go through the traders to reach consumers – I assume manufacturers do not sell to them directly. Although it is distance between suppliers and traders that will matter rather than an international border per se, I take the upstream markets to be foreign source countries and traders to be importers for the purposes of the empirical application in the following sections.

To focus attention on the key ideas, I begin by considering trade in a single sector, sourced from a single country. In Appendix B, I extend consideration to multiple source countries and sectors in order to take the model to data.

4.1.1 Demand

Home country consumers are endowed with income E and demand differentiated goods, $i \in \Omega$. Each good has the potential for unlimited vertical improvements, as in the quality ladders model of Grossman and Helpman (1991). Updates arrive in increments of size g via a Poisson innovation process with an arrival rate λ_i that may differ across goods. For want of a better term, I refer to a given version of a good as the "style". The initial style for all goods is $z_{i0} = 1$, and the frontier style of good i at time t is $z^*_{it} = g^{J_{it}}$ where J_{it} is the total number of innovations that have arrived. The actual style available to consumers in the home market at time t is denoted z_{it}, and may not be at the frontier, so that $z_{it} = g^{J_{is}}$ for some $s \leq t$.

Consumers prefer more up-to-date styles of each good, and have a taste for variety. A representative consumer maximizes flow utility given by:

$$U_t = \left(\int_{i \in \Omega} (\bar{z}_{it} \psi_i)^{\frac{1}{\sigma}} q_{it}^{\frac{\sigma-1}{\sigma}} di \right)^{\frac{\sigma}{\sigma-1}}$$

where q_{it} is the quantity of good i consumed at time t, Ω is the measure of goods, $\sigma > 1$ is the elasticity of substitution, ψ_i is a time-invariant taste shifter for good i, and $\bar{z}_{it}$ is a normalized measure of the style available to consumers at time t relative to expectations about the frontier.[11]

4.2 Firm technology

4.2.1 Foreign producers

Firms in the source country produce goods using a constant returns to scale technology with marginal cost c.[12] Many firms produce each good.

Blueprints for improved styles of a good arrive exogenously at a Poisson rate λ_i, and independently across goods. Each firm producing a good receives the most recent blueprint with exogenous probability θ_1, which is i.i.d. with respect to each arrival - blueprints are not cumulative so it makes no difference whether the firm had the

[11] Normalized style is defined $\bar{z}_{it} = (J_{is} - \mathbb{E}_t J_{it}) \ln g + D$ where D is a constant. This normalization delivers analytical tractability while preserving intuitive demand forces as discussed in Appendix B.

[12] Allowing for variation in the cost of production across goods, so that marginal cost for good i is $\frac{c}{\varphi_i}$, does not change the theoretical results. However, it is not separately identified from the taste shifter ψ_i in the estimation, and so I abstract from it for simplicity of exposition.

second-to-lastblueprint. Innovations do not change the cost of production. I assume that more than one firm always has the latest blueprint, and that these firms engage in Bertrand competition.

4.2.2 Traders

There is a fixed measure of traders based in the home country, and each deals in a single unique good i, which is sourced from foreign producers and resold without transformation to consumers in the home market. Traders make purchases at discrete points in time, and continuously sell down stocks in the home market between purchases. The style of goods held as stock are fixed at the time of purchase, while innovations continue to arrive in the upstream market.

Purchasing goods from foreign suppliers incurs a per-shipment fixed cost F_i and a per-unit cost c_i (each of which may vary by trader and sourcing method, as described below), and an iceberg trade cost τ to bring the goods to the home market (which is the same for all goods).[13] The total cost of sourcing a quantity x is therefore $C_i = c_i \tau x + F_i$.

Because each trader deals in one unique good, I refer interchangeably to a trader or a good i characterized by $\{\lambda_i, \psi_i, F_i\}$ where λ_i is the rate of product innovation, ψ_i is a taste shifter, and F_i is the fixed cost of sourcing. Traders may exit if they do not find it profitable to trade.

Assume for the moment the existence of a constant aggregate price index P, which allows us to derive demand for each good at time t, and variable profits as a function of prices and demand. (In Section 4.4, I will derive the price index.) Regardless of the sourcing strategy, optimal pricing is a constant markup over variable cost, while demand is increasing in the vintage and taste parameter.

$$p_{it} = \left(\frac{\sigma}{\sigma - 1}\right) c_i \tau \qquad \forall\, t \tag{1}$$

$$q_{it} = \bar{z}_{it} \psi_i \left(\frac{p_{it}}{P}\right)^{-\sigma} \frac{E}{P} \tag{2}$$

[13] Traders may have other fixed costs of operation that are not purchase-specific (e.g. rent for a shop), which are irrelevant to the choice between traveling and ordering, but I assume they have no additional marginal costs.

$$\pi_{it} = \frac{1}{\sigma} p_{it} q_{it} \tag{3}$$

Note that demand is time-varying but deterministic, so that a trader can perfectly anticipate the total quantity she will sell over a period of length T. Demand falls as the vintagee held in stock depreciates relative to the expected frontier, declining at rate $\frac{\partial q}{\partial t} = -\lambda_i \psi_i \ln g \left(\frac{p_{it}}{P}\right)^{-\sigma} \frac{E}{P}$. Intuitively, goods for which the frontier evolves more quickly (higher λ_i) have faster effective depreciation rates. It will be convenient for the next section to represent flow profit as a function of unit cost and current style at time t, which are determined by traders' choices, while holding exogenous parameters fixed: $\pi\left(\bar{z}_{it}, c_i\right)$.

4.3 Optimal sourcing

When sourcing goods from foreign producers, traders face search and contracting problems. To maximize their profits in the presence of these information frictions, they choose both the method—whether to travel to the source market in person, or order remotely—and the frequency of restocking purchases.

I solve for the optimal restocking period and sourcing method as a function of good-specific and market-level parameters. I begin by finding the optimal frequency of restocking, conditional on choosing to travel or order. Then, I solve for the choice to travel, order, or exit the market, depending on which yields the highest total profits. I assume that all parameters are constant over time, so that choosing the importing strategy that maximizes per period profits over a single restocking cycle is equivalent to maximizing lifetime profits for an infinitely-lived firm.

4.3.1 Optimal stocking with travel

Traveling to buy in person is assumed to completely solve information problems associated with distance. Optimal stocking with travel is therefore the same as in a world without information frictions, except that traders must incur a fixed cost in order to achieve the "frictionless" outcome.

When traveling to buy, a trader pays a fixed cost $F_i = F_i^{tr} + G$. The first component is costs incurred when the trader travels, which includes explicit costs of travel such as plane fare, but may also capture the opportunity cost of the trader's time, like

or dislike for travel, and so on. The second component, G, represents other fixed costs of importing that would be incurred regardless of whether the trader travels, and which are assumed not to vary individually. Once the fixed cost is paid and the trader travels to the source country, she perfectly observes the frontier style, z^*_{it}, and the suppliers who are able to provide this style (i.e. who have the most recent blueprint). Because there are always multiple sellers of the frontier style and they produce at constant marginal cost and engage in Bertrand competition, the trader is able to buy any quantity x at a price equal to the marginal cost of production c.

The trader must choose how frequently to travel to make purchases. On the one hand, more frequent purchases mean selling goods that are more up-to-date. On the other, the fixed cost of restocking will be incurred more frequently. A trader chooses the length of time T between purchases to maximize total average flow profits:[14]

$$max_T \left\{ \frac{1}{T} \left[\int_{t=0}^{T} \pi \left(\bar{z}_{it}, c \right) dt - F_i^{tr} - G \right] \right\}$$

Variable profits decline over the restocking period as style depreciates. Solving the profit maximization problem yields a square-root rule for the optimal period of time to wait between purchases when traveling, T_i^{tr}:

$$T_i^{tr} = \sqrt{\frac{2 \left(F_i^{tr} + G \right)}{\lambda_i \psi_i \pi_{iz}}} \tag{4}$$

where $\pi_{iz} = \frac{1}{\sigma} \ln g \left(\frac{p_i}{P} \right)^{1-\sigma} E$. Intuitively, Equation 4 says that the optimal period to wait between restocking weighs the opportunity cost of selling out-of-date stock (in the denominator) against the fixed cost of making purchases (in the numerator).[15] The optimal stocking period is therefore shorter for varieties that evolve more frequently (higher λ_i) and are in greater demand (higher ψ_i), and for traders with idiosyncratically lower fixed costs of travel (smaller F_i^{tr}).

[14] Consumer expectations update after restocking and the aggregate price index is stationary, so optimal choices starting from any time t will be equivalent to the profit maximizing choice starting from time 0.

[15] Including physical inventory costs would yield this same rule, with a larger denominator (representing the cost of waiting) leading to a shorter restocking period.

4.3.2 Optimal stocking with ordering

When ordering remotely, traders face both contracting and search frictions. The contracting friction arises from the lack of third-party contract enforcement, requiring all arrangements to be self-enforcing. The search friction arises from the fact that products in the source market evolve, but the latest version is not immediately visible to traders who are searching remotely rather than in person. In this section, I specify the mechanisms underlying both frictions.

When ordering, a trader pays only the G component of the fixed cost of sourcing. Once this fixed cost is paid, the trader can source goods, but can only access a vintage $z_{it}^{or} \leq z_{it}^{*}$ and pays a unit cost $\tilde{c}_i \geq c$. The differences between the vintage and unit cost available when ordering versus traveling are determined by the size of search and contracting frictions. The total cost of an order of x units is $C_i^{or} = G + \tilde{c}_i \tau x$.

Contracting

In contrast to most trade models, I do not assume that agreements between buyers and sellers are costlessly enforced. Instead, buyers and sellers play a repeated game of moral hazard. Consistent with my data, I focus on equilibria in which cheating is off path and buyers pay before receiving goods and must induce honest behavior from sellers. The solution takes the form of a "contracting premium", analagous to an efficiency wage. Traders pay a premium over marginal cost, and cut off the relationship if the supplier ever cheats. The premium must be large enough that the supplier prefers to behave honestly and receive the expected stream of future rents, rather than cheating once and never receiving that buyer's business again. Buyers pay a unit cost $\tilde{c}_i(T) = \frac{c}{\delta_i(T)}$ where the contracting premium, $\frac{1}{\delta_i(T)} > 1$, will differ based on the product and the endogenous period of time between orders as derived below.

Infinitely-lived traders and suppliers play a repeated game. There is no access to third-party contract enforcement, although I relax this assumption to allow for costly enforcement in a counterfactual scenario in Section 6.[16] Traders buy from suppliers at discrete moments, with an endogenous period of length T between purchases. Recall

[16] The lack of third-party contract enforcement is consistent with the context in Nigeria as perceived by traders, who feel they do not have any recourse when cheated by suppliers, and is broadly similar to the operating environment for small businesses in many developing countries. As noted in the literature on relational contracting, transaction partners may rely on relational incentives like those modeled here on the margin, even when there is the potential to fall back on the legal system.

that there are always multiple suppliers able to supply a given style of each good. At the time of each purchase, a trader can buy from a single supplier, or has outside options to travel or to not make any purchase. Each supplier may sell to multiple traders, but these transactions are separable, so that their outside option is zero if they do not sell to a particular buyer. I assume that before transacting a supplier can observe a buyer's "type" as defined by the triple $\{\lambda_i, \psi_i, F_i\}$. After paying the fixed cost of importing, a buyer observes all sellers, and makes a take-it-or-leave-it offer to one of them. If the supplier accepts the offer, the buyer sends payment, and the seller then decides whether or not to send the goods as agreed.[17] After receiving the goods, the buyer sells down stock in the home market over a period of time T, and then repeats the same sourcing game. (Details of the stage-game and timing assumptions are shown in Appendix B.)

There is no public information about current or past transactions—both traders and suppliers know only the history of bilateral relationships in which they are involved, and nothing about other relationships. This rules out the influence of public reputation (Shapiro (1983)) or multilateral punishment strategies (Greif (1993)), and is close to the empirical reality observed in this setting.[18]

Traders and their suppliers perceive some probability that they will do business with the same partner at the time of the next purchase, which causes them to effectively discount the expected future surplus from the relationship. This uncertainty about whether future payouts will be realized comes from two sources. First, a given seller may not have the blueprint for the most recent style in the future. Second, the match may be permanently broken up due to exogenous factors.

I focus on sub-game perfect equilibria in which the equilibrium path features honest behavior, consistent with the fact that non-delivery is empirically rare. The trader offers a price $\tilde{c}^*$ and continues to buy from the same supplier so long as that supplier has not cheated, has the latest blueprint, and the match has not been exogenously broken up. If any of those events occur,, the trader never buys from that supplier

[17] Other timing is possible, but there is full prepayment by Lagos traders in 97% of observed transactions in the data. Therefore, I proceed by reducing the problem to a principal-agent framework in which the buyer moves first and has to induce cooperation from an opportunistic supplier.

[18] When asked who they would inform if they were cheated by a supplier, half of traders said they would tell no one, and most of the rest said they would tell only family and friends—only 16% said they would tell other traders or a market association. This is sensible when viewed in light of traders' incentives for information sharing – they have little interest in sharing potentially valuable information about suppliers with competitors.

again. When looking for a new supplier, the trader randomly selects from among the set of potential suppliers who have the latest blueprint and have not cheated that trader in the past. If the offer to the initial supplier is rejected, the trader continues to make the same offer sequentially to other suppliers. If the offer is never accepted, she defaults to her outside option. The supplier accepts any offer, but cheats if the payment is $\tilde{c} < \tilde{c}^*$ and sends the goods if $\tilde{c} \geq \tilde{c}^*$.

The probability that a match-breaking event occurs will depend on the length of time between purchases, T. Recall that new styles arrive at Poisson rate λ_i, and a given supplier receives the blueprint for each style with probability θ_1. Shocks leading to random match break-up are also Poisson, arriving at rate θ_2. Therefore, the total probability that the match continues in the next restocking cycle if the supplier has behaved honestly is:

$$\delta_i(T) = e^{-\theta_2 T}\left[1 - (1-\theta_1)\left(1 - e^{-\lambda_i T}\right)\right] \tag{5}$$

where $e^{-\theta_2 T}$ is the probability that the match has *not* been broken up exogenously, and $(1-\theta_1)\left(1-e^{-\lambda_i T}\right)$ is the probability that at least one new style has arrived but the current supplier does not have the blueprint. Faster changing products (with higher λ_i) have a lower probability of match continuation, all else equal, because it is more likely that at least one new style will have arrived that the supplier may not be able to produce.

The unit cost offers that induce good behavior in the SPE described above are determined by an incentive compatibility constraint that ensures the supplier would rather behave honestly than cheat:

$$\frac{(\tilde{c}_i - c)x_i(T)}{1-\delta_i(T)} \geq \tilde{c}_i x_i(T)$$

where the lefthand-side is the expected value of the stream of future profits from continued trade, and the righthand-side is the value of cheating now and never receiving the trader's business again. There are many unit costs that satisfy the incentive compatibility constraint. However, Bertrand competition between suppliers implies that the equilibrium outcome will be the minimum unit cost that satifies the constraint, $\tilde{c}_i(T) = \frac{c}{\delta_i(T)}$. The trader will take the dependence of the cost on the time between purchases into account when choosing the optimal ordering frequency.

Search

Incomplete information about the set of products available in source countries is represented by a lag between the time when a new blueprint arrives in the upstream market and the time when it is observable by traders in the downstream market. In contrast to traveling, which allows the trader to observe the true frontier style, when ordering the trader only observes the style that was on the frontier at time $t-\alpha$. This is a simple way of representing a search problem, but captures the key force relating information about product characteristics to distance and travel. When traders travel to the source market, they can observe the characteristics of goods actually available at that moment in the source market. When they do not travel, they use remote communication methods such as the internet or word of mouth, and new products may take some time to spread in this way.

Ordering frequency

The trader's profit maximization problem takes the same form when ordering as when traveling, but now accounts for the fact that unit cost will depend on the frequency of ordering, and that the latest style found will be α periods behind the frontier:

$$max_T \left\{ \frac{1}{T} \left[\int_{t=0}^{T} \pi \left(\bar{z}_{it}^{or}, \tilde{c}_i \left(T \right) \right) dt - G \right] \right\}$$

where $\bar{z}_{it}^{or} = \ln \left(\frac{z_{i,s-\alpha}^*}{\mathbb{E}_t z_{it}^*} \right) + D$ and $\tilde{c}_i \left(T \right) = \frac{c}{\delta_i(T)}$.
The solution for $T_{i,}^{or}$is implicitly defined by the first-order condition:

$$0 = \frac{\pi_c \left(\bar{z}_{it}^{or}, \tilde{c} \left(T \right) \right)}{\bar{z}_{it}^{or}} \tilde{c}' \left(T \right) \ln g \left(\lambda_i \left(-\alpha - \frac{1}{2} T \right) + \frac{D}{\ln g} \right) - \frac{1}{2} \lambda_i \ln g \frac{\pi \left(\bar{z}_{it}^{or}, \tilde{c} \left(T \right) \right)}{\bar{z}_{it}^{or}} + \frac{G}{T^2} \tag{6}$$

where π_c is the derivative of the profit function with respect to the unit cost of goods. As $\{\theta_1, \theta_2\} \rightarrow \{1, 0\}$ so that the total hazard of match break-up goes to zero, the contracting premium also goes to zero, and the solution to the contracting problem approaches the perfect enforcement solution achieved when traveling. The optimal restocking period will be shorter than the efficient restocking period under perfect enforcement. Intuitively, this is because the contracting premium increases as the restocking period gets longer, increasing the costs of waiting to restock, while the gains (waiting an additional period to incur the fixed cost) have remained the same.

As in the travel problem, faster arrival of new styles (higher λ_i) and larger sales volumes (higher ψ_i) lead to a shorter optimal restocking period. However, the effect of λ_i on stocking frequency is magnified in the ordering problem, because it increases not only the opportunity cost of selling out-of-date stock but also the contracting premium, because opportunitistic suppliers correctly perceive that the probability of doing business with the same trader in the future is lower.

The main effect of the search friction α is to shift the overall level of profits, because traders start with a less up-to-date style at the beginning of the restocking period. The only effect on the optimal restocking period is indirect, via the derivative of profits with respect to unit cost – traders are less willing to incur marginal increases in unit cost when the quantity they sell is higher. The search friction interacts with the speed of product change λ_i in an intuitive way: the opportunity cost of a given time lag is higher when the number of improvements that arrive in that period is higher.

4.3.3 Choice to travel, order, or exit

The previous section shows optimal restocking solutions, conditional on market participation and either traveling or ordering. Traders will choose to travel if the loss in variable profits due to the search friction and relational contracting premium is sufficiently high compared to the fixed cost of travel.

I denote maximized average total profits per period of time under travel and order strategies as Π^{tr} and Π^{or}:

$\Pi_i^{tr} = \frac{1}{T_i^{tr}} \left[\int_{t=0}^{T_i^{tr}} \pi\left(\bar{z}_{it}, c\right) dt - F_i^{tr} - G \right]$

$\Pi_i^{or} = \frac{1}{T_i^{or}} \left[\int_{t=0}^{T_i^{or}} \pi\left(\bar{z}_{it}^{or}, \tilde{c}\left(T\right)\right) dt - G \right]$

The trader chooses to trade if $max\left\{\Pi_i^{tr}, \Pi_i^{or}\right\} \geq 0$ and travels if $\Pi_i^{tr} \geq \Pi_i^{or}$, where the former is a zero profit condition for market participation.

4.4 Aggregation and welfare

Closing the model requires solving for the aggregate price index and the measure of varieties that will be traded. Still restricting attention to varieties within a single country-sector group, I define a style-weighted aggregate price index.

In order to consider a stationary equilibrium, I assume that the timing of initial market entry and therefore the timing of restocking periods is i.i.d. across the distribution of

goods. At any given time, the actual style of a variety can be replaced by its average over the restocking cycle, so that the aggregate stationary price index is:

$$P = \left(\int_{i \in \Omega} E_t \left[\bar{z}_{it} \right] \psi_i p_i^{1-\sigma} di \right)^{\frac{1}{1-\sigma}} \tag{7}$$

Consumer spending is fixed at E, and so consumer welfare is simply $U = \frac{E}{P}$. Changes in parameters can affect consumer welfare through two extensive margins – traders entering or exiting, and traders switching between traveling and ordering – and two intensive margins – changes in how up-to-date product styles are and changes in price.

To demonstrate the role of each margin, consider a change in a component of the travel cost that affects all traders, such as a fall in plane fares. This will have three effects on consumer welfare. First, traders who travel both before and after the change will choose to travel more frequently due to the reduction in fixed cost, which will make goods sold by those traders more up-to-date. Second, some traders who did not previously find it profitable to trade will choose to enter the market and travel, which increases the set of varieties available to consumers. Third, some traders will shift from ordering to traveling. The net effect of this last term on consumer welfare is ambiguous – prices of goods sold by those traders will fall (due to the elimination of the contracting premium), but they may be more or less up to date, depending on whether the elimination of the remote search lag is larger or smaller than the decrease in the frequency of importing.

In order to bring this model to data in which consumers buy products from many sectors sourced from many countries, I nest the framework above in Cobb-Douglas utility function over country-sector specific composites of differentiated goods (details provided in Appendix B). Consumers spend a constant share of their income on goods from each country-sector, and all of the parameters and equilibrium objects previously defined are country-sector specific.

5 Estimation

In order to quantify search and contracting frictions and understand their effects on consumer welfare, I structurally estimate the model presented in the previous section. I estimate thirteen parameters for each source country k and sector j, following a two-step procedure. In the first step, I estimate four parameters that can be inferred

directly from the data, separately from the rest of the system. In the second step, I make parametric assumptions about the form of underlying product/trader heterogeneity, simulate trader decisions conditional on the first stage parameters using the decision rules from the model, and then estimate the remaining parameters using a simulated method of moments approach. The main objects of interest are the search friction for each country-sector, α_{kj}, and the contracting premiums paid by traders when they order remotely, $\frac{1}{\delta_{ikj}}$.

5.1 Sample definition

The estimation uses data from the Lagos Trader Survey, and takes each trader's purchases from a single source country as the unit of observation.[19] I estimate parameters for 12 country-sector combinations with a sufficiently large sample of traders, with European and West African countries each aggregated into groups. These combinations account for 69% of total expenditure in the data. I exclude observations in which a co-owner lives in the source country (8.2% of the total).[20]

5.2 First step: Estimating trade costs, elasticities, and supplier match parameters

I begin by estimating four parameters for each country-sector group that can be extrapolated directly from the data, without using the structure of rest of the model: iceberg trade costs τ_{kj}, the elasticity of substitution across varieties σ_{kj}, the probability of the current supplier having the most recent blueprint θ_{1kj}, and the hazard of exogenous match breakup with the current supplier θ_{2kj}.

Iceberg trade costs

Physical and regulatory trade costs τ_{kj} are calculated directly from reported expenditures on transportation, shipping agents' fees, and clearing the port in Nigeria. I calculate the ad valorem equivalent for each shipment by dividing total transportation and regulatory expenditures by the value of goods being shipped, and then averaging across all shipments for a given country-sector group in 2013-14.

[19] A trader who sources shirts from both China and India is treated in the estimation as two traders dealing in two varieties, treating importing decisions as separable across countries.

[20] In this case, someone is able to conduct transactions with suppliers in person and so it is not clear how to treat the contracting problem. The migration decision is consistent with the forces modeled, and could be considered as another extensive margin choice (along with traveling, ordering, or exiting), but is omitted here due to a lack of detailed data on the migration choice.

Column (1) of Table 5 shows the values calculated for goods from each country in each sector. These range from a low of 5% for apparel from other West African countries to a high of 33% for apparel from the US. The ranges are consistent with costs of oceanic shipping found in the literature (e.g. Hummels (2007)), and follow generally sensible patterns (lower for nearby countries, higher for high weight-to-value goods such as toiletries).

Elasticity of substitution across goods

CES demand implies that the elasticity of substitution between varieties within a country-sector group is related to the markup charged to consumers: $p = \frac{\sigma}{\sigma-1} c\tau$. Since p, c, and τ are all observed,[21] I calculate the trade cost-adjusted markup $\frac{p}{c\tau}$ for each shipment, and then average across all shipments for a given country sector to estimate σ_{kj}.

Estimates are shown in column (2) of Table 5, and range from a high for toiletries ($\sigma_{kj} = 6.78$) from the UAE to a low for apparel from China ($\sigma_{kj} = 2.22$), with most groups falling between 2 and 5. This is largely consistent with the literature - for instance, estimates in Ossa (2014) range from 2 to 4 for most categories including non-food consumer goods.

Supplier match parameters

The repeated game with foreign suppliers presented in Section 4 implies that the probability of switching suppliers when ordering is a function of time between purchases, T_{ikjt}, and unobserved parameters θ_{1kj} (the probability of a supplier having the most recent blueprint), θ_{2kj} (hazard of match break-up), and λ_{ikj} (rate of product innovation):

$$\delta_i(T) = e^{-\theta_2 T} - e^{-\theta_2 T}(1-\theta_1)\left(1 - e^{-\lambda_i T}\right)$$

Since the realization of the product innovation process is observed, I can eliminate the dependence on λ_i and instead write two conditional probabilities:

$Pr\left(new\ supplier \,|\, new\ style\right) = 1 - \theta_1 e^{-\theta_2 T_{tikj}}$

$Pr\left(new\ supplier \,|\, old\ style\right) = 1 - e^{-\theta_2 T_{tikj}}$

If no new styles have arrived since the last purchase, a trader only switches to a new

[21] This assumes marginal costs other than stock and trade costs are minimal. For wholesale businesses with few employees that don't engage in transformation of goods, this is a reasonable approximation.

Table 5: First stage parameter estimates

Country	Sector	Iceberg trade cost τ	Elasticity of substitution σ	Pr. supplier has style θ_1	Hazard rate of supplier breakup θ_2
China	Apparel	1.19	2.22	0.96	0.02
China	Electronics	1.10	4.24	0.93	0.01
China	Beauty	1.32	5.03	1.00	0.02
China	Hardware	1.23	3.88	0.97	0.02
China	Homewares	1.30	2.90	0.96	0.02
Europe	Apparel	1.16	2.78	0.92	0.04
Turkey	Apparel	1.28	2.78	0.91	0.01
UAE	Apparel	1.22	2.46	0.94	0.02
UAE	Electronics	1.08	2.61	0.98	0.03
UAE	Beauty	1.26	6.78	0.86	0.02
US	Apparel	1.33	2.29	0.95	0.00
West Africa	Apparel	1.05	4.88	0.94	0.02

Notes: Iceberg trade costs are defined as one plus the sum of reported goods transportation and regulatory costs divided by the value of goods for each shipment, and averaged across shipments.

supplier if the match with the current supplier has been exogenously broken up in the meantime (which occurs at hazard rate θ_2). If a new style has arrived, the trader buys again from the same supplier so long as that supplier has the most recent blueprint (which occurs with probability θ_1) and has not been hit with a match break-up shock.

I estimate θ_1 and θ_2 jointly using a maximum likelihood routine for each country-sector group. Estimates are shown in columns (3) and (4) of Table 5. Lower θ_1 and higher θ_2 imply that the risk of a supplier relationship breaking is greater. This leads the contracting premium associated with a given restocking period to be higher, which is taken into account in trader's choice of the optimal restocking period and the choice of whether to order or travel. The average premium actually paid by traders who order in a given sector from a given source country will depend on θ_1 and θ_2, but also on equilibrium selection into importing strategies.

5.3 Second stage estimation

In the second stage, I use a simulated method of moments approach to estimate the remaining nine parameters for each country-sector. Given parameter estimates from the first step, parametric assumptions about the distribution of underlying prod-

uct/trader heterogeneity, and a starting parameter vector, I simulate the profit maximizing decisions of each trader, and calculate a set of aggregate moments. I then iterate over choices of parameters and select the parameter vector that minimizes the distance between the simulated moments and their data analogues.

In order to pin down the search and contracting frictions, I need to fully characterize the features of all traders and their equilibrium choices – whether they order or travel, how frequently they import, what profits they earn, and the frequency with which they find new product styles. Traders are heterogenous along three dimensions, $\{\lambda_i, \psi_i, F_i^{tr}\}$, none of which are directly observable in the data (although they will be strongly related to variation in outcomes - the probability of finding new styles, the level of variable profits, and the frequency of travel, respectively). Therefore, I make parametric assumptions about the distribution of characteristics. I take each element of $\{\lambda_i, \psi_i, F_i^{tr}\}$ to be log-normally distributed and ex-ante independent of one another.[22] In practice, I will estimate the means of the λ and F distributions, $\left\{\mu_{kj}^{\lambda}, \mu_{kj}^{F}\right\}$ (but not μ_{kj}^{ψ}, which is normalized to zero because the average is pinned down by parameters of the profit function described below), and the standard deviations $\left\{\sigma_{kj}^{\lambda}, \sigma_{kj}^{\psi}, \sigma_{kj}^{F}\right\}$.

Note that while the underlying distributions of trader characteristics are independent, those of active traders will be correlated due to endogenous exit of traders with unprofitable combinations of characteristics. All else equal, traders with faster changing varieties (higher λ_i), those facing less demand (lower ψ_i), and those with idiosyncratically higher costs of travel (higher F_i^{tr}) will earn lower profits and be more likely to exit. In equilibrium, for instance, we should expect high λ varieties to only be traded if there is high demand for them or a trader finds travel unusually cheap.

Given the distribution of trader characteristics and the parameters estimated in the first stage, there are four remaining parameters to be estimated: the maximum and derivative of variable profits per period with repect to style (π_0 and π_z), the fixed cost of ordering (G), and the search lag associated with ordering (α). I draw a sample of simulated "traders" and for each one find the profit maximizing import frequency conditional on ordering or traveling, given by Equations (5) and (7). I calculate maximized average monthly variable profits and average monthly total profits conditional

[22] With sufficient data, the λ_i could in theory be estimated for every trader i directly from observables $\mathbf{1}_{ti}^{prod}$ and T_{ti} via the relationship $Pr\left(\mathbf{1}_{ti}^{prod} = 1\right) = 1 - e^{\lambda_i T_{ti}}$. However, I do not observe enough transactions for most traders to make this practically implementable.

on ordering or traveling:

$$\bar{\pi}_{ikj} = \begin{cases} \psi_{ikj}\left(\pi_{0kj} - \frac{1}{2}\pi_{zkj}\lambda_{ikj}T_{ikj}^{tr}\right) & if\ \mathbf{1}_{ikj}^{tr} = 1 \\ \psi_{ikj}\left(\delta_{ikj}\left(T_{ikj}^{or}\right)\right)^{\sigma_{kj}-1}\left(\pi_{0kj} - \pi_{zkj}\lambda_{ikj}\alpha_{kj} - \frac{1}{2}\pi_{zkj}\lambda_{ikj}T_{ikj}^{or}\right) & if\ \mathbf{1}_{ikj}^{tr} = 0 \end{cases}$$

$$\bar{\Pi}_{ikj} = \begin{cases} \bar{\pi}_{ikj} - \frac{G_{kj}+F_{ikj}^{tr}}{T_{ikj}^{tr}} & if\ \mathbf{1}_{ikj}^{tr} = 1 \\ \bar{\pi}_{ikj} - \frac{G_{kj}}{T_{ikj}^{tr}} & if\ \mathbf{1}_{ikj}^{tr} = 0 \end{cases}$$

where $\delta_{ikj}\left(T_{ikj}^{or}\right)$ is as defined in Equation (6), and π_{0kj} and π_{zkj} are constants to be estimated for each country-sector group without separately identifying the constituent parts defined in Section 4. I then allow each simulated trader to choose whether to travel, order, or exit the market, depending on $max\left\{\Pi_{ikj}^{tr}, \Pi_{ikj}^{or}, 0\right\}$.

To estimate the nine parameters in the second stage (five parameters of trader heterogeneity distributions and π_0, π_z, G, and α), I match twelve aggregate simulated and data moments for every country-sector: the mean and variance of months between purchases for travelers and orderers (4), the mean and variance of monthly profits for travelers and orderers (4), the correlation between monthly profits and time between purchases (1), the fraction of purchases involving new styles for travelers and orderers (2), and the overall fraction of traders who travel (1).

Simulated method of moments selects parameters $\hat{\Psi}_{kj}$ to minimize the distance between simulated and empirical moments, using the criterion function:

$$\hat{\Psi}_{kj} = \min(m_{kj,data} - m_{kj,sim}\left(\Psi_{kj}\right))\mathrm{W}(m_{kj,data} - m_{kj,sim}\left(\Psi_{kj}\right))'$$

where $m_{kj,data}$ is the vector of empirical moments for country k and sector j and $m_{kj,sim}\left(\Psi_{kj}\right)$ is the vector of simulated moments calculated at Ψ_{kj}.[23]

5.4 Structural results

In this section, I present the results of second stage estimation, and compare the results to a variety of non-targeted data moments and plausibility considerations.

[23] I use the diagonal of the variance-covariance matrix of the moments as the weighting matrix W rather than the full variance-covariance matrix, due to concerns about bias. I use numerical methods to find gradients used in calculating the standard errors, taking the first-stage estimates as given.

Table 6 shows a subset of the data moments used in the estimation and their simulated analogues (the remainder are shown in Table 1 in Appendix C).

The estimated search frictions imply that goods found via remote search are, on average, the equivalent of 1.7 months behind the frontier. The estimated contracting premiums imply that traders who choose to order pay an average cost premium of 15.5%. The latter can be compared to the regression relating travel to unit costs in Section 3. Although unit cost differences were not targeted by the estimation procedure, the coefficient on travel in column (6) of Table 4 implies a premium of 14%, and the structural estimate is within the confidence interval.

The full set of estimated parameters is shown in Table 2 in Appendix C. To aid in interpretation, Tables 7 and 8 show three estimates with clear economic interpretations and comparisons to non-targeted moments available in the data. Table 7 shows the average number of new styles expected to arrive per year for goods that are traded in equilibrium and the search friction (α) by sector, averaged across source countries. The former follows intuitively plausible patterns: apparel is the sector with the fastest-changing goods, followed by electronics and homewares, while beauty products and hardware change only about a third as often as apparel and half as often as electronics. This is also largely consistent with traders' subjective rating of how frequently new products become available in their line of business, which follows the same ranking across sectors (except for beauty products, which traders feel change more often than the estimates suggest). The main moments in the data that drive identification of the rates of product change are the frequency of importing and whether or not the product was the same as something previously purchased. Conditional on how long it has been since the last purchase, finding new styles more frequently implies higher λ_i.

The second half of Table 7 shows the average search friction by sector. The estimates imply that the search friction is largest for electronics and beauty products (on average across source countries), with an effective lag of more than two months in the goods found via remote search versus in-person search. Recall that the lag, measured in units of time, also interacts with the rate of product change. Apparel products are more affected by the search friction than the lag alone suggests because the rate of product evolution is higher – a 1.7 month lag in apparel means getting goods that are on average a half style arrival behind the frontier, while the similarly sized search lag in beauty products only implies goods that are the equivalent of 0.1 arrivals behind.

Table 6: Data moments and simulated moments

Country	Sector	Moment:		Order			Travel		
			Travel %	**New style %**	**Mean var. profit / month**	**Mean months btwn. ship.**	**New style %**	**Mean var. profit / month**	**Mean months btwn. ship.**
China	*apparel*	**Data**	**0.42**	**0.63**	**1,160**	**6.72**	**0.55**	**1,870**	**7.42**
		Sim	0.40	0.64	1,009	7.14	0.71	2,528	8.24
China	*electronic*	**Data**	**0.34**	**0.33**	**511**	**7.21**	**0.54**	**1,004**	**8.94**
		Sim	0.33	0.40	509	9.01	0.53	1,096	10.89
China	*beauty*	**Data**	**0.26**	**0.25**	**825**	**7.87**	**0.37**	**603**	**6.13**
		Sim	0.22	0.30	422	5.66	0.59	767	14.87
China	*hardware*	**Data**	**0.15**	**0.18**	**769**	**7.08**	**0.18**	**1,750**	**11.80**
		Sim	0.16	0.17	778	7.12	0.35	1,622	11.79
China	*homeware*	**Data**	**0.24**	**0.36**	**1,099**	**8.17**	**0.44**	**1,067**	**11.36**
		Sim	0.27	0.34	773	8.32	0.60	1,706	11.41
Europe	*apparel*	**Data**	**0.52**	**0.54**	**131**	**5.45**	**0.61**	**935**	**6.70**
		Sim	0.52	0.40	485	4.38	0.56	1,115	8.51
Turkey	*apparel*	**Data**	**0.34**	**0.58**	**1,148**	**6.71**	**0.72**	**391**	**5.55**
		Sim	0.34	0.58	765	6.24	0.95	1,406	13.18
UAE	*apparel*	**Data**	**0.57**	**0.51**	**193**	**8.95**	**0.72**	**866**	**6.90**
		Sim	0.58	0.41	237	8.34	0.70	708	7.59
UAE	*electronic*	**Data**	**0.56**	**0.38**	**407**	**8.63**	**0.54**	**879**	**9.13**
		Sim	0.54	0.41	283	6.70	0.56	1,009	9.80
UAE	*beauty*	**Data**	**0.38**	**0.26**	**464**	**7.01**	**0.34**	**422**	**10.02**
		Sim	0.38	0.12	472	3.66	0.51	532	10.92
USA	*apparel*	**Data**	**0.62**	**0.70**	**561**	**4.63**	**0.79**	**1,101**	**10.51**
		Sim	0.54	0.80	729	6.41	0.73	1,073	9.17
W. Africa	*apparel*	**Data**	**0.88**	**0.57**	**843**	**2.47**	**0.79**	**559**	**5.97**
		Sim	0.90	0.57	417	1.98	0.79	743	6.14

Notes: This table shows seven of the twelve moments matched in the estimation. The remaining five are the variances of profits and months between shipments for orders and travel (4) and the correlation between profits and the months between shipments (1); the full table is provided in Appendix C.

Table 7: Second stage estimates by sector

Innovation rates – Expected improvement arrivals per year					
Sector	*Apparel*	*Electronics*	*Homewares*	*Hardware*	*Beauty*
Estimate	3.16	0.83	0.83	0.40	0.52
Non-targeted data (rank)	1	3	4	5	2
Search friction – Months behind frontier					
Estimate	1.67	2.42	.96	0.28	2.31

Notes: The non-targeted data is the relative rank of each sector according to traders' subjective assessment of how frequently new styles become available in their line of business Estimates of arrival rates and search frictions are averages across countries.

The estimates of α are driven by the combination of the selection of different λ_i into travel versus ordering (i.e. the difference in the frequency with which travelers versus orderers find new styles) and the differences in variable profits across traveling and ordering.

Table 8 shows the fixed cost of travel paid by traders who select into travel in equilibrium, averaged by source country. These can be compared with the travel costs reported in the data, which were not targeted by the estimation. The estimated fixed costs include any implicit costs of travel, positive or negative, such as the opportunity cost of traders' time and any taste or distaste for travel to a particular location, while the costs in the data include only explicit costs such as plane tickets and visas. The estimated costs of travel are substantially higher than the reported costs for Turkey and the UAE, and similar for West Africa, China, and Europe. They are lower than reported costs in the United States, where it is plausible that this reflects positive utility associated with the travel itself. The lower estimated cost of travel reflects the fact that traders travel more frequently than the reported costs of travel would suggest they should, given the differences in variable profits and the probability of finding a new style associated with travel versus ordering.

5.5 Welfare and market structure under frictionless counterfactual

In order to understand the effect of information problems on welfare and market structure, I consider how equilibrium outcomes would change if in-person and remote search were equally effective (i.e. $\alpha = 0$) and costless third-party contract enforcement

Table 8: Second stage estimates by country

Travel cost – $US per trip

Country	*China*	*Europe*	*Turkey*	*UAE*	*US*	*W. Africa*
Estimate	$2,237	$2,060	$2,793	$2,548	$1,211	$135
Non-targeted data	$2,236	$1,988	$1,631	$1,380	$2,030	$59

The non-targeted is the average cost per trip reported by traders, including plane fares, visas, hotels, ground transportation and any other travel-related expenditures. Values reported in various currencies are converted to $US at the average mid-market exchange rate in the month the trip was taken. Structural estimates of travel costs are averages across sectors.

were introduced.[24] In the absence of both search and contracting frictions, welfare from the modeled consumer goods sectors increases by 14%. The gain from eliminating the search problem alone is 4.5%, and the gain from perfect contract enforcement alone is 7% – the whole is greater than the sum of the parts due to an interaction through the discrete choice to travel. Shutting down information frictions has welfare effects through four channels: 1) prices, 2) the vintage of goods, 3) the set of varieties, and 4) firm profits.[25]

Removing information frictions also changes features of market structure, as illustrated in Figure 4. Average firm profits are higher in the frictionless world, but average firm size is actually smaller because information problems push the smallest firms out of the market. The set of varieties available to consumers increases, and the new varieties are those for which styles or technologies change rapidly or for which total demand is small. Perhaps counterintuitively, the effect of information problems on the average vintage of goods available is ambiguous – the search problem always makes available goods more out-of-date, but this is in some cases offset by the fact that the contracting game induces traders to buy more frequently than is efficient in order to keep the unit cost premium down.

[24] I now need to separate out the price index that was previously imbedded in the profit function constants π_0 and π_z. To do so, I normalize relative to the price index at baseline, and restate parameters in terms of the ratio of the counterfactual and baseline price indices, $\hat{R}$ so that $\hat{\pi}_0 = \pi_0 \hat{R}^{\sigma-1}$ and $\hat{\pi}_z = \pi_z \hat{R}^{\sigma-1}$. I solve for firm choices under the counterfactual, calculate the implied price index and $\hat{R}$, and find a fixed point for $\hat{R}$.

[25] I assume that increases in trader profits are rebated lumpsum to consumers.

Figure 4: Changes in market structure without information frictions

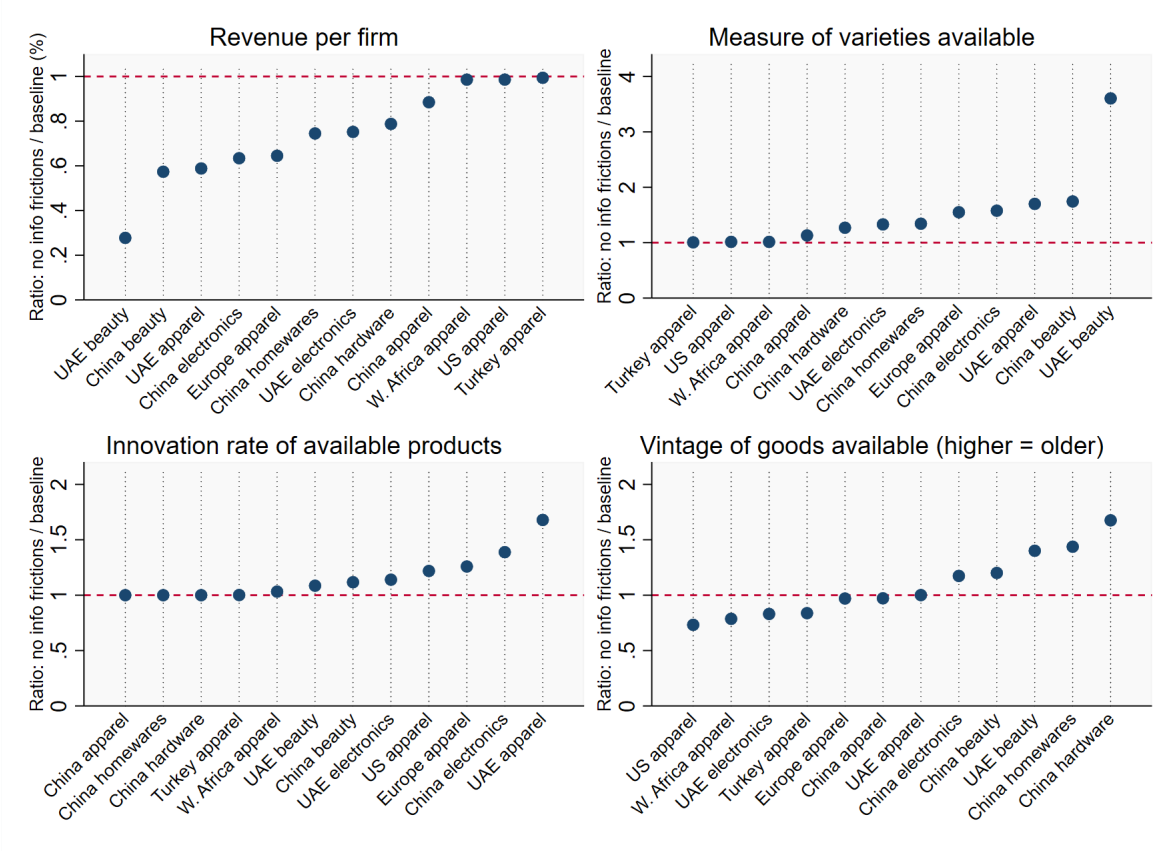

6 Counterfactuals and policy applications

What does the quantitative importance of information frictions imply for trade facilitation policy, and for our understanding of the relationship between trade costs and development? In this section, I use the estimated model parameters to consider counterfactual scenarios of three types: 1) changes in market features that vary at different levels of economic development; 2) policies that target information problems in trade indirectly, by affecting the price of coping strategies; and 3) policies that target information problems directly.

6.1 Developed country expenditure and firm size

The availability of fixed cost strategies for addressing search and contracting problems, such as travel, raises the question of how economies of scale influence the welfare effects of these frictions. In a standard trade model with monopolistic competition and CES

preferences, an increase in expenditure per capita affects consumer welfare through two channels: the value of increased consumption, and an increase in the variety of goods available. The latter occurs because it becomes profitable to produce and trade some varieties when the demand at a given price increases relative to the fixed cost. In the model with search and contracting frictions, an increase in per capita expenditure effects welfare through two additional channels: first, traders will import more frequently, which makes average stock more up-to-date and lowers prices via the contracting premium; and second, some traders switch from ordering to traveling.

I consider how moving Nigeria to US-level expenditure per capita changes the relative welfare losses due to search and contracting frictions, via the implied increase in the potential revenue associated with selling a particular variety (which is analogous to firm size in this model). The fraction of total consumer spending that goes to non-food consumer goods is roughly the same in the United States as in Nigeria – 15.9% versus 16.7%. However, because the United States is far wealthier, per capita spending on these goods in 2014 was approximately 20 times higher.

I re-estimate the model holding the baseline estimated parameters fixed but increasing expenditure to US levels. The average "firm size" (or, revenue stream associated with a single variety) increases from roughly $27,000 per year to almost $325,000 per year. I then re-estimate the model with the same (US level) expenditure, but with no search friction and perfect contract enforcement. Recall that in the baseline, removing search and contracting frictions yields a 14% welfare increase. At US levels of expenditure, the comparable gain is only 6%. Intuitively, the losses due to search and contracting frictions are smaller in a context in which firm scale is such that more varieties are imported via travel, or imported frequently enough that the equilibrium contracting premium is small.

This is a stylized way of representing the relationship between income levels, firm size, and information costs – I hold other parameters fixed, and do not model differences in market structure or import penetration that might also arise with income differences – but it provides insight into how the results presented in the previous section might generalize with respect to levels of economic development. When the potential revenue associated with a product is larger, firms will opt into fixed cost coping strategies and high purchasing frequencies, and so the relative gains to consumers from elimination of information frictions will be smaller, even when the underlying frictions themselves are the same.

6.2 Reductions in travel costs

A large literature looks at how reductions in the costs of transporting goods affects market integration and gains from trade (summarized in Donaldson (2015)). One implication of this paper is that costs of movement for people may affect welfare as much as the cost of moving goods in some circumstances. Some policy tools affect both – such as investments in transportation infrastructure relevant to both passenger and freight traffic – but others that are not typically thought of in the context of trade facilitation may be unexpectedly important through their effect on the cost of solving information problems. Air travel regulation and visa policies fall in this category. International travel is both expensive and often restricted for Nigerians, along with residents of many African and other developing countries. For instance, Nigeria is ranked 92nd out of 104 countries on the 2016 Henley and Partners Visa Restrictions Index, and 68th out of 75 countries in the ratio of international flight costs per 100km traveled to GDP per capita (based on the 2016 Kiwi Aviation Price Index).

The restrictedness and cost of mobility is clearly influenced by policy decisions. Air service agreements (ASAs) are international treaties that govern air travel. ASAs may regulate the number of airlines that can fly between two countries, which routes they can fly and how frequently, what prices they can charge, and the total capacity or type of airplanes that can be used. The economic effects of improvements in air services and visa policies have been considered, but primarily through the lens of gains from tourism or attracting capital or multinationals (e.g. Campante and Yanagizawa-Drott (2018)). Here, I consider whether important economic effects may actually come indirectly through the price and availability of goods, especially in developing countries.

To investigate the effect of air travel liberalization on gains from trade, I consider what would happen if Nigeria and China signed an Open Skies Agreement (OSA), which implies fully liberalized air travel, with unrestricted access for airlines in each country. The ASA signed by Nigeria and China in 2014 allows for only one airline from each country to land in one city in the other (Guanzhou in China and Lagos in Nigeria), and requires schedules and rates to be approved by the regulating bodies of both countries. A small empirical literature has estimated the effects of liberalizing ASAs on flight traffic and pricing. I borrow from Cristea, Hummels and Roberson (2012), who estimate that moving to a full OSA decreases a quality-adjusted flight

price index by 32% on average across routes between the United States and a range of partner countries.

To implement this scenario in the model, I calculate the average reduction in flight cost ($372, or 32% of the average $1162.50 cost of a plane ticket from Nigeria to China reported in the LTS data), subtract it from the fixed cost of travel faced by each simulated trader buying from China under the baseline estimates from Section 5, and allow traders to reoptimize in response to the new cost. I assume that the policy change has no effect on ticket pricing to other destinations, and do not calculate the gains to any travelers or businesses other than the modeled consumer goods importers. The travel cost reduction operates on welfare through four channels: 1) traders who travel earn higher total profits due to fixed cost savings, 2) traders travel more frequently and provide more up-to-date goods, 3) some varieties which were not profitable to trade when travel costs were higher become available, and 4) some traders switch from ordering to traveling, which lowers prices but has an ambiguous effect on vintage. In total, the fraction of traders who travel when importing from China increases by 18%, and welfare from the modeled sectors increases by a little over 0.55% (including sourcing from other countries where there are no reductions in travel cost).

6.3 Financial services that improve contracting

It is reasonable to ask is whether and how terms of payment have arisen to cope with the contracting problem traders face. One might think, for instance, that traders would have "on deposit" or "open account" arrangements with suppliers, sending some or all of the payment only after goods have been received and confirmed to be satisfactory. In fact, this is extremely rare in this context. More than 97% of total payment value is made before goods are received. This is consistent with the hypothesis that buyers and sellers actually face a two-sided contracting problem – sellers in source countries have no more reason to trust Nigerian traders to send payments once goods have been received than traders do to trust sellers to send goods once payment has been received.

Another option is to use a financial intermediation service provided by a third party. Although the Lagos Trader Survey does not measure use of bank intermediation directly, it appears to be rare based both on the near universal pre-payment and the fact that it is almost never mentioned by traders in response to questions about how

they guarantee the quality of imported products (in contrast, traveling, hiring an agent, and having a long relationship with a supplier are all commonly mentioned). This is consistent with what we know about trade finance from other contexts – Niepmann and Schmidt-Eisenlohr (2017) report that less than 15% of global trade by value involves bank intermediation (via letter of credit or documentary collection) and that it is less common for trade involving sub-Saharan African countries.

Formal intermediation services are available to Nigerian traders, and so it is useful to consider what the model implies about demand for these services and their potential effect on the market. To do so, I introduce an escrow service into the model as a potential third importing "method" available to traders. This service closely mirrors the actual terms offered (at the time covered by the data) by Alibaba.com's secure payment service. Alibaba is a large online business-to-business trading platform that is based in China but has both sellers and buyers all over the world, housing $464 billion in third-party transactions in fiscal year 2016. For a limited set of suppliers based in China, the service allows buyers to make a full payment up front, but holds it in escrow until delivery of the goods is confirmed. The service costs 5% of the value of the transaction.

In the model, I assume that a supplier who attempts to cheat by collecting payment without having sent goods gets "caught" with some probability when the buyer has used the escrow service and is unable to keep the money. Although cheating is still off the equilibrium path, this alters the expected return to cheating, and therefore reduces the contracting premium needed to satisfy the supplier's incentive compatibility constraint. I assume that both traders and their suppliers share the same perception of the likelihood of accurate enforcement in the event of cheating.

Figure 5 shows the fraction of traders who demand the secure payment service at varying beliefs about the accuracy of enforcement. If enforcement were completely accurate – i.e. both sides knew that any supplier attempting to cheat would be caught 100% of the time – then demand for the service would be quite high, at 88% of all traders buying from suppliers in China. As the perceived accuracy of enforcement declines, however, demand drops off steeply. I do not know the fraction of Lagos traders who have actually used the escrow service, but do observe that only 31% of traders in the LTS data who buy from China say they have ever used Alibaba for sourcing in any way. Taken as a strict upper bound on the set who choose to use the escrow service, this implies that they perceive the probability of

accurate enforcement in the event of a dispute to be worse than a coin flip. While this is a rough, back-of-the-envelope calculation, it is consistent with both the limited ability that a third-party has to adjudicate facts in a small, semi-formal, cross-border transaction, and the subjective perception among many traders that fraud is common in online transactions and that service providers will be biased toward Chinese sellers. In contrast, a less expensive escrow service would be in demand even at relatively low perceived enforcement accuracy. Figure 5 also shows demand for an escrow service charging 1% of transaction value, which remains above 50% even when perceived enforcement accuracy is below 20%. Agents and financial services do not offer a silver bullet for solving contracting problems in trade. They simply transfer the contracting problem onto another party, and consequently will be used only when they have access to some technology or relationship that allows the contracting problem to be solved at lower cost than the other options available, including traveling.

Figure 5: Demand for escrow service at cost of 5% and 1% of transaction value

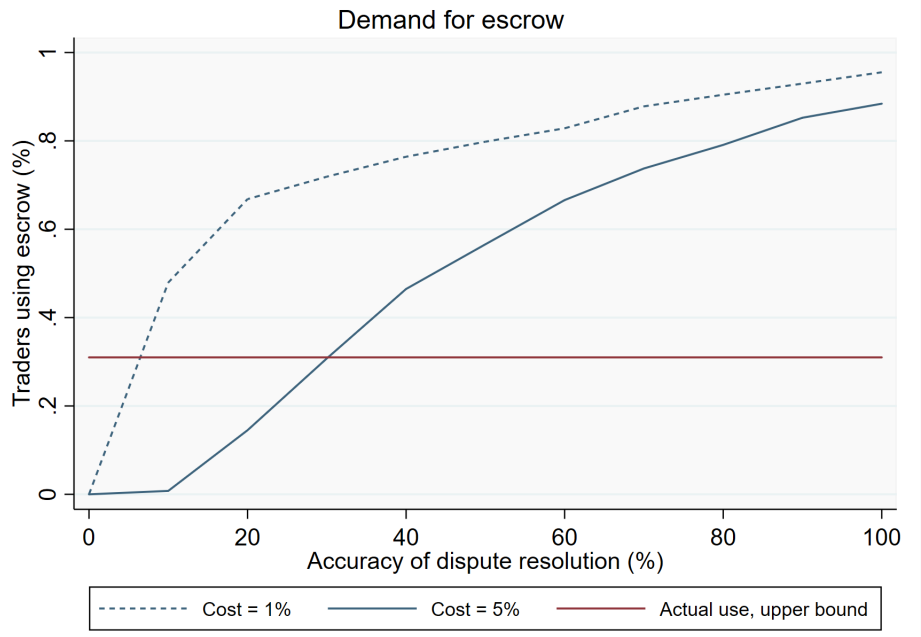

7 Conclusion

Search and contracting frictions in trade can have a large impact on consumer welfare. Using unique data that documents the process of trade in Nigeria at a transaction

level, I find that traveling to do business with suppliers in person is a key strategy for coping with information problems. To account for empirical relationships between travel and features of importing, I build a model of trade with search and contracting problems. Importers endogenously select into a mode and frequency of trade based on the characteristics of the goods they trade. The structure of the model combined with the detailed observables available in the data allow me to estimate both the total value of solving information problems implied by selection into travel, and to separately identify the roles of search and contracting.

Although the estimates are context-specific, the underlying mechanisms are general, and offer predictions about the features of markets in which search and contracting problems will have a large effect on welfare. While it is clear that effective contract enforcement and inexpensive information technology are likely to reduce the costs of search and contracting problems, the model suggests some less obvious relationships. For instance, search and contracting frictions will matter more for trade in products that change frequently, where the costs of traveling (or other fixed cost coping strategies) are high, and when firms are small. They will also matter more in markets where there is more "churning" in the sets of potential buyers and sellers, perhaps due to high entry and exit rates, or to a low probability that the best match in one period will also be the best match in future periods..

The influence of churning and firm size offer particularly interesting avenues for future research in light of evidence that the market power of intermediaries plays an important role in the prices faced by consumers (Atkin and Donaldson (2015), Bergquist and Dinerstein (2020)). On the one hand, a smaller set of large, stable traders should find it less costly to solve information problems (via strategies like travel, regular communication, long relationships, and reputation and coordination mechanisms). On the other, this market structure may yield higher markups and lower pass-through of reductions in information costs. It is ambiguous ex-ante which effect would dominate in consumer surplus.

This evidence on both the size of information frictions and the underlying mechanisms suggests a new range of trade facilitation policies and services, which may be particularly important to trade in and between developing countries. Interventions that attack information problems may be an effective way of achieving trade cost reductions, particularly in environments where tariffs are already low, or where improvements in regulation or infrastructure are very costly.

References

Ahn, JaeBin, Amit Khandelwal, and Shang-Jin Wei. 2011. "The role of intermediaries in facilitating trade." *Journal of International Economics*, 84(1): 73–85.

Aker, Jenny. 2010. "Information from Markets Near and Far: Mobile Phones and Agricultural Markets in Niger." *American Economic Journal: Applied Economics*, 2(3): 46–59.

Allen, Treb. 2014. "Information Frictions in Trade." *Econometrica*, 82(6): 2041–2083.

Anderson, James E., and Eric van Wincoop. 2004. "Trade Costs." *Journal of Economic Literature*, 42(3): 691–751.

Antras, Pol, and Arnaud Costinot. 2011. "Intermediated Trade." *The Quarterly Journal of Economics*, 126(3): 1319–1374.

Antras, Pol, and C. Fritz Foley. 2015. "Poultry in Motion: A Study of International Trade Finance Practices." *Journal of Political Economy*, 123(4): 853–901.

Atkin, David, and Dave Donaldson. 2015. "Who's Getting Globalized? The Size and Implications of Intra-national Trade Costs." NBER Working Paper.

Banerjee, Abhijit V., and Esther Duflo. 2000. "Reputation Effects and the Limits of Contracting: A Study of the Indian Software Industry." *Quarterly Journal of Economics*, 115(3): 989–1017.

Bardhan, Pranab, Dilip Mookherjee, and Masatoshi Tsumagari. 2013. "Middlemen Margins and Globalization." *American Economic Journal: Microeconomics*, 5(4): 81–119.

Bergquist, Lauren Falcao, and Michael Dinerstein. 2020. "Competition and Entry in Agricultural Markets: Experimental Evidence from Kenya." *American Economic Review*, 110(12): 3705–3747.

Biglaiser, Gary. 1993. "Middlemen as Experts." *The RAND Journal of Economics*, 24(2): 212–223.

Campante, Filipe, and David Yanagizawa-Drott. 2018. "Long-Range Growth: Economic Development in the Global Network of Air Links." *The Quarterly Journal of Economics*, 133(3): 1395–1458.

Chaney, Thomas. 2014. "The Network Structure of International Trade." *American Economic Review*, 104(11): 3600–3634.

Cristea, Anca. 2011. "Buyer-seller relationships in international trade: Evidence from U.S. States' exports and business-class travel." *Journal of International Economics*, 84(2): 207–220.

Cristea, Anca, David Hummels, and Brian Roberson. 2012. "Estimating the Gains from Liberalizing Services Trade: The Case of Passenger Aviation."

Crozet, Matthieu, Guy Lalanne, and Sandra Poncet. 2013. "Wholesalers in international trade." *European Economic Review*, 58: 1–17.

Donaldson, Dave. 2015. "The Gains from Market Integration." *Annual Review of Economics*, 7(1): 619–647.

Grant, Matthew, and Meredith Startz. 2020. "Cutting Out the Middleman: The

Structure of Chains of Intermediation."

Greif, Avner. 1993. "Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders' Coalition." *American Economic Review*, 83(3): 525–548.

Grossman, Gene, and Elhanan Helpman. 1991. "Quality Ladders and Product Cycles." *Quarterly Journal of Economics*, 106(2): 557–586.

Hummels, David. 2007. "Transportation Costs and International Trade in the Second Era of Globalization." *Journal of Economic Perspectives*, 21(3): 131–154.

Jensen, Robert. 2007. "The Digital Provide: Information (Technology), Market Performance, and Welfare in the South Indian Fisheries Sector." *Quarterly Journal of Economics*, 122(3): 879–924.

Krishna, Kala, and Yelena Sheveleva. 2016. "Wheat or Strawberries? Intermediated Trade with Limited Contracting." *American Economic Journal: Microeconomics*.

Leke, Acha, R. Fiorini, R. Dobbs, F. Thompson, A. Suleiman, and D. Wright. 2014. "Nigeria's renewal: Delivering inclusive growth." McKinsey & Company Global Institute.

Macchiavello, Rocco. 2010. "Development Uncorked: Reputation Acquisition in the New Market for Chilean Wines in the UK." SSRN 1559654, Rochester, NY.

Macchiavello, Rocco, and Ameet Morjaria. 2015. "The Value of Relationships: Evidence from a Supply Shock to Kenyan Rose Exports." *American Economic Review*, 105(9): 2911–2945.

MacLeod, W. Bentley. 2007. "Reputations, Relationships, and Contract Enforcement." *Journal of Economic Literature*, 45(3): 595–628.

McMillan, John, and Christopher Woodruff. 1999. "Interfirm Relationships and Informal Credit in Vietnam." *Quarterly Journal of Economics*, 114(4): 1285–1320.

Niepmann, Friederike, and Tim Schmidt-Eisenlohr. 2017. "International trade, risk and the role of banks." *Journal of International Economics*, 107: 111–126.

Ossa, Ralph. 2014. "Trade Wars and Trade Talks with Data." *American Economic Review*, 104(12): 4104–4146.

Poole, Jennifer. 2010. "Business travel as an input to international trade."

Portes, Richard, and Helene Rey. 2005. "The determinants of cross-border equity flows." *Journal of International Economics*, 65(2): 269–296.

Rauch, James, and Vitor Trindade. 2002. "Ethnic Chinese Networks in International Trade." *Review of Economics and Statistics*, 84(1): 116–130.

Shapiro, Carl. 1983. "Premiums for High Quality Products as Returns to Reputations." *Quarterly Journal of Economics*, 98(4): 659–679.

Steinwender, C. 2018. "Real Effects of Information Frictions: When the States and the Kingdom Became United." *American Economic Review*, 108(3): 657–696.

Taylor, Curtis, and Steven Wiggins. 1997. "Competition or Compensation: Supplier Incentives under the American and Japanese Subcontracting Systems." *American Economic Review*, 87(4): 598–618.