

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

Meredith Startz[†]

November 27, 2016 - [Most recent version](#)

Abstract

Distance between buyers and sellers can create search and contracting problems: how to find out what goods are available in far away places, and ensure they are actually delivered? Traveling to do business in person is one way of dealing with both, transforming a remote transaction into one that is face-to-face. I estimate the magnitude of search and contracting frictions in a developing country context by exploiting the fact that travel is a common, costly, and easily observable strategy for coping with them. I collect transaction-level panel data from Nigerian importers of consumer goods that combines the “what” of trade (e.g. products, quantities) with variables describing “how” trade is conducted (e.g. travel, payment terms). To account for patterns inconsistent with a full information environment, I build and estimate a model that embeds a search problem and a repeated game with moral hazard into a monopolistically competitive trade framework. Welfare from imported consumer goods would be 29% higher in the absence of both frictions. I decompose the total barrier into parts attributable to search and to contracting, and show why the effects will be larger in markets with low consumer spending, high firm entry/exit rates, and frequently changing products. The results suggest that greater attention to market integration policies beyond transportation and tariffs could have large welfare effects, particularly in developing countries. In counterfactual scenarios, I show that deregulation of air travel between Nigeria and China would yield gains in Nigeria on the order of \$650 million per year through consumer goods trade alone, while existing financial services do little to mitigate frictions because they do not offer a better contract enforcement technology than travel or repeated interaction.

*Thanks to Mark Rosenzweig, Chris Udry, Sam Kortum, Dean Karlan, and David Atkin for their guidance and support, as well as Chris Woodruff, Pete Schott, Costas Meghir, Dan Keniston, Giovanni Maggi, Shelby Grossman, Nick Ryan, Nancy Qian, Shelly Lundberg, Penny Goldberg, Matthew Grant, Dick Startz, Jeff Weaver, Sharat Ganapati, Glynis Startz, Ana Reynoso, Chiara Margaria, Tomasso Porzio, and Gabriella Santangelo for innumerable helpful conversations and comments, and to the participants of the Yale Development and International Trade seminars and the PEDL and FREIT networks. This project was made possible by incredible research assistance from Hakeem Bishi, Dylan Groves, Yemi Ogundare, Nicole Wilson, Anifat Ibrahim, Mayor Alaja, Aluko Olanrewaju, 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.

[†]Ph.D. candidate, Department of Economics, Yale University. Email: meredith.startz@yale.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 currently available with his own eyes, and exchange goods and money on the spot. But in China, an immense variety of shoes are produced at low cost. How can the Nigerian wholesaler—who cares about characteristics such as style, material, and fit—find out what varieties of shoes are currently available in Guangzhou? How can he be sure that a supplier in Guangzhou will actually deliver the shoes after taking his money? Even if he solves these problems once, he will face them again the next time he buys, because the set of available shoes will have changed and sellers will still be opportunistic.

One solution is to simply travel to China to do business in person. Travel solves these problems by eliminating the distance between the buyer and seller, effectively transforming a remote transaction into the type that would be conducted at home in Lagos. However, travel is not obligatory. Lagos wholesalers can search for products using phones, internet, and word-of-mouth, and opportunistic 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 can be used to quantify distance-related information problems, and to identify the underlying mechanisms relating them to firm and product characteristics. Although I make use of data that captures international trade and travel, the situation is analagous to transactions over distance within a country.

As the example suggests, distance between potential transaction partners creates both a search problem (the difficulty of finding out what goods are available) and a moral hazard problem (the difficulty of ensuring that goods or money actually get sent), which I refer to together as “information costs”. This may help explain a persistent empirical puzzle: the estimated trade costs implied by flows of goods are consistently much higher than directly observable costs such as transportation 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 perfect contract enforcement, a special case in which a search friction can be inferred from spatial price dispersion ([Jensen \(2007\)](#); [Aker \(2010\)](#); [Allen \(2014\)](#); [Steinwender \(2014\)](#)).

I show that search and contracting frictions in differentiated goods trade are large, and can have a substantial impact on welfare in developing countries. My approach is to quantify information

problems that are not directly observable by focusing on the easily observable strategies that importers use to cope with those problems, particularly travel. Unique panel data I collected from Nigerian importers allows me to do this by connecting traditional trade data (e.g. the type and value of goods traded) to 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 non-food consumer goods such as clothing, electronics, and furniture, who were randomly sampled from a census I conducted of over 50,000 shops in commercial districts of Lagos. I capture every import transaction over a two year period, totaling 3,907 purchases from 34 source countries and over a thousand foreign suppliers.

I document four empirical patterns that motivate the structure of a model relating travel and importing. First, travel is common but not universal—used by 62% of traders and in 32% of all shipments—and is less likely when importing from countries that are costly to reach. Second, travel expenditures are large, at nine percent of the value of imported goods—equivalent to amounts spent on transportation and regulatory costs of importing combined. Third, travel is persistent over time, and does not decline significantly in experience with particular countries or suppliers, suggesting motives beyond initial matching or learning. Finally, transactions that involve travel look different from those that don't. In purchases conducted in-person, traders pay lower unit costs, charge higher markups, and are more likely to buy new product varieties and to buy from new suppliers. These patterns are difficult to account for in a world without quantitatively important information frictions.

I build a model that accounts for these patterns by embedding a search problem and a repeated game with moral hazard into a [Melitz \(2003\)](#) style framework. Differentiated goods are produced in a source market and improve stochastically over time, as in the quality ladders model of [Grossman and Helpman \(1991\)](#). Traders source these goods from foreign suppliers, and resell them to consumers at home. They 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 products 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 sub-game perfect equilibria that feature full pre-payment by traders and honest behavior by suppliers along the equilibrium path. The solution is analagous to an efficiency wage: the trader pays a cost 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 trade based on underlying heterogeneity in how quickly the products they sell evolve over time and how popular they are. A key departure from most trade models is that the frequency of trade matters for welfare; it determines both the vintage of goods available to consumers and the unit cost (because smaller, more frequent purchases reduce the temptation for suppliers to renege), and therefore mediates the effect of search and contracting problems on consumers. Strategies are not symmetric across products (e.g. faster-changing products are more likely to be traded frequently and via travel), but can be aggregated into a stationary equilibrium with a constant price index. 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, and firms earn lower profits.

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 switching products, the frequency of purchases, and variable profits. The size of the contracting cost premium is pinned down by the observed probability that the trader buys from the same supplier again. In contrast, the search friction is related to the observed probability of finding new product styles and the difference in variable profits from traveling versus ordering. If we observe a large fraction of traders finding new products but never switching suppliers, all else equal, travel must be driven more by search motives than contracting. If finding new products is rare, but supplier switching is common, we infer that contracting motives are strong relative to search.

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 11 different source country-sector pairs. On average, importing without traveling yields goods that are 2.5 months behind the frontier available in the source country (the search friction) and requires paying an 11.9% price premium to induce good behavior from suppliers (the contracting friction). Removing both frictions increases welfare in the traded consumer goods sector by 29.2%—roughly half of the gains from eliminating physical and regulatory trade costs. This estimate is large, but plausible considering that it represents the complete elimination of a barrier in a context in which trade costs are high compared to most rich countries, and applies only 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 16.3%, and the gains from contracting alone would be 9.0%—the whole is greater than the sum of the parts due to an interaction through the discrete choice to travel. These totals reflect sensible underlying variation across countries and sectors—for instance, consumer welfare losses due to information frictions are larger when importing apparel from China than hardware (because apparel changes more often), and than apparel from Benin (because Benin is cheaper to travel to than China).

Information frictions also distort market structure. In the presence of search and contracting problems, average firm profits are lower than in a frictionless world, but average firm size is actually larger because the smallest firms are pushed out of the market. The set of varieties available to consumers is reduced, and the varieties that are lost are those that change rapidly or for which total demand is small. Surprisingly, 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.

By fully specifying the relationship between information frictions and underlying product and firm characteristics, this 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 travel costs are low, information problems will have little effect on welfare. Similarly, the effects should be smaller in 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 game toward what would be achieved under perfect enforcement. In the first of three counterfactual scenarios, I increase consumer spending to match that in the United States, and show that when revenue is higher, the same fundamental search and contracting frictions yield much smaller welfare costs. Intuitively, fixed cost strategies for addressing information problems in trade with China are inconsequential relative to revenue for Walmart; not so for small traders in Lagos.

This evidence on information frictions directs attention to a range of trade facilitation policies and services beyond 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 substantially, and the resulting gain in surplus would be equivalent to increasing Nigerian consumer spending by \$645 million. In a third counterfactual, I show why existing financial services have not gone further in ameliorating the contracting problem: they are too expensive. I add to the model the option for traders to purchase an escrow service similar to that currently offered by Alibaba.com, and find that demand is low at actual prices and plausible expectations about the effectiveness of dispute resolution. Agents and financial services simply transfer 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.¹

¹Financial services with better enforcement technology or lower prices could have substantial welfare effects. In ongoing work in Lagos, I am conducting a field experiment involving an escrow service based on reputation and default rules. By combining the theoretical framework in this paper with experimental data, I will be able to further explore the potential for financial services to address contracting problems, as well as tracing out the mechanisms through which reputation may substitute for formal enforcement.

The idea that information problems contribute to trade costs is an old one. Estimates of gravity models often include proxies for information problems between countries, 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 in distant locations, often in developing countries, and typically shows 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 \(2014\)](#)).

I go beyond this work by microfounding and quantifying both imperfect contract enforcement and search over products with differentiated characteristics. I build on ideas introduced in [Rauch and Trindade \(2002\)](#), who 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 often opt out of this relational contracting mode, choosing to use costly monitoring strategies like travel 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 large theoretical literature speculates about what middlemen do, and most of the proposed mechanisms have to do with solving matching or contracting problems (e.g. [Biglaiser \(1993\)](#); [Antras and Costinot \(2011\)](#); [Dasgupta and Mondria \(2012\)](#); [Bardhan, Mookherjee and Tsumagari \(2013\)](#)). 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 recent work by [Atkin and Donaldson \(2015\)](#) shows that they are important in determining the actual prices received by consumers. While I do not endogenize market structure or pass-through by traders in

this paper, modeling and providing empirical evidence on the production function of intermediaries themselves is an important step toward understanding their influence on prices and welfare.

2 Context and data

I study traders operating in Lagos, Nigeria, who import consumer goods from source countries around the world. In this section, I describe the environment in which those traders operate, the original survey data used in this project, and some basic facts about the traders and their businesses.

In 2015, Nigeria was the 6th most populous country in the world, with 182 million people, and the largest economy in Africa. Lagos is the commercial capital of Nigeria, and is the main port of entry for imported goods.² With a population of approximately 21 million people, the city of Lagos is roughly the size of the Scandinavian countries combined, and has a larger economy than Kenya.³

Nigerian production is heavily focused in agriculture and raw materials, and almost all manufactured goods, from shoes to heavy machinery, are imported. The non-food 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, a proportion which has been rising with incomes in recent years (though food still accounts for over half of all consumer spending).⁴ Although the Nigerian economy is often thought of as revolving entirely around oil extraction, a rebasing exercise by the National Bureau of Statistics in 2013 revealed that the wholesale and retail trading sector is actually the second largest contributor to GDP after agriculture, and was the largest contributor to GDP growth from 2010 to 2013, by a wide margin (Leke et al. (2014)). These facts suggest that gaining a better understanding of market failures in the trading sector in developing countries is economically important due both to its potential effects on consumers and on productivity in the trading sector itself.

2.1 Data: Lagos Trader Survey

The empirical evidence in this paper is based on original survey data I collected from traders located in Lagos, Nigeria. In this section, I describe the construction of the sample frame and the structure and content of the survey, referred to below as the Lagos Trader Survey, or LTS. Additional details are given in the Data Appendix.

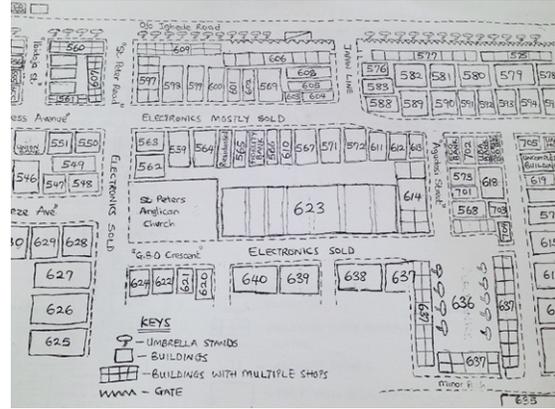
²Nigerian Ports Authority, 2016

³McKinsey Global Institute Report, calculations based on Lagos Bureau of Statistics figures for 2010.

⁴Euromonitor Passport database, consumer spending data for 2010 - 2015.



(a) Balogun Market, Lagos Island



(b) Mapping for shop census

Sample frame

Traders were identified through a census of 52,830 shops in commercial areas of Lagos conducted between October 2014 and April 2015.⁵ Since no business register or firm census exists, we began with a list of markets and plazas (indoor markets, typically located on privately-owned land) provided by LAWMA, the Lagos State waste collection agency, and added a number of market areas located on Nigerian federal government land. The goal of the listing was to identify wholesale and retail businesses that were likely importers of final goods, and it therefore focuses on commercial and wholesaling areas of the city, and does 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.

A team of research assistants covered all markets and plazas in the listing on foot, enumerating individual shops by location and the type of products being sold. The listing notes, for example, that on the second floor of Queen’s Plaza on Imamm Ligali Street in an area called Lagos Island, there were four shops selling clothing, nine selling shoes or bags, two selling luggage, and one canteen. The count included any business in permanent physical premises, but excluded mobile vendors and hawkers. The result is not a census of the entire city, but comes as close as possible to capturing the universe of shops in the key commercial and wholesaling areas identified through the listing.

After excluding offices, services, warehouses, and vacant or inaccessible shops, a simple random sample was taken from the remaining 24,159 shops.⁶

⁵The Lagos Trader Project, including both the census and survey, is a joint fieldwork effort with Shelby Grossman.

⁶Because the listing did not identify individual shops, the sampling took the form of selecting how many shops in a given location selling a given product type should be interviewed. For instance, the sampling might dictate that we interview the owner of one of the five shops selling shoes on the first floor of a plaza. In order to ensure that the selection of shops within that group was not biased, enumerators were given randomized “directions” to a shop – to

Data structure and survey content

Interviews for the Lagos Trader Survey were conducted between April and August 2015 with the owners of 1,179 shops, a response rate of 82%.⁷ The data contains information about 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 Lagos Trader Survey data combines information that might be found in detailed official customs data with a novel set of variables describing the actual process of firm-to-firm trade at a transaction level. 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 available in customs data from a limited set of countries, such as the date of individual transactions (rather than total annual or quarterly flows), buyer and seller identities, the price at which goods are sold in the importing country (which enables me to calculate markups), and physical and regulatory trade costs (including shipping, tariffs, and costs to clear the port in Lagos). I add to this a unique set of variables that describe how transactions are implemented, such as whether the buyer traveled to the source country, travel costs, whether the product was exactly the same as something that had been purchased previously (as opposed to a new style, size, 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. To my knowledge, no other data includes this level of information about the process of trade. Financial variables were reported by traders in a variety of currencies, and are converted to US dollars using the average mid-market exchange rate in the month in which the transaction occurred.⁸

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 (trader demographics, number and type of employees, revenue, use of financial services, and others), 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 (including countries it does not currently buy from), and whether the trader has imported or traveled in the past even if they do

enter the designated floor of the plaza, turn in a randomly selected direction (left or right), and count to the Xth shop (randomly selected out of the total number) selling the designated product.

⁷Of the 18% total of non-respondents, 12% could not be located due to changes in market areas (e.g. several large markets on Lagos Island burned down in January), shops that had closed, or listing errors. The remaining 6% were located but the interview could not be conducted, either due to the respondent refusing to participate, or inability to make contact with the owner of the business (as opposed to an employee or apprentice).

⁸This introduces some error since plane tickets may not have been purchased in the month in which the transaction occurred, and reported sale prices of goods are the average over the actual period in which stock was sold. However, inflation in Nigeria was relatively low over the relevant period and the exchange rate between the naira and dollar was quite stable. The period covered by the data ends just as a major exchange rate crisis hits Nigeria in December 2014.

not do so in 2013/14. It also asks about the full history with any suppliers they have purchased from in the past two years, including the year when the relationship was originally established and how they came to know the supplier. 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), 2,481 international shipments, and 3,907 individual product import purchases over a two year period. Traders report importing from 32 different source countries, and transacting with 1,073 suppliers.

In order to obtain information at this level of detail, several strategies were employed that introduce some limitations. First, individual suppliers are tracked over time and connected to particular transactions using unique identifiers, but these identifiers are often nicknames that will not allow the same supplier to be identified across buyers or in other data sources. This was necessary to elicit information from respondents, who are extremely secretive about the sources of their goods (which is in itself suggestive about the value of information in this market). Second, most respondents were asked for details about every shipment, but to avoid unmanageably long interviews, traders who reported more than six trips or orders (6% of respondents) were only asked for details about a randomly selected six of each. A small number of traders with more than 12 trips (/orders) (4% of respondents) were instead asked about a “typical” trip to (/order from) each source country during the recall period. These observations are reweighted in all analysis.

2.2 Context: Importing consumer goods in Lagos, Nigeria

Who are Lagos traders?

Both importing and traveling to make import purchases are very common in the sample of traders surveyed. Almost all traders are selling items produced outside Nigeria, but I use the word “importer” to refer to those who make purchases directly from suppliers abroad. Table 1 shows the fraction of traders who imported or traveled internationally for business during 2013 or 2014 or in the past (pre-2013). All analysis in the remainder of the paper focuses on the 620 traders who imported at least once in the 2013/2014 period. Of those who imported directly, 62% traveled for business at least once.

Traders are owner-operators of wholesale import firms. Table 2 shows summary statistics describing trader demographics and the basic characteristics of their importing businesses. The typical trader is male, middle-aged, and educated, with more than half having some post-secondary education and 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 themselves, instead sourcing ready-made products abroad and reselling them in Lagos. A few do some branding or customization of products, but this is not common. Due to the nature of the business, these firms have a high ratio of working capital to labor inputs. They look large relative to developing country businesses in some respects, with average annual revenue of \$107,581 from directly imported stock (with a long right tail; the median is \$27,113). In others, they look more typical of small and medium enterprises in developing countries—the average trader has one shop and one worker (including paid workers, apprentices, and family employees). Most specialize in one category of products, such as women’s shoes, home appliances, or cosmetics, and the majority of business is wholesale, although a significant minority of reported revenue is from retail sales.

Imports and sourcing methods

The imports represented in the Lagos Trader Survey data cover a wide variety of products and are sourced from more than thirty different countries. Figure 1 shows the top twelve countries, by number of traders sourcing there in Panel A, and by the total value of imports observed in Panel B. By far the most common and largest source in terms of import value is China, but the United Arab Emirates (specifically, Dubai), Turkey, Hong Kong, Benin, India, and perhaps surprisingly, the United States and United Kingdom are also common. A number of these source countries appear to serve as entrepot locations—clothing from Benin and electronics from Dubai are unlikely to have been originally manufactured in those locations. Indeed, 70% of suppliers are described as upstream wholesalers rather than manufacturers, but this reflects substantial variation across countries. Nearly all West African suppliers are other wholesalers, while buying directly from a manufacturer is much more common in Europe and East Asia. In the analysis below, I do not make any distinctions based on the type of supplier, but it is worth noting that the transactions modeled here may be only one step in a supply chain featuring similar information frictions at other stages, both up- and down-stream.

For the purposes of analysis, I group products into six categories: apparel (which includes shoes, bags, and textiles), electronics, beauty, hardware, home goods, and miscellaneous other products. Table 3 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 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 the 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 (e.g. car parts and electrical components) changes the least frequently.

Traders' explanations of why they travel rather than simply ordering over the phone or internet highlight search and contracting problems consistently and explicitly. Survey responses to this question almost universally focused on these issues, with comments such as "So that I can ensure the quality and determine if there is [sic] new designs in the market", "Designs change so one needs to know what the market is selling", "It is very risky to order products from abroad because of low quality of the good [and] dubious suppliers", and "To get information about new products, to get sincere suppliers to deal with and also to [get] cheaper products for profit". Traders also note that they get better prices when they travel, which gives them a competitive edge: "For my profit margin to come up [a] little, and sell cheaper, at least a bit lesser than what others [charge ...] will help sale[s] move".

Explanations of why or why not to travel also emphasize the costliness of this strategy, and the necessity to achieve a certain scale in the business overall or the size of the purchase to make it worthwhile: "Sales are slow now 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", or "For me to travel, I have to consider time and money spend to get visa, airfare and also to buy in large quantity".

3 Empirical patterns in trade and travel

In this section, I describe four facts about the empirical relationship between importing and travel. These facts are difficult to account for in a world without quantitatively important information frictions, and guide the structure of the model in Section 4.

Fact 1 - Travel is common but not universal, and is less likely when importing from destinations that are costly to reach

Traveling to the source country to buy goods in person when importing 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 2013/2014, 53% always ordered remotely, and the remaining 13% did some of both.

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.

Fact 2 - Travel is costly relative to the value of goods imported, and comparable to the costs of transportation and clearing the port combined

Figure 3 shows the average value of goods imported in a single shipment, broken out by whether the trader traveled or ordered. The average ordered shipment contains \$14,856 worth of goods, and involves \$1,619 in transportation and regulatory costs (10.9% of the shipment value).⁹ The average shipment involving travel was larger, at \$18,427 worth of goods, but with a similar 9.4% paid in transport costs. Travel expenditures averaged \$1,622, or 8.8% of the value of goods.¹⁰ This number includes the cost of plane tickets, visas, and any 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 like or dislike for travel in itself. In the structural estimation that follows, I will allow the fixed cost of travel to vary freely and capture these indirect costs as well. However, explicit travel costs alone, aggregated across all countries and sectors, are still revealing. First, travel is sufficiently costly relative to the value of goods to be of quantitative importance. Second, travel costs are on the same order of magnitude as the transportation and regulatory costs that are the most studied components of overall trade costs.¹¹

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), network formation in Chaney (2014), and learning about demand in Eaton et al. (2014)), a reasonable starting hypothesis would be that travel facilitates learning or initial matching, 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

⁹Transportation and regulatory costs are the total reported by traders in response to three questions: 1) "How much did you pay in total for transporting/shipping for ALL the products you bought on this trip [/in this order]?" 2) "How much did you pay in total to clear the port in Nigeria for everything you bought on this trip [/in this order]? Please include any tariffs, agent fees, and tips." 3) "How much did you pay this agent?" where "this agent" refers to an earlier question about any paid agent who inspected the goods in an order.

All potential regulatory costs are bundled together under the total cost of "clearing the port" both because some traders simply pay a total to an agent who specializes in getting goods through the port, and may not actually know what portion goes to tariff payments versus other costs, and because clearing the port likely involves bribes and potential misreporting of the type of goods, and questions that attempt to separate out these components would be sensitive and unlikely to elicit honest responses.

Payments to agents in the source country are included in this amount because shipping and inspection services are often bundled as one service and traders are unable to separate out the two components. This raises the possibility that some part of this expenditure in fact also reflect a strategy for solving information problems. I take a conservative approach by attributing the entire amount to transportation and regulatory costs.

¹⁰Travel expenditures are the total reported by traders in response to three questions: 1) "How much money did you spend on visas for this trip?" 2) "How much money did you spend on airfare for this trip?" and 3) "How much money did you spend on other travel expenses for this trip, such as ground transportation, hotels, hiring agents, or money exchange?"

¹¹The 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. A survey question that asked directly about carrying home goods in luggage provides further evidence on this point. Traders report that they carried any goods home in their luggage 28% of trips, and even among those that did, the average value of goods carried was less than one-sixth of the total value of goods purchased.

early in the relationship. This would suggest that the probability of travel declines in experience; however, that is empirically not the case. Table 4 shows linear probability model regressions relating whether or not an importer traveled when making a purchase to the importing firm’s age, years of experience in a particular country, and the length of the relationship with a particular supplier. Column (2) includes trader fixed effects, and column (3) includes trader-country fixed effects. No specification provides evidence that travel is decreasing in experience with a particular supplier. Column (2) suggests there may be 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 variation in travel.

Overall, the data provides no evidence to support learning-driven travel or a model in which information problems dissipate as a trader finds out what products are available in a country or identifies inherently honest or reliable suppliers. Instead, in the model that follows, I will take new products to be constantly arriving in the source market, and all potential suppliers to be opportunistic, so that past experience does not reduce the need to travel.

Fact 4 – Traveling is associated with lower costs, higher markups, new styles, and switching suppliers

The most important motivation for the model developed in the following section is that purchases made while traveling look systematically different from those made via remote ordering. Table 5 shows the relationship between traveling and four key transaction features: 1) log unit cost of goods purchased (paid to suppliers), 2) markups (unit sale price in Lagos divided by unit cost), 3) whether or not the product was precisely the same as something previously purchased, and 4) whether or not it was the first time buying from the supplier. Columns (1) - (4) 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 31% lower unit costs, but are sold at higher markups back in Lagos (an additional 34% markup over cost). Traders are 7 percentage points more likely to buy a new product (in the broad sense that it is not exactly the same as something they previously purchased – whether it is a completely different product, or simply a new model or color) when they travel, and 5 percentage points more likely to buy from a new supplier for the first time.

These patterns reflect both selection into travel and any causal effect of travel on transaction outcomes. Although I do not interpret the coefficients as reflecting 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. skill and taste for travel). Columns (5) - (8) show the same four outcomes in regressions with trader-country fixed effects (all traders deal in products in a single sector, so this is equivalent to trader-country-sector fixed effects), which yield results qualitatively similar to the first specification.¹²

¹²The coefficients in columns (5) - (8) are identified off the subset of traders who both travel and order from the

Taken together, the four 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 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 are difficult to account for in a world without substantial search and contracting frictions. However, to gauge the size of these frictions and their effects on welfare, it is necessary to disentangle 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, and is able to account for all the empirical patterns documented in this section.

4 Model

In this section I develop a model of importing in the presence of search and contracting frictions, with differentiated goods that improve over time. Traders based in a “home” country 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 must search for new goods and face a contract enforcement problem every time they do so. Two sourcing methods are available: traveling to make purchases, or ordering remotely. The model below describes how forward-looking traders will choose the optimal sourcing method, endogenously selecting into travel based on heterogenous product characteristics.

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

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 in response to small changes costs or benefits will be those who had smaller marginal gains from travel in the first place. Of course, these traders may also pursue different strategies when buying different products, although I show in the Appendix that a specification with trader-country-product fixed effects has similar results. Either explanation is consistent with the model I show in the following section, in which selection into travel is driven by differences in product types.

go through the traders to reach consumers – manufacturers cannot 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 with many differentiated product varieties, sourced from a single country. Later, in order to bring the model to data, I will extend consideration to many source countries and sectors.

4.1.1 Demand

Home country consumers are endowed with income E and demand many varieties from a fixed measure of differentiated goods, $i \in \Omega$. Each good has the potential for unlimited vertical improvements, as in the quality ladders model of [Grossman and Helpman \(1991\)](#). Improvements to the style or quality of the good arrive in increments of size g via a Poisson innovation process, for which the arrival rate λ_i may differ across goods. For the purposes of this model, the rate λ_i is taken as a feature of the type of good, rather than the outcome of endogenous investment in innovation. For instance, cotton athletic socks might have an inherently slow rate of change, while party dresses might be expected to have more frequent fashion changes. 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 to have more up-to-date styles of each good, and have a taste for variety. A representative consumer maximizes flow utility at time t 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 .¹³

¹³Normalized style is defined $\bar{z}_{it} = \ln\left(\frac{z_{it}}{E_t z_{it}^*}\right) + D$ where D is a nuisance constant that ensures the style measure is positive over the relevant range. This normalization delivers analytical and empirical tractability; I discuss in the Appendix how it simplifies the estimation and captures the key forces of a more general demand specification. Most crucially, the normalization reflects the fact that λ_i is a fixed aspect of the type of good. Consumers may prefer more recent versions of cotton socks to older ones, but they don't buy fewer socks simply because new styles of party dresses arrive more frequently. Without this normalization (or a more complex overall representation of demand, discussed in the Appendix), slow-changing goods (e.g. socks) would systematically fall out of the market, as they

4.2 Firm technology

4.2.1 Foreign producers

Firms in the source country produce all goods using a constant returns to scale technology with marginal cost c .¹⁴ Multiple firms produce each good.¹⁵

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 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 previously had the second most recent blueprint. As in [Grossman and Helpman \(1991\)](#), 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

Traders are 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 the frontier in the upstream market may continue to move forward.

Purchasing goods from foreign suppliers incurs a per-transaction fixed cost F_i and a per-unit cost of purchasing goods from a supplier 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).¹⁶ Thus, the total cost of sourcing a quantity x is $TC_i = c_i\tau x + F_i$.

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 overseas sourcing. Prospective traders draw these characteristics from a common

mechanically became less desirable relative to types of goods that change more frequently (e.g. dresses), which does not accord with common sense about demand for these products. I assume that consumers update their expectations about the status of the style arrival process each time traders restock, so that $\mathbb{E}_t z_{it} \geq z_{is} \forall t \geq s$ because what is available in the source market is known to be at least as good as what is available in the home market.

¹⁴Allowing for variation in the cost of production across goods, so that marginal cost for good i is $\frac{c}{\varphi_i}$, does not change any of the theoretical results or estimation (as I show in the Appendix). 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.

¹⁵It does not matter whether a given firm produces a single variety or multiple varieties.

¹⁶Traders may have other fixed costs of operation that are not purchase-specific (e.g. shop rental), which are irrelevant to the choice between traveling and ordering, but I assume they have no additional marginal costs.

underlying distribution on entering the market, and may exit immediately if they do not find it profitable to trade. .

Assume for the moment the existence of a constant aggregate price index P . In Section 4.4, I will derive the exact price index and show that a stationary equilibrium exists in which traders' expectations are correct. This allows me to derive demand for each good at time t , and variable profits as a function of prices and demand. Optimal pricing is a constant markup that does not depend directly on the style level or taste parameter.

$$q_{it} = \bar{z}_{it} \psi_i \left(\frac{p_{it}}{P} \right)^{-\sigma} \frac{E}{P} \quad (1)$$

$$\pi_{it} = \frac{1}{\sigma} p_{it} q_{it} \quad (2)$$

$$p_{it} = \left(\frac{\sigma}{\sigma - 1} \right) c_i \tau \quad \forall t \quad (3)$$

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 , $x_i(T) = \int_{t=s}^{s+T} q_{it} dt$. Demand falls as the style of goods 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 profits as a functions of unit cost and current style at time t , which are determined by traders' choices, while holding other parameters fixed: $\pi(\bar{z}_{it}, c_i)$.

4.3 Optimal sourcing

When sourcing goods from foreign producers, traders face search and contracting problems that affect the style and price of the goods. 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.

In this section, 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 a trader with given characteristics choosing to travel or choosing to order. Then, I solve for the choice to travel, order, or exit the market, depending on which extensive margin

choice 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 do business in person is assumed to completely solve information problems associated with distance between buyers and sellers. Optimal stocking with travel is therefore the same as optimal stocking 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 $F_i = F_i^{tr} + G$. F_i^{tr} includes explicit costs of travel such as plane fare, but may also represent idiosyncratic components such as the opportunity cost of the trader’s time, like or dislike for travel, and so on. G represents other fixed costs of purchasing and importing 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. have the most recent blueprint). Because there are always multiple sellers of the frontier style and they engage in Bertrand competition, the trader is able to buy any quantity x at a price equal to the marginal cost of production c .

Conditional on travel, the trader must choose how frequently to make purchases. On the one hand, more frequent purchases mean that goods held in stock will on average be less out of date. On the other, it means that the fixed cost of restocking will be incurred more frequently.¹⁷ A trader chooses the length of time T between purchases to maximize total average per period profits:

$$\max_T \left\{ \frac{1}{T} \left[\int_{t=0}^T \pi(\bar{z}_{it}, c) dt - F_i^{tr} - G \right] \right\} \quad (4)$$

All parameters stay the same from one restocking cycle to the next, and so maximizing average profits in one cycle is equivalent to maximizing lifetime profits by repeating the same optimal cycle.¹⁸

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} , familiar from the literature on inventory management (and the Baumol-Tobin cash demand model):

¹⁷I abstract from discounting and physical inventory costs; adding them does not change the nature of the problem.

¹⁸Because consumer expectations update after each restocking and the aggregate price index is stationary, optimal choices starting from any time t will be equivalent to the profit maximizing choice starting from time 0.

$$T_i^{tr} = \sqrt{\frac{2(F_i^{tr} + G)}{\lambda_i \psi_i \pi_z}} \quad (5)$$

where $\pi_z = \ln g \left(\frac{p}{P}\right)^{1-\sigma} E$. Intuitively, 5 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).¹⁹ 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}).

4.3.2 Optimal stocking with ordering

When ordering remotely from a foreign supplier, traders face both contracting and search frictions. The contracting friction arises from the lack of legal 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 (e.g. via the internet or existing contacts) 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. Once this fixed cost is paid, the trader can source goods, but can only access a style $z_{it}^{or} \leq z_{it}^*$ (where the superscript denotes ordering, as opposed to traveling) and pays a unit cost $\tilde{c}_i \geq c$. As I will show in this section, the differences between the style 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 $TC_i^{or} = G + \tilde{c}_i \tau x$.

Contracting

In contrast to most firm-level trade models, I do not assume that agreements between buyers and sellers are costlessly and automatically 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 the equilibrium path and buyers pay before receiving goods and must induce honest behavior from sellers. The solution takes the form of a “contracting premium”, analogous to an efficiency wage. Traders pay a premium over marginal cost, but cutting off the relationship if the supplier ever cheats. The premium must be large enough that the supplier would prefer to behave honestly and continue to receive the expected stream of future rents, rather than cheating once and never receiving that

¹⁹Including discounting or inventory costs would yield this same rule, with a larger denominator (representing the cost of waiting) leading to a shorter restocking period.

buyer’s business again. Buyers will pay a unit cost $\tilde{c}_i(T) = \frac{c}{\delta_i(T)}$ where the contracting premium, $\frac{1}{\delta_i(T)} \geq 1$, will differ based on the product type and depends on the endogenously chosen period of time between orders. In this section, I show the derivation of this equilibrium solution.

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. Traders buy from suppliers at discrete moments, with an endogenous period of length T between purchases. Recall that there are always multiple suppliers able to supply a given style of each good.²⁰ 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 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: 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.

The timing of the stage game played between a trader and potential supplier is as follows:

1. A trader decides whether or not to pay the fixed cost G to place an order. If he pays the fixed cost, he learns the identity of all suppliers who have the most recent blueprint.²¹
2. The trader makes a take-it-or-leave-it offer to a supplier to provide a quantity x at a unit cost \tilde{c} .
3. The supplier chooses to accept or reject the offer.
4. If the first supplier accepts the offer, the trader sends a payment $x\tilde{c}$.²² If the first supplier

²⁰Firms in the upstream market are identical ex ante – the only difference between them is that at a particular point in time they will (exogenously) have different vintages of blueprints. Therefore, the game below is driven only by moral hazard, and not learning about supplier types.

²¹Due to the search friction introduced in the following section, the most recent style that is observable when ordering remotely may not actually be the frontier style.

²²I focus consideration on equilibria in which the trader sends payment before goods are received. Other timing is possible – suppliers could send goods first and receive payment after, or traders could send partial payment as a deposit and the rest after receipt of goods. However, I ignore these possibilities for two reasons. First, given the assumptions, the trader has all the bargaining power, and so it would be difficult to construct a cooperative equilibrium that prevented the trader from simply taking the goods and never sending payment. Second, it is consistent with the data, which shows full prepayment by Lagos traders in XX% of observed transactions. 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.

- rejects the offer, the trader starts again from Step 2 with a new supplier at no additional cost.
5. After receiving payment, the supplier decides whether to behave honestly by sending quantity x to the trader, or to cheat by keeping the money and sending nothing.
 6. The trader either receiving the goods or does not, and therefore perfectly observes the choice made by the supplier in Step 5.²³
 7. Steps 1 - 6 are instantaneous; I abstract from the time involved in negotiations or shipping. After goods arrive, the trader sells down stocks in the home market and Nature moves over a period of length T , during which new styles arrive in the source market and buyer-seller relationships are broken up exogenously.

I focus on sub-game perfect equilibria of the repeated game in which the equilibrium path features honest behavior. 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 the supplier has ever cheated, the trader never buys from that supplier 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 never 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, he defaults to his outside option. The supplier accepts any offer, but cheats if the payment is $\tilde{c} < \tilde{c}^*$ and sends the goods if $\tilde{c} \geq \tilde{c}^*$.

Neither traders nor suppliers discount future profits per se, but perceive some probability less than one that they will do business with the same partner at the time of the next purchase, which causes them to effectively discount the expected stream of future surplus from the relationship. This uncertainty about whether future payouts will actually 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 broken up due to other factors taken as exogenous. Traders cannot commit to buy from the same supplier in the future, and will not do so if either event occurs.

The probability that a match-breaking event occurs will depend on the length of time between purchases, T . I assume that suppliers know the demand function faced by the trader, and therefore the amount of time T it will take the trader to sell a purchase of size x . This allows them to correctly anticipate the probability that the match will continue into the next stage, which is key in solving for the equilibrium offer \tilde{c}^* that induces good behavior. Recall that new styles arrive at Poisson rate

²³It is possible to imagine a game in which traders face some uncertainty about whether the supplier cheated, perhaps because non-delivery of goods could be blamed on the shipping company or seizure at the port. I abstract from this, because empirically only 2.4% of importers report that they ever paid for goods that did not arrive in the previous year.

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:

$$\delta_i(T) = e^{-\theta_2 T} \left[1 - (1 - \theta_1) \left(1 - e^{-\lambda_i T} \right) \right] \quad (6)$$

where $e^{-\theta_2 T}$ is the probability that the match has *not* been broken up exogenously, and $(1 - \theta_1) (1 - e^{-\lambda_i T})$ is the probability that at least one new style has arrived but the current supplier does not have the blueprint. Therefore, $\delta_i(T)$ is the total probability that a relationship is *not* broken up over a period of length T . 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 satisfies 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 after paying the fixed cost, when ordering the trader only observes the style that was on the frontier at time $t - \alpha$. This is 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 only observe what can be seen via 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 best style found will be α periods behind the frontier:

$$\max_T \left\{ \frac{1}{T} \left[\int_{t=0}^T \pi(\bar{z}_{it}^{or}, \tilde{c}(T)) dt - G \right] \right\} \quad (7)$$

where $\bar{z}_{it}^{or} = \ln \left(\frac{z_{i,s-\alpha}^*}{\mathbb{E}_t z_{it}^*} \right) + D$ and $\tilde{c}_i(T) = \frac{c}{\delta_i(T)}$.

The solution for T_i^{or} is implicitly defined by the first-order condition:

$$0 = \frac{\pi_c(\bar{z}_{it}^{or}, \tilde{c}(T))}{\bar{z}_{it}^{or}} \tilde{c}'(T) \ln g \left(\lambda_i \left(-\alpha - \frac{1}{2}T \right) + \frac{D}{\ln g} \right) - \frac{1}{2} \lambda_i \ln g \frac{\pi(\bar{z}_{it}^{or}, \tilde{c}(T))}{\bar{z}_{it}^{or}} + \frac{G}{T^2} \quad (8)$$

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 under relational contracting 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 a higher contracting premium, because opportunistic 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 at that cost 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} :

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

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

The trader chooses to trade if $\max\{\Pi_i^{tr}, \Pi_i^{or}\} \geq 0$ 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:

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

In order to consider a stationary equilibrium, in which the price index is constant – validating the perception of individual traders – I assume that the timing of initial market entry and therefore the timing of restocking periods is i.i.d. across the distribution of entrants. At any given time, the actual style distance of a variety can be replaced by its average over the restocking cycle, so that the aggregate stationary price index is:

$$P = \left(\int_{i \in \Omega} E_t[\bar{z}_{it}] \psi_i p_i^{1-\sigma} di \right)^{\frac{1}{1-\sigma}} \quad (9)$$

Note that the full measure of goods Ω includes firms that travel, order, and exit, and that those that exit make zero contribution to the price index.

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 the average style distance and changes in price.

To demonstrate the role of each margin, consider a decomposition of the effect of a change in 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 next effect of this 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.

4.5 Many countries and products

In order to bring the model to data in which consumers buy products from many sectors sourced from many countries, I embed the framework developed above in a simple demand specification. Let the utility of consumers in the home country at time t be Cobb-Douglas over composites of differentiated goods from sectors $j = 1, 2, \dots, J$ and source countries $k = 1, 2, \dots, K$ and a homogenous outside good denoted $j = 0$:

$$U = \prod_{j=0}^J \prod_{k=1}^K U_{kj}^{\varphi_{kj}}, \quad \sum_{j=0}^J \sum_{k=1}^K \varphi_{kj} = 1,$$

where U_{kj} is utility from consumption of goods in sector j from country k and φ_{kj} is the expenditure share for each sector and source country. Allowing for parameter variation across country-sector groups, the price index previously described is more properly the country-sector specific price index P_{kj} , and E_{kj} is the consumer expenditure on sector j consumption from country k implied by the Cobb Douglas share.

5 Estimation

In order to quantify search and contracting frictions and understand their effects on market structure and consumer welfare, I structurally estimate the model presented in the previous section. I estimate

thirteen parameters separately for goods from each source country and sector (denoted with 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 (McFadden (1989)). 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. A trader who sources shirts from both China and India is treated in the estimation as two traders dealing in two unique varieties. This implies that importing decisions are separable across countries. Given that the suppliers, travel, shipping, and so on are clearly country-specific, this is reasonable with respect to the decisions being modeled.

I estimate parameters for 11 country and sector combinations with a sufficiently large sample of traders, shown in Table 6; these account for 83% of total expenditure observed in the data. I exclude observations in which a co-owner lives in the source country (8.2% of the total), because a residual claimant is able to conduct transactions with suppliers in person. 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 beyond the scope of this paper. Second, I exclude traders who both traveled and ordered from a particular country during the study period (11.4% of the total).

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. 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. 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 and 2014.

Column (1) of Table 6 shows the values calculated for goods from each country in each sector. These range from a low of 6% for apparel from other West African countries to a high of 38% for toiletries and beauty products from Dubai. Although values should vary based on shipping distance and weight-to-value ratios, which I do not measure directly, 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 identified by the markup charged to consumers in the home country over marginal cost: $p = \frac{\sigma}{\sigma-1}c\tau$. Since p , c , and τ are all observed, 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 6, and range from highly substitutable apparel ($\sigma_{kj} = 6.2$) from West African neighbors to a low for apparel from China ($\sigma_{kj} = 2.16$), with most groups falling between 2 and 5. This is also 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 observed time since last purchase, 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:

$$Pr(\text{new supplier} \mid \text{new style}) = 1 - \theta_1 e^{-\theta_2 T_{tikj}}$$

$$Pr(\text{new supplier} \mid \text{old style}) = 1 - e^{-\theta_2 T_{tikj}}$$

If no new styles have arrived since the last purchase, a trader only switches to a new 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. The ML estimator of $\hat{\theta}_{1kj}$ and $\hat{\theta}_{2kj}$ maximizes the log-likelihood function:

$$\begin{aligned} \ln \mathcal{L}(\theta_{1kj}, \theta_{2kj} \mid \mathbf{1}_{tikj}^{sup}, \mathbf{1}_{tikj}^{prod}, T_{tikj}) &= \sum_{t=1}^M \sum_{i=1}^N \mathbf{1}_{tikj}^{prod} \mathbf{1}_{tikj}^{sup} \ln(1 - \theta_{1kj} e^{-\theta_{2kj} T_{tikj}}) + \mathbf{1}_{tikj}^{prod} (1 - \mathbf{1}_{tikj}^{sup}) \ln(\theta_{1kj} e^{-\theta_{2kj} T_{tikj}}) \\ &\quad + (1 - \mathbf{1}_{tikj}^{prod}) \mathbf{1}_{tikj}^{sup} \ln(1 - e^{-\theta_{2kj} T_{tikj}}) + (1 - \mathbf{1}_{tikj}^{prod}) (1 - \mathbf{1}_{tikj}^{sup}) \ln(e^{-\theta_{2kj} T_{tikj}}) \end{aligned}$$

where $\mathbf{1}_{tikj}^{sup}$ is an indicator function equal to one when the transaction involves switching suppliers and $\mathbf{1}_{tikj}^{prod}$ is an indicator equal to one when the transaction involves buying a new product style. I estimate at the purchase level, with the subscript t indicating that I observe multiple purchases per trader.

Estimates are shown in columns (3) and (4) of Table 6. Lower θ_1 and higher θ_2 imply that the risk of a supplier relationship breaking is greater, all else equal. 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 equilibrium selection into importing strategies.

5.3 Second stage estimation

In the second stage, I use simulated method of moments to estimate the remaining nine parameters for each country-sector. Given parameter estimates from the first step, parametric assumptions about the distribution of underlying product/trader heterogeneity, and a starting parameter vector, I simulate the profit maximizing decisions of each trader, and calculate a set of aggregate moments describing these decisions. I then iterate over new choices of parameters and select the parameter vector that minimizes the distance between the simulated aggregate 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).²⁴ Therefore, I begin by making 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:

$$\lambda_{ikj} \sim \ln N \left(\mu_{kj}^\lambda, \sigma_{kj}^\lambda \right)$$

$$\psi_{ikj} \sim \ln N \left(\mu_{kj}^\psi, \sigma_{kj}^\psi \right)$$

$$F_{ikj}^{tr} \sim \ln N \left(\mu_{kj}^F, \sigma_{kj}^F \right)$$

where μ_{kj}^ψ is set to zero for all k and j because the average is pinned down by parameters of the profit function described below. Note that while the distributions of trader characteristics are independent ex-ante, the characteristics of actually observed traders will be correlated due to endogenous exit of traders with unprofitable combinations of characteristics. All else equal, traders with faster changing product varieties (higher λ_i), those facing less demand for their products (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, so that in equilibrium, 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 with respect 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 “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 on ordering or traveling:

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

²⁴With 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 every trader to make this practically implementable.

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

where $\delta_{ikj} \left(T_{ikj}^{or} \right)$ is as defined in Equation (6), and $\pi_{0kj} = \left(\frac{\sigma_{kj}}{\sigma_{kj}-1} c_{kj} \tau_{kj} \right)^{1-\sigma_{kj}} E_{kj} P_{kj} (D + \ln g)$ and $\pi_z = \left(\frac{\sigma_{kj}}{\sigma_{kj}-1} c_{kj} \tau_{kj} \right)^{1-\sigma_{kj}} E_{kj} P_{kj} \ln g$ but can be treated as constants to be estimated for each country-sector group without separately identifying all the constituent parts. I then allow each 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 of monthly profits and the covariance between monthly profits and time between purchases for travelers and orderers (3), the mean and variance of the fraction of purchases involving new styles for travelers and orderers (4), and the overall fraction of traders who travel (1).

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

$$\hat{\Psi}_{kj} = \min(m_{kj,data} - m_{kj,sim}(\Psi_{kj}))J(m_{kj,data} - m_{kj,sim}(\Psi_{kj}))'$$

where $m_{kj,data}$ is the vector of empirical moments for country k and sector j and $m_{kj,sim}(\Psi_{kj})$ is the vector of simulated moments calculated at Ψ_{kj} . I use the diagonal of the variance-covariance matrix of the moments as the weighting matrix J rather than the optimal full variance-covariance matrix, due to concerns about bias raised in [Altonji and Segal \(1996\)](#). I use numerical methods to find gradients used in calculating the standard errors.²⁵

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. Table 7 shows the data moments used in the estimation and their simulated analogues, and Table 8 shows the estimated parameters for each country-sector group. The estimated search frictions imply that goods found via remote search are, on average, the equivalent of 2.5 months behind the frontier. The estimated contracting

²⁵Currently reported standard errors take the parameters from the first step as given—standard errors that account for noise from the first stage estimates are in process.

premiums imply that traders who choose to order pay an average cost premium of 11.9%. 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 (5) of Table 5 implies a premium of 14.9%, and the structural estimate is within the confidence interval.

To aid in interpretation, Tables 9 and 10 show three estimates with clear economic interpretations and useful comparisons to non-targeted moments available in the data. Table 9 shows the average number of new styles expected to arrive per year for goods that are traded (i.e. the average λ_i across traders who do not exit 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 (with new styles arriving a little over twice per year), followed by electronics, 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), and was not targeted by the estimation. 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 search, finding new styles more frequently implies higher λ_i .

The second half of Table 9 shows the average search friction by sector. The estimates imply that the search friction is largest for apparel (on average across countries), with an effective lag of almost four months in the goods found via remote search versus in-person search. The second largest search friction is in homewares (e.g. furniture, appliances) with an average remote search lag of about 2.5 months. 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 relative to homewares than the lag alone suggests because the rate of product evolution is higher – a four month lag in apparel means getting goods that are a full style arrival behind the frontier (dividing average arrivals per year by the lag), while the 2.5 month lag in homewares only implies goods that are the equivalent of 0.2 arrivals behind. 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 10 shows the average 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 close to the non-targeted reported costs for West

Africa, China, and the UAE. They are lower than reported costs in the United States and United Kingdom, where it is plausible that this reflects positive utility associated with the travel itself – for instance, many traders who buy in the United States or United Kingdom may be visiting family members at the same time. These differences reflect 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 perfect third-party contract enforcement were introduced. Individual firm responses will add up to changes in the aggregate price index (which individual firms will, in turn, respond to), and so I now need to separate out the price index that was previously imbedded in the estimated 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 $\hat{\pi}_z = \pi_z \hat{R}^{\sigma-1}$. Starting from some value of \hat{R} , I solve for firm choices under the counterfactual, calculate the implied price index and \hat{R} , and iterate until I find a fixed point for \hat{R} .

In the absence of both search and contracting frictions, welfare in the consumer goods sector increases by 29.2%. The gains from eliminating the search problem alone would be 16.3%, and the gains from contracting alone would be 9.0% – 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 available, 3) the set of varieties available, and 4) firm profits.

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. Surprisingly, 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.

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 problem directly.

6.1 Developed country expenditure and firm size

The availability of fixed cost strategies for addressing search and contracting problems, such as travel, immediately raise the question of how economies of scale influence the welfare effects of these frictions. The relevant scale in this context is not firm size, per se, but rather than revenue flow associated with trading a particular product from a particular country. Perhaps surprisingly, 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 (GDP per capita in 2014 was \$54,398 in the United States and \$3,203 in Nigeria), per capita spending on these goods in 2014 was approximately 17 times higher. In this section, I consider how moving to US-level expenditure changes the relative welfare losses due to search and contracting frictions.

In a standard [Melitz \(2003\)](#) type trade model with monopolistic competition and CES preferences, an increase in expenditure affects 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 quantity of demand at a given price increases relative to the fixed cost. This could be considered a firm size effect if each firm sells one variety, but is more generally due to the potential revenue stream associated with each variety. 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; second, some traders switch from ordering to traveling.

The losses due to search and contracting frictions will be smaller in a context in which more varieties are already imported via travel, or imported frequently enough that the equilibrium contracting premium is small. To quantify this intuition, 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 \$36,000 per year to almost \$500,000 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 29.2% welfare increase. At US levels of expenditure, the comparable gain is only 6.9%.

This is a stylized way of representing the relationship between income levels, firm size, and information costs (for instance, I hold other parameters fixed, and do not model possible differences in market structure or import penetration), but it provides useful insight into how the results presented in the previous section might generalize with respect to levels of economic development. When revenue streams are larger, firms will opt in to fixed cost 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 cost of moving goods affects market integration and gains from trade (summarized in [Donaldson \(2015\)](#)). One implication of this paper is that freedom 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 road quality or rail access – 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 92 out of 104 countries on the 2016 Henley and Partners Visa Restrictions Index, and 68 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, typically between a pair of countries. 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 restrictive air service regulations and visa policies have been considered, but primarily through the lens of gains from tourism or attracting multinational corporations. This paper suggests that their more important economic effects may actually come indirectly through the price and availability of goods, especially in developing countries.

To investigate the effect that air travel liberalization might have 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 to any cities in each country, at any frequency, capacity, and price they choose to set. The most recent ASA signed by Nigeria and China in 2014 allows for only one airline from each country to land in one city in the other (Guangzhou 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 counterfactual policy 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 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 15%, and welfare increases by a little over 1% across the entire consumer goods sector – a gain equivalent to a \$645 million increase in consumer spending.

6.3 Financial services that improve contracting

The immediate question that arises given the presence of a serious moral hazard problem is whether and how terms of payment have arisen to cope with this issue. 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 to motivate good behavior. 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 the remainder of payments once goods have been received than traders do to trust sellers to send goods once payment has been received.

Another option given a two-sided contracting problems is to use an 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, asking a friend to check goods for you, and having a long relationship with a supplier are all commonly mentioned). This is consistent with the little we know about trade finance from other contexts – for instance, [Antras and Foley \(2015\)](#) report that only 16.2% of transactions by a major U.S. poultry exporter with customers involve direct bank intermediation (via letter of credit or documentary collection). In contrast, over 40% of transactions in Antras and Foley are “open account” (i.e. full post-payment, after goods are received), which they argue reflects the quality of institutional contract enforcement in the receiving country.

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 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. It offers a secure payment service on wholesale orders from a limited set of suppliers based in China. The buyer makes a full payment up front, but the total amount is held in escrow until delivery of the goods is confirmed. The service costs 5% of the value of the transaction.

When a trader in the model purchases the secure payment service, a supplier that attempts to cheat by collecting payment without having sent goods gets “caught” with some probability. 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 70% of all traders buying from suppliers in China. As the perceived accuracy of enforcement declines, however, demand drops off steeply.

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 because they simply transfer the contracting problem onto another party; they 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.

7 Conclusion

In this paper, I show that search and contracting frictions in differentiated goods trade can have a large impact on welfare in developing countries. Using unique data that documents the process of firm-to-firm 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. New facts about the relationship between travel and importing – that travel is extremely costly, persistent across firms of varying experience, and associated with key differences in prices and products – are difficult to reconcile with full information conceptions of trade. To account for these facts, I build a parsimonious model of trade with search and contracting problems, and show why the frequency of trade matters, determining both the prices and vintage of goods available to consumers. 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 I obtain are specific to importing differentiated non-food consumer goods in Nigeria, the underlying mechanisms are general, and offer predictions about the features of markets in which search and contracting problems are likely to have a large effect on welfare and those in which they will not. This is a step toward unpacking the black box of trade costs, both within and between countries. In addition to the straightforward – that places with effective contract enforcement institutions and inexpensive access to effective information technology will suffer less from search and contracting problems – the model suggests a variety of less obvious relationships between market parameters and information costs. Search and contracting frictions will matter more for trade in products that change frequently, in markets where there is more “churning” in the sets of potential buyers and sellers (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), when firms are small (or the demand for each variety is low), and the costs of traveling are high. 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 \(2015\)](#)). On the one hand, a smaller set of large, stable traders should find it less costly to solve information problems (via regular communication, long relationships, and reputation and coordination mechanisms as in [Greif \(1993\)](#), [Rauch and Casella \(2003\)](#), and others). 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. Future work could evaluate the impacts of specific policies and services related to trade finance, information coordination (e.g. through trade fairs or international credit ratings), or reductions in travel cost.

References

- Ahn, JaeBin, Amit K. Khandelwal, and Shang-Jin Wei.** 2011. “The role of intermediaries in facilitating trade.” *Journal of International Economics*, 84(1): 73–85.
- Aker, Jenny C.** 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.
- Altonji, Joseph G., and Lewis M. Segal.** 1996. “Small-Sample Bias in GMM Estimation of Covariance Structures.” *Journal of Business & Economic Statistics*, 14(3): 353–366.
- 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.
- Arkolakis, Costas.** 2010. “Market Penetration Costs and the New Consumers Margin in International Trade.” *Journal of Political Economy*, 118(6): 1151–1199.
- Atkin, David, and Dave Donaldson.** 2015. “Who’s Getting Globalized? The Size and Implications of Intra-national Trade Costs.” National Bureau of Economic Research Working Paper 21439.
- Banerjee, Abhijit V., and Esther Duflo.** 2000. “Reputation Effects and the Limits of Contracting: A Study of the Indian Software Industry.” *The 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.
- Biglaiser, Gary.** 1993. “Middlemen as Experts.” *RAND Journal of Economics*, 24(2): 212–223.
- Chaney, Thomas.** 2008. “Distorted Gravity: The Intensive and Extensive Margins of International Trade.” *The American Economic Review*, 98(4): 1707–1721.
- Chaney, Thomas.** 2014. “The Network Structure of International Trade.” *The American Economic Review*, 104(11): 3600–3634.

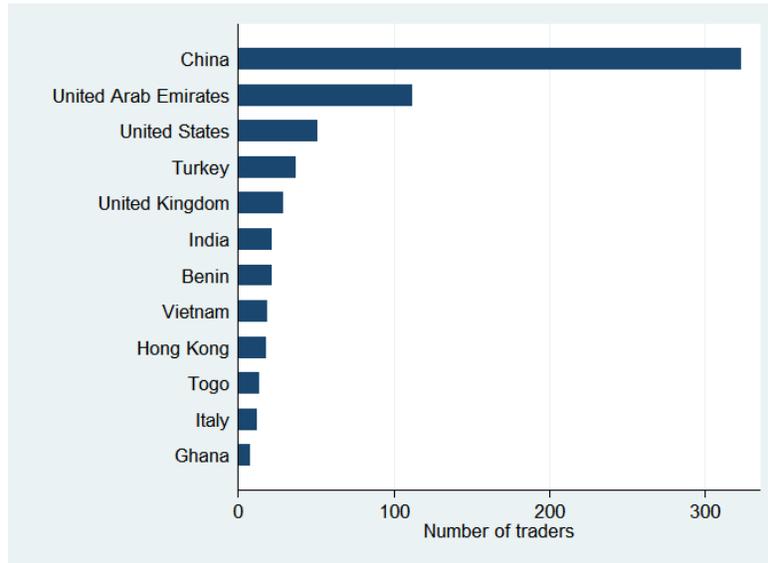
- Cristea, Anca D.** 2011. “Buyer-seller relationships in international trade: Evidence from U.S. States’ exports and business-class travel.” *Journal of International Economics*, 84(2): 207–220.
- Cristea, Anca D., David Hummels, and Brian Roberson.** 2012. “Estimating the Gains from Liberalizing Services Trade: The Case of Passenger Aviation.”
- Crozet, Matthieu, Guy Lalanne, and Sandra Poncet.** 2013. “Wholesalers in international trade.” *European Economic Review*, 58: 1–17.
- Dasgupta, Kunal, and Jordi Mondria.** 2012. “Quality Uncertainty and Intermediation in International Trade.” University of Toronto, Department of Economics Working Paper.
- Donaldson, Dave.** 2015. “The Gains from Market Integration.” *Annual Review of Economics*, 7(1): 619–647.
- Eslava, Marcela, James Tybout, David Jenkins, C. Krizan, and Jonathan Eaton.** 2015. “A Search and Learning Model of Export Dynamics.” Society for Economic Dynamics 2015 Meeting Paper 1535.
- Greif, Avner.** 1993. “Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders’ Coalition.” *The American Economic Review*, 83(3): 525–548.
- Grossman, Gene M., and Elhanan Helpman.** 1991. “Quality Ladders and Product Cycles.” *The Quarterly Journal of Economics*, 106(2): 557–586.
- Heise, Sebastian, Justin R. Pierce, Georg Schaur, and Peter K. Schott.** 2015. “Trade Policy and the Structure of Supply Chains.”
- Hummels, David.** 2007. “Transportation Costs and International Trade in the Second Era of Globalization.” *The 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.” *The Quarterly Journal of Economics*, 122(3): 879–924.
- Leke, Acha, Reinaldo Fiorini, Richard Dobbs, Fraser Thompson, Aliyu Suleiman, and David 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.” Social Science Research Network SSRN Scholarly Paper ID 1559654, Rochester, NY.

- Macchiavello, Rocco, and Ameet Morjaria.** 2015. “The Value of Relationships: Evidence from a Supply Shock to Kenyan Rose Exports.” *The American Economic Review*, 105(9): 2911–2945.
- MacLeod, W. Bentley.** 2007. “Reputations, Relationships, and Contract Enforcement.” *Journal of Economic Literature*, 45(3): 595–628.
- Mailath, George J., and Larry Samuelson.** 2006. *Repeated Games and Reputations: Long-Run Relationships*. Oxford University Press, USA.
- McFadden, Daniel.** 1989. “A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration.” *Econometrica*, 57(5): 995–1026.
- McMillan, John, and Christopher Woodruff.** 1999. “Interfirm Relationships and Informal Credit in Vietnam.” *The Quarterly Journal of Economics*, 114(4): 1285–1320.
- Melitz, Marc J.** 2003. “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity.” *Econometrica*, 71(6): 1695–1725.
- Monarch, Ryan, and Tim Schmidt-Eisenlohr.** 2016. “Learning and the Value of Relationships in International Trade.” CESifo Group Munich CESifo Working Paper Series 5724.
- Nigerian Ports Authority, Port Cargo Handling and Storage Services.** n.d.. “Cargo Throughput (Tonnes): 2009-2014.”
- Ossa, Ralph.** 2014. “Trade Wars and Trade Talks with Data.” *The American Economic Review*, 104(12): 4104–4146.
- Poole, Jennifer.** 2010. “Business travel as an input to international trade.” *UC Santa Cruz, mimeo*.
- Portes, Richard, and HÅ©IÅšne Rey.** 2005. “The determinants of cross-border equity flows.” *Journal of International Economics*, 65(2): 269–296.
- Rauch, James E., and Alessandra Casella.** 2003. “Overcoming Informational Barriers to International Resource Allocation: Prices and Ties*.” *The Economic Journal*, 113(484): 21–42.
- Rauch, James E., 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.” *The Quarterly Journal of Economics*, 98(4): 659–679.
- Steinwender, Claudia.** 2014. “Information Frictions and the Law of One Price: When the States and the Kingdom became United.”

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

Figure 1: Imports by source country

(a) Top 12 source countries by number of traders



(b) Top 12 source countries by value of imports

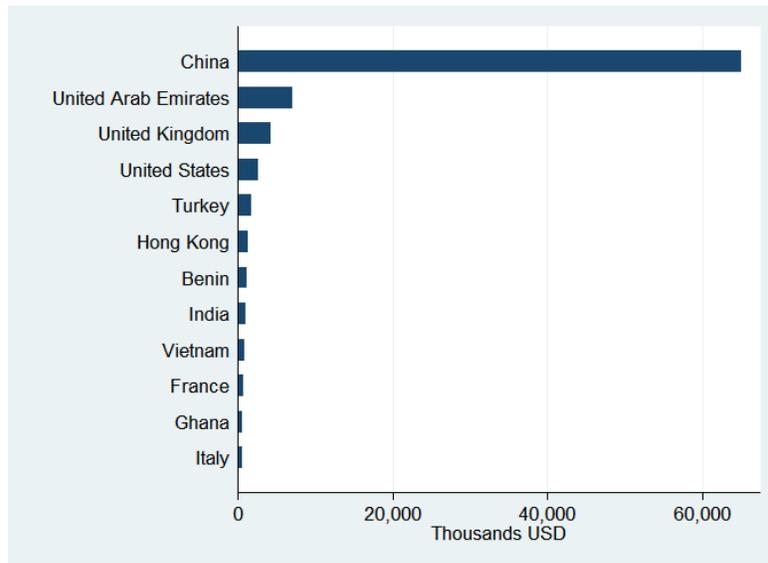


Figure 2: Likelihood of traveling and travel cost

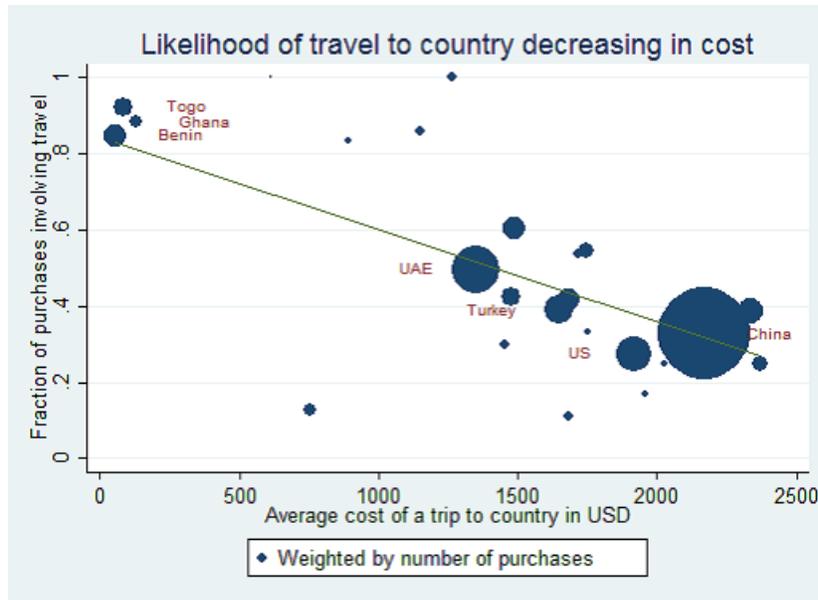


Figure 3: Travel, transport, and tariff expenditures relative to goods value

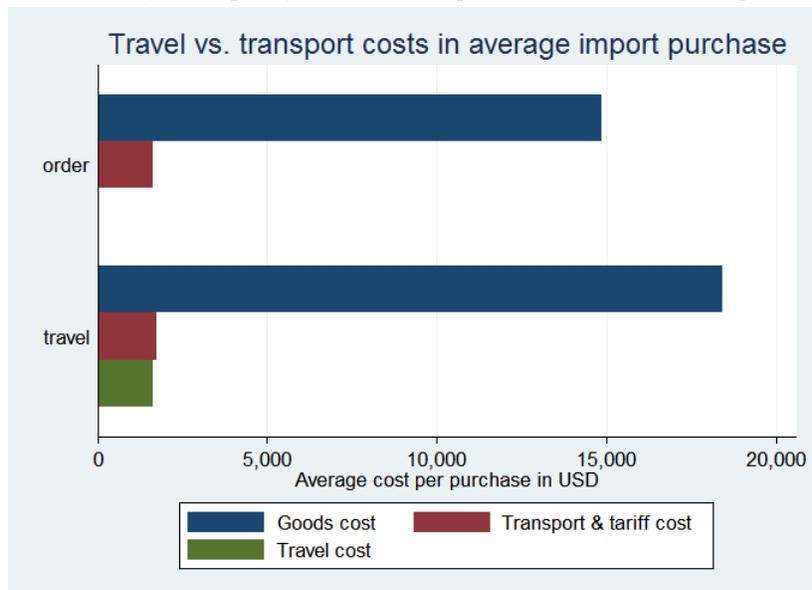


Figure 4: Changes in market structure without information frictions

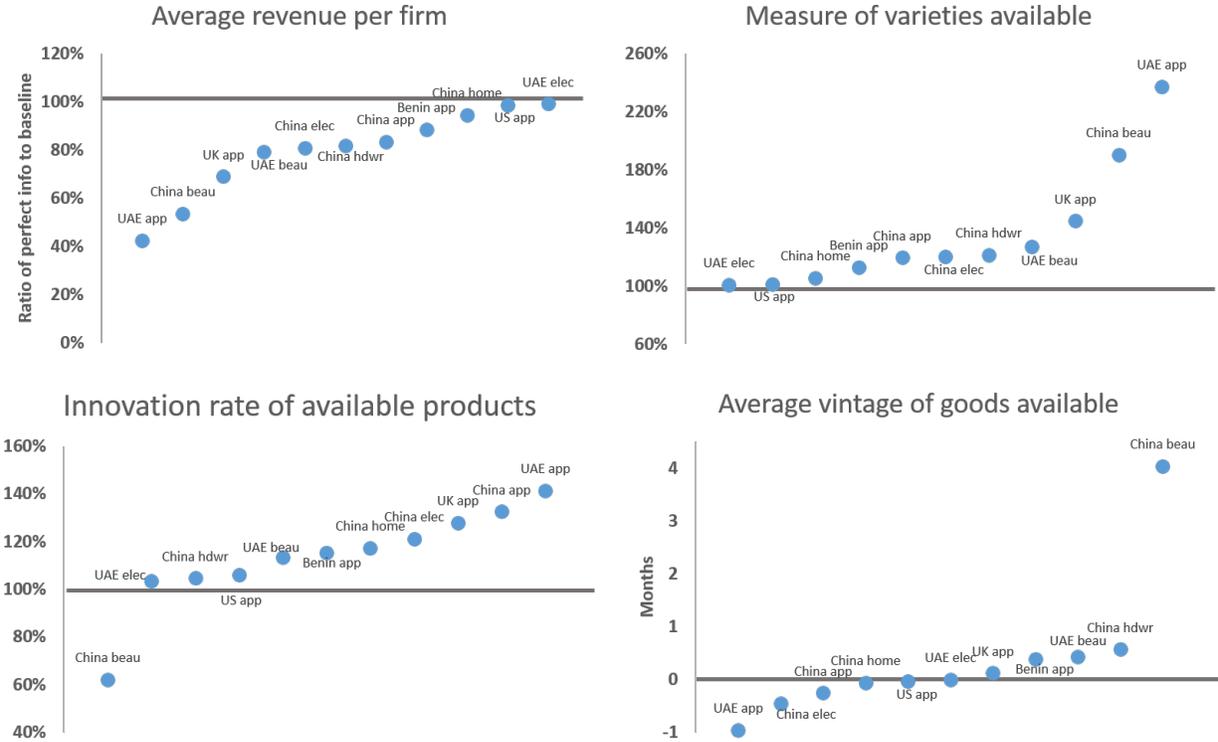


Figure 5: Demand for escrow service at prices of 5% and 1% of transaction value

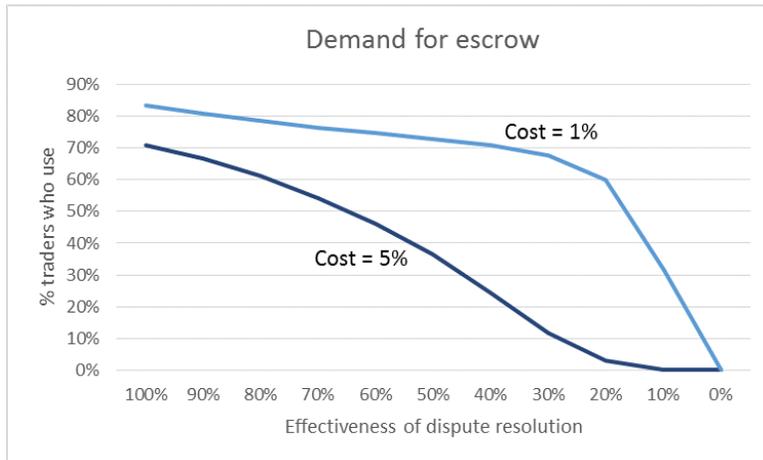


Table 1: Import and travel prevalence

	Imported	Traveled internationally
Never	44%	62%
Pre-2013 only	3%	5%
In 2013/2014	53%	33%

Notes: N=1,179 traders.

Table 2: Summary statistics about importing traders

	Mean	Median
<i>Trader</i>		
Male	74%	
Age	38.7	37
Schooling beyond secondary	54%	
Uses smartphone	88%	
<i>Business</i>		
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

Notes: N=620 traders.

Table 3: 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)
<i>Apparel</i>	Polo shirt, sandals, French lace	319	83%	8.63
<i>Electronics</i>	iPad, memory card, phone charger	68	76%	8.26
<i>Beauty</i>	Lipstick, hair relaxer, earring	61	75%	7.53
<i>Hardware</i>	Brake pad, electric cable, halogen bulb	50	35%	6.73
<i>Homewares</i>	Gas cooker, home theater, dining chair	59	59%	7.59
<i>Miscellaneous</i>	Diaper, umbrella, tennis ball	63	53%	7.55

Notes: N=620 traders.

Table 4: Probability of traveling when importing

Dependent variable:	(1)	(2)	(3)	(4)
	Traveled	Traveled	Traveled	Traveled
Business age	0.011*** (0.003)	0.006 (0.005)		
Years buying from country		0.010 (0.007)	-0.016** (0.007)	
Years buying from supplier		-0.004 (0.006)	0.005 (0.004)	0.004 (0.004)
Observations	3696	3035	3037	3213
Sector x country FEs	yes	yes		
Trader and country FEs			yes	
Trader x country FEs				yes
Means of independent variables				
Business age	10.64			
Years buying from country	5.37			
Years buying from supplier	3.88			

Notes: Observations are at the import purchase level, i.e. a product bought from a particular supplier located in another country in a given month. All columns are linear probability models, where the dependent variable “traveled” is a binary indicator for whether or not that particular purchase was conducted while the trader had traveled to the supplier’s location in a country other than Nigeria. Standard errors clustered at the trader level are shown in parentheses.

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

Table 5: Relationship between travel and transaction features

Panel A				
Dependent variable:	(1)	(2)	(3)	(4)
	Log unit cost	Markup	New style	New supplier
Traveled	-0.30** (0.119)	0.34*** (0.109)	0.07** (0.032)	0.05* (0.028)
Observations	2741	2614	3536	3354
Mean of outcome	1.90	2.07	0.51	0.20
Sector x country FEs	yes	yes	yes	yes
Trader x country FEs				
Panel B				
Dependent variable:	(5)	(6)	(7)	(8)
	Log unit cost	Markup	New style	New supplier
Traveled	-0.13** (0.054)	0.09* (0.048)	0.02 (0.044)	0.12*** (0.046)
Observations	2647	2513	3431	3259
Mean of outcome	1.90	2.07	0.51	0.20
Sector x country FEs				
Trader x country FEs	yes	yes	yes	yes

Notes: Observations are at the import purchase level, i.e. a product bought from a particular supplier located in another country in a given month. Columns 3, 4, 7, and 8 are linear probability models. Log unit cost is the price paid to a supplier abroad, and all currencies are converted to \$US using the mid-market exchange rate in the month in which the transaction took place. Markup is defined as the price at which the product was sold by the trader in Lagos divided by the unit cost. New style is a binary outcome equal to one if the product was not exactly the same as something the trader had previous purchased (including in color, size, etc). New supplier is a binary outcome equal to one if the transaction was the first time the trader had ever bought from a particular supplier. Columns 1 - 4 include sector by source country fixed effects (e.g. apparel from China), while columns 5 - 8 include trader by country fixed effects (all traders deal in products within a single sector, so this is equivalent to trader by country by sector fixed effects). Standard errors clustered at the trader level are shown in parentheses.

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

Table 6: 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	1.17	2.16	0.92	0.014
China	Electronics	1.14	3.01	0.95	0.005
China	Beauty	1.27	4.53	1.00	0.030
China	Hardware	1.24	3.90	0.98	0.018
China	Homewares	1.25	3.43	0.99	0.010
UAE	Apparel	1.25	2.37	0.95	0.014
UAE	Electronics	1.09	2.40	1.00	0.042
UAE	Beauty	1.38	2.91	0.75	0.010
UK	Apparel	1.19	4.56	0.83	0.046
US	Apparel	1.31	2.29	0.95	0.003
West Africa	Apparel	1.05	6.21	0.91	0.011

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

Table 7: Data moments and simulated moments

Moment:	Fraction travel		Probability of new style		Variable profit / month		Months between shipments							
	All		Order		Order		Order							
	Data	Sim	Data	Sim	Data	Sim	Data	Sim						
<i>China</i>	0.40	0.40	0.64	0.43	0.58	0.58	\$1,738	\$1,445	\$2,586	\$2,719	7.21	6.09	9.42	8.74
<i>China</i>	<i>0.32</i>	<i>0.37</i>	<i>0.34</i>	<i>0.42</i>	<i>0.51</i>	<i>0.55</i>	<i>\$525</i>	<i>\$449</i>	<i>\$690</i>	<i>\$837</i>	<i>11.12</i>	<i>10.28</i>	<i>11.37</i>	<i>11.86</i>
<i>China</i>	0.25	0.36	0.26	0.16	0.21	0.35	\$635	\$608	\$687	\$853	10.05	4.30	14.40	26.71
<i>China</i>	<i>0.16</i>	<i>0.23</i>	<i>0.18</i>	<i>0.23</i>	<i>0.36</i>	<i>0.34</i>	<i>\$909</i>	<i>\$1,602</i>	<i>\$1,890</i>	<i>\$1,778</i>	<i>7.90</i>	<i>5.28</i>	<i>15.00</i>	<i>14.03</i>
<i>China</i>	0.21	0.22	0.27	0.30	0.50	0.49	\$1,584	\$585	\$1,010	\$891	6.86	6.27	18.00	23.08
<i>UAE</i>	<i>0.53</i>	<i>0.55</i>	<i>0.48</i>	<i>0.41</i>	<i>0.72</i>	<i>0.74</i>	<i>\$179</i>	<i>\$186</i>	<i>\$868</i>	<i>\$647</i>	<i>12.00</i>	<i>9.65</i>	<i>11.76</i>	<i>11.38</i>
<i>UAE</i>	0.57	0.59	0.39	0.33	0.49	0.47	\$273	\$1,256	\$1,189	\$1,531	13.09	2.64	12.00	10.76
<i>UAE</i>	<i>0.39</i>	<i>0.44</i>	<i>0.30</i>	<i>0.30</i>	<i>0.50</i>	<i>0.47</i>	<i>\$440</i>	<i>\$338</i>	<i>\$198</i>	<i>\$391</i>	<i>11.48</i>	<i>9.83</i>	<i>14.00</i>	<i>19.67</i>
<i>UK</i>	0.78	0.83	0.60	0.32	0.70	0.76	\$36	\$77	\$325	\$353	0.98	2.45	15.27	12.65
<i>USA</i>	<i>0.65</i>	<i>0.65</i>	<i>0.68</i>	<i>0.66</i>	<i>0.80</i>	<i>0.81</i>	<i>\$711</i>	<i>\$656</i>	<i>\$1,391</i>	<i>\$1,427</i>	<i>7.30</i>	<i>7.79</i>	<i>16.42</i>	<i>14.16</i>
<i>W. Africa</i>	0.86	0.95	0.25	0.37	0.81	0.38	\$658	\$379	\$2,253	\$2,370	1.78	1.61	2.45	3.26

Notes: This table shows seven of the twelve moments matched in the estimation. The remaining five are the variance of profits and months between shipments for orders and travel (4) and the covariance between profits and the months between shipments (1).

Table 8: Second stage parameters estimates

Parameter	π_0	π_z	α	μ_λ	σ_λ	G	μ_F	σ_F	σ_ψ	J-statistic
Definition	Max. profit (log \$US/month)	Deriv. profit wrt style (log)	Search lag (months)	Style arrival (log)	SD style arrival (log)	Order cost (\$US)	Travel cost (log \$US/)	SD travel cost (log \$US/trip)	Taste shifter	
Country	Sector									
China	Apparel	7.15 (0.22)	6.67 (0.21)	1.59 (1.41)	-1.42 (0.26)	0.15 (0.12)	8.48 (1.01)	0.24 (1.42)	-0.01 (0.05)	27.02
China	Electronics	5.83 (0.01)	5.79 (0.06)	1.64 (0.05)	-0.92 (0.08)	0.08 (0.00)	7.81 (0.02)	-0.23 (0.01)	-0.20 (0.01)	25.12
China	Beauty	3.71 (0.03)	4.18 (0.11)	1.99 (0.15)	-2.41 (0.16)	0.00 (0.03)	9.13 (0.25)	0.00 (0.17)	0.43 (0.01)	207.11
China	Hardware	6.84 (0.54)	7.04 (0.63)	0.61 (4.59)	-2.69 (0.68)	-0.08 (0.44)	10.61 (2.95)	0.99 (0.07)	0.04 (0.08)	451.52
China	Homewares	6.24 (0.08)	5.52 (0.53)	2.37 (1.15)	-2.75 (0.47)	0.02 (0.05)	8.61 (0.09)	0.39 (0.48)	-0.07 (0.05)	32.32
UAE	Apparel	5.15 (0.14)	4.46 (0.26)	10.39 (0.34)	-1.37 (0.32)	-0.25 (0.09)	8.71 (0.15)	0.14 (0.09)	0.06 (0.03)	5.68
UAE	Electronics	7.10 (0.15)	6.14 (0.82)	0.48 (2.48)	-2.28 (0.77)	0.05 (0.30)	7.04 (0.15)	0.59 (1.98)	-0.26 (0.04)	2589.37
UAE	Beauty	5.28 (0.30)	5.17 (0.74)	1.40 (0.85)	-2.51 (0.81)	0.04 (0.04)	8.81 (1.23)	1.22 (0.44)	0.02 (0.08)	40.28
UK	Apparel	4.39 (0.18)	3.59 (0.85)	4.17 (1.31)	-0.22 (0.94)	0.21 (0.13)	6.84 (0.20)	-1.05 (0.55)	0.15 (0.05)	118.81
US	Apparel	6.91 (0.15)	5.05 (0.30)	3.05 (1.31)	-1.84 (0.37)	0.13 (0.13)	7.08 (0.17)	-0.73 (0.24)	-0.27 (0.03)	18.60
W. Africa	Apparel	6.00 (0.02)	5.20 (0.15)	0.25 (0.28)	-1.00 (0.18)	0.30 (0.23)	3.80 (0.55)	0.00 (0.21)	0.60 (0.00)	453.21

Notes: Standard errors are shown in parentheses below each parameter estimate. Current standard errors take first stage estimates as given. The μ and σ estimates are the location and scale parameters of log normal distributions, not the mean and variance of the distribution.

Table 9: Second stage estimates by sector

Innovation rates – Expected improvement arrivals per year					
Sector	<i>Apparel</i>	<i>Electronics</i>	<i>Homewares</i>	<i>Hardware</i>	<i>Beauty</i>
Estimate	3.38	1.49	0.97	0.78	0.59
Non-targeted data (rank)	1	3	4	5	2
Search friction – Months behind frontier					
Estimate	3.89	1.06	2.37	0.61	1.70

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, i.e. apparel styles change the most frequently, hardware the least frequently

Table 10: Second stage estimates by country

Travel cost – \$US per trip					
Country	<i>China</i>	<i>UAE</i>	<i>UK</i>	<i>US</i>	<i>W. Africa</i>
Estimate	\$1,736	\$1,382	\$850	\$1,036	\$59
Non-targeted data	\$2,154	\$1,370	\$2,326	\$2,174	\$66

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 begun.