Who's Getting Globalized?

The Size and Implications of Intra-national Trade Costs*

David Atkin† and Dave Donaldson‡

July 2015

Abstract

How large are the intra-national trade costs that separate consumers in remote locations of developing countries from global markets? What do those barriers imply for the intra-national incidence of the gains from falling international trade barriers? We develop a new methodology for answering these questions and apply it to newly collected CPI micro-data from Ethiopia and Nigeria (as well as to the US). In order to overcome three well-known challenges that arise when using price gaps to estimate trade costs, we: (i) work exclusively with a sample of goods that are identified at the barcode-level (to mitigate bias due to unobserved quality differences over space); (ii) collect novel data on the origin location of each product in our sample (to focus only on the pairs of locations that actually identify trade costs); and (iii) use estimates of cost pass-through to correct for mark-ups that potentially vary over space (to extract trade costs from price variation in an environment with potentially oligopolistic intermediaries). Without these corrections, we find that our estimates of the cost of distance would be biased downwards by a factor of approximately four. Our preferred estimates imply that the effect of log distance on trade costs within Ethiopia or Nigeria is four to five times larger than in the US. We also use our pass-through estimates to calculate the incidence of surplus increases due to falling world prices. We find that intermediaries capture the majority of the surplus, and that their share is even higher in distant locations, suggesting that remote consumers see only a small part of the gains from falling international trade barriers.

*We thank Rohit Naimpally, Guo Xu, Fatima Aqeel and Max Perez Leon for excellent research assistance, and Alvaro González, Leonardo Iacovone, Clement Imbert, Horacio Larreguy, Philip Osafo-Kwaako, John Papp, and the World Bank Making Markets Work for the Poor Initiative for assistance in obtaining segments of the data. We have benefited greatly from many discussions with Glen Weyl, as well as conversations with Treb Allen, Pol Antràs, Ariel Burstein, Arnaud Costinot, Michal Fabinger, Penny Goldberg, Seema Jayachandran, Marc Melitz, and David Weinstein and from comments made by numerous seminar participants. Finally, we thank the International Growth Centre in London for their generous financial support and the Kilts-Nielsen Data Center at The University of Chicago Booth School of Business for making available to us their Nielsen Consumer Panel Data.

†MIT and NBER. E-mail: atkin@mit.edu

‡Stanford and NBER. E-mail: ddonald@stanford.edu

1 Introduction

Recent decades have seen substantial reductions in the barriers that impede trade between nations, a process commonly referred to as "globalization". But trade does not start or stop at national borders. The trading frictions faced by many firms and households include not only the international trade costs that have fallen in recent times, but also the *intra*-national trade costs that separate these agents from their nearest port or border. Many commentators have argued that these intra-national trade costs are especially large in developing countries, potentially limiting the gains from globalization for remote regions (see, for example, WTO (2004)). Yet we lack rigorous estimates of the size and implications of these costs, particularly in data-scarce regions of the world such as sub-Saharan Africa. In this paper we develop a new methodology for estimating trade costs and apply it to newly collected micro-data from Ethiopia and Nigeria (as well as to the United States, for purposes of comparison). In addition, we explore the implications of our estimates for the geographic incidence of globalization within these countries.

To fix ideas, consider a product that is imported from abroad. Suppose this product enters a country through port of origin o where it sells to domestic traders at the wholesale price P_o. These traders then sell the product at a destination location d for P_d, where these prices reflect the identity

$$P_d - P_o = \tau(\mathbf{X}_{od}) + \mu_d. \tag{1}$$

This equation states that the spatial price gap $P_d - P_o$ reflects both the intra-national trade costs $\tau(\mathbf{X}_{od})$ over a route that has cost-shifting characteristics $\mathbf{X}_{od}$ (such as distance) as well as the mark-up μ_d charged by traders. In common with a voluminous existing literature, we seek to estimate how $\tau(\mathbf{X}_{od})$ depends on $\mathbf{X}_{od}$ by drawing inferences from the equilibrium distribution of prices over space. But we make progress with respect to this literature by using new data and new tools to overcome three well-known challenges that plague such inferences:

1. *Spatial price gaps may reflect differences in unobserved product characteristics (such as quality) across locations.* Clearly one cannot hope to apply equation (1) by making comparisons across non-identical products. In recognition of this point, and following the pioneering work of Broda and Weinstein (2008) for the US, we have compiled what we believe to be the first dataset on the geography of prices of products defined at a barcode-equivalent level within a developing country.

2. *Spatial price gaps are only rarely directly informative of trade costs.* It is standard in the literature to assume that trading is perfectly competitive ($\mu_d = 0$).[1] Under this assumption—which we relax shortly—equation (1) states that price gaps identify trade costs: $P_d - P_o = \tau(\mathbf{X}_{od})$. But this method is only applicable when the researcher knows which locations are origins and destinations. Our paper provides the first widespread attempt to learn the precise origin locations for each product in our sample. Failure to incorporate this new information

[1] As we discuss further below, a related and commonly used assumption is that preferences and market structure belong to the special case in which mark-ups are positive but do not vary across locations and hence mark-ups do not bias estimates of how $\tau(\mathbf{X}_{od})$ depends on $\mathbf{X}_{od}$.

would cause us to underestimate intra-national trade costs by a factor of two in Nigeria and Ethiopia.

3. *Spatial price gaps may reflect varying mark-ups across locations as well as trade costs.* If $\mu_d \neq 0$ then, as equation (1) makes clear, spatial price gaps $P_d - P_o$ cannot be used to identify how $\tau(\mathbf{X}_{od})$ depends on $\mathbf{X}_{od}$ because unobserved mark-ups may also depend on $\mathbf{X}_{od}$ (for example, if markups respond to marginal costs which are in turn a function of trade costs).[2] However, we demonstrate that, in a general model of oligopolistic trading, estimates of the extent to which shocks to P_o pass through into P_d act as a sufficient statistic for how mark-ups respond to trade costs or any other cost-shifter. The pass-through rate, which we estimate separately for each location and product in our sample, therefore allows us to purge price variation of mark-up variation in a flexible manner, without having to estimate demand or supply relations. Applying this correction we find that existing approaches would underestimate trade costs by an additional factor of two. This finding has the auxiliary implication that mark-ups are lower in relatively remote locations in our sample countries. Our results suggest that this is predominantly because demand curves, for nearly all products and locations in our sample, are more elastic at higher prices. Since remote locations pay higher prices due to trade costs, intermediaries choose to charge lower markups in these locations (despite remote locations appearing to be less competitive).

In summary, we estimate that the cost of distance within our two African sample countries is approximately four to five times larger than in the US. In addition, we find that this relatively high burden of distance in our African sample countries remains once we adjust our distance metric for the availability and quality of roads. Reassuringly, our results on the costs of intra-national trade in Africa relative to the US line up with direct evidence from trucker surveys in Teravaninthorn and Raballand (2009).

The preceding discussion has centered on our new method for estimating $\tau(\mathbf{X}_{od})$ but two additional results follow, both concerning how the surplus created by international trade varies across locations within a country. Remote locations, especially those in sub-Saharan Africa, pay high trade costs. This naturally implies that, all else equal, remote locations enjoy less surplus from foreign goods. Our first result provides evidence consistent with this by demonstrating that the probability of a product not being found by price enumerators is higher in remote locations. In a second result, we estimate the relative shares of surplus that would accrue to consumers versus traders following a change in the port price such as would be caused by a change in inter-national trade barriers. That is, we estimate the incidence of a global price change. Using an extension of the analysis in Weyl and Fabinger (2011), we show how the pass-through rate in a market can be used to construct a sufficient statistic for the distribution of surplus in that market. Our estimates imply that the incidence of the gains from globalization is skewed in favor of intermediaries and

[2] The distinction between marginal costs and mark-ups is important here, even beyond the usual reasons grounded in policy and distributional consequences, since the incidence of global price changes hinges on the relative importance of marginal costs and mark-ups in intra-national trade.

deadweight loss (relative to consumers), and increasingly so in remote locations. This is to be expected if, as our data suggest, remote locations are less competitive.

A notable feature of the above results is that they require no data on the *quantities* of products traded or sold—data that would not be available for narrowly-defined products in most developing countries. Nor do they require us to estimate price elasticities or mark-ups, functions of first-order derivatives which are difficult to estimate even with quantity data. Instead, the pass-through rate, a substantially easier object to estimate, uncovers the key second-order demand derivatives that determine both how mark-ups vary and how surplus is distributed over space.[3] The need to understand intra-national trade costs in the absence of such quantity data is a primary motivation for the methodologies we propose in this paper.

The work in this paper relates to a number of different literatures. First, this paper complements a recent literature that extends models of international trade to accommodate intra-national trade costs (see, for example, Ramondo, Rodriguez-Clare, and Saborio-Rodriguez (2012), Redding (2012), Du, Wei, and Xie (2013), Agnosteva, Anderson, and Yotov (2014), Fajgelbaum and Redding (2014), and Cosar and Fajgelbaum (2015)). The implications of these models depend crucially on the size of intra-national trade costs, and we provide a methodology for estimating such costs using spatial price gaps.

Since our methodology is applicable to the measurement of trade costs more generally, the paper relates to a voluminous literature surveyed in Fackler and Goodwin (2001) and Anderson and van Wincoop (2004) that uses spatial price dispersion in order to identify trade costs. Various segments of the literature have dealt with each of the three previously described challenges in isolation, although we believe that our work is unique in attempting to circumvent all of them together. In terms of the first challenge, Broda and Weinstein (2008), Burstein and Jaimovich (2009) and Li, Gopinath, Gourinchas, and Hsieh (2011) draw on proprietary retailer or consumer scanner datasets from the US and Canada in order to compare prices of extremely narrowly identified goods (that is, goods with unique barcodes) across space. In terms of the second challenge, Eaton and Kortum (2002), Simonovska (forthcoming), Donaldson (forthcoming) and Simonovska and Waugh (2014) argue that spatial arbitrage is free to enter and hence that spatial price gaps place lower bounds on the costs of trade, where these lower bounds are binding among pairs of locations that do trade. A central obstacle in this literature has been the need to work with narrowly defined products and yet also know which location pairs are actually trading those narrowly defined products. Our approach exploits unique data on the location of production of each product in our sample that allows us to do both. In terms of the third challenge, Feenstra (1989), Goldberg and Knetter (1997), Atkeson and Burstein (2008), Goldberg and Hellerstein (2008), Burstein and Jaimovich (2009), and Li, Gopinath, Gourinchas, and Hsieh (2011) consider, as we do, the possibility that producers or intermediaries have market power and hence that firms may price to market. In particular, this literature has placed emphasis on the extent of exchange rate pass-through and

[3] In contrast to price elasticities that require both price and quantity data, the estimation of pass-through rates requires only price data. Additionally, estimating either typically requires exogenous cost shocks which directly identify pass-through rates but require instrumental variables techniques (with their well known finite-sample shortcomings) to identify price elasticities.

its implications for estimating market power. We instead apply a similar logic to the market for each product and location in our sample with the goal being to infer how intermediaries' market power and equilibrium mark-ups vary across locations, as well as how variable mark-ups over space cloud inference about how the costs of trading vary over space.

The paper also relates to work that considers the distribution of gains from trade in the presence of markups and intermediation. A recent literature explores the interaction between the gains from trade and variable mark-ups, although with a focus that is very much on producers rather than intermediaries with market power. (See, for example, Melitz and Ottaviano (2008), Feenstra and Weinstein (2010), Edmond, Midrigan, and Xu (2011), Arkolakis, Costinot, Donaldson, and Rodriguez-Clare (2012), De Loecker, Goldberg, Khandelwal, and Pavcnik (2012), and Cosar, Gieco, and Tintelnot (forthcoming)). The growing literature on intermediation in trade includes Chau, Goto, and Kanbur (2009), Ahn, Khandelwal, and Wei (2011), Antràs and Costinot (2011), and Bardhan, Mookherjee, and Tsumagari (2013). This work aims to understand when trade is conducted via intermediaries rather than by producers directly. Our work addresses the consequences of intermediation—by traders who potentially possess market power—for the magnitude of intra-national barriers to trade and the incidence of globalization.

Finally, the elasticity of the slope of inverse demand plays a key role in determining the variation in markups across space and the incidence of globalization. In the paper we show that for almost every location and product, consumer demands become more elastic at higher prices. Recent work has highlighted the importance of this elasticity in determining the impacts of trade in the presence of imperfect competition (see, for example, Zhelobodko, Kokovin, Parenti, and Thisse (2012), Neary and Mrazova (2013) and Mayer, Melitz, and Ottaviano (2014)). Given the paucity of empirical estimates of this critical elasticity, an additional contribution of this paper is to provide guidance on realistic parameter values.

The remainder of this paper proceeds as follows. Section 2 describes the new dataset that we have constructed for the purposes of measuring and understanding intra-national trade costs in our sample of developing countries. Section 3 outlines a theoretical framework in which intra-national trade is carried out by intermediaries who potentially enjoy market power, as well as how we use this framework to inform empirical work that aims to estimate the size of intra-national trade costs. Section 4 discusses the empirical implementation of this methodology and presents our findings. Section 5 describes how pass-through rates provide a sufficient statistic for identifying the distribution of the gains from trade between consumers and intermediaries and implements the procedure. Section 6 concludes.

2 Data

The above Introduction details three challenges faced by researchers hoping to uncover trade costs from spatial price gaps. A core component of this paper is the creation of a dataset that allows us to overcome these challenges. The methodology we propose requires data of two types. First, we require retail price data that document the price of narrowly defined (i.e. barcode-level) products, at many points in space within a group of developing and developed countries, at high

frequency over multiple years. Second, we need to know the location at which each product in each country in our sample is produced or imported. Here we briefly describe these two types of data and their construction. Appendix A provides more detail.

2.1 Retail price data

We work with the CPI microdata from two sub-Saharan African (SSA) countries—Ethiopia and Nigeria—because of their large geographic sizes and because they were particularly forthcoming in making price data available to researchers. Enumerators in these countries visit pre-specified sample outlet locations within a market town or city, typically many times per month, and obtain price quotes for a pre-specified list of precisely defined products. We obtained data in digital form spanning the period from September 2001 to June 2010 in the case of Ethiopia and January 2001 to July 2010 in the case of Nigeria.

Both Ethiopia and Nigeria report a CPI that is based on price observations in urban areas. In Ethiopia we obtained a sample of 103 urban market places and in Nigeria we obtained a sample of 36 state capitals (one for each state). These locations are shown on the maps of Ethiopia and Nigeria depicted in panels A and B of Figure 1, respectively. (These maps also depict major and minor roads, as discussed in Section 4, as well as the production locations for each product in our sample, as discussed in Section 2.2.)

Both Ethiopia and Nigeria base their CPI on a set of products that did not change substantially during our sample period. These products are designed to span the typical consumption basket. Of the many products that are covered, the vast majority refer to activities—such as a "man's haircut"—or goods—such as "rice"—whose very nature means that the products cannot be precisely codified. Because concerns of spatially-varying unobserved quality differences have appeared prominently in the literature (see, for example, Broda and Weinstein (2008)), we work instead with the sub-sample of products that we consider to be particularly narrowly defined. In practice, this involved a restriction to brandname products with detailed product descriptions. Because these descriptions appear to be as precise as those linked to unique barcodes in US data, we refer to these products as products that are defined at the barcode level. Note that, while the original sample of products in our sample is designed to be representative of consumer spending, our restriction to a sample of products with brandnames is not likely to be representative in this regard. However, to the extent that the technology used to trade important barcode-level products in the CPI basket is similar to that used to trade other products, the resulting sample will still be representative of the cost of trading goods within Ethiopia and Nigeria. The resulting sample contains 15 products in Ethiopia and 19 products in Nigeria that were broadly available across both the locations and years in our sample (examples of which include "Titus Sardines (125 g)", "Bedele Beer (300 cc)" and "Lux Toilet Soap (90 g)"; a full list is given in Appendix Table C.1).

In order to provide a basis of comparison for our Ethiopian and Nigerian estimates, we seek similar estimates for the United States. Following Broda and Weinstein (2008), we use data from the Nielsen Consumer Panel (NCP) due to its extensive geographic and product coverage. The NCP incentivizes sample households to use hand-held barcode scanners to scan all products pur-

chased by the household and enter the price that they paid for each product. From the resulting price observations we use each household's county of residence to aggregate up to a dataset that contains the average price paid, in each of the 2,856 counties and each month in 2004-2009, for the 1.4 million unique barcodes purchased by NCP households. In order to obtain a sample similar in nature to the SSA samples, we work with a relatively small sub-sample of barcodes that are the leading product in 230 Nielsen "product modules" (examples of which include frozen pot pies or chilli sauce).

Our main analysis uses a cleaned sample of price data, obtained by applying a simple cleaning algorithm to the raw data. However, we also report (in Table 3) results obtained from the original, uncleaned data set; these are similar in terms of both the magnitudes and statistical significance of the resulting estimates. In addition, in order to arrive at estimates of trade costs that are in real terms, we require all prices to be similarly expressed in terms that correspond to a common year (which we take to be 2001) and a common currency unit (US dollars). We therefore apply a simple correction, based on inflation rates calculated from prices at origin locations and the prevailing bilateral exchange rates, to all prices used in the analysis.[4]

2.2 Production origin location data

To identify origin locations, we conducted telephone interviews with the firms that produce each product, asking for the precise location(s) of production that serve markets in each country in each year. For the case of imported products we have contacted distributors to learn the port of entry of each imported product in each country (and year) in our sample. From these two sources we obtain the latitude and longitude of the production location(s) or port of entry for every good in our sample.

For Ethiopia and Nigeria, we were able to locate the factory that produced every product in our sample.[5] The US posed additional difficulties since firms were less willing to respond for confidentiality reasons, and in many cases could not easily find out where a particular barcode was made. Of the 230 leading products, we were able to successfully find factory locations for 88 products.[6]

In a handful of cases the product was reported to be produced at multiple locations, in which case we pair each destination location (for a given product) with its closest origin location. In no cases was a product's origin location(s) reported to have changed over the sample period so in our analysis below each destination has a fixed and unique origin location for each product. We refer to one of these unique origin-destination pairs as a "trading pair", i.e. a pair where one of the locations is the destination where it is sold and the other is either the factory location for that product or the port of entry.

[4] Appendix A contains further details. Note that we would obtain even larger differences in trade costs between the US and Sub-Saharan Africa if we were to convert trade costs into real consumption quantities using PPP-adjusted exchange rates as price levels are higher in the US.

[5] For goods imported into Nigeria, we take Lagos as the origin price since all the goods enter through the port there. For Ethiopia, which is landlocked, we take the town of Kombolcha, the first major stop on the trucking route from the port in Djibouti (which conducts more than 90 percent of Ethiopia's trade, Assefa (2013)) to Addis Ababa.

[6] In 101 cases we were told that the information was confidential, in 5 cases the firm claimed that it did not know, in 7 cases the product was made abroad and the port of entry could not be provided, in 10 cases there were too many factory locations to list, and in 19 cases we could not contact the firms.

Due to the requirements of our procedure for estimating pass-through rates (as explained in Section 4.2 below), we omit any barcode for which the price at the product's origin location is observed in six or fewer months.[7] This final procedure removes no products from the Ethiopian sample and one product from the Nigerian sample. Unlike in SSA, where enumerators are told to seek out specific products, the NCP only records products actually purchased by sample households. Since many of the US factories are located in counties with small populations, 42 products from the US sample are lost through this restriction.

Our final sample contains 15 products in Ethiopia, 18 products in Nigeria and 46 products in the US (coming from 8, 6 and 36 unique origin locations respectively, with the exact factory locations shown in Figure 1). The average origin prices, price gaps and origin-destination distances are reported in Table 1 while Appendix Table C.1 lists the products themselves. One may be concerned about the representativeness of our trade costs estimates if few people live in remote areas. Despite the fact that, in Ethiopia and Nigeria, our origin locations are typically in the major cities of Addis Ababa and Lagos, there is no systematic tendency for population at destination locations to be higher close to the origin locations. Appendix Figure C.1 plots the distribution of the population at each destination location for all trading pairs (except from the origin to itself) across all products. As is evident, many people in these countries live in locations where goods are traveling long distances to reach them.

2.3 Distance data

Our primary measure of distance is simply the geodesic distance (i.e. the shortest distance along the Earth's surface) between two locations. In Section 4.5, we try to explain differences across countries by disentangling the quantity and quality of their transportation networks. To control for quantity differences, we use a distance measure that is based on the distance between origin and destination locations while following the quickest road route as calculated by Google Maps.[8,9] The approximate road network used, for each country, in these alternative distance calculations is shown in Figure 1.[10] Because measures of road quality are unavailable for our African countries, we work with a measure of quality based purely on travel speeds. We again extract these from Google Maps (the time taken to travel along the quickest route).[11]

[7] For the origin price, we take prices in the closest location to the factory where the product is observed. If the closest location is more than 100 miles away from the factory we omit this product.

[8] Our assumption that intermediaries follow the quickest route from any origin location to any destination location rules out the possibility of economies of scale in trading which could give rise to (for example) hub-and-spoke trading networks. The complexity of such networks puts a full treatment of this possibility beyond the scope of this paper (and the literature to date).

[9] These quickest road route distances are based on Google Map data from the year 2012—the earliest year for which these data were available—and so may not be strictly equal to the actual road distances in earlier sample years.

[10] The road network is only approximate because the underlying road network data used by Google Maps is not publicly accessible. The roads illustrated in Figure 1 therefore come from public data sources (Digital Chart of the World, via DivaGIS; and http://infrastructureafrica.org) that closely match the road networks shown in Google Maps.

[11] To provide a sense of the mapping between time and distance in the three countries, we picked 5 random points in each country and calculated the minutes/mile on the nearest national highway, nearest secondary road and nearest tertiary road. For Ethiopia, these were 1.2, 1.4 and 1.9 minutes/mile respectively; for Nigeria these were 1.2, 1.6 and 2.6 minutes/mile; and for the US these were 0.8, 1.2 and 1.2 minutes/mile.

3 Theoretical Framework

In this section, we first describe a model of intra-national trade carried out by intermediaries who potentially enjoy market power. We then go on to discuss how this framework can be used to inform estimates of the size of intra-national trade costs, as well as the distribution of the gains from trade between consumers and intermediaries.

3.1 Model environment

We begin by considering a single product which is potentially sold in multiple markets indexed by d. For the sake of simplicity we assume that there are no interactions across locations. In each market the (inverse) demand for the product is given by $P(Q_d; D_d)$ where Q_d is the total quantity consumed and D_d parameterizes (without loss of generality) the demand curve in location d. Section 3.4 introduces the multiple products and time periods that enter our empirical analysis.

The product in question could be either domestically produced or imported from abroad. Domestically produced products are made at a single factory location indexed by o, and imported products are imported into the country through a port or border crossing at location o. Regardless of whether the products are made at home or abroad, the domestic 'origin' of the product is therefore location o. We assume that the product is bought and sold on wholesale markets at the origin (i.e. factory gate or port) location for a price P_o. This product is then traded from location o to the destination location d by domestic intermediaries. These intermediaries specialize in the activity of purchasing a product in bulk at a wholesale market, transporting the product to a destination location, and finally selling the product to consumers at that location. Again for the sake of parsimony, we assume that these distribution and retail activities are bundled into the actions of one 'intermediary' sector.

Intermediaries incur total costs $C_d(q_d)$ while trading q_d units from location o to location d. (To simplify notation we use d rather than od subscripts here to denote an origin-destination trading pair since we only consider a single source location for the product. When we introduce multiple products in Section 3.4 originating from potentially different source locations we reintroduce the od notation.) These total costs include both a fixed cost, F_d, and a marginal cost, c_d. We assume that the marginal cost is both 'specific' (i.e. charged per unit of product shipped) and constant (i.e. independent of q_d).[12] Marginal costs in the intermediary sector are the sum of the cost of buying the product at the origin location (which is simply the origin price, P_o) and the marginal costs of trading, denoted by $\tau = \tau(\mathbf{X}_d)$, which depend on a vector of potential cost-shifters $\mathbf{X}_d$.

We refer to $\tau(\mathbf{X}_d)$, the focus of our analysis, as 'trade costs' for short. Our goal is to estimate the extent to which these costs depend on cost shifters—that is, we aim to recover the derivative $\frac{\partial \tau(\mathbf{X}_d)}{\partial x_d}$ for some particular cost-shifter x_d. A leading example of such a cost-shifter in the empirical literature is the distance from location o to location d. An analysis of distance forms the bulk of our empirical analysis in Section 4, but other potentially important cost-shifters might include intra-

[12] We believe that this assumption, that marginal trade costs are specific rather than *ad valorem*, is realistic in the setting considered here. However, if the true trade costs contain an *ad valorem* component, this assumption will lead us to underestimate trade costs, as discussed in footnote 22 below.

national borders, differences in language or ethnicity, or roadblocks at which formal regulatory burdens or even requests for bribes might be encountered by traders. Note that, following the previous literature, we take an all-encompassing view of intra-national trade costs, $\tau(\mathbf{X}_d)$, such that they consist of the entire marginal cost of buying a product in location o and selling it to consumers in location d (such as, for example, destination-specific local retail costs).[13]

Summarizing these assumptions about costs we have:

Assumption 1 [Costs]. *The cost to an intermediary of selling q_d units of the product in location d, when sourcing the product from location o is given by:*

$$C(q_d) = [P_o + \tau(\mathbf{X}_d)]\, q_d + F_d.$$

Let there be m_d identical intermediaries who buy the product at location o and sell it at location d. We assume that intermediaries maximize profits by choosing the quantity of the product to sell (as seems reasonable in this setting since intermediates must purchase goods at the origin before selling them at the destination). Let Q_d denote the total amount sold in location d by all intermediaries. The essential strategic interaction across intermediaries here is the extent to which an intermediary's quantity choice q_d affects other intermediaries' profits through the aggregate quantity Q_d. We follow the 'conduct parameter' approach to modeling oligopolistic interactions (e.g. Seade, 1980) and assume that this relationship is summarized by the parameter $\theta_d \equiv \frac{dQ_d}{dq_d}$. We allow this parameter to take any value; however, as is well known, a number of prominent assumptions about market structure are encapsulated in distinct values of the parameter θ_d. For example, the case of Cournot oligopoly corresponds to $\theta_d = 1$, the case of a pure monopolist corresponds also to $\theta_d = 1$, the case of collusion corresponds to $\theta_d = m_d$, and the case of perfect competition corresponds to $\theta_d \to 0$. Finally, we define the ratio $\phi_d \equiv \frac{m_d}{\theta_d}$ as the 'competitiveness index' (since it rises in the number of intermediaries m_d and falls in these intermediaries' perceived individual influence on aggregate supply θ_d) and assume that this ratio is fixed within a location. Summarizing our assumptions about market structure we have:

Assumption 2 [Market structure]. *m_d identical intermediaries selling the product in location d choose supply q_d to maximize profits subject to the perceived response of other firms summarized by the parameter $\theta_d \equiv \frac{dQ_d}{dq_d}$. The competitiveness index $\phi_d \equiv \frac{m_d}{\theta_d}$ is fixed within a location d.*

It is important to note that, in our empirical analysis below, we will not need (nor be able) to separately identify m_d or θ_d. These two variables matter only to the extent that they shift the competitiveness index ϕ_d. Further, note that while we assume that the parameter ϕ_d is fixed within each location, it can vary freely across locations and hence be arbitrarily correlated with the cost shifters x_d—this flexibility is important given that the key variation for identifying trade costs is in the cross-sectional comparisons of prices across locations. Focusing on the Cournot case, $\theta_d = 1$, for example, Assumption 2 implies that entry is fixed, an assumption that we believe is plausible (over the high-frequency, monthly-data time periods in our application) given credit constraints,

[13] In Section 4.4 we discuss attempts to separate destination-specific costs from origin-to-destination distance.

reputation issues and ethnic trading traditions. In addition, Section 4.4 presents a number of empirical extensions designed to explore the sensitivity of our results to Assumption 2.

3.2 Identifying intra-national trade costs from price gaps

Recall that our goal is to estimate $\frac{\partial \tau(\mathbf{X}_d)}{\partial x_d}$, the extent to which trade costs $\tau(\mathbf{X}_d)$ depend on some particular cost-shifter x_d. Following a large previous literature, we estimate this relationship using information contained in spatial price gaps, $P_d - P_o$. Using Assumption 1, and the intermediaries' first-order conditions, spatial price gaps can be written as:

$$P_d - P_o = \tau(\mathbf{X}_d) + \mu\left(c_d, \phi_d, D_d\right), \tag{2}$$

where $\mu\left(c_d, \phi_d, D_d\right)$ is the mark-up charged by intermediaries in location d.[14] Without loss of generality, the mark-up is a function of the intermediaries' marginal costs c_d, the competitive environment faced by the intermediaries (summarized by the competitiveness index ϕ_d), and the demand conditions D_d.

To see how $\frac{\partial \tau(\mathbf{X}_d)}{\partial x_d}$ can be identified empirically, we consider how a small change in x_d would alter equation (2). This is empirically analogous to a comparison between two destination locations with a small difference in the value of their cost shifter x_d. Such a perturbation to equation (2) would satisfy:[15]

$$\frac{d(P_d - P_o)}{dx_d} = (1 + \frac{\partial \mu_d}{\partial c_d})\frac{\partial \tau(\mathbf{X}_d)}{\partial x_d} + \frac{\partial \mu_d}{\partial \phi_d}\frac{\partial \phi_d}{\partial x_d} + \frac{\partial \mu_d}{\partial D_d}\frac{\partial D_d}{\partial x_d} \equiv \rho_d \frac{\partial \tau(\mathbf{X}_d)}{\partial x_d} + \frac{\partial \mu_d}{\partial \phi_d}\frac{\partial \phi_d}{\partial x_d} + \frac{\partial \mu_d}{\partial D_d}\frac{\partial D_d}{\partial x_d}. \tag{3}$$

In this expression the quantity ρ_d is known as the (short-run) pass-through rate, defined as the effect of a firm's marginal cost on the price it charges while holding competitiveness (and hence entry) fixed, as per Assumption 2. That is, $\rho_d \equiv \frac{\partial P_d}{\partial c_d} = 1 + \frac{\partial \mu_d}{\partial c_d}$.[16] Note that while ρ_d is defined such that ϕ_d is held fixed, equation (3) is fully general and allows for the possibility that ϕ_d is correlated with x_d. In the remainder of this paper we refer to the short-run pass-through rate ρ_d as simply *the* pass-through rate. It is straightforward to show that, in general, pass-through takes the form:

$$\rho_d = \left[1 + \frac{(1 + E_d(P_d))}{\phi_d}\right]^{-1}, \tag{4}$$

where $E_d(P_d) \equiv \frac{Q_d}{\left(\frac{\partial P_d}{\partial Q_d}\right)} \frac{\partial\left(\frac{\partial P_d}{\partial Q_d}\right)}{\partial Q_d} \lessgtr 0$ is the elasticity of the slope of inverse demand. As equation (4)

[14] That is,

$$\mu_d \equiv P_d - c_d = c_d[\frac{\phi_d}{-\eta(D_d, c_d, \phi_d)} - 1]^{-1},$$

where η is the elasticity of inverse demand. We assume that both the second-order and stability conditions shown in Seade (1980) hold, namely that $\frac{Q_d}{\left(\frac{\partial P_d}{\partial Q_d}\right)} \frac{\partial\left(\frac{\partial P_d}{\partial Q_d}\right)}{\partial Q_d} > -2\phi_d$ and $\frac{Q_d}{\left(\frac{\partial P_d}{\partial Q_d}\right)} \frac{\partial\left(\frac{\partial P_d}{\partial Q_d}\right)}{\partial Q_d} > -\phi_d - 1$.

[15] In writing equation (3) we have assumed that ϕ_d and D_d depend on x_d in a continuous manner. As we describe below, this assumption is not necessary for our empirical analysis but we make it here for simplicity.

[16] As is typical in the Industrial Organization literature, and as in Weyl and Fabinger (2013), we define the pass-through rate ρ_d as the effect of marginal costs on prices (i.e. $\frac{\partial P_d}{\partial c_d}$) rather than the *proportional* effect of marginal costs on prices (i.e. $\frac{\partial \ln P_d}{\partial \ln c_d}$). Note that the well-known case of monopolistic competition with CES preferences and a continuum of firms delivers proportional pass-through equal to one but $\rho_d > 1$.

makes clear, pass-through depends on only two market characteristics: the competitiveness of the market (i.e. ϕ_d) and the second-order curvature of the demand curve (i.e. $E_d(P_d)$, the elasticity of the slope of demand).

The expression in equation (3) highlights the challenges involved when attempting to identify the object of interest here—the way that trade costs depend on a cost-shifter x_d, or $\frac{\partial \tau(\mathbf{X}_d)}{\partial x_d}$—from the extent to which price gaps vary across locations with different values of the cost shifter (i.e. from $\frac{d(P_d - P_o)}{dx_d}$). To interpret equation (3), begin by observing that if the mark-up did not depend on the cost-shifter (i.e. $\frac{d\mu_d}{dx_d}$=0) then variation in spatial price gaps would identify $\frac{\partial \tau(\mathbf{X}_d)}{\partial x_d}$ in a straightforward manner since $\rho_d = 1$ in such a setting and the last two terms would be zero. As discussed in the Introduction, this has been the case assumed in virtually all of the existing literature on estimating trade costs from price gaps. In what follows, we relax this assumption.

The first term in equation (3) describes the most obvious concern when mark-ups are variable, namely that the pass-through rate $\rho_d \neq 1$. That is, when an oligopolist's marginal costs increase (e.g. because $\frac{\partial \tau(\mathbf{X}_d)}{\partial x_d} > 0$), the oligopolist will find it optimal to increase its price either by less than (i.e. $\rho_d < 1$) or more than (i.e. $\rho_d > 1$) the marginal cost increase. All that can be said in general is that some of the marginal cost will be passed through into prices (i.e. $\rho_d > 0$). The extent of imperfect pass-through (i.e. the deviations of ρ_d from unity) governs, in an important manner, the extent to which spatial price gaps provide a biased estimate of trade costs. As we describe shortly, this observation forms the core of our empirical strategy for estimating trade costs in imperfectly competitive settings. In summary, a cost shifter such as x_d has both a direct effect (via the marginal cost, i.e. $\frac{\partial \tau(\mathbf{X}_d)}{\partial x_d}$) and an indirect effect (via the mark-up) on the price charged; but in any case, ρ_d is a sufficient statistic for the magnitude of the indirect effect, which allows us to uncover the direct effect, the object of interest here.

The second and third terms in equation (3) describe a source of bias that is conceptually distinct from that in the first term. These terms capture the natural possibility that mark-ups vary across locations not just because, whenever $\rho_d \neq 1$, mark-ups vary with marginal costs and marginal costs vary across locations, as captured in the first term, but simply because competitive conditions ϕ_d vary across locations or because preferences D_d vary across locations. (In fact, as we discuss in Section 5, long-run entry decisions within our framework will naturally lead to less competition in remote locations.) Our empirical strategy for dealing with this source of bias is, as described in detail below, based on attempts to control for these two terms. While this would be challenging in general, because competitiveness ϕ_d and preferences D_d are not observable, we will be helped by the fact that the pass-through rate ρ_d is observable and, as equation (4) suggests, knowledge of the pass-through rate reveals a great deal about competitiveness and preferences.

3.3 The case of constant pass-through demand

Equation (3) above described three sources of bias that may arise when estimating trade costs from spatial price gaps in settings of imperfect competition: incomplete pass-through, variation in competitiveness, and variation in preferences. However, an estimate of the pass-through rate ρ_d can be used to avoid all three of these sources of bias. The methodology we propose here therefore

proceeds in two steps: in a first step we estimate ρ_d, and in a second step we use this estimate of ρ_d to correct for bias due to variable mark-ups.

In order to do so as parsimoniously as possible, we make an additional assumption: that the pass-through rate ρ_d is constant over quantities. As equation (4) makes clear, a sufficient condition for the pass-through rate to be constant over quantities (given Assumption 2, which holds ϕ_d constant over quantities) is that consumer preferences are such that the elasticity of the slope of inverse demand, $E_d(P_d)$, is constant at all prices P_d. Bulow and Pfleiderer (1983) prove that a necessary and sufficient condition for E_d to be constant is that demand belongs to the following class:

Assumption 3 [Bulow-Pfleiderer demand]. *Consumer preferences take the constant pass-through, Bulow-Pfleiderer inverse demand form such that price P_d depends on total demand Q_d in the following manner:*

$$P_d(Q_d) = \begin{cases} a_d - b_d \left(Q_d\right)^{\delta_d} & \text{if } \delta_d > 0 \text{ and } a_d > 0, b_d > 0, 0 < Q_d < \left(\frac{a_d}{b_d}\right)^{\frac{1}{\delta_d}} \\ a_d - b_d \ln(Q_d) & \text{if } \delta_d = 0 \text{ and } a_d > 0, b_d > 0, 0 < Q_d < \exp\left(\frac{a_d}{b_d}\right) \\ a_d - b_d \left(Q_d\right)^{\delta_d} & \text{if } \delta_d < 0 \text{ and } a_d \geq 0, b_d < 0, 0 < Q_d < \infty. \end{cases} \tag{5}$$

For this demand form we have $E_d = \delta_d - 1$; that is, by design, E_d is pinned down by a (constant) model parameter, but this parameter is free to vary. Hence, from equation (4), equilibrium pass-through under Assumption 3 is equal to

$$\rho_d = \left[1 + \frac{\delta_d}{\phi_d}\right]^{-1}. \tag{6}$$

Equilibrium pass-through can be 'incomplete' (i.e. $\rho_d < 1$) with $\delta_d > 0$ and 'more than complete' (i.e. $\rho_d > 1$) with $\delta_d < 0$. Nothing in this class of preferences restricts how pass-through ρ_d varies across locations d within a country; the only restriction is that pass-through does not change within a location in response to the quantity Q_d supplied there in equilibrium. Finally, note that, whatever the demand parameters, the state of competitiveness (summarized by ϕ_d) matters for equilibrium pass-through; in particular, if competition were perfect (i.e. $\phi_d \to \infty$) then equilibrium pass-through is 'complete' (i.e. $\rho_d = 1$) for any value of the demand parameter δ_d.

From an empirical perspective, there are a number of attractive features of the constant pass-through demand class. The first is that, as we describe in the next subsection, this demand class leads to a parsimonious empirical strategy for estimating trade costs despite the presence of variable mark-ups. A second attraction of the approach embodied in Assumption 3 is its flexibility. This demand class nests prominent special cases such as linear, quadratic and isoelastic demand, but whereas those special cases restrict the pass-through demand parameter δ_d to take a particular value, the constant pass-through demand class of Assumption 3 allows δ_d to take on any value.[17]

[17] While in principle demand could be such that even the second-order curvature parameter E_d varies with the quantity demanded, in such a case our estimates will still provide a local approximation to the pass-through rate around the equilibrium market quantity. For example, in Section 5 we explore the incidence of a small change in international tariffs, an exercise for which locally constant pass-through estimates are sufficient for calculating local incidence.

Assumption 3, together with the intermediaries' first-order conditions, implies that

$$P_d - P_o = \rho_d \tau(\mathbf{X}_d) + (1 - \rho_d)(a_d - P_o). \tag{7}$$

Recall from equation (3) that a challenge in estimating $\frac{\partial \tau(\mathbf{X}_d)}{\partial x_d}$ arises because unobserved preferences and market structure may co-vary with x_d. Equation (7) highlights that, in the Bulow-Pfleiderer case, two variables—the pass-through rate ρ_d and the demand-shifter a_d—are sufficient to control for these two sources of omitted variable bias. Naturally, the a_d component is a demand-side parameter, but what is useful empirically is that the other demand parameters, b_d and δ_d, do not enter equation (7) directly. Instead, b_d does not enter at all and the effect of δ_d is subsumed by the presence of ρ_d since, as per equation (6), ρ_d depends on δ_d. Likewise, equation (7) suggests that ρ_d acts as a sufficient statistic for the competitiveness of a location.[18]

Equation (7) will form the bedrock of our empirical strategy for correcting for the three biases laid out in the previous section and estimating $\frac{\partial \tau(\mathbf{X}_d)}{\partial x_d}$ from variation in price gaps $P_d - P_o$ across locations d with differing levels of the cost-shifter x_d. In order to describe this strategy, which draws on data spanning many locations, products and time periods, we first introduce our notation (and the additional assumptions required in a dynamic environment) for incorporating such variation.

3.4 From theory to estimation

We now extend the discussion above, which pertained to a single product in a given destination market d, to a setting in which we observe multiple products $k \in K$ selling in locations $d \in N$ at multiple time periods $t \in T$. However, for simplicity we continue to assume that there are no interactions across locations, products or time periods.[19] We therefore simply allow all variables and parameters from the previous subsection to vary freely across products, destination locations

[18] To see this formally, note that taking the derivative of equation (7) with respect to x_d implies

$$\frac{d(P_d - P_o)}{dx_d} = \rho_d \frac{\partial \tau(\mathbf{X}_d)}{\partial x_d} + \frac{d\rho_d}{dx_d}\left(\frac{P_d - a_d}{\rho_d}\right) + \frac{\partial a_d}{\partial x_d}(1 - \rho_d),$$

where the second and third terms then correspond to $\left(\frac{\partial \mu_d}{\partial \phi_d}\frac{\partial \phi_d}{\partial x_d} + \frac{\partial \mu_d}{\partial D_d}\frac{\partial D_d}{\partial x_d}\right)$ in equation (3) for the Bulow-Pfleiderer case. However, it is more straightforward to transform equation (7) to obtain

$$\frac{P_d - \rho_d P_o}{\rho_d} - \frac{(1 - \rho_d)}{\rho_d} a_d = \tau(\mathbf{X}_d).$$

This implies that a perturbation with respect to x_d in the left-hand side identifies $\frac{\partial \tau(\mathbf{X}_d)}{\partial x_d}$, the object of interest. The procedure we describe below provides an estimate of ρ_d for each location. In addition, we propose a procedure to control for the demand-shifter a_d. Put together this implies that the left-hand side is identified at each location and hence its variation across locations with distinct values of x_d identifies $\frac{\partial \tau(\mathbf{X}_d)}{\partial x_d}$.

[19] This assumption involves three restrictions on the economic environment. First, we abstract from general equilibrium considerations that would introduce interactions in factor markets across or within locations; however, our empirical approach below will introduce various fixed effects that are likely to control for any bias due to such interactions. Second, we do not model explicitly the possibility that the demand curve for a given product-location is dependent on the price of other products, or on the income of consumers in that location. While these effects are not modeled explicitly, we do allow the level a^k_{dt} and slope b^k_{dt} of inverse demand in equation (5), where we would expect the bulk of income and cross-price substitution effects to play out, to vary freely across locations, products and time. (We also allow the second-order demand curvature parameter, δ^k_{dt}, to vary freely across location and products and—to a lesser extent, as discussed below—across time.) Finally, we assume that intermediaries' marginal costs are sufficiently low (relative to consumers' travel costs) that consumer always prefer to buy goods locally from an intermediary rather than traveling themselves to other locations to make their purchases.

and time periods. As products are made in different origin locations, we must now keep track of the source for each product in each destination and so we replace d subscripts with od subscripts.

We now discuss our proposed strategy for estimating $\frac{\partial \tau(\mathbf{X}^k_{odt})}{\partial x^k_{odt}}$ from variation in price gaps $P^k_{dt} - P^k_{ot}$ across locations d with differing levels of the cost-shifter x^k_{odt}. As discussed briefly above, our strategy relies on using an estimate of ρ^k_{odt} to correct for the possibility that intermediaries charge differential mark-ups at locations with different values of the cost-shifter x^k_{odt}. To implement this strategy we will therefore first need to estimate ρ^k_{odt}. However, there is no hope of estimating a separate value of ρ^k_{odt} for each time period t, as there would then be as many values of ρ^k_{odt} to estimate as there are price observations. We therefore proceed first with the extreme assumption that ρ^k_{odt} is constant over the entire sample period of length T (but still free to vary across products k and destinations d). However, we relax this assumption in Section 4 by estimating separate values of $\rho^k_{od\widetilde{T}}$ in various periods of length $\widetilde{T} < T$ (where $\widetilde{T} \geq 2$ is a minimum requirement for identification). Summarizing this discussion, we have:

Assumption 4 [Static Pass-Through]. *The pass-through rate ρ^k_{odt} is free to vary across products k and destination locations d, but it is fixed across time periods within a product-destination. That is, $\rho^k_{odt} = \rho^k_{od}$ for all $t \in T$.*

Recalling that $\rho^k_{odt} = \left(1 + \frac{\delta^k_{dt}}{\phi^k_{odt}}\right)^{-1}$, two natural sufficient conditions for Assumption 4 arise. The first is that the demand curvature parameter δ^k_{dt} is constant across time periods. We believe this to be a reasonable restriction on preferences that is considerably weaker than in the existing literature. Recall that, while this sufficient condition restricts δ^k_{dt} to be constant across time periods, we have placed no restrictions at all on the level of inverse demand a^k_{dt} or the slope of inverse demand b^k_{dt}. The second condition is that the competitiveness parameter ϕ^k_{odt} is constant across time periods, which would be the case if both the number of intermediaries m^k_{odt} and those intermediaries' conduct parameter θ^k_{odt} do not change across time periods (since $\phi^k_{odt} \equiv \frac{m^k_{odt}}{\theta^k_{odt}}$). While the constancy of the conduct parameter θ^k_{odt} is a natural restriction, holding m^k_{odt} constant amounts to assuming that entry across time periods doesn't respond to changes in the economic environment. This restriction is clearly more plausible over short time spans, the length of which is unknown; we therefore find it reassuring that our results are robust to using different lengths of time $\widetilde{T}$ over which entry is assumed to be fixed.[20]

Combining Assumptions 1-4 and equation (7) then immediately implies the following:

$$P^k_{dt} = \rho^k_{od} P^k_{ot} + \rho^k_{od} \tau(\mathbf{X}^k_{odt}) + (1 - \rho^k_{od}) a^k_{dt}. \tag{8}$$

This equation forms the core of our empirical analysis. The immediate challenge in taking equation (8) to the data is that, while the variables P^k_{dt}, P^k_{ot} and $\mathbf{X}^k_{odt}$ are directly observable, the objects $\tau(\cdot)$, ρ^k_{od} and a^k_{dt} are not. Our approach will be to estimate the parameters of interest, $\tau(\cdot)$ and ρ^k_{od}, while treating a^k_{dt} as unobserved heterogeneity (i.e. an econometric error term, the properties of

[20]The two-step strategy that we describe below estimates ρ^k_{od} using monthly time variation, so Assumptions 2 and 4 are consistent with the common argument that high-frequency time-series analysis is likely to reveal short-run responses (and cross-sectional analysis long-run responses); see, e.g., Houthakker (1965).

which we discuss at length in Section 4). While it would be possible, in principle, to estimate $\tau(\cdot)$ and ρ^k_{od} in this equation directly, for ease of exposition, and to reduce computational burden, we instead describe an unbiased two-step procedure that achieves the same goal. We describe this two-step procedure briefly here, and in more detail in Section 4.

Step 1: Recover estimates of pass-through ρ^k_{od}: Equation (8) implies that a regression of destination prices P^k_{odt} on origin prices P^k_{ot}, conditional on controls for both trade costs $\tau(\mathbf{X}^k_{odt})$ and local demand shifters a^k_{dt}, would reveal the equilibrium pass-through rate ρ^k_{od} inherent to each destination market and product.[21] While both trade costs and local demand shifters are unobservable to researchers—indeed, if these were observable then answers to the questions we pose in this paper would be immediately available—in Section 4.2 we propose an empirical strategy that aims to control for these variables and hence to provide unbiased estimates of the equilibrium pass-through rate ρ^k_{od} prevailing in each destination location d and product k.

Step 2: Recover estimates of intra-national trade costs $\tau(\cdot)$: Suppose, with Step 1 complete, that an unbiased estimate of ρ^k_{od} is available; denote this estimate $\widehat{\rho^k_{od}}$. Then we can write equation (8) as

$$\frac{P^k_{dt} - \widehat{\rho^k_{od}} P^k_{ot}}{\widehat{\rho^k_{od}}} = \tau(\mathbf{X}^k_{odt}) + \frac{(1 - \widehat{\rho^k_{od}})}{\widehat{\rho^k_{od}}} a^k_{dt}. \tag{9}$$

In contrast to the spatial price gap $P^k_{dt} - P^k_{ot}$ that has featured prominently in the existing literature on trade costs, we refer to the left-hand side of equation (9) as the 'adjusted price gap'. Equation (9) suggests that, once the left-hand side is written in terms of the adjusted price gap rather than the price gap, the object of interest, $\frac{\partial \tau(\mathbf{X}^k_{odt})}{\partial x^k_{odt}}$, can be traced out empirically from variation in x^k_{odt}.[22]

Note that if mark-ups did not exist, or did not vary across locations, then we would be in the case for which $\widehat{\rho^k_{od}} = 1$—exactly the case in which the adjusted price gap would be equal to the price gap and the methods used in the existing literature would be valid. Away from this knife-edge case, however, pass-through $\widehat{\rho^k_{od}}$ would not equal one. Indeed, in Step 1 we find, as is consistent with many previous estimates of pass-through rates, that $\widehat{\rho^k_{od}}$ often differs substantially from one. Our approach is therefore designed to provide unbiased estimates of trade costs for any value of $\widehat{\rho^k_{od}}$. The only complication—as suggested by equation (9)—is that the unobserved demand-

[21] A potential concern here is that the variation used to identify ρ^k_{od}, variation in origin prices, may be correlated with shocks to the price of some other product (say k') produced at the same origin location. To the extent that products k and k' are substitutes/complements, the resulting changes in $P^{k'}_{dt}$ could affect the demand for product k in such a way as to affect (if they were to affect the second-order curvature of the demand curve, E^k_{dt}) the pass-through rate of interest, ρ^k_{od}. A related concern is that, if trade costs $\tau(\mathbf{X}^k_{odt})$ contain a component that is common to both products k and k', then the pass-through from trade costs into prices will affect the price of both products, which again may affect the mark-up charged on product k to the extent that the products are substitutes/complements. We have ruled out such cross-product general equilibrium effects in this section by assumption, but because the bulk of consumption is non-traded, and because many products come from separate origin locations, we feel this assumption offers a useful approximation for the questions posed here.

[22] Returning to the possibility that trade costs are *ad valorem* rather than specific, as in Assumption 1, we can sign the bias. Consider the case where $\tau(\mathbf{X}^k_{odt}) = P^k_{ot}(\eta_0 + \eta_1 x^k_{odt})$. Equation (9) then becomes $\frac{P^k_{dt} - \rho^k_{od} P^k_{ot}}{\rho^k_{od}} = \frac{1}{1+\eta_0+\eta_1 x^k_{odt}} \tau(\mathbf{X}^k_{odt}) + (\frac{1}{\rho^k_{od}} - \frac{1}{1+\eta_0+\eta_1 x^k_{odt}}) a^k_{dt}$ and we will underestimate $\frac{\partial \tau(\mathbf{X}^k_{odt})}{\partial x^k_{odt}}$ since $\frac{1}{1+\eta_0+\eta_1 x^k_{odt}}$ is concave in distance.

shifter a_{dt}^k must be controlled for (and multiplied by a term involving pass-through, $\frac{(1-\widehat{\rho_{od}^k})}{\widehat{\rho_{od}^k}}$). In Section 4.3 we propose an empirical strategy that does exactly this.

4 Estimating Intra-national Trade Costs

In this section we provide estimates of how intra-national trade costs depend on distance in Ethiopia, Nigeria and the United States. That is, we provide estimates of $\frac{\partial \tau(\mathbf{X}_{odt}^k)}{\partial x_{odt}^k}$ for a particular set of cost-shifters x_{odt}^k that are based on metrics of distance. Our estimates rely on a two-step empirical procedure, described in Sections 4.2 and 4.3 below, which aims to provide unbiased estimates of intra-national trade costs from price gaps across locations even when mark-ups vary across those locations. However, we begin in Section 4.1 with a simpler first look at spatial price gaps in order to facilitate a comparison with the existing literature.

4.1 A first look at spatial price gaps

For benchmarking purposes we begin by imposing the restriction that $\rho_{od}^k = 1$, such that mark-ups do not vary across locations. This has been the dominant approach in the existing literature on estimating trade costs, and would hold if trading were perfectly competitive. Under this restriction, equation (9) then implies

$$P_{dt}^k - P_{ot}^k = \tau(\mathbf{X}_{odt}^k). \tag{10}$$

As is clear from equation (10), in the case where $\rho_{od}^k = 1$ trade costs can be easily inferred from price gaps. This is intuitive: if mark-ups don't vary across locations then prices vary across locations only because of trade costs. Unfortunately, the assumption of constant mark-ups is directly refuted in our data, as we show in Section 4.2 below. Nevertheless, it is instructive to consider what the spatial price gaps in our data imply for estimates of intra-national trade costs were we to (erroneously) assume that $\rho_{od}^k = 1$.

While the methodology we develop in this paper could be applied to any vector of cost-shifters $\mathbf{X}_{odt}^k$, in practice we work primarily with one variable, the natural logarithm of the distance between location o and location d. We denote this variable x_{od}. We begin with this simple distance metric because of its prominence in the literature, but we explore additional distance variables in Section 4.5 such as those that adjust for road quality and availability. To highlight this emphasis on x_{od}, consider the following decomposition,

$$\tau(\mathbf{X}_{odt}^k) = f(x_{od}) + \chi_{odt}^k, \tag{11}$$

where $f(\cdot)$ is a nonparametric function that captures how log distance x_{od} affects trade costs and χ_{odt}^k embodies any component of trade costs that does not depend on log distance. This decomposition, along with our decision to work with the log of distance, holds without loss of generality due to the fact that we place no restrictions on $f(\cdot)$.

Our results throughout this Section present nonparametric estimates of the function $f(\cdot)$. In all cases we normalize our estimate of $f(\cdot)$ such that normalized trade costs are zero at the most prox-

imate destination location (approximately 50 miles from the source) in each country. The absolute level of the reported relationships is therefore not meaningful (nor is it identified in the more general models that we estimate below). In practice, we use locally weighted polynomials (with an Epanechnikov kernel of bandwidth = 0.5) to estimate $f(\cdot)$. We also include a fixed effect for each product-time period interaction (where a time period is a month-year pair) to control for systemic product-specific shocks to trade costs.[23] To estimate nonparametric regressions with fixed effects we follow the procedure in Baltagi and Li (2002). Finally, the reported 95 percent confidence intervals are obtained from block-bootstrapping 100 times at the product-destination level following the procedure in Deaton (1997).

Estimating trade costs (when $\rho^k_{od} = 1$) using spatial price gaps between all pairs of locations

Before presenting results on $\frac{\partial \tau(\mathbf{X}^k_{odt})}{\partial x^k_{odt}}$ derived from equation (10), we begin with an additional intermediate step that is designed to connect our results to the existing literature. To do so, we estimate equation (10) among a sample that includes all pairs of locations. The key distinction here is that this construction pays no regard to which locations are actually origin and destination locations for any given product. Without this knowledge there is no reason to expect that an analogous equation (i.e. $P^k_{it} - P^k_{jt} = \tau(\mathbf{X}^k_{ijt})$) would apply within a sample of all locations i and j.[24] Nevertheless, we present estimates of the effects of distance obtained from a sample of all locations because these estimates speak to what we would conclude from our sample if, as is the case in the prior literature, we proceeded without knowledge of which pairs were trading pairs.

Our estimate is displayed as the dotted line in Figure 2.[25] In each of our three sample countries, there is a strictly positive relationship between (the absolute value of) intra-national price gaps $P^k_{it} - P^k_{jt}$ and log distance x_{ij}, when this relationship is estimated across a sample of all pairs of locations i and j. However, it is surprising that this relationship has a similar slope in two countries (Ethiopia and the US) with seemingly distinct trading environments, and yet such substantial dif-

[23] With a balanced panel of observations the inclusion of product-time period fixed effects would have no bearing on our estimates of the purely cross-sectional $f(\cdot)$. Their inclusion here has the advantage of controlling for any sample selection concerns in which the availability of a product in a given time period is correlated with distance. Note also that the dependent variable, $P^k_{dt} - P^k_{ot}$, contains a component, P^k_{ot}, that would be perfectly correlated with product-time period fixed effects were it not for the small number of products that are sourced from multiple locations.

[24] Instead we expect $P^k_{it} - P^k_{jt} = \tau(\mathbf{X}^k_{oit}) - \tau(\mathbf{X}^k_{ojt})$, which implies that variation in price gaps across pairs of locations i and j separated by cost-shifters x^k_{ijt} does not identify the object of interest, $\frac{\partial \tau(\mathbf{X}^k_{ijt})}{\partial x^k_{ijt}}$, as $\tau(\mathbf{X}^k_{ijt})$ does not appear on the right-hand side. Note that even if arbitrage were perfect *across* all locations (something we have ruled out in Section 3) then this would imply that $P^k_{it} - P^k_{jt} \leq \tau(\mathbf{X}^k_{ijt})$ and hence $\frac{\partial \tau(\mathbf{X}^k_{ijt})}{\partial x^k_{ijt}}$ would still not be identified (and while a lower-bound could be placed on $\tau(\mathbf{X}^k_{ijt})$, the same cannot be said for $\frac{\partial \tau(\mathbf{X}^k_{ijt})}{\partial x^k_{ijt}}$).

[25] These all-pair estimates use the absolute value of the price gap as the dependent variable because, in the absence of any knowledge about the trading status and direction among locations i and j, one would not know whether to expect $P_i > P_j$ or $P_i < P_j$. Our sample here consists of all unique pairs of locations (so as not to double-count pairs) for which $i \neq j$. The resulting all-pairs sample is sufficiently large that we faced two computational limitations. First, the US sample is so large that local polynomial estimation on the full sample was infeasible; we therefore work with a random 10 percent sub-sample of locations. Second, in all three countries the sample was too large to compute confidence intervals via a bootstrap routine; we therefore instead display the 95-percent confidence interval from the asymptotically normal conditional variance of the local polynomial estimator. These limitations do not apply to our smaller sample of trading pairs that underpin all other estimates in this paper (including our preferred estimates).

ferences in slope between two countries (Ethiopia and Nigeria) whose trading environments one might expect to be relatively similar. These counter-intuitive results are perhaps less surprising when we remember that, as stressed above, the economic basis for the all-pairs analysis conducted here—and more generally for variation in price gaps $P_{it}^k - P_{jt}^k$ across all sample locations to identify $\tau(\cdot)$, even under perfect competition—is unclear. We now move on to further estimates that we believe provide a closer connection to the structural object of interest here, the magnitude of $\frac{\partial \tau(\mathbf{X}_{ijt}^k)}{\partial x_{ij}}$.

Estimating trade costs (when $\rho_{od}^k = 1$) using spatial price gaps among trading pairs only

We now turn to the estimation of equation (10), which holds as stated only among trading pairs—that is, among pairs of locations for which one location is the origin (for the particular product under consideration) and the other is the destination.[26] In contrast to the estimates from the all-pairs sample used above, we now expect our estimate—because it focuses only on trading pairs—to identify $\frac{\partial \tau(\mathbf{X}_{odt}^k)}{\partial x_{od}}$ in the case of perfect competition (that is, when mark-ups don't vary across locations, or $\rho_{od}^k = 1$). These estimates are displayed, again for each country separately and following the nonparametric estimation procedure described above, in the dashed line in Figure 2.

For two countries in our sample, Ethiopia and Nigeria, the dashed line (which uses only trading pairs) is about twice as steep as the dotted line (which uses all location pairs). Despite the simplicity of the bias-correction procedure we employ here, which simply requires data on the location of production/importation for each product in our sample, we are not aware of prior work that documents systematically the difference between the all-pairs and trading pairs approaches.[27] Our results suggest that the all-pairs approach can dramatically underestimate trade costs.

We turn now to the third country in our sample, the US, for which the estimated relationship between trade costs and distance is non-monotonic. This finding is challenging (though not impossible) to explain if equation (10) is taken literally, such that the estimate in Figure 2 is truly an estimate of how the costs of trading rise with (log) distance.[28] However, in a more general environment (such as that formalized in Section 3 above) in which spatial price gaps reflect both marginal costs of trading and mark-up differences across locations (i.e. $\rho_{od}^k \neq 1$), a price gap that falls with distance is entirely possible, as we describe below.

This finding from the US sample highlights the fact that, while we believe the results in Figure 2 are useful for illuminating the difference between the all-pairs and the trading pairs approaches to estimating trade costs from price gaps, both of these approaches have little to say about trade costs in environments that depart from perfect competition (that is, those where $\rho_{od}^k \neq 1$). In Section 4.2 below we go on to estimate ρ_{od}^k for all products k and locations d and then, upon finding that we can almost always reject the null that $\rho_{od}^k = 1$, pursue an estimate of trade costs in Section

[26] In this and all ensuing analysis we omit the trivial location pair for which $o = d$.

[27] This echoes Cosar, Grieco, and Tintelnot (forthcoming) who, in contemporaneous work, document that, for trade in German and Danish wind turbines, the border effect (i.e. the additional cost incurred when trade crosses the international border) estimated on the all-pairs sample is smaller than that on the trading pairs sample.

[28] One possible explanation would be that local distribution costs—costs such as retail factor prices that are paid in the destination location regardless of a good's origin—are higher at locations near to the origin than at locations further afield. However, in Section 4.4 below we explore the plausibility of such spatial variation in local distribution costs and find little evidence for variation of this type.

4.3 that is robust to the presence of imperfect competition.

4.2 Step 1: Estimating pass-through rates

We now move on to estimate the pass-through rate ρ_{od}^k for each location d and for each product k in our sample. This is the first step of the two-step procedure for estimating trade costs that is outlined in Section 3.4. The pass-through estimates we obtain here are also useful for identifying the incidence of a world price change, as described in Section 5.2 below.

Recall from equation (8) above that pass-through (ρ_{od}^k) relates to the extent to which exogenous origin prices (P_{ot}^k) affect endogenous destination prices (P_{dt}^k) in the following manner:

$$P_{dt}^k = \rho_{od}^k P_{ot}^k + \rho_{od}^k \tau(\mathbf{X}_{odt}^k) + (1 - \rho_{od}^k) a_{dt}^k, \tag{12}$$

where a_{dt}^k represents a shifter of the location of the inverse demand curve from equation (5) above.

Estimating the pass-through rate (ρ_{od}^k) from equation (12), using variation in origin prices (P_{ot}^k), requires controls for the two other terms on the right hand side of the equation: the cost of trading (i.e. $\tau(\mathbf{X}_{odt}^k)$) and local demand shifters (i.e. a_{dt}^k). Unfortunately, both the cost of trading and local demand shifters are unobservable to researchers—indeed, if these were observable then answers to the question posed in this paper would be immediately available. In the absence of such controls we use a fixed effects approach, which requires that the product-specific variation in trade costs and local demand shifters within destinations *over time* is orthogonal to the variation in the origin price over time. Formally, we assume that $\tau(\mathbf{X}_{odt}^k) = \beta_{1od}^k + \beta_{2od}^k t + \zeta_{odt}^k$, such that $\tau(\mathbf{X}_{odt}^k)$ can be decomposed into local but time-invariant (β_{1od}^k), local but trend-like ($\beta_{2od}^k t$), and residual (ζ_{odt}^k) factors.[29] Analogously, we assume that destination market additive demand shocks a_{dt}^k from equation (5) above can be decomposed as follows: $a_{dt}^k = \alpha_{1d}^k + \alpha_{2d}^k t + \nu_{dt}^k$. Note that while this assumption places certain restrictions on how the additive demand shifter, a_{dt}^k, varies across locations, time and products, we place no restrictions on the multiplicative demand shifter, b_{dt}^k, from equation (5). Combining equation (12) with these assumptions we estimate pass-through rates ρ_{od}^k by location and product by estimating the following specification,

$$P_{dt}^k = \rho_{od}^k P_{ot}^k + \gamma_{od}^k + \gamma_{od}^k t + \varepsilon_{dt}^k, \tag{13}$$

where P_{dt}^k is the destination price, P_{ot}^k is the origin price, γ_{od}^k is a product-destination fixed effect, $\gamma_{od}^k t$ is a product-destination linear time trend, and $\varepsilon_{dt}^k = \rho_{od}^k \zeta_{odt}^k + (1 - \rho_{od}^k)\nu_{odt}^k$ is an unobserved error term. The computational advantage of such a specification is that we can estimate pass-through rates from separate regressions for each product-destination pair (recall that each product-destination pair has a unique origin location in any period, and since we have no products where we were told the factory had moved over our sample period, γ_{od}^k is equivalent to a product-destination fixed effect). However, we explore the sensitivity of our results to substantially weakening our identification assumptions through the inclusion of either year-destination γ_{dy} or time-destination γ_{dt} fixed effects in Section 4.4.

[29] Recall that in addition to these fixed effects, all prices are inflation-adjusted using the procedure described in Appendix A and hence are purged of changes in national price levels over time.

For OLS estimates of ρ_{od}^k in equation (13) to be unbiased we require the additional assumption that the origin price P_{ot}^k is not correlated with the time-varying and local (that is, destination location d-specific) shocks to trade costs $\tau(\mathbf{X}_{odt}^k)$ or demand shifters a_{dt}^k.[30] If origin prices are set abroad (in the case of imported goods), or are pinned down by production costs at the factory gate (in the case of domestic goods), or are set on the basis of demand shocks at the origin location (locations that we omit from our analysis), then this orthogonality restriction seems plausible. But a nation-wide demand shock for product k (above and beyond the nation-wide or local demand shock for all products that can be controlled for by the addition of a time-destination γ_{dt} fixed effect that we introduce below) would violate this assumption and lead us to overestimate ρ_{od}^k. Because it is plausible that demand shocks are spatially correlated, we assess the possibility of this bias in Section 4.4 below by exploring how our estimates change when we restrict our sample to only those destination locations d that are beyond a given distance threshold from the origin.

Appendix Figure C.2 displays our estimates of the pass-through rate ρ_{od}^k for all products k and locations d, plotted against the relevant log origin-to-destination distance x_{od} (in addition, we display the estimated nonparametric relationship between ρ_{od}^k and x_{od}). These estimates span a substantial range, though recall that the only restriction that theory places on the pass-through rate is that it be positive, a restriction that very few of our estimates violate.[31] Importantly, our procedure in Step 2, which draws on the estimates here, will adjust for the sampling variation in these estimates via a block-bootstrap procedure. Despite the wide range spanned by these estimates, some general tendencies are noteworthy. One feature of these estimates is that, regardless of the country, the pass-through rate is lower, on average, at destinations that are further distances from the product's source.[32] A second feature is that most of these pass-through estimates are below one, often considerably so. The average estimated pass-through rate in our sample is approximately 0.5.[33] Hence, our estimates suggest that pass-through below one is commonplace in our sample countries, especially since these estimates should, if anything, be biased upwards if product-specific demand shocks at the origin are correlated with those at the destination. This is important because, as discussed above, incomplete pass-through is *prima facie* evidence for imperfect competition.

Recall that the primary motivation for estimating the pass-through rate ρ_{od}^k is that it enters Step 2, our procedure for estimating $\frac{\partial \tau(\mathbf{X}_{odt}^k)}{\partial x_{od}}$, as well as our incidence analysis in Section 5.2. However, it is worth noting that the pass-through rate is a measure of interest in its own right, since it measures the extent to which cost shocks at a distant origin location feed through into equilib-

[30] Formally, this requires $\mathbb{E}\left[P_{ot}^k \zeta_{odt}^k | \gamma_{od}^k, \gamma_{od}^k t\right] = 0$ and $\mathbb{E}\left[P_{ot}^k \nu_{odt}^k | \gamma_{od}^k, \gamma_{od}^k t\right] = 0$ since $\varepsilon_{dt}^k = \rho_{od}^k \zeta_{odt}^k + (1 - \rho_{od}^k)\nu_{odt}^k$ is the composite error term (and ρ_{od}^k is a scalar within each product-destination cell).

[31] The percentage of ρ_{od}^k estimates lying below zero is 8, 17 and 2 in Ethiopia, Nigeria and the US, respectively. Footnote 36 describes how we treat these inadmissible values in our baseline Step 2 estimation procedure, and Section 4.4 discusses the robustness of these estimates to alternative procedures.

[32] This is confirmed by significant negative coefficients (-0.0449, -0.101 and -0.0112 with t-statistics of -2.95, -1.98 and -2.95 for Ethiopia, Nigeria and the US respectively) from the regression of pass-through estimates on log origin-to-destination distance.

[33] For Ethiopia, the mean estimate of ρ_{od}^k is 0.58, with a standard deviation of 0.44 and 89 percent of these estimates below 1 (58 percent significantly so at the 5 percent level). Similarly, for Nigeria the mean is 0.39 (SD =0.66) with 92 percent below 1 (65 percent significantly so), and for the US the mean is 0.77 (SD =0.36) with 78 percent below 1 (31 percent significantly so).

rium retail prices at a destination location. As discussed above, and shown in Appendix Figure C.2, remote locations have, on average, lower estimated pass-through rates; that is, retail prices in more remote locations respond relatively weakly to a given cost shock at the origin. Despite the simplicity of this exercise, our nonparametric finding that pass-through rates are lower in remote locations is, to the best of our knowledge, new in the literature.

There are a number of alternative ways to estimate pass-through rates and it is important to explore the sensitivity of our results to these alternative modeling and econometric assumptions. However, because our primary interest is not in estimates of ρ_{od}^k *per se* but in how estimates of ρ_{od}^k affect estimates of $\frac{\partial \tau(\mathbf{X}_{odt}^k)}{\partial x_{od}}$, we postpone this sensitivity analysis to Section 4.4, after reporting our main estimates of $\frac{\partial \tau(\mathbf{X}_{odt}^k)}{\partial x_{od}}$ in Section 4.3.

4.3 Step 2: Using pass-through adjusted price gaps to measure the effect of distance on trade costs

In Section 4.1 above, we detailed how the price gaps among pairs of trading locations increased with (log) distance. However, this positive relationship is not driven solely by the fact that the costs of trading increase with distance. In addition, intermediaries charge mark-ups, and our model clarifies that the size of the mark-up may be related to distance for three distinct reasons. For the empirically relevant case of incomplete pass-through, remote locations (1) face higher trade costs, which reduces the mark-up that intermediaries choose to charge; (2) may be served by less competitive routes, which raises mark-ups; and (3) may have different levels of the demand shifter also generating a correlation between mark-ups and distance. Our aim in this section is to estimate the true costs of distance—$\frac{\partial \tau(\mathbf{X}_{odt}^k)}{\partial x_{od}}$, that is, how a cost shifter x_{od} such as distance affects trade costs, $\tau(\mathbf{X}_{odt}^k)$—by correcting for these three biases using a procedure that draws on the pass-through estimates reported in Step 1 above. The theoretical framework outlined above offers guidance.

In the case of constant pass-through demand preferences, recall that equation (9) states that

$$\frac{P_{dt}^k - \rho_{od}^k P_{ot}^k}{\rho_{od}^k} = \tau(\mathbf{X}_{odt}^k) + \frac{(1-\rho_{od}^k)}{\rho_{od}^k} a_{dt}^k. \tag{14}$$

Dividing the price gap $P_{dt}^k - P_{ot}^k$ by the pass-through rate purges the price gap of the first bias—spatial variation in markups due to trade costs. However, as in Step 1 above, an identification challenge is posed by the presence of the unobserved demand-shifter a_{dt}^k on the right-hand side of this estimating equation. We again pursue a fixed effects approach and assume that a_{dt}^k can be decomposed as $a_{dt}^k = \alpha_t^k + \alpha_d + u_{dt}^k$ and that the variation in additive demand shifters a_{dt}^k across destination locations (i.e. the variation u_{dt}^k that remains after macro-level time-product effects, α_t^k, and destination effects, α_d, are removed) is uncorrelated with shifters to the cost of trading across locations, $\mathbf{X}_{odt}^k$.[34] Finally, as in Step 1 above, it is important to note that we require no restrictions on the multiplicative demand shifters, b_{dt}^k, from equation (5) above.

With this assumption in place, we obtain our main estimating equation for identifying the

[34] Formally, this requires that $\mathbb{E}\left[\mathbf{X}_{odt}^k u_{dt}^k | \alpha_t^k, \alpha_d\right] = 0$. Note that the assumptions made about demand here are different from—but not contradictory to—those made in Section 4.2 above.

extent to which distance affects trade costs:

$$\frac{P_{dt}^k - \widehat{\rho_{od}^k} P_{ot}^k}{\widehat{\rho_{od}^k}} = \tau(\mathbf{X}_{odt}^k) + \alpha_t^k \left(\frac{1 - \widehat{\rho_{od}^k}}{\widehat{\rho_{od}^k}} \right) + \alpha_d \left(\frac{1 - \widehat{\rho_{od}^k}}{\widehat{\rho_{od}^k}} \right) + \tilde{\varepsilon}_{dt}^k, \tag{15}$$

where $\widehat{\rho_{od}^k}$ is a consistent estimator of the pass-through rate ρ_{od}^k obtained in Step 1 above, α_t^k is a product-time fixed effect, α_d is a destination fixed effect and $\tilde{\varepsilon}_{dt}^k$ is a composite error term.[35] In Section 4.4, we explore the sensitivity of our results to weakening our identification assumptions by replacing the destination fixed effect with either a year-destination or time-destination fixed effect.

Equation (15) implies that this relationship between distance and trade costs is revealed, despite the potential presence of market power in the trading sector, by simply using 'adjusted price gaps' (i.e. $\frac{P_{dt}^k - \widehat{\rho_{od}^k} P_{ot}^k}{\widehat{\rho_{od}^k}}$) as the dependent variable, rather than price gaps (i.e. $P_{dt}^k - P_{ot}^k$) as has been prominent in the literature. Transforming the price gap by replacing P_{dt}^k with $P_{dt}^k / \widehat{\rho_{od}^k}$, as well as including the product-time and destination fixed effects interacted with $\frac{1-\widehat{\rho_{od}^k}}{\widehat{\rho_{od}^k}}$, explicitly controls for the fact that mark-ups may vary over space due to different marginal costs, different levels of competition, or the possibility that the demand shifter varies over time, products and destinations.

In principle, the joint effects of all elements of the vector $\mathbf{X}_{odt}^k$ on the trade costs function $\tau(\mathbf{X}_{odt}^k)$ can be estimated non-parametrically using the method above. But as discussed in Section 4.1, we apply the decomposition $\tau(\mathbf{X}_{odt}^k) = f(x_{od}) + \chi_{odt}^k$ and focus on the nonparametric relationship between trade costs and log distance x_{od} embodied in $f(\cdot)$.

The solid line in Figure 3 displays our baseline nonparametric estimate of how $\tau(\mathbf{X}_{odt}^k)$ depends on log distance—that is, the function $f(\cdot)$—when using the adjusted price gap specification suggested by equation (15).[36] We follow the same procedure as in Section 4.1 above, estimating equation (15) by replacing the term $\tau(\mathbf{X}_{odt}^k)$ with $f(x_{od})$ and employing the procedure in Baltagi and Li (2002). We report the 95 percent confidence interval that obtains when block-bootstrapping 100 times at the product-destination level following the procedure in Deaton (1997)[37], and we normalize the plotted line so as to pass through zero at the most proximate location. For comparison

[35] That is, $\tilde{\varepsilon}_{dt}^k \equiv \frac{(1-\rho_{od}^k)}{\rho_{od}^k} u_{dt}^k - \alpha_t^k \left(\frac{1-\widehat{\rho_{od}^k}}{\widehat{\rho_{od}^k}} - \frac{1-\rho_{od}^k}{\rho_{od}^k} + 1 \right) - \alpha_d \left(\frac{1-\widehat{\rho_{od}^k}}{\widehat{\rho_{od}^k}} - \frac{1-\rho_{od}^k}{\rho_{od}^k} \right)$. As with all previous specifications, we also control for product-time fixed effects α_t^k when estimating equation (15). As $\widehat{\rho_{od}^k}$ appears on both sides of equation 15, there will be correlated measurement error from any finite sample bias in estimating ρ_{od}^k. However, since the mean of the bootstrap estimates is close to the OLS estimate, in practice this appears not to be an issue. Note also that correlation between $\frac{(1-\rho_{od}^k)}{\rho_{od}^k}$ (in $\tilde{\varepsilon}_{dt}^k$) and the dependent variables $\mathbf{X}_{odt}^k$ is not problematic as we condition on $\frac{1-\widehat{\rho_{od}^k}}{\widehat{\rho_{od}^k}}$ by including a full set of α_t^k fixed effects interacted with this object.

[36] One potential concern with implementing the procedure in equation (15) is that it requires division by $\widehat{\rho_{od}^k}$, so our results could in principle be sensitive to our treatment of estimated $\widehat{\rho_{od}^k}$s close to zero. In our baseline results we therefore winsorize all pass-through rate estimates $\widehat{\rho_{od}^k}$ that fall below 0.2. Our results are robust to this procedure, however, as we discuss in Section 4.4 below.

[37] We bootstrap the confidence interval in order to mitigate concerns about generated regressor bias given that some regressors depend on the estimated values of $\widehat{\rho_{od}^k}$ obtained in Step 1. The fact that the estimates of $\widehat{\rho_{od}^k}$ vary at the product-destination level motivates our choice of blocks at which to block-bootstrap, but we explore alternatives below.

we also report (with a dashed line) the estimate reported in Figure 2, which is what we obtain when using trading pairs only (as with the solid line) but when assuming that $\widehat{\rho_{od}^k}$=1. In all three countries, the solid line (which relies on the estimates of $\widehat{\rho_{od}^k}$ from Step 1 above) is considerably steeper than the dashed line (which sets $\widehat{\rho_{od}^k} = 1$ by assumption), implying that procedures that ignore the possibility of mark-up variation over space would considerably understate intra-national trade costs in our sample. Notably, the solid line (in contrast to the dashed line) for the U.S. is weakly upward-sloping everywhere as might be expected from a cost curve. And equally, as one might expect, the solid line estimates for Ethiopia and Nigeria seem similar, at least relative to that for the U.S. which is considerably shallower (note that the scale of the x-axis for the U.S. is smaller than that of the other countries).

Because the estimated nonparametric relationships in Figure 3 appear relatively linear, and to facilitate our analysis of robustness checks below, we now consider a parametric version of the estimation of equation (15) in which $\tau(\mathbf{X}_{odt}^k)$ is assumed to be a linear function of log distance; that is, we set $f(x_{od}) = \mu x_{od}$ in what follows. Estimates of the parameter μ are reported in Table 2. Columns 1-3 refer to Ethiopia, and we describe these specifications in turn here; columns 4-6 and 7-9 then contain entirely analogous specifications for Nigeria and the US, respectively. By way of comparison, we start in columns 1 and 2 with the linear analog of the two first-pass estimates discussed in Section 4.1. Column 1 shows how (the absolute value of) price gaps vary with log distance across all location pairs, while column 2 shows how price gaps vary with log distance across trading pairs only. Finally, column 3 reports the estimated coefficient from a parametric version of equation (15) based on a regression of the 'adjusted price gap' on log distance. All specifications in Table 2 include product-time fixed effects and report standard errors that are clustered at the product-time period level. We also report block-bootstrapped standard errors in column 3 to mitigate generated regressor bias concerns, setting blocks either at the product-time period or the product-destination level. All estimates are statistically significant. Consistent with the nonparametric estimates in Figure 3, the estimates based on adjusted price gaps in column 3 are considerably larger than those based on simple price gaps (columns 1 and 2)—and the same is true for Nigeria (columns 4-6) and the US (columns 7-9). Further, the extent to which the coefficient changes when correcting for variable mark-ups (that is, in the Ethiopian case, when comparing column 2 to column 3) is similar—by a factor of approximately two—in all three countries. This again suggests that, regardless of the country we examine, assuming that mark-ups are constant (i.e. $\rho_{od}^k = 1$) can lead to substantial bias.

What does differ across countries in Table 2 is the magnitude of the coefficient estimates. To interpret these magnitudes, consider the following. The least remote locations in our sample are approximately 50 miles (or 3.9 log miles) away from the source of production.[38] The most remote locations in our African countries are approximately 500 miles (6.2 log miles) away from the source of production.[39] The estimates in columns 3, 6 and 9 of Table 2 then imply that the additional trade

[38]This distance represents approximately the fourth percentile of the distribution of route lengths in Ethiopia and Nigeria, and the third percentile in the US.

[39]This travel time falls within the 99th percentile of route lengths in Ethiopia, the 77th percentile in Nigeria, but

costs incurred when trading goods to the most remote compared to the least remote African sample locations (a difference of 2.3 log miles) is 9 cents in Ethiopia and 13 cents in Nigeria, but the same distance in the US would cost only 2 cents. The mean product observation in our Ethiopia sample costs 43 cents. So the *ad valorem* equivalent of this relative cost of remoteness is 20 percent. The equivalent calculation for our Nigeria sample (mean product cost of 1.03 dollars) suggests a relative cost of remoteness of 12 percent, and the equivalent calculation for the US (mean product cost of 61 cents) is 4 percent. The estimates in Table 2 therefore suggest that the intra-national trade costs imposed by distance are considerable in our African countries, and are substantially smaller in the US. Section 4.5 provides further discussion of the interpretation of the coefficient estimates.

4.4 Robustness checks

We now evaluate the sensitivity of the baseline estimates presented in Section 4.3, to a number of alternative empirical assumptions. These alternative estimates are presented in Table 3. We begin, in row 1, by re-stating the baseline estimates of interest—the parametrically estimated effect of log distance on trade costs, corresponding to columns 3, 6 and 9 from Table 2.

We first examine various alternatives for estimating pass-through rates. These estimates are not a focus of this paper but, rather, a crucial input for Step 2, in which we estimate intra-national trade costs in oligopolistic settings. As discussed in Section 4.2, our baseline pass-through estimates are sometimes close to zero or (in less than 4 percent of cases, as discussed in footnote 31) actually negative. Our baseline estimates winsorized all pass-through estimates, bottom-coding them as $\widehat{\rho_{od}^{k}} = 0.2$, but the estimates in row 2 suggest that this is inconsequential as we obtain similar results when using all of the raw pass-through estimates. The estimates in rows 3 and 4 relax Assumption 4, that the pass-through rate (while free to vary across products and locations) is constant throughout our sample time period. Row 3 splits our 10-year sample into two 5-year periods and estimates a separate pass-through rate within each of these shorter time periods; row 4 does the same for each of four 2.5-year time periods.[40] Despite allowing the pass-through rate to vary freely across these shorter time periods, we obtain similar estimates to those in our baseline. Although it is the short-run pass-through rate that provides a sufficient statistic for the competitive conditions prevailing across locations at any given point in time, rows 5 and 6 explore the sensitivity of our estimates to using longer-term pass-through rates; this has the advantage of allowing for the possibility of sticky price adjustment. Row 5 estimates the pass-through rate by regressing destination prices on origin prices, but also on three lagged origin price terms (as well as the fixed effects in equation 13); the pass-through rates that then enter our Step 2 analysis in row 5 are the sum of these four origin price coefficients. Evidently, allowing longer-term pass-through rates (typically higher than short-run pass-through rates) to inform our Step 2 estimates does not have a large effect on the coefficient in any country. In a similar vein, the results in row 6 aggregate the data used to estimate pass-through rates up to the quarterly level and proceeds as before with this aggregated data. This acts to reduce the coefficients (relative to baseline) in all three countries,

only the 45th percentile in the US, a much larger country.

[40] For the US sample, where we only have 6 years of data, we divide the sample into 3- and 1.5-year time periods respectively in rows 3 and 4.

but the relative effect of distance on trade costs in each country is very similar to those in rows 1-5.

A natural concern with our pass-through estimation procedure is that a shock to the origin price is correlated with an unobserved macro shock that could affect trade costs and hence affect locations differently. One possible candidate for such a shock is the price of oil, which would raise the price of gasoline and hence the relative cost of accessing remote locations, leading our procedure to over-estimate the pass-through rate.[41] In row 7 we therefore present results obtained while estimating pass-through rates using regressions that control for the world oil price (a simple average of Brent, Dubai and West Texas spot prices converted into local currency for each country and then inflation-adjusted) in each of the product-destination pass-through regressions; on average this has only a minor impact on our estimates.

We now turn to an extension to our baseline procedure for carrying out Step 2 above. Row 8 estimates equation (15) while directly including fixed effects for each destination location d (recall from equation (15) that in our baseline procedure these fixed effects were only included as interactions with $\frac{(1-\rho_{od}^k)}{\rho_{od}^k}$). This is important because we suspect that many components of trade costs (such as local distribution or retail costs, the non-tradable components of traded retail goods) may co-vary with distance to the origin locations. It is therefore noteworthy that our coefficient estimates change very little after including destination fixed effects, though the standard errors rise substantially in some cases.[42] This suggests that spatial variation in location-specific trade costs is not correlated with spatial variation in the distance between average trading pairs.

The next set of extensions alters our entire estimation procedure (that is, both Steps 1 and 2). Rows 9 and 10 increase the set of fixed effects included in Step 1 and 2 to account for more general demand and trade costs structures which could confound our main estimates. Recall from equation (15) that, once one works with adjusted price gaps as the dependent variable, the only demand shocks that could potentially bias estimates of trade costs (those which shift the inverse demand curve, i.e. a_{dt}^k) enter equation (15) interacted with a term involving the pass-through rate, $\left(\frac{1-\widehat{\rho_{od}^k}}{\widehat{\rho_{od}^k}}\right)$. In rows 9 and 10 we therefore interact two more exhaustive sets of fixed effects, year-location γ_{yt} and time-location fixed effects γ_{dt}, respectively, in lieu of the destination fixed effects in equation (15). Additionally, these two specifications also include the same fixed effects in Step 1, adding either year-location or time-location fixed effects to estimating equation (13). The estimates change very little for either the year-destination fixed effect specification (row 9) or the time-destination

[41] Given that oil is also produced and processed close to some of the origin locations in two of our sample countries (Nigeria and the US) a related concern is that the world price for oil affects factor prices in destination locations relatively near to oil-engaged origin locations, which then affect transport costs to these destination locations; alternatively factor price effects could increase demand in these destination locations.

[42] Because the main regressor of interest here—log origin-destination distance—is fixed over time, the only variation that remains after controlling for destination fixed effects in row 8 is that generated by the fact that products come from distinct origin locations. So for a country like Nigeria, where (as is apparent in Figure 1), most products are produced in the same (or extremely proximate) origin locations, it is not surprising that the standard error in row 8 is considerably higher than in the baseline (row 1). In fact, in the US case there is insufficient variation to include destination-specific fixed effects and simultaneously cluster the standard errors at the product-time period level; the reported standard error in row 8 of Table 3 for the US therefore reports robust standard errors that are not clustered (and similarly for row 12 that we will discuss shortly). Our main concern here is with the coefficient estimate, however, which is unaffected by these considerations.

specification (row 10). (Unfortunately, the latter of these two specifications is computationally infeasible for the US since it would entail estimating a regression with several million fixed effects.)

Rows 11 and 12 continue to examine the possibility that our estimates are biased due to covariance between trade costs and unobserved demand shocks. As discussed above, both Step 1 and Step 2 required us to assume that demand shocks at the origin location are uncorrelated with those at the destination location (conditional on demand shocks that are common across products). Because unobserved demand shocks (be they due to taste, migration or income shocks) are likely to be spatially correlated, our identifying assumption seems increasingly plausible for destination locations that are increasingly remote from origin locations. In row 11 we therefore repeat our entire estimation procedure (that is, both Steps 1 and 2) for a sample that contains destination locations that are more than 100 miles from the origin; row 12 performs the analogous check with this exclusion raised to 200 miles. That the estimated coefficients do not change substantially as we increase this exclusion band from 0 miles (row 1) to 100 miles (row 11) to 200 miles (row 12) therefore suggests that our identifying assumption is plausible.

Finally, we explore two miscellaneous extensions to our main analysis. First, in row 13 we report an estimate that is based on an entirely raw data sample. This suggests that the various data cleaning procedures described in Section 2.1, while potentially important *a priori*, do not appear to matter much for our central parameter estimate. Lastly, in rows 14 and 15 we report specifications that interact log origin-destination distance x_{od} with the log weight (per unit) of the product k in question; row 14 reports the level effect of log distance and row 15 reports the coefficient on the interaction term (the level effect of the product's weight is subsumed by the product-time period fixed effects that we use throughout). While a full exploration of how trade costs differ across products is beyond the scope of this paper, we find it reassuring that our procedure picks up, in each of our three sample countries, the common-sense notion that heavier goods are costlier to trade, and statistically significantly so.

All told, the general message of Table 3 is that, regardless of the precise form that a number of important identifying assumptions take, or the exact implementation of a number of empirical details, our baseline parameter estimates are relatively stable. We therefore go on, in the next subsection below, to a deeper discussion of our baseline results.

4.5 Discussion of Results

We now describe additional context behind our baseline results from Section 4.3 above. We first explore some possible explanations for relatively high costs of distance in our African countries, and then discuss how our results compare to other estimates in the literature.

Unpacking the effect of distance on trade costs

Figure 3 and Table 2 document how the impact of a given unit of (log) distance on intranational trade costs is approximately 4-5 times higher in our African sample countries (Ethiopia and Nigeria) than in the US. One potential explanation for our findings is that there are simply more roads in the US than in Nigeria. We therefore re-estimate the costs of distance using the quickest-route distance measure described in Section 2.3. If part of the US's trade cost advantage

lies in the omnipresence of roads, we should see this advantage fall relative to Ethiopia and Nigeria when we use the road distance metric. Table 4 reports how our estimate of the cost of (log) distance in each of our African sample countries (Ethiopia and Nigeria) changes as we shift from a distance metric that is based on road distance (along the quickest route) rather than straight-line (i.e. great circle) distance.

In the case of Ethiopia we find that the relative coefficients fall but by a small amount—from 3.53 times higher cost of (log) distance in Ethiopia relative to the US, to 3.19 times higher. But in the case of Nigeria the relative coefficient is actually slightly higher (5.26 to 5.40). This suggests that adjusting for the mere presence of roads does little to explain why the costs of distance are so much higher in Ethiopia or Nigeria than in the US.

A natural continuation of this explanation is to adjust for the quality of roads in the US relative to our African sample countries—we certainly expect the US to be relatively endowed with high-quality roads. To adjust for road quality we use the quickest-route travel time measure described in Section 2.3. As reported in Table 4, when we use travel time (along the quickest route) from origin to destination as our 'distance' metric, the effect of (log) distance on trade costs is now 2.46 times higher in Ethiopia than in the US, and 4.01 times higher in Nigeria than in the US. So adjusting—to the best of our ability—for speed of travel along the roads in our three countries reduces the gap between our African countries and the US, but a considerable gap remains. This suggests that even the cost per (log) hour of travel is substantially smaller in the US than in our African countries, a finding that is surprising given that we expect higher wages (or other factor prices) to lead to higher costs of time in the US.

Following from this, it is important to recall that, in looking at how retail prices vary over space, our approach has identified all-encompassing trade costs—the full cost of getting a good from its origin o to its destination d. As a consequence, our estimates of the cost of distance include any systematic variation in local retail and distribution costs between origin and destination locations. If our sample countries have differing spatial gradients of local retail/distribution costs (perhaps due to differing rent or wage gradients) or local retail productivity, our results could be picking up these differences. However, as we have seen (in Table 3) the inclusion of destination location fixed effects does very little to change our baseline point estimates. Spatial variation in factor prices does not seem to be strongly correlated with (log) origin-destination distances in any of our three sample countries, and hence this is not a successful explanation for our finding that the costs of distance appear to be considerably higher in Ethiopia or Nigeria than in the US.

There remain a range of other factors that could potentially contribute to higher intra-national trade costs within African countries yet not operate through road distance, travel time or location-specific factors. For example, commentators on African transport (e.g., Teravaninthorn and Raballand (2009)) have highlighted inferior technology, both through the use of old truck fleets that are fuel-inefficient, terrible road conditions that necessitate frequent truck repairs, and poor logistics. Fuel costs are also generally higher in Africa, due in part to a lack of local refining capacity. Additionally, many routes are characterized by low payload utilization and a low total number of miles traveled per trip, reducing any economies of scale. Finally, there are long waiting times for loading

and unloading as well as frequent checkpoints (whose time costs are not accounted for in Google Maps algorithms for calculating travel times in Africa) often accompanied by bribe demands.

Comparisons to existing literature

We now discuss how our trade costs estimates compare to those in the existing literature. We are not aware of any previous work that has estimated how origin-destination price gaps depend on distance (or any other cost-shifter x^k_{odt}), nor any work that has attempted to purge such price gap inferences of spatially-varying mark-ups. Numerous studies have documented how spatial price gaps co-vary with distance but, as is clear from Figures 2 and 3, our new data on origin locations suggest that, within our sample at least, spatial price gaps can provide misleading estimates when not restricted to trading pairs or when failing to correct for variation in mark-ups over space.

An alternative method for estimating trade costs, however, is to simply ask transport firms what they would charge for a shipment (or transportation-using firms what they pay for such services). The well-known difficulty with this method is that it measures the price that traders charge their customers rather than those traders' marginal costs. In addition, there are supplementary costs of trading (such as regulatory barriers or local distribution costs) that are not borne by the surveyed transport firms, as well as aspects of quality (time in transit, uncertainty, damage or loss of goods in transit) that are difficult for surveyors to measure. But a relative comparison of such estimates across countries may still be meaningful if the price-cost margin charged by traders, and the proportion of costs that are unobserved, is similar across countries. A particularly relevant study containing such results is Teravaninthorn and Raballand (2009). These authors survey trucking companies in the US and along major transport corridors within sub-Saharan Africa. They report an average of costs per unit distance in the US, as well as along one West African trucking corridor (Bamako-Accra) and one East African corridor (Mombasa-Nairobi). While these main trucking arteries do not pass through our sample countries, our hope is that estimates based on these journeys nevertheless provide a useful comparison. As shown in Table 4, the survey estimates, expressed as ratios of East and West African costs to US costs, are similar to—though somewhat lower than—our trade costs estimates for Nigeria and Ethiopia respectively.[43] Furthermore, both our estimates and those in Teravaninthorn and Raballand (2009) imply that costs in West Africa are around twice those in East Africa.

How markups vary over space

We now turn to a discussion of what our results imply for the variation in mark-ups over space. Our finding that the effect of distance on trade costs is larger once we account for mark-up variation across space implies that remote locations feature lower mark-ups. At first glance this is surprising. Appendix Figure C.2 shows that pass-through rates tend to be lower in more remote locations, suggesting that these locations are less competitive and mark-ups there should

[43] One possible explanation for the fact that the Teravaninthorn and Raballand (2009) estimates are lower than our estimates is that they focus on some of the regions' relatively high-quality transport corridors rather than the full set of roads, major and minor, that underpin our estimates.

be higher rather than lower.[44] However, there are two other channels at work that can generate the decline in mark-ups with distance. Most directly, because we estimate that pass-through rates are below one for the vast majority of product-location pairs, the higher trade costs $\tau(\mathbf{X}^k_{odt})$ in remote locations move consumers almost everywhere onto a more elastic portion of their demand curve, leading intermediaries to charge lower mark-ups in remote locations. In addition, remote locations may be poorer and so have lower demand shifters a^k_{dt}, also lowering mark-ups.

Figure 4 explores the extent to which these two possibilities are driving our result. Here we plot three lines. First, with a solid line we plot the empirical derivative of mark-ups with respect to distance, as recovered from the gap between the two curves in Figure 3. Second, with a dotted line we plot the empirical derivative of a^k_{dt} with respect to distance, as recovered from the demand-shifter fixed effects $\alpha^k_t + \alpha_d$ in specification 15. And finally, with a dashed line we plot the empirical function $-\tau(\mathbf{X}^k_{odt})$, with respect to distance (as presented earlier, but without the minus sign, in Figure 3). As can be seen, there is limited evidence that demand shifters are systematically lower in more remote locations. Hence, our finding of lower mark-ups in more remote locations is predominantly driven by the fact that the high trade costs of reaching those locations moves remote consumers onto more elastic parts of their demand curve.

5 Implications for the Incidence of Globalization

As tariffs fall and international transportation and logistics improve—events often termed 'globalization'—the port price of an imported product will fall. A natural question then arises: Who captures the benefits of these port price changes? Within a given location, do the gains accrue to consumers or producers? And how does this vary across locations? In a perfectly competitive economy with no internal trade barriers, it is clear that all consumers, regardless of their locations, enjoy the benefits, and do so equally. But given the evidence for both high internal trade costs and imperfect competition we have seen in Section 4 above, the answers to these questions in our setting are more nuanced.

We provide here a two-part investigation of these questions. First, we consider the distribution of surplus across locations purely due to the extensive margin. Because of high costs of intra-national trade, it is entirely possible that remote locations never import the product in question. That is, remote locations may enjoy no consumer or producer surplus from the product in question. Second, we consider the distribution of surplus between producers and consumers (and deadweight loss), and how the relative share of surplus accruing to consumers changes across locations. Our findings below suggest that relatively remote consumers are also relatively separated from the gains created by globalization: they live in locations with less total social surplus and capture a smaller share of the total social surplus that does exist in these locations. In this sense, our findings on internal trade costs in Section 4 imply an uneven distribution of the gains from globalization, both across locations and between consumers and intermediaries.

[44] To see this, note that under Assumptions 1-4 the size of the mark-up, $\mu^k_{odt} = (1-\rho^k_{od})(a^k_{dt} - \tau(\mathbf{X}^k_{odt}) - P^k_{ot})$, is a function of the pass-through rate and the gap between the demand-shifter a^k_{dt} and the intermediary's total cost, $\tau(\mathbf{X}^k_{odt}) + P^k_{ot}$.

Before going forward, it is important to keep two caveats in mind. First, for our African sample countries no data on consumption *quantities* are available for the narrowly defined products that we study here. For this reason we speak only about the relative share of surplus accruing to different agents rather than absolute amounts. Second, while our central interest here is in the incidence of globalization (the distribution of surplus from an imported good) the estimates so far were obtained from a sample of largely domestic goods.[45] However, because much domestic goods production takes place in main port cities, and because the same intermediaries who distribute domestic goods will also distribute foreign goods, we believe that the estimates in Section 4 are useful for the questions posed here.

5.1 Implication 1: Remoteness and the size of the surplus

The estimates in Figure 3 (solid line) suggest that intra-national trade costs rise substantially with distance in Ethiopia and Nigeria. This implies that remote consumers in these countries, living far from a major port, would pay substantial trade costs—and also substantially higher prices, as also seen in Figure 3 (dashed line)—to purchase foreign goods. We therefore expect these consumers to consume, all else equal, lower quantities of imported products and hence obtain less consumer surplus from them. In the extreme, the high cost of distance may even lead to the imported product not being available in remote locations.

Without data on the quantity of products sold we are limited in our ability to estimate the magnitude of this effect. But we are able to provide suggestive evidence that this extensive margin affects remote locations by exploiting the fact that, in Ethiopia and Nigeria, a missing price observation indicates that the CPI enumerators were unable to find that product (in any given location and month), which indicates that consumption was likely to be zero or minimal.[46] Figure 5 displays how this—admittedly imperfect—proxy for zero consumption relates to distance in our two sample African countries. The y-axis reports an indicator variable for whether the enumerator found the product in that location for a given month.[47] It is clear that remote locations, those that are far from the origin location of a given product, are considerably less likely to consume that product. Columns 1 and 6 of Table 5 confirms this negative relationship using a linear probability model that regresses an indicator variable for product availability on distance from the origin location and a time-product fixed effect. For example, in Ethiopia the probability of enumerators locating a product in a particularly remote location (500 miles away) is 22 percent lower than in one of the most proximate locations (50 miles away); the corresponding figure for Nigeria is 11 percent.

5.2 Implication 2: Remoteness and the distribution of surplus

We now consider a second question: If there is social surplus from an imported product in a location, how is that surplus distributed among market participants (consumers and intermediaries)? That is, we consider the incidence (across agents) of a change in the port price for the

[45] As shown in Appendix Table C.1, 5 of the 15 Ethiopian goods, and 3 of the 18 Nigerian goods, are imported.

[46] Note that such an exercise is not possible for the US sample since no enumerators are sent to search for products but instead a small number of households in each county are sampled.

[47] We restrict attention to products that were found in at least one location in the country in that month.

product within a given remote location. Because this question concerns *shares* of surplus, our lack of quantity data is less of a constraint, as we discuss below.

At the outset it is important to note that there are *a priori* reasons to expect that remote locations may see larger shares of surplus accruing to intermediaries. Consider the theoretical framework from Section 3 above, where to simplify notation we revert temporarily to the setting with just one location, product and time period. One can formally show that a sufficient condition for profits Π_d to be decreasing in the trade cost-shifter x_d is $\rho_d < 1$. Since the vast majority of our estimated pass-through rates satisfy this condition, we expect there to be lower profits in relatively remote locations. These lower profits will lead to less entry in the long run thereby reducing the level of competition among intermediaries and thus, all else equal, increasing the share of surplus accruing to these intermediaries.

Naturally, it is challenging to identify the shares of surplus accruing to consumers and intermediaries. Fortunately, based on an extension of the logic in Weyl and Fabinger (2011),[48] in the theoretical framework we have outlined there exists a simple connection between pass-through and the division of surplus, which allows an answer to this question without the need for data on consumption quantities.

As we show formally in Appendix B, under Assumptions 1-3 we can express the ratio of intermediary variable profits, IS_d,[49] and consumer surplus, CS_d, as

$$\frac{IS_d}{CS_d} = \frac{1}{\rho_d} + \frac{1-\phi_d}{\phi_d}, \tag{16}$$

which means that the share of surplus enjoyed, in equilibrium, by intermediaries relative to consumers is a simple function of just the equilibrium pass-through rate ρ_d and the competitiveness index ϕ_d prevailing in that market. Conditional on obtaining estimates of ρ_d and ϕ_d, therefore, equation (16) provides a direct estimate of the division of surplus.[50] To obtain some intuition for why the distribution of surplus can be captured by such a simple expression, consider a monopolist intermediary ($\phi_d = 1$) facing a small change in the port price. By the envelope theorem the change in intermediary profits is simply equal to the quantity Q_d. Meanwhile the change in consumer surplus is equal to the quantity consumed multiplied by the change in the destination price as a result of the changing port price, or simply $Q_d\rho_d$. Hence the ratio of intermediary to consumer surplus is equal to the inverse of the pass-through rate as in equation (16) when $\phi_d = 1$.

In order to implement the calculation suggested by equation (16) we therefore require esti-

[48] This result extends the monopoly analysis in Weyl and Fabinger (2011) to the case of oligopoly. Subsequent versions (e.g. Weyl and Fabinger (2013)) contain a similar formulation derived independently.

[49] We work with a notion of surplus defined on total variable profits because nothing in our dataset can be used to estimate the fixed costs intermediaries pay.

[50] An additional result that holds under Assumptions 1-3 has been generously brought to our attention by Glen Weyl: the ratio of deadweight loss (DWL) to intermediaries' surplus in this environment is given by

$$\frac{DWL_d}{IS_d} = (1-\rho_d) + \rho_d\phi_d - \left(\frac{\rho_d\phi_d}{(1-\rho_d)+\rho_d\phi_d}\right)^{\frac{\rho_d}{1-\rho_d}}(\rho_d\phi_d + 1).$$

This result, together with that in equation (16) allows for a straightforward decomposition of total social surplus into that accruing to consumers (CS_d), intermediaries (IS_d) and to deadweight loss (DWL_d).

mates of both the pass-through rate ρ_d and the competitiveness index ϕ_d. We therefore revert our notation to that of the empirical setting, with multiple products k, origin-destination trading pairs od, and time periods t (but where, following Assumption 4, the pass-through rate ρ^k_{od} and competitiveness index ϕ^k_{od} are constant over time; this then implies that $\frac{IS^k_{od}}{CS^k_{od}}$ will also be constant over time). Recall that ρ^k_{od} and ϕ^k_{od} are connected through the formula:

$$\rho^k_{od} = \left(1 + \frac{\delta^k_d}{\phi^k_{od}}\right)^{-1}, \tag{17}$$

where δ^k_d is the preference parameter that captures the curvature of the demand curve (the elasticity of the slope of demand) in the constant pass-through demand class from Assumption 3. Equation (17) suggests that estimates of ρ^k_{od} could be used to estimate ϕ^k_{od}. Unfortunately, in general (that is, without further restrictions on δ^k_d) there is no unique mapping between ρ^k_{od} and ϕ^k_{od}. We therefore assume (in Assumption 5 below) that the variation in the demand-side determinants of pass-through (i.e. the parameters δ^k_d) and the supply-side determinants of pass-through (i.e. the parameters ϕ^k_{od}) are sufficiently orthogonal over destination markets d and products k as to allow data on the pass-through rate (i.e. an estimate of ρ^k_{od}) to identify ϕ^k_{od}. However, the particular assumption made in Assumption 5, while easy to interpret, is overly sufficient since it restricts there to be only $N + T$ unknown parameters to be estimated from NT pass-through estimates.[51]

Assumption 5 [Identification of δ and ϕ]. *The demand parameter δ^k_d is constant over destination locations d but can vary freely across products k; that is, δ^k_d=δ^k for all d. Similarly, the competitiveness index parameter ϕ^k_{od} is constant over products k but is free to vary across destinations d; that is, ϕ^k_{od}=ϕ_{od} for all k.*

This is a particularly stark assumption, but one that is perhaps not implausible. Because of economies of scale it seems reasonable that the same intermediary will supply multiple goods to a given location. Consequently, the essential variation in the number of intermediaries and their competitive conduct (and hence the overall competitiveness index ϕ^k_{od}), is primarily across locations rather than across products within a location. Likewise, while we allow the additive and multiplicative shifters of demand (i.e. a^k_{dt} and b^k_{dt}) to vary across locations, products and time, it seems plausible that the second-order curvature parameter δ^k_{dt}, the unique demand-side parameter that governs pass-through, is constant across locations and time.

Assumptions 1-5 therefore imply that a consistent estimator of the competitiveness index at a destination (i.e. ϕ_{od}), up to a scalar, can be obtained by estimating the following regression by OLS

$$\Xi^k_{od} = \gamma_{od} + \gamma^{1k} + \gamma^{2k}\zeta^k_{od} + \varepsilon^k_{od}. \tag{18}$$

In this expression: letting $\widehat{\rho^k_{od}}$ be a consistent estimator of the equilibrium pass-through rate obtained in Step 1, if $\widehat{\rho^k_{od}} < 1$ then $\Xi^k_{od} \equiv \ln(\frac{1}{\widehat{\rho^k_{od}}} - 1)$, and if $\widehat{\rho^k_{od}} > 1$ then $\Xi^k_d \equiv \ln(1 - \frac{1}{\widehat{\rho^k_{od}}})$; γ_{od} is a destination fixed effect; γ^{1k} and γ^{2k} are product-specific fixed effects; ζ^k_{od} is a dummy variable that equals 1 if $\widehat{\rho^k_{od}} < 1$ and zero otherwise; and ε^k_{od} is an error term (due to the finite sample variance

[51] Recall, N is the total number of destinations and T the total number of time periods.

in our estimates $\widehat{\rho_{od}^k}$). This procedure effectively projects the estimated pass-through rates on location and product fixed effects. We normalize the estimated ϕ_{od} values such that the lowest in the sample (within each country) is $\phi_{od} = 1$, which amounts to a normalization of the least competitive location to a monopolistic scenario. Subject to this normalization, a consistent estimator of the competitiveness index in any location d is $\widehat{\phi_{od}} \equiv e^{-\widehat{\gamma_{od}}}$.

A natural question to ask is how well these estimates agree with external proxies for competitiveness. One possible proxy is the number of intermediaries who are active in a given location, which we expect to be positively correlated with the competitiveness index ϕ_{od} (since, by the definition in Assumption 2, $\phi_{od} \equiv \frac{m_{od}}{\theta_{od}}$, where m_d is the number of intermediaries in location d). Indeed, if intermediaries compete with Cournot conjectures then $\theta_{od} = 1$, and the competitiveness index ϕ_{od} is simply a measure of the number of intermediaries m_{od}. In Ethiopia we have data from two cross-sectional surveys collected in 2001 and 2008 that report an empirical correlate of m_d, namely the number of wholesale trading firms that are active in each location d. Panel A of Figure 6 illustrates how this empirical proxy for m_d correlates with our estimate of $\widehat{\phi_{od}}$ for each destination location d in Ethiopia. The positive and statistically significant correlation between these two variables suggests that our procedure, and Assumption 5 in particular, generates plausible estimates.

Panel B of Figure 6 shows nonparametric plots of how the estimated competitiveness index $\widehat{\phi_{od}}$ varies with (log) distance to the main commercial city (Addis Ababa or Lagos) in each of our two African sample countries.[52] There is clearly a downward-sloping relationship, and the parametric equivalent of this descriptive relationship—presented as columns 2 and 7 of Table 5—suggests that the negative slope is statistically significant in each country. That is, more remote locations have a less competitive intermediary sector serving them. This is consistent with the argument offered above: as long as $\rho_{od}^k < 1$, profits Π_{odt}^k will be falling with a cost-shifter such as distance and hence we should expect the number of entrants to be similarly falling with distance.

We now go on to use our estimates of pass-through $\widehat{\rho_{od}^k}$ and competitiveness $\widehat{\phi_{od}}$ in equation (16) to provide a consistent estimate of the ratio of intermediaries' surplus to consumer surplus. We calculate $\frac{IS_{od}^k}{CS_{od}^k}$ for each location d and product k and present in Figure 7 a nonparametric regression of this statistic on (log) distance from origin to destination. We use an analogous calculation to report the share of consumer surplus in total surplus and how this varies by location. The general relationship is remarkably similar in each of our sample countries: Intermediaries capture a larger share of the surplus (relative to consumers) in remote locations relative to more proximate locations. The total share of surplus enjoyed by consumers (that is, not going to intermediaries or to deadweight loss) is relatively low in remote locations. Again, we present descriptive regressions in Table 5 that capture the parametric equivalents of these figures. These estimates imply that the additional share of surplus going to consumers in the least remote locations (50 miles away) compared to the most remote locations (500 miles away) is 4 percent in Ethiopia, 13 percent in Nigeria and 1 percent in the US.

[52] Such an exercise is less sensible for the US since there no single city that so dominates economic activity such that distance from it is a good measure of remoteness. Column 11 of Table 5 presents estimates based on distance to Chicago (an important production and distribution hub); there is a positive but statistically insignificant relationship.

Overall, these findings suggest that consumers in remote locations are doubly harmed by their remoteness. First, as argued in Section 5.1 above, consumers who are relatively remote from origin locations face a considerably lower chance of finding an imported good available for consumption at all. This suggests that the total quantity of surplus available to interior residents from a reduction in the world price of a product is lower in relatively remote locations. Second, as we have seen in this subsection, relatively remote consumers capture a lower share of the surplus that does exist in their location. In both senses, the incidence of 'globalization' is relatively unfavorable to consumers in remote locations.

6 Conclusion

In this paper we have set out to answer the question, *How large are the intra-national trade costs that separate consumers in remote locations of developing countries from global markets?* We find that the effect of distance on intra-national trade costs is substantially underestimated by standard spatial price gap methods used to infer trade costs. That is, the cost of distance approximately doubles when we discard uninformative price gaps, those price gaps for which neither of the pairs is a source location for the good in question. And the cost of distance approximately doubles again when we use a sufficient statistic, pass-through rates, to adjust spatial price gaps for spatial variation in mark-ups (implying that markups are lower in remote areas, a result of high trade costs placing remote consumers on more elastic parts of their demand curves). Our main finding is that the costs of intra-national trade are approximately 4 to 5 times larger in our sub-Saharan Africa sample countries (Ethiopia and Nigeria) than in the US. This has obvious implications for consumer welfare in these developing countries, particularly for those consumers whose relative remoteness—their location far from a country's major port, for example—means that as consumers they are connected to world markets only via these high intra-national trade costs.

Not only do these consumers therefore necessarily pay relatively more for imported goods, which reduces the amount of potential surplus these consumers can derive from foreign goods. But in addition—as we have documented using a methodology in which pass-through rates are again a sufficient statistic—of the surplus that remains once foreign products do arrive at remote locations, a relatively smaller fraction of that surplus accrues to consumers (instead of intermediaries and deadweight loss) than in locations near to the port of entry.

Like much research in developing countries, we are hampered by a lack of detailed consumption micro-data for the barcode-level products that we study here. Our goal therefore has been to develop a methodology for estimating and understanding intra-national trade costs that draws only on widely available price data. Yet we have shown how one can nevertheless make progress—despite embracing a general environment with minor restrictions on tastes, technology and market structure—because equilibrium prices and the extent to which they respond to cost shocks contain essential information about the marginal costs and benefits of the decisions intermediaries make, and those decisions reveal exactly what is needed to understand the size and the implications of intra-national trade costs.

References

AGNOSTEVA, D. E., J. E. ANDERSON, AND Y. V. YOTOV (2014): "Intra-national Trade Costs: Measurement and Aggregation," Working Paper 19872, National Bureau of Economic Research.

AHN, J., A. K. KHANDELWAL, AND S.-J. WEI (2011): "The Role of Intermediaries in Facilitating Trade," *Journal of International Economics*, 84(1), 73–85.

ANDERSON, J., AND E. VAN WINCOOP (2004): "Trade Costs," *Journal of Economic Literature*, 42(3), 691–751.

ANTRÀS, P., AND A. COSTINOT (2011): "Intermediated Trade," *The Quarterly Journal of Economics*, 126(3), 1319–1374.

ARKOLAKIS, C., A. COSTINOT, D. DONALDSON, AND A. RODRIGUEZ-CLARE (2012): "The Elusive Pro-Competitive Effects of Trade," Discussion paper.

ASSEFA, E. (2013): "Ethiopia spending two million USD per day for transit costs," *Addis Standard*, July 16th.

ATKESON, A., AND A. BURSTEIN (2008): "Pricing-to-Market, Trade Costs, and International Relative Prices," *The American Economic Review*, 98(5), 1998–2031.

BALTAGI, B. H., AND D. LI (2002): "Series Estimation of Partially Linear Panel Data Models with Fixed Effects," *Annals of Economics and Finance*, 3(1), 103–116.

BARDHAN, P., D. MOOKHERJEE, AND M. TSUMAGARI (2013): "Middlemen Margins and Globalization," *American Economic Journal: Microeconomics*, 5(4), 81–119.

BRODA, C. M., AND D. E. WEINSTEIN (2008): "Understanding International Price Differences Using Barcode Data," Discussion paper.

BULOW, J. I., AND P. PFLEIDERER (1983): "A Note on the Effect of Cost Changes on Prices," *The Journal of Political Economy*, 91(1), 182–185.

BURSTEIN, A., AND N. JAIMOVICH (2009): "Understanding Movements in Aggregate and Product-Level Real Exchange Rates," Discussion paper.

CHAU, N. H., H. GOTO, AND R. KANBUR (2009): "Middlemen, Non-Profits and Poverty," Discussion paper.

COSAR, A. K., AND P. D. FAJGELBAUM (2015): "Internal Geography, International Trade, and Regional Specialization," *American Economic Journal: Microeconomics*, forthcoming.

COSAR, A. K., P. L. E. GIECO, AND F. TINTELNOT (forthcoming): "Borders, Geography, and Oligopoly: Evidence from the Wind Turbine Industry," *Review of Economics and Statistics*.

COSAR, A. K., P. L. GRIECO, AND F. TINTELNOT (forthcoming): "Bias in Estimating Border-and Distance-related Trade Costs: Insights from an Oligopoly Model," *Economics Letters*.

DE LOECKER, J., P. K. GOLDBERG, A. K. KHANDELWAL, AND N. PAVCNIK (2012): "Prices, Markups and Trade Reform," Working Paper 17925, National Bureau of Economic Research.

DEATON, A. (1997): *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy*. Johns Hopkins University Press, Baltimore.

DONALDSON, D. (forthcoming): "Railroads of the Raj: Estimating the Impact of Transportation Infrastructure," *The American Economic Review*.

DU, Q., S.-J. WEI, AND P. XIE (2013): "Roads and the Real Exchange Rate," Working Paper 19291, National Bureau of Economic Research.

EATON, J., AND S. KORTUM (2002): "Technology, Geography, and Trade," *Econometrica*, 70(5), 1741–1779.

EDMOND, C., V. MIDRIGAN, AND D. XU (2011): "Competition, Markups, and the Gains from International Trade," Discussion paper.

FACKLER, P. L., AND B. K. GOODWIN (2001): "Spatial Price Analysis," *Handbook of Agricultural Economics*, 1, 971–1024.

FAJGELBAUM, P., AND S. J. REDDING (2014): "External Integration, Structural Transformation and Economic Development: Evidence from Argentina 1870-1914," Working Paper 20217, National Bureau of Economic Research.

FEENSTRA, R. (1989): "Symmetric pass-through of tariffs and exchange rates under imperfect competition: an empirical test," *Journal of International Economics*, 27(1-2), 25–45.

FEENSTRA, R. C., AND D. E. WEINSTEIN (2010): "Globalization, Markups, and the U.S. Price Level," Working Paper 15749, National Bureau of Economic Research.

GOLDBERG, P. K., AND R. HELLERSTEIN (2008): "A Structural Approach to Explaining Incomplete Exchange-Rate Pass-Through and Pricing-to-Market," *The American Economic Review*, 98(2), 423–429.

GOLDBERG, P. K., AND M. M. KNETTER (1997): "Goods Prices and Exchange Rates: What Have We Learned?," *Journal of Economic Literature*, 35(3), 1243–1272.

HOUTHAKKER, H. S. (1965): "New Evidence on Demand Elasticities," *Econometrica*, 33(2), 277–288.

LI, N., G. GOPINATH, P. GOURINCHAS, AND C. HSIEH (2011): "International Prices, Costs and Markup Differences," *The American Economic Review*, 101(6), 2450–2486.

MAYER, T., M. J. MELITZ, AND G. I. P. OTTAVIANO (2014): "Product Mix and Firm Productivity Responses to Trade Competition," Discussion paper.

MELITZ, M., AND G. OTTAVIANO (2008): "Market size, trade, and productivity," *Review of Economic studies*, 75(1), 295–316.

NEARY, P., AND M. MRAZOVA (2013): "Not so Demanding: Preference Structure, Firm Behavior, and Welfare," Discussion paper.

RAMONDO, N., A. RODRIGUEZ-CLARE, AND M. SABORIO-RODRIGUEZ (2012): "Scale Effects and Productivity Across Countries: Does Country Size Matter?," Working Paper 18532, National Bureau of Economic Research.

REDDING, S. J. (2012): "Goods Trade, Factor Mobility and Welfare," Working Paper 18008, National Bureau of Economic Research.

SEADE, J. (1980): "On the Effects of Entry," *Econometrica*, pp. 479–489.

SIMONOVSKA, I. (forthcoming): "Income Differences and Prices of Tradables: Insights from an Online Retailer," *Review of Economic Studies*.

SIMONOVSKA, I., AND M. E. WAUGH (2014): "The elasticity of trade: Estimates and evidence," *Journal of International Economics*, 92(1), 34 – 50.

TERAVANINTHORN, S., AND G. RABALLAND (2009): *Transport prices and costs in Africa: a review of the main international corridors*. World Bank Publications.

WEYL, E. G., AND M. FABINGER (2011): "A Restatement of the Theory of Monopoly," Discussion paper.

WEYL, E. G., AND M. FABINGER (2013): "Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition," *Journal of Political Economy*, 121(3), pp. 528–583.

WTO (2004): *World Trade Report* pp. 114–148. World Trade Organization, Geneva.

YOUNG, A. (2012): "The African Growth Miracle," *Journal of Political Economy*, 120(4), pp. 696–739.

ZHELOBODKO, E., S. KOKOVIN, M. PARENTI, AND J.-F. THISSE (2012): "Monopolistic Competition: Beyond the Constant Elasticity of Substitution," *Econometrica*, 80(6), 2765–2784.

Figure 1: Maps of sample locations

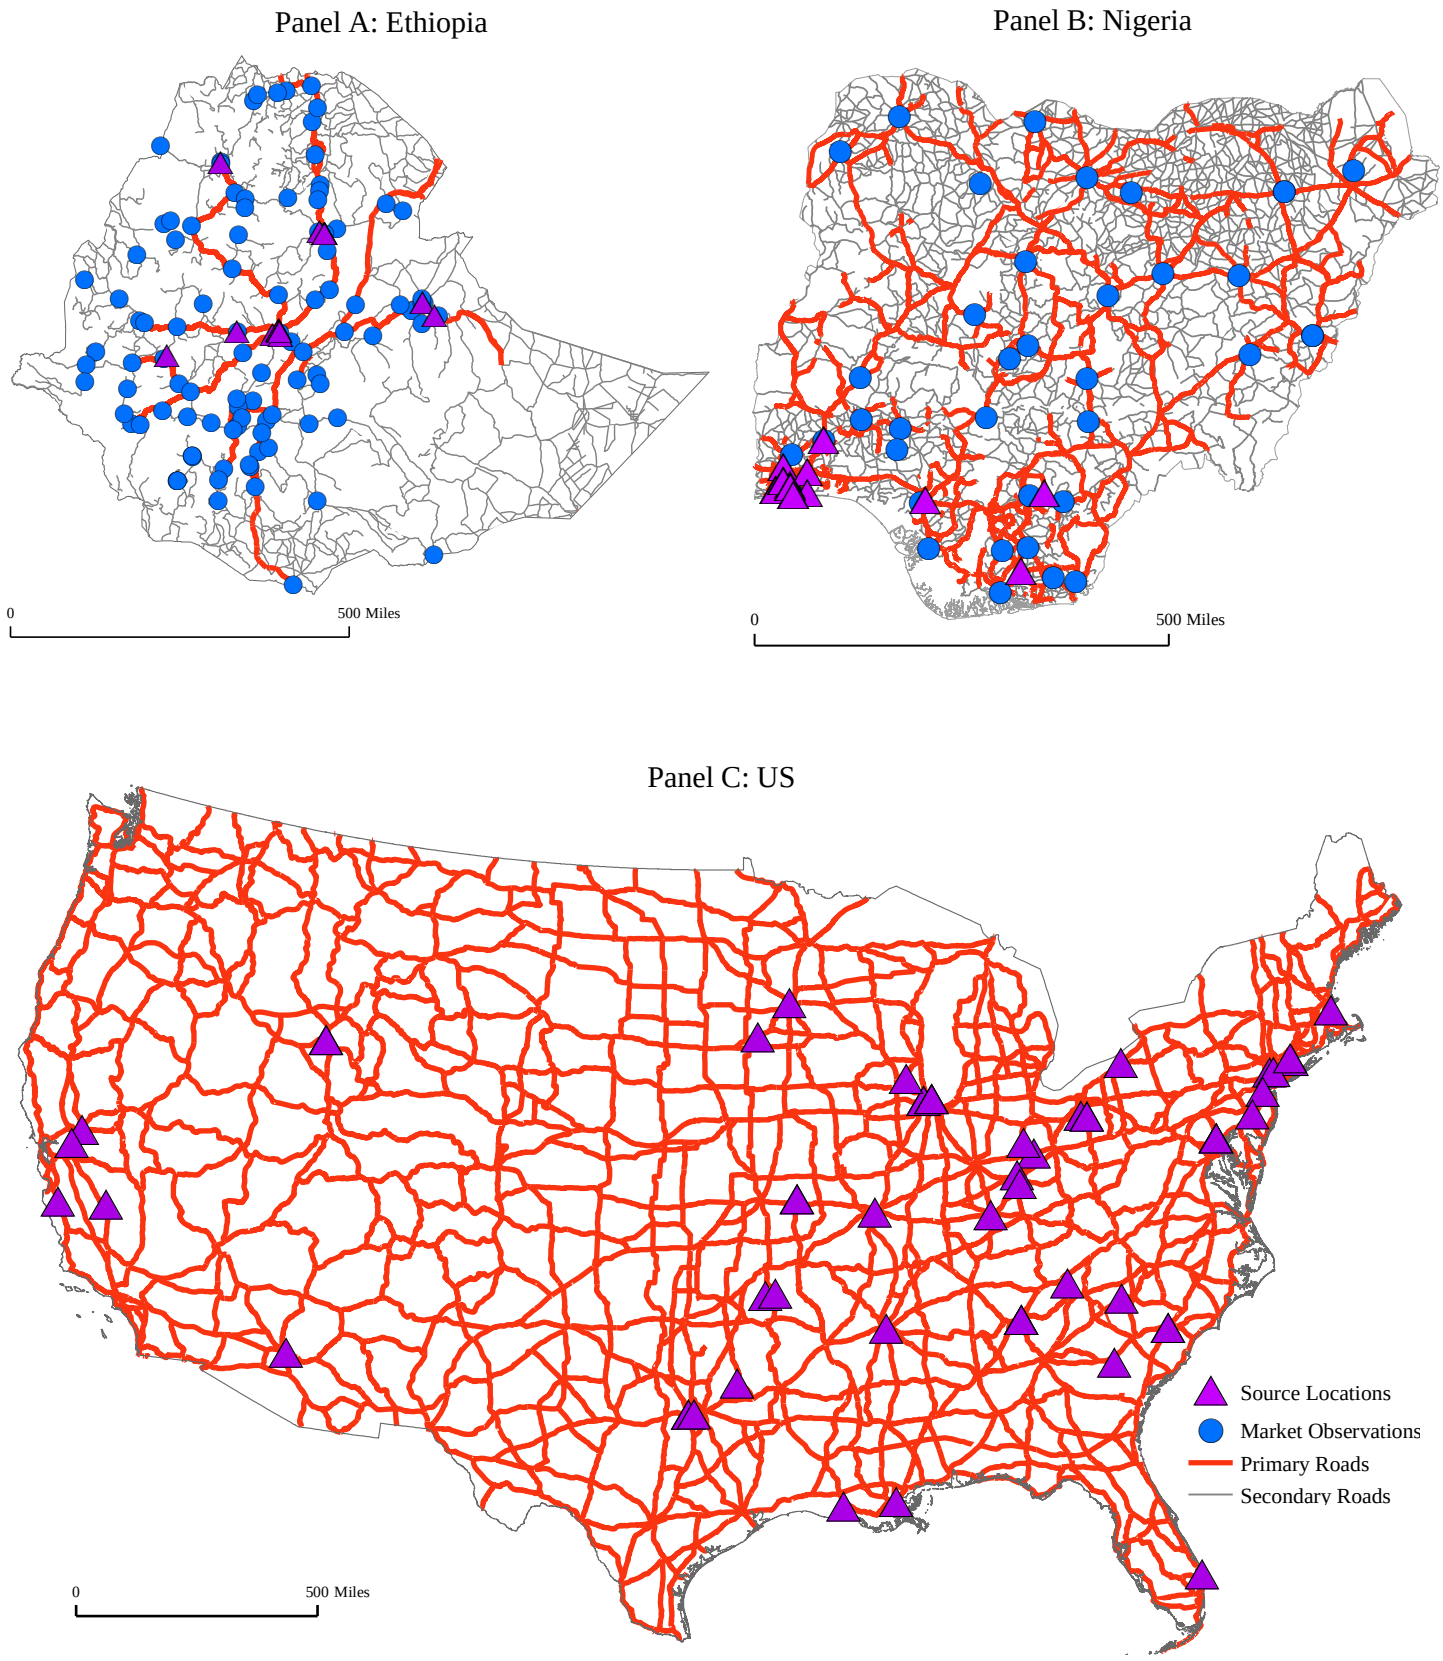

Notes: Market observations for the US are available for each of 1,881 counties so for clarity these are not illustrated.

Figure 2: Price gaps and distance

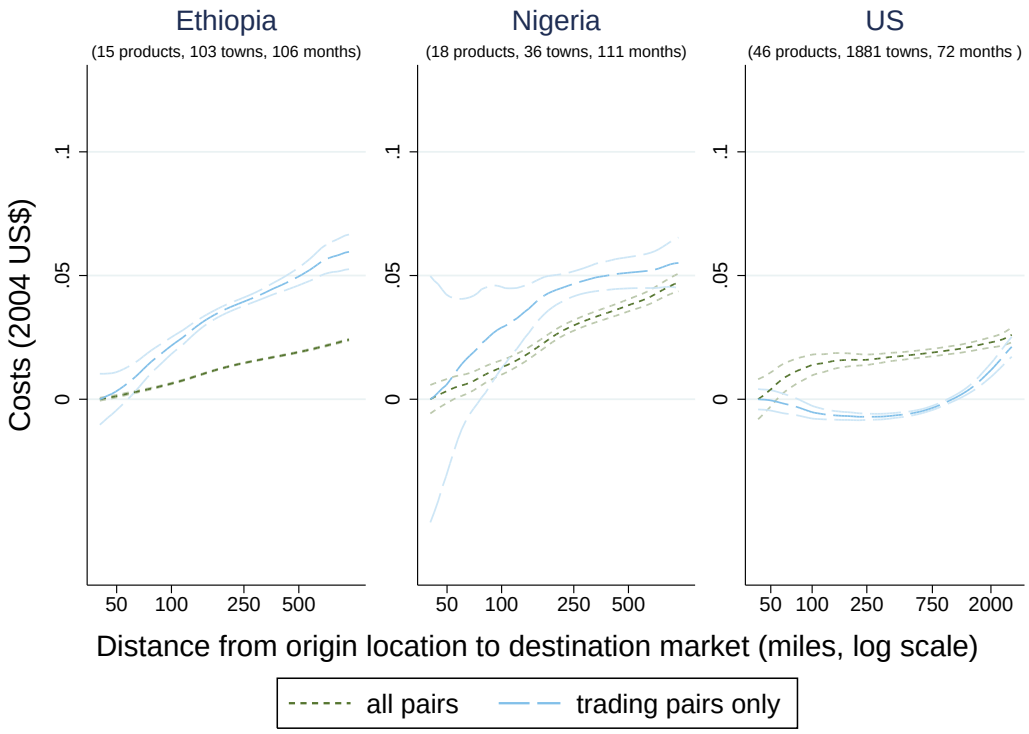

Notes: Locally weighted polynomial (Epanechnikov kernel, bandwidth=0.5). 95% confidence intervals shown (block bootstrapped at product-destination level). All plots control for product-time fixed effects. US plot uses compressed x-axis scale.

Figure 3: The effect of distance on intra-national trade costs

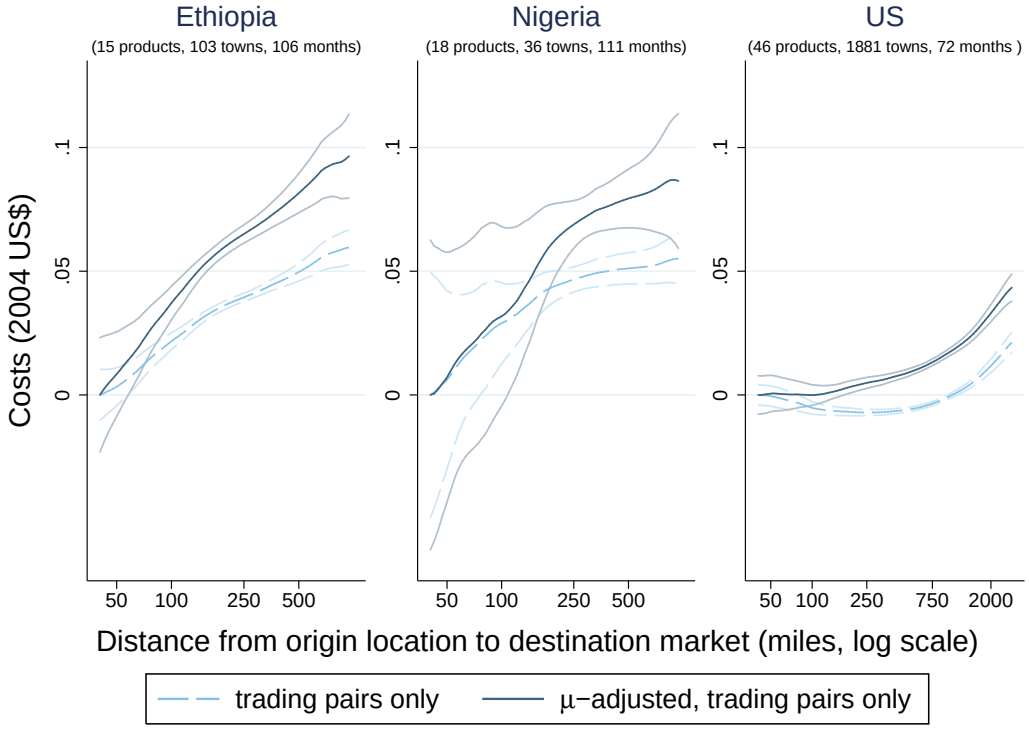

Notes: Locally weighted polynomial (Epanechnikov kernel, bandwidth=0.5). 95% confidence intervals shown (block bootstrapped at product-destination level). All plots control for product-time fixed effects. Markup-adjusted ("μ-adjusted") plot controls for interactions between pass-through and fixed effects as described in text. US plot uses compressed x-axis scale.

Figure 4: Variation in mark-ups across space

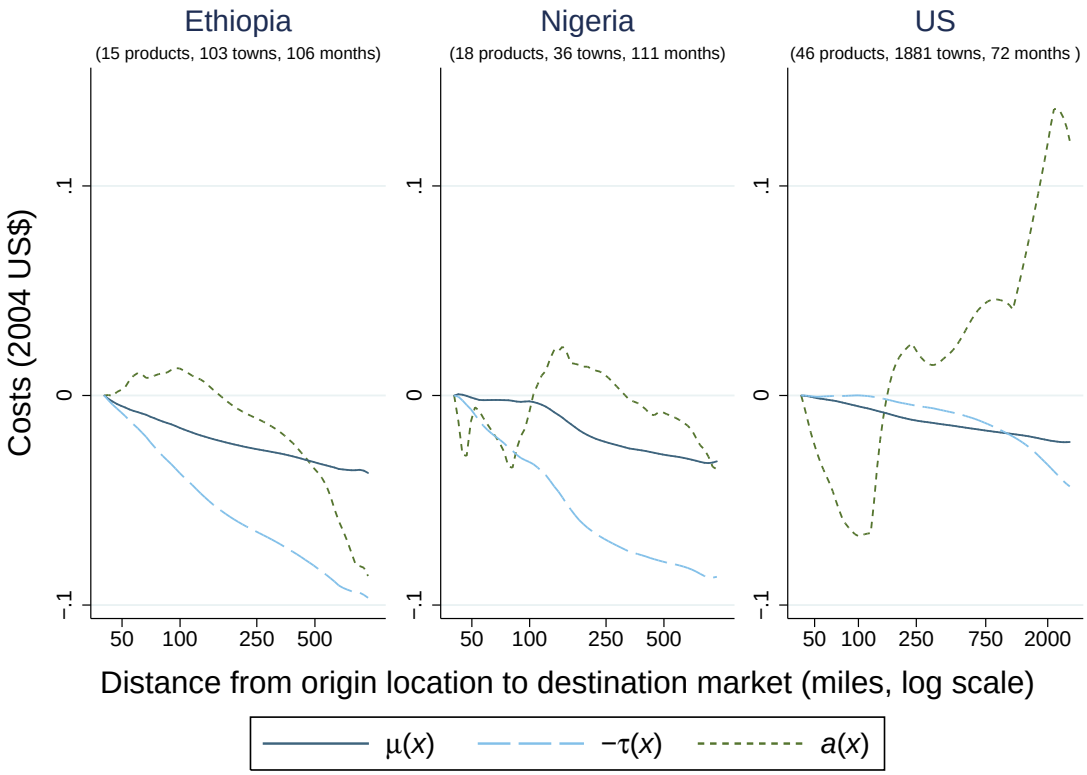

Notes: Locally weighted polynomial (Epanechnikov kernel, bandwidth=0.5). Semiparametric plots of demand shifters and trade costs from the adjusted price gap regression (i.e. equation 15). US plot uses compressed x-axis scale.

Figure 5: Product availability

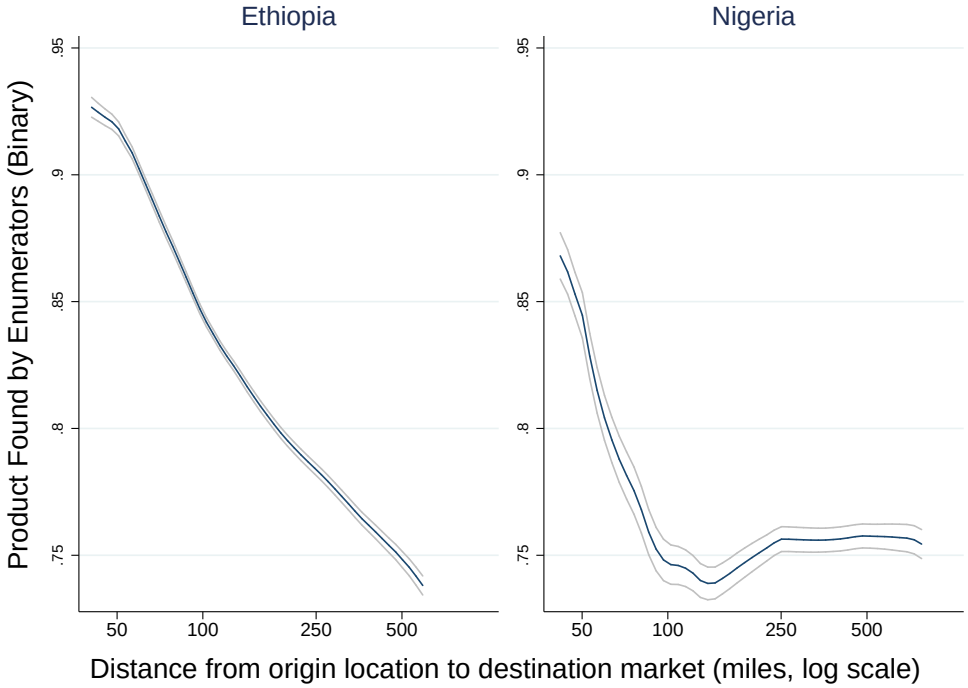

Notes: Locally weighted polynomial (Epanechnikov kernel, bandwidth=0.5). 95% confidence intervals shown. Linear probability model. Sample restricted to time-product pairs where product found in at least one location.

Figure 6: Competitiveness of intermediaries and distance

Panel A: Competitiveness index and Ethiopia Distributive Trade Surveys (2001 and 2008)

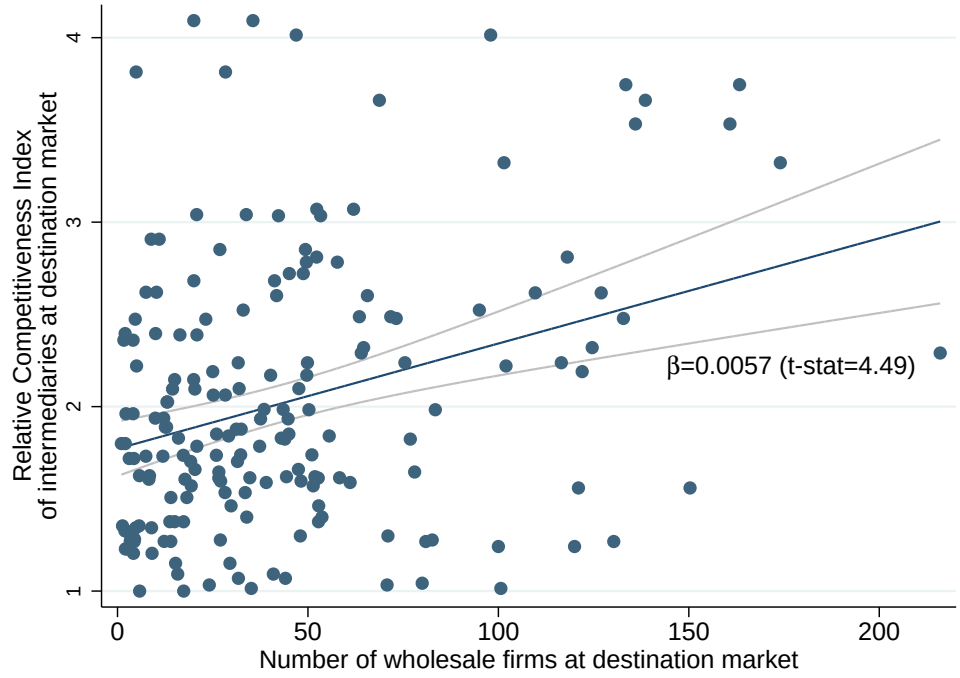

Notes: Plots of competitiveness index against number of firms (2001 and 2008 Ethiopian Distributive Trade Survey). 95% confidence interval shown.

Panel B: Relative competitiveness index and distance

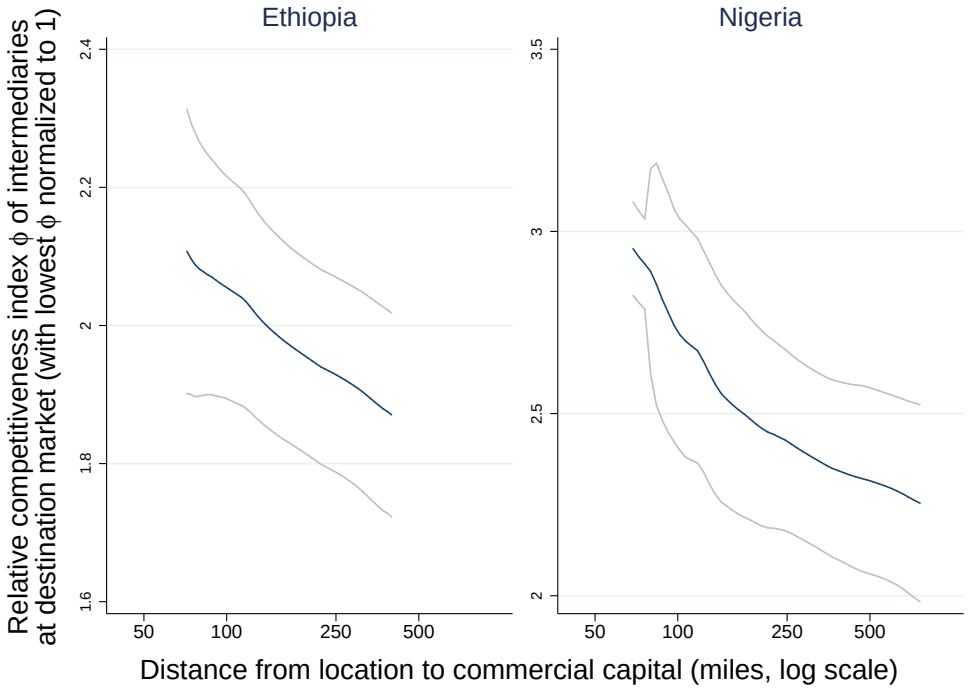

Notes: Locally weighted polynomial (Epanechnikov kernel, bandwidth=0.5). 95% confidence intervals shown.

Figure 7: Ratio of intermediary profit/deadweight loss to consumer surplus and distance

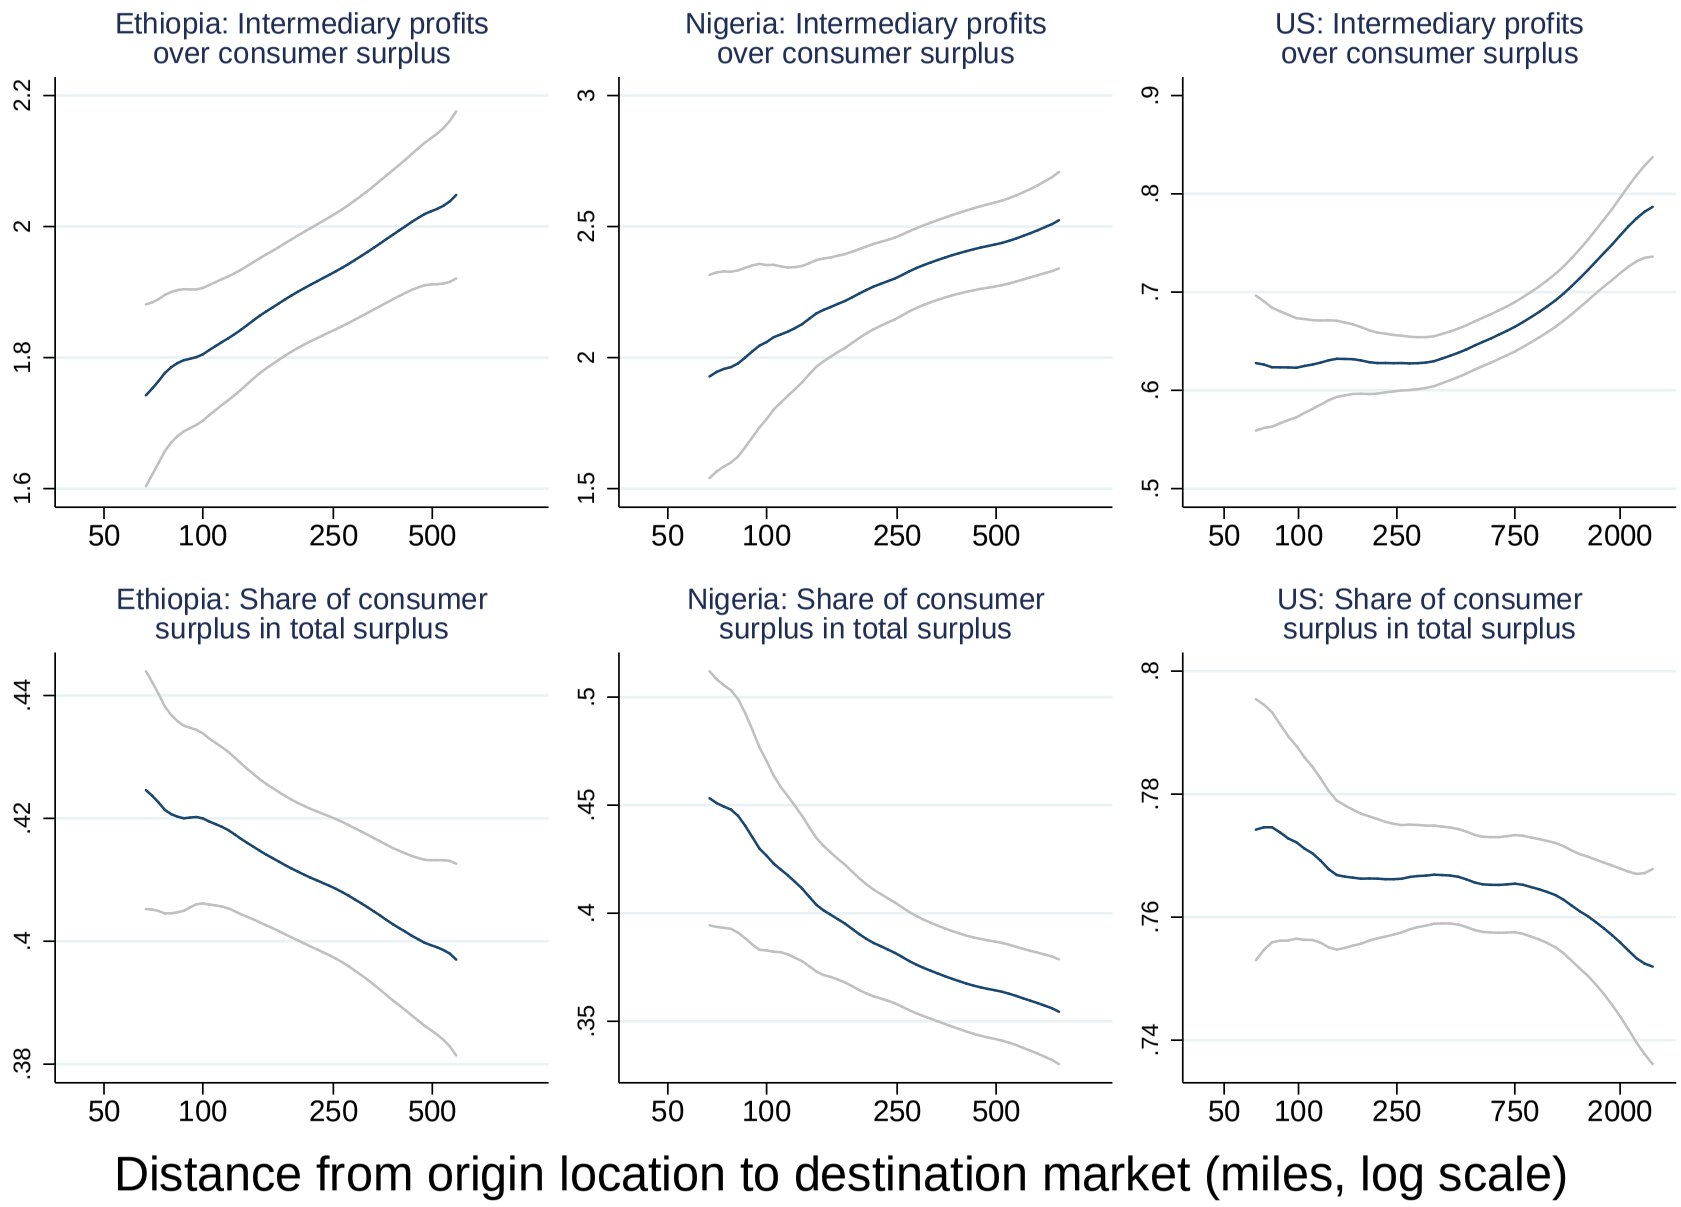

Notes: Locally weighted polynomial (Epanechnikov kernel, bandwidth=0.5). 95% confidence intervals shown. US plot uses compressed x-axis scale.

Table 1: Descriptive statistics

	(1) Ethiopia	(2) Nigeria	(3) US
Origin Price	0.400	0.968	0.618
	(0.477)	(1.267)	(0.819)
Absolute Price Gap (All Pairs)	0.055	0.131	0.094
	(0.124)	(0.260)	(0.149)
Price Gap (Trading Pairs)	0.025	0.062	-0.006
	(0.120)	(0.272)	(0.210)
Distance between locations (All Pairs, Miles)	269.23	300.20	775.76
	(137.90)	(149.36)	(533.94)
Distance to origin location (Trading Pairs, Miles)	221.16	350.48	674.01
	(107.23)	(170.88)	(546.70)
Distance to origin location (Trading Pairs, Miles, Pop. weighted)	188.19	300.91	952.17
	(86.30)	(191.68)	(696.35)
Log distance to origin location (Trading Pairs, Log Miles)	5.24	5.69	6.14
	(0.66)	(0.66)	(1.00)
Median Weight (Grams)	350	250	369
Observations (Trading Pairs)	100,761	26,025	175,782
Observations (All Pairs)	4,130,923	395,762	28,098,179

Notes: Row 1 reports the mean price at the factory location for that product or the port of entry. Row 2 reports absolute price gaps using data from all location pairs. Row 3 reports price gaps only using data from "trading pairs", i.e. pairs where one of the locations is either the factory location for that product or the port of entry. Prices are deflated by the average of the proportional price change for each good at its origin location. Real prices are converted into US Dollars using the prevailing exchange rate during the base period (January 2001). Rows 4-6 report distances in miles between locations for all location pairs, for trading pairs, and for population-weighted trading pairs while row 7 reports the distance in log miles between trading pairs. Standard errors in parentheses.

Table 2: Estimating the effect of distance on intra-national trade costs

Dependent variable:	Absolute Price Gap (All Pairs)	Price Gap (Trading Pairs)	Adjusted Price Gap (Trading Pairs)
	(1)	(2)	(3)
		Ethiopia	
Log distance to origin (miles)	0.0115***	0.0248***	0.0374***
	(0.000439)	(0.00125)	(0.00223)
			{0.00370}
			[0.00594]
Observations	4,130,923	100,761	100,761
R-squared	0.356	0.252	0.932
	(4)	(5)	(6)
		Nigeria	
Log distance to origin (miles)	0.0210***	0.0254***	0.0558***
	(0.00237)	(0.00437)	(0.00759)
			{0.0115}
			[0.0221]
Observations	395,762	26,025	26,025
R-squared	0.500	0.500	0.964
	(7)	(8)	(9)
		US	
Log distance to origin (miles)	0.00684***	0.00437***	0.0106***
	(0.000480)	(0.000731)	(0.00100)
			{0.00142}
			[0.00140]
Observations	28,098,179	175,782	175,782
R-squared	0.432	0.408	0.928
Time-Product FE	Yes	Yes	Yes
Time-Product $\times \frac{1-\hat{\rho}}{\hat{\rho}}$	No	No	Yes
Destination $\times \frac{1-\hat{\rho}}{\hat{\rho}}$	No	No	Yes

Notes: Columns 1, 4 and 7 use data on the absolute price gap between all location pairs. Columns 2, 5 and 8 use data on the actual price gap between "trading pairs", e.g. destination price minus origin price for pairs where one of the locations is either the factory location for that product or the port of entry. Columns 3, 6 and 9 use the adjusted price gap $(P^k_{dt} - \widehat{\rho^k_{od}} P^k_{ot})/\widehat{\rho^k_{od}}$ and additionally include time-product and destination fixed effects multiplied by $(1 - \widehat{\rho^k_{od}})/\widehat{\rho^k_{od}}$ in order to control for omitted variable bias due to the level of market power covarying with distance. Prices are deflated by the average of the proportional price change for each good at its origin location. Real prices are converted into US Dollars using the prevailing exchange rate during the base period (January 2001). All regressions include time-product fixed effects. Time-product clustered standard errors in round parentheses. Time-product block bootstrapped standard errors in curly parentheses, product-destination block bootstrapped standard errors in square parentheses. * significant at 10 percent level, ** at 5 percent and *** at 1 percent.

Table 3: Estimating the effect of distance on intra-national trade costs: robustness checks

	(1)	(2)	(3)
	Coefficient from regression of adjusted price gap on log distance to origin (miles)		
	Ethiopia	Nigeria	US
1. Baseline Specification	0.0374***	0.0558***	0.0106***
	(0.00223)	(0.00759)	(0.00100)
2. Not Winsorizing Pass Through Rates	0.0321***	0.0388***	0.00983***
	(0.00254)	(0.00619)	(0.00103)
3. ρ Estimated Every 5 Years	0.0329***	0.0859***	0.00943***
	(0.00200)	(0.0101)	(0.00113)
4. ρ Estimated Every 2.5 Years	0.0394***	0.0581***	0.0162***
	(0.00265)	(0.00928)	(0.00207)
5. ρ Estimated Using 3 Lags	0.0322***	0.0587***	0.00731***
	(0.00193)	(0.00749)	(0.00117)
6. ρ Estimated Using Quarterly Pass Through Rates	0.0150***	0.0279***	0.00465***
	(0.00163)	(0.00819)	(0.000921)
7. Controls for Oil Price in ρ	0.0389***	0.0583***	0.0215***
	(0.00236)	(0.00791)	(0.00298)
8. Destination Fixed Effects	0.0270***	0.0490	0.0117***
	(0.00173)	(0.0376)	(0.00100)
9. Destination-Year Interactions and Pass Through Controls	0.0438***	0.0637***	0.00925***
	(0.00138)	(0.00537)	(0.00101)
10. Destination-Time Interactions and Pass Through Controls	0.0471***	0.0764***	
	(0.00153)	(0.00605)	
11. Removing Locations <100 Miles Away	0.0668***	0.0720***	0.0159***
	(0.00439)	(0.0118)	(0.00170)
12. Removing Locations <200 Miles Away	0.0723***	0.0563***	0.0194***
	(0.00671)	(0.0164)	(0.00160)
13. Not Cleaning Price Data	0.0349***	0.0724***	0.0117***
	(0.00182)	(0.0104)	(0.00108)
14. Interaction with Weight: Log distance coef.	-0.0803***	-0.296***	0.000265
	(0.00816)	(0.0256)	(0.00384)
15. Interaction with Weight: Log distance × log weight coef.	0.0229***	0.0528***	0.00191***
	(0.00190)	(0.00450)	(0.000703)

Notes: Each cell reports (for a particular row specification and column sample country) the main coefficient on log distance from a regression of price gaps on log distance using the adjusted price gap methodology described in Section 4.3 of the main text. Each of the 15 robustness specifications is described in Section 4.4 of the main text. All regressions include time-product fixed effects. Time-product clustered standard errors in parentheses except for the US estimates in rows 8 and 12 which are unclustered due to computational limits. * significant at 10 percent level, ** at 5 percent and *** at 1 percent.

Table 4: Comparisons across distance metrics and existing estimates

	(1)	(2)
	Ratio relative to US marginal cost of distance	
	Ethiopia	Nigeria
Great circle distance	3.53	5.26
Road distance	3.19	5.40
Travel time	2.46	4.01
	(3)	(4)
	Ratio relative to US trucking costs	
	East Africa (Mombasa-Nairobi)	West Africa (Bamako-Accra)
Per km for one truckload	1.88	3.28

Notes: Columns 1 and 2 report the ratio of the coefficient on log distance for Ethiopia and Nigeria respectively, each compared to the coefficient on log distance for the US, from regressions of the adjusted price gap $(P_{dt}^k - \widehat{\rho_{od}^k} P_{ot}^k)/\widehat{\rho_{od}^k}$ on log distance and time-product and destination fixed effects multiplied by $(1 - \widehat{\rho_{od}^k})/\widehat{\rho_{od}^k}$ in order to control for omitted variable bias due to the level of market power covarying with distance. Ratios are presented for estimates calculated using three different distance metrics: row 1 uses geodesic (i.e. as the crow flies) distance, row 2 uses quickest route road distance as calculated by Google Maps, and row 3 uses quickest route travel time again calculated by Google Maps. The raw coefficients on log distance using these alternate distance metrics are reported in Appendix Table C.2. Columns 3 and 4 report ratios of trucking costs along major East and West African transport corridors relative to US trucking costs from Teravaninthorn and Raballand (2009) (all calculated per km for one truckload through surveys of trucking firms).

Table 5: Regressing product availability, competitiveness and surplus on distance

	(1)	(2)	(3)	(4)	(5)
			Ethiopia		
Dependent variable:	1=Price Record, 0=No Price Record (Time-Good-Location)	Competitiveness Index of Intermediaries (All Locations)	Ratio of Intermediary Profits to Consumer Surplus (Good-Location)	Ratio of Deadweight Loss to Consumer Surplus (Good-Location)	Consumer's Share of Total Surplus (Good-Location)
Log distance between origin and destination	-0.0959***		0.229***	0.0368***	-0.0185**
	(0.00309)		(0.0598)	(0.00908)	(0.00763)
Log distance between location and Addis Ababa		-0.344***			
		(0.127)			
Observations	145,682	100	1,418	1,418	1,418
R-squared	0.14	0.059	0.009	0.009	0.004

	(6)	(7)	(8)	(9)	(10)
			Nigeria		
Dependent variable:	1=Price Record, 0=No Price Record (Time-Good-Location)	Competitiveness Index of Intermediaries (All Locations)	Ratio of Intermediary Profits to Consumer Surplus (Good-Location)	Ratio of Deadweight Loss to Consumer Surplus (Good-Location)	Consumer's Share of Total Surplus (Good-Location)
Log distance between origin and destination	-0.0490***		0.347***	0.0833***	-0.0546***
	(0.00410)		(0.113)	(0.0138)	(0.0155)
Log distance between location and Lagos		-0.444***			
		(0.114)			
Observations	35,406	36	526	526	526
R-squared	0.298	0.168	0.019	0.063	0.023

		(11)	(12)	(13)	(14)
				US	
Dependent variable:		Competitiveness Index of Intermediaries (All Locations)	Ratio of Intermediary Profits to Consumer Surplus (Good-Location)	Ratio of Deadweight Loss to Consumer Surplus (Good-Location)	Consumer's Share of Total Surplus (Good-Location)
Log distance between origin and destination			0.0623***	-0.000350	-0.00422
			(0.0108)	(0.000299)	(0.00372)
Log distance between location and Chicago		25.33			
		(54.71)			
Observations		2,865	9,614	9,564	9,564
R-squared		0.000	0.003	0.000	0.000

Notes: Columns 1 and 6 regress product availability on log distance at the monthly time-product-location level by ordinary least squares. Sample restricted to monthly time-product pairs for which the product is observed in at least one location. Both columns include time-product fixed effects and time-product clustered standard errors are reported in parentheses. Columns 2 to 5, 7 to 10 and 11 to 14 regress estimates of the competitiveness index, the ratio of intermediary profit to consumer surplus, the ratio of deadweight loss to consumer surplus and the share of consumer surplus in total surplus on log distance. Since the competitiveness index is location not location-product specific, distance to the commercial capital is used in columns 2, 7 and 11 rather than origin-destination distance. Standard errors in parentheses. * significant at 10 percent level, ** at 5 percent and *** at 1 percent.

A Data Appendix

The methodology proposed in the main text requires retail price data for narrowly defined products, at many points in space within a group of developing and developed countries, at high frequency for many periods. Here we supplement Section 2 with a more detailed description of the data sources we use to obtain these price data and how we construct our retail price dataset. (In the interest of readability and completeness, we repeat certain details that are also contained in the main text.) We apply our methodology to a sample from two sub-Saharan African countries (Ethiopia and Nigeria) and from the United States. The construction of our African and US samples differ, so we describe them separately below.

A.1 Ethiopia and Nigeria

Because particularly little is known about the magnitude and implications of intra-national trade costs in sub-Saharan Africa (SSA), we began the data collection exercise underpinning this paper with a search for SSA countries that collect high quality retail price data satisfying the following requirements. First, a substantial number of sampled products must be extremely narrowly defined (that is, defined at a level equivalent to the barcode level). Second, the price of these products prevailing at retail markets must be recorded at many locations throughout the country. And finally, these price records must be recorded at high frequency (e.g. on a monthly basis) and for at least six years (the length of our US sample, as described below). While the official data collection infrastructure in this region is generally weak (Young, 2012), it is relatively strong for the purposes of collecting the official consumer price index via surveys of the prices of products at retail establishments. Fortunately, therefore, many SSA countries (and developing and developed countries more widely) collect exactly the type of data that is required for our analysis. We work with two SSA countries in particular—Ethiopia and Nigeria—because of their large geographic sizes and because they were particularly forthcoming in making available to researchers microdata on unit-level CPI observations. We obtained data in digital form spanning the period from September 2001 to June 2010 in the case of Ethiopia and January 2001 to July 2010 in the case of Nigeria.

Selection of Locations: Both Ethiopia and Nigeria report a CPI that is based on price observations in urban areas. In Ethiopia we obtained a sample of 103 urban market places[1] and in Nigeria we obtained a sample of 36 state capitals (one for each state).[2] These locations are shown on the maps of Ethiopia and Nigeria depicted in panels A and B of Figure 1 of the main text, respectively. Reassuringly, the sample locations provide fairly even geographic coverage of these two countries.

Survey Procedure: Enumerators visit pre-specified sample outlet locations within a market town or city, typically many times per month.[3] They are instructed to find the precise product described

[1] Urban market places are defined as localities with 2,000 or more inhabitants with exceptions made for administrative capitals. The urban market places in the survey were chosen to be nationally representative and were either a capital of an administrative region, an industrial market place and/or a market place for trade and commercial activities.

[2] For larger states in Nigeria, our price observations are actually an average of prices in the state capital and other major cities within the state as discussed further in footnote 5, below.

[3] In Ethiopia, prices are surveyed every day for the first fifteen days of the month; whereas in Nigeria, most of the

in each product's description and to report a missing value if that product was unavailable (that is, substitutes are not permitted). Once found, the enumerator's job is to learn at what price the product would typically be sold on the given date and at the given location.[4] The sample frame of outlets is intended to comprise a representative sample of open markets, kiosks, groceries, butcheries, pharmacies, super markets, etc.. In many cases, multiple outlet observations are obtained.[5]

Selection of Products: Both Ethiopia and Nigeria base their CPI on a set of products that did not substantially change during our sample period of 2001-2010. These products are designed to span the typical consumption basket (among both goods and services) in each country. Of the many products that are covered, the vast majority refer to activities—such as a "man's haircut" or a "one km taxi journey"—or goods—such as "rice" or "bread"—whose very nature means that the products cannot be precisely codified. To avoid concerns of spatially-varying unobserved quality differences we work with the sub-sample of products that we consider to be "narrowly defined." In practice, this involved a restriction to products whose producer's brandname was provided in the product description. Most products for which a brandname was provided contained additional details, such as a description including the product size, which were designed to be sufficiently precise to allow enumerators to locate the exact product.[6] Because these descriptions appear to be as precise as those linked to unique barcodes in the US Nielsen Consumer Panel data described below, we refer to these products as products that are defined at the "barcode level." The resulting sample contains 15 products in Ethiopia and 19 products in Nigeria that were broadly available across both the locations and years in our sample (examples of which include "Titus Sardines" (125gm), "Bedele Beer" (300cc) and "Lux Toilet Soap" (90gm)).

Additional Processing: As with all micro-level price data, our data contain observations that appear to be mis-recorded.[7] Accordingly, our main analysis uses a cleaned sample of price data, obtained by applying a simple cleaning algorithm to the raw data.[8] However, we also report (in

products in our sample are surveyed once per week throughout the entire month. In both cases, the data available to us are an average over these multiple temporal observations per month.

[4]The precise means by which these instructions are conveyed to enumerators varies across the two countries. In Ethiopia, enumerators do not actually purchase the items in question but instead interview traders and consumers about prices typically paid (after bargaining). In Nigeria, enumerators are instructed to make actual (bargained) purchases where thought necessary. The Nigerian enumerators' manual advises that the enumerator "must not be too sophisticated in dressing as this may make the respondent arbitrarily raise the price of his goods".

[5]In Ethiopia, price observations are collected at 1-3 outlets per location, with the exception of three larger cities (Addis Ababa, Dire Dawa and Harar) in which case the city is first divided into a number of separate major market areas and then 1-3 outlets are surveyed per market area. In Nigeria, two outlets are selected in each state capital, and also in each of up to two other urban locations within the state; the precise locations of these non-capital locations change from from month to month and are not known to us so we therefore take the location to be simply the state capital's location. The data available to us comprise an average over these multiple spatial observations per date.

[6]Exceptions were the "Singer sewing machine" and the "Philips electric iron" in Nigeria; and the "National tape recorder (2 speaker)", the "Croft Men's Leather Shoes", the "Philips 3 band Radio" and the "Philips 21 inch color TV" in Ethiopia. We drop these products because their descriptions do not appear to be precise enough to determine the exact product being surveyed.

[7]In the majority of these cases it appears that the enumerator in a particular location obtained the price for a different size or specification of product (contrary to instructions). While any such substitutions are meant to be recorded and flagged by enumerators to avoid such errors, we conjecture that such information was lost in the reporting and digitization process.

[8]The algorithm is as follows. First, we remove price quotes that lie more than 10 standard deviations away from the log mean price of a product. Second, we eyeball the time series of prices for each product in each location. If a price

Table 3 of the main text) results obtained from the original, uncleaned data set; these are similar in terms of both the magnitudes and statistical significance of the estimates. In addition, in order to arrive at estimates of trade costs that are in real terms, we require all prices to be similarly expressed in terms that correspond to a common year (which we take to be 2001) and a common currency unit (US dollars). We therefore apply a simple procedure to correct for inflation, based on prices at origin locations, and then further convert all prices used in the analysis into 2001 US dollars using the prevailing exchange rate at that time.[9]

A.2 United States

In order to provide a basis of comparison for our Ethiopian and Nigerian estimates, we seek similar estimates for the United States. Following Broda and Weinstein (2008), we use data from the Nielsen Consumer Panel (NCP, formerly known as the AC Nielsen Homescan Consumer Panel) due to its extensive geographic and product coverage.[10] We have access to data spanning the period from January 2004 to December 2009. The NCP incentivizes sample households (of which there were approximately 40,000 between 2004 and 2007 and 60,000 in the last two years of our sample) to use hand-held barcode scanners, at home, to scan all products purchased by the household on the date of purchase and, where necessary (that is, when a purchase is made at a smaller retail establishment that does not participate in Nielsen's "ScanTrak" system), enter the price that they paid for each product purchased. From the resulting price observations over space, products and time (of which there are over 300 million) we use each household's county of residence to aggregate up to a dataset that contains the average price paid, in each of the 2,856 counties and each month in 2004-2009, for the 1.4 million unique barcodes purchased by NCP households.

Unlike the samples for Ethiopia and Nigeria which are collected in order to calculate the CPI, the NCP includes many products that are neither widely available across the US nor important components of the average household budget. In order to obtain a sample similar in nature to the SSA samples, we work with relatively small sub-sample of barcodes that are the leading product in each of Nielsen's 230 "product modules" (examples of which include frozen pot pies, chili sauce, or women's shaving cream).[11]

---

quote seems unusually high or low, we verify whether nearby prices for that same product in that same period were also unusually high. If they were not, we remove this outlier.

[9] This procedure is as follows. For every origin location-product-month price observation used in our analysis, we calculate the proportional change in price over the previous month. We then calculate the (unweighted) mean of all these proportional price changes across all products in our sample for each month and use this as a measure of monthly inflation relative to the previous month with which to convert all prices into real prices in the base period. In the rare case when a product price is not observed at the source location in a given month we use the pro-rated multi-month change between available monthly observations. The normalized prices are converted into 2001 US dollars using the prevailing exchange rate during the first month of the sample.

[10] We obtained the NCP data from the Kilts-Nielsen Data Center at The University of Chicago Booth School of Business (http://research.chicagobooth.edu/nielsen).

[11] We define the leading product as the barcode that is represented in the greatest number of counties per year, averaged across 2004-2009. In order to maximize spatial and temporal coverage, we omit certain barcodes (those observed in fewer than 100 counties per year in each sample year, or in fewer than four years) prior to calculating the leading barcode. . We also omit barcodes of supermarket own-brand items since the NCP anonymizes the supermarket name for confidentiality reasons hence these barcodes do not uniquely identify a particular product. After these omissions, 230 product categories contained at least one barcode.

B Theoretical Appendix

In this appendix we formally prove the claim in equation (16) that the share of surplus enjoyed, in equilibrium, by intermediaries relative to consumers is a simple function of just the equilibrium pass-through rate ρ_d and the competitiveness index ϕ_d prevailing in that market. We revert to the notation of Section 3 here, considering just one product in one time period in one destination d.

Consider first the calculation of the amount of consumer surplus generated by the presence of the product in destination market d in any partial equilibrium market setting (that is, where the prices in all other markets are held constant). Consumer surplus when Q_d is supplied to the market is defined as $CS_d(Q_d) \equiv \int_{\psi=0}^{Q_d}[P_d(\psi) - P_d(Q_d)]d\psi$, where $P_d(\psi)$ is the consumers' inverse demand curve evaluated at quantity ψ and Q_d is the total amount consumed in equilibrium in the market. Now consider a marginal change in the exogenous origin price, P_o. The marginal change in consumer surplus is then, by an envelope argument, $\frac{dCS_d(Q_d)}{dP_o} = -Q_d(P_d(P_o))\rho_d(P_o)$, where, in a slight abuse of notation, $Q_d(P_d(P_o))$ indicates that the quantity consumed depends on the destination price P_d, but the destination price is itself a function of the level of the origin price: $P_d(P_o)$. Similarly, $\rho_d(P_o)$ is the equilibrium pass-through rate when the origin price is P_o. Integrating such small changes from the original value of the origin price, P_o, to the highest value the origin price can take, infinity, the total consumer surplus available in this market (whose origin price is P_o) is:

$$CS_d(P_o) = \int_{v=P_o}^{\infty} Q_d(P_d(v))\rho_d(v)dv. \tag{19}$$

This is an alternative expression for the level of consumer surplus in the market that highlights its essential connection to the pass-through rate.

We now seek a similar expression for the amount of intermediaries' surplus in this market, $IS_d(P_o)$. The equilibrium level of profits for one intermediary is $\Pi_d(Q_d) = [P_d(Q_d) - P_o - \tau(X_d)]$. Differentiating this with respect to P_o yields

$$\frac{d\Pi_d(Q_d)}{dP_o} = [\rho_d(P_o) - 1]q_d + [P_d(Q_d) - P_o - \tau(X_d)]\frac{dQ_d}{dP_o}\frac{1}{m_d} \tag{20}$$

in a symmetric equilibrium. Applying the chain rule to $\frac{dQ_d}{dP_o}$ and noting that the intermediary's first-order condition requires that $P_d(Q_d) - P_o - \tau(X_d) = -\frac{dP_d}{dQ_d}\theta_d q_d$ we have

$$\frac{d\Pi_d(Q_d)}{dP_o} = q_d\left[\rho_d(P_o)\left(\frac{\phi_d - 1}{\phi_d}\right) - 1\right], \tag{21}$$

where $\phi_d \equiv \frac{m_d}{\theta_d}$ is the 'competitiveness index' introduced in Section 3.1 of the main text. Integrating these small changes to P_o from its starting value, P_o, to the highest value that it can take, and again summing over the m_d symmetric intermediaries serving market d, the total intermediaries' surplus available in this market is:

$$IS_d(P_o) = \int_{v=P_o}^{\infty} Q_d(P_d(v))dv + \left(\frac{1-\phi_d}{\phi_d}\right)\int_{v=P_o}^{\infty} Q_d(P_d(v))\rho_d(v)dv. \tag{22}$$

Applying equations (19) and (22) it is then straightforward to show that the ratio of interme-

diaries' surplus $IS_d(P_o)$ to consumer surplus $CS_d(P_o)$ in the market at destination location d when the origin price is P_o is given by

$$\frac{IS_d}{CS_d}(P_o) = \frac{1}{\overline{(\rho_Q)}_d} + \frac{1-\phi_d}{\phi_d}, \tag{23}$$

where $\overline{(\rho_Q)}_d$ is a quantity weighted average of the pass-through rate, defined as

$$\overline{(\rho_Q)}_d \equiv \frac{\int_{v=P_o}^{\infty} Q_d(P_d(v))\rho_d(v)dv}{\int_{v=P_o}^{\infty} Q_d(P_d(v))dv}. \tag{24}$$

The result in equation (23), which is derived for a completely general demand structure, highlights the close connection between pass-through and the division of surplus in a general oligopolistic setting. However, pass-through enters this formula always as a weighted average $\overline{(\rho_Q)}_d$ of pass-through values at different quantities. Unfortunately, in our setting the weights in this weighted average formula (consumption quantities $Q_d(P_d(v))$ are not observed, nor is there any hope of credibly estimating the demand structure so as to estimate these weights because consumption quantities are not observed. However, in the case of the constant pass-through class of demand (i.e. that described in Assumption 3 of the main text), pass-through is constant across quantities (that is, $\rho_d(v) = \rho_d$ for all v) and hence the weights in equation (24) need not be observed. That is, under Assumption 3 we have

$$\frac{IS_d}{CS_d}(P_o) = \frac{1}{\rho_d} + \frac{1-\phi_d}{\phi_d}, \tag{25}$$

the expression in equation (16) of the main text.

C Additional Figures and Tables

Figure C.1: The distribution of population across space

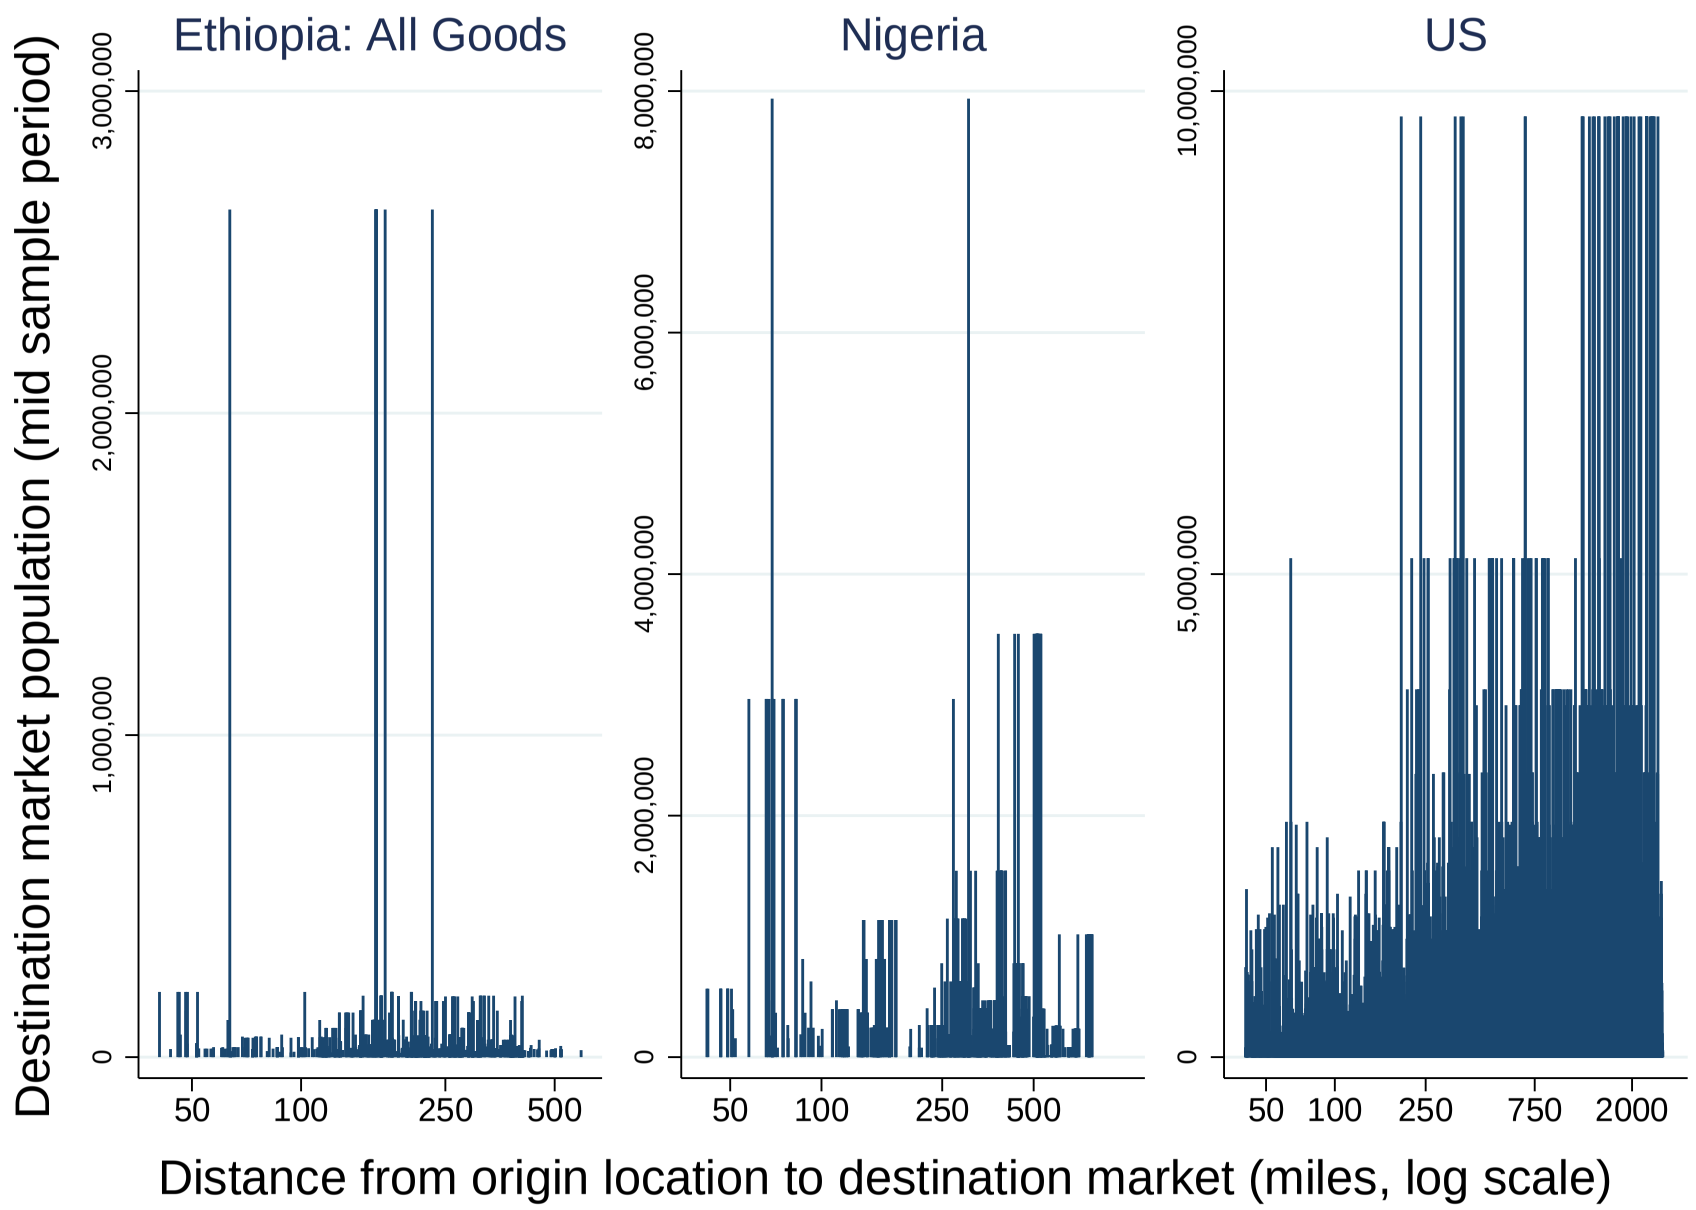

Notes: Distance from origin location plotted against the population in the destination location for every product-location pair. US plot uses compressed x-axis scale.

Figure C.2: Estimated pass-through rates for all goods and distance

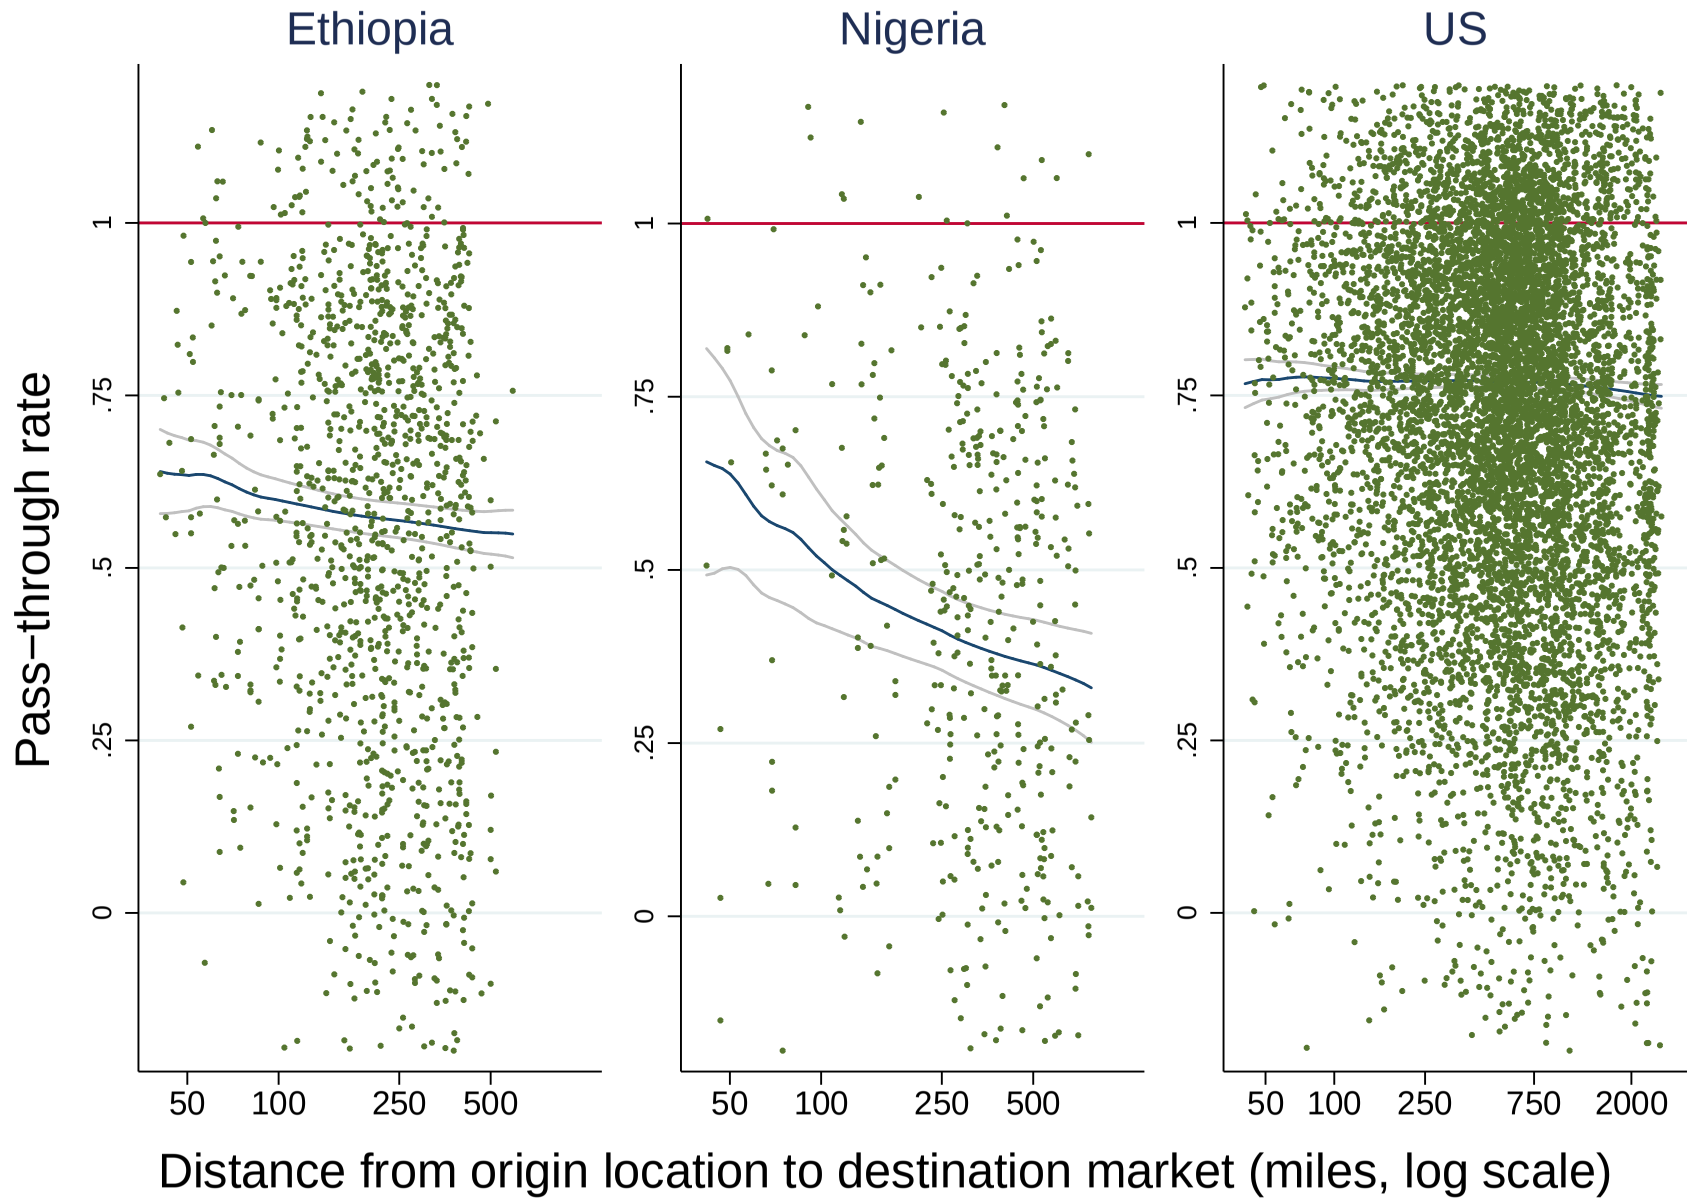

Notes: Each dot is an estimated pass-through rate for a product-location pair $\widehat{\rho_{od}^{k}}$. The solid line plots a locally weighted polynomial (Epanechnikov kernel, bandwidth=0.5) of these individual pass-through rates. 95% confidence intervals shown. US plot uses compressed x-axis scale.

Table C.1: Product list

(1) Ethiopia	(2) Nigeria	(3) US
*Batteries - Eveready Drycell No - 40 Gm	Cereal - Cerelac - 400 Gm Tin	Nutritional Supplements - Airborne - 10 Tablets
*Cigarettes - Rothmans - 20 Pack	Beer - Guinnes Stout - 30 Cl	Wafers & Toast & Bread Sticks - Ak Mak Whole Wheat and Sesame Crackers - 4.75 Oz
*Detergents - Zahira - 50 Gm	Cement - Elephant - 50 Kg	Contact Lens Solution - Alcon Opti Free Express - 12 Oz
Water Bottled - Ambo Mineral Water - 500 Cc	Cement - Nigecem - 50 Kg	Eye Drops & Lotions - Alcon Systane Lubricant Eye Drops - 0.5 oz
Beer - Bedele - 300 Cc	Cigarettes - Benson & Hedges - 30 Gm	Dinners Frozen - Banquet Salisbury Steak & Gravy - 9.5 Oz
Soft Drinks - Coca Cola Fanta - 300 Cc	Detergent - Omo - 1 Packet	Pot Pies Frozen - Banquet Chicken - 7 Oz
Beer - Harar - 330 Cc	Drink - Bournvita - 450 Gm	Lighters - Bic Disposable Butane Lighters - 5 Ct
Beer - Meta Abo - 330 Cc	Soft Drinks - Coca Cola - 35 Cl Btl	Lip Remedies Remaining - Blistex Medicated Lip Ointment - 0.21 Oz
Soft Drinks - Pepsi Cola Mirinda - 300 Cc	Soft Drinks - Coca Cola - 35 Cl Can	Lip Remedies Solid - Blistex Medicated Lip Balm - 0.15 Oz
*Electric Bulb - Philips 40/60 Watt No - 20 Gm	Drink - Milo - 450 Gm	Vitamins Multiple - Centrum Silver Multivitamin for Adults - 100 Ct Btl
Hair Oil - Zenith Non-Liquid Form - 330 Cc	Drugs - Panadol - 1 Packet	Abrasive Cleansers Powdered - Comet Powdered Cleaner with Bleach - 21 Oz
Motor Oil - Mobil Lt - 1.2 Kg	*Evaporated Milk - Carnation - 170 Gm	Precut Fresh Salad - Dole Fresh Favorites Classic Iceberg Salad Mix - 16 Oz
Pen Ball Point - Bic England No - 20 Gm	Evaporated Milk - Peak - 170 Gm	Gum Bubble - Dubble Bubble B/G PC - 16 Oz
*Toilet Soap - Lux - 90 Gm	Margarine - Blue Band - 1 Tin	Ground And Whole Bean Coffee - Folgers Classic Roast - 39 Oz
Wine - Saris Normal - 750 Cc	Powdered Milk - Peak - 400 Gm	Dip Canned - Frito Lay Original Flavor Bean - 9 Oz
	*Sardines - Titus - 125 Gm	Hand Cleaners And Hand Sanitizer - Germ X Original Hand Sanitizer- 1.5 Oz
	*Sugar - St Louis - 1 Packet	Disposable Diapers - Huggies Pull Ups Goodnites Underpants - 40-75 Lbs
	Tea - Lipton - 1 Box	Salad & Potato Toppings Dry - Hormel Real Bacon Bits - 3 Oz
		Cooking Sauce - Hunt'S Manwich Original Sloppy Joe Sauce - 15.5 Oz
		Oriental Sauces - Kikkoman Soy Sauce - 10 Oz
		Bakery Breakfast Cakes/Sw - Little Debbie Pastries Honey Buns - 10.5 Oz
		Bakery Cakes Fresh - Little Debbie Swiss Cake Rolls - 13 Oz
		Bakery Doughnuts Fresh - Little Debbie - 6 Ct
		Pizza Refrigerated - Mama Rosa's Original Cheese Medium Pizza - 17 Oz
		Salad Dressing Refrigerated - Marie's Blue Cheese - 12 Oz
		Gravy Mixes Packaged - McCormick Brown Gravy Mix - 0.87 Oz
		Sauce & Seasoning Mix Remaining - McCormick Beef Stew Seasoning Mix - 1.5 Oz
		Sauce Mix Taco - McCormick Mexican Original - 1.25 Oz
		Seasoning Mix Chili - McCormick Original Chili Seasoning Mix - 1.25 Oz
		Candy Hard Rolled - Necco Assorted Wafers - 2.02 Oz
		Denture Cleansers - Polident 3 Minute Tablets - 84 Ct
		Bakery Remaining Fresh - Rich's New York Style Eclair, Chocolate Iced, Bavarian Cream Filled - 8 Oz
		Shave Creams Women's - Skintimate Raspberry Rain Moisturizing Shave Gel - 2.75 Oz
		Soup Frozen Refrigerated - Skyline Original Recipe Chili - 13.25 Oz
		Fruit Spreads - Smucker's Squeeze Strawberry Spread - 20 Oz
		Soap Liquid - Softsoap Soothing Aloevera - 7.5 Oz
		Abrasive Cleansers Liquid - Soft Scrub Anti-bacterial with Bleach - 24 Oz
		Frozen Hors D' Oeuvres & Snacks - Superpretzel Soft Pretzels - 13 Oz
		Tabasco/Pepper Sauce - Tabasco Original Flavor Pepper Sauce - 2 Oz
		Candy Non Chocolate Miniatures - Tootsie Roll Midges Candy - 12 Oz
		Candy Lollipops - Tootsie Roll Pops - 10.125 Oz
		Entress Poultry Food Frozen - Tyson Crispy Chicken Strips, Chicken Breast Strip Fritters with Rib Meat - 28 Oz
		Frozen Food Entrees Italian - Weight Watchers Smart Ones Classic Favorites Three Cheese Ziti Marinara - 9 Oz
		Frozen Food Entrees Mexican - Weight Watchers Smart Ones Classic Favorites Santa Fe Style Rice & Beans - 9 Oz
		Jelly - Welch's Concord Grape Jelly - 32 Oz
		Sandwiches Refrigerated Frozen - White Castle Microwaveable Cheeseburgers - 6 Ct

Notes: Asterisks before a product name denote that that product is imported.

Table C.2: Estimating the effect of distance on intra-national trade costs: other distance metrics

	(1)	(2)	(3)
	Ethiopia		
Dependent variable:	Adjusted Price Gap (Trading Pairs)		
Distance metric:	Great Circle Distance (miles)	Google Road Distance (miles)	Google Travel Time (mins)
Log distance to origin	0.0374***	0.0383***	0.0411***
	(0.00223)	(0.00232)	(0.00246)
Observations	100,761	100,761	100,762
R-squared	0.932	0.932	0.933

	(5)	(6)	(7)
	Nigeria		
Dependent variable:	Adjusted Price Gap (Trading Pairs)		
Distance metric:	Great Circle Distance (miles)	Google Road Distance (miles)	Google Travel Time (mins)
Log distance to origin	0.0558***	0.0648***	0.0669***
	(0.00759)	(0.00844)	(0.00877)
Observations	26,025	26,020	26,020
R-squared	0.964	0.958	0.958

	(9)	(10)	(11)
	US		
Dependent variable:	Adjusted Price Gap (Trading Pairs)		
Distance metric:	Great Circle Distance (miles)	Google Road Distance (miles)	Google Travel Time (mins)
Log distance to origin	0.0106***	0.0120***	0.0167***
	(0.00100)	(0.00120)	(0.00159)
Observations	175,782	174,599	165,616
R-squared	0.928	0.933	0.940

Notes: Regressions of the adjusted price gap $(P^k_{dt} - \widehat{\rho^k_{od}} P^k_{ot})/\widehat{\rho^k_{od}}$ on log distance additionally including time-product and destination fixed effects multiplied by $(1 - \widehat{\rho^k_{od}})/\widehat{\rho^k_{od}}$ in order to control for omitted variable bias due to the level of market power covarying with distance. Estimates calculated using three different distance metrics: columns 1, 5 and 9 use geodesic (i.e. as the crow flies) distance, columns 2, 6 and 10 use quickest route road distance as calculated by Google Maps, and columns 3, 7 and 11 use quickest route travel time again calculated by Google Maps. All regressions include time-product fixed effects. Time-product clustered standard errors in parentheses. * significant at 10 percent level, ** at 5 percent and *** at 1 percent.