CROSS-BORDER PRODUCT ADOPTION: INDIVIDUAL IMPORTS, MIGRANT NETWORKS, AND DOMESTIC RETAILERS

By

David Argente, Esteban Mendez and Diana Van Patten

April 2026

COWLES FOUNDATION DISCUSSION PAPER NO. 2522

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS

YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

http://cowles.yale.edu/

Cross-Border Product Adoption: Individual Imports, Migrant Networks, and Domestic Retailers
David Argente, Esteban Méndez, and Diana Van Patten
NBER Working Paper No. 35104
April 2026
JEL No. D22, D8, E02, F1, F14, F15, F6, L2, O1, O12, O3, O30, O47

ABSTRACT

This paper studies how new varieties enter markets and become locally available. We provide causal evidence of demand externalities that operate in two steps. First, information about new varieties diffuses directly through real-world social ties among consumers. Second, early purchases generate an indirect spillover to firms: local retailers learn from "pioneer" consumers which new varieties are most likely to succeed and adjust their product offerings accordingly. We study this process in the context of direct-to-consumer imports. Using customs records on individuals' purchases matched to population-wide social networks, international migrant links, and retailer catchment areas, we document economically meaningful demand externalities. Product-specific demand shocks abroad transmit through migrant networks and shift which varieties consumers purchase. Leveraging these shocks as a plausibly exogenous source of local demand variation, we show strong peer effects: prior purchases by close neighbors, coworkers, or friends increase an individual's likelihood of purchasing the same variety, especially for premium and visible goods. We leverage this result to identify an indirect spillover from consumers to firms: retailers are more likely to add a variety when it becomes popular among consumers in their catchment area. Combining the instrument with linked consumer--retailer data and a self-conducted retailer survey, we show that this response reflects learning about latent demand for varieties not yet stocked locally. Together, social diffusion and retailer learning generate demand multipliers that reshape local product availability and expand access to global variety.

David Argente
Yale University
Yale School of Management
Economics
and NBER
david.argente@yale.edu

Diana Van Patten
Yale University
Yale School of Management
and NBER
diana.vanpatten@yale.edu

Esteban Méndez
Central Bank of Costa Rica
mendezce@bccr.fi.cr

1 Introduction

In recent decades, developing countries around the world experienced a supermarket revolution that shifted consumption from local bodegas and kirana shops to mass-market retailers (Atkin et al., 2018; Reardon et al., 2003). Today, a new transformation is underway: e-commerce is redefining how individuals source their consumption (Bai et al., 2020). Yet even as retail modernizes, consumers in small, remote, and developing economies often face limited local variety—not because products do not exist globally, but because many varieties are effectively missing from domestic shelves. One increasingly important way households bridge this gap is through direct-to-consumer imports: consumers purchase varieties available abroad before local retailers carry them. For example, in Latin America, consumer online purchases from overseas retailers account for more than 50% of e-commerce in Bolivia, Costa Rica, Ecuador, the Dominican Republic, El Salvador, Panama, and Paraguay, with an average of 39% in the region (EBANX, 2023).

This paper leverages this setting to study how new varieties enter markets: when a small number of consumers discover and try new products first, does that information remain confined to those shoppers, or can it spread and ultimately reshape what is available locally? We provide causal evidence of economically meaningful demand externalities that operate in two steps. First, information about new varieties diffuses *directly* through real-world social ties among consumers. Second, early purchases of new varieties generate an *indirect* spillover to firms: local retailers learn from "pioneer" consumers which new varieties are most likely to succeed and adjust their assortments accordingly. Together, these forces create demand multipliers, so that small shocks to "who tries what" ripple through networks and, via retailer responses, translate into broader changes in local product availability.

The mechanism we study is not conceptually specific to importing: it reflects a general process of how product ideas diffuse across networks of people and how local firms learn from early adopters in their catchment areas. We focus on direct-to-consumer imports because they provide a uniquely clean laboratory for measuring and identifying these effects. First, direct consumer imports often involve varieties that are not (yet) supplied locally, making "missing varieties" observable. Second, foreign demand shocks can be traced through migrant networks, providing plausibly exogenous variation that helps identify the causal effect of social exposure on adoption.

Third, customs records capture the universe of these purchases at the transaction level, allowing us to observe purchases and diffusion at scale rather than through a single platform or retailer.

The externalities we document are particularly relevant in environments where both sides of the market face severe information frictions. Consumers considering unfamiliar varieties must form beliefs about quality, fit, and reliability—often with limited ability to return a poor match. Retailers face a parallel problem: they must decide which *new* products to stock without directly observing local demand, and searching over a vast product space is costly and risky (Bai et al., 2020; Juhász and Steinwender, 2018; Startz, 2016). These barriers can be particularly binding for small retailers, who disproportionately shape the assortments available to lower-income households (Faber and Fally, 2022). In such settings, consumers' early purchases are informative signals: each purchase reveals product-specific information that can spread through social ties and, in turn, shift both what others buy and what firms choose to carry.

We study direct externalities among peers in two complementary steps. First, we test whether product-specific expenditure shocks transmit across borders through international migrant networks: when a product becomes popular in a foreign location, does this trigger imports from individuals connected to it via relatives who migrated? Second, we leverage these foreign shocks—mapped through migrant links to local networks—as plausibly exogenous sources of domestic demand variation to construct an instrument. This design allows us to identify how an individual's likelihood of importing a given foreign product responds to prior imports by peers in their local network and to product characteristics.

The third and main part of our analysis focuses on an indirect spillover from consumers to firms. We ask whether local retailers can learn about *latent demand* for new varieties—those not yet available domestically—by observing which goods are imported by consumers in their catchment area. Isolating this channel is challenging because the relevant signal is precisely demand for *missing* products: retailers do not observe local sales for varieties they do not yet stock, standard data rarely link consumers' purchases of goods not available locally to the retailers in their vicinity, and even with the data in hand, isolating the externality requires exogenous variation in exposure to products. Accordingly, the existing literature has focused on learning and search frictions along the supply chain rather than learning from final consumers (Bai et al. 2020; Startz 2016; Allen 2014, Fernandes and Tang 2014). Combining our instrument with linked consumer–retailer data and direct survey evidence, we show that retailers treat pioneer consumers' early purchases as informative signals, update their sourcing decisions, and expand assortments toward those varieties—generating multipliers in variety discovery that reshape local product availability.

Taken together, the results highlight a two-step mechanism for how new varieties enter and spread within markets: social diffusion influences which varieties are discovered, and spillovers through firm learning determine whether those discoveries scale. When both steps operate, modest shocks to initial adoption can translate into market-wide changes, changing what ends up on local shelves and who gains access to global varieties.

Our context is the adoption of foreign goods by Costa Rican individuals and firms. Costa Rica is a small open economy where many goods available online are not sold domestically, and international returns are rare.[1] Measuring how consumption of varieties diffuses and whether retailers learn from consumers is non-trivial. First, there are data constraints, as the analysis requires information on consumption and networks. On the consumption front, we leverage novel administrative data on imports by *individual consumers*; for instance, the data would record a purchase at a U.S. online retailer that was then delivered by mail to Costa Rica. Each record includes details like date, price, weight, product code up to the HS-10 level, and country of origin. Using this previously unstudied data, we document new facts on individual imports, including main products imported, origins, and gravity parameters. On the networks front, previous research has mainly identified networks using shared characteristics, such as race or cohort, or has focused on a single network definition.[2] We directly observe networks of relatives, neighbors, coworkers, and friends; and can identify Costa Ricans living abroad and *link them back* to their networks still in Costa Rica. These local networks are computed for the *entire adult population* using information on family trees, employer-employee records, a new measure of friendships developed

[1]In developing markets, the costs of returning an item are often high; anecdotally, consumers usually absorb the cost of internationally shipping back the item, the cost of processing the return, and suffer the delays of international shipping. In fact, in the data, as little as 0.01% of individual imports are returned.

[2]E.g., Agarwal et al. (2018); Bandiera et al. (2009); Charles et al. (2009); Conley and Udry (2010); Kuhn et al. (2016); Maurer and Meier (2008).

for this paper (leveraging the universe of local peer-to-peer money transfers), and both local and foreign residences. Crucially, individual-level data on networks and imports can be linked. Finally, to speak to retailers' responses, we use administrative data for all formal retailers, including product-specific imports and catchment areas based on their customers' location. We also conducted a large-scale survey spanning 4% of retailers to complement our analysis and delve deeper into the mechanisms at play.

Second, identifying demand externalities is often challenging (Brock and Durlauf, 2001; Manski, 1993; Moffitt, 2000), as it is hard to distinguish a true network effect from correlated shocks and common characteristics. To overcome these challenges, we propose a new instrument which leverages several aspects of our context and our data, and that is based on the following idea:

> Suppose individuals L and N live in Costa Rica, and L has a sister living in Los Angeles (LA) while N has a sister living in New York City (NYC). *If product i becomes more popular in LA as compared with NYC in period t, then L is more likely than N to import product i in period $t+1$.*

The spirit behind this instrument is how, anecdotally but also intuitively, information on products is transmitted to developing countries after relatives migrate to developed countries, where more products are available. Theoretically, we derive a microfoundation of this instrument based on the idea that individuals receive noisy signals about shocks through relatives abroad. Empirically, the instrument exploits that (i) we can identify Costa Ricans living abroad and where they reside in the U.S. (1% of the Costa Rican population), (ii) we can link these emigrants to their family network still living in Costa Rica (5% of the population), (iii) we collected regional data from several sources to follow product-specific dynamics across the U.S., (iv) consumer trends in the U.S. do not respond to local conditions in Costa Rica, and (v) we can track Costa Rican product-specific foreign purchases at a daily frequency. This strategy also has the large advantage that the analysis can be run using product-level individual consumption, as opposed to total consumption, which aids in separating the consumption network effect from income shocks, and will allow us to explore heterogeneous effects. The instrument strongly predicts product-specific imports by individuals with close relatives in the U.S.[3] Therefore, the first insight

[3] This holds while focusing on *residualized* expenditures, which are uncorrelated over

from the instrument, and the paper's first contribution, is to show how expenditure shocks propagate across international migrant networks. Among those with relatives abroad, a one standard deviation increase in exposure to a product leads to an 11-15% higher probability of importing it within a quarter.[4]

We then use the individual-product-time variation from the instrument to examine if, after a Costa Rican with a relative abroad increases her exposure to a product, others in her network (neighbors, coworkers, or friends), but who do *not* have relatives living in the U.S., become more likely to import *the same* product. The 2SLS finds that a 10 percentage point (pp) increase in the share of network members with a relative abroad who import product p leads to a 3.5 to 5 pp higher probability of importing this product for individuals without relatives abroad.[5] To grasp the aggregate effect of the direct externalities, suppose there is a \$100 increase in U.S. per capita spending on a product. Then, total Costa Rican consumer imports—per individual with a direct connection to the U.S.—would increase by 40 cents.[6] This diffusion channel also has implications for inequality, as families with relatives abroad have, on average, 12% higher incomes than those without such connections. Thus, the demand externality tends to enable lower-income Costa Rican families to benefit indirectly from migration.

We find significant heterogeneity across products in how demand shocks diffuse through local networks. Given exposure, some products exhibit a strong diffusion, i.e., many people import them, while for others diffusion is weak. Consistent with information being relevant, products diffuse more strongly if they are premium varieties, visible goods, or belong to more dynamic categories. Suggestively, goods also diffuse more if their initial importer is well-connected and has high centrality.

Finally, our third and main contribution documents a new channel by which retailers learn about the local demand for foreign products. The limited evidence on the determinants of retailers' sourcing choices—and on whether they learn from their

---

time and capture the entry and exit of new product brands and varieties in the U.S.

[4]A one standard deviation is equivalent to increasing per capita spending by \$1.5 in the U.S. region where the migrant is located.

[5]As a robustness exercise, we push our data further and construct a distance-3 nodes instrument which extends the methodology of De Giorgi et al. (2019); results hold and remain statistically equal to those of our baseline approach.

[6]The multiplier is in per capita terms based on those with a direct U.S. connection, as they are the bridge through which demand propagates; if the countries were more connected via migrants, the effect would be larger, which is why a per capita estimate is useful.

consumers—largely stems from measurement challenges. Even with ideal data, identification remains non-trivial due to reverse causality and confounding factors such as common demand shocks. Moreover, even if retailer learning from consumers were causally identified, the mechanisms underlying retailers' responses would remain an open question. To make progress, we again leverage our instrument to test whether retail firms with exogenous exposure to a foreign product via their customers are more likely to import that product. In doing so, an added challenge is defining a retailer's catchment area, as retailers may serve multiple neighborhoods. We address this by estimating retailer-specific catchment areas, which we call *retailer gravity zones,* using sales receipt data with customer details (available for roughly two-thirds of retailers) and by proposing a method to approximate the catchment areas for all remaining firms, a strategy that can be replicated in settings lacking customer location data.

With the instrument and gravity zones, we show that retailers respond to the exogenous exposure of their customers to a product. A one standard deviation increase in the share of individuals with relatives abroad who import a product raises the likelihood that nearby retailers import the same product by 9%. This supply response is large and mainly driven by small retailers, who face higher search costs—making consumer insights especially valuable—and may be more receptive to local consumers' needs. Moreover, households who import directly have higher income than those who do not; as retailers introduce varieties of foreign products, the indirect externality expands products access for lower-income households (Faber and Fally, 2022). Thus, while the initial benefits of migration and direct imports are concentrated among the wealthier, diffusion via retail ultimately contributes to the diffusion of these varieties across income groups.

We also explore heterogeneous supply responses based on the product's popularity among individuals. The key is that the peer effects-based analysis conducted earlier allows us to understand the strength of diffusion by product to examine if, conditional on being exposed to a product, retailers become more likely to begin importing it if it featured strong diffusion and *less likely to import it* if the product was not popular among locals. We find that when individuals display *strong* local demand for a foreign good—evidenced by robust diffusion—local retailers respond by importing and selling that product domestically. Conversely, if a product experiences *weak* diffusion after its initial import, indicating low local demand, domestic retailers become *less* *likely* to import it.[7] Thus, retailers are responsive to the revealed preferences among individuals, in line with learning about local demand for particular foreign goods. Reassuringly, and aligned with the notion that retailers now serve the local market, individual imports of a product decline once retailers begin selling it domestically.

We then design a survey to validate the channels uncovered in the administrative data. Our large-scale survey spans 700 retail firms—approximately 4% of retailers in the country. First, over 80% of respondents receive feedback from their customers on what products to stock. Second, in line with our finding that retailers respond to individual imports, 60% of retailers noted that observing customers importing new products would make them more likely to start importing and selling them domestically. Third, the survey confirms that small retailers rely more heavily on direct customer imports to gauge the local demand for a potential new product. Fourth, how do retailers gauge demand? Around half of those who are responsive to individual imports gain insights into which imported products interest their customers when they physically visit the store and ask about the availability of these goods. Taken together, our findings point to an indirect externality such that retailers—particularly small ones—learn about the local demand for foreign goods by observing the degree of interest for different imported goods among their customers.[8]

Finally, we combine the estimated effects to grasp their overall impact: individual imports, both due to migrant connections and domestic peer effects, and imports of local retailers. Suppose there is a $100 increase in U.S. per capita spending on a product. In this case, we find total Costa Rican imports of this product—per individual with a relative abroad—would increase in $6, given the degree of interconnectedness

[7]Other test in support of firms learning about local demand derives from the notion that employees can be exposed to products where they live and pass information to employers. However, if they live outside the retailer's gravity zone, where preferences might differ, their insights should be less informative about latent demand. Indeed, we find that retailers are unresponsive to the exposure of employees who live outside their catchment area, underscoring the importance of local demand knowledge. This finding aligns with survey results, where retailers are five times more likely to gather insights from employees living close-by.

[8]An example helps tie the exercises together. In 2017, fidget spinners became particularly popular in Houston. In Costa Rica, we observe a corresponding increase in first-time imports in the relevant product category among individuals with relatives in Houston and their peers. Shortly thereafter, retailers in areas highly exposed via migrant networks are more likely to import the product for the first time. One such retailer is in our survey; its manager states that customer feedback shapes stocking decisions and that the likelihood of importing and selling a new product rises when customers begin purchasing it online from abroad.

across countries and the strength of the demand propagation.[9] This sizable effect can be decomposed into additional imports due to the direct externality and individuals' responses—7% of the effect—and the indirect externality and imports of domestic retailers—the remaining 93%. Such magnitudes underscore why accounting for the newly documented supply-side indirect effect when estimating the full response is key.

Related Literature The paper contributes to the literature on trade and information frictions by providing the first direct evidence on how retailers learn from final consumers to inform their sourcing choices and overcome search barriers. This literature has studied price search barriers (Allen, 2014; Steinwender, 2018), search frictions along the supply chain (Chaney 2014; Bai et al. 2020), and product search (Startz 2016; Juhász and Steinwender 2018). Our work relates to Startz (2016), who provides insights into how Nigerian sellers overcome search frictions by traveling to find products, Juhász and Steinwender (2018), who show information technology improvements are valuable for conveying product characteristics, Felbermayr et al. (2015), who find reductions in the informational component of trade costs enhance efficiency, and Wei et al. (2021), who document information-discovery externalities when firms first enter an export market. Our work is also related to Bai et al. (2020), who document the importance of information barriers in cross-border e-commerce, and connects to findings on market integration through online shopping, including Gorodnichenko and Talavera (2017) and Duch-Brown et al. (2021). We complement this work by documenting a new channel through which retailers can overcome search barriers and presenting evidence on the mechanism with self-collected survey data. The analysis also estimates retailer catchment areas using novel data on customer residences. While Batch et al. (2024) use credit card data to partition the U.S. into consumer zones, our gravity zones are retailer-specific. Moreover, as customer location data are observed alongside employer–employee links, we propose an approximation method that yields estimates with a correlation above 0.98 relative to those based on customer locations and that can be applied in other settings.

Beyond international trade, our paper relates to work examining retail behavior (Bronnenberg and Ellickson, 2015; Hortaçsu and Syverson, 2015) and to the literature

[9] Note that this multiplier is again calculated in per capita terms, where the denominator in Costa Rica are individuals with a relative abroad, as they act as the bridge through which demand propagates. A more connected set of countries would face an overall stronger propagation, which is why a per capita estimate is informative.

in marketing and operations on retail demand forecasting (e.g., Fildes et al., 2022). These literatures typically study how retailers predict demand for products already in the assortment using historical sales, prices, promotions, and other store-level information. We complement this work by studying how retailers learn about latent demand for new varieties that are not available locally. In our setting, consumers' direct imports provide a signal about which foreign varieties are likely to succeed domestically, and we provide direct evidence—from a retailer survey and from linked consumer-retailer data—that retailers use these signals to inform sourcing choices and overcome search barriers.

The paper also contributes to the literature on international trade and consumption, which has examined how the availability of retail outlets and online platforms impact consumer welfare (Atkin et al., 2018; Couture et al., 2021; Dolfen et al., 2023) and how small firms shape the assortment available to low-income consumers (Faber and Fally, 2022). We focus on direct consumer imports, a new channel through which trade externalities can influence the varieties available locally and thereby shape consumer gains (Broda and Weinstein, 2006). To the best of our knowledge, this is the first paper to leverage customs data on individual imports. With the direct-to-consumer market rapidly expanding, and expected to accelerate further with increased global internet penetration and improved transport and logistics, this topic is fertile ground for future research, with contemporary work already studying the increasing value of tariff exemptions on individual imports in the U.S. (Fajgelbaum and Khandelwal, 2024).

Moreover, we study direct externalities as a building block in our retailer analysis. Our peer effects study uses a battery of network definitions for the full population—including new measures of close neighbors (via voting records) and friends (via money transfers)—allowing us to assess economy-wide impacts across multiple product categories. Peres et al. (2010) emphasize that empirical research on individual-level diffusion remains limited because it is rarely possible to observe real social networks, track individuals' exposure to products, and measure adoption decisions at scale. In our setting, we can do all three simultaneously. Moreover, the product variation is key in our identification strategy and allows for heterogeneous effects across characteristics, while previous work has yet to fully examine heterogeneous and economy-wide impacts. De Giorgi et al. (2010) make progress with measures based on household income, but observe only total consumption. Bailey et al. (2022) also make progress by using Facebook data to define networks, but rely on a single product. More broadly, there is a longstanding literature studying how social interactions can impact behavior, both in developed and developing countries.[10] Unlike previous work, we propose an interaction between direct and indirect externalities, which leads to retailers learning from consumers' experimentation.

Finally, our findings relate to the migration–trade literature documenting that immigrant networks are associated with greater international exchange (e.g., Peri and Requena-Silvente, 2010), and to studies on the transmission of ideas via migrant networks and other potential benefits for countries of origin (Acosta et al., 2008; Agarwal et al., 2018; Beine et al., 2008). Also related, McCully et al. (2024) leverages Homescan panel data and a structural approach to study how consumption is affected in destination countries. We propose a method to link migrants to their foreign city and home networks, and use these links to translate foreign product-specific demand shocks into plausibly exogenous shifts in domestic exposure—a key ingredient in our identification strategy. Beyond this methodology, we also show that migrants facilitate the diffusion of new products to their origin locations.

The rest of the paper is organized as follows. Section 2 describes the data used in our analysis. Section 3 presents stylized facts on individual imports. We describe our estimation framework and results on direct externalities in Section 4. Section 5 describes our results for retail firms, and Section 6 concludes.

2 Data

We now describe the battery of administrative datasets used in our analysis. Notably, while the data is anonymized, variables *across* datasets can be linked via unique (pseudonymous) identifiers at the individual level.

Customs Data We leverage customs records from 2014 to 2019. Each import includes up to a 10-digit HS code, along with information on the amount transacted, the quantity traded, the arrival date, and the country of origin.[11] As in other coun-

[10]E.g., Bandiera et al. (2009); Bandiera and Rasul (2006); Beaman et al. (2021); Bertrand et al. (2000); Carter et al. (2021); De Giorgi et al. (2010); Duesenberry (1948); Duflo and Saez (2003); Mas and Moretti (2009); Veblen (1899).

[11]For some categories, an HS-10 classification does not exist, so the HS-8 or HS-6 code is the narrowest classification. We use the most disaggregated category available. While

tries, customs records are available for firm-level imports. In addition, if an *individual* imports a good (for instance, if she bought an item from an online retailer in the U.S. which was then shipped to Costa Rica), this transaction is also recorded.[12] The median value of these individual imports is $30. To the best of our knowledge, this is the first paper to leverage this type of customs records despite the fact that, far from an unusual practice, cross-border shopping has become prevalent.

Networks Data Within Costa Rica This paper undertakes an effort to combine several reference groups and paint a relatively complete picture of the network of each Costa Rican individual. In particular, we define networks in three different ways.

Networks of neighbors: First, we assume that an individual's network consists of those who live in close proximity. Networks of neighbors are constructed from official records maintained by the National Registry. While records include district of residence, with 488 districts in total, they also detail the voting center which is closest to each citizen's residence *for each adult citizen.* With 2,028 voting centers in total, the median number of adults assigned to each voting center is 586.[13] We propose the latter voting center information as a novel way to get a precise measure of close-by neighbors.

Networks of coworkers: Second, we assume that the relevant network is composed of contemporaneous coworkers.[14] Matched employer-employee data was obtained from the Registry of Economic Variables of the Central Bank, which tracks the universe of formal employment between 2015 and 2019. These data allows us to recover networks of coworkers which change at a monthly frequency, as people change their employers.

Networks of "friends": Third, we create a novel measure of social networks which connects pairs of individuals who have sent money to each other *bilaterally.*[15] We use

HS codes are not barcodes, this can be seen as an advantage in our setting; a person might learn about a new type of flask bottle from a peer, but order a blue one instead of a green one, which would typically be in the same HS code but would not have the same barcode.

[12]While individual imports could potentially also include imports from informal sellers, it will become clear that this would only make our estimates a lower bound: if a person imports a product and informally sells it domestically, then the incentives for others to import it decrease, which would attenuate our coefficient of interest. Moreover, nearly all individuals import each product only once within our sample period (2015-2019), which suggests these are not informal sellers who use this method to stock, and dropping those who import more regularly does not meaningfully change our results.

[13]For details on the distribution of voting centers, see Méndez and Van Patten (2022).

[14]This is in line with De Giorgi et al. (2019), who identify coworkers as a good reference group given the large share of the day spent with them, among other reasons.

[15]For instance, if user A has only sent money to user B, we would not record this rela-

data on comprehensive transactions on Sinpe Móvil, an application that since 2015 allows Costa Ricans to make peer-to-peer money transfers via their mobile phones (Alvarez et al., 2023). It has been adopted by nearly 70% of all adults and processes the equivalent of 22% of GDP in transactions each year. We construct a time-invariant measure, as follows: we start at the *end* of the sample period and retrospectively ask: which pairs of peers have made transfers to each other bilaterally? These pairs are considered friends, which has the advantage of eliminating transfers to, for instance, a nanny or a housekeeper. This method allows us to proxy for networks of first-degree friends for each individual which are usually infeasible to recover; more details are available in Appendix D. Moreover, we exclude individuals whose number of friends lies above the 99th percentile of the sample distribution to avoid the large connected components commonly observed in social network data and to focus on close relationships, as explained in Appendix E.5.

To the best of our knowledge, the breadth of these networks spans more ground than any previous work, enabling us to compare the impact of demand externalities across different networks. Table B.1 presents summary statistics for each network. Networks of neighbors are fewer in number but larger in size, while friend networks are the most numerous and have the lowest median number of members.

Retailer Location and Gravity Zone We leverage data on corporate income tax returns spanning the universe of formal firms in the country. The data span 2015 to 2019, and includes typical balance sheet variables along with details on each firm's sector and location. Section 5 and Appendix J.1 discuss how we construct retailer-specific catchment areas leveraging details on their customers' location, available for a majority of retailers from electronic receipts data.

The instrumental variables strategy proposed in this paper requires three additional data sources, which are described below.

Family Networks We construct *nationwide* family networks based on information from Costa Rica's National Registry. This novel data includes official information to build each person's family tree based on existing records and without relying on name-matching. The data is dynamic and at a monthly frequency.

tionship as a friendship. If, however, both A and B have sent money to each other, then their relationship is classified as a friendship.

Networks of International Migrants In Costa Rica, voting is mandatory, and it is among the countries with the highest number of migrants residing abroad who register to vote at their corresponding Costa Rican consulate (approximately 51%).[16] This registration results in a record of migrants' foreign residences maintained by the Supreme Electoral Tribunal. The data maps citizens residing abroad to the consulate closest to their residence from 2014 to 2022. Large countries, such as the U.S., have multiple consulates, typically in cities with a high mass of Costa Ricans.[17] While this information is available in other countries, to the best of our knowledge, this is the first paper to leverage it to recover international migrant networks.

Consumer Dynamics in the U.S. We obtain U.S. consumer trends by product from two alternative sources, which complement each other. First, we rely on the Consumer Expenditure Survey (CEX), which includes quarterly data by Metropolitan Statistical Area (MSA) for 700 categories of products between 2015 and 2022.[18] Although the MSAs for which estimates are produced do not cover the entirety of U.S. territory, they include every city that hosts a Costa Rican consulate, corresponding to the main destinations where Costa Ricans reside abroad. In fact, according to the American Community Survey, over 82% of Costa Ricans living in the U.S. reside in one of these cities during our sample period. The CEX data (UCC codes) can then be mapped to HS codes using the concordance developed by Furman et al. (2017); also used by Hottman and Monarch (2020) and Borusyak and Jaravel (2021). The variation from this mapping is mainly at the HS-4 or HS-6 level (see Table C.3), as CEX categories are often more aggregated than customs' HS codes.

While our main results are based on the CEX, we leverage a second source of data on consumer trends by product, which aims to complement the CEX, precisely by providing variation for narrower product codes. The logic behind this second source is the following: In the U.S., many tradable products are imported. Thus, U.S. expenditures on these products by region should co-move with the imports of

[16] For instance, the equivalent share of migrants residing in the U.S. and registered to vote in their home country in Mexico is 1.5%, and the median in Latin America is 17.6%.

[17] The cities with Costa Rican consulates with ratified voting centers are: Atlanta, Chicago, Houston, Los Angeles, Miami, New York, and Washington D.C. There are also honorary consulates in Minneapolis, Puerto Rico, and Tucson. These consulates partition the U.S., and the area serviced by each consulate is well-defined.

[18] Details on geographic coverage are available in the BLS website (link).

these products in those areas.[19] Following this idea, we use HS-10 level quarterly imports by customs districts in the U.S. from the Census Bureau, which include over 20 thousand product codes, to obtain variation at the HS-10 level (see Table C.3). Conveniently, while HS-10 codes do not necessarily coincide across different nations, U.S. being Costa Rica's main trading partner, they do align for these two countries.[20]

The U.S. has 47 customs districts; instead of assuming a product is consumed in the same customs district it is imported into, we follow Acosta and Cox (2019) and allow for movements of imports within the U.S. using data from the Department of Transportation's Freight Analysis Framework (FAF), which provides estimates of where imported goods travel once they enter into U.S. borders across 132 FAF zones.[21]

Reassuringly, Appendix B.1 presents evidence—both in levels and in *changes*—of a strong correlation between expenditures in the CEX and the one-quarter lagged value of imports by product code and by region. This lag is intuitive, since goods crossing the U.S. border take time to reach retailers and, in turn, to be purchased by households and recorded in the CEX. Accordingly, throughout the paper, we use one-quarter lagged U.S. imports as a proxy of contemporary expenditures on those products. In line with this strong correlation, we show that our main results are statistically equal regardless of whether we measure U.S. expenditures via the CEX or via the U.S. imports data.

To further validate the CEX, Appendix B.1 also leverages transaction-level data on debit card expenditures by region and by Merchant Category Codes (MCCs) in the U.S., with over 10 million cards between 2017 and 2020.[22] While imperfect, as MCCs tend to be designed for financial transaction tracking, it is reassuring that—just as in the case of customs data—the correlation between CEX and card expenditures by region and product code is strong both in levels and in changes.

[19] Note how it is helpful that we will ultimately rely only on variation in (internationally) tradable products, not on changes in expenditures on non-tradables.

[20] Specifically, we manually check that the definition of each HS-10 code that is imported by individuals in Costa Rica is the same as in the U.S.'s Harmonized Tariff Schedule (HTS).

[21] While FAF zones are of moderate size, it will become clear later that, for our purposes, it is not crucial to pinpoint the precise location where a good was consumed; instead, we are interested in the broad area within the U.S. where consumption took place.

[22] These data come from Facteus, a provider of financial data for business analytics.

3 Stylized Facts on Individual Imports

While the determinants of trade between firms have been largely studied, the same is not true for goods imported directly by consumers. This section documents new patterns that govern the decision to import by individuals, and compares them with those in overall trade. These stylized facts, in turn, will be useful to understand the role of peers in individual imports in the next section of the paper.

Which goods are being imported? We first document which are the types of goods that are most commonly imported by individuals. Table A.1 shows a top-10 ranking which results from collapsing imports from HS-10 to HS-4 categories to be more informative.[23] As shown, the most popular category by far is 6204, which includes several types of women's or girls' clothing items. The top categories also include some types of motor vehicles, bags, men's and boys' clothing, toys, and household items. These top categories imported by individuals are very different from the top codes of final goods imported in general at the country-level, and have a modest overlap with the top codes imported by retailers, as reported in Table A.2. Individual imports of these goods are sizable and represent 6% of retail sales, 1% of total household consumption, and 0.6% of GDP. Moreover, the HS-4s imported by individuals represent 65% of the total value sold by retailers (and 54% of the total number of HS-4s sold by retailers).

Which origins? Table 1 reports the top origins of imports by individuals, both according to volume of imports and value of imports. The U.S. is the origin of most goods imported directly by consumers, followed by China and Japan. Thus, while the rest of the paper will focus on imports from the U.S., it is worth highlighting that this already captures a large share of the relevant imports. The U.S. is also the main origin for retailers' imports of final goods (35%), followed by China (14%). During our time period, Costa Rica had both a de-minimis regime and a Free Trade Agreement with the U.S., which meant zero tariffs for most U.S. imports.[24]

[23] Without such a collapse, most top 10 imports would belong to the 6204 category. Within this HS-4, the most popular HS-6 categories for women are cotton trousers (620462), trousers of synthetic fibers (620463), and dresses of synthetic fibers (620443).

[24] The main U.S.-based e-commerce platform visited by Costa Ricans during this time period (2015-2019) is Amazon.com.

Table 1: Top Origins of Individual Imports

Volume		Value	
(1)	(2)	(3)	(4)
Origin	Percentage	Origin	Percentage
United States	61%	United States	49%
China	17%	Japan	14%
Japan	2%	China	9%
Bolivia	2%	South Korea	7%
Mexico	2%	Canada	4%

Notes: The table reports the main origins of imports by individuals, both by volume of imports and imports' value. Data spans 2015-2019.

Gravity and individual imports We now investigate if individual imports are governed by gravity. For comparison, Table 2 reports the results of estimating a traditional gravity equation on individual imports, on imports of final goods, and on total imports. Remarkably, the roles of origin GDP and distance for individual imports in column (1) are both statistically equal to the role of the same variables in other columns *and* to the role of these variables in the literature (Head and Mayer, 2014). There are, however, a few differences between typical gravity variables in individual imports and those variables in column (2) and in the broader literature. Namely, contiguity seems to play a much more important role for individual imports, and common language, which tends to lead to more total imports, is insignificant (and has a negative coefficient) in the case of individual imports.

Table 2: Gravity Equation for Individual, Final Goods, and Total Imports (PPML)

	(1) Individual Imports	(2) Imports of Final Goods	(3) Total Imports
Origin GDP	1.358	1.340	1.247
	(0.113)***	(0.083)***	(0.067)***
Distance	-0.889	-1.110	-1.901
	(0.237)***	(0.138)***	(0.126)***
Contiguity	2.968	0.616	0.981
	(0.462)***	(0.501)	(0.424)**
Common Language	-0.381	0.584	0.721
	(0.331)	(0.391)	(0.296)**
Colonial Dependency	0.303	-0.086	-0.137
	(2.710)	(0.366)	(0.277)
Time FE	Yes	Yes	Yes
Pseudo R^2	0.853	0.927	0.922
Observations	925	925	925

Notes: The table reports the coefficients resulting from a gravity equation with individual imports as a dependent variable in column (1), with imports of final goods in column (2), and all imports in column (3). The estimation relies on PPML. GDP considers PPP from the World Development Indicators database. Other variables are obtained from CEPII (Conte et al., 2022). Standard errors are clustered by origin country. Data is collapsed to the annual level and spans 2015-2019.

Frequency of importing We also document that, in stark contrast with how retailers import recurrently, nearly all individuals import each HS-10 category only once, also indicating that imports from informal retailers are not widespread. As shown in Figure A.1, repeated imports of the same product are rare (Panel (a)) and almost all imported values are below $100 (Panel (b)).

4 Direct Externalities in Individual Imports

We provide a simple theoretical microfoundation in Appendix C.1 to guide the construction of our exposure measure.[25] Our analysis of *direct* externalities then turns to two empirical questions. First, we study whether expenditure shocks propagate through international migrant networks, generating demand shocks in Costa Rica. Second, we examine whether, once an individual imports a product, the probability that others in her local network import the same product changes, and how these effects vary across products and network types.

4.1 Empirical Strategy

Our empirical strategy leverages several aspects of our setting and data. In particular, we construct an instrument based on the following idea, which we will formalize in the next section: *Suppose both L and N live in Costa Rica, and L has a sister living in Los Angeles (LA) while N has a sister living in New York City (NYC). If product i becomes more popular in LA as compared with NYC in period t, then L is more likely than N to import product i in period $t + 1$.*

Panel A of Figure 1 provides a general summary of the instrument. Following the figure's notation, suppose a family in Costa Rica has a member who migrated to a U.S. city on the West Coast, in blue, and a specific product p' becomes more popular in this West Coast city as compared with other cities in the U.S. Then, relatives of this migrant in Costa Rica become more exposed to this product and more likely to import it than other Costa Ricans with relatives in different U.S. cities.

The spirit behind this instrument is how, anecdotally but also intuitively, information on product dynamics is transmitted to developing countries once relatives migrate to developed countries, where more products are available. The instrument

[25] As we elaborate below, the model rationalizes the exposure measure by formalizing the idea that individuals receive noisy signals about shocks through relatives abroad.

Figure 1: Instrument Relying on International Family Networks and Exogenous Product Trends

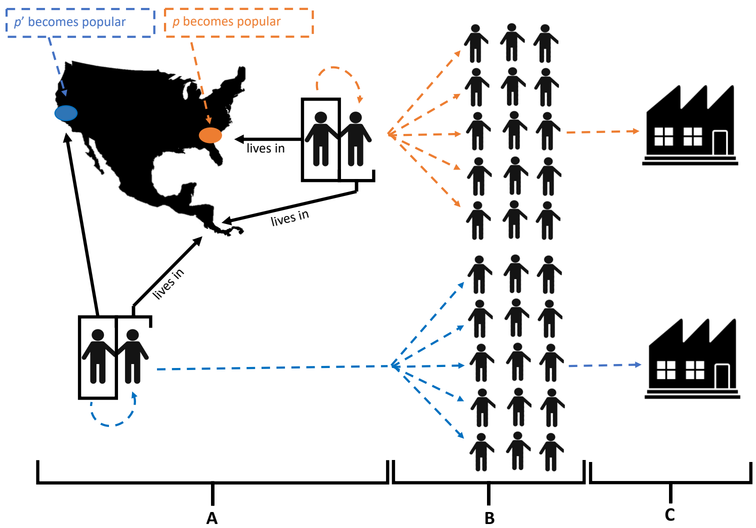

Notes: The figure summarizes the idea behind our main instrument and analysis. The instrument (Panel A) leverages information on the family networks of migrants to different U.S. cities, along with variation in product trends across these cities. Panel B represents our second stage, in which we measure the effect of exogenous exposure to a product on the probability of importing the same product for people who share a network. Panel C represents our study of the supply response after individuals import a product.

exploits (i) that we can identify Costa Rican citizens who are living abroad along with the location where they reside in the U.S. (1% of the population), (ii) that we are able to link these migrants to their close relatives who still reside in Costa Rica (5% of the population), (iii) that we have data at the MSA-level and customs district-level to follow product-specific dynamics across the U.S. over time, (iv) that product-specific consumer trends in the U.S. do not respond to local conditions in Costa Rica, as Costa Rican migrants are a small share of the U.S. population, and (v) the availability of individuals' product-specific imports at a high frequency. Moreover, the instrument serves a dual purpose: in the *first stage*, it enables us to test whether product-specific expenditure shocks propagate across international migrant networks, generating demand shocks in Costa Rica that we leverage in our *second stage* to explore peer effects (from individuals with relatives in the U.S. to those without such connections).

Expenditure Shocks in the U.S. We construct a measure of product-specific expenditure shocks in the U.S., which can vary across time, cities, and products, but that we can purge from business cycles in the U.S., differential level effects, and national product trends. More rigorously, let s denote a U.S. city, p a product, t a quarter, and c a Costa Rican consulate in the U.S. Consider the following specification:

$$\ln E_{spt} = \underbrace{\lambda_{sp}}_{\text{level}} + \underbrace{\mu_{st}}_{\text{local business cycles}} + \underbrace{\phi_{pt}}_{\text{national product trends}} + \ln \widetilde{E}_{spt}, \tag{1}$$

where E_{spt} are expenditures in city s on product p at time t. Let $\ln \widetilde{E}_{spt}$ be the residuals of this regression, which would capture the differential product trends across U.S. cities. Note that the fixed effects would prevent $\ln \widetilde{E}_{spt}$ from varying (i) because people in a city are more prone to buy a certain product, for example, those in Chicago buying more winter coats (level effect); (ii) because a particular region had a positive or negative income shock (local cycles); or (iii) because a product became more or less popular (national product trend). Table C.4 performs variance decomposition to better grasp the relevance of each fixed effect in this residualization.

To better understand the underlying variation, Appendix C.3 draws on detailed microdata to show that the residuals are driven by local product dynamics. Across product categories, movements in the residuals closely track the entry and exit of stores and product brands in each location. We also provide illustrative examples of how residual dynamics coincide with local shocks in Appendix C.4.

As Costa Rican consulates can span several cities, we aggregate our measure to the consulate-level weighting by population shares in each city:

$$\ln \widetilde{E}_{cpt} = \sum_{s \in c} \left(\frac{CR_s}{CR_c} \right) \ln \widetilde{E}_{spt}, \tag{2}$$

where CR_s/CR_c is the share of Costa Ricans in consulate c who reside in city s.[26] Section 4.5 uses randomization inference in support of these shocks being unconditionally randomly assigned across consulates.[27] Furthermore, as reported in Appendix C.2, we conduct a battery of tests, all of which reject serial correlation in $\widetilde{E}_{cpt}$.

[26]Costa Ricans by city are obtained from the American Community Survey. The share is time-invariant as the average Costa Rican residing in the U.S. by 2019 migrated in 1994; i.e., movements abroad are rare, and could lead to selection which we prefer to shut down. Thus, we fix these shares in 2014, one year before our sample period starts. The exposure measure remains virtually unchanged if we first aggregate the expenditure data to the consulate level, and then run the two-way saturated fixed effects specification, as shown in Figure C.4.

[27]Moreover, as we show in Section 4.5, a "recentered" version of our exposure measure à la Borusyak and Hull (2023) delivers results which align with those of our baseline exposure.

4.2 Expenditure Shocks Propagate Across Migrant Networks

The first contribution of our paper is to show how product-specific expenditure shocks propagate across international migrant networks. First, we show that this holds at the individual level. Then, we construct an instrument for the first stage of a series of 2SLS analyses to study direct and indirect demand responses.

4.2.1 Individual-Level Analysis

We consider an individual-level specification, which examines if Costa Ricans with relatives in the U.S. respond to product-specific shocks where their relatives reside. Appendix C.1 provides a microfoundation for this specification, interpreting the consulate–product–time exposure as a latent expenditure shock to a given product in a given U.S. location and period. Individuals receive noisy signals about this shock through their relatives abroad and update their beliefs accordingly. Under a standard Gaussian learning framework, posterior expectations move proportionally with the realized shock, delivering a reduced-form first stage in which importing responds linearly to log exposure. Therefore, we consider:

$$\text{Import}_{ipt}^{US,\,direct} = \beta_0 \ln \widetilde{E}_{cp,t-1} + \gamma_{i\tilde{p}} + \gamma_{it} + \varepsilon_{ipt}, \tag{3}$$

where $\text{Import}_{ipt}^{US,\,direct}$ equals one if individual i *with a relative in U.S. consulate* c imports an HS-10 product p at time t for the first time.[28] A "first-time" import is such if the individual has not imported the HS-10 product since 2005.[29] Term

[28] Throughout the paper, we define a "relative" as including parents, siblings, own children, partner, and partner's parents, siblings, and children. Furthermore, while most people with relatives in the U.S. have all relatives living in the same consulate, there are a few individuals with relatives in different consulates (95.8% of them are connected to a single consulate and 99.9% to two consulates). For the few cases connected to multiple consulates, our calculation considers a weighted sum of relatives. Namely, $\sum_c s_{ic} \ln \widetilde{E}_{cp,t-1}$, where s_{ic} denotes i's relatives who reside in consulate c as a share of all her relatives who live in the U.S. Results are robust to dropping individuals with relatives in multiple consulates.

[29] Throughout the paper, whenever the dependent variable is defined as importing product p "for the first time" at time t, the corresponding unit–product pair remains in the panel after adoption and is coded as zero in all subsequent periods. Hence, first adoption is absorbing in the outcome definition but not in sample inclusion: adopters are not removed from the estimating sample after their first import. Likewise, $\text{ShareImporters}_{bpt}$ is computed each period using the full relevant network population in the denominator, rather than only those units that have not yet imported p. Hence, the dynamic estimates are identified from the timing of new first-import events within a stable panel, and any changes in denominators

$\gamma_{i\tilde{p}}$ is an individual-product fixed effect, which controls for either HS-4 codes for the case of the CEX or HS-6 for the case of U.S. imports data. The latter is quite demanding, as most of the variation in the CEX is at the HS-4 level (see Table C.3). While equation (1) residualizes based on HS-10 categories, further saturating this specification (and subsequent ones) with finer product fixed effects often results in imprecise estimates.[30] Finally, γ_{it} denotes an individual-time fixed effect—note that a product-time fixed effect would not alter the estimates, as all product×time variation was already removed from the exposure in equation (1). Standard errors are clustered at the individual-product (HS-4) level.

Results are presented in Table E.2. The regressions, based on over 700 million individual-product-time observations, find that a one standard deviation increase in the exposure to a product leads to a 12% higher probability of importing that product next quarter, compared to other Costa Ricans with relatives in the U.S. The table reports a similar effect using U.S. imports by customs districts to construct exposure.

4.2.2 Network-Level Instrument and First Stage

We leverage this result as a building block in setting up the first stage of a 2SLS. Importantly, note that any instrument *must* be defined at the network level. To illustrate why, consider an individual in the second stage *without* relatives abroad. Our goal is to determine whether her exposure via others in her network—with relatives abroad—affects her importing behavior. But to do so, we must account for the exposure of *all* individuals in her network with relatives in U.S. cities, aggregating across all cities. Therefore, we now construct a measure to summarize exposure by network to understand if those without relatives abroad react to the exposure of their peers.

Network-Level Exposure We consider a linear-in-means exposure mapping. Appendix C.1 derives a microfoundation where we model peers as allocating limited attention across potential sources of information (coming from those with relatives in different consulates), where deviating from baseline source availability is costly in relative-entropy (KL-divergence) terms. The resulting optimal aggregation delivers

arise only from genuine changes in network membership (e.g., coworker turnover), not from past adoption of product p.

[30] To avoid redundancy, and as the network-level result is the one used as a building block in other sections, details on timing, clustering, and the role of each fixed effect will be discussed at length in the next subsection.

a linear-in-means exposure mapping, which is standard in the network-interactions literature (Bramoullé et al., 2020; Goldsmith-Pinkham and Imbens, 2013; Manski, 1993). Therefore, we construct the exposure mapping of a Costa Rican in network b to product p via imports as follows:

$$\ln \widetilde{E}_{bpt} = \sum_{c} s_{bc} \ln \widetilde{E}_{cpt}, \tag{4}$$

where s_{bc} is the share of U.S.-connected individuals in network b whose relative is assigned to consulate c.[31] As discussed in Section 4.5 and Appendix C.1, results are robust to alternative aggregations, microfounded by distinct attention costs, including a hard-max and a smooth-max.

First Stage Let $\text{ShareImporters}_{bpt}^{US,\,direct}$ denote the share of individuals within network b who have relatives in the U.S. and import product p for the first time at time t.[32] Both this variable and our instrument defined in equation (4) vary at the network–product–time level. We estimate the model separately by network type on an individual×product×time panel, assigning the corresponding network-level variables to each individual observation.[33] We then estimate

$$\text{ShareImporters}_{i,bpt}^{US,\,direct} = \beta_1 \ln \widetilde{E}_{i,bp,t-1} + \gamma_i + \gamma_{b\tilde{p}} + \gamma_{bt} + \varepsilon_{i,bpt}. \tag{5}$$

Estimating this equation on the stacked individual-level panel aligns the first-stage projection with the fixed-effects structure used throughout the paper. Accordingly, the number of observations reported in the first-stage table corresponds to the number

[31] Note that s_{bc} is fixed over time; Costa Ricans living abroad and all networks are set to 2014, a year before the start of our analysis, except for coworkers where the average turnover per year is approximately 40%. For this network only, we allow individuals to move across firms and recompute consulate-share weights over time using each firm's contemporaneous composition, while the underlying migrant network remains fixed. Appendix H.4 further uses a proxy of the historical 2005 migrant network.

[32] As detailed in Appendix E.5, those who use the payments app *and* have a relative in the U.S. are more likely than average to import. Thus, to obtain a strong instrument, for the first stage of this network only, we focus on first-time imports *within sample* (2015-2019), but do not force them to be first-time imports ever—which we do for other networks and *all* second stages, as our panel starts in 2005, when individual imports were virtually zero.

[33] Therefore, we redefine $\text{ShareImporters}_{i,bpt}^{US,\,direct} \equiv \text{ShareImporters}_{bpt}^{US,\,direct}$ and $\ln \widetilde{E}_{i,bpt} \equiv \ln \widetilde{E}_{bpt}$, although these variables do not vary across individuals within a (b, p, t) cell.

of individual–product–time observations in the estimation sample, while the identifying variation is at the network–product–time level.

The terms $\gamma_{b\tilde{p}}$ and γ_{bt} denote network–HS-4 product code and network–time fixed effects, respectively, and γ_i are individual fixed effects. As in equation (3), $\tilde{p}$ defines products at either the HS-4 or HS-6 level. Including product–time fixed effects would not affect the estimation, since all HS-10×time variation is already absorbed in equation (1). The regression is estimated separately for each network type, so that $b \in B$, where B denotes a network type—neighborhoods, firms, or friendship networks. Standard errors are clustered at the network–product (HS-4) level. Because different networks can load on the same consulate×product×time shocks, Appendix E.1 reports alternative clustering schemes, including shifter-level clustering by HS-10×time; the resulting standard errors are very similar.[34]

A few remarks are in order. A few remarks are in order. First, the variables $\text{ShareImporters}_{bpt}^{US,\text{direct}}$ and $\ln \widetilde{E}_{bp,t-1}$ are constructed exclusively from Costa Rica residents with a relative living in the U.S., and are then mapped to each individual i based on network b, product p, and time t. Second, the left-hand-side variable is conservative as it includes only first-time imports.[35] Third, the battery of fixed-effects strengthens identification. Namely, addresses the "friendship paradox" (Aronow and Samii, 2017), and manages interference in network settings (Borusyak and Hull, 2023), as such interference is product-invariant and inherent to the network. Additionally, γ_{bt} controls for correlated shocks (Manski, 1993), underscoring how multiple products help resolve challenges unique to single-product settings. On its part, $\gamma_{b\tilde{p}}$ controls for the relevant product level variation and for network taste and characteristics, for instance, addressing if a rich neighborhood tends to import a product; more generally, it would account for exogenous effects (Manski, 1993). While we saturate the regression with fixed effects, significant variation remains; to visualize it, we compute $\ln \widetilde{E}_{bp,t}$ *netted of fixed effects* in equation (5), and calculate its variance for each network-product pair. Figure C.6 shows this variation across network types and

[34] Appendix E.1 also discusses our clustering choice and reports clustering by network×HS-10 and by HS-10. These alternatives yield standard errors that are very similar to, and in some cases smaller than, those obtained under our baseline network–product clustering.

[35] Note that measurement error on the left-hand-side variable would, in general, not bias this coefficient. This result holds as long as the exposure's residual is uncorrelated with the measurement error, which in our case is likely to occur.

products. Finally, the timing of equation (5) is guided by local projection exercises and survey data on typical import durations (Appendix E.6); these results shed light on the dynamics of the effect and motivate the lag structure in our specifications.

The results of this first-stage are shown in Table 3.[36] Instruments are strong for every network, as reflected by the F-statistics. Moreover, results are remarkably similar across network types—neighbors, coworkers, and friends: A one standard deviation increase in exposure to a product leads to a 11-15% higher share of individuals *with relatives abroad* importing this product next quarter.[37] Table E.5 displays analogous results relying on U.S. imports by customs districts to construct the instrument; reassuringly, they are statistically equal to the baseline results based on CEX data.

Table 3: First-Stage Regressions

Dep. variable: Share of importers of product p
with a relative in the U.S. and who belong to network b at time t

	%Δ w.r.t. mean import probability		
	Neighbors (1)	Coworkers (2)	Friends (3)
$\ln \widetilde{E}_{i,bp,t-1}$	10.632	15.050	10.708
	(1.851)***	(4.639)***	(2.708)***
F-statistic	33.00	10.52	15.63
Observations	289,340,892	299,920,162	260,952,672
Clusters	200,308	236,804	4,568,240
Mean import prob. $[i, bpt]^{US}$	0.001	.0003	0.001
Mean import prob. $[bt]^{US}$	0.150	0.114	0.447
$b\tilde{p}$, bt, i FE	Yes	Yes	Yes

Notes: The table shows our first stage results. Robust standard errors, adjusted for clustering by network-product (HS-4), are in parentheses. Coefficients report the percent change in the mean import share implied by a one–standard-deviation increase in exposure. Exposure is standardized. We include network×product (HS-4), network×time, and individual fixed-effects. Percentage mean import probabilities are reported. Appendix E.5 presents details on the sample used in each regression. Data is quarterly and spans 2015-2019.

4.3 Importing Externalities Across Domestic Networks

Relying on the instrument above, we want to understand if people in a network *who are unrelated to migrants in the U.S.* increase their probability to import a particular

[36] Appendix E.5 provides details on the samples of products and networks used in each regression.

[37] A one standard deviation is equivalent to an increase of $1.5 per capita in the average U.S. consulate. Note that regressions control for network-time fixed effects, thus, for instance, network size would not affect our coefficients.

product after being exposed to it via their peers who do have relatives abroad (Panel B of Figure 1). Thus, from our IV's first stage, we leverage predicted values for the shares of importers of a product *with relatives in the U.S.* as explanatory variable. Our dependent variable would instead depend on the probability of importing a particular product for people in the network without relatives in the U.S., as follows:

$$\text{Import}_{i,bpt} = \beta_2 \, \widehat{\text{ShareImporters}}^{US,direct}_{bp,t-1} + \gamma_i + \gamma_{b\tilde{p}} + \gamma_{bt} + \varepsilon_{i,bpt}, \tag{6}$$

where $\text{Import}_{i,bpt}$ equals one if individual i in network b, who does not have relatives in the U.S., imports product p for the first time at time t. The corresponding excluded instrument is exposure at $t-2$, and the set of fixed effects coincides with that in equation (5).[38] Just as for the first stage, we consider three network definitions: neighbors, coworkers, and friends, and run independent regressions for each of them. Each network has complementary strengths. Networks of neighbors span all the population and allow us to study the role of indirect demand externalities in triggering a supply-side response. Networks of coworkers may exhibit less correlated spatial shocks but cover only the formally employed. Networks of friends are a novel way of measuring connections beyond observables, but the analysis is limited to those people who have adopted the mobile payment app. Overall, utilizing all networks paints a better and more robust picture of the role of demand propagation in product adoption.[39]

The baseline results of the 2SLS estimations are shown in Table 4.[40] The magnitudes of the 2SLS coefficients are similar across network types. We find that a

[38] As there are millions of Costa Ricans *without* relatives abroad and a 2SLS with billions of observations is unfeasible to run, we use a random sample as detailed in Appendix E.5.

[39] Because individuals may belong to multiple social environments, it is useful to clarify the interpretation of β_2. We conduct estimations separately for each network definition, rather than stacking alternative network types into a single regression. Within a given specification, the endogenous peer variable and the instrument are network$\times$product$\times$time objects that are assigned to individuals in that network so that the first and second stages share the same fixed-effects structure. Under the maintained exclusion restriction, β_2 therefore identifies the causal effect of an instrument-induced increase in adoption within the particular reference group under study. Since social ties overlap, a given adoption event may affect exposure under more than one network definition across separate specifications. Coefficients across network definitions are therefore not additive and should not be interpreted as a structural decomposition of total diffusion; rather, each specification quantifies diffusion for a particular definition of the relevant reference group.

[40] Appendix E.5 describes the samples of products and individuals used in each regression.

Table 4: 2SLS: Propagation Within Network

Dep. variable: Prob. importing product p for individual i without relatives in the U.S. and who belongs to network b at time t

	%Δ w.r.t. mean import probability		
	Neighbors (1)	Coworkers (2)	Friends (3)
$\widehat{\text{ShareImporters}}^{US\,direct}_{bp,t-1}$	13.897	18.069	14.679
	(5.193)***	(8.154)**	(5.027)***
F-stat first stage	33.00	10.52	15.63
Observations	289,340,892	299,920,162	260,952,672
Clusters	200,308	236,804	4,568,240
Mean import prob. $[i, bpt]^{US}$	.0002	.0001	.0004
Mean import prob. $[bt]^{US}$	0.044	0.052	0.156
$b\tilde{p}$, bt, i FE	Yes	Yes	Yes

Notes: The table displays the results of running our 2SLS. Robust standard errors, adjusted for clustering by network-product (HS-4), are in parentheses. Coefficients are percentages of mean outcome for a one standard deviation change in the instrumented share. Regressions control for network×product (HS-4), network×time, and individual fixed-effects. Percentage mean import probabilities are reported. Appendix E.5 details the sample for each regression. Data is quarterly and spans 2015-2019.

one–standard deviation increase in the instrumented share of importers with U.S. relatives increases next-quarter importing by 14–18% relative to the mean.[41] Put differently, a 10 pp increase in the share of those in a neighborhood with relatives abroad who import a product leads to a 4.7 pp higher probability of importing this product for individuals in this neighborhood without relatives abroad, with the corresponding effects for coworkers and friends being 3.5 pp and 4.7 pp, respectively—see Table E.4 for results without normalizations and Table H.1 for results with less stringent fixed effects. Furthermore, as network types vary in size (Table B.1), a 10 pp increase in the share of those with relatives abroad means different *absolute* changes in the number of people importing.[42] In fact, one additional friend who imports a new product has an impact on the importing probability of those in her network equivalent to that of 33 coworkers and 91 neighbors. In other words, a friend is more influential than the average coworker, and even more so than the average neighbor.[43] Appendix

[41] Throughout the paper, we scale effects by the standard deviation of the instrumented share, so that a one standard deviation change corresponds to the typical exogenous shift in exposure used to identify the 2SLS estimate.

[42] While the median network of neighbors has 781 people—and therefore a 10 pp increase spans several individuals—the median close friends network has 7 individuals.

[43] While the proportion might seem large, perhaps the following perspective is helpful: if we take a random sample of 87 people from an individual's neighborhood, only one of them is likely to be close enough to her to exert a peer effect comparable to that of a close friend.

F also compares the magnitudes of the direct and indirect exposure; a back-of-the-envelope calculation delivers that an indirectly exposed individual (like a neighbor) is less effective in generating an import than a relative abroad. Reassuringly, results when relying on U.S. imports data to construct our instrument are very similar (and statistically equal) to the baseline results based on the CEX, as shown in Table E.5.[44]

4.4 Aggregate Effect of the Direct Externalities

Mean import probabilities in each regression are small *by design*: They represent the probability that an *individual* in a given network orders *a particular product from the U.S. for the first time* on a determined *quarter*, as the left-hand-side variable is individual-product-time-specific and "turns on" only with first time U.S. imports.[45] To understand the economic relevance of the estimates, one has to consider the increase in import probabilities across millions of individuals, networks, and products, which we do as a back-of-the-envelope exercise.

Combining results from our first and second stages in equations (5) and (6), and defining $Importers^{US,direct}$ and $Importers^{US,indirect}$ as the total number of individuals in the directly connected group and the indirectly exposed group, respectively, we can obtain a back-of-the-envelope estimate of how U.S. spending on a product would influence total Costa Rican individual imports of that product. Namely, to translate the extensive-margin individual responses into aggregate trade values, we scale the estimated change in the probability of first importing a product by the average dollar value of a first import. This mapping provides an approximate conversion from adoption probabilities to expected import flows.[46] Specifically, we compute the following

[44] Appendix E.3 discusses the ordinary least squares (OLS) and the reduced form results.

[45] While we have run all these regressions at the network-product-quarter level, delivering much larger import probabilities, individual-level regressions are cleaner, more transparent, and better identified. Results at the network level are consistent with individual-level results and available upon request.

[46] Implicitly, the calculation assumes that the average value of imports per adopting consumer is stable and that substitution across products or adjustments along the intensive margin are limited over the range of variation considered. The resulting dollar multipliers should therefore be interpreted as a back-of-the-envelope measure of the trade implications of the extensive-margin responses we estimate.

objects following a change in U.S. spending on a product:

$$A = \underbrace{\Delta ShareImporters^{US,direct}}_{\beta_1 \Delta \widetilde{E}_{bpt}} \times Importers^{US,direct} \times AvgPrice^{Individual} \tag{7}$$

$$B = \underbrace{\Delta Prob(Import_i)}_{\beta_2 \Delta ShareImporters^{US,direct}} \times Importers^{US,indirect} \times AvgPrice^{Individual}, \tag{8}$$

where $AvgPrice^{Individual}$ is the average value of an individual import. A in equation (7) would then speak to the value of new imports of directly exposed individuals, as it is the product of the change in the share that imports, the total importers with direct connections, and the average price of an individual import. B in equation (8) captures the value of new imports of those indirectly exposed via their network. Using networks of neighbors as our reference, a comparison of B/A in the above equations, along with the estimation from column (1) of Table 4, imply that a new exogenous individual import generated by the directly exposed consumer in her network would lead to approximately 13 new imports by indirectly exposed people in her network.

As an example, suppose that $\Delta \widetilde{E}_{bpt}$ corresponds to a \$100 increase in U.S. per capita spending on a product. In this case, this calculation implies that Costa Rican consumer imports per individual with a direct connection to the U.S. would increase by 40 cents, given the degree of interconnectedness between the two countries and the strength of the demand propagation across individuals.[47] As will be studied in Section 5.3, while this multiplier is non-negligible, the impact of imports by individuals (i.e., $A + B$) will represent only a modest fraction (about 8%) of the total impact on local demand, as the bulk of the increase will emerge from retailers' responses.

4.5 Alternative Specifications and Robustness

It is worth spelling out the exclusion restriction of our instrument. Our identification strategy requires that the likelihood of buying product p of a Costa Rican—without

[47]Note that this multiplier is calculated in per capita terms, where the relevant denominator in Costa Rica is the number of individuals with a relative abroad. We consider this to be the relevant parameter, as these individuals with relatives in the U.S. act as the bridge through which demand propagates. A more connected pair of countries (with more directly exposed individuals) would likely face a stronger total degree of propagation, which is why a per capita estimate is informative.

relatives abroad—in a network connected to a U.S. city via family ties co-moves with *changes* in expenditures on p in this U.S. city only through the relatives' influence.[48] Arguably, our main specification, which is saturated with a battery of fixed effects, addresses most first-order concerns related to this assumption.

To further probe this restriction, we pursue complementary approaches. First, we construct an alternative instrument. Second, we test the implication that the same U.S. shocks should not predict adoption in networks that lack ties to the corresponding U.S. locations, using placebo exposure measures that break the mapping between migrant networks and U.S. consulates. Third, we leverage a historical migrant network that predates our sample by 10 years. Finally, we estimate more stringent specifications.

Instrument Using Distance-3 Nodes We can push our data to construct an alternative instrument which can rule out lingering alternative hypotheses. The instrument exploits that we have information on both coworkers and spouses, and that spouses who work at different firms can be seen as a bridge between sets of coworkers that are otherwise disjoint; an observation in De Giorgi et al. (2019), and that we extend to include product-level variation and an instrument for the 3-nodes-away exposure. The approach relies on the notion that *if the coworker of the spouse of my coworker* has a relative in the U.S. and becomes exogenously more likely to import a product—controlling for common shocks experienced at my firm—this should not influence my probability of importing this exact same product directly, only indirectly though peer effects. Figure 2 presents a diagram to make this notion more clear; it considers an individual i working at firm D. The individual's exposure to a particular product p depends on the exposure to p faced by the spouses of her coworkers (D1, D2, D3) at their firms, which in turn depends on the family ties that employees of those firms have with people residing in different U.S. cities, and on how expenditures of product p evolve in those cities. We then consider the following regression for individual i, which depends on product p at time t:

$$Import_{i,Dpt} = \delta_0 + \theta \widehat{ShareImporters}_{ip,t-1} + \delta_{D\tilde{p}t} + \delta_i + \varepsilon_{i,Dpt}, \tag{9}$$

[48]While, given our fixed effects, this statement only has to hold in changes, Appendix C.6 also presents evidence in support of this statement in levels. We find balanced observables (age, gender, wage) among Costa Rican migrants to different U.S. cities.

Figure 2: Diagram of Instrument Using Distance-3 Nodes

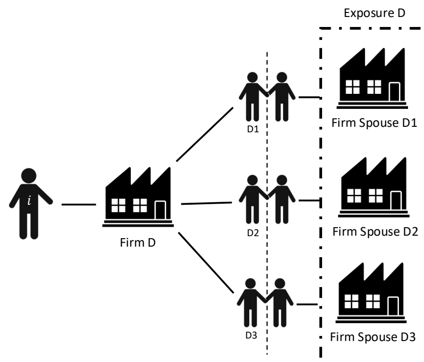

Notes: The figure shows the idea behind our instrument, where the relevant exposure is product-specific, time-varying, and depends on exogenous consumer trends, as described in Section 4.1.

where the dependent variable, $Import_{i,Dpt}$, is equal to one if individual i at firm D imports product p at time t for the first time. On the right-hand side of the regression, $\widehat{ShareImporters}_{ip,t-1}$ is instrumented by the mean exposure of firms employing the coworkers' spouses.[49] $\delta_{D\tilde{p}t}$ are own-firm×HS-4 product code×time fixed-effects; these fixed-effects are key, as they force most of the identifying variation to come from *differences* between the shocks at coworkers' spouses firms and shocks directly affecting firm D. Finally, δ_i are individual fixed-effects. This specification is very demanding; however, there is still significant variation left across individual-product pairs, summarized in Figure H.1. We also list the product codes which are among the top 10 codes of this subsample but were not in the top 10 codes in the full sample; remarkably, all of these have to do with work-related products. More details are available in Appendix H.1.

Key advantages of this instrument The distance-3 nodes instrument is immune to several identification concerns. As an example, consider correlated preferences among people in a network and their relatives abroad, i.e., people with relatives in NYC have different product-time demands than those with relatives in Houston, and their friends in Costa Rica show the same differential demand patterns. The concern would be selection into locations, so people go to NYC because the city's preferences are correlated with theirs and because they have Costa Rican friends/colleagues who are

[49]Note that De Giorgi et al. (2019) rely on characteristics of the firms employing the coworkers' spouses; in our case, this would imply using imports directly. Instead, we opt for a more demanding specification that instruments for those imports.

similar to New Yorkers.[50] Given this scenario—potentially the worst possible for our baseline instrument—there are two possibilities: (a) there is *not* assortative matching in the marriage market along lines which influence product demands, in which case the instrument based on Figure 2 would deliver a correct estimate; or (b) there *is* assortative matching in the marriage market along lines which influence product demands, in which case, the own-firm-product-time fixed effect in equation (9) would co-move with our instrument and would prevent θ from being identified from such assortative matching; again, this approach would deliver a correct estimate, as the fixed effects force the identifying variation to come from cases in which there is *not* perfect assortative matching along lines that influence migration to particular cities—cases in which migrant destinations at both firms are not aligned.

We find evidence in support of demand externalities, even under this more demanding specification. As shown in column (2) of Table 5, a one–standard deviation increase in the instrumented share of coworkers with relatives in the U.S. who import a product increases the probability that a coworker without U.S. connections starts importing the same product the next quarter by about 21% over the mean import probability. Moreover, reassuringly, we cannot reject equality of the estimated coefficients across alternative network definitions, including between the distance-3 specification and column (1) of Table 5, which is estimated on the same subsample of individuals. This stability suggests that the results are not driven by any particular way of partitioning the underlying social graph.[51]

Placebo Exposure Measures and Recentering A key identifying assumption is that the residualized U.S. city/consulate × product × time shocks $\widetilde{E}_{cpt}$ are randomly assigned, and affect Costa Rican adoption only through migrant-mediated social connections. To probe this exclusion restriction, we construct placebo exposure measures that break the link between the U.S. shocks and the migrant networks through which information could plausibly travel.

[50] Note that this would have to occur while keeping migrant observables balanced across locations (see Appendix C.6).

[51] However, columns (1) and (2) of Table 5 are not directly comparable, as the distance-3 specification implies exposure to different products. Note that the first stage F-statistic in column (2) is large, indicating that our foreign exposure measure strongly predicts imports by spouses' coworkers. This reflects the relevance of the instrument in the distance-3 specification.

Table 5: 2SLS: Individual Imports and Distance-3 Exposure

Dependent variable: $Import_{i,Dpt}$ *(Prob. of individual i of importing product p at time t)*	%Δ w.r.t. mean import probability	
	Baseline IV	Distance-3 IV
	(1)	(2)
$\widehat{ShareImporters}_{ip,t-1}$	19.724	21.514
	(9.247)**	(3.644)***
F-stat first stage	11.56	10,519
Observations	310,967,666	392,785,925
Clusters	557,863	17,331,982
Mean import prob. $[i, Dpt]^{US}$	.0003	.0003
Mean import prob. $[Dt]^{US}$	0.051	0.051
$D\tilde{p}$, Dt, and i FE	Yes	No
$D\tilde{p}t$ and i FE	No	Yes

Notes: Both estimations are constructed based on the same sample and run at the individual level. Robust standard errors, adjusted for clustering by network-product (HS-4) in column (1) and individual-product (HS-4) in column (2), are in parentheses. Coefficients are percentages of mean outcome for a one standard deviation change in the instrumented share. Column (1) controls for own-firm×HS-4, own-firm×time, and individual fixed effects, while column (2) controls for own-firm×HS-4 product code×time, and individual fixed effects. Percentage mean import probabilities are reported. Appendix E.5 presents details on the sample used in each regression.

Specifically, we randomly permute the shocks $\widetilde{E}_{cpt}$ across U.S. consulates, within HS-4 product categories and time periods. This reassignment preserves the distribution of shocks but severs the mapping between each Costa Rican network and the U.S. consulate to which it is socially connected. As a result, the permuted shocks generate placebo exposure for networks that have no systematic ties to the corresponding U.S. locations. We then re-estimate the reduced-form regressions using these placebo exposure measures and repeat this exercise 100 times. Thus, the permutation directly tests whether the same U.S. shocks predict adoption in networks that lack the corresponding migrant connections; they do not. Figure H.2 plots the distribution of placebo coefficients together with the coefficient obtained using the true exposure measure (red vertical line). The actual coefficient lies far in the tails of the placebo distribution, indicating that the relationship we estimate is not driven by mechanical correlations in product-time shocks or other common trends.

Further, following Borusyak and Hull (2020), we construct a "recentered" version of the exposure measure by subtracting the expectation of the treatment under the randomized distribution from the original exposure. As predicted by the theory, this procedure yields a slightly larger but very similar coefficient (Panel (c) of Figure

H.2 for neighbor networks), reinforcing that the results are not driven by spurious correlations in the underlying shock structure.

Historical Migrant Network Our baseline instrument uses the 2014 migrant network, measured one year before the start of the sample period, both because official records begin in 2014 and to ensure that the network shares entering the instrument are predetermined relative to the 2015–2019 shocks. Because the shares are fixed before the sample begins, the identifying variation cannot be generated by migrants moving across U.S. consulates in anticipation of the short-run shocks underlying the instrument. A remaining concern is that historical selection into particular consulates may be correlated with subsequent product-specific demand shocks in Costa Rica.[52] To address this possibility, we construct a proxy for the 2005 migrant network, a decade before our sample window. This alternative measure relies on historical settlement patterns and substantially reduces the role of migration and recent selection dynamics.[53] Reassuringly, estimates based on the 2005 network in Table H.3 are statistically indistinguishable from our baseline results. These findings are difficult to reconcile with a selection-based account in which consulate assignment is spuriously capturing later product-time demand shocks in Costa Rica.

Other Alternative Specifications Adding fixed-effects to our specification can be a powerful tool to rule-out alternative hypotheses. We first consider a *district*×*HS-4 product code*×*time* fixed-effect, and re-run the analysis defining networks as neighborhoods. Including this fixed-effect limits us to consider variation *within* small areas.[54] Results remain largely unchanged, as shown in column (1) of Table H.2. This is useful, for example, to rule out explanations in which sellers target specific areas of the country with marketing campaigns or content exposure related to a particular product category, as well as explanations driven by platform logistics or spatial differences

[52] Recall, however, that Appendix C.6 shows that observable migrant characteristics are balanced across U.S. consulates

[53] To do so, we identify migrants who were abroad in 2014 and at least 18 years old in 2005, and then use monthly employer–employee records for 2006–2014 to exclude any individual who appears employed in Costa Rica during that period. This procedure yields a proxy for the 2005 migrant network that excludes likely migrants from 2006–2014 (29 percent of the total).

[54] Each district has four neighborhoods on average.

in product availability.[55]

In a similar spirit, we add a *network×HS-2 product code×time* fixed effect, restricting identification to variation *within* relatively narrow product categories. As column (2) of Table H.2 shows, effects are again largely unchanged. Like the exercise using distance-3 nodes, this result speaks against people from a certain network having a preference for a product category, and thus moving to cities where this category is trendy.[56] Also, like the analysis with distance-3 nodes, this control would take care of sector-level trends in particular cities. Furthermore, a selection-based threat to identification would require pre-sample sorting across U.S. consulates to be systematically correlated with subsequent product-specific demand shocks in Costa Rica. The fact that the estimates are unchanged when we replace the baseline 2014 network with a proxy for the 2005 network strongly mitigates this concern. Overall, the distance-3 and related robustness exercises discussed above suggest that confounding factors that may jointly affect individuals' migration destinations and correlated tastes within their networks cannot fully explain the results; nevertheless, they cannot be completely ruled out. In this sense, our estimate should be interpreted as an upper bound of the effect.

To assess clustering by U.S. destination, we compute Herfindahl–Hirschman Indices (HHIs) for each network type. Across neighborhoods, two randomly selected migrants share the same U.S. consulate 28% of the time (27% for coworkers and 25% for friends), partly reflecting that over 40% of Costa Rican migrants reside in the New York consulate. IExcluding the New York consulate reduces this share to roughly 19% across neighborhoods, coworkers, and friends (18.8%, 18.6%, and 18.7%, respectively). Reassuringly, if we re-sample only among networks linked via relatives

[55]We note that, even if the effect we document were mediated by an online platform or marketplace, the mechanism would still reflect an information externality. Suppose individuals A and B live in the same neighborhood, and A purchases a product on an online platform. The seller then targets B with advertising for that product, and B subsequently purchases it. When A bought the product, she did not internalize that her purchase would provide information to the seller about the preferences of nearby consumers. In the absence of A's purchase, the seller would not have targeted B, and B would not have learned about that product. In this sense, the purchase by A generates an information externality that affects B, even if the transmission occurs through the platform's targeting algorithm. The same logic applies if targeting occurs, for example, via social media.

[56]The logic behind the fixed-effect is that, for example, a person might move to NYC because she likes fashion (HS-2), but is unlikely to move because she likes female trousers made of wool.

to consulates *other than New York*, the point estimates of our 2SLS specification remain statistically equal to our baseline estimates and almost identical in magnitude, as shown in Table H.4.

Finally, Table C.2 and Figure C.3 present results using alternative constructions of the exposure measure. Rather than defining network exposure as an average (equation equation (4)), we consider alternative functional forms, including maximum exposure by product and period. Estimates remain statistically indistinguishable across these specifications.

4.5.1 A Remark on Observed Networks

It is not possible to observe all the connections that each Costa Rican has with the U.S. We cannot observe Costa Ricans who are not registered at a U.S. consulate, but beyond this, people might know U.S. residents and communicate with them, even if they are not relatives. This challenge is pervasive in the networks literature (Goldsmith-Pinkham and Imbens, 2013; Manski, 1993). Nonetheless, we believe we have enough evidence to show that this is not a major concern. First, note that as long as the measured shares s_{bc} are proportional to the true shares, this concern should *not* impact our first stage results, so our first finding on the propagation across international migrant networks would remain unchanged. To see why, note that s_{bc} in equation (4) would remain unchanged as long as the overall shares by consulate remain the same; for instance, if people have unobserved connections in the same U.S. cities where they have relatives. Another (arguably less likely) possibility lies at the other extreme: the observed migrant connections in the U.S. are located in cities whose shocks are orthogonal to those affecting the unobserved connections. In this case, the exposure measure constructed from the observed network would be an imperfect proxy for the true exposure. Such measurement error does not bias the 2SLS estimator itself, but it attenuates the first-stage relationship and reduces the effective strength of the instrument. Hence, if anything, this type of measurement error would bias the analysis toward finding weaker first-stage relationships and noisier second-stage estimates rather than mechanically generating the patterns we document. Second, while missed connections could affect the interpretation of our second-stage estimates, the distance-3 nodes instrument includes own-*network*$\times$*HS-4 product code*$\times$*time* fixed effects, making it largely immune to this concern. Reassuringly, this specification delivers estimates that are statistically indistinguishable from those in our main specification, suggesting that any resulting bias is likely small.[57] Third, as will become clear, our results on retailers depend on (i) our instrument as constructed in the first stage and (ii) the overall *ranking* of product-specific propagation in the second stage. Thus, any bias that rescales the 2SLS coefficients but preserves the propagation ranking would be inconsequential for the results on retailers.

4.6 Heterogeneity in Demand Propagation

We now explore the determinants of how strongly a product propagates within a network after it is imported, and assess whether the evidence is consistent with an information-diffusion mechanism. To do so, we estimate specifications with instrumented interaction terms. Accordingly, in tables with multiple endogenous regressors (e.g., a main regressor and its interaction), we report the conditional first-stage F-statistics of Sanderson and Windmeijer (2016) for each endogenous regressor. We also report the corresponding Stock–Yogo weak-identification critical values (7.03 and 4.58 for 10% and 15% maximal IV size, respectively).[58]

Dynamic Product Categories We start by comparing dynamic product categories with established ones, as one would expect the information channel to be particularly relevant for categories with more dynamism. We rely on Business Dynamics Statistics (BDS) data, which tracks dynamics on establishments with paid employees. These measures are available for the entire U.S. economy, and by industrial sector, 4-digit NAICS, state, and MSA.[59] Specifically, we use data on the creation of jobs by new establishments and on the entry of new establishments by product category to classify a product as "dynamic" ("established") if its creation of jobs by new establishments and entry of new establishments is above (below) the median within

[57]Further, recall that our first stage includes network-time fixed-effects, which would prevent biases which are constant across products. Moreover, we can relax this condition further, as results hold controlling for *network*$\times$*HS-2 product code*$\times$*time* fixed-effects, which would prevent biases that are constant *within* product categories.

[58]These values correspond to the case with two endogenous regressors and two excluded instruments, which is the relevant setting for our interaction specifications.

[59]The BDS is created from the Longitudinal Business Database (LBD), a confidential database used by qualified researchers via secure Federal Statistical Research Data Centers.

our sample (2015-2019).[60] Table I.1 shows our results, which are consistent regardless of the definition of dynamic products. We document a stronger propagation of products in more dynamic categories, as shown by the positive coefficients in the (instrumented) interaction terms. This result aligns with demand externalities aiding in relaxing information frictions, which might be larger in more dynamic categories.

Further, our instrument depends on residuals and therefore has mean zero by construction. In line with this result on dynamic categories, we can then explore if the impact of positive and negative changes of this residual is symmetric.[61] Indeed, Table I.2 documents that positive values of exposure have an impact up to ten times larger than negative values (in the opposite direction) in our first stage. The latter also aligns with results from Table C.2, in which large positive values of exposure appeared as drivers of the effect.

Importer Characteristics Products might propagate more if they are initially imported by someone more connected to others. To explore this, we create a measure of degree centrality, which depends on how many friends a person has using our app-based definition of friendship.[62] We then consider how diffusion in a network depends on the average centrality of its members *with relatives abroad*. Results in Table I.3 suggest that the more central the importers in the first stage, the stronger the diffusion across the neighborhood in the second stage. While these results are indicative, note that the interaction term is noisy; this aligns with recent findings from Akbarpour et al. (2023), which document that the choice of optimal seeds can have limited impact on diffusion within a social network. Furthermore, an importer's demographics also impact the strength of propagation. Appendix I.3 reports how propagation is stronger if the initial importer is high income or female.

[60]The two variables constructed are: (i) employment gains from new establishments, which equals the share of jobs created by new establishments to total employment in the product category; and (ii) entry of establishments, which is the share of new establishments over total establishments in a product category. We then define the variable $Dynamic_p$ used in columns (1)-(3)—definition (i)—and columns (4)-(6)—definition (ii)—of Table I.1.

[61]For instance, intuitively, people might be more likely to transmit information about novelties than about products for which individuals lose interest; in such a case, one would expect the impact of positive changes in exposure to be stronger than the one of negative changes.

[62]Degree centrality is one of the simplest centrality measures; a node's degree is a count of its friend connections, and the degree centrality for a node is just its degree. For instance, a node with 4 friends would have a degree centrality of 4.

Goods' Visibility Intuitively, demand shocks for more visible goods should propagate more easily; either because others are more likely to inquire about a visible good or gather information on it, or because conspicuous consumption may amplify the propagation of demand shocks for such goods. In any case, the more visible a product is, the more likely it is to propagate (and vice versa for non-visible goods). To test for this force, we rely on the product-specific visibility index developed by Charles et al. (2009). We then construct an indicator variable equal to one if the good is *below* the median visibility score. As shown in column (2) of Table I.3, there is a dramatic difference in diffusion between goods depending on their visibility; for non-visible goods, the instrumented variable has a 32 p.p. weaker effect on the probability of importing compared to visible goods.

Premium Products We finally explore if results are heterogeneous between types of goods depending on whether they are *premium* or not. Intuitively, information on goods that are more expensive may be more valuable or trigger conspicuous consumption more strongly. We define premium goods as those whose average price per kilogram is above the median of their HS-4 product category.[63] We focus on networks of neighbors and later also explore if retailers are also more responsive to these more expensive varieties. Table I.4 shows the results of interacting a product-specific premium dummy with our exposure measure. Column (1) shows that the direct externality is about six times as large for premium goods. The latter aligns, for instance, with information spreading more for products which are more expensive (and therefore riskier to import ex-ante).

5 Learning from Consumers: Retailers' Imports

So far, we documented direct demand externalities. First, we showed how expenditure shocks propagate across international migrant networks. We then documented how after an individual imports a product, others in her local network become more likely to import it. We now explore an indirect externality. We study if, once individuals in a network decide to import a product, there might be useful information about the local demand for this product which becomes available to domestic retail firms. For instance, retailers might be more eager to start importing a product the more locals

[63]Results are robust to different cutoffs as long as they are over the 50th percentile.

are willing to acquire it, i.e., the stronger the observed propagation after an individual import. In contrast, retailers might be more cautious about importing products that show weak demand propagation among consumers (Panel C of Figure 1).

Causally documenting this learning from customers is not simple; for instance, it could arise from common shocks or correlated preferences. This challenge has led to limited evidence on this topic in the firm dynamics literature. To test these forces, we leverage both the first *and* second stage results of the previous section. Namely, we begin by considering the following regression, which leverages the instrument based on U.S. connections:

$$\text{Import}_{f,bpt} = \beta_3 \widehat{\text{ShareImporters}}^{US,direct}_{fp,t-2} + \gamma_f + \gamma_{b\tilde{p}} + \gamma_{bt} + \varepsilon_{f,bpt}, \tag{10}$$

where $\text{Import}_{f,bpt} = 1$ if retail firm f in neighborhood b imports product p at time t *for the first time*, γ_f are retailer fixed effects, and other right-hand-side variables are defined as in equation (5).[64] Ideally, the variable $\widehat{\text{ShareImporters}}^{US,direct}_{fpt}$ would depend on retailer f's catchment area—i.e., its area of influence where its customers are—as retailers' decisions are likely to be influenced by their clients. Therefore, we construct retailer-specific catchment areas, that we refer to as *retailer gravity zones* to then create a weighted-average of the exposure faced by *each retailer's own customers*. To do so, we use information on the residence of each retailer's *customers*, which is available for a majority of retailers from data on electronic vouchers. Moreover, we propose a method to approximate these gravity zones for *all* retailers based on employees' residences, which can be applied in other contexts where the customer-specific location data is unavailable and which delivers exposure measures correlated almost perfectly (0.98) with those leveraging customers' location. Appendix J.1 includes details on the gravity zones and exposure construction.

Then, to study whether retailers' responses depend on a product's propensity to propagate across individuals, we classify products according to individuals' responses following an exogenous import shock. To do so, we exploit an exogenous proxy for the extent to which goods spread across individuals. Specifically, as shown in Section

[64]Our panel on imports begins in 2005, while our regression sample runs from 2015 to 2019. We consider an import as a "first-time" import if the retailer has not imported the HS-10 product since 2005. Relaxing this strict constraint to be the first import within the 2014-2019 period does not substantially change the results.

4.6, less visible products propagate less. We therefore interact our main variable of interest with the indicator $LowVisibility_p$ to shed light on the mechanism underlying retailers' responses.

5.1 Results on Retailers and Mechanism

Table 6 reports the results of the analysis. First, in column (1), we find that retailers respond to increases in local adoption of a product. A one standard deviation increase in the instrumented share of individuals with relatives abroad who import product p raises the probability that a local retailer imports that same HS-10 product for the first time by 9% of the mean.[65] Thus, we document that retailers respond to the observed local demand for foreign goods by importing those items.[66]

We now turn to the mechanism underlying this result by testing tests whether products that are more likely to propagate across individuals are also more likely to elicit a supply response from retailers, measured by a higher probability of importing those products.[67] We also ask whether retailers become *less likely* to import products with low propagation. Such an asymmetric response would be consistent with retailers learning about the *level of local demand* for a particular product, rather than merely responding to a product-discovery channel. If a foreign product propagates strongly, individuals in the firm's catchment area may inquire about it at local stores, prompting retailers to stock it in response to perceived demand. Conversely, if individuals learn that a product is poorly matched to local tastes or unpopular within their network, this information may also reach retailers and reduce their propensity to import it compared to retailers in other locations that did not receive this insight. In particular, column (2) of Table 6, which in line with the results of Section 4.6 uses low-visibility goods as a proxy for low-propagation goods, supports this interpretation: the effect for less visible goods is small and negative, whereas the effect for more visible goods is positive and economically large. Consistent with retailers subsequently serving this demand locally, local projections show that individual imports

[65] The retailer estimates are identified from within-HS-10-product variation in local consumer adoption and therefore do not rely on overlap between the aggregate import baskets of households and retailers.

[66] Results in columns (1), (4), and (5) are robust to alternatively consider total imports by individuals as the endogenous variable, as reported in Table J.7.

[67] Table J.8 shows that this result is statistically indistinguishable when retailers are allowed to import from any origin country, rather than only from the U.S.

Table 6: Supply Response from Retailers

Dependent variable: Prob. of retailer f importing product p at time t

	%Δ w.r.t. mean import probability			
	All Retailers		Small Retailers	Large Retailers
	(1)	(2)	(3)	(4)
$\widehat{\text{ShareImporters}}^{US\,direct}_{fp,t-2}$	9.045	12.366	9.777	6.266
	(0.722)***	(0.819)***	(0.819)***	(1.436)***
$LowVisibility_p \times \widehat{\text{ShareImporters}}^{US\,direct}_{fp,t-2}$		-14.804		
			(1.712)***	
F-stat first stage	1035.5	1035.5	944.1	519.6
SW F – interaction		1020.6		
SW F		1154.2		
Stock-Yogo 10% critical value		7.03		
Stock-Yogo 15% critical value		4.58		
Observations	97,500,332	97,500,332	92,579,877	4,917,311
Clusters	2,187,504	2,187,504	2,113,143	115,253
Mean dependent variable	0.03	0.02	0.02	0.09
$b\tilde{p}$, bt, f FE	Yes	Yes	Yes	Yes

Notes: Robust standard errors, adjusted for clustering by retailer-product (HS-4), are in parentheses. Coefficients are percentages of mean outcome for a one standard deviation change in the instrumented share. Regressions control for neighborhood×product (HS-4), neighborhood×time, and retailer fixed-effects. Percentage mean import probabilities are reported. The value of the Sanderson and Windmeijer (2016) conditional first-stage F-statistics (SW F) for the validity of the instruments is reported in columns (2) and (3), along with the Stock-Yogo critical values for a perfectly identified model with two endogenous variables. Appendix E.5 presents details on the sample used in each regression. Data is quarterly and spans 2015-2019.

decline once retailers start selling the product domestically (Appendix E.6).

Appendix J.3 provides an additional test in support of firms learning about the level of local demand by leveraging the imperfect overlap between employer-employee networks and the residential location of employees. The idea behind this exercise is that employees can be exposed to foreign products in their neighborhoods and transmit information about the existence of these products to their employers, which would be relevant under a product-discovery story. However, if employees reside in areas far away from the retailer's catchment area, outside its gravity zone, their insights should be less informative about the particular level of the local demand that their employer would face. Indeed, we find that retail firms do not respond to the exposure faced by employees who live far away from their catchment area, highlighting the relevance of local demand knowledge.

We further investigate the mechanism behind the result in column (1), by exploring which retailers are more likely to respond. Comparing columns (4) and (5), we find that supply effects are mainly driven by small retailers.[68] This is consistent with the notion that small retailers, due to lower productivity, are less able to pay the search costs associated with identifying new foreign products, which larger retailers can afford. Thus, small retailers are more likely to take advantage of information from consumers when choosing which products to import. In fact, as shown in column (2) of Table I.4, the effect is much larger for premium, more expensive products, for which importing for the first time without a notion of local demand might entail a greater risk. In addition, small retailers have a more direct connection with local consumers and might be more responsive to their requests and needs when choosing which products to source.[69] While Table 6 considers *all* retailers, and therefore relies on a proxy of each gravity zone (see Appendix J.1), Table J.6 presents estimations with the subsample for which customers' locations are available. The results and the narrative are consistent between tables—in fact, most results are statistically equal.[70]

Taken together, our empirical investigation then reveals that retail firms learn from their local consumers about which products to source from abroad. This channel is particularly relevant for small retailers, who take advantage of the revealed local demand to identify the "preferred" varieties that align with local customer taste.

5.2 Survey Evidence on the Mechanism

We conduct a nationally representative survey spanning 700 retail firms. Given Costa Rica's size and number of retailers, this is a large-scale survey reaching about 4% of retail firms.[71] Respondents were required to have a relatively deep knowledge about the

[68] We define a retailer as small if its annual employment is below 30, which corresponds to the 95th percentile of the firm-size distribution. Using the median as the cutoff would be too restrictive, since the median retailer in our sample has only two employees; a median-based definition would therefore focus only micro-enterprises.

[69] For instance, if consumers frequent a store inquiring about a product, a small retailer whose manager is at the shop might be more likely to react to these inquiries.

[70] This stability alleviates the concern that the retailer results are driven by labor-residence proximity, neighborhood demographics, or migrant ties embedded in the employee proxy; instead, the findings are robust to measuring catchment areas purely from realized customer origins.

[71] For proportions, this sample size implies a margin of error of approximately ±3.7 percentage points at the 95% confidence level under simple random sampling.

business.[72] The large-scale survey was conducted in partnership with CID-Gallup.[73] All interviewers underwent a one-week training on survey administration, and the instrument was refined through an initial pilot phase. The survey was conducted primarily by telephone during the second and third quarters of 2024; in exceptional cases, face-to-face interviews were carried out when no alternative data collection methods were feasible. Survey responses were linked to administrative firm characteristics, so that we can explore how responses vary across, for instance, firm sizes.

The survey was designed to examine the mechanisms previously documented. Retailers were asked about the role of their customers in determining their choice of which goods to import. We also designed the survey to investigate if results in columns (4) and (5) of Table 6 reflect actual experiences of retailers, and to understand *how* would a small retailer learn about what its customers import and if they liked the item; i.e., the mechanism behind responses in columns (2) and (3). The survey also aimed to confirm that the timing of equation (10) was reasonable, and to acquire a survey-based counterpart of the results based on employees outside the retailer's gravity zone presented in Appendix J.3. The survey instrument and summary statistics on the product assortment by surveyed retailers are reported in Appendix J.5.

Results The survey results closely align retailers' real-world experiences with our empirical findings. Starting with the broader result, we document that 81% of respondents reported receiving customer feedback on product selection, underscoring the influence of consumer preferences on inventory decisions. Additionally, aligned with column (1) of Table 6, 60% of retailers noted that observing potential customers buying new products from abroad would make them more likely to start importing and selling those products locally. This suggests a dynamic local market adaptation to global consumer behaviors informed by direct consumer imports.

Furthermore, consistent with columns (4) and (5) of Table 6 and as reported in panel (a) of Figure 3, the survey documents that small retailers rely more heavily on direct customer imports to gauge the local demand for a potential new product than large firms. Aligned with the empirical analysis, these results indicate a distinct

[72]Respondents were categorized based on their roles within the company, as owner, manager or director, employee involved with sales, or employee in a capacity unrelated to sales.

[73]CID-Gallup has more than 40 years of experience conducting market research and surveys in Latin America. It was established in 1977 in Costa Rica.

strategic approach based on company size, with smaller domestic firms being more responsive to local consumer needs via our mechanism.

Figure 3: Survey Results: Learning from Consumer Imports

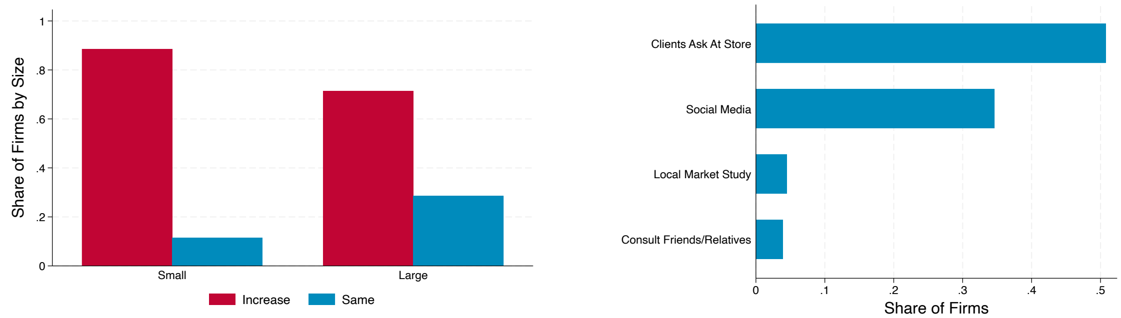

(a) Small vs Large Retailers

(b) Learning Channels

Notes: Panel (a) shows the share of firms which answered "Would increase" to question 2 (*Suppose there is a product that is not available in Costa Rica. If your potential customers start buying the product from abroad over the internet, the likelihood that your company will start importing and selling that product locally)*, by size. In line with Table 6, firms are defined as small if they have 30 employees or less. Also aligned with the table, with includes retailer fixed-effects, we condition on firms who have imported any product in the past. Panel (b) shows the main (number 1) mechanism listed by retailers in Question 4 (i.e., about the *main mechanisms by which the company would become aware that consumers are excited about a new product that is for sale abroad but is not yet available in the country*). In both panels, the sample includes only firms which sell goods to final consumers.

We next leverage the survey to delve into understanding the mechanism behind the results. Columns (2) and (3) of Table 6 indicate that retailers are more responsive to products that propagate widely among individuals, but how do retailers learn about the products consumers are importing and whether they like them? Panel (b) of Figure 3 illustrates the responses. Around half of the retailers gain insights into which imported products interest their customers when people physically visit the store and ask about the availability of these products. This primary channel for gathering information is followed by social media, local market studies, and consultations with family and friends. Also related to the mechanism, and in line with the results in Appendix J.3, retailers are over five times more likely to consult employees who reside near the store than employees who reside further away. In summary, this survey evidence supports our empirical analysis and reinforces the narrative that retail firms, especially smaller ones, rely on local consumers' importing experiences to inform their decisions on which products to source. The key channel to gain these insights is direct customer inquiries made in their physical stores.

5.3 Total Effect and Distributional Implications

We now bring together the effects we documented, both in terms of direct demand externalities across individuals and indirect demand effects from customers to retailers, to understand their impact on the local demand for new varieties. To do so, let's recall that, in Section 4.4, we estimated the impact of an increase in exposure from the U.S. on individual imports via direct externalities. Similarly, the back-of-the-envelope approximation of the increase in retailers' imports of a particular good is given by the product of:

$$C = \underbrace{\Delta Prob(Import_f)}_{\beta_3 \Delta ShareImporters^{US,direct}} \times Retailers \times AvgValue^{Retailer}, \tag{11}$$

where *Retailers* is the total number of retailers and $AvgValue^{Retailer}$ is the average value of a retailer's shipment (\$4,986). Then, equation (11) along with equation (7) imply that local retailers' imports of a new product would increase in \$191 for every dollar exogenously imported by an individual directly connected with the U.S. (i.e., C/A). Appendix J.4 benchmarks this implied gain against findings from the advertising literature.

What is the total multiplier given the demand propagation from the U.S., to individuals with relatives, to others in their network, and finally to retailers? To illustrate, suppose $\Delta\widetilde{E}_{bpt}$ corresponds to a \$100 increase in U.S. per capita spending on a product. Then, Costa Rican imports of this product per individual with a relative abroad would increase in approximately \$6, given the degree of interconnectedness between the two countries and the strength of the direct and indirect demand propagation.[74]

This is a total sizable effect, which can be decomposed into additional *new* imports of consumer products due to the direct externality and individuals' responses—7% of the effect—and the indirect externality and imports of domestic retail firms—the remaining 93% of the effect. These magnitudes underscore how accounting for this new supply-side indirect externality is key when estimating a full response.[75] In line

[74]Like in Section 4.4, this multiplier is in per capita terms, where the denominator in Costa Rica is the number of individuals with a relative abroad, as these individuals act as the bridge through which demand propagates. A more connected set of countries (with more directly exposed individuals) would face an overall stronger propagation.

[75]While restricting attention to first-time imports by product category strengthens identi-

with this sizable effect, Appendix G uses cross-country data in a gravity framework to show that individual import flows are particularly responsive to measures of social connectedness, suggesting that similar forces likely operate beyond our setting.

Finally, the documented externalities have distributional implications. First, the diffusion channel may enable lower-income families to benefit indirectly from migration. Those with relatives abroad have, on average, 12% higher incomes, and consequently, are more likely to import in response to foreign expenditure shocks.[76] Second, those who import directly have, on average, 49% higher income than those who do not. As retailers learn from consumer imports and subsequently introduce products, the indirect externality leads to new product varieties available locally for lower-income households. These forces indicate that the diffusion of new products via retailers can lead to more evenly distributed variety gains across income groups, provided that the newly introduced varieties do not crowd out those disproportionately consumed by low-income households.

6 Concluding Remarks

This paper investigates the role of direct and indirect externalities in propagating demand shocks—both across individuals and from individuals to retailers. Our analysis makes three key contributions. First, we demonstrate how product-specific expenditure shocks propagate through international migrant networks, suggesting that migration policies can influence global product diffusion. Second, we leverage that these foreign expenditure shocks generate local demand shocks in Costa Rica through migrant networks to quantify how an individual's first import increases the probability that others in her network import the same product—a force that varies across network types and product characteristics. Third, and most importantly, we use this individual-level propagation as a building block to show that retailers learn about local demand for foreign products by observing consumer imports: when imports signal strong local demand, retailers respond by importing the product themselves.

A decomposition of these effects reveals that while individual responses (direct

fication, it may lead us to underestimate the size of demand spillovers. Conversely, retailers may initially overstock newly introduced products, which would overstate the true increase in underlying consumer demand within a given period.

[76] Incomes are calculated based on wage income from social security records.

externalities) matter, retailer responses (indirect externalities) account for 92% of the overall impact, underscoring their central role in expanding access to product varieties. Notably, although direct imports are concentrated among higher-income households, the diffusion of product varieties via retail channels allows these varieties to reach lower-income families. Whether and how retailers substitute incumbent varieties with new ones—and the resulting welfare effects across income groups—represents a promising avenue for future research.

The presence of these externalities implies that gains from trade may be larger than previously documented, generating a multiplier effect for policies that stimulate foreign product demand, such as lower tariffs or relaxed import requirements. This insight is especially relevant for developing markets and is central to debates over tariff exemptions for individual imports. For example, in the U.S., the White House announced the suspension of the "de minimis" provision of Section 321 of the Tariff Act of 1930—which waives tariffs for shipments valued under $800 in 2025. Because information created by early individual importers spills over to peers and can induce retailer adoption, these policies can have multiplier effects beyond the set of direct importers. Conversely, lowering the fixed costs of individual importing—especially in dynamic product categories with high new-variety turnover—can facilitate variety discovery and accelerate diffusion. Since retailer adoption disproportionately expands access to these new varieties among lower-income consumers, these frictions may have important distributional implications.

Finally, this paper is the first to study the determinants of individual imports. Although historically uncommon, the rapid expansion of the direct-to-consumer market is only expected to accelerate due to increased internet penetration, improved logistics, and globalization. For instance, Temu, a Chinese app that allows for foreign direct-to-consumer purchases, was Apple's most downloaded free app in the U.S. for 2023, and low-value imports represented about 15% of all imports from China in 2021 according to the U.S. Customs Border Protection.[77] Thus, this paper serves as an initial contribution to what promises to be fertile ground for future research.

[77] Note that, while large platforms can reduce search costs (e.g., via reviews), they also present an overwhelming menu of options. Personal networks across and within borders provide trusted, high-signal filters and thus remain relevant despite platform aggregation. Moreover, social media and communication platforms have recently lowered the cost of information flows across borders, further reinforcing the mechanisms we study.

References

Acosta, M., Cox, L., 2019. The Regressive Nature of the U.S. Tariff Code: Origins and Implications. Technical Report. Working Paper, Columbia University.

Acosta, P., Calderon, C., Fajnzylber, P., Lopez, H., 2008. What is the impact of international remittances on poverty and inequality in Latin America? World Development 36, 89–114.

Agarwal, A., Singh, R., Toshniwal, D., 2018. Geospatial sentiment analysis using twitter data for UK-EU referendum. Journal of Information and Optimization Sciences 39, 303–317.

Akbarpour, M., Malladi, S., Saberi, A., 2023. Just a Few Seeds More: The Inflated Value of Network Data for Diffusion. Working Paper.

Allen, T., 2014. Information Frictions in Trade. Econometrica 82, 2041–2083.

Alvarez, F., Argente, D., Lippi, F., Méndez, E., Van Patten, D., 2023. Strategic Complementarities in a Dynamic Model of Technology Adoption: P2P Digital Payments.

Aronow, P.M., Samii, C., 2017. Estimating average causal effects under general interference, with application to a social network experiment .

Atkin, D., Faber, B., Gonzalez-Navarro, M., 2018. Retail Globalization and Household Welfare: Evidence from Mexico. Journal of Political Economy 126, 1–73.

Bai, J., Chen, M., Liu, J., Mu, X., Xu, D.Y., 2020. Search and information frictions on global e-commerce platforms: Evidence from Aliexpress. Technical Report. National Bureau of Economic Research.

Bailey, M., Johnston, D.M., Kuchler, T., Stroebel, J., Wong, A., 2022. Peer Effects in Product Adoption. American Economic Journal: Applied Economics .

Bandiera, O., Barankay, I., Rasul, I., 2009. Social Connections and Incentives in the Workplace: Evidence From Personnel Data. Econometrica 77, 1047–1094.

Bandiera, O., Rasul, I., 2006. Social Networks and Technology Adoption in Northern Mozambique. Economic Journal 116, 869–902.

Batch, A., Bridgman, B., Dunn, A., Gholizadeh, M., 2024. Consumption zones. Journal of Economic Geography , lbae035URL: https://doi.org/10.1093/jeg/lbae035, doi:10.1093/jeg/lbae035.

Beaman, L., BenYishay, A., Magruder, J., Mobarak, A.M., 2021. Can network theory-based targeting increase technology adoption? American Economic Review 111, 1918–1943. doi:10.1257/aer.20190772.

Beine, M., Docquier, F., Rapoport, H., 2008. Brain Drain and Human Capital Formation in Developing Countries: Winners and Losers. The Economic Journal 118, 631–652.

Bertrand, M., Luttmer, E.F., Mullainathan, S., 2000. Network Effects and Welfare Cultures. The Quarterly Journal of Economics 115, 1019–1055.

Borusyak, K., Hull, P., 2020. Non-Random Exposure to Exogenous Shocks: Theory and Applications. Working Paper 27845. National Bureau of Economic Research. URL: http://www.nber.org/papers/w27845, doi:10.3386/w27845.

Borusyak, K., Hull, P., 2023. Nonrandom Exposure to Exogenous Shocks. Econometrica 91, 2155–2185.

Borusyak, K., Jaravel, X., 2021. The Distributional Effects of Trade: Theory and Evidence from the United States. Technical Report. National Bureau of Economic Research.

Bramoullé, Y., Djebbari, H., Fortin, B., 2020. Peer effects in networks: A survey. Annual Review of Economics 12, 603–629.

Brock, W.A., Durlauf, S.N., 2001. Discrete Choice with Social Interactions. The Review of Economic Studies 68, 235–260.

Broda, C., Weinstein, D.E., 2006. Globalization and the Gains from Variety. The Quarterly Journal of Economics 121, 541–585.

Bronnenberg, B.J., Ellickson, P.B., 2015. Adolescence and the path to maturity in global retail. Journal of Economic Perspectives 29, 113–134. URL: https://www.aeaweb.org/articles?id=10.1257/jep.29.4.113, doi:10.1257/jep.29.4.113.

Caplin, A., Dean, M., 2015. Revealed preference, rational inattention, and costly information acquisition. American Economic Review 105, 2183–2203.

Carter, M.R., Laajaj, R., Yang, D., 2021. Subsidies and the african green revolution: Direct effects and social network spillovers of randomized input subsidies in mozambique. American Economic Journal: Applied Economics 13, 206–231. doi:10.1257/app.20190396.

Chaney, T., 2014. The Network Structure of International Trade. American Economic Review 104, 3600–3634.

Charles, K.K., Hurst, E., Roussanov, N., 2009. Conspicuous Consumption and Race. The Quarterly Journal of Economics 124, 425–467.

Conley, T.G., Udry, C.R., 2010. Learning about a New Technology: Pineapple in Ghana. American Economic Review 100, 35–69.

Conte, M., Cotterlaz, P., Mayer, T., 2022. The CEPII Gravity Database .

Couture, V., Faber, B., Gu, Y., Liu, L., 2021. Connecting the Countryside via E-Commerce: Evidence from China. American Economic Review: Insights 3, 35–50.

De Giorgi, G., Frederiksen, A., Pistaferri, L., 2019. Consumption Network Effects. The Review of Economic Studies 87, 130–163. doi:10.1093/restud/rdz026.

De Giorgi, G., Pellizzari, M., Redaelli, S., 2010. Identification of Social Interactions through Partially Overlapping Peer Groups. American Economic Journal: Applied Economics 2, 241–75.

Dolfen, P., Einav, L., Klenow, P.J., Klopack, B., Levin, J.D., Levin, L., Best, W., 2023. Assessing the gains from e-commerce. American Economic Journal: Macroeconomics 15, 342–370.

Duch-Brown, N., Grzybowski, L., Romahn, A., Verboven, F., 2021. Are online markets more integrated than traditional markets? evidence from consumer electronics. Journal of International Economics 131, 103476.

Duesenberry, J.S., 1948. The consumption function: A study of relations between income and consumer expenditures. University of Michigan.

Duflo, E., Saez, E., 2003. The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment. The Quarterly Journal of Economics 118, 815–842.

EBANX, B.B..., 2023. Digital payments connecting businesses and people in rising economies: an overview of online commerce in Latin America and Africa, 2023. Online report. Accessed: 2023-09-13.

Faber, B., Fally, T., 2022. Firm Heterogeneity in Consumption Baskets: Evidence from Home and Store Scanner Data. The Review of Economic Studies 89, 1420–1459.

Fajgelbaum, P.D., Khandelwal, A., 2024. The Value of De Minimis Imports. Technical Report. National Bureau of Economic Research.

Felbermayr, G., Grossmann, V., Kohler, W., 2015. Migration, international trade, and capital formation: Cause or effect?, in: Chiswick, B.R., Miller, P.W. (Eds.), Handbook of the Economics of International Migration. Elsevier. volume 1B. chapter 18, pp. 913–1025. doi:10.1016/B978-0-444-53768-3.00018-7.

Fernandes, A.P., Tang, H., 2014. Learning to export from neighbors. Journal of International Economics 94, 67–84. URL: https://www.sciencedirect.com/science/article/pii/S0022199614000865, doi:https://doi.org/10.1016/j.jinteco.2014.06.003.

Fildes, R., Ma, S., Kolassa, S., 2022. Retail forecasting: Research and practice. International Journal of Forecasting 38, 1283–1318.

Furman, J., Russ, K., Shambaugh, J., 2017. US tariffs are an arbitrary and regressive tax. VoxEU blog .

Goldsmith-Pinkham, P., Imbens, G.W., 2013. Social Networks and the Identification of Peer Effects. Journal of Business & Economic Statistics 31, 253–264.

Gorodnichenko, Y., Talavera, O., 2017. Price setting in online markets: Basic facts, international comparisons, and cross-border integration. American Economic Review 107, 249–282.

Head, K., Mayer, T., 2014. Gravity Equations: Workhorse,Toolkit, and Cookbook, in: Handbook of international economics. Elsevier. volume 4, pp. 131–195.

Hortaçsu, A., Syverson, C., 2015. The ongoing evolution of US retail: A format tug-of-war. Journal of Economic Perspectives 29, 89–112. URL: https://www.aeaweb.org/articles?id=10.1257/jep.29.4.89, doi:10.1257/jep.29.4.89.

Hottman, C.J., Monarch, R., 2020. A matter of taste: Estimating import price inflation across U.S. income groups. Journal of International Economics 127, 103382.

Jordà, , 2005. Estimation and Inference of Impulse Responses by Local Projections. American Economic Review 95, 161–182.

Juhász, R., Steinwender, C., 2018. Spinning the web: Codifiability, information frictions and trade. NBER working paper 18.

Kuhn, K., Galloway, T., Collins-Williams, M., 2016. Near, far, and online: Small business owners' advice-seeking from peers. Journal of Small Business and Enterprise Development .

Manski, C.F., 1993. Identification of Endogenous Social Effects: The Reflection Problem. The Review of Economic Studies 60, 531–542.

Mas, A., Moretti, E., 2009. Peers at Work. American Economic Review 99, 112–45.

Matějka, F., McKay, A., 2015. Rational inattention to discrete choices: A new foundation for the multinomial logit model. American Economic Review 105, 272–298.

Maurer, J., Meier, A., 2008. Smooth it Like the 'Joneses'? Estimating Peer-Group Effects in Intertemporal Consumption Choice. The Economic Journal 118, 454–476.

McCully, B.A., Jaccard, T., Albert, C., 2024. Immigrants, imports, and welfare: Evidence from household purchase data. Technical Report. RF Berlin-CReAM Discussion Paper Series.

Méndez, E., Van Patten, D., 2022. Voting on a Trade Agreement: Firm Networks and Attitudes Toward Openness. Technical Report. National Bureau of Economic Research.

Moffitt, R.A., 2000. Welfare Benefits and Female Headship in U.S. Time Series. American Economic Review 90, 373–377.

Peres, R., Muller, E., Mahajan, V., 2010. Innovation diffusion and new product growth models: A critical review and research directions. International Journal of Research in Marketing 27, 91–106. doi:10.1016/j.ijresmar.2009.12.012.

Peri, G., Requena-Silvente, F., 2010. The trade creation effect of immigrants: Evidence from the remarkable case of spain. Canadian Journal of Economics 43, 1433–1459. doi:10.1111/j.1540-5982.2010.01620.x.

Reardon, T., Timmer, C.P., Barrett, C.B., Berdegué, J., 2003. The Rise of Supermarkets in Africa, Asia, and Latin America. American Journal of Agricultural Economics 85, 1140–1146.

Sanderson, E., Windmeijer, F., 2016. A weak instrument F-test in linear IV models with multiple endogenous variables. Journal of Econometrics 190, 212–221.

Sims, C.A., 2003. Implications of rational inattention. Journal of monetary Economics 50, 665–690.

Startz, M., 2016. The value of face-to-face: Search and contracting problems in Nigerian trade. Available at SSRN 3096685 .

Steinwender, C., 2018. Real effects of information frictions: When the states and the kingdom became united. American Economic Review 108, 657–696.

Veblen, T., 1899. The Preconceptions of Economic Science. The Quarterly Journal of Economics 13, 121–150.

Wei, S.J., Wei, Z., Xu, J., 2021. On the market failure of “missing pioneers”. Journal of Development Economics 152, 102705. doi:10.1016/j.jdeveco.2021.102705.

Online Appendix for

Cross-Border Product Adoption: Individual Imports, Migrant Networks, and Domestic Retailers

March 5$^{\text{th}}$, 2026

David Argente
Yale University and NBER

Esteban Méndez
Central Bank of Costa Rica

Diana Van Patten
Yale University and NBER

Table of Contents

A Details on New Stylized Facts

Table A.1: Top HS-4 Codes Imported by Individuals

Code	Description	Share
6204	Women's or girls' suits, ensembles, jackets, dresses, skirts, divided skirts, trousers, bib and brace overalls, breeches, and shorts.	10%
8703	Motor cars and other motor vehicles designed to transport people (other than those of heading 8702, including station wagons and racing cars).	5%
6206	Women's or girls' blouses, shirts, and shirt-blouses.	2%
4202	Trunks, suitcases, vanity cases, executive cases, briefcases, school satchels, and similar containers; traveling bags, backpacks, handbags, and similar products.	2%
8708	Parts and accessories of the motor vehicles of headings 8701 to 8705.	2%
6205	Men's or boys' shirts.	2%
6110	Sweaters, pullovers, sweatshirts, waistcoats (vests), and similar articles, knitted or crocheted.	2%
3926	Articles of plastics and articles of other materials of headings 3901 to 3914 (includes plastic parts or accessories, but can vary greatly).	2%
9503	Tricycles, scooters, pedal cars, and similar wheeled toys; dolls' carriages; dolls and other toys; reduced-size ("scale") models and similar recreational models.	2%
6302	Bed linen, table linen, toilet linen, and kitchen linen.	2%

Notes: The table documents the top HS-4 codes imported by individuals in Costa Rica. This ranking results from collapsing imports from HS-10 to HS-4 categories to be more informative; otherwise most top 10 imports would belong to the 6204 category. Data spans 2015-2019.

Table A.2: Top HS-4 Codes for Imported Final Goods: All and by Retailers

All Final Goods		Retailers' Imports	
(1)	(2)	(3)	(4)
Code	**Share**	**Code**	**Share**
8536	3%	8708*	4%
7318	3%	6204*	3%
3926*	2%	3926*	2%
3923	2%	8536	2%
4016	2%	7318	2%
8544	2%	4202*	2%
7326	1%	8421	2%
8481	1%	4016	1%
8302	1%	8481	1%
8482	1%	6206*	1%

Notes: The table documents the top HS-4 codes for *all* final goods imported in Costa Rica (column (1)) and for final goods imported by retail firms (column (3)), with the respective shares. This ranking results from collapsing imports from HS-10 to HS-4 categories and does not weight by product value, and rather reports the most commonly imported goods. We denote with an asterisk the top codes which coincide with the top categories imported by individuals and reported in Table A.1. Data spans 2015-2019.

Figure A.1: Imports by Individuals

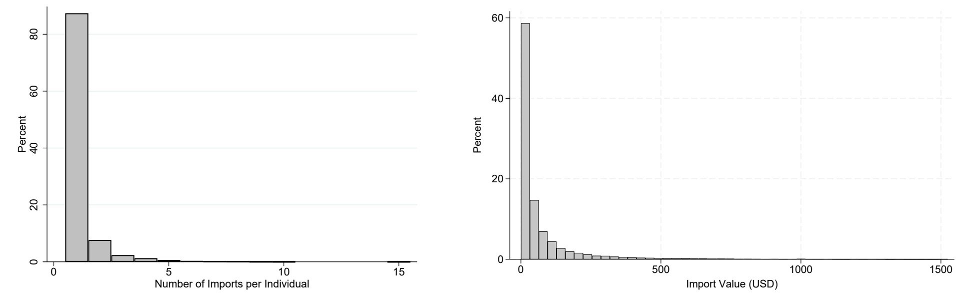

(a) Imports by HS-10 Code (b) Distribution of Import Values

Notes: Panel (a) shows a histogram with the total number of imports by HS-10 code by individual. As shown, most product categories are imported by each person only once. Panel (b) shows a histogram with the distribution of values of individual imports. For visual purposes, the figure omits individual imports of motor vehicles and boats (HS-codes 87-89). The data spans 2015-2019.

B Setup: Additional Results

B.1 CEX vs. Other Expenditures Data

Relationship Between CEX and Imports Data In the U.S., many tradable products are imported. Thus, expenditure shares for these products in the CEX by region should co-move with the imports of these products in these areas. Following this idea, we use data on imports by customs districts in the U.S., adjusted using Freight Analysis Framework (FAF) data from the Department of Transportation as explained in Section 2, to assess the representativeness of the CEX at narrowly-defined categories and geographic areas. This notion follows Acosta and Cox (2019), who show that these customs districts data closely matches aggregate patterns in the CEX. Figure B.1 shows a strong correlation between expenditures in the CEX and expenditures based on customs districts data, when defining products as 4- or 6-digit HS codes, regions as Primary Sampling Units (PSUs), and time as quarters between 2015 and 2019. The correlation is strong both in levels, as shown by Panels (A)-(B), and in changes, as shown in Panels (C)-(D).

Figure B.1: Expenditure Shares in the CEX vs. Customs Districts

(A) Levels: HS-4

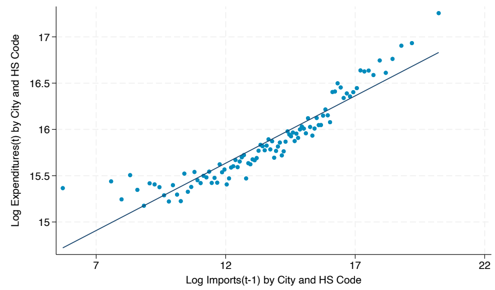

(B) Levels: HS-6

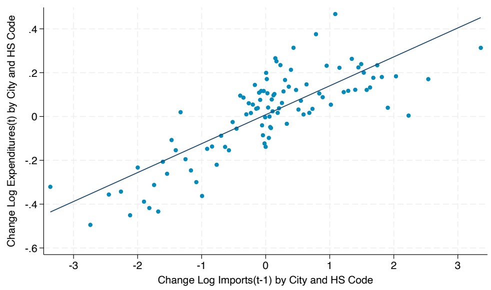

(C) Changes: HS-4

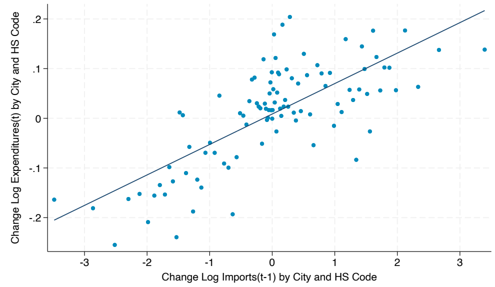

(D) Changes: HS-6

Notes: The figures show the relation between expenditures in the CEX (vertical axis) and expenditure based on customs districts data (horizontal axis), when defining products as HS-4 or HS-6 product codes, regions as PSUs, and time as quarters for the period 2015-2019. Panels (A) and (B) show the correlation in levels for products as HS-4 and HS-6 codes, respectively. Panels (C) and (D) show the correlation for the same definition of products in changes. We trim the top and bottom one percent.

Relationship Between CEX and Debit Card Data We use data on debit card transactions by region and by type to further validate the CEX. This data comes from Facteus, a provider of financial data for business analytics. The data contain information on total expenditures by category at the zip-code level and with daily frequency. Approximately 10 million debit cards are included. The debit cards in the Facteus panel are issued by "challenger banks." The dataset spans from 2017 to 2019 and includes information on more than 200 Merchant Category Codes (MCCs), one per transaction, corresponding to the MCC standard by Visa and Mastercard. We manually create a bridge between MCC and Standard Industrial Classification (SIC) codes.[78] Figure B.2 shows a strong correlation between expenditures in the CEX and expenditures based on card transactions data, when defining products as

[78]This bridge was created in parallel by two independent teams of RAs, then cross-checked, and finally revised by the authors. MCCs were derived from SIC codes; however, MCCs and SIC codes do not always correspond: in some cases, several SIC codes are consolidated into one MCC, while in other cases, such as for "T&E and direct marketing merchants," MCCs do not have corresponding SIC code.

HS-4 or HS-6 codes, regions as PSUs, and time as years. As in the case of customs data, the correlation is strong both in levels and in changes.

Figure B.2: Expenditure Shares in the CEX vs. Card Transactions

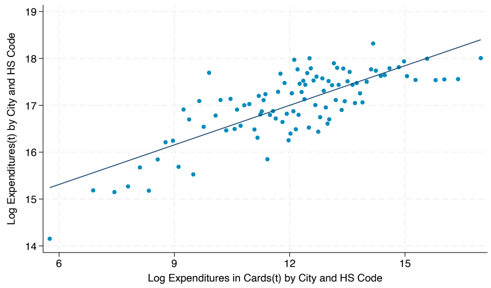

(A) Levels: HS-4

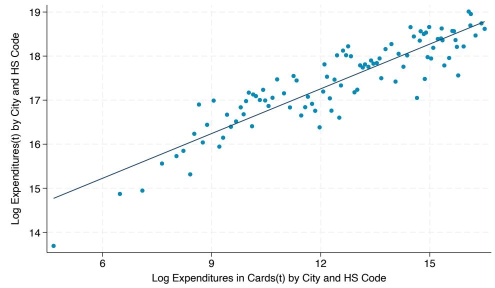

(B) Levels: HS-6

(C) Changes: HS-4

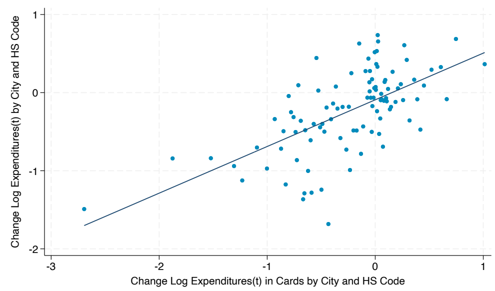

(D) Changes: HS-6

Notes: The figure shows the relation, both in levels and in changes, between expenditures based on the CEX (vertical axis) and on card transactions (horizontal axis), when defining products as HS-4 or HS-6 product codes, regions as PSUs, and time as years; we trim the top and bottom one percent.

B.2 Network Descriptive Statistics

Table B.1: Network Summary Statistics

Network type	Total number of networks (1)	Median individuals per network (2)	25th percentile individuals per network (3)	75th percentile individuals per network (4)
Neighbors	1,681	781	352	2100
Coworkers	11,803	12	4	37
Friends	109,438	7	4	12

Notes: The table shows the number of distinct networks per network type, along with the median, 25th and 75th percentiles of people who compose each network. These are only networks which have at least one person with a relative abroad. See Appendix D for details on why network of friends are so numerous.

C Details on Exposure Measures

C.1 Microfoundation of Exposure Measures

Our specifications rely on two exposure measures. Equation (3) considers exposure at the individual level. We then aggregate exposures by network to use in the 2SLS in equation (4). We now provide a microfoundation for both specifications.

A Learning Microfoundation for Individual Exposure. Let $e_{cpt} \equiv \ln \widetilde{E}_{cpt}$ denote the product-specific expenditure shock in U.S. consulate c at time t. Suppose that each individual i (in Costa Rica) who is connected to consulate c receives a noisy signal about e_{cpt}:

$$z_{i,cpt} \; = \; e_{cpt} \; + \; u_{i,cpt}, \tag{12}$$

where $u_{i,cpt}$ is i.i.d. mean-zero noise. Assume a Gaussian prior and Gaussian signal noise:

$$e_{cpt} \sim \mathcal{N}(0, \sigma_e^2), \qquad u_{i,cpt} \sim \mathcal{N}(0, \sigma_u^2), \qquad e_{cpt} \perp u_{i,cpt}.$$

Under these assumptions, the posterior mean of e_{cpt} given the signal $z_{i,cpt}$ is linear:

$$\mathbb{E}[e_{cpt} \mid z_{i,cpt}] \; = \; \frac{\sigma_e^2}{\sigma_e^2 + \sigma_u^2} \, z_{i,cpt} \; = \; \frac{\sigma_e^2}{\sigma_e^2 + \sigma_u^2} \, (e_{cpt} + u_{i,cpt}) \, . \tag{13}$$

Equation (13) above provides a learning-based interpretation of the individual-level exposure shifter used in equation (3) in the paper. For a directly connected individual i whose relative resides in consulate c, the regressor $e_{cp,t-1}$ can be interpreted as the realized payoff-relevant state about which the individual forms posterior beliefs based on a noisy signal. Under Gaussian priors and additive signal noise, individuals act on the posterior mean $\mathbb{E}[e_{cpt} \mid z_{i,cpt}]$, which is linear in the observed signal. In particular, averaging over mean-zero signal noise implies $\mathbb{E}[\mathbb{E}[e_{cp,t-1} \mid z_{i,cp,t-1}] \mid e_{cp,t-1}] = \frac{\sigma_e^2}{\sigma_e^2+\sigma_u^2} \, e_{cp,t-1}$. As a result, the belief-based payoff shifter that governs individual behavior moves proportionally with the realized shock $e_{cp,t-1}$ in expectation. Consequently, equation (3) is consistent with an underlying learning model in which individual behavior responds to posterior beliefs.

An Information-Theoretic Microfoundation for Network Exposure. To set up our 2SLS, we need to summarize exposure by network to understand if those without relatives abroad react to the exposure of their peers. We now provide a unified microfoundation for network exposure that nests the linear-in-means mapping used in equation (4). The key additional ingredient is the possibility that network-level adoption incentives are driven primarily by the *most salient* source among those reaching the network when information-processing capacity is limited.

Since s_{bc} is the share of individuals in network b connected to consulate c, it can be interpreted as a baseline distribution over sources: it summarizes how frequently information from each c can enter network b. We model a representative agent in network b as allocating attention across consulates via a probability distribution π_{bcpt} satisfying $\pi_{bcpt} \geq 0$ and $\sum_c \pi_{bcpt} = 1$. We assume that concentrating attention away from baseline availability s_{bc} is costly, and measure this cost using the Kullback–Leibler (KL) divergence,

$$D_{\mathrm{KL}}(\pi_{bpt} \,\|\, s_b) \;\equiv\; \sum_c \pi_{bcpt} \log\left(\frac{\pi_{bcpt}}{s_{bc}}\right). \tag{14}$$

KL divergence is a standard nonnegative measure from information theory: it equals zero when $\pi_{bcpt} = s_{bc}$ for all c and increases as attention becomes more concentrated relative to s_{bc}. In rational-inattention models, such relative-entropy costs capture limited information-processing capacity (e.g. Caplin and Dean, 2015; Matějka and McKay, 2015; Sims, 2003). The representative agent chooses attention weights to maximize expected payoff net of the KL cost:

$$\max_{\{\pi_{bcpt}\}_c} \sum_c \pi_{bcpt}\, e_{cpt} \;-\; \tau \sum_c \pi_{bcpt} \log\left(\frac{\pi_{bcpt}}{s_{bc}}\right) \quad \text{s.t.} \quad \pi_{bcpt} \geq 0,\; \sum_c \pi_{bcpt} = 1, \tag{15}$$

where $\tau > 0$ governs the marginal cost of concentrating attention (higher τ implies more diffuse attention and less responsiveness to extreme sources).

The first-order condition implies the optimal attention rule

$$\pi_{bcpt} = \frac{s_{bc} \exp(e_{cpt}/\tau)}{\sum_{c'} s_{bc'} \exp(e_{c'pt}/\tau)}.$$

Substituting the optimal π_{bcpt} back into equation (15) yields the maximized value

$$\tau \log\left(\sum_c s_{bc} \exp(e_{cpt}/\tau)\right),$$

the familiar *log-sum-exp* (smooth-max) operator. Substituting back $e_{cpt} \equiv \ln \widetilde{E}_{cpt}$, this delivers a smooth-max network exposure index,

$$\ln \widetilde{E}_{bpt}^{\text{smooth}} \;\equiv\; \tau \log\left(\sum_c s_{bc} \exp\left(\frac{\ln \widetilde{E}_{cpt}}{\tau}\right)\right). \tag{16}$$

The parameter τ governs the degree of attention concentration: larger τ keeps attention closer to baseline availability s_{bc}, while smaller τ places disproportionate weight on the largest e_{cpt}. Importantly, as $\tau \to \infty$, attention becomes diffuse and equation (16) *converges to the linear-in-means* mapping, providing a microfoundation for equation (4) in the paper. To see this, note that as $\tau \to \infty$,

$$\tau \log\left(\sum_c s_{bc} \exp\Big(\ln \widetilde{E}_{cpt}/\tau\Big)\right) \;=\; \sum_c s_{bc} \ln \widetilde{E}_{cpt} + o(1),$$

which coincides with equation (4). Note also that, as $\tau \to 0$, equation (16) converges to the hard-max exposure

$$\ln \widetilde{E}_{bpt}^{\max} \equiv \max_{c \in \mathcal{C}_b} \ln \widetilde{E}_{cpt} \tag{17}$$

corresponding to the limiting case in which attention concentrates on the single most salient source. Although this hard-max exposure is computationally convenient, it can be statistically fragile. If $\ln \widetilde{E}_{cpt}$ contains measurement error, taking a maximum tends to amplify noise and outliers, since the maximum is disproportionately influenced by extreme realizations. As a result, our preferred specification is the linear-in-means measure, corresponding to the limit $\tau \to \infty$. This linear-in-means exposure mapping is standard in the network-interactions literature (e.g., Manski 1993, Goldsmith-Pinkham and Imbens 2013, Bramoullé et al. 2020), where—as in our microfoundation above—each agent's expected payoff from adopting a new variety is a weighted sum of peers' shocks under a linear best-response or Bayesian-updating framework. To assess robustness, we explore how results change based on the hard-max exposure (corresponding to the limit $\tau \to 0$) as opposed to our baseline exposure (the limit $\tau \to \infty$). As shown in Table C.2, results are statistically equal. Figure C.3 presents a more general comparison while defining networks as neighborhoods; for values of $\tau \to \infty$ (baseline), $\tau \to 0$ (hard-max), $\tau = 0.1$ (smooth-max), $\tau = 1$, and $\tau = 2$; the estimated coefficient remains stable and statistically equal.

Table C.2: Main Results Instrumenting with Maximum Exposure

Dep. variable: Prob. importing product p for individual i without relatives in the U.S. and who belongs to network b at time t

	%Δ w.r.t. mean import probability		
	Neighbors (1)	Coworkers (2)	Friends (3)
$\widehat{\text{ShareImporters}}_{bp,t-1}^{US\,direct}$	16.573	18.954	13.792
	(5.818)***	(10.330)*	(4.860)***
F-stat first stage	22.59	6.80	15.75
Observations	289,340,892	299,920,162	260,952,672
Clusters	200,308	236,804	4,568,240

Notes: The table shows the results of running equation (6) using an alternative instrumental variable. The instrument considers the maximum exposure per network for each product and period. Robust standard errors, adjusted for clustering by network-product (HS-4), are in parentheses. Coefficients are percentages of mean outcome for a one standard deviation change in the instrumented share. Regressions control for network×product (HS-4), network×time, and individual fixed-effects.

Figure C.3: Point Estimate Based on Different Values of τ to Construct Exposure

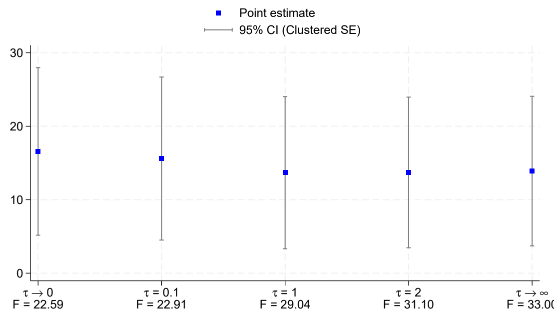

Notes: The figure compares our baseline 2SLS estimate from Table 4 with alternatives based on distinct exposure measures. Networks correspond with neighborhoods. Point estimates are denoted in blue along with 95% confidence intervals. The horizontal axis displays the value of τ associated with each point estimate and the corresponding first-stage F-statistic.

C.2 Variation and Serial Correlation Tests

Table C.3: Variation for Exposure Measures
(% of products with underlying variation at each HS-code level)

Panel (a): CEX				*Panel (b): U.S. imports*			
HS-4	HS-6	HS-8	HS-10	HS-4	HS-6	HS-8	HS-10
91.01	6.35	2.12	0.53	0	**59.95**	21.63	18.42

Notes: The table shows the percentage of products in our sample whose underlying variation is at each HS-code level, which tells us the level at which the exposure measures for our IV strategy are varying, depending on the source from which we obtain expenditures by product code, region and time in the U.S. Panel (a) shows that most of the variation is at the HS-4 level when using the CEX. Panel (b) shows that most of the variation is at the HS-6 level when relying on U.S. imports data by customs districts.

Figure C.4: Alternative Aggregated Exposure vs. Baseline

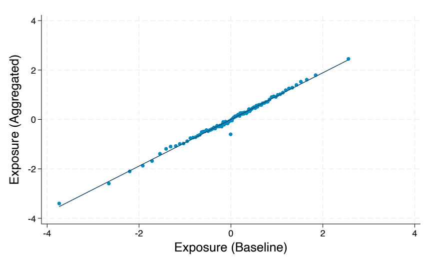

Notes: The figure compares our baseline consulate-level exposure with an alternative measure constructed by aggregating the underlying expenditure data (by product) directly to the consulate level and estimating a two-way saturated fixed-effects specification to obtain the corresponding expenditure shocks.

Table C.4: Residual Variance of Log Expenditures

Specification	Variance	Specification	Variance
No controls	4.424		
Product	2.566	City × Time	3.117
Time	4.414	Product × Time	2.456
City	3.166	City × Product, City × Time	1.095
City × Product	1.144	City × Product, Product × Time	1.033
City × Product, Time	1.132	City × Product, City × Time, Product × Time	0.996

Notes: The table reports the variance of the residuals from regressions of log expenditures on various sets of fixed effects. Each row shows the remaining variance after absorbing the listed fixed effects.

Serial Correlation Tests To assess the presence of serial correlation in $\widetilde{E}_{cpt}$, we conducted the Wooldridge test (Wooldridge, 2002) for autocorrelation in panel data. We find no evidence of serial correlation (i.e., the null hypothesis of no first-order autocorrelation could not be rejected). We also conducted unit root tests to determine the stationarity of $\widetilde{E}_{cpt}$. Specifically, we employed both the Levin-Lin-Chu test (Levin et al., 2002) and the Im-Pesaran-Shin (IPS) test (Im et al., 2003), two tests that are appropriate for panel data with a large number of cross-sections and a smaller number of time periods. The LLC test assumes a common unit root process, while the IPS test allows for individual unit root processes across panels. After allowing for several lags to account for possible autocorrelation and dynamic effects within the data, we corroborate the stationarity of $\widetilde{E}_{cpt}$. Table C.5 presents the results.

Table C.5: Serial Correlation Tests for $\widetilde{E}_{cpt}$

	Wooldridge Test (1)	Levin-Lin-Chu Test (2)	Im-Pesaran-Shin Test (3)
H_0	No serial autocorrelation	Panels contain unit roots	All panels contain unit roots
Criteria	Prob> $F = 0.2730$	One lag: P-val = 0.000 Two lags: P-val = 0.000	No lags: P-val = 0.000 Two lags: P-val = 0.000

Notes: The table shows different tests for serial correlation with their null hypothesis (H_0) and p-value, which corroborate the stationarity of $\widetilde{E}_{cpt}$. The first column shows results for the Wooldridge test (Wooldridge, 2002) for autocorrelation in panel data. We also use the Levin-Lin-Chu (Levin et al., 2002) and the Im-Pesaran-Shin (IPS) tests (Im et al., 2003) in columns (2) and (3), respectively, which are appropriate for panel data with a large number of cross-sections and a smaller number of time periods.

C.3 Drivers of Variation in the Residuals

We examine the determinants of the variation in our residuals (equation (1)), asking whether it is driven by local brand entry and exit within narrow product categories.

Data To explore this question, we leverage microdata on the dynamics of different narrowly defined products, brands, and retailers across U.S. regions and time, for several product categories. Our source is *Consumer Insights* (CI), which provides consumer survey data for durable goods; items like home appliances, power tools, and electronics are all examples of goods that fit into this category.[79] This Durable IQ survey is run quarterly and surveys 600,000 households in the United States each year. The survey asks consumers a series of questions about the comprehensive purchases that they have made within the past 90 days, and is highly specific about features of each product. The data spans 2015-2022 and includes product type, brand, outlet (i.e., retailer, including Lowe's, Home Depot, or Walmart), online vs. in-store channel, demographics of users, and city where the purchase took place.

Strategy All the cities in the CEX are also available in the CI data. We then match product-brand pairs in the CI data to their corresponding categories in the CEX. In some cases, products in both data sets coincide one-to-one. In other cases, CI classifies products differently from how they are described in the CEX category, or does not include products in some CEX categories.[80] Therefore, we focus on items available in both data sets and that roughly correspond with each other. Since our purpose is to study the drivers of the variation in the residuals using the detailed CI microdata, we also focus on CEX categories whose residualized expenditures significantly correlate with those of the CI data. Thus, we consider five main categories which include electronics, small appliances, major appliances, and tools.[81]

Results Column (1) of Table C.6 shows that the log expenditures in the CI data correlate well with those in the CEX after including the battery of fixed effects; i.e., residualized expenditures are correlated. To explore the drivers of the variation in the residuals, we define products in the CI data as product type-brand pairs; for instance, a JBL mini speaker. Column (2) shows significant correlation between the residuals in expenditures in the CEX and changes in the number of product-brand pairs in the CI data; in fact, as shown in column (4), changes in brands within product categories are the main drivers of changes in CI expenditures, with a correlation of 0.87. Columns (3) and (5) show that the number of retailers in a location also correlates with changes in expenditures in both datasets, although not as strongly as with product-brand pairs.

[79] While the Nielsen IQ data would have provided information on more categories, its access is restricted whenever a co-author is not tenured or tenure-track faculty, a PhD student or postdoc. Thus, it is incompatible with using Costa Rican administrative data.

[80] For instance, a CEX category includes: car stereo, CB radio, mixer, speakers, stereo, amplifier, clock radio, receiver, turntable, walkie-talkie, equalizer, compact disc players, short-wave radio, stereo system, tuner, satellite radio, audio cassette players/recorders. In this case, the CI data has the following products which match the category in the CEX reasonably well: home theater in box, mini speaker, mini/shelf stereo system, separate receiver, soundbar, traditional speaker.

[81] Namely, the categories are: "Small electric kitchen appliances," "Stereos, radios, speakers, and sound components including those in vehicles," "Clothes washer or dryer (owned

Table C.6: Correlation: Categories in the CEX vs. Consumer Insights (CI)

	Log Exp. (CEX)			Log Exp. (CI)	
	(1)	(2)	(3)	(4)	(5)
Log Exp. (CI)	0.3651***				
	(0.094)				
Log Products (CI)		0.3765**		0.8696***	
		(0.152)		(0.115)	
Log Retailers (CI)			0.2564*		0.5074***
			(0.147)		(0.097)
Observations	2,299	2,300	2,300	2,299	2,999
R-squared	0.673	0.671	0.671	0.983	0.979
sp, *st*, *pt* FE	Yes	Yes	Yes	Yes	Yes

Notes: The table shows correlations between the *residuals* of (log) products and retailers and (log) expenditures by product category per city in the CEX and CI data. Products are defined as type-brand pairs based on UCC codes in the CI data, and as HS-4 categories in the CEX data. Robust standard errors are in parentheses. Regressions control for city×product, city×time, and product×time fixed-effects.

We then zoom into the five main categories described above, and ask whether the dynamics of entry and exit of brands *within* narrowly defined products in each category can explain movements in the residualized expenditures. Figure C.5 plots these residuals across time and many different U.S. cities. First, in line with the tests in Table C.5, these residuals do not display serial correlation. Second, in line with Table C.6, the residualized expenditures (solid black lines) strongly co-move with the product type-brand residuals (dashed red line). In other words, the number of brands *within* narrowly defined products in each category seem to drive changes in the residuals. Finally, also in line with Table C.6, the dynamics of retailers also co-move with the other residuals, albeit more weakly than for product brands. Overall, the analysis suggests that the dynamics of the residuals are greatly driven by the differential entry and exit of products of different brands across space and time.

Specific Examples The CI data allow us to identify whether surges in residuals coincide with the localized introduction of brands. For example, in the category Small Electric Kitchen Appliances in Minneapolis–St. Paul–Bloomington, MN–WI, we observe a sharp increase in residualized expenditures in 2018q2, coinciding with the introduction of new brands such as Instant Pot and Frigidaire in the food-processor segment. In the same category, Denver–Aurora–Lakewood, CO, shows a similar spike in 2017q1, corresponding with the first appearance of Kitchensmith and Oster in both toaster and can-opener products. Likewise, in the category Clothes Washers and Dryers, Washington–Arlington–Alexandria, DC–VA–MD–WV, exhibits a surge in 2016q2 following the entry of Asko and Speed Queen, while Dallas–Fort Worth–Arlington, TX, shows a spike in 2015q1 with the introduction of Panda and Magic Chef.

home)," "Portable heating and cooling equipment," and "Power tools."

C.4 Examples in the CEX Data

We provide examples of positive and negative (residualized) expenditure shocks in specific U.S. cities. These shocks reflect localized retail expansions, product launches, and region-specific events that were corroborated with publicly available information:

– *Massachusetts and New Hampshire, 2019q1.* A large positive shock in winter-sport equipment, including skis, bindings, and skates (HS 950611–950699), coincides with the exceptionally strong 2018–2019 ski season in New England, when abundant snowfall boosted regional sales of winter-sport gear.[82]

– *Massachusetts and New Hampshire, 2019Q2–Q3.* A spending surge in home and office furnishings (HS 83040000) coincides with the announcement and subsequent opening of Wayfair's first brick-and-mortar store in Natick, MA.[83]

– *Hawaii, 2018q4.* A sharp increase in video-recording equipment (HS 852110) coincides with the launch of the GoPro HERO7 and record tourist arrivals, both of which spurred local purchases of action cameras for snorkeling and surfing.[84]

– *Texas, 2018q1.* A spike in household furniture (HS 940350) aligns with the opening of the new IKEA Grand Prairie store in December 2017, which generated a wave of local furniture purchases in the following quarter.[85]

– *Georgia, 2017q1–q2.* Apparel expenditures (HS 610413, 6204110000, 620413) decline sharply following Macy's and JCPenney store closures across the state.[86]

– *Wisconsin, 2018q2–q4 and Illinois, 2018q3.* Spending on men's suits and knitwear (HS 620331–39, 610431–39) falls after the liquidation of the Bon-Ton/Carson's department-store chains and concurrent Sears closures.[87]

– *Alaska, 2019q4.* Apparel purchases drop markedly (HS 610431–39, 620431–39) amid the state's warmest year on record, which reduced demand for cold-weather gear.[88]

[82] See Valley News, "Plenty of Snow Translated Into Strong New England Ski Season," 2019, at vnews.com/2019/06/08/plenty-of-snow-translated-into-strong-new-england-ski-season-26139635/.

[83] See "Wayfair Expands to Physical Retail with First Brick-and-Mortar Store," 2019, at investor.wayfair.com/news/news-details/2019/Wayfair-Expands-to-Physical-Retail-with-First-Brick--Mortar-Store.

[84] See GoPro Inc. Press Release, "GoPro Reports Record Month-One Sell-Through for HERO7 Black," 2018, at investor.gopro.com/press-releases/press-release-details/2018/GoPro-HERO7-Black-Achieves-Strongest-Week-One-Unit-Sell-Thru-in-Company-History.

[85] See IKEA U.S. Newsroom, "IKEA Grand Prairie to Open on December 2017" at www.ikea.com/us/en/newsroom/corporate-news/ikea-grand-prairie-to-open-on-december-13-2017-pubbde4b667/.

[86] See Athens Banner-Herald, "Macy's at Georgia Square Mall to Close," 2017 at www.onlineathens.com/story/news/state/2017/01/05/macy-s-georgia-square-close/15442941007.

[87] See Milwaukee Journal Sentinel, "Bon-Ton Identifies Stores It Plans to Close," 2018, at www.jsonline.com/story/money/business/2018/01/31/bon-ton-identifies-stores-plans-close/1082868001; NBC Chicago, "Carson's Closing All Illinois Stores," April 2018, at https://www.nbcchicago.com/news/local/carsons-bon-ton-closing-illinois-stores/45032/.

[88] See NOAA, "State of the Climate: National Climate Report for December 2019" at www.ncei.noaa.gov/news/national-climate-201912.

Figure C.5: Residuals per Product Category and Microdata on Brand Dynamics

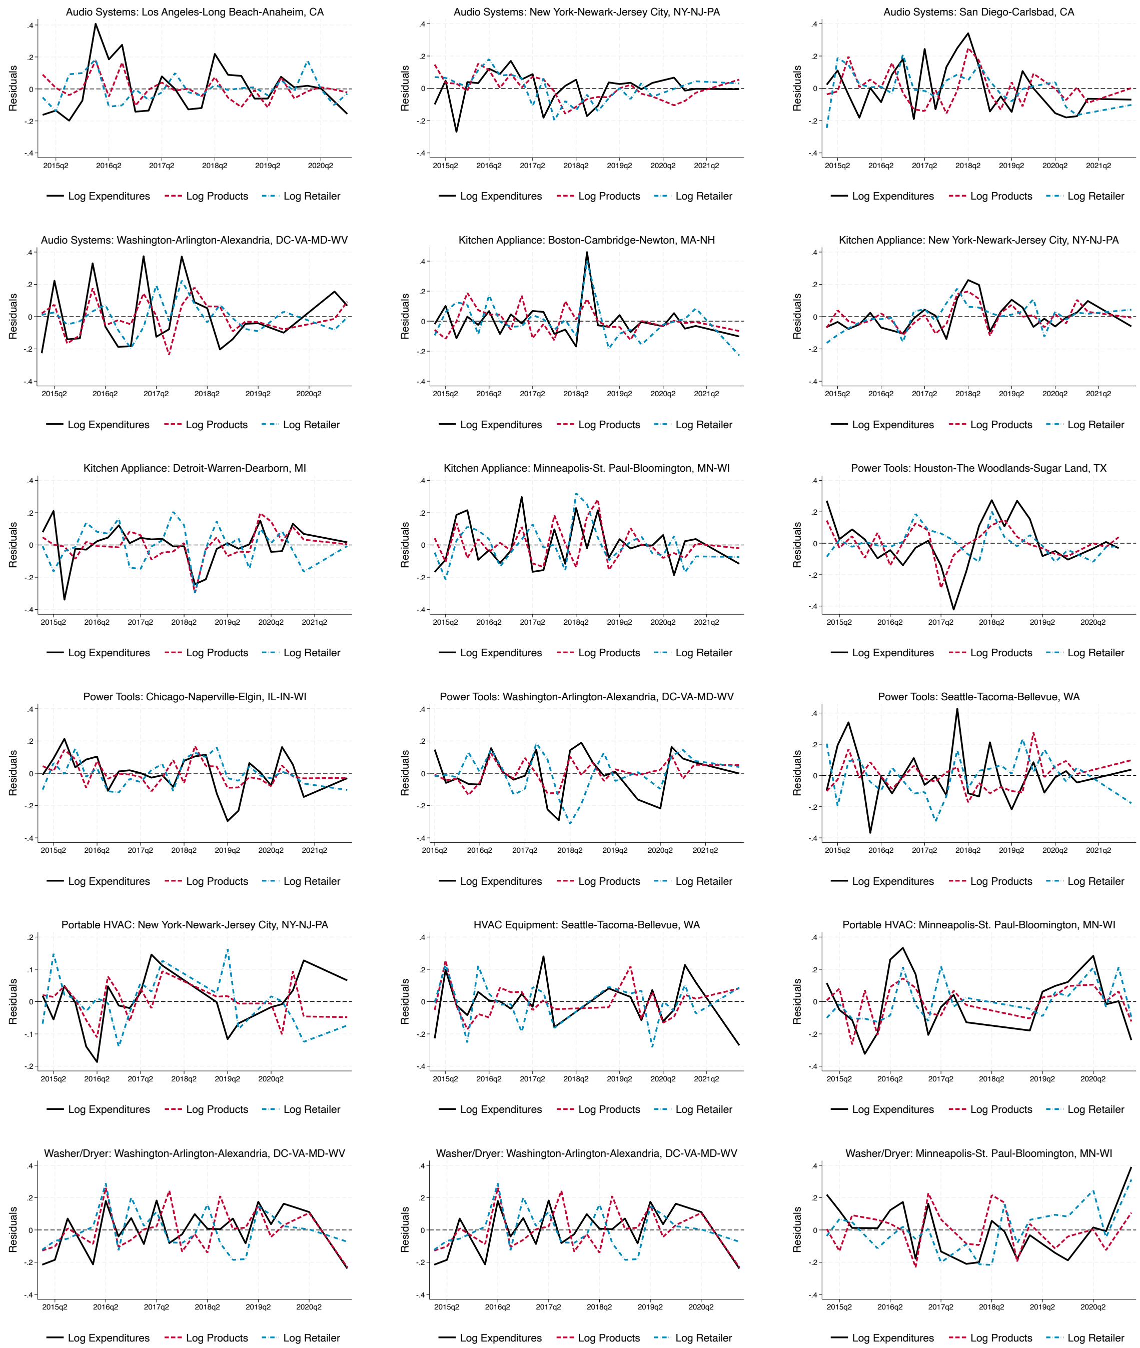

Notes: The figures plot *residuals* of (log) expenditures, (log) products, and (log) retailers by product category across different regions relying on the Consumer Insights (CI) data. The corresponding categories and regions are labeled on top of each plot. Products are defined as product type-brand pairs.

C.5 Identifying Variation by Network-Product Pair

Figure C.6: Identifying Variation by Network-Product Pair

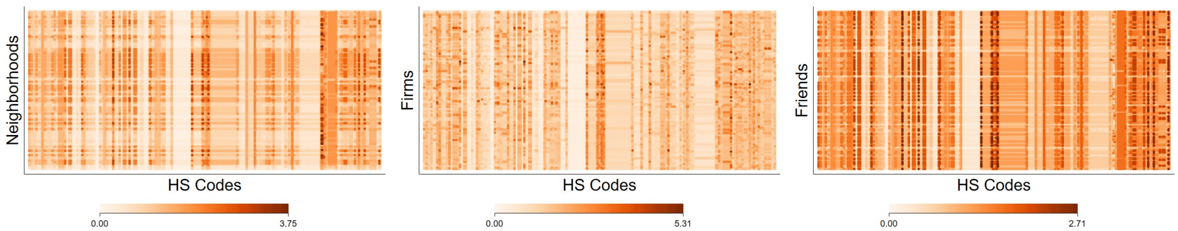

(a) Neighbors (b) Coworkers (c) Friends

Notes: The figure displays differences in the identifying variation across network-product pairs. We compute the term $\ln \widetilde{E}_{bp,t}$ netted of fixed-effects and calculate the variance of this term for each network-product pair. Each panel shows this variance for a network type. Given the large number of networks, for visual purposes panels (b) and (c) are based on a random sample; details in Appendix E.5.

C.6 Balance Test for Migrants to Different U.S. Consulates

Our instrument exploits variation in consumer trends for specific products across the U.S., and links it to people in Costa Rica based on relatives across different U.S. consulates. While we remove the levels from the relevant variation that we use to construct our instrument in equation (1), we want to verify that the observable characteristics of Costa Rican migrants to different consulates across the U.S. balance. To do so, we calculate normalized differences for different characteristics following Imbens and Wooldridge (2009), namely, for individuals in consulate c, one would calculate the following for observable characteristic X:

$$\frac{\bar{X}_c - \bar{\mu}_{-c}}{\sqrt{S_c^2 + S_{-c}^2}},$$

where $\bar{X}_c$ (S_c) is the mean value (standard deviation) of X for people migrating to c and $\bar{\mu}_{-c}$ (S_{-c}) is the mean value (standard deviation) of X for people migrating to a consulate other than c. The rule of thumb is that an absolute value of the normalized difference exceeding 0.25 indicates strong imbalances.

Table C.7: Characteristics of Migrants and Normalized Differences

Main consulate	Total	Age (years)		Female (=1)		Wages	
in the U.S.	N	Mean	Norm. diff.	Mean	Norm. diff.	Mean	Norm. diff.
Atlanta	2,605	39.14	-0.07	0.47	0.05	456	-0.01
Houston	1,771	40.11	-0.00	0.48	0.06	566	0.11
Los Angeles	3,080	42.11	0.12	0.53	0.14	575	0.12
Miami	3,458	41.59	0.09	0.51	0.10	526	0.07
New York	9,785	39.77	-0.04	0.38	-0.16	343	-0.22
Washington	1,860	39.05	-0.07	0.44	-0.01	604	0.14
Chicago	883	40.26	-0.12	0.46	0.02	662	0.14

Notes: Mean wage is in thousands of (real) Costa Rican currency. Data is monthly and spans 2015-2019.

Table C.7 shows these normalized differences for age (in years), gender, and wages for the main Costa Rican consulates in the U.S. While the first two observables are available for all migrants from National Registry data, the last one is only available for migrants who were formally employed before migrating, and whose employment took place at least during one month between 2006 and 2019. As shown, the balance in characteristics of migrants to different U.S. consulates is remarkable; *all of the normalized differences are close to zero and well below 0.25* (in absolute value).

D Details on Networks of Friends

As briefly explained in Section 2, we use data on comprehensive transactions on *Sinpe Móvil*, an application that allows Costa Ricans to make peer-to-peer money transfers using their mobile phones, to construct networks of friends. Over 60% of all adults in the country are users of this technology to send money to their peers (Alvarez et al., 2023). First, we leverage information on bilateral transactions across users, and their unique identifiers, to identify which pairs of people have sent money to each other in the past. Second, we want to clear this mapping from people who used to app to make a payment (for instance, a parent transferring money to a nanny). Thus, we focus only on pairs of individuals who have sent money to each other *bilaterally*, and use this to construct our proxy of "friends." For instance, if user A has only sent money to user B, we would not record this relationship as a friendship. If, however, both A and B have sent money to each other at some point in time, then their relationship is classified as a friendship. While imperfect, this allows us to proxy for networks of friends which are usually impossible to recover.

In the first stage for neighborhood and coworker networks we focus on directly exposed individuals (those with at least one relative in the U.S.) and examine whether the probability of importing a given product depends on the exposure of their relative to that product in the U.S. city where the relative lives. We implement the same first-stage construction for friendship networks (restricting attention to individuals for whom we observe at least one friendship tie).

Analyzing the second stage presents a greater level of complexity. Figure D.1 shows an example. Suppose A and B are friends. Panel (a) is a diagram showing A's friends, and panel (b) depicts B's friends. Moreover, orange circles represent friends who have relatives living in the U.S. (i.e., they are directly exposed), while blue circles denote friends who are not directly exposed. Focus on panel (a): A only has one exposed friend. Now, is B's exposure coming from this one friend only? Just observing panel (a), it might be tempting to answer positively; however, as shown in panel (b), this is not necessarily the case. Note that this is not an issue for networks of neighbors or coworkers, because they are partitions within each time period.

This example illustrates the rationale behind our decision to define networks of

Figure D.1: Networks of Friends: Example

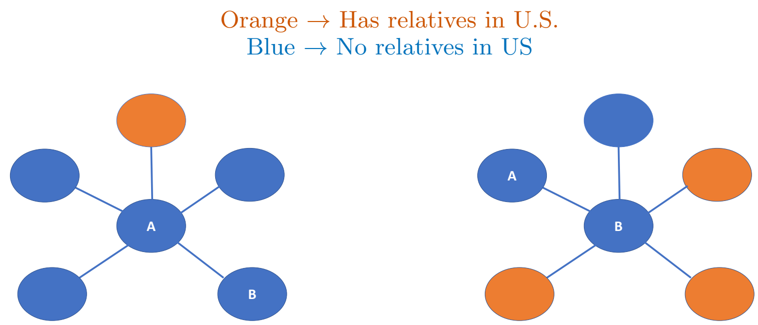

(a) A's Friends (b) B's Friends

friends *on an individual-specific basis* (i.e., A has three friends, each friend has her friends...). The example also shows why there are as many networks as users of the app with at least one friend, and why for the second stage involving the *friends* network, our dependent variable includes only imports of the *centroid* of the network (i.e., A's imports when considering A's network). We also make the assumption, in line with our exposure mapping for other networks, that only direct effects (i.e., friends "one link away") enter the exposure mapping's functional form.

E Main Analysis: Additional Results

E.1 Note on Clustering

Our exposure mapping in equation (4) is built from U.S. consulate×product×time shocks, $\ln \widetilde{E}_{cpt}$, using fixed network shares s_{bc}, and is therefore similar in spirit to a shift-share. In the terminology of Adão et al. (2019), the U.S. shocks are the primitive "shifters," and the s_{bc} are exposure shares, and inference would ideally be robust to arbitrary dependence among all observations that load on the same shifters.

In our setting, however, the first-stage and second-stage specifications in equation (5) and equation (6) are estimated on stacked individual–network–product–time data and saturated with fixed effects that absorb the main low-dimensional components through which the U.S. shocks can generate dependence in Costa Rica. In particular, we include individual fixed effects γ_i, network×HS-4 (or HS-6) product fixed effects $\gamma_{b\tilde{p}}$, and network×time fixed effects γ_{bt}; upstream HS-10 product×time effects are already partialled out in equation (1). Intuitively, γ_{bt} removes shocks common to all products within a network at a given time, $\gamma_{b\tilde{p}}$ removes time-invariant network–product heterogeneity, and product×time effects remove shocks common to all networks for a given product and time.

Baseline clustering. After controlling for this rich set of fixed effects, the identifying variation in the instrument and the endogenous variables comes from *residual* differences across network–product cells over time, rather than from co-movements at the network or product level. This suggests treating each network–product pair $(b, \tilde{p})$ (where $\tilde{p}$ is HS-4/HS-6) as the relevant "panel unit" and modelling the remaining error correlation as primarily within $(b, \tilde{p})$ over time and across individuals. Formally, letting $\varepsilon_{i,bpt}$ denote the residual from equation (5) or equation (6), our maintained assumption is that, conditional on the fixed effects and on the slowly-varying s_{bc},

$$E\big[Z_{bpt}\, \varepsilon_{i,bpt}\, Z_{jqs}\, \varepsilon_{k,jqs}\big] = 0 \quad \text{whenever } (b, \tilde{p}) \neq (j, \tilde{q}),$$

while allowing for arbitrary correlation across individuals and over t within a given $(b, \tilde{p})$ cell, where Z_{bpt} denotes our instrument. Under this correlation structure, clustering standard errors at the network–product (HS-4/HS-6) level is the appropriate cluster-robust procedure: it allows for arbitrary within-cell and serial correlation, while treating different network–product cells as asymptotically independent. Throughout the paper, we therefore report standard errors clustered at the network–product level.

Robustness checks. A complementary perspective emphasizes that dependence can arise mechanically from common shifter realizations. In our design, the primitive shocks vary at the consulate×product×time level, and the exposure measure aggregates these shocks using fixed shares. An AKM-style implementation therefore motivates clustering together observations that load on the same shifter realizations in a given period. A natural and parsimonious choice is clustering by product×time: this allows for arbitrary dependence across individuals and networks jointly exposed to the same period-specific product shock component.

Whether one should additionally allow for dependence *across time* at the shifter level depends on the extent of serial correlation in the shifters. Table C.5 assesses this directly for our residual exposure measure; we find no evidence of first-order serial correlation across various tests. This evidence suggests that any dependence induced by the shifters is primarily contemporaneous rather than persistent, making product×time clustering a relevant AKM-style robustness exercise in our application. Table E.1 re-estimates the main 2SLS specifications across networks and for retailers reporting standard errors clustered by product×time and by product. These alternatives yield very similar, and sometimes smaller, standard errors than those obtained under our baseline network–product clustering. We also report clusters by network×HS-10, again, inference is not materially altered; in fact, this alternative yields errors that are smaller than our baseline.

Table E.1: 2SLS: Network and Retailer Adoption with Alternative Clustering

Dep. variable: Prob. importing product p for individual i without relatives in the U.S. and who belongs to network b at time t

	%Δ w.r.t. mean import probability			
	Neighbors	Coworkers	Friends	Retailers[a]
	(1)	(2)	(3)	(4)
$\widehat{\text{ShareImporters}}^{US\,direct}_{bp,t-1}$	13.897	18.069	14.679	9.045
Cluster by network×HS-4 (baseline)	(5.193)***	(8.154)**	(5.027)***	(0.722)***
Cluster by network×HS-10	(4.106)***	(7.354)**	(3.379)***	(0.672)***
Cluster by HS-10×time	(4.716)***	(8.554)**	(6.340)***	(1.833)***
Cluster by HS-10	(4.112)***	(8.599)**	(5.664)***	(2.565)***
F-stat first stage (baseline)	33.00	10.52	15.63	1035.5
Observations	289,340,892	299,920,162	260,952,672	97,500,332
Clusters (network×HS-4)	200,308	236,804	4,568,240	2,187,504
Clusters (network×HS-10)	698,578	826,102	16,070,571	7,599,302
Clusters (HS-10×time)	7,626	7,713	7,626	7,198
Clusters (HS-10)	432	432	432	432
$b\tilde{p}$, bt, i FE	Yes	Yes	Yes	-
$b\tilde{p}$, bt, f FE	-	-	-	Yes

Notes: The table displays our 2SLS results. Robust standard errors under various clustering choices are reported. [a]For retailers, column (4), clusters by network×HS-4 and network×HS-10 correspond with retailer×HS-4 and retailer×HS-4 clusters. Coefficients are percentages of mean outcome for a one standard deviation change in the instrumented share. Regressions in columns (1)-(3) control for network×product (HS-4), network×time, and individual fixed-effects. The regression in column (4) controls for neighborhood×product (HS-4), neighborhood×time, and retailer fixed-effects. Appendix E.5 includes more details on the sample used in each regression.

E.2 Propagation Across Migrant Networks: Individual-Level

Table E.2: Individual Imports and Relatives' Exposure to Products Abroad

Dependent variable: Prob. importing product p for individual i with a relative in the U.S. in consulate c at time t

	Instrument based on	
	CEX	U.S. imports
	(1)	(2)
$ln\widetilde{E}_{cp,t-1}$	12.387	22.203
(%Δ w.r.t. mean import probability)	(2.005)***	(4.417)***
F-stat first stage	38.16	25.27
Observations	709,806,755	710,177,605
Clusters	11,793,674	28,576,891
$i\tilde{p}$, it FE	Yes	Yes

Notes: Column (1) shows results when constructing exposure measures based on the CEX, while column (2) relies on imports by U.S. customs districts. Robust standard errors, adjusted for clustering by individual-product, are in parentheses and define product as HS-4 for column (1) and HS-6 for column (2). Coefficients report the percent change in the mean import share implied by a one–standard-deviation increase in exposure. Exposure is standardized. Dependent variables are the probability that an *individual* imports a *specific* product code in a particular quarter *and* from the U.S., thus, by design, the percentage mean import probability of a product is small; .001 and .004 for each column. Regressions control for individual-product and individual-time fixed effects. These regressions use the full sample of individuals and products.

E.3 Reduced Form and OLS Results

Table E.3: OLS and Reduced Form Regressions

Dep. variable: Prob. importing product p for individual i without relatives in the U.S. and who belongs to network b at time t

Panel (a): Reduced Form

	%Δ w.r.t. mean import probability		
	Neighbors (1)	Coworkers (2)	Friends (3)
$ln\widetilde{E}_{bp,t-1}$	14.151	18.162	15.539
	(4.773)***	(6.127)***	(3.589)***
Adjusted-R^2	0.003	0.008	0.022
Observations	289,340,892	299,920,162	260,952,672
Clusters	200,308	236,804	4,568,240
$b\tilde{p}$, bt, i FE	Yes	Yes	Yes

Panel (b): OLS

	%Δ w.r.t. mean import probability		
	Neighbors (1)	Coworkers (2)	Friends (3)
$\text{ShareImporters}_{bp,t-1}^{US\,exposure}$	16.371	-0.928	23.401
	(5.501)***	(0.412)**	(27.711)
Adjusted-R^2	0.003	0.008	0.022
Observations	289,340,892	299,920,162	260,952,672
Clusters	200,308	236,804	4,568,240
$b\tilde{p}$, bt, i FE	Yes	Yes	Yes

Notes: The table shows the reduced form and OLS results in panels (a) and (b), respectively. Robust standard errors, adjusted for clustering by network-product (HS-4), are in parentheses. Reduced form coefficients report the percent change in the mean import share implied by a one–standard-deviation increase in (standardardized) exposure. We include network×product (HS-4), network×time, and individual fixed-effects. Appendix E.5 presents details on the samples. Data is quarterly and spans 2015-2019.

Table E.4: Second-Stage Regressions - No Normalizations

Dep. variable: Prob. importing product p for individual i without relatives in the U.S. and who belongs to network b at time t

	Neighbors (1)	Coworkers (2)	Friends (3)
$\widehat{\text{ShareImporters}}_{bp,t-1}^{US\,exposure}$	0.474	0.351	0.471
	(0.177)***	(0.158)**	(0.161)***
F-statistic	33.00	10.52	15.63
Observations	289,340,892	299,920,162	260,952,672
Clusters	200,308	236,804	4,568,240
$b\tilde{p}$, bt, i FE	Yes	Yes	Yes

Notes: The table shows our second stage results for diffusion among peers. Robust standard errors, adjusted for clustering by network-product (HS-4), are in parentheses. We include network×product (HS-4), network×time, and individual fixed-effects. Appendix E.5 presents details on the sample used in each regression. Data is quarterly and spans 2015-2019.

E.4 Results with Exposure Based on U.S. Customs Data

Table E.5: Results Relying on U.S. Customs Data (Neighbors)

Dep. variable: Prob. importing product p for individual i without relatives in the U.S. and who belongs to network b at time t

First Stage		Second Stage	
	(1)		(2)
$ln\widetilde{E}_{bp,t-1}$	26.926	$\widehat{\text{ShareImporters}}^{US\,exposure}_{bp,t-1}$	45.122
	$(5.964)^{***}$		$(14.959)^{***}$
F-statistic first stage	20.38		20.38
Observations	274,933,487		274,933,487
Clusters	484,377		484,377
Mean import prob. $[i, bpt]^{US}$	.0003		.0003
$b\tilde{p}$, bt, i FE	Yes		Yes

Notes: The table displays the results of running our 2SLS for networks of neighbors while relying on imports by U.S. customs districts to estimate the exposure to U.S. expenditure shocks. Robust standard errors, adjusted for clustering by network-product (HS-6), are in parentheses. First stage coefficients report the percent change in the mean import share implied by a one–standard-deviation increase in exposure. Exposure is standardized. Second stage coefficients are percentages of mean outcome for a one standard deviation change in the instrumented share. Regressions control for network×product (HS-6), network×time, and individual fixed-effects. Percentage mean import probabilities are reported. Appendix E.5 presents details on the sample in each regression. Data is quarterly and spans 2015-2019.

E.5 Samples Across Regressions

Regressions are run at the individual level. Regressions are also run by product—with hundreds (thousands) of products in the CEX (U.S. customs districts data). Moreover, regressions are run at the quarterly level from 2015 to 2019. For regressions involving only Costa Rican residents *with* relatives in the U.S., it is possible to use the entire sample of individuals (e.g., Table E.2).

However, 2SLS regressions, which involve everyone *without* relatives in the U.S. in the second stage, pose a challenge. We would be dealing with approximately 4 million adults, so that stacking the full individual–product–quarter panel would imply nearly 80 billion potential observations. This figure corresponds to the full Cartesian product of individuals, products, and quarters. In practice, the estimating sample is restricted to product–quarter cells with positive exposure. However, this dimensionality, in addition to the heavy battery of fixed-effects and the need to run a 2SLS, would make regressions computationally unfeasible. To make the 2SLS computationally feasible, we draw a 1.35% random sample of adult individuals without relatives abroad within each network. After imposing the same product and exposure restrictions used in the main specification, this sampling yields approximately 300 million observations, which allows to conduct computations with a large battery of fixed effects and for the possibility of adding interaction terms. For these randomly selected samples, we then conduct the first stage, which considers the *entire* set of individuals in their network *with* relatives abroad. The actual number of observations in each regression varies, as each network type has exposure to distinct products, and for instance, some networks might not have exposure at all to a product in a given quarter. Table E.6 summarizes the samples used in each table and figure of the main paper.

For all main regressions, which rely on the CEX to construct exposure measures, *all networks use the entire sample of products,* except for Table 5, which uses a 50% random sample and still includes over 300 million observations, as given its construction, which instruments using the exposure of *every* coworkers' spouse's firm (see Appendix H.1), each individual is exposed to more products than in the baseline coworkers regression. Using the full sample of products for estimations in which we construct our instrument based on U.S. imports by customs districts is not feasible. The reason is that there are 2,443 narrowly defined product codes in the U.S. imports data which are imported by individuals in CR; this would make most regressions have over one *billion* observations, even based on the subsample of individuals we described above. Therefore, throughout the paper, results relying on U.S. imports are also based on random samples. The size of these random samples is chosen to exactly match the total number of products in the CEX.

Table E.6: Product Samples in Each Table and Figure of the Main Paper

Table (1)	Network (2)	Sample of HS-10 Products (3)	Unit of Observation (4)
Tables 3, 4	Neighbors	100% sample	individual$\times p \times t$
Tables 3, 4	Coworkers	100% sample	individual$\times p \times t$
Tables 3, 4	Friends	100% sample	individual$\times p \times t$
Table 5	Distance-3 nodes	50% random sample	individual$\times p \times t$
Table 6	Retailers	100% sample	retailer$\times p \times t$

Notes: Whenever the exercise does not include all products, the subsample is chosen at random. Table 5 has half of the sample of products as, given its construction (Appendix H.1), each individual is exposed to more products than in the coworkers regression, which substantially expands the number of observations.

For visual purposes, Figure C.6 uses a random sample of networks equal to the total number of networks of neighbors on the vertical axes.[89] Results and robustness checks in the appendix follow a similar pattern as described above; for individual-level regressions based on the CEX, we always use the entire sample of products, and for estimates based on U.S. imports, we use a random sample. Table 6 relies on the

[89]Recall networks of coworkers and friends are more numerous (albeit smaller) than networks of neighbors (see Table B.1).

entire sample of products. Retailers are defined as ISIC Rev.4 codes 45-47 and all regressions regarding retailers include the entire sample of retailers in Costa Rica.

Additional details on friends network Finally, some individuals have an exceptionally large number of friendships (i.e., hundreds of friends), which goes against the intent of our measure: to capture relatively close relationships. Therefore, we trim the sample by excluding observations above the 99th percentile. The latter also aids in making computations manageable, as individuals with a very large number of friends pose a challenge in this regard. The sample for friends networks differs from other networks in one more aspect: people who use the payments app and who have a relative in the U.S. are more likely than average to import. Thus, if we only consider first-time imports in the first stage, the instrument becomes weak (F-stat 6.97), even though the second stage coefficient is significant and statistically equal to the one reported in Table 4 (16.056 with a standard error of 7.759). Therefore, to obtain a strong instrument, for the first stage of this network only, we restrict importing events to be first-time imports *within sample* (2015-2019), but we will not force them to be first-time imports ever—which we did for other networks and for *all* second stages, as we have the full panel starting from 2005, when online shopping and individual imports were almost zero.

E.6 Timing of the Specifications: Local Projections

We use Jordà (2005) local projections to better understand the timing of the propagation after a network is exposed to a product.[90] In particular, we consider the following set of panel local projections:

$$y_{bp,t+h} = \beta^h \ln \widetilde{E}_{bpt} + \lambda^h x_{bpt} + \gamma^h_{b\tilde{p}} + \gamma^h_{bt} + \varepsilon_{bp,t+h}, \tag{18}$$

where $h = 0, 1, 2, 3$ and x_{bpt} is a vector of controls with lags of the outcome variable and the shock. Using this specification, we first run the first stage where the dependent variable corresponds with $ShareImporters^{US\,direct}$. Results for the first stage use one lag of the outcome variable and the shock, as indicated by an Akaike information criterion, and as including further lags results in a weak instrument. We also follow equation (18) to run reduced form regressions for our second stage and retailers' responses, where the dependent variables are first-time imports of those without relatives abroad in neighborhood b and first-time imports of retailers in b, respectively. The reduced form results use three lags of the outcome variable and of the shock;

[90]Local projections are based on sequential regressions of the endogenous variable shifted several steps ahead (Jordà, 2005). They are able to accommodate IV estimations (Jordà et al., 2020), and they can robustify inference and simplify the computation of standard errors (Montiel Olea and Plagborg-Møller, 2021).

pre-trends are controlled for by our lag specification. Results are similar with less stringent specifications on the number of lags.

We now present impulse responses from an exogenous increase in exposure estimated using local projections. Panel (a) of Figure E.1 reports the first stage. One quarter after the increase in exposure, the import probability starts falling (vertical dashed line). Thus, we include exposure with one lag in equation (5), as it is the first period that would capture the full effect.

Figure E.1: Local Projections: Impulse Response

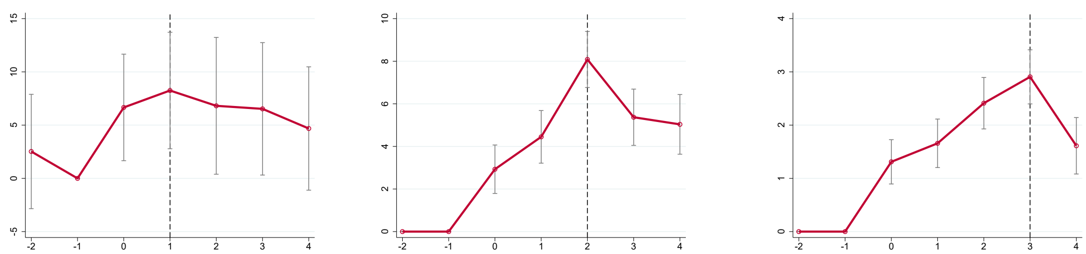

(a) First Stage (b) Peers, Reduced Form (c) Retailers, Reduced Form

Notes: Estimations follow equation (18). Panel (a) shows impulse responses for our first stage specification. Panel (b) shows impulse responses of the probability of importing for people without relatives abroad to an increase in network exposure from abroad. Panel (c) shows impulse responses of the probability of importing for retailers to an increase in network exposure.

Panels (b) and (c) of Figure E.1 show the impulse response of the reduced form regressions for the second stage—concerning of those without relatives abroad—and retailers, respectively. There is a clear peak in the impulse response in periods 2 and 3 in each panel, which corresponds with the timing used in our analysis.

We then implement a similar local projection for the IVs. First, we consider imports of people in a network who are unrelated to migrants abroad, as follows:[91]

$$Imports_{bp,t+h} = \beta^h \widehat{ShareImporters}^{US\,direct}_{bp,t-1} + \lambda^h x_{bpt} + \gamma^h_{b\tilde{p}} + \gamma^h_{bt} + \varepsilon_{bp,t+h}. \quad (19)$$

Panel (a) of Figure E.2 reports the results. The figure shows that the impulse response peaks on impact, which is in line with the timing of our second stage given equation (19), as its main independent variable is lagged one period. Lastly, we study the response of retailers using an IV specification that aligns with our baseline:

$$Imports^F_{bp,t+h} = \beta^h \widehat{ShareImporters}^{US\,direct}_{bp,t-2} + \lambda^h x_{bpt} + \gamma^h_{b\tilde{p}} + \gamma^h_{bt} + \varepsilon_{bp,t+h}. \quad (20)$$

Panel (b) of Figure E.2 shows that retailers import on impact, and the impulse response is then flat, which aligns with the timing in equation (10) as there is a two-lag difference between the dependent and independent variables in equation (20).

[91]For examples of IV applications using local projections, see Jordà et al. (2020).

Figure E.2: Local Projections IV: Impulse Response

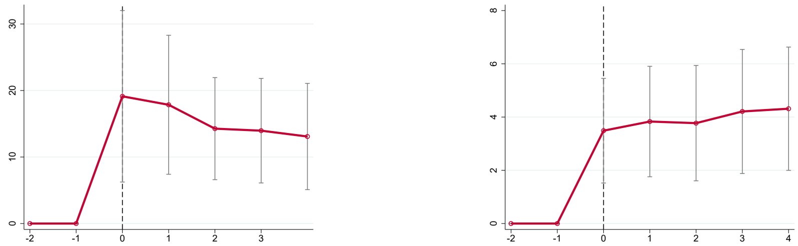

(a) Individuals' Response (Second Stage)   (b) Retailers' Response

Notes: Panel (a) shows the impulse response resulting from the local projection in equation (19), which concerns imports of individuals without relatives abroad. Panel (b) shows the impulse responses of the probability of importing for retailers in equation (20).

Reassuringly, the timing of the local projections for retailers' responses aligns with the timing reported by respondents in our survey (see Section 5.2 for survey details). Responses to the question: "*If you decided to start importing a new product, how long [in quarters] would it take from the moment you make the decision to having the product for sale?*" are presented in Figure E.3. From the moment they observe the local demand, it takes most retailers a quarter to respond (i.e., two periods after those *with* relatives import, and three periods after the U.S. receives the shock).

Figure E.3: Typical Time to Import After a Rise in Observed Demand

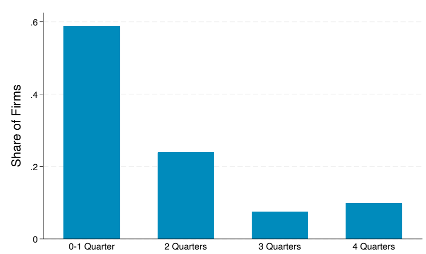

Notes: The figure shows the share of responses of retailers by quarter to the question: "*If you decided to start importing a new product, how long [in quarters] would it take from the moment you make the decision to having the product for sale?*."

F Back-of-the-Envelope Comparison of Exposures

An individual can be exposed directly and indirectly. We compare both exposures in a reduced form exercise. The sample considers individuals with relatives abroad, and the regressors are their direct exposure and the average exposure of neighbors who have relatives abroad, which is constructed as a leave-out measure that does not consider the consulate where the individual has a relative. We present the results in Figure F.1, and do so based on two batteries of peer effects (panels B1 and B2).

Figure F.1: Direct vs. Indirect Exposure

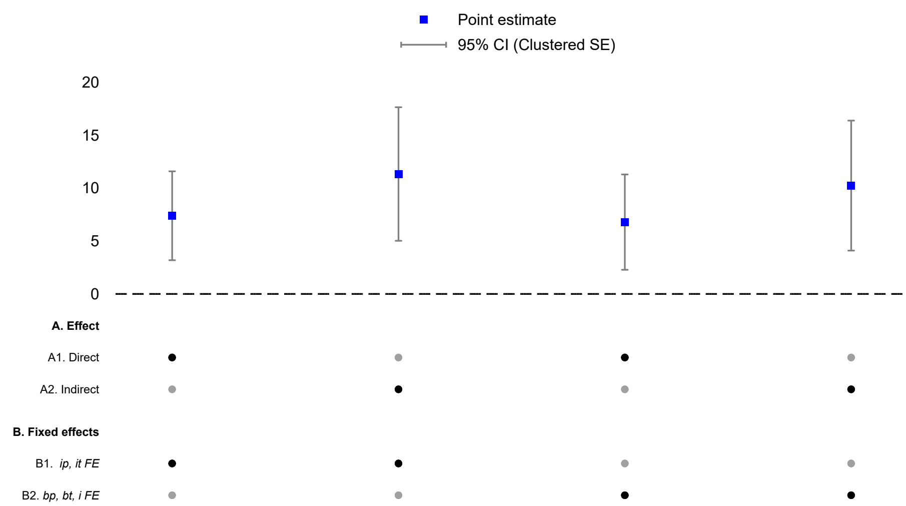

Notes: The figure shows the point estimates associated with regressing an individual's imports on her direct and indirect exposure measures. Indirect exposures are constructed as leave-out measures, which do not consider the consulate where the individual has a relative. Networks are defined as neighborhoods. Regressions have robust standard errors, adjusted for clustering by network-product (HS-4). Vertical bars denote 95% confidence intervals. The independent variables are standardized. Regressions based on panel B1 control for individual×product (HS-4) and individual×time fixed-effects. Regressions based on panel B2 control for network×product (HS-4), network×time, and individual fixed-effects.

These results are reassuring in verifying that both effects coexist. However, although the effects are statistically equal, this does not imply that the influence of a relative abroad is comparable to that of a single neighbor. While 96% of individuals who are directly exposed have one relative residing in the U.S., the indirect exposure results from averaging the exposure of, on average, 85 directly exposed neighbors. Therefore, a per-person comparison must adjust the exposure by the number of people exerting it. A back-of-the-envelope per-person calculation implies that the indirect impact of a neighbor is about 3% of that of a direct relative.

G Cross-Country Aggregate Evidence

We build on Bailey et al. (2021), who use Facebook data to construct a measure of pairwise social connectedness between 170 countries. They show that this measure captures not only standard determinants of trade, but also forces such as migration. Consistent with our findings, they document at the macro level that two countries trade more when they are more socially connected, especially for product categories (i.e., HS2 codes) where information frictions are likely to be large. Bailey et al. (2021) define the Social Connectedness Index ($SCI_{i,j}$) as the total number of Facebook friendship links between individuals in locations i and j, divided by the product of the number of Facebook users in those locations. Then, they estimate gravity regressions showing that trade flows increase strongly with social connectedness, even after controlling for geographic distance, common borders, common official language, colonial ties, and importer and exporter fixed effects.

Based on their aggregate data, Bailey et al. (2021) argue that social connectedness potentially facilitates trade by reducing information frictions: that is, by lowering search costs through interpersonal networks that transmit information about products. Using the insight of Rauch (1999), they show that the elasticity of trade with respect to social connectedness is larger for goods not traded on organized exchanges, where information frictions can be more severe. This macro evidence aligns with our findings, which would microfound these aggregate patterns that are found to hold among a large set of countries.

We complement their analysis by focusing on goods that are more likely to be imported *directly by individuals*—those where peer effects and product discovery through social networks may be particularly important. These include consumer-oriented goods such as perfumes, clothing, and jewelry.[92] Namely, we construct a dummy variable for these products and estimate a modified gravity specification that allows the elasticity of trade with respect to social connectedness to vary for this subset of products.

As shown in Table G.1 and consistent with our microlevel estimates, the elasticity between trade and social connectedness is significantly larger for goods that are more likely to be imported by individuals—those where peer effects and product information diffusion play a more prominent role. This suggests that social connectedness helps reduce information frictions in global trade, particularly for goods frequently imported by individuals and in a manner consistent with the mechanisms we identify in the Costa Rican context.

[92]The full list of HS2 codes likely to be imported directly by individuals, based on importing events that occurred in Costa Rica, includes: 33 (essential oils, perfumes, and cosmetics), 42 (leather articles such as handbags and travel goods), 43 (fur skins and artificial fur), 49 (printed books and newspapers), 57 (carpets and other textile floor coverings), 61 and 62 (apparel, knitted and not knitted), 63 (other made-up textile articles such as curtains and bedding), 64 (footwear), 65 (headgear), 66 (umbrellas and walking sticks), 67 (artificial flowers and hair articles), 71 (pearls and jewelry), 82 (cutlery and small tools), 90 (optical, medical, and photographic instruments), 91 (clocks and watches), 92 (musical instruments), 94 (furniture, lighting, and prefabricated structures), 95 (toys, games, and sports goods), 96 (miscellaneous manufactured articles such as combs and pens), and 97 (works of art and collectibles).

Table G.1: Gravity Regressions - Goods Trade Heterogeneity - Individual Imports

Dependent variable: Product-Specific Exports

	(1)	(2)	(3)
log(SCI)	0.275	0.267	0.277
	(0.027)***	(0.028)***	(0.023)***
log(SCI) × Individual		0.053	0.075
		(0.032)*	(0.027)***
Common official language	-0.048	-0.048	0.011
	(0.074)	(0.074)	(0.059)
Common border	0.317	0.317	0.436
	(0.087)***	(0.087)***	(0.081)***
Common colonizer	-0.051	-0.053	0.001
	(0.145)	(0.147)	(0.095)
In colonial relationship	-0.180	-0.185	-0.031
	(0.242)	(0.242)	(0.196)
Observations	2,263,574	2,263,574	2,263,574
Pseudo R^2	0.932	0.932	0.946
Other Gravity Controls	Yes	Yes	Yes
Origin Country × HS-2 FE	Yes	Yes	Yes
Destination Country × HS-2 FE	Yes	Yes	Yes
Distance Group × HS-2 FE	No	No	Yes
log(Distance) × HS-2	Yes	Yes	No

Notes: The dependent variable is exports of product category k from country i to country j in 2017. Product-level trade data are aggregated up to the first 2 digits of the HS-code product classification. Other gravity controls include a common border dummy, a common official language dummy, a dummy indicating whether the two countries had a common colonizer post-1945, and a dummy indicating whether the pair of countries was in a colonial relationship post-1945. We also separately control for the logarithm of distance interacted with product categories in columns (1) and (2), and for distance groups (dummies for percentiles of the distance distribution) interacted with product categories in column (3). The specification is based on Bailey et al. (2021). Individual is a dummy variable indicating products more likely to be imported directly by individuals. All specifications include fixed effects for the importer and exporter country interacted with product categories. Standard errors are clustered by exporter and importer country. The data include 165 countries and 96 product categories, which amounts to 2,597,760 observations. Observations that are fully explained by the fixed effects are dropped before the PPML estimation.

H Robustness Exercises

H.1 Instrument Using Distance-3 Nodes

As explained in Section 4.5, one of our robustness exercises considers equation (9), where θ is our parameter of interest. To isolate this effect, De Giorgi et al. (2019) exploit that social relationships are established along two lines: at the family level (e.g., husband and wife) and at the firm level. The idea is that shocks at the firm of a coworker's spouse are a valid instrument for the household's consumption changes.[93]

Remarks on Data Construction We identify couples in our sample where both spouses are employed. We then exclude couples who work at the same firm, and also

[93] De Giorgi et al. (2019) run a regression in first differences. We depart from this approach as our dependent variable is an indicator, and product-variation allows for fixed-effects.

coworkers whose spouses work at the same firm, to avoid feedback effects. Information transmission, we assume, occurs across the remaining couples in the sample.

Figure H.1: Distance-3 Nodes: Identifying Variation by Individual-Product Pair

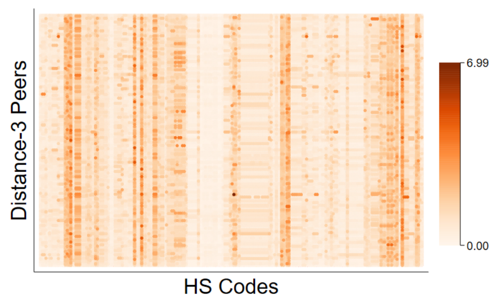

Distinct top codes:
Wooden furniture of a kind used in offices (940330)
Pens and marker parts (960899)
Women's suits of synthetic fibers (610413)

Notes: The figure summarizes the differences in the identifying variation across individual-product pairs. Namely, we compute the term $\widehat{Import^{d}_{ip,t}}$, netted of the fixed-effects in equation (9), and calculate the variance of this term for each individual-product pair. We also display the product codes which were both among the top 10 codes of this sample and were not included in the top 10 codes in the full sample.

H.2 Placebo Exposures and Recentering

Figure H.2: Placebo vs. Actual Coefficients and Recentered Results for Neighbors

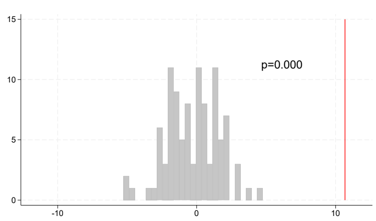

(a) First Stage

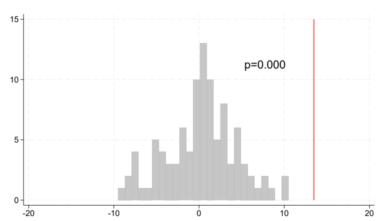

(b) Reduced Form

	Dep. Var: $Import_{i,bpt}$
$\widehat{\text{ShareImporters}}^{recenter}_{bp,t-1}$	14.214
	(5.124)***
F-stat first stage	35.23
Observations	289.3M
Clusters	200,308
$b\tilde{p}$, bt, i FE	Yes

(c) Recentered 2SLS

Notes: Panels (a) and (b) plot the distribution of placebo coefficients based on placebo exposure measures defining networks as neighborhoods and after 100 randomizations. The red vertical lines plots the actual first stage (panel a) and reduced form coefficient (panel b). The p-values are the share of the placebo coefficients that are larger than the coefficient for the actual first stage or reduced form. Panel (c) reports the 2SLS based on a "recentered" version of exposure, by subtracting the expectation of the treatment value under the randomized distribution from our original exposure.

H.3 Results with Different Fixed-Effects

Table H.1: Results Relaxing Fixed-Effects for Neighbors

Dep. variable: Prob. importing product p for individual i without relatives in the U.S. and who belongs to network b at time t

	%Δ w.r.t. mean import probability		
	(1)	(2)	(3)
$\widehat{\text{ShareImporters}}_{bp,t-1}^{US\,direct}$	14.967	14.312	19.738
	(5.092)***	(5.027)***	(5.813)***
%Δ w.r.t. baseline	+7.7%	+3.0%	+42.0%
F-stat first stage	44.02	41.49	41.05
Observations	289,341,228	289,341,228	289,341,228
Clusters	200,644	200,644	200,644
$b\tilde{p}$ FE	No	No	No
bt FE	Yes	No	Yes
sp FE in residualization	Yes	Yes	No

Notes: The table shows the results of running equation (6) defining networks as neighborhoods and removing certain controls. Column (1) removes $\gamma_{b\tilde{p}}$ fixed effects from the baseline specification in equation (6), while column (2) also removes the γ_{bt} fixed effects from the same equation. Column (3) removes both the level effect from the regression equation (6) and the region×HS-10 (sp) fixed effect from the residualization of the instrument in equation (1). Robust standard errors, adjusted for clustering by network-product (HS-4), are in parentheses. Coefficients are percentages of mean outcome for a one standard deviation change in the instrumented share. All regressions control for individual fixed-effects. Appendix E.5 presents details on the sample used in each regression. Data is quarterly and spans 2015-2019.

Table H.2: Results with More Demanding Specifications for Neighbors

Dep. variable: Prob. importing product p for individual i without relatives in the U.S. and who belongs to network b at time t

	%Δ w.r.t. mean import probability	
	(1)	(2)
$\widehat{\text{ShareImporters}}_{bp,t-1}^{US\,direct}$	16.342	15.721
	(5.724)***	(5.472)***
F-stat first stage	39.46	37.92
Observations	289,339,460	289,324,490
Clusters	200,235	199,535
District×$\tilde{p}$×t FE	Yes	No
Network×HS-2×t FE	No	Yes

Notes: The table shows the results of running equation (6) defining networks as neighborhoods and adding more demanding controls. Column (1) includes District×$\tilde{p}\times t$ fixed effects, while column (2) includes network×HS-2×t fixed effects. Robust standard errors, adjusted for clustering by network-product (HS-4), are in parentheses. Coefficients are percentages of mean outcome for a one standard deviation change in the instrumented share. All regressions control for network×product (HS-4), network×time, and individual fixed-effects. Appendix E.5 presents details on the sample used in each regression. Data is quarterly and spans 2015-2019.

H.4 Results Based on Alternative Migrant Networks

Table H.3: Results Based on the 2005 Migrant Network

Dep. variable: Prob. importing product p for individual i
without relatives in the U.S. and who belongs to network b at time t

	%Δ w.r.t. mean import probability		
	Neighbors	Coworkers	Friends
	(1)	(2)	(3)
$\widehat{\text{ShareImporters}}^{US\,direct}_{bp,t-1}$	13.555	16.686	13.666
	(5.269)***	(7.939)**	(6.341)**
F-stat first stage	26.03	10.21	6.16
Observations	286,671,727	294,837,860	265,362,703
Clusters	193,986	221,050	4,655,580

Notes: The top panel shows the results of running equation (6) based on the 2005 (proxy) migrant network. Robust standard errors, adjusted for clustering by network-product (HS-4), are in parentheses. Coefficients are percentages of mean outcome for a one standard deviation change in the instrumented share. The regression controls for network×product (HS-4), network×time, and individual fixed-effects. Data is quarterly and spans 2015-2019.

Table H.4: Results Excluding the New York Consulate

Dep. variable: Prob. importing product p for individual i
without relatives in the U.S. and who belongs to network b at time t

	%Δ w.r.t. mean import probability		
	Neighbors	Coworkers	Friends
	(1)	(2)	(3)
$\widehat{\text{ShareImporters}}^{US\,direct}_{bp,t-1}$	18.496	18.443	12.585
	(6.063)***	(19.457)	(4.909)***
F-stat first stage	32.00	1.08	13.90
Observations	290,557,732	123,518,595	209,186,860
Clusters	176,842	114,251	3,535,108

Notes: The table shows the results of running equation (6) and re-sampling only among networks linked via relatives to consulates *other than New York*. Robust standard errors, adjusted for clustering by network-product (HS-4), are in parentheses. Coefficients are percentages of mean outcome for a one standard deviation change in the instrumented share. The regression controls for network×product (HS-4), network×time, and individual fixed-effects. Data is quarterly and spans 2015-2019.

I Determinants of Product Propagation

I.1 Dynamic vs. Established Products

Table I.1: Strength of Externalities According to Products' Dynamism

Dep. variable: Prob. importing product p for individual i
without relatives in the U.S. and who belongs to network b at time t

	%Δ w.r.t. mean import probability					
	Jobs by New Establishments			Entry of Establishments		
	Neighbors	Coworkers	Friends	Neighbors	Coworkers	Friends
	(1)	(2)	(3)	(4)	(5)	(6)
$\widehat{Dynamic_p \times \text{ShareImp}}^{US\,direct}_{bp,t-1}$	23.171	21.925	116.196	21.572	30.037	35.216
	(10.440)**	(18.950)	(271.858)	(11.852)*	(17.773)*	(16.842)**
$\widehat{\text{ShareImporters}}^{US\,direct}_{bp,t-1}$	-3.225	7.539	-95.266	-4.907	1.509	-15.589
	(7.186)	(7.888)	(267.024)	(10.317)	(8.526)	(15.728)
SW F – interaction	28.62	10.56	0.21	45.48	10.75	7.79
SW F	8.03	7.75	0.15	6.43	6.12	2.17
Stock-Yogo 10% critical val.		7.03			7.03	
Stock-Yogo 15% critical val.		4.58			4.58	
Observations	286.8M	297.2M	259.4M	286.8M	297.2M	259.4M
Clusters	195,639	231,103	4.5 M	195,639	231,103	4.5M

Notes: The table shows the results of running equation (6), where the IV is interacted with an indicator equal to one if the good is classified as dynamic; the interaction is then instrumented. Columns (1)-(3) classify a product as dynamic if the creation of jobs by new establishments is above the median of the sample, while columns (4)-(6) classify a product as dynamic if the entry of new establishments is above the median. Robust standard errors, adjusted for clustering by network-product (HS-4), are in parentheses. The independent variable is standardized. The value of the Sanderson and Windmeijer (2016) conditional first-stage F-statistics (SW F) for the validity of the instruments is also reported, along with the corresponding Stock-Yogo critical values for a perfectly identified model with two endogenous variables.

I.2 Asymmetric Response to Positive and Negative Shocks

Table I.2: Effect of Positive and Negative Changes in Exposure

Dep. var.: Prob. importing product p in b at t for importers with relatives in the U.S.

	%Δ w.r.t. mean import probability		
	Neighbors	Co-workers	Friends
	(1)	(2)	(3)
$Positive_{bp,t-1} \times \ln \widetilde{E}_{bp,t-1}$	33.527	55.784	47.125
	(6.903)***	(71.127)	(9.974)***
$\ln \widetilde{E}_{bp,t-1}$	-3.413	28.178	-9.170
	(2.721)	(24.768)	(3.038)***
F-stat first stage	20.31	5.97	11.60
Observations	289,340,892	299,920,162	260,952,672
Clusters	200,308	236,804	4,568,240

Notes: The table shows our first-stage regression interacting exposure with an indicator equal to one if the residual is positive. Robust standard errors, adjusted for clustering by network-product (HS-4), are in parentheses. Mean import probabilities are reported. Regressions control for network×product (HS-4), network×time, product×time fixed-effects. Appendix E.5 details the sample of products per regression.

I.3 Centrality, Demographics, Visibility, and Premium Goods

Column (1) of Table I.3 displays results according to the centrality of importers with relatives in the U.S. Centrality is defined as degree centrality using our app-based friendship measure. The indicator $\widehat{Centrality}_{bt} = 1$ if the average degree centrality in b is above the median across network types.[94] While noisy, results suggest that the more central the importers in the first stage, the stronger the propagation across the network in the second stage. Column (2) reports heterogeneous results based on whether the product is visible or not. Networks correspond with neighborhoods.

Table I.5 leverages demographics. In column (1), we construct an indicator $Low\,Income_{bt} = 1$ if those with a relative in the U.S. in neighborhood b have a wage income below the sample median. As shown, and consistent with our discussion on distributional effects, products are more likely to propagate in a high income environment. Column (2) considers an indicator equal to one if the share of those with a relative in the U.S. in neighborhood b who are male is above the sample median. Consistent with how popular women's clothing is among imported products, we find significantly more diffusion in neighborhoods where directly connected individuals include more *women*. We construct a similar indicator but depending on age, which equals one if the neighborhood has relatively young directly individuals. The estimation is noisy as there is limited variation across neighborhoods.[95]

[94] Recall that networks of friends are time-invariant (Appendix D), and so is the centrality measure per individual; the measure per neighborhood changes over time as people move.

[95] Coefficients of variation for Age, $Female$, and $Income$ are 16%, 38%, and 68%, respectively. We also explored if there was a stronger response in the last quarter of the year, and found that, conditional on the import levels, the effect is not stronger in Q4.

Table I.3: Strength of Externalities, Importer's Centrality, and Visibility

Dep. variable: Prob. importing product p for individual i without relatives in the U.S. and who belongs to network b at time t

	%Δ w.r.t. mean import probability	
	Centrality	Visibility
	(1)	(2)
$Centrality_{bt} \times \widehat{\text{ShareImporters}}^{US\,direct}_{bp,t-1}$	17.241	
	(10.260)*	
$Non\text{-}visible_p \times \widehat{\text{ShareImporters}}^{US\,direct}_{bp,t-1}$		-31.534
		(14.112)**
$\widehat{\text{ShareImporters}}^{US\,direct}_{bp,t-1}$	2.996	18.978
	(7.338)	(5.839)***
SW F – interaction	37.95	34.36
SW F	20.72	40.91
Stock-Yogo 10% critical val.	7.03	
Stock-Yogo 15% critical val.	4.58	
Observations	289,340,892	289,340,892
Clusters	200,308	200,308

Notes: The table runs equation (6) interacting the main independent variable with average degree centrality (column (1)) and visibility (column (2)); interactions are then instrumented. Networks are defined as neighborhoods. Robust standard errors, clustered by network-product (HS-4), are in parentheses. Coefficients are percentages of mean outcome for a one standard deviation change in the instrumented share. Regressions control for network×product (HS-4), network×time, and individual fixed-effects. We report the Sanderson and Windmeijer (2016) conditional first-stage F-statistics (SW F) for the validity of the instruments and the Stock-Yogo critical values for a perfectly identified model with two endogenous variables.

Table I.4: Strength of Externalities and Premium Products

Dependent variable: Prob. importing product p in network b at time t for non-relative i (column 1) and for retailer r (column 2)

	%Δ w.r.t. mean import probability	
	Neighbors	Retail Firms
	(1)	(2)
$Premium \times \widehat{\text{ShareImporters}}^{US\,direct}_{bp,t-k}$	15.719	3.006
	(6.369)**	(0.995)***
$\widehat{\text{ShareImporters}}^{US\,direct}_{bp,t-k}$	2.446	6.892
	(5.729)	(0.969)***
SW F – interaction	75.36	1963.39
SW F	15.96	391.81
Stock-Yogo 10% critical val.	7.03	
Stock-Yogo 15% critical val.	4.58	
Observations	289,340,892	97,500,332
Clusters	200,308	2,187,504
$b\tilde{p}$, bt, i FE	Yes	-
$b\tilde{p}$, bt, r FE	-	Yes

Notes: The table runs equation (6) (in column (1)) and equation (10) (in column (2)), interacting the independent variable with a premium dummy; the interaction is then instrumented. A premium product has a price per kg above the median of its HS-4 category. Subscript $k = 1$ for column (1) and $k = 2$ for column (2). Robust standard errors are clustered by network-HS-4 in column (1) and by retailer-HS-4 in column (2). Coefficients are percentages of mean outcome for a one standard deviation change in the instrumented share. All regressions control for network×HS-4 and network×time fixed-effects, column (1) also has individual fixed-effects and column (2) retailer fixed-effects. We report the value of the Sanderson and Windmeijer (2016) conditional first-stage F-statistics (SW F) for the validity of the instruments and the Stock-Yogo critical values for a perfectly identified model with two endogenous variables.

Table I.5: Heterogeneity in Propagation by Demographics

Dep. variable: Prob. importing product p for individual i without relatives in the U.S. and who belongs to network b at time t

	%Δ w.r.t. mean import probability		
	Low Income	Gender	Age
	(1)	(2)	(3)
$\widehat{\textit{Low Income}_{b,t-1} \times \text{ShareImporters}^{US\,direct}_{bp,t-1}}$	-18.245		
	(8.859)**		
$\widehat{\textit{Male}_{b,t-1} \times \text{ShareImporters}^{US\,direct}_{bp,t-1}}$		-16.275	
		(8.833)*	
$\widehat{\textit{Young}_{b,t-1} \times \text{ShareImporters}^{US\,direct}_{bp,t-1}}$			-3.554
			(10.202)
$\widehat{\text{ShareImporters}}^{US\,direct}_{bp,t-1}$	15.332	18.613	15.503
	(5.582)***	(6.994)***	(7.671)**
SW F – interaction	19.55	37.84	33.94
SW F	31.99	27.78	17.08
Stock-Yogo 10% critical val.		7.03	
Stock-Yogo 15% critical val.		4.58	
Observations	289,340,892	289,340,892	289,340,892
Clusters	200,308	200,308	200,308

Notes: The table runs equation (6) interacting the main independent variable with measures of income (column (1)), share of women (column (2)), and age (column (3)) of those with a relative in the U.S. Networks are defined as neighborhoods. Robust standard errors, adjusted for clustering by network-HS-4, are in parentheses. Coefficients are percentages of mean outcome for a one standard deviation change in the instrumented share. Regressions control for network×product (HS-4), network×time, and individual fixed-effects. The value of the Sanderson and Windmeijer (2016) conditional first-stage F-statistics (SW F) for the validity of the instruments is also reported, along with the corresponding Stock-Yogo critical values for a perfectly identified model with two endogenous variables.

J Retailers' Response: Additional Results

J.1 Retailer-Specific Gravity Zones

Equation (10) includes $ShareImporters^{US,\,direct}_{fpt}$ as an explanatory variable, which ideally should consider imports (and exposure) of individuals with relatives in the U.S. *and who reside within firm f's catchment area.* We proceed in steps to understand which neighborhoods belong to each retailer's catchment area, which we call the retailer's *gravity zone*.

First, we leverage information from electronic invoices to estimate retailer-specific gravity zones. In Costa Rica, electronic invoices are digital documents used to record sales transactions in compliance with tax regulations. Businesses are required to issue electronic invoices for all taxable transactions to simplify tax reporting and reduce evasion. Retailers, in particular, issue one invoice per sale. This invoice includes the retailer's unique ID, and for a significant share of all sales, it also includes the unique ID of the *final customer who purchased the good*—recall that these IDs are all pseudonymous, but they can be linked across datasets. Not all invoices include this detail, but about one-fifth do, as it allows for better tracking of transactions and detailed records for both businesses and consumers. Thus, businesses encourage customers to provide their ID for each sale, and customers have the added value of keeping an electronic record and invoice of each transaction in their email. Second, each customer's ID in this dataset is mapped to the corresponding neighborhood where she resides.

Third, $ShareImporters_{fpt}^{US,direct}$ is constructed as a weighted average across all neighborhoods where retailer f's customers reside, with the weights based on the proportion of the retailer's total sales to customers living in each location. Similarly, it is instrumented based on an exposure measure which depends on a customer-weighted average across neighborhoods.

Approximation for *All* Retailers The above procedure has the advantage of using observed data on customers' locations. However, it faces a challenge: electronic-invoice data is only available *after* 2020, while the sample period for the estimation is 2015 to 2019.[96] Therefore, there are about 40% firms which appear in our estimation sample, but for which there are no electronic voucher data with customer IDs—this could happen, for example, because they exited or because none of their clients provided their ID. Not to lose these firms, we explore an alternative procedure which relies on an approximation of customers' location to identify gravity zones: we rely on *employees'* neighborhoods of residence to construct our weighted averages. This proxy is remarkably good; in fact, the correlation in the exposure measures constructed via customers' residences and via employees' residences is **0.98**. Given this high correlation, it should not be surprising that estimations with either measure are statistically equal, as discussed in Section 5 and as reflected in Tables 6 and J.6.

Multi-Establishment Retailers Multi-establishment retailers can be naturally accommodated in our estimation framework. The $ShareImporters_{fpt}^{US,,direct}$ for these firms is calculated as an average across all its gravity zones—which, as mentioned above, can depend on the neighborhoods where the retailer has customers or where its employers live. We do not force these zones to be joint; they can consist of disjoint or geographically distant areas, thereby not posing a challenge for multi-establishment firms. Note that establishment-level estimations would not be useful, as customs data indicate a firm's total imports but do not disaggregate those imports by establishment. Another decision that must be made for multi-establishment retailers in equation (10) is the selection of the neighborhood for the fixed effects. We opt to use the neighborhood where the firm has the most employees, arguably its largest establishment—a choice that is innocuous under an assumption of uniform assortment.

[96]In particular, we rely on 2023 electronic-invoice data, as it is the year (prior to 2024) for which the largest share of invoices include a customer ID.

Table J.6: Supply Response—Gravity Zones Based on Customers' Location

Dependent variable: Prob. of retailer f importing product p at time t

	%Δ w.r.t. mean import probability			
	All Retailers		Small Retailers	Large Retailers
	(1)	(2)	(3)	(4)
$\widehat{\text{ShareImporters}}^{US\,exposure}_{bp,t-2}$	7.493	10.084	8.074	6.136
	(0.861)***	(0.870)***	(1.023)***	(1.531)***
$LowVisibility_p \times \widehat{\text{ShareImporters}}^{US\,exp}_{bp,t-2}$		-12.830		
			(2.376)***	
F-stat first stage	750.36		641.14	344.00
SW F – interaction		880.08		
SW F		954.94		
Stock-Yogo 10% critical value		7.03		
Stock-Yogo 15% critical value		4.58		
Observations	52,490,096	52,490,096	48,478,231	4,010,988
Clusters	998,496	998,496	937,521	87,828
Mean dependent variable	0.03	0.03	0.02	0.09
$b\tilde{p}$, bt, f FE	Yes	Yes	Yes	Yes

Notes: Retailer-specific gravity zones are constructed based on the residence of each retailer's customers. Robust standard errors, adjusted for clustering by retailer-product (HS-4), are in parentheses. Coefficients are percentages of mean outcome for a one standard deviation change in the instrumented share. Regressions control for neighborhood×product (HS-4), neighborhood×time, and retailer fixed-effects. Percentage mean import probabilities are reported. The value of the Sanderson and Windmeijer (2016) conditional first-stage F-statistics (SW F) for the validity of the instruments is reported in columns (2) and (3), along with the Stock-Yogo critical values for a perfectly identified model with two endogenous variables. Appendix E.5 presents details on the sample used in each regression. Data is quarterly and spans 2015-2019.

J.2 Complementary Figures and Tables

Table J.7: Retailers' Response, Total Imports as the Endogenous Variable

Dependent variable: Prob. of retailer f importing product p at time t

	%Δ w.r.t. mean import probability		
	All Retailers	Small Retailers	Large Retailers
	(1)	(2)	(3)
$\widehat{\text{TotalImports}}^{US}_{fp,t-1}$	9.039	9.803	6.163
	(0.670)***	(0.760)***	(1.414)***
F-stat first stage	9,892	9,150	1,557
Observations	97,492,988	92,572,530	4,917,304
Clusters	2,188,011	2,113,638	115,261
Mean dependent variable	0.03	0.02	0.09
$b\tilde{p}$, bt, f FE	Yes	Yes	Yes

Notes: Robust standard errors, adjusted for clustering by retailer-product (HS-4), are in parentheses. The endogenous variable considers total imports by individuals, instrumented with shocks emerging from the migrants network. This independent variable is standardized. Regressions control for neighborhood×product (HS-4), neighborhood×time, and retailer fixed-effects. Percentage mean import probabilities are reported.

Table J.8: Retailers' Imports from Any Country and Exposure

Dependent variable: Prob. of retailer r importing product p in neighborhood b at time t

	%Δ w.r.t. mean import probability
$\overbrace{\text{ShareImporters}}^{US\,exposure}_{bp,t-2}$	10.234
	(0.919)***
F-stat first stage	171.1
Observations	97,500,332
Clusters	2,187,504
Mean dependent variable	0.147
$b\tilde{p}$, bt, r FE	Yes

Notes: Robust standard errors, adjusted for clustering by retailer-product (HS-4), are in parentheses. The independent variable is standardized. Percentage mean import probabilities are reported. The regression includes neighborhood×product (HS-4), neighborhood×time, and retailer fixed-effects.

J.3 Mechanism: Gravity Zones and Employer-Employee Data

We can leverage our estimated retailer-specific gravity zones and the employer-employee data to better understand the mechanism behind retailers' response to individual-level imports. The idea behind this exercise is that employees can be exposed to foreign products in their neighborhoods and transmit information about the existence of these products to their employers. However, if employees live in areas which are relatively far away from the retailer and outside of the gravity zone where its customers live, they should not be able to speak about the particular level of the local demand that their employer will face. This strategy exploits that there is an imperfect overlap between a retailer's catchment area and the residence of its employees.

We construct measures of exogenous exposure by employees depending on the exposure faced in the neighborhoods where they reside, and proceed by constructing two separate measures. First, one focusing on employees who live outside of the retailer's estimated gravity zone (denoted by GZ; see Appendix J.1 for details on how we construct each catchment area) and a second one focusing on employees who reside in *municipalities* (denoted by M) other than where their employer is located. Namely, the exposure of firm f is an employee-weighted average across the neighborhoods where its employees reside, restricting attention to those located *outside* the firm's gravity zone GZ_f or its municipality M_f.:

$$\widetilde{E}^{L\notin GZ}_{f,bpt} = \sum\nolimits_{h\notin GZ_f} \frac{L_{fht}}{\sum_{j\notin GZ_f} L_{fjt}} \widetilde{E}_{hpt} \quad \text{and} \quad \widetilde{E}^{L\notin M}_{f,bpt} = \sum\nolimits_{h\notin M_f} \frac{L_{fht}}{\sum_{j\notin M_f} L_{fjt}} \widetilde{E}_{hpt},$$

where b is the neighborhood of retailer f, h denotes employees' residential neighborhoods, L_{fht} is the number of employees of f who reside in neighborhood h at time t, and $\widetilde{E}_{hpt}$ denotes exposure. The resulting variable therefore captures the exposure to product p faced by employees of retailer f who reside "far away" from the firm's catchment area. We then consider, for imports of p by retailers in b at time t,

$$\text{Import}_{f,bpt} = \delta + \kappa \widetilde{E}^{L,far}_{f,bpt-2} + \zeta \widetilde{E}^{US\,direct}_{f,bpt-2} + \gamma_{b\tilde{p}} + \gamma_{bt} + \gamma_f + \varepsilon_{f,bpt},$$

where $\text{Import}_{f,bpt} = 1$ if retailer f in neighborhood b imports product p at t for the first time. We include as an independent variable the firm's exposure from their employees who reside "far away," which can be defined alternatively as explained above. We also control for firm exposure and include a battery of fixed effects.

As shown in Table J.9, retailers do not show a meaningful response to the exposure of employees living far away, regardless of whether we define outsiders based on gravity zone or municipality. The latter aligns with firms learning about the level of the local demand for a product, as opposed to just a product-discovery story. In line with these results, retailers in out survey were five times more likely to gather insights from employees living close by than far away.

Table J.9: Retailers' Imports and Exposure of Employees Living Far Away

Dependent variable: Prob. of retailer r importing product p in neighborhood b at time t

	%Δ w.r.t. mean import probability	
	Gravity Zones (1)	Municipalities (2)
$\widetilde{E}^{emp,far}_{bp,t-2}$	0.672	0.601
	(1.185)	(1.021)
$\widetilde{E}^{US\,exposure}_{bp,t-2}$	7.450	9.111
	(0.825)***	(0.670)***
F-statistic	42.37	93.09
Observations	52,748,612	98,044,144
Clusters	999,887	2,192,746
Mean dependent variable	0.029	0.147
$b\tilde{p}$, bt, f FE	Yes	Yes

Notes: Robust standard errors, adjusted for clustering by retailer-product (HS-4), are in parentheses. Coefficients are percentages of mean outcome for a one standard deviation change in the instrumented share. Percentage mean import probabilities are reported. The regression includes neighborhood×product (HS-4), neighborhood×time, and retailer fixed-effects.

J.4 Benchmarking Retailer Learning Against Advertising

The information and advertising channels are conceptually distinct. Our "information channel" operates through retailers learning which products are most successful among consumers. By contrast, the "advertising channel" operates by informing or persuading consumers about products that retailers already offer. While both channels can increase sales, they do so through different mechanisms and require retailers to undertake very different activities. Moreover, whereas advertising typically involves firms promoting products over which they have some form of ownership or control, retailers benefiting from the information channel generally do not have exclusivity over the products they learn about. Thus, retailers have limited incentives to subsidize individual imports or experimentation for a product, as other retailers can free-ride on the information revealed through the network.

We first quantify the magnitude of the retailer information channel to assess its economic relevance and benchmark it against advertising. From equation (11), we infer that local retailers' imports of a new product increase by \$191 for every dollar exogenously imported by an individual directly connected to the U.S. Customs data further indicate that the median product imported by an individual is valued at \$30. Combining these figures, the median exogenous importation event increases retailer sales by 191×30, where we assume the import response translates one-for-one into incremental retail sales. Then, let m denote the retailer's net profit margin, defined as profits divided by sales. We set $m = 0.2$, consistent with Broda and Weinstein (2006) and with the median net margin in the distribution of Costa Rican retailers. We estimate this margin using administrative data covering the universe of retail invoices, which report both retail prices and wholesale costs at the HS-10 level. Although available only after 2020, these data provide the most accurate measure of retail net margins in Costa Rica and allow us to focus on the subset of goods imported by individuals. Under this definition, the implied increase in retailer profits from a median exogenous importation event equals

$$\Delta\Pi = m \times (191 \times 30).$$

This calculation implies retailer profits of approximately \$1,146 for the median exogenous importation event. This is an upper bound on what a retailer would be willing to pay one of these "influencers" with U.S. connections for product discovery. As shown in the main text, this implied willingness to pay is driven by smaller retailers. Why, then, do such payments not commonly occur in practice? Retailers typically lack property rights over the information generated by experimentation: the information is non-excludable, giving rise to a free-rider externality across retailers.

To benchmark this estimate against advertising costs, we distinguish between traditional advertising (e.g., TV, radio, magazines) and influencer-based advertising, where brands collaborate with individuals to promote their products.

J.4.1 Traditional Advertising

We begin by recalling three key facts from the recent literature. First, advertising and marketing are expensive activities for firms. Argente et al. (2025) study advertising behavior in the U.S. retail sector and show that relatively few firms advertise, and those that do are positively selected (for example, only 12% of brands advertise and they account for 57% of sales). Argente et al. (2025) show that advertising is concentrated among firms above the 80th percentile of the size distribution. Second, Argente et al. (2025) also show that advertising is often intermittent, suggesting that advertising is not necessary for sales and that there are important fixed costs of advertising. Third, conditional on advertising, the elasticity of sales to advertising is low. Lodish et al. (1995) document an average advertising elasticity of 0.05 for established products. More recent estimates by Shapiro et al. (2021) are lower. They report the mean (median) of the distribution of long-run own advertising elasticities is approximately 0.023 (0.014), and a negative marginal ROI for over 80% of brands.

Consistent with the idea that advertising is available mainly to large, successful brands, a recent eMarketer report finds that two companies (Walmart and Amazon) accounted for 84.2% of retail-media digital ad spending in 2024.[97] The report also estimates that only 0.7% of retail ad spending is allocated to physical retail channels typically used by small retailers. By contrast, our information channel delivers its largest gains for small retailers, who are largely absent from these advertising markets.

To clarify and compare economic magnitudes between the information channel and the advertising channel, suppose a retailer captures \$1,146 per year in profit from the identification of a successful product (the "information channel"). Reproducing the same profit via advertising requires additional annual sales $\Delta\text{Sales} = \frac{\$1{,}146}{m}$, where m denotes the retailer's net margin. Let ε be the advertising elasticity (the percent change in sales induced by a 1% increase in advertising exposure). The required increase in advertising intensity is

$$\%\Delta\text{Adv} = \frac{\%\Delta\text{Sales}}{\varepsilon} \qquad \text{and} \qquad \%\Delta\text{Sales} = \frac{\Delta\text{Sales}}{\text{Sales}} \times 100\%.$$

Using elasticities from the literature ($\varepsilon \approx 0.023$ and $\varepsilon \approx 0.05$), a small retailer with annual sales of \$200,000 and a 20% margin, reproducing a \$1,146 profit gain would require an increase in ad exposure of approximately 57–125% (for $\varepsilon = 0.05$ and $\varepsilon = 0.023$, respectively). For a larger retailer with annual sales of \$1,000,000, the corresponding increases are approximately 11–25%. Even under high margins, these are substantial increases in advertising intensity for the typical small retailer.

Taken together, these calculations show that the retailer information channel we document is economically meaningful—particularly for smaller retailers that cannot cheaply replicate the benefits of product discovery via advertising. Advertising is an important activity for large, successful firms; they can use it to boost sales and, consistent with our evidence, are therefore less dependent on the information channel.

J.4.2 Influencer-Based Advertising

The empirical literature on influencer-based advertising is considerably less developed than the work on traditional media. A first challenge is measurement: many influencers are compensated partly or entirely in kind (free products, exclusive access, etc.), so the relevant costs are not fully recorded in firms' advertising accounts. This makes it difficult to construct standard measures of advertising intensity and to estimate sales elasticities in the spirit of Lodish et al. (1995) or Shapiro et al. (2021). Survey evidence in Hub (2023) further shows that compensation modes are

[97] See https://www.emarketer.com/content/what-advertisers-retailers-need-know-about-retail-media-2025.

heterogeneous and that a substantial share of campaigns involve non-cash payments. Moreover, the effectiveness of a given dollar depends heavily on the influencer's follower network, so that a single "influencer elasticity" is often not well defined. As a result, existing studies frequently measure outcomes in terms of impressions, engagement, or click-through rates rather than directly estimating sales or profit elasticities.

Recent academic work begins to fill this gap, but it also emphasizes how little is known about the full return on influencer spending. For instance, Beichert et al. (2024) review the prior marketing literature and note that most studies do not jointly account for both the cost of influencer campaigns and the revenue they generate. Their main contribution is to propose and implement a "return on influencer spend" (ROIS) metric that explicitly relates campaign-induced revenues to influencer compensation. Even in that setting, however, ROIS estimates are highly heterogeneous across brands and influencer types, and systematic evidence on profit-based ROIs—comparable to the long-run advertising elasticities discussed above—remains scarce.

Influencer marketing is particularly attractive to small brands and direct-to-consumer (DTC) firms, for whom large-scale traditional advertising is often impractical or too costly. Consistent with this, Hub (2023) report that 43% of brands in their sample spend less than $10,000 per year on influencer marketing. Influencer campaigns operate by activating a word-of-mouth mechanism: retailers pay individuals with access to consumer networks to talk about their products, much as individuals in our setting transmit information about products to relatives and acquaintances.

To benchmark the magnitude of our estimated information channel, we ask how much paid influencer spending would be required to generate the same profit gain as a single exogenous importation event. A related practical reason to frame our benchmark in event terms is that, as explained earlier, influencer marketing is typically organized and evaluated at the level of discrete campaigns or posts, rather than as a continuously adjusted advertising "intensity." More precisely, Beichert et al. (2024) define return on influencer spend as ROIS = (Revenue − Cost)/Cost and report a mean ROIS of 8.55. This implies that one dollar of influencer spending is associated with $1 + \text{ROIS} = 9.55$ dollars in revenue. Under a constant net margin m, expected net profit per dollar of influencer spending is therefore $m \times 9.55 - 1$. With $m = 0.2$, this equals 0.91. Thus, to generate the same profit gain as the information channel, an influencer-based strategy would require influencer spending of approximately

$$\text{Spend}^{\text{influencer}} \approx \frac{\Pi^{\text{info}}}{m \times 9.55 - 1}.$$

Using our calibration of $\Pi^{\text{info}} \approx \$1{,}146$ for the median product, this implies roughly $1,260 in influencer expenditures to match the profit gain from one exogenous import event. By contrast, the information channel delivers the same incremental profit without any direct advertising outlay by retailers.

In sum, advertising involves firms paying to push information about a chosen product toward consumers, aiming to shape awareness and demand. In our setting, information flows in the opposite direction: retailers learn which products are most valued by observing consumers' purchases and use this information to adjust their assortments. Crucially, this learning process does not require additional monetary outlays by retailers, and the non-excludable nature of the information discourages such spending, because any one retailer's investment would be easily free-ridden upon by others. By contrast, influencer campaigns require product selection ex ante and explicit, typically recurring expenditures. Both mechanisms can increase sales, but they do so through distinct informational channels. Quantitatively, our calibration implies that a single exogenous importation event by a migrant-connected consumer generates an economically meaningful profit gain of about $1,146 for the median product. Replicating this gain through either traditional or influencer advertising would require sizable spending and, for typical small retailers, very large percentage increases in advertising exposure. Taken together with the non-excludable nature of the information, these magnitudes help explain why we rarely observe retailers directly paying migrants for product experimentation and underscore that the demand externalities we document are quantitatively important.

J.5 Survey of Retailers

Table J.10: Product Assortment by Retailer

Assorted Products (13 digits)				Assorted Categories (HS-2)			
Average (1)	25th Percentile (2)	Median (3)	75th Percentile (4)	Average (5)	25th Percentile (6)	Median (7)	75th Percentile (8)
271	9	29	334	15	3	8	27

Notes: The table shows details on the distribution of products assorted by the retailers in our survey's sample. Columns (1)-(4) consider the most disaggregated product category available in our data—13-digit codes available from electronic vouchers emitted by each retailer (see Appendix J.1 for details on these data). Columns (5)-(8) display statistics on assortment by HS-2 product category.

Survey Instrument We report the (translated) questions to retailers.

1. When deciding which products to have for sale, have you ever received feedback from your customers about which products to stock?

   ( ) Yes ( ) No

   ( ) My company does not sell physical products to consumers

   ( ) Don't know/No response

2. Suppose there is a product that is not available in Costa Rica. If your potential customers start buying the product from abroad over the internet, the likelihood that your company will start importing and selling that product locally

   ( ) Would increase ( ) Would decrease

   ( ) Would not change ( ) Don't know/No response

3. If you decided to start importing a new product, how long would it take from the moment you make the decision to having the product for sale?

   ( ) 0-1 quarter ( ) 2 quarters

   ( ) 3 quarters ( ) 4 quarters or more

   ( ) Don't know/No response

4. List from 1 to 3 the main mechanisms by which the company would become aware that consumers are excited about a new product that is for sale abroad but is not yet available in the country. List a maximum of three options, with 1 being the most important.

   ____ Customers ask about the product in the store

   ____ Market study conducted by the company at the local level

   ____ Market study conducted by the company at the national level

   ____ Consultation with friends, neighbors, family members living near the store

   ____ Consultation with employees living near the store

   ____ Consultation with employees living far from the store

   ____ Social media

   ____ Don't know/No response

Supplementary References

Acosta, M. and Cox, L. (2019). The Regressive Nature of the U.S. Tariff Code: Origins and Implications. Technical report, Working Paper, Columbia University.

Adão, R., Kolesár, M., and Morales, E. (2019). Shift-Share Designs: Theory and Inference. *The Quarterly Journal of Economics*, 134(4):1949–2010.

Alvarez, F., Argente, D., Lippi, F., Méndez, E., and Van Patten, D. (2023). Strategic Complementarities in a Dynamic Model of Technology Adoption: P2P Digital Payments.

Argente, D., Fitzgerald, D., Moreira, S., and Priolo, A. (2025). How do firms build market share? the role of demand frictions. *Amer. Econ. Rev.: Insights*.

Bailey, M., Gupta, A., Hillenbrand, S., Kuchler, T., Richmond, R., and Stroebel, J. (2021). International trade and social connectedness. *Journal of International Economics*, 129:103418.

Beichert, M., Bayerl, A., Goldenberg, J., and Lanz, A. (2024). Revenue generation through influencer marketing. *Journal of Marketing*, 88(4):40–63.

Broda, C. and Weinstein, D. E. (2006). Globalization and the Gains from Variety. *The Quarterly Journal of Economics*, 121(2):541–585.

De Giorgi, G., Frederiksen, A., and Pistaferri, L. (2019). Consumption Network Effects. *The Review of Economic Studies*, 87(1):130–163.

Hub, I. M. (2023). The state of influencer marketing. *Benchmark Report. The State of Influencer Marketing Benchmark Report 2023*.

Im, K. S., Pesaran, M. H., and Shin, Y. (2003). Testing for unit roots in heterogeneous panels. *Journal of Econometrics*, 115(1):53–74.

Imbens, G. W. and Wooldridge, J. M. (2009). Recent Developments in the Econometrics of Program Evaluation. *Journal of Economic Literature*, 47(1):5–86.

Jordà, Ò., Schularick, M., and Taylor, A. M. (2020). The effects of quasi-random monetary experiments. *Journal of Monetary Economics*, 112:22–40.

Jordà, (2005). Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review*, 95(1):161–182.

Levin, A., Lin, C.-F., and Chu, C.-S. J. (2002). Unit root tests in panel data: asymptotic and finite-sample properties. *Journal of Econometrics*, 108(1):1–24.

Lodish, L. M., Abraham, M., Kalmenson, S., Livelsberger, J., Lubetkin, B., Richardson, B., and Stevens, M. E. (1995). How T.V. Advertising Works: A Meta-Analysis of 389 Real World Split Cable T.V. Advertising Experiments. *Journal of Marketing Research*, 32(2):125–139.

Montiel Olea, J. L. and Plagborg-Møller, M. (2021). Local Projection Inference Is Simpler and More Robust Than You Think. *Econometrica*, 89(4):1789–1823.

Rauch, J. E. (1999). Networks versus markets in international trade. *Journal of international Economics*, 48(1):7–35.

Sanderson, E. and Windmeijer, F. (2016). A weak instrument F-test in linear IV models with multiple endogenous variables. *Journal of Econometrics*, 190(2):212–221.

Shapiro, B. T., Hitsch, G. J., and Tuchman, A. E. (2021). Tv advertising effectiveness and profitability: Generalizable results from 288 brands. *Econometrica*, 89(4):1855–1879.

Wooldridge, J. M. (2002). *Econometric Analysis of Cross Section and Panel Data*. MIT Press. Cambridge, MA.