Retail Globalization and Household Welfare: Evidence from Mexico

David Atkin

Massachusetts Institute of Technology and National Bureau of Economic Research

Benjamin Faber

University of California Berkeley and National Bureau of Economic Research

Marco Gonzalez-Navarro

University of California Berkeley

The arrival of global retail chains in developing countries is causing a radical transformation in the way households source their consumption. This paper draws on a rich collection of Mexican microdata to estimate the effect of foreign supermarket entry on household welfare and decomposes this effect into six channels. We find that foreign entry causes large welfare gains for the average household predominantly driven by a reduction in the cost of living—both through price reductions at domestic stores and through the direct consumer gains from foreign stores. These gains are, on average, positive for all income groups but are regressive.

I. Introduction

A radical transformation is occurring in the way households in developing countries source their consumption. A key driver of this so-called "su-

Zakia Akanda, Fatima Aqeel, Christopher Balette, Sakshi Kumar, Shruti Maken, Ramsha Mugal, and Fern Ramoutar provided excellent research assistance. We are grateful to Hector Guzman at Banco de México and Mauricio Varela at the University of Arizona for pro-

Electronically published January 16, 2018
[*Journal of Political Economy*, 2018, vol. 126, no. 1]

permarket revolution" has been the arrival of global retail chains in developing countries (Reardon et al. 2003; Humphrey 2007; Bronnenberg and Ellickson 2015).[1] This process of retail globalization has led to heated policy debates. Those against foreign retailers point to the large share of employment in the traditional retail sector, while those in favor emphasize potential benefits from lower consumer prices.

Importantly, these debates have also led to stark differences in policies toward retail FDI across developing countries. While some countries such as Argentina, Brazil, Mexico, and most of eastern Europe chose to fully liberalize retail FDI at the beginning of the 1990s, several developing countries including India continue to severely restrict foreign retail entry and others such as Indonesia, Malaysia, and Thailand reimposed regulatory barriers on foreign retailers after initially allowing entry (Dufey, Grieg-Gran, and Ward 2008; Wrigley and Lowe 2010).[2] These policy differences matter because retail is a key sector of the economy in terms of both employment and consumption, on average accounting for 15–20 percent of total employment, 10–15 percent of total GDP, and more than 50 percent of total household expenditure in developing countries.[3]

Despite the rapid globalization of retail in the developing world and widespread policy interest, the existing literatures in trade and develop-

viding us with the retail openings data. We also thank Rafael Posse and Javier Romero at INEGI and José Antonio Murillo at Banco de México for providing us access to the non-public-access microdata of the Mexican consumer price index (CPI). Etienne Gagnon at the Fed Board in Washington, DC, kindly shared his data on the public-access part of the Mexican CPI. Natalia Volkow at INEGI was pivotal in gaining access to the confidential microdata of the Mexican retail census. We are also grateful to the Mexico City office of a large international market research company for supporting an academic collaboration and granting us access to its uncensored bar code–level consumption microdata and to Enrique Seira Bejarano and Paul Gertler, who were instrumental in making that happen. Finally, we thank Costas Arkolakis, Dominick Bartelme, Jonathan Eaton, Jessie Handbury, Erik Hurst, Seema Jayachandran, Dan Keniston, Dmitri Koustas, David Lagakos, Natalia Ramondo, Steve Redding, Andrés Rodríguez-Clare, and participants at multiple seminars for helpful comments. This research has been supported by funding from the Ontario Work-Study Program. Data are provided as supplementary material online.

[1] Between 2000 and 2012, foreign affiliate sales of the 250 largest global retailers grew by more than 400 percent, reaching US$1 trillion (Deloitte Consulting 2014), with much of the growth in developing countries: the developing country share of world inward retail foreign direct investment (FDI) doubled to 25 percent since the 1990s and the stock grew more than 20-fold (UNCTAD World Investment Report 2013). Over the same period, i.e., since 2000, the share of traditional retail in developing country grocery expenditures decreased from 80 percent to 57 percent (Bronnenberg and Ellickson 2015).

[2] For example, it took India's Congress Party until 2012 to finally approve foreign entry into multibrand retail. Several Indian states subsequently blocked foreign entry, and most recently the Bharatiya Janata Party government announced it would move back to an outright nationwide ban of foreign retailers.

[3] Figures are based on developing country samples from the 2013 International Labour Organization (ILOSTAT) Database (employment), 2012 UN National Accounts data (GDP), and the World Bank Living Standards Measurement Study household consumption surveys from 2000–2010 (retail expenditures).

ment have so far paid relatively little attention to this facet of international integration. This paper seeks to fill that gap. We bring together a new and uniquely rich collection of microdata to assess the consequences of retail FDI in the context of Mexico, a country whose retail landscape underwent a dramatic transformation as foreign retailers came to dominate its market over the last 20 years. Our analysis coincides with the major wave of foreign store expansion in Mexico. Over our sample period, January 2002 to March 2014, the number of foreign supermarkets close to quadrupled from 365 to 1,335 stores. This rapid expansion provides an ideal empirical setting to study the impact of retail globalization.

This paper aims to contribute to our understanding of three central questions: (1) What is the effect of retail FDI on average household welfare in the municipality of entry? (2) What are the channels underlying this effect? (3) To what extent do the gains from retail FDI differ across the preexisting income distribution? In answering these questions, the paper also makes two methodological contributions to the literature that focuses on quantifying the gains from trade and FDI. The first is that rather than imposing structure ex ante to limit the data requirements to a set of readily available cross-country moments, we instead exploit newly available and extremely detailed microdata that allow us to estimate a very general expression for the welfare gains from retail FDI. In particular, data on bar code–level consumer prices and consumption quantities, worker-level incomes, and store-level profits allow us to capture all major components of household welfare without shutting down any potential channels—such as gains from variety or procompetitive effects on prices in domestic stores—ex ante. The second contribution is that, rather than relying on cross-sectional moments that may or may not capture the causal effects of integration, we propose an event study design to credibly identify the moments we feed into the welfare expression.

At the center of the analysis lies the construction of a new collection of microdata. We combine data on all foreign-owned supermarket locations and opening dates over the period 2002–14 with five additional data sets: (i) monthly store-level consumer prices at the bar code equivalent level (e.g., a 16-pill package of Bayer Aspirin with a 300-milligram dosage) from the confidential microdata of the Mexican consumer price index (CPI); (ii) daily household-by-store-level data on consumption quantities and prices at the bar code equivalent level from the consumer panel of the Mexican operation of a large international market research company; (iii) store-level revenues, costs, and profits for the universe of urban retail establishments from two cross sections of the confidential microdata of the Mexican retail census; (iv) quarterly worker-level incomes by occupation and sector from Mexico's urban employment and occupation surveys; and (v) household-level income shares by occupation and sector matched to consumption shares across products and store formats from Mexico's biannual income and expenditure surveys.[4]

The analysis proceeds in three steps. In step 1, we write down a general expression for the effect of retail FDI on household welfare in the municipality of entry. We decompose the total effect into six distinct effects: three effects on household cost of living (the price index) and three effects on household nominal incomes. On the cost-of-living side, we distinguish between the effect on consumer prices at preexisting domestic retailers (the procompetitive price effect), the effect due to exit of domestic retailers (the procompetitive exit effect), and the direct price index effect that encompasses all the consumer gains derived from being able to shop at the foreign store itself, including different prices for preexisting products, new product variety, as well as different store amenities. The nominal income effect comprises a retail labor income effect (from employment in either traditional or modern retail), a retail profits effect for domestic store owners, and an indirect effect on other sources of household income from nonretail sectors of the local economy.

In step 2, we estimate the empirical moments required to quantify the six effects that underlie the total household gains from retail FDI. To tackle the procompetitive price effect, we estimate how prices change in domestic stores in response to the first entry of a foreign supermarket in the municipality. The first empirical challenge is that the composition of goods and stores changes over time. To address this, we use the Mexican CPI microdata to construct a 12-year time series of monthly prices for bar code equivalent products sold in a particular retail outlet in a particular municipality. The second empirical challenge is nonrandom entry of foreign retailers across municipalities and over time. We propose an event study design that allows us to transparently and nonparametrically test whether foreign retailers targeted store openings toward municipalities with preexisting price trends. The store opening data suggest that, over our period of study, foreign retailers operated under the objective of rapidly establishing store presence across all of urban Mexico. If so, the timing of opening will be determined by the speed of obtaining zoning permits and building delays and be uncorrelated with location-specific changes in prices or incomes. We test this identifying assumption by estimating a full set of monthly treatment effects beginning in the years before the store opening event and continuing for several years after. In support of our assumption, we find no evidence of pretrends in these monthly treatment effects.

[4] Note that we refer to all retail establishments as "stores" in this paper even though the data include all types of retail units (e.g., street markets, convenience stores, and supermarkets).

While our data allow us to observe the consumer price changes within continuing domestic shopping outlets in order to estimate the procompetitive price effect, the two remaining cost-of-living effects are more difficult to quantify. The issue is that the implicit changes in the price index that result from either the arrival of a new foreign store (the direct price index effect) or the exit of domestic stores (the procompetitive exit effect) are inherently unobservable. To quantify the cost-of-living implications of these changes in the available consumer choice set, we require further assumptions about consumer demand to pin down "virtual prices"—the price at which demand would be zero—for foreign stores before they entered and domestic stores after they exited.

To this end, we use two different approaches. Our preferred approach is an exact estimation of the cost-of-living effect under a multitier constant elasticity of substitution (CES) preference structure in which, within a broad product group, consumers have asymmetric CES preferences over stores or, more precisely, over store-level consumption aggregates. These aggregates are themselves optimal bundles chosen from the specific products available in each store. The direct price index effect under this approach requires information on the ex post household expenditure shares on foreign stores in combination with estimates of the elasticity of substitution across local stores (in both cases at the household income group and product group levels). To obtain these estimates, we exploit the uncensored consumer panel microdata that contain prices and household consumption quantities at the bar code equivalent level matched to individual retailer identities. For the supply-side variation needed to identify the elasticity of substitution, we exploit the fact that local prices in supermarket chains are driven, at least in part, by national and regional supply shocks and pricing rules.

This CES approach has several benefits. First, as shown by Anderson, De Palma, and Thisse (1992), these preferences generate the same demands as would be obtained from aggregating many consumers who make discrete choices over which store to shop in. Second, this approach has the appeal of being widely used in the trade literature starting with Feenstra (1994), in part because it yields a very parsimonious expression for the welfare gain from new products (or stores in our case). Third, the CES approach allows us to relate our results to the recent quantitative literature on the gains from trade as the expression for our direct price index effect is identical to the well-known import share sufficient statistic for the gains from trade in Arkolakis, Costinot, and Rodríguez-Clare (2012) and extended to horizontal FDI by Ramondo and Rodríguez-Clare (2013).

While the assumption of CES preferences has its virtues, it also imposes a particular structure on household demand. As an alternative approach, we also estimate a first-order approximation of the cost-of-living effect that is solely based on observable price changes due to foreign entry.

The advantage of this alternative approach is that it yields a simple and transparent Paasche price index that approximates the consumer gains that arise from foreign store entry. The disadvantage is that, since this approach essentially assumes that the foreign stores were always present and simply lowered their prices at the time of entry, we are necessarily abstracting from any gains due to the new product and store variety provided by foreign stores or the fact they may provide different amenities to shoppers. For this reason, the difference between the direct price index effect under CES and the first-order approach is also informative as it provides an approximate estimate of the proportion of the gains that come from these variety and amenity channels as opposed to lower prices on preexisting products.

To estimate the effects on nominal household incomes, we construct a quarterly time series of individual income, occupation, sector, and employment status using the Mexican employment and occupation surveys. The identification issues are similar to those we address in the price regressions, and we follow a similar event study approach. To capture the effects on retail profits for owners of local domestic stores, as well as the effects on store exit, we complement these data with the confidential microdata on store counts and profits from the Mexican retail census.

In step 3, we combine the estimated effects on consumer prices, consumption quantities, and nominal incomes from step 2 with the theoretical framework in step 1 in order to quantify the household welfare effects of foreign entry. To do so, we require pre-entry household consumption shares across various product groups and store types, as well as labor and business income shares from various occupations and sectors. We obtain this information from the Mexican income and expenditure surveys, which allow us to estimate a predicted welfare change for each household in the sample depending on their particular work and consumption patterns (restricting attention to locations without a foreign store at the time of the survey).

We find that foreign supermarket entry causes large and significant welfare gains for the average household in the municipality of entry, equal to 6 percent of initial household income. The majority of this effect is driven by a significant reduction in the cost of living. While there is a 0.7 percent increase in the cost of living due to procompetitive exit effects, this is more than compensated by a reduction of 1.6 percent due to procompetitive effects on consumer prices charged by preexisting domestic stores and a reduction of 5.5 percent due to the direct price index effect (i.e., foreign supermarkets offering cheaper prices, new varieties, and different shopping amenities to consumers).[5] The relatively large direct effect is

[5] Note that these price index changes refer to the entirety of household consumption, accounting for the fact that retail is, on average, half of household consumption during the estimation period. We show that nonretail prices do not respond to foreign retail entry.

consistent with raw moments in the data that we present as motivating evidence: foreign retailers charge, on average, 12 percent lower prices for an identical bar code in the same location and time, offer five times the number of products compared to modern domestic stores, and after entering, capture more than one-third of total household retail spending on average. The first-order approximation of the direct price index effect is 40 percent of the size of the CES estimate, suggesting that just under half of the direct effect can be accounted for by the cheaper prices at foreign stores, with the remainder due to the additional benefits from product and store variety and differences in foreign store amenities. The nominal income effects are small in comparison. We find no effect on average municipality-level household incomes or employment rates. We do, however, find evidence of store exit and adverse effects on store profits, employment, and labor incomes for the traditional retail sector. While these adverse income effects are sizable, they affect only a fraction of households and so are swamped in the aggregate by reductions in the cost of living that benefit all households.

We also quantify the distribution of the gains from retail FDI. While, on average, all household income groups experience significant gains from foreign entry, the richest income groups gain about 50 percent more than the poorest. We find that the key driver is the fact that the richest households substitute over 50 percent of their retail consumption to foreign stores, while the poorest substitute less than 15 percent. Since the elasticity of substitution is broadly similar for both income groups, these market share differences imply that wealthier households in Mexico value the consumption choices on offer at foreign stores significantly more than poorer ones (e.g., as a result of the rich placing a higher value on foreign brands, high-quality varieties, and large pack sizes or on store amenities such as parking, car accessibility, wide aisles, security, and hygiene).

Finally, we investigate two remaining questions. First, we ask to what extent the estimated welfare gains are specific to foreign supermarket entry (FDI) rather than being driven by the entry of modern store formats more generally. We do not find procompetitive price effects or comparable direct price index effects when running identical specifications for the entry of domestic retailers with similar big-box formats to the foreign entrants. We also find that the procompetitive and direct price index effects of foreign entry are large and significant in locations with preexisting domestic big-box stores, as well as in those without. Together, these pieces of evidence suggest that the large consumer gains we estimate for Mexico are specific to retail FDI. Second, we explore the causes of the price reductions in domestic stores that we have labeled the procompetitive price effect. A back-of-the-envelope decomposition shows substantial reductions in markups but also reductions in marginal costs, suggesting the presence of local spillovers from foreign entry (e.g., domestic stores adopting the better management practices and logistics used by foreign retailers or price reductions among local suppliers).

The paper closely relates to a small body of work that explores the economic consequences of foreign supermarkets in developing countries (Javorcik and Li 2013; Iacovone et al. 2015).[6] Relative to these papers that have exclusively focused on the spillover effects on domestic suppliers, this paper instead focuses on the consequences for consumers, workers, and business owners located in the municipality where the foreign store entry occurs. To the best of our knowledge, this is the first paper to provide empirical evidence on these first-order effects of retail globalization. We note that in order to do so convincingly, this paper's focus is on quantifying the effects of foreign retail entry on local household welfare within the municipality of entry. This focus allows us to credibly estimate impacts of foreign entry by comparing the municipality of entry to other locations that did not experience a foreign store opening in the same period. The limitation of such an approach is that it is silent on potentially interesting national-level effects such as changes in manufacturing productivity that are absorbed by the time fixed effects in our empirical setting. Our work also relates to the study by Lagakos (2016), who emphasizes the role of endogenous store format choices in explaining cross-country differences in retail sector total factor productivity. Consistent with our finding of much larger gains for richer households, he finds that car ownership rates are significantly related to the adoption of modern store formats.

The paper is also closely related to the recent literature that estimates the gains from international integration for developing countries and the distribution of those gains (Porto 2006; Goldberg and Pavcnik 2007; Topalova 2010; Atkin 2013; Faber 2014; Fajgelbaum and Khandelwal 2014; Donaldson, forthcoming). Relative to the existing literature, we focus on the consequences of retail globalization, a channel of integration that has received relatively little attention. Methodologically, this paper differs in its careful empirical evaluation of all major components of household welfare and, in particular, the cost of living. Rather than relying on state-level price deflators (e.g., Deaton and Tarozzi 2000; Topalova 2010), household consumption surveys combined with simulated price changes at the product group level (e.g., Deaton 1989; Porto 2006), or cross-country trade flows (e.g., Caron, Fally, and Markusen 2014; Fajgelbaum and Khandelwal 2014), this paper draws on price and consumption data at the level of individual households, bar code equivalent products, and stores to provide a more precise and complete estimate of changes in the price index.

[6] Varela (2016) uses Walmart's local entry decisions in Mexico to estimate a structural model of diseconomies of scale in outlet expansion.

The paper is also related to the trade literature that estimates the gains from new imported product variety (Feenstra 1994; Broda and Weinstein 2006; Feenstra and Weinstein 2017). As well as drawing on these tools to estimate the cost-of-living gains from a new foreign retailer, the richness of our data allows us to directly observe foreign production shares in individual household consumption baskets at the level of disaggregated product groups. To the best of our knowledge, this is the first time such a match has been possible in order to quantify the consumer gains from international integration.

Finally, since Walmart de México is the major foreign retailer expanding during our estimation period, the paper relates to an extensive literature on the effects of Walmart in the United States (e.g., Basker 2005a; Hausman and Leibtag 2007; Jia 2008; Holmes 2011). This paper offers two main innovations relative to the existing literature. First, studying a developing country allows us to shed light on the impact of exposing a largely traditional retail environment to what is arguably the world's technological frontier in retailing. Second, in contrast to the piecemeal approach adopted by the literature to date, this paper is the first to set up a unified empirical framework and estimate the effect of store entry on both cost of living and nominal incomes.

We structure the remainder of the paper as follows. Section II describes the Mexican context and provides motivating evidence. Section III presents the theoretical framework. Section IV describes the six data sets. Section V presents the empirical strategy and estimation results. Section VI draws on these estimation results in combination with estimates of household demand parameters to quantify the gains from retail FDI. Section VII presents conclusions.

II. Background and Motivating Evidence

A. *Background*

Prior to the 1980s, retail FDI into Mexico had to be approved on a case-by-case basis and generally required a minimum of 51 percent Mexican ownership. These restrictions were gradually relaxed in the 1980s with foreign companies able to own up to 49 percent of a Mexican firm without explicit authorization. The 1993 FDI law allowed foreign firms full ownership rights and full freedom to repatriate profits. FDI was further protected with the North American Free Trade Agreement (NAFTA) third-party dispute resolution mechanisms starting in 1994.

The first significant retail FDI into Mexico was the US company Safeway's purchase of 49 percent of Casa Ley (a regional retailer in northern Mexico). More transformative was Walmart's decision to enter the Mexican market in the early 1990s as NAFTA was being negotiated. Walmart initially entered via a joint venture with the Mexican retailer Cifra, a chain from Mexico City with around 100 supermarket units at the time. In 1997, Walmart bought out Cifra and in 2000 changed the name of the company to Walmart de México (Walmex). In contrast to the United States, Walmex focused heavily on food retail and targeted relatively affluent Mexican consumers. In the ensuing years, Walmex and its multiple supermarket brands (Walmart, Sam's Club, Superama, Aurrera, and Bodega Aurrera) became the largest retail chain in Mexico, as well as Mexico's largest employer, with over 210,000 employees in January 2014.[7] Although Walmart has been the most notable foreign entrant, two large French supermarket chains also entered and subsequently left the market (Auchan and Carrefour), while several other US firms continue to operate in Mexico (Costco, HEB, S-Mart, Smart and Final, and Waldo's).

The expansion of Walmart and other foreign supermarket chains proceeded relatively slowly during the second half of the 1990s, predominantly serving the main metropolitan centers of Mexico. As depicted in figure 1, the number of foreign supermarkets in Mexico expanded from 204 stores at the end of 1995 to 365 stores at the end of 2001. In both periods, the presence of foreign stores was heavily concentrated in a handful of major cities. Between 2002 and 2014, the sample period of our empirical analysis, the number of foreign retailers increased by a factor of four, from 365 to 1,335 supermarkets. As is apparent in figure 1, this period saw the expansion of foreign supermarkets beyond the large metropolitan areas of Mexico to smaller second- and third-tier cities. At the start of our sample in 2002, foreign stores were present in 96 of Mexico's 2,438 municipalities. In contrast, by 2014 foreign stores were present in 461 municipalities.

B. *Motivating Evidence*

How do foreign-owned supermarkets differ from the domestic retailers that they compete with after they enter? In this subsection, we use the consumer panel microdata and the administrative records of the Mexican National Retail Association (ANTAD)—both described in Section IV below—to document a set of stylized facts about how these stores differ.

Column 1 of table 1 regresses log prices on a dummy for whether the store is foreign-owned and on municipality-by-bar-code-by-month fixed effects. On average, foreign stores charge 12 percent lower prices for

[7] In this paper we consider only foreign entry in supermarket retail—which we take to be stores of 10,000 square feet and above—and so exclude smaller foreign-owned store formats such as convenience stores (in part because of data constraints; see n. 18). When considering impacts on domestic stores, we include all types of store formats including supermarkets, traditional stores, and street markets.

FIG. 1.—Foreign store presence at the end of 1995 (*top*), end of 2001 (*middle*), and end of 2013 (*bottom*). Municipalities in gray indicate foreign store presence at the end of 1995 (*top*, 204 stores), 2001 (*middle*, 365 stores), and 2013 (*bottom*, 1,335 stores). The data come from annual publications of the Mexican National Retail Association (ANTAD). For the period after 2006, we complement these data with annual retailer reports, press releases, and store location lists from retailer websites. See the data section for further details.

identical bar codes compared to domestic retailers in the same municipality during the same month. Interestingly, the sign of this difference is reversed and its magnitude doubles when we replace the municipality-by-bar-code-by-month fixed effects by municipality-by-product-group-by-

TABLE 1
How Do Foreign Owned Supermarkets Differ Ex Post?

	Dependent Variable			
	Log Price (1)	Log Price (2)	Log Number of Bar Codes (3)	Log Floor Space (4)
Foreign store dummy	−.118***	.249***	1.612***	1.911***
	(.00913)	(.0160)	(.0671)	(.0416)
Municipality-by-year fixed effects	✓	✓	✓	✓
Municipality-by-product-by-month fixed effects	✓	✓	X	X
Municipality-by-bar-code-by-month fixed effects	✓	X	X	X
Observations	18,659,777	18,659,777	10,393	11,113
R^2	.923	.368	.139	.302
Number of municipalities	151	151	151	499

Note.—The table reports the coefficient from regressing log prices, log number of bar codes, or log floor space on a foreign store dummy. Columns 1 and 2 are based on the Mexican consumer panel microdata for the years 2011–14 and compare foreign-owned supermarkets to all types of domestic retail establishments (traditional and modern) either with or without bar code–level fixed effects. Column 3 also uses the consumer panel to compare the number of bar code products sold in foreign stores to modern domestic retailers (domestic supermarkets or big-box stores). Column 4 is based on ANTAD data on establishment floor space and again compares foreign-owned supermarkets to modern domestic retailers. Regressions using the consumer panel data are weighted by household survey weights. Standard errors are clustered at the municipality level and reported in parentheses.

* Significant at the 10 percent level.
** Significant at the 5 percent level.
*** Significant at the 1 percent level.

month fixed effects (col. 2). Thus, foreign stores appear to offer a product mix that is significantly higher quality—where quality is proxied by price—and/or larger pack sizes within product groups (with anecdotal evidence suggesting both are true).

Foreign stores also sell a much larger set of product varieties. Column 3 uses the consumer panel data to regress the log of the count of different bar codes consumed by all households at a particular retailer in a particular year on a foreign store dummy and municipality-year fixed effects. We restrict this foreign-owned versus domestic comparison to modern store formats (i.e., supermarkets), leaving aside the smaller traditional domestic stores. Even with this restriction, a foreign-owned store offers approximately five times as many bar code products. Foreign stores offer not just more varieties but also different varieties. Using data on purchases in the consumer panel, on average, 60 percent of the bar code varieties offered by foreign stores in a given year and municipality are not offered by any domestic retailer (modern or traditional) in the locality. This difference in consumer choice is also clear when comparing the floor space records using the ANTAD data. Column 4 of table 1 shows that the average foreign-owned store is approximately six times the size of a domestic retailer that is also a member of ANTAD.

Finally, there are a number of differences in the shopping amenities—the store's environment, location, and so forth—offered by foreign-owned supermarkets compared to domestic retailers. In terms of positive amenities, foreign-owned supermarkets are typically more hygienic; offer greater security, more parking spaces, and better car accessibility; and display and organize their products more attractively. In addition, households may value American or European supermarket brands more than domestic ones for aspirational reasons. In terms of negative amenities, foreign-owned stores tend to be located farther away from the town center because of both their larger size and their later entry into the market. Given significant differences in car ownership rates across the Mexican income distribution as well as potential differences in proximity to supermarkets across rich and poor neighborhoods, this differential accessibility will play an important role when estimating heterogeneity in the gains from foreign entry. In the theoretical framework that follows, a revealed preference approach captures these different amenities through income and product group–specific taste shifters across retailers that generate observable differences in the post-entry market shares of foreign stores.

To summarize, foreign stores differ substantially on a number of key dimensions: they charge lower prices, offer higher-quality products and larger pack sizes, sell a much larger variety of products, and offer a different set of amenities. The size of these differences is substantial, certainly compared with the differences between big-box stores and preexisting retailers in the United States (e.g., Hausman and Leibtag 2007; see further discussion in Sec. VI.D.3). Such differences come in part from the fact that these foreign retailers pioneered the use of big-box store formats, regional distribution centers, cutting-edge logistics such as cold chains for fresh products, and lean global supply chains (Biles 2008). Essentially, our empirical setting captures the entry of global retail chains at the world technological frontier in retailing into local retail markets that are largely dominated by traditional store formats, street markets, and small regional supermarket chains.

III. Theoretical Framework

In this section, we derive a general expression for assessing the impact of foreign supermarket entry on local household welfare as a function of various observable moments in our rich collection of microdata. In order to calculate the change in welfare due to the entry of a foreign supermarket, we consider the compensating variation for household h, the change in exogenous income required to maintain utility when a foreign retailer arrives between period 1 and period 0, with periods denoted by superscripts:[8]

$$CV_h = \underbrace{\left[e\left(\mathbf{P}^1, u_h^0\right) - e\left(\mathbf{P}^0, u_h^0\right)\right]}_{\text{Cost-of-living effect (CLE)}} - \underbrace{\left(y_h^1 - y_h^0\right)}_{\text{Nominal income effect (IE)}}, \tag{1}$$

where $\mathbf{P}^t$ is the vector of prices faced by the household in period t, u_h^t is the household's utility, and y_h^t is its nominal income (and where, for notational parsimony, changes between periods 0 and 1 denote the causal changes induced by foreign retail entry).

The first term is the cost-of-living effect, the welfare change due to the price changes induced by the arrival of the foreign retailer. The second term is the nominal income effect, the welfare change due to any changes in household income that result from the arrival of the foreign retailer. In the next two subsections we decompose the cost-of-living effect and nominal income effect into six distinct channels and express these as functions of observable moments in our microdata.

A. *Estimating the Cost-of-Living Effect*

While, at least in principle, the nominal income effect can be empirically estimated without imposing additional structure, this is not the case for the cost-of-living effect. While we can observe the vector of price changes $\mathbf{P}_{dc}^1 - \mathbf{P}_{dc}^0$ for products sold in domestic continuing stores indexed by dc, that is, those that are present in both periods, there are two sets of price changes that are inherently unobservable: the price changes $\mathbf{P}_f^1 - \mathbf{P}_f^0$ at entering foreign retailers indexed by f and the price changes $\mathbf{P}_{dx}^1 - \mathbf{P}_{dx}^0$ at domestic exiting retailers indexed by dx. In particular, foreign retailers' prices are not observed prior to their entry, and exiting domestic retailers' prices are not observed after exit. As first noted by Hicks (1940), we can replace these two unobserved price vectors with "virtual" price vectors, the price vectors that would set demand for these stores equal to zero given the vector of consumer prices for other goods and services.

To see this more clearly, note that the cost-of-living effect in (1) can be rewritten by dividing it into three quite distinct subcomponents, one for each of the three sets of price changes above: a direct price index effect due to the implicit price changes at foreign stores (i.e., the gains enjoyed by consumers shopping at the new foreign store), a procompetitive price effect due to continuing domestic retailers changing prices as a result of foreign retail entry, and a procompetitive exit effect due to the implicit price changes at exiting domestic stores (i.e., the losses suffered by customers of closing stores):

[8] This approach follows earlier work by Hausman (1997) and Hausman and Leonard (2002).

$$\begin{aligned}\text{CLE} = {} & \underbrace{\left[e\left(\mathbf{P}_f^1, \mathbf{P}_{dc}^1, \mathbf{P}_{dx}^{1*}, u_h^0\right) - e\left(\mathbf{P}_f^{1*}, P_{dc}^1, \mathbf{P}_{dx}^{1*}, u_h^0\right)\right]}_{\text{(1) Direct price index effect (DE)}} \\ & + \underbrace{\left[e\left(\mathbf{P}_f^{1*}, \mathbf{P}_{dc}^1, \mathbf{P}_{dx}^{1*}, u_h^0\right) - e\left(\mathbf{P}_f^{0*}, \mathbf{P}_{dc}^0, \mathbf{P}_{dx}^{0*}, u_h^0\right)\right]}_{\text{(2) Procompetitive price effect (PP)}} \\ & + \underbrace{\left[e\left(\mathbf{P}_f^{0*}, \mathbf{P}_{dc}^0, \mathbf{P}_{dx}^{0*}, u_h^0\right) - e\left(\mathbf{P}_f^{0*}, \mathbf{P}_{dc}^0, \mathbf{P}_{dx}^0, u_h^0\right)\right]}_{\text{(3) Procompetitive exit effect (PX)}},\end{aligned} \tag{2}$$

where asterisks denote virtual prices: $\mathbf{P}_f^{0*}$ and $\mathbf{P}_f^{1*}$ are pre- and post-entry prices in foreign stores that would set demand equal to zero, while $\mathbf{P}_{dx}^{0*}$ and $\mathbf{P}_{dx}^{1*}$ are the same for domestic exiting stores. Since virtual prices are inherently unobservable, they must be estimated, which requires a demand function or at least an approximation to one. Below we propose two approaches: an exact estimation under CES demand and a first-order approximation. Finally, we note that although we label the price changes captured by the second term as "procompetitive," they may derive from either reductions in markups or increases in productivity at domestic stores (distinctions that do not matter on the cost-of-living side but would generate different magnitudes of profit and income effects that we capture from observed changes on the nominal income side). We further investigate this distinction in Section IV.

1. Exact Estimation under CES Demand

We propose a three-tier demand system. In the upper tier there are Cobb-Douglas preferences over product groups $g \in G$ (e.g., beverages), in the middle tier there are asymmetric CES preferences over local retailers selling that product group $s \in S$ (e.g., Walmex, a foreign retailer; Soriana, a domestic retailer in modern retail; or a mom-and-pop store in the traditional retail sector), and in the final tier there are preferences over the individual products within the product groups $b \in B_g$ (e.g., a product such as a 330-milliliter Coca-Cola can) that we can leave unspecified for now:

$$U_h = \prod_{g \in G} (Q_{gh})^{\alpha_{gh}}, \tag{3}$$

$$Q_{gh} = \left[\sum_{s \in S_g} \beta_{gsh} q_{gsh}^{(\eta_{gh}-1)/\eta_{gh}}\right]^{\eta_{gh}/(\eta_{gh}-1)}, \tag{4}$$

where α_{gh} and β_{gsh} are (potentially household- or income group–specific) preference parameters that are fixed across periods; Q_{gh} and q_{gsh} are product group and store–product group consumption aggregates with associated price indices P_{gh} and r_{gsh}, respectively; and η_{gh} is the elasticity of substitution across local retail outlets.[9] For each broad product group, consumers choose how much to spend at different stores based on the store-level price index r_{gsh} (which itself depends on the products they anticipate buying in each store given its product mix and product-level prices).

This structure seems reasonable given that stores often specialize in certain product groups and, at least within a quarter, consumers often shop at several stores and choose different sets of products in each store. While the demand system is homothetic, we capture potential heterogeneity across the income distribution by allowing households of different incomes to differ in their expenditure shares across product groups (α_{gh}), their preferences for consumption bundles at different stores within those product groups (β_{gsh} and the preference parameters that generate q_{gsh}), as well as their elasticity of substitution across local stores (η_{gh}).[10]

This approach has several advantages. First, as shown by Anderson et al. (1992), these preferences generate the same demands that would be obtained from aggregating many consumers who make discrete choices over which store to shop in. This mapping is appealing, particularly since in estimating price elasticities our unit of observation will be household income groups (observed separately for each location, period, and product group).

Second, it has the appeal of being widely used in the trade literature starting with Feenstra (1994), in part because it yields a very parsimonious expression for the welfare gains from new products (or stores in our case).[11] Building on Feenstra (1994), the following expression provides the exact proportional cost-of-living effect under this demand system:

[9] This structure imposes that the elasticity of substitution between two modern retail stores is the same as across a modern and traditional store. We later explore the sensitivity of our estimates to relaxing this assumption by placing modern and traditional stores in different nests.

[10] While convenient for empirical tractability, this ad hoc treatment of nonhomotheticity shuts down a second-order price index effect. Large first-order effects of foreign entry on incomes may push some households across income groups and thereby change their preference parameters as defined above. Since we will allow preferences to differ across seven broad income groups, it is reasonable to think that few households are shifted in this manner.

[11] Note that even the translog specification as in Feenstra and Weinstein (2017) requires knowledge of the causal effect of foreign entry on Herfindahl indices of the local retail market.

$$
\begin{aligned}
\frac{\mathrm{CLE}}{e\left(\mathbf{P}_f^{0*}, \mathbf{P}_{dc}^0, \mathbf{P}_{dx}^0, u_h^0\right)} &= \frac{e\left(\mathbf{P}_f^1, \mathbf{P}_{dc}^1, \mathbf{P}_{dx}^{1*}, u_h^0\right)}{e\left(\mathbf{P}_f^{0*}, \mathbf{P}_{dc}^0, \mathbf{P}_{dx}^0, u_h^0\right)} - 1 \\
&= \prod_{g\in G}\left[\left(\frac{\sum_{s\in S_g^{dc}}\phi_{gsh}^1}{\sum_{s\in S_g^{dc}}\phi_{gsh}^0}\right)^{1/(\eta_{gh}-1)} \prod_{s\in S_g^{dc}}\left(\frac{r_{gsh}^1}{r_{gsh}^0}\right)^{\omega_{gsh}}\right]^{\alpha_{gh}} - 1,
\end{aligned} \tag{5}
$$

where S_g^{dc} denotes the set of continuing domestic retailers within product group g, $\phi_{gsh}^t = r_{gsh}^t q_{gsh}^t / \Sigma_{s\in S_g} r_{gsh}^t q_{gsh}^t$ is the expenditure share for a particular retailer in product group g, and the ω_{gsh}'s are ideal log change weights.[12]

For each product group g, the expression has two components. The $P_{gh} \equiv \prod_{s\in S_g^{dc}}(r_{gsh}^1/r_{gsh}^0)^{\omega_{gsh}}$ term is a Sato-Vartia (i.e., CES) price index for price changes in continuing domestic stores that forms the *procompetitive price effect*.[13] The price terms r_{gsh}^t are themselves price indices of product-specific prices p_{gsb}^t within domestic continuing stores, which, in principle, could also account for new product varieties using the same methodology. Empirically, we find no evidence of such effects in response to foreign retail arrival and so abstract from this possibility in the exposition.[14]

The term

$$
\left(\frac{\sum_{s\in S_g^{dc}}\phi_{gsh}^1}{\sum_{s\in S_g^{dc}}\phi_{gsh}^0}\right)^{1/(\eta_{gh}-1)}
$$

captures the gains to customers of the foreign store in the numerator, the *direct price index effect*, and domestic store exit in the denominator, the *procompetitive exit effect*. For expositional purposes, consider the simple case in which there are no procompetitive effects (such as when firms are monopolistically competitive as in Krugman [1980]):

$$
\frac{\mathrm{CLE}}{e\left(\mathbf{P}_f^{0*}, \mathbf{P}_{dc}^0, \mathbf{P}_{dx}^0, u_h^0\right)} = \prod_{g\in G}\left[\left(\sum_{s\in S_g^{dc}}\phi_{gsh}^1\right)^{1/(\eta_{gh}-1)}\right]^{\alpha_{gh}} - 1. \tag{6}
$$

[12] In particular,

$$
\omega_{gsh} = \left(\frac{\tilde{\phi}_{gsh}^1 - \tilde{\phi}_{gsh}^0}{\ln\tilde{\phi}_{gsh}^1 - \ln\tilde{\phi}_{gsh}^0}\right) \Big/ \sum_{s^{dc}\in S_g^{dc}}\left(\frac{\tilde{\phi}_{gsh}^1 - \tilde{\phi}_{gsh}^0}{\ln\tilde{\phi}_{gsh}^1 - \ln\tilde{\phi}_{gsh}^0}\right),
$$

which in turn contain expenditure shares of different retailers within product groups where the shares consider only expenditure at continuing retailers $\tilde{\phi}_{gsh}^t = r_{gsh}^t q_{gsh}^t / \Sigma_{s\in S_g^{dc}} r_{gsh}^t q_{gsh}^t$.

[13] Notice that the assumption of CES preferences does not imply the absence of procompetitive effects as we do not impose additional assumptions about market structure (e.g., monopolistic competition).

[14] In particular, we find no evidence in the CPI microdata that foreign retail entry increases the propensity for product additions or replacements among domestic retailers. We report these regressions in online app. table A.1.

The welfare gain from a new store is a function of the market share of that store after entry and the elasticity of substitution across stores. The revealed preference nature of this approach is clear. If consumers greatly value the arrival of the new store—be it because the store offers low prices p^1_{gsb}, more product variety that reduces r^1_{gsh}, or better amenities captured by a large β_{gsh}—the market share is higher and the welfare gain greater. Hence, these market share changes capture all the potential benefits of shopping in foreign stores outlined in the motivating evidence of Section II. How much greater depends on the elasticity of substitution. Large foreign market shares will imply small welfare changes if consumers have a high elasticity of substitution between stores and large welfare changes if they are inelastic. A similar logic applies to the exit of domestic stores, where a large period 0 market share means large welfare losses, again tempered by the elasticity of substitution.[15]

Equation (6) also makes clear a third benefit of this approach. The CES assumption allows us to relate our estimation results to the recent quantitative literature on the gains from trade and FDI since the expression of the cost-of-living effect in the absence of procompetitive effects is identical to the well-known import share sufficient statistic of Arkolakis et al. (2012). Thus, our welfare expression allows us to shed light on the importance of procompetitive effects, to separately estimate effects on nominal incomes and household cost of living, and to quantify the distribution of the gains from FDI through the household-level heterogeneity we incorporate.

In our quantification, we decompose the welfare gains into their constituent parts. Accordingly, we add and subtract terms to equation (5) to separate the cost-of-living effect into the direct price index effect and the two procompetitive effects described in the preceding paragraphs:

$$\begin{aligned}\frac{\mathrm{CLE}}{e\left(\mathbf{P}_f^{0*}, \mathbf{P}_{dc}^0, \mathbf{P}_{dx}^0, u_h^0\right)} = &\underbrace{\left\{\prod_{g\in G}\left[\left(\frac{\sum_{s\in S_g^{dc}}\phi^1_{gsh}}{\sum_{s\in S_g^{dc}}\phi^0_{gsh}}\right)^{\frac{1}{\eta_{gh}-1}} P_{gh}\right]^{\alpha_{gh}} - \prod_{g\in G}\left[\left(\frac{1}{\sum_{s\in S_g^{dc}}\phi^0_{gsh}}\right)^{\frac{1}{\eta_{gh}-1}} P_{gh}\right]^{\alpha_{gh}}\right\}}_{\text{(1) Direct price index effect (DE)}} \\ &+ \underbrace{\left[\prod_{g\in G}\left(P_{gh}\right)^{\alpha_{gh}} - 1\right]}_{\text{(2) Procompetitive price effect (PP)}} + \underbrace{\left\{\prod_{g\in G}\left[\left(\frac{1}{\sum_{s\in S_g^{dc}}\phi^0_{gsh}}\right)^{\frac{1}{\eta_{gh}-1}} P_{gh}\right]^{\alpha_{gh}} - \prod_{g\in G}\left(P_{gh}\right)^{\alpha_{gh}}\right\}}_{\text{(3) Procompetitive exit effect (PX)}},\end{aligned} \tag{7}$$

where recall that $P_{gh} \equiv \prod_{s\in S_g^{dc}}\left(r^1_{gsh}/r^0_{gsh}\right)^{\omega_{gsh}}$.

[15] We note that such a revealed preference approach would not capture cultural losses from the closure of traditional stores to the extent that those are not internalized by consumers. Similarly, such an approach captures only the health consequences of changes in shopping behavior that are internalized by consumers.

2. First-Order Approach Using Observed Price Differences

While the assumption of CES preferences has its virtues, it also imposes a particular structure on household demands. As an alternative approach, we exploit the richness of the store price data to estimate a first-order approximation of the cost-of-living effect that is solely based on observable price changes due to foreign entry.

We take a first-order Taylor expansion of the expenditure function around period 1 prices and apply Shephard's lemma. Focusing on the sales and price changes in the set of domestic stores continuously selling product b across both periods (for which we can observe price changes) provides us with the *procompetitive price effect*:

$$\mathrm{PP}' \approx \sum_{b} \sum_{s \in S_b^{dc}} \left[q_{bsh}^1 \left(p_{bs}^1 - p_{bs}^0 \right) \right], \tag{8}$$

where q_{bsh}^t is the quantity consumed of product b in store s by household h in period t and S_b^{dc} is the set of domestic stores continuously selling product b across both periods. Rewriting the PP′ in proportional terms,

$$\frac{\mathrm{PP}'}{e\left(\mathbf{P}_f^1, \mathbf{P}_{dc}^1, \mathbf{P}_{dx}^{1*}, u_h^0 \right)} \approx \sum_{b} \sum_{s \in S_b^{dc}} \left[\phi_{bsh}^1 \left(\frac{p_{bs}^1 - p_{bs}^0}{p_{bs}^1} \right) \right], \tag{9}$$

where ϕ_{bsh}^1 is the household expenditure share spent on the product in period 1. To a first-order approximation, the procompetitive effect is simply a Paasche price index of the product-level price changes at continuing domestic stores due to foreign entry multiplied by the period 1 share of total expenditure captured by that store-product pair. Since the first-order approach explicitly assumes that no stores exited between periods 0 and 1, there are no separate exit and price effects.

For the *direct price index effect*, we focus on the sales and price changes at foreign stores in the Taylor expansion around period 1 prices:

$$\frac{\mathrm{DE}'}{e\left(\mathbf{P}_f^1, \mathbf{P}_{dc}^1, \mathbf{P}_{dx}^{1*}, u_h^0 \right)} \approx \sum_{b} \sum_{s \in S_b^{f}} \left[\phi_{bsh}^1 \left(\frac{p_{bf}^1 - p_{bds}^0}{p_{bf}^1} \right) \right], \tag{10}$$

where S_b^f is the set of foreign stores selling product b in period 1. Hence, the direct price index effect corresponds to a Paasche price index of the product-level price differences between foreign stores in period 1, p_{bf}^1, and domestic stores in period 0, p_{bds}^0, multiplied by the period 1 share of total expenditure captured by foreign stores for that particular product. Essentially, this approach is equivalent to assuming that the foreign stores were always present and always selling the same set of products and providing the same amenities, but in period 0 they charged the pre-entry prices charged by domestic stores. In this sense, we abstract from unobserved gains to product variety, store variety, and amenity and solely focus on observable price changes.

The benefits of this approach are clear. It yields a transparent Paasche price index that approximates the consumer gains from foreign entry purely on the basis of observable moments in the price microdata. Essentially, we are multiplying the post-entry foreign market share by the observed price differences between foreign and domestic stores for the direct price index effect and the post-entry domestic market share by the price changes at domestic stores for the procompetitive effect. The disadvantages are equally clear. Since we implicitly assume that the foreign stores were always present, we miss any gains that arise from the existence of another store as well as the greater product variety and the amenities provided by foreign stores. Given that these variety and amenity differences are substantial, a fact we highlighted in Section II, we prefer the exact approach under CES for our baseline estimates.

Finally, reporting both approaches provides an additional benefit. For the reasons discussed above, the difference between the direct price index effect under CES (the first term in eq. [7]) and the direct price index effect under the first-order approach (eq. [10]) provides us with an approximate estimate of the proportion of the direct consumer gains from the new foreign store that come from new product variety, new store variety, and different store amenities as opposed to from lower prices alone.[16]

B. *Estimating the Nominal Income Effect*

The nominal income effect in equation (1) can also be separated into distinct subcomponents. We divide the household's income sources into three groups: Households obtain labor income from working in retail, business income from owning and operating their own retail outlet, and both labor and business income from other sectors (i.e., nonretail) indexed by o. For labor and business incomes in retail we additionally distinguish between the traditional retail segment (mom-and-pop stores, street

[16] The comparison is approximate since the first-order direct price index effect (that is due only to price gaps) is biased upward: the Paasche weights ignore the substitution away from foreign stores if they charged domestic store prices in the pre-entry period. We can avoid this bias by using the CES preference structure described above along with knowledge of the η_{gh} parameters to estimate counterfactual first-period expenditure weights that take account of this substitution and then calculate an unbiased Sato-Vartia price index. In contrast, the advantage of the Paasche approach is that we do not require unobserved counterfactual market shares. Given this trade-off, the main text reports the Paasche approach and n. 40 reports the Sato-Vartia approach. Note that using observed foreign store budget shares before and after entry to construct a price index using observed price gaps would estimate neither the full cost-of-living effect (which requires us to compute virtual prices consistent with zero consumption before entry) nor the effect of price changes alone (which requires counterfactual pre-entry budget shares as we discuss above).

stalls, etc.), indexed by τ, and modern store formats (supermarkets, chain stores, etc.), indexed by μ:

$$y_h = \sum_{i\in\{\tau,\mu\}} l_{ih} + \sum_{i\in\{\tau,\mu\}} \pi_{ih} + \sum_{i\in\{o\}} (l_{ih} + \pi_{ih}), \tag{11}$$

where l_{ih} and π_{ih} denote labor income and business income from sector i, respectively. Taking a first difference and dividing through by initial income, we obtain three nominal income effects:

$$\begin{aligned}\frac{\mathrm{IE}}{e\left(\mathbf{P}_f^{0*}, \mathbf{P}_{dc}^0, \mathbf{P}_{dx}^0, u_h^0\right)} \approx & -\underbrace{\sum_{i\in\{\tau,\mu\}}\left[\theta_{ilh}^0\left(\frac{l_{ih}^1 - l_{ih}^0}{l_{ih}^0}\right)\right]}_{\text{(4) Retail labor income effect}} \\ & -\underbrace{\sum_{i\in\{\tau,\mu\}}\left[\theta_{i\pi h}^0\left(\frac{\pi_{ih}^1 - \pi_{ih}^{i0}}{\pi_{ih}^0}\right)\right]}_{\text{(5) Retail profit effect}} \\ & -\underbrace{\sum_{i\in\{o\}}\left[\theta_{ilh}^0\left(\frac{l_{ih}^1 - l_{ih}^0}{l_{ih}^0}\right) + \theta_{i\pi h}^0\left(\frac{\pi_{ih}^1 - \pi_{ih}^{i0}}{\pi_{ih}^0}\right)\right]}_{\text{(6) Other income effect}},\end{aligned} \tag{12}$$

where θ_{ilh}^0 and $\theta_{i\pi h}^0$ are the period 0 shares of total income that come from labor and business income in sector i, respectively.

Foreign retail entry may change labor incomes in both the traditional and modern retail sectors, the *retail labor income effect*. Foreign entry may also affect the profits of domestic store owners, the *retail profit effect*. Finally, foreign entry may give rise to general equilibrium effects on labor and business incomes in other sectors of the local economy or affect incomes for households producing goods sold at local retailers, the *other income effect*. Each of these nominal income effects can occur at both the intensive and extensive margins. At the intensive margin, foreign entry can affect earnings of individuals who remain active in a given sector and occupation. At the extensive margin, foreign entry may lead households to reallocate across sectors and occupations.[17]

IV. Data

The theoretical framework outlined in the previous section allows us to express the gains from foreign retail entry in equations (7) and (12) as

[17] To see this more clearly, we can decompose each of the three terms in expression (12) into three mutually exclusive margins: an intensive margin (e.g., in the case of labor income, $\theta_{ilh}^0[(l_{ih}^1 - l_{ih}^0)/l_{ih}^0]$ if $l_{ih}^0 > 0$ and $l_{ih}^1 > 0$), a job loss margin (e.g., $\theta_{ilh}^0(-1)$ if $l_{ih}^0 > 0$ and $l_{ih}^1 = 0$), and a new jobs margin (e.g., l_{ih}^1/y_h^0 if $l_{ih}^0 = 0$ and $l_{ih}^1 > 0$). Note that we do not attempt to value changes in household leisure time, which we implicitly assume to be fixed.

a function of (i) causal effects on consumer prices, consumption quantities, and household nominal incomes; (ii) household demand parameters that govern the elasticity of substitution across retail outlets; and (iii) household expenditure shares across product groups and store types within product groups, and income shares across sectors and occupations. This section describes the data sources we draw on to obtain these estimates. Online appendix table A.2 contains descriptive statistics for the key variables in each data set.

Store opening dates and locations.—Our main regressor of interest is the first entry of a foreign-owned supermarket in a municipality. To generate this variable, we obtain data on store locations and dates of opening from Mexico's national association of retail businesses ANTAD (Asociación Nacional de Tiendas de Autoservicio y Departamentales). All major national and regional retailers in Mexico are part of ANTAD, comprising more than 34,000 retail units with close to 25 million square meters of retail space. Between 2002 and 2006, ANTAD collected detailed data from its members about the location and date of opening of every establishment. For subsequent periods we obtained foreign-owned supermarket openings directly from retailers' annual reports. If these were not available, we used their store locations as of March 2014 (listed on their websites) and then obtained opening dates from local newspaper coverage of store openings or by calling them.[18] Throughout the paper, we define a retailer as foreign if 49 percent or more of the firm is held by a foreign retailer (the cap on foreign ownership before FDI liberalization in 1994). With the exception of Safeway, which owns 49 percent of Casa Ley, all foreign retailers own majority stakes.

In our empirical work it will be important to control for trends using comparable municipalities. Since foreign retailers rarely open in rural areas or small towns, we exclude these from our analysis by restricting attention to the 608 municipalities with at least one chain store (i.e., ANTAD member) in at least one year of our ANTAD data. Unsurprisingly, these municipalities are larger (a median population of 63,000 compared to a median of 8,000 for the remaining 1,848 municipalities) and exclusively urban. By the end of our sample in 2014, 76 percent of these sample municipalities contain a foreign retailer whereas only 16 percent did at the start of our sample in 2002.

In both the theoretical framework and empirical analysis, we distinguish between traditional and modern retail store formats. The distinction

[18] Our focus on foreign supermarket entry comes in part from the fact that this strategy was not feasible for smaller store formats, where the number of units per chain goes from hundreds (supermarkets) to thousands (convenience stores). In the 2002–6 data, smaller-format foreign entry patterns are uncorrelated with foreign entry in supermarket retail. We also thank Mauricio Varela for providing data on Walmart store openings between 2002 and 2006 (when Walmart was not a member of ANTAD).

that we can observe most consistently across data sets is between mom-and-pop stores, street stalls, and independent specialist stores (e.g., butchers and hardware stores) on the traditional side and supermarkets, chain specialist and convenience stores, and department stores on the modern side.

CPI microdata.—To estimate the procompetitive price effect, we rely on the monthly microdata that are used to construct the Mexican CPI. These data consist of retail price quotes collected by Mexico's national statistics agency, Instituto Nacional de Estadística y Geografía (INEGI) (since July 2011), and Mexico's Central Bank (prior to July 2011). Every month INEGI enumerators obtain price quotes (inclusive of any promotions/sales and value-added tax) for around 83,500 items covering 315 product categories in 141 urban municipalities. These individual price quotes are made publicly available on a monthly basis in an official government gazette.[19] Additional details are provided in online appendix B.1. Because computing the CPI requires prices of identical products in the same retail outlet over time, these data are ideally suited to estimate price changes at surviving domestic retail establishments.

In addition to these publicly accessible data, we also obtain access to the confidential data columns of the Mexican CPI. These crucially allow us to observe the municipality in which the price quote was taken, as well as store format type and retailer name. The latter information allows us to explore heterogeneity across traditional and modern stores as well as to remove foreign stores from the estimation of the procompetitive price effect.

As the price sampling is designed to be representative of Mexican household consumption, these data have a number of useful features. First, the price quotes are collected from not only modern stores but also traditional stores (including street stalls). Second, the quotes cover not just retail product groups but also services such as health, education, housing, and transport. Third, within a given product group, the products and stores sampled are chosen to match the consumption patterns of urban households in the Encuesta Nacional de Ingresos y Gastos de los Hogares (ENIGH) consumption surveys discussed below.

When comparing prices of the same product over time, we will focus on the subset of goods that are identified by their brand, pack size, and variety (e.g., fresh whole milk Alpura brand 1-liter carton). These bar code equivalent products constitute more than one-third of all price quotes in the CPI microdata, and product groups that predominantly contain bar code equivalent products account for, on average, 40 percent of household retail expenditure. Focusing on these products allows us to ensure there are no changes in product characteristics over time that may confound estimates of the procompetitive price effect. The final esti-

[19] We thank Etienne Gagnon for access to the data he assembled directly from the gazette.

mation sample consists of roughly 3.3 million monthly store-price observations in 120 product categories across 76 urban municipalities over the period 2002–14. In our main analysis, we assume that the price changes due to foreign entry that we estimate from bar code equivelent products for the reasons outlined above are also applicable to non–bar code equivalent products within the same product group. We also report robustness results in which we make alternative assumptions about price changes for non–bar code equivalent items.

Consumer panel microdata.—The estimation of the direct price index effect requires data on the post-entry retail market shares of foreign supermarkets as well as estimates of the elasticity of substitution across local stores (both across product groups and across different levels of household income). For this purpose, we use the consumer panel microdata of a large international market research company.[20] This Mexican consumer panel covers the years 2011–14 and is similar in nature to the home scanner data that market research companies collect in the United States. The panel consists of approximately 6,000 urban households classified by seven income groups and distributed across 156 municipalities. Households are visited biweekly to obtain complete consumption diary information about all products purchased by the household. As with the CPI data, these data are at the bar code equivalent level, with enumerators carefully noting the brand, variety, and pack size. The household sample is updated annually to be representative of all cities over 50,000 once the provided survey weights are taken into account. These microdata comprise roughly 24 million transaction-level observations between January 2011 and June 2014. Importantly, in contrast to the academic use versions of similar US data sets, we have retailer identities for every transaction in a household's consumption basket. Thus, these data are ideally suited to observe retailer market shares by household, as well as to estimate elasticities of substitution across stores.

Retail census microdata.—For the purpose of estimating the effect of foreign entry on retail profits and domestic store exit, we use the confidential version of the Economic Census microdata for the years 2003 and 2008 (Censos Economicos 2004 and 2009 from INEGI). The Economic Census records establishment-level information for the universe of urban retail establishments. The restricted access version of the data we use allows us to separately observe the number of modern retail stores (supermarkets, chain stores, and department stores) and traditional retail stores (the remaining stores), as well as store-level revenues and costs from which we compute profits. Additional details are provided in online appendix B.2. The resulting data set contains 1.3 million retail establishments across our 608 sample municipalities in 2003 and 1.5 million in 2008.

[20] These data were made available to us through an academic collaboration with its Mexico City office under the condition that the firm's name remained anonymous.

Employment and occupation survey microdata.—To estimate the effect of foreign entry on nominal incomes and employment, we require high-frequency survey data. To this end, we use INEGI's National Urban Employment Surveys between 2002 and 2005 and its successor, the National Employment and Occupation Surveys, between 2005 and 2012. These surveys are rotating panels of households in which a given household is followed over five quarters. The survey tracks sector, occupation, and income similarly to the ENIGH data set described below but has the advantage of a much larger sample size: every quarter more than 100,000 individual residences are surveyed with the details of each working-age household member recorded. The resulting sample comprises roughly 5 million person-quarter observations across 273 urban municipalities.

Household income and expenditure survey microdata.—In our quantification exercise, in order to calculate welfare effects across the income distribution, we need to know the expenditure shares of households across various product groups matched to income shares from various occupations and sectors. For this purpose, we use the Mexican National Income and Expenditure Surveys (ENIGH), which are administered biannually by INEGI between 2006 and 2012. These data allow us to observe the income and sources of income for each household as well as their expenditure shares across all retail and nonretail product groups. Additionally—starting in 2006; hence our use of only the more recent rounds—these data record for each product group the proportions sold at different types of stores (supermarket, street stall, etc.). Given the welfare expression derived in the previous section, we require pre-entry income and expenditure shares, and so we restrict attention to the 12,293 households residing in 240 urban municipalities that had not yet experienced foreign retail entry at the time the ENIGH survey was conducted. We match the product modules covered by the consumer panel data above to the household income and expenditure surveys at the level of 12 broad product groups. As discussed above, the match between ENIGH and the CPI price data is straightforward, as both are based on the same product groups.

V. Estimating the Effects of Foreign Retail Entry

This section draws on the microdata described in the previous section to estimate the effect of foreign retail entry on local consumer prices, retail market shares, store exit, household labor and business incomes, and employment. As well as being of interest in their own right, these estimates enter into our cost-of-living and nominal income expressions, equations (7) and (12), and hence form the basis of the quantification exercise in Section VI. We focus on the effect of the first foreign store entry as we wish to capture the impact of both the initial entry and any subsequent foreign entry induced by the initial entry. In practice, this choice makes little difference as the vast majority of urban municipalities—more than 85 percent of entry events in our sample—receive just one foreign store during our estimation period that ends in 2014.

A. *Effect on Consumer Prices*

1. Effect on Consumer Prices in Domestic Retail Stores

Empirical strategy.—To estimate the effect of foreign supermarket entry on consumer prices in domestic retail stores, we combine information on the universe of foreign store locations and opening dates with monthly panel data on local bar code–level prices from the 2002–14 CPI microdata. Since foreign stores are not randomly allocated, the obvious identification concern is that store openings are correlated with preexisting price trends. There are several possible scenarios. First, it could be the case that foreign retailers target municipalities with higher preexisting price growth or time their opening in a way that is correlated with positive local retail price shocks. Both of these scenarios would lead to an upward-biased estimate of the treatment effect of foreign entry on domestic store prices. Alternatively, foreign stores may target faster-growing municipalities whose retail environments are also becoming more competitive so that store prices could be on a preexisting downward trajectory. Finally, rather than targeting a particular set of municipalities at particular points in time, foreign retailers may have expanded rapidly between 2002 and 2014 with the aim of establishing store presence across the whole of urban Mexico as quickly as possible. In this final scenario, we would not expect substantial bias as, at least over this period, neither the selection of municipalities in our urban estimation sample nor the timing of opening would be strongly correlated with preexisting price trends.

We use the microdata to explore which of these scenarios is relevant by estimating the following baseline event study specification:

$$\ln p_{gsbmt} = \sum_{j=-13}^{36} \beta_j I(\text{Months since Entry}_{mt} = j) + \delta_{gsbm} + \eta_t + \epsilon_{gsbmt}, \qquad (13)$$

where $\ln p_{gsbmt}$ is the log price of bar-code-product b in product group g, individual store s, in municipality m and month t; $I(\cdot)$ is an indicator function; and Months since Entry_{mt} counts the months since the first foreign entry for each municipality m at time t (with negative values counting months before entry, positive values counting months after entry, and zero being the month the first foreign store enters a municipality).[21]

[21] We take $j = -6$ as the (omitted) reference category and define the indicator variable $I(\text{Months since Entry}_{mt} = 36)$ to take the value one for all $j \geq 36$, and similarly

Since the procompetitive price effect in equation (7) relates only to price changes at domestic stores, we remove foreign stores from the sample. The β_j parameters capture the effect of foreign entry on domestic store prices for each of j months before and after the opening event, δ_{gsbm} is a bar-code-by-store fixed effect, and η_t is a month fixed effect.

By estimating the treatment effect in the 12 months leading up to the opening event, this approach allows us to test for the presence and slope of trends in the run-up to the foreign store opening event in a transparent way and without imposing parametric structure. The absence of preexisting trends would suggest that the two troubling scenarios outlined above are not an issue, while if there are trends, the event study design allows us to sign and quantify the bias.

To estimate the event study on a fully balanced sample of municipalities both before and after the store opening, we exclude municipalities in which the first foreign store opened in the first 12 months of our data set (July 2002–June 2003) and municipalities in which the first foreign store opened in the last 36 months of our data set (April 2011–March 2014) or later. Balancing the panel is important to alleviate selection concerns when exploring the time path of treatment effects. Given the need to balance, there is a clear trade-off between a longer event study and a smaller, less representative, sample. Our choice of window is guided by the fact that we lose only 6 percent of our store price observations through this restriction (although, along with other robustness checks, we show results with an extended window below).

Estimation results.—Panel A of figure 2 presents the event study graph. Prices are flat (and not significantly different from zero) in the lead-up to the store entry event, start falling as soon as entry occurs, and level off approximately 24 months after entry at a negative and significant 3 percentage points. As evidenced by the treatment effect estimated for after 36 months (labeled "≥36" in the figure), this procompetitive price effect appears to be permanent. When we parametrically test for trend breaks, we find a precisely estimated flat price trend before foreign entry, a significant negative trend break at the time of foreign entry, and a return to a flat price trend about 2 years after foreign entry.[22] Note that since the CPI samples products and stores in proportion to their weight in the

---

I(Months since Entry$_{mt}$ $= -13$) to take the value one for all $j \leq -13$. As discussed below, unlike the points in between, these two periods before and after the event study cannot be estimated using a balanced municipality sample.

[22] We regress log bar code prices on a post-entry dummy, a dummy for 24 months or more after entry, and their interactions with Month since Entry$_{mt}$ in addition to δ_{gsbm} and η_t. The point estimates on Month since Entry$_{mt}$ are $-.00011$ (standard error [SE] = .00036) before entry, $-.00116$ (SE = .00046) for the interaction of Month since Entry$_{mt}$ with a post-entry dummy, and $-.00010$ (SE = .00043) for the interaction of Month since Entry$_{mt}$ with the dummy for 24 months or more after entry.

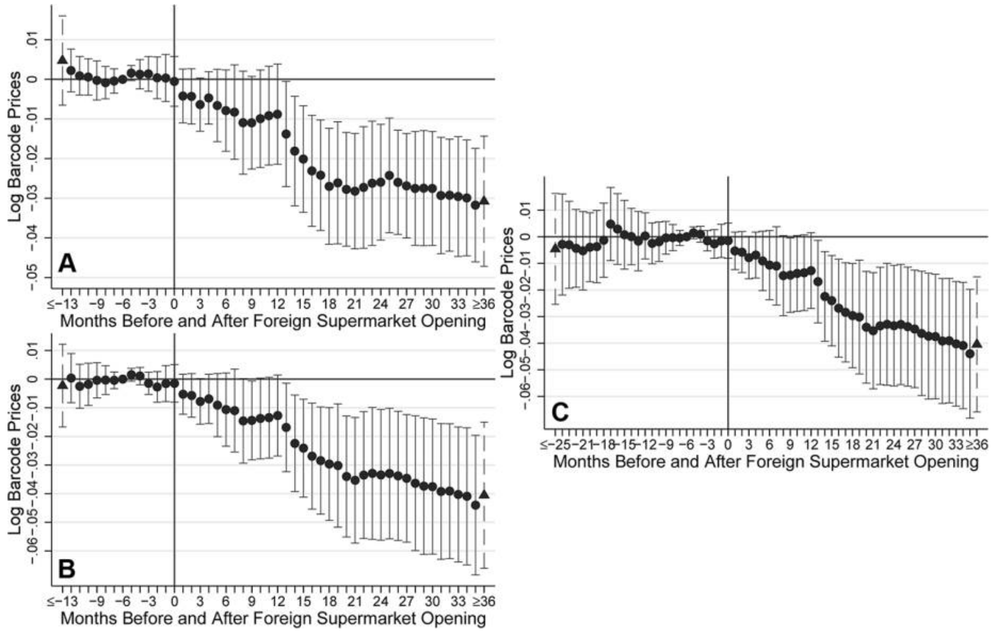

Fig. 2.—Effect on the consumer prices at domestic retailers: monthly event study. Each circle corresponds to a treatment coefficient from a regression of log prices on the 48 (or 60) monthly treatment effects in addition to bar-code-by-store fixed effects and month fixed effects. The triangles denote the coefficients on a dummy for the pre-event study period (≤13 or ≤25) and post-event study period (≥36) included in the same regression. Panel A presents the baseline without additional controls. Panels B and C include additional store-type-by-product-group-by-month as well as region-by-month and municipality-size-by-month fixed effects. Panel C also extends the pre-event study period. The data come from the Mexican CPI price microdata over the period 2002–14 covering 120 retail product groups in 76 urban municipalities, with a sampling probability designed to reflect nationally representative urban consumption weights. The reference category in each graph is bar code prices 6 months before foreign entry. The graphs depict 95 percent confidence intervals based on standard errors clustered at the municipality level.

consumption basket of a representative household, the regressions are implicitly weighted by base period expenditure weights.[23] Hence, these point estimates indicate that foreign retail entry significantly lowers the average household retail price index when using a first-order Laspeyres approximation. The finding that the coefficients fall gradually in the first 2 years after opening rather than immediately suggests that local consumers adjust their shopping behavior gradually, as recently found to be the case for US retailers (Einav, Levin, and Klenow 2015).

In addition to the baseline event study specification, we present two additional event studies that serve as robustness checks. First, in case our results are driven by more granular trends not captured by the month fixed effects, we replace the 141 month fixed effects with 33,516 store-type-by-product-group-by-month fixed effects, 705 region-by-month fixed effects, and 705 municipality-size-by-month fixed effects (panel B of fig. 2).[24] Second, to address any concerns that longer preexisting trends may not be detected in our 12 premonths event study, we also extend the event study to include treatment effects for the 24 months before the opening event (panel C of fig. 2). The coefficient patterns across the three panels are remarkably similar, and the point estimate increases from a 3 percentage point reduction to 4 percentage points after including the additional controls. Columns 1–3 of table 2 present the coefficients in table form (in quarterly rather than monthly bins for compactness).[25]

The absence of preexisting trends in prices relative to locations not receiving their first foreign store, and the subsequent leveling off 2 years after entry, provide no evidence in support of the hypothesis that foreign retailers targeted municipalities on the basis of preexisting price trends during our sample period or entered in response to changing economic conditions before entry. Instead, the results are consistent with a scenario in which foreign retailers rapidly expanded their store networks in order to achieve a store presence across urban Mexico as quickly as possible

[23] We confirm the accuracy of these implicit weights by rerunning specification (13) after weighting each price quote observation so that the total weight of each product group matches the consumption weights in the 2002 ENIGH household expenditure surveys. The point estimates reported in online app. table A.3 are virtually identical.

[24] Store types refer to either modern store formats or traditional retail outlets. Mexican regions are defined by five contiguous geographical zones according to the Instituto Federal Electoral. Similarly, for the size-by-month fixed effects we assign each municipality in our sample to one of five population quintiles that we define over the population in the year 2000. Note that we do not include bar-code-by-month fixed effects since the product descriptions we use to define bar codes are recorded consistently within stores over time, but not necessarily across stores or municipalities.

[25] We also confirm that the results are predominantly driven by single-store entry events as claimed at the beginning of this section. Online app. table A.4 restricts the estimation sample to municipalities where only one foreign store entered over our sample period. Consistent with the vast majority of foreign entry events involving a single store, the estimates are similar in sign, size, and statistical significance.

TABLE 2
EFFECT ON THE PRICES OF DOMESTIC RETAILERS

	Dependent Variable				
	Log Price (1)	Log Price (2)	Log Price (3)	Log Price (Nonretail) (4)	Log Price (Nonretail) (5)
Foreign entry: more than 4 quarters before (unbalanced)	.00378	−.00311	−.00412	.00254	.00112
	(.00612)	(.00719)	(.00716)	(.00512)	(.00689)
Foreign entry: 4 quarters before	.000274	−.00217	−.00343	.00403	.00974
	(.00231)	(.00358)	(.00325)	(.00338)	(.00640)
Foreign entry: 3 quarters before	−.00146	−.00119	−.00208	.00841**	.00971*
	(.00227)	(.00217)	(.00200)	(.00408)	(.00489)
Foreign entry: 2 quarters before (omitted)	0	0	0	0	0
	(0)	(0)	(0)	(0)	(0)
Foreign entry: 1 quarter before	−.000276	−.00271	−.00257	−.00108	.00210
	(.00199)	(.00195)	(.00192)	(.00371)	(.00493)
Foreign entry: 1 quarter after	−.00396	−.00504	−.00495	.00733	.00971
	(.00290)	(.00308)	(.00310)	(.00509)	(.00836)
Foreign entry: 2 quarters after	−.00684**	−.00871**	−.00877**	.00975*	.00882
	(.00311)	(.00402)	(.00405)	(.00533)	(.00622)
Foreign entry: 3 quarters after	−.00999*	−.0129*	−.0133*	.00266	.00160
	(.00566)	(.00690)	(.00696)	(.00456)	(.00326)
Foreign entry: 4 quarters after	−.0110*	−.0147**	−.0152**	−.000294	.00101
	(.00589)	(.00685)	(.00689)	(.00594)	(.00668)
Foreign entry: 5 quarters after	−.0145**	−.0181**	−.0190**	.00739	.00717
	(.00624)	(.00743)	(.00732)	(.00634)	(.00890)
Foreign entry: 6 quarters after	−.0234***	−.0273***	−.0283***	.00636	.00232
	(.00654)	(.00870)	(.00844)	(.00694)	(.00812)
Foreign entry: 7 quarters after	−.0278***	−.0320***	−.0330***	−.00182	−.00684
	(.00708)	(.0101)	(.00986)	(.00732)	(.00747)
Foreign entry: 8 quarters after	−.0281***	−.0347***	−.0354***	−.00145	−.00450
	(.00738)	(.0109)	(.0107)	(.00921)	(.0107)

Foreign entry: 9 quarters after	−.0263***	−.0341***	−.0341***	.0103	.00635
	(.00687)	(.0112)	(.0110)	(.0101)	(.0124)
Foreign entry: 10 quarters after	−.0282***	−.0369***	−.0362***	.0144	.00508
	(.00694)	(.0115)	(.0112)	(.0108)	(.0114)
Foreign entry: 11 quarters after	−.0296***	−.0393***	−.0385***	.00981	−.00168
	(.00704)	(.0117)	(.0113)	(.0114)	(.0111)
Foreign entry: 12 quarters after	−.0313***	−.0425***	−.0414***	.0111	.00117
	(.00727)	(.0119)	(.0115)	(.0133)	(.0135)
Foreign entry: more than 12 quarters after (unbalanced)	−.0316***	−.0412***	−.0392***	.0248	.0121
	(.00835)	(.0127)	(.0119)	(.0162)	(.0160)
p-value (post 3 years = 12 quarters after)	.941	.744	.548	.0077	.057
Month fixed effects	✓	✓	✓	✓	✓
Bar-code-by-store fixed effects	✓	✓	✓	✓	✓
Product-group-by-store-type-by-month fixed effects	X	✓	✓	X	✓
Region-by-month fixed effects	X	✓	✓	X	✓
Municipality-size-by-month fixed effects	X	✓	✓	X	✓
Control for local government expenditure	X	X	✓	X	✓
Observations	3,228,544	2,850,238	2,560,558	1,581,115	1,321,733
R^2	.996	.996	.996	.997	.997
Bar-code-by-store cells	149,273	124,466	114,207	42,715	34,877
Store-type-by-product-by-month cells	33,516	33,516	32,790	8,186	8,027
Region-by-month cells	705	705	690	705	690
Municipality-size-by-month cells	705	705	690	705	690
Municipality clusters	76	76	65	90	65

Note.—The table reports regressions of log prices at domestic stores on 16 quarters-since-foreign-entry treatment effects and two pre- and post-event study dummies. Data come from the CPI price microdata for the period 2002–14 covering 120 retail product groups and 76 urban municipalities, with a sampling probability designed to reflect nationally representative urban consumption weights. The dependent variable is the log price for unique bar-code-by-store combinations. Foreign entry indicates the first presence of a foreign supermarket in the municipality. Regressions include different sets of fixed effects and additional controls as indicated in the table. Columns 4 and 5 show results on nonretail CPI price quotes including transportation, housing, education, health, and other services. Bar-code-by-store fixed effects in those columns refer to item-purveyor identifiers that enumerators track over time (e.g., a haircut from the same hairdresser or a taxi ride on the same route from the same taxi company). Standard errors are clustered at the municipality level and reported in parentheses.

* Significant at the 10 percent level.
** Significant at the 5 percent level.
*** Significant at the 1 percent level.

(and hence variation in opening times is driven by local planning approvals and building delays).

The remaining endogeneity concern is that foreign retailers anticipate breaks in local economic trends. For example, foreign retailers may anticipate local road or other infrastructure investments and target entry to coincide with these investments. We should be clear what would constitute a concern in this context: The local infrastructure investment must be placed at random, in the sense that it is uncorrelated with pre or post trends in prices; it must induce a trend break in prices that lasts only 2 years since prices return to trend after that; and it must be anticipated by the foreign retailer, yet the foreign retailer must always precisely preempt its arrival since we see no drop in prices before entry. Taken individually, each of these three conditions appears unlikely, particularly the last one given the stochastic nature of delays both in opening a new store and in investing in infrastructure. And of course, even if all three conditions are satisfied, it is not obvious why foreign retailers explicitly target places with anticipated negative price shocks.

Nevertheless, we present two additional robustness checks that serve to address these concerns directly. First, we add direct controls for local government expenditures reported by INEGI at the municipality-year level. As presented in table 2, the event study coefficients are virtually unchanged after controlling for municipality-year variation in local public expenditure, providing reassurance that our effects are not driven by foreign stores targeting infrastructure investments. Second, we also estimate the baseline event study specification, equation (13), on the nonretail CPI microdata. These data include price time series for consumer expenditures on, for example, the same local haircut, taxi ride, cleaning service, apartment rent, or medical procedure. This serves two purposes. If we do not expect nonretail prices to respond, it serves as a placebo falsification test as omitted variables that change retail price trends would likely also change nonretail prices. Conversely, if we think that nonretail prices may respond to foreign retail entry through indirect channels, the size of the response is needed for the quantification exercise. As shown in columns 4 and 5 of table 2, the point estimates of the event study specifications run on nonretail prices are a series of precisely estimated zeros. Once again, this placebo result provides reassurance that the timing of foreign entry is exogenous in our specification.

Heterogeneity.—In our theoretical welfare expression, we allow for heterogeneity in price changes across product groups and store formats (i.e., modern vs. traditional).[26] We can estimate these moments since

[26] We also test for heterogeneity across Mexican-US border states (Baja California, Sonora, Chihuahua, Coahuila, Nuevo León, and Tamaulipas) relative to nonborder states and find no significant differences, with the point estimate on the border state interaction very close to zero (0.002 with SE = 0.018).

the confidential version of the CPI microdata provides store formats and retailer names in addition to product groups. We estimate the following specification:

$$\ln p_{gsbmt} = \sum_{gi} \beta_{gi} (\text{Foreign Entry}_{mt} \times \text{Product}_{gi}) + \delta_{gsbm} + \eta_{git} + \theta_{rt} + \phi_{zt} + \epsilon_{gsbmt}, \quad (14)$$

where Foreign Entry$_{mt}$ is an indicator that takes the value of one if there is a foreign store in municipality m in period t and Product$_{gi}$ is an indicator variable that takes the value of one if the retail price quote belongs to product group g and store type i (i.e., modern or traditional). The δ_{gsbm} are bar-code-by-store fixed effects. As before, we also include product-group-by-store-type-by-month (η_{git}), region-by-month (θ_{rt}), and municipality-size-by-month (ϕ_{zt}) fixed effects to control for time-varying product-group-by-store-type-specific shocks to prices and shocks that affect regions or municipality types differently. The β_{gi} estimates capture the effect of foreign entry on domestic retail prices for product group g and store type i. As we are not interested in the very short-run impacts of foreign retail entry, we exclude price data for the first 24 months after entry to capture medium-run price adjustments (recall that the price coefficients in the event study level off 24 months after entry).[27] This specification is subject to similar identification concerns discussed above, and we rely on the lack of pretrends in the previously reported event study.

Table 3 reports the results for food and nonfood items, for modern versus traditional stores, and for the cross of the two. We find that foreign supermarket entry reduces domestic store prices similarly across both food and nonfood products. In contrast, the price reductions appear to be larger for modern domestic store formats than for traditional store formats.

2. Post-entry Price Gaps between Foreign and Domestic Stores

Empirical strategy.—As shown in equation (10), empirical estimates of the post-entry price differences between foreign and domestic stores can be used to estimate a simple and transparent approximation to the direct price index effect (in combination with the pre-post price changes in domestic stores calculated in the previous subsection).

To estimate these post-entry price differences, we compare prices of identical bar codes in the same municipality and month using the

[27] Formally, we cannot reject that the effect 24 months after entry is different from the effect more than 24 months after entry using the specifications in cols. 1–3 of table 2 (p-values of .74, .31, and .56, respectively).

TABLE 3
EFFECT ON THE PRICES OF DOMESTIC RETAILERS: HETEROGENEITY

	DEPENDENT VARIABLE			
	Log Price (1)	Log Price (2)	Log Price (3)	Log Price (4)
Foreign entry	−.0373***			
	(.0119)			
Foreign entry × food		−.0395***		
		(.0137)		
Foreign entry × nonfood		−.0362**		
		(.0154)		
Foreign entry × traditional store			−.0235	
			(.0198)	
Foreign entry × modern store			−.0526***	
			(.0169)	
Foreign entry × food × traditional store				−.00425
				(.0162)
Foreign entry × nonfood × traditional store				−.0287
				(.0231)
Foreign entry × food × modern store				−.0559***
				(.0169)
Foreign entry × nonfood × modern store				−.0497*
				(.0255)
Bar-code-by-store fixed effects	✓	✓	✓	✓
Product-by-store-type-by-month fixed effects	✓	✓	✓	✓
Region-by-month fixed effects	✓	✓	✓	✓
Municipality-size-by-month fixed effects	✓	✓	✓	✓
Observations	2,790,780	2,790,780	2,790,780	2,790,780
R^2	.996	.996	.996	.996
Number of bar-code-by-store cells	123,937	123,937	123,937	123,937
Number of product-by-store-type-by-month cells	33,516	33,516	33,516	33,516
Number of region-by-month cells	705	705	705	705
Number of municipality-size-by-month cells	705	705	705	705
Number of municipality clusters	76	76	76	76

NOTE.—The table reports regressions of log prices at domestic stores on an indicator for foreign entry interacted with indicators for product groups and store types. The data come from the Mexican CPI price microdata over the period 2002–14 covering 120 retail product groups and 76 urban municipalities, with a sampling probability designed to reflect nationally representative urban consumption weights. The basic specification is the same as that reported in col. 2 of table 2 except that the foreign entry effect is averaged across quarters and the estimation sample excludes an adjustment period of 24 months after entry. Columns 2–4 interact foreign entry with indicators for different groups of product and store types. Standard errors are clustered at the municipality level and reported in parentheses.

* Significant at the 10 percent level.
** Significant at the 5 percent level.
*** Significant at the 1 percent level.

consumer panel and the following specification:

$$\ln p_{gsbmt} = \beta_{gi}\text{Domestic Store}_{sgi} + \delta_{gbmt} + \epsilon_{gsbmt}, \tag{15}$$

where Domestic Store$_{sgi}$ is a dummy that takes the value of one if the retailer is not a foreign-owned store and δ_{gbmt} is a bar-code-by-municipality-by-month fixed effect. As in the previous subsection, to account for potential heterogeneity we also allow the coefficient on Domestic Store$_{sgi}$ to vary by food and nonfood categories g and by modern and traditional stores i. Here, and in all our household-level data sets, we weight our regressions using the household survey weights provided to ensure that our results are representative.

Estimation results.—Table 4 presents the estimation results. As reported in Section II, foreign stores charge approximately 12 percent lower prices for identical bar code items compared to domestic stores in the same location during the same month. In terms of heterogeneity, the price advantage of foreign stores is most pronounced compared to traditional domestic retailers (a 17 percent price difference), but the difference remains both economically and statistically significant when comparing foreign stores to modern domestic supermarkets (a 4 percent price difference). In terms of heterogeneity across product groups, the price differences appear to be most pronounced for food relative to nonfood product groups.

B. Effect on Consumption Quantities

1. Post-entry Market Shares of Foreign Retailers

Empirical strategy.—To calculate the direct price index effect in expression (7), we require estimates of the effect of foreign supermarket entry on the retail expenditure shares of foreign stores (broken down by both product group and household income group). To obtain these estimates we again turn to the consumer panel microdata and estimate the following specification:

$$\sum_{s\in S^f_{gmt}} \phi_{gshmt} = \beta_{gh} + \epsilon_{ghmt}, \tag{16}$$

where $\Sigma_{s\in S^f_{gmt}} \phi_{gshmt}$ are retail expenditure shares spent at foreign stores by individual household h in product group g. We estimate the average post-entry expenditure shares at foreign stores, β_{gh}, separately for each of 12 product groups and seven household income groups in the consumer panel data described in Section IV. As above, we restrict our sample to focus only on expenditure shares of foreign stores in locations where a foreign retailer had been open for 24 months or more (recall that the price effects leveled off at 24 months, suggesting that consumer

TABLE 4
Post-entry Price Differences for Identical Bar Codes

	Dependent Variable			
	Log Price (1)	Log Price (2)	Log Price (3)	Log Price (4)
Domestic store	.118***			
	(.00913)			
Domestic store × food		.124***		
		(.00979)		
Domestic store × nonfood		.0744***		
		(.00765)		
Domestic store × traditional			.173***	
			(.00874)	
Domestic store × modern			.0397***	
			(.0113)	
Domestic store × food × traditional				.174***
				(.00942)
Domestic store × nonfood × traditional				.170***
				(.0108)
Domestic store × food × modern				.0431***
				(.0124)
Domestic store × nonfood × modern				.0189***
				(.00713)
Municipality-by-bar-code-by-month fixed effects	✓	✓	✓	✓
Observations	18,659,777	18,659,777	18,659,777	18,659,777
R^2	.923	.923	.923	.923
Number of municipalities	151	151	151	151

Note.—The table reports regressions of log prices at both domestic and foreign stores on an indicator for whether the price is recorded at a domestic store. The data come from the Mexican consumer panel microdata over the period 2011–14. The dependent variable is log bar code prices, and the reference category in all columns is bar code prices in foreign-owned retailers. Columns 2–4 report price differences across different product groups and store types as indicated. Regressions are weighted by household survey weights. Standard errors are clustered at the municipality level and reported in parentheses.
* Significant at the 10 percent level.
** Significant at the 5 percent level.
*** Significant at the 1 percent level.

shopping habits are stable by that time). Thus, this specification estimates medium-run market shares of foreign stores.

Estimation results.—Figure 3 presents the results. On average, foreign stores capture more than 30 percent of total household retail expenditure after entering. Via revealed preference these results provide prima facie evidence of substantial consumer gains from retail FDI in a developing country context.

The consumer panel includes locations that received their first foreign stores prior to the start of our estimation sample in 2002. These early receivers may differ from the locations where foreign entrants arrived during our 2002–14 sample period or may have experienced subsequent shocks that altered their demand for foreign retailers. As a robustness

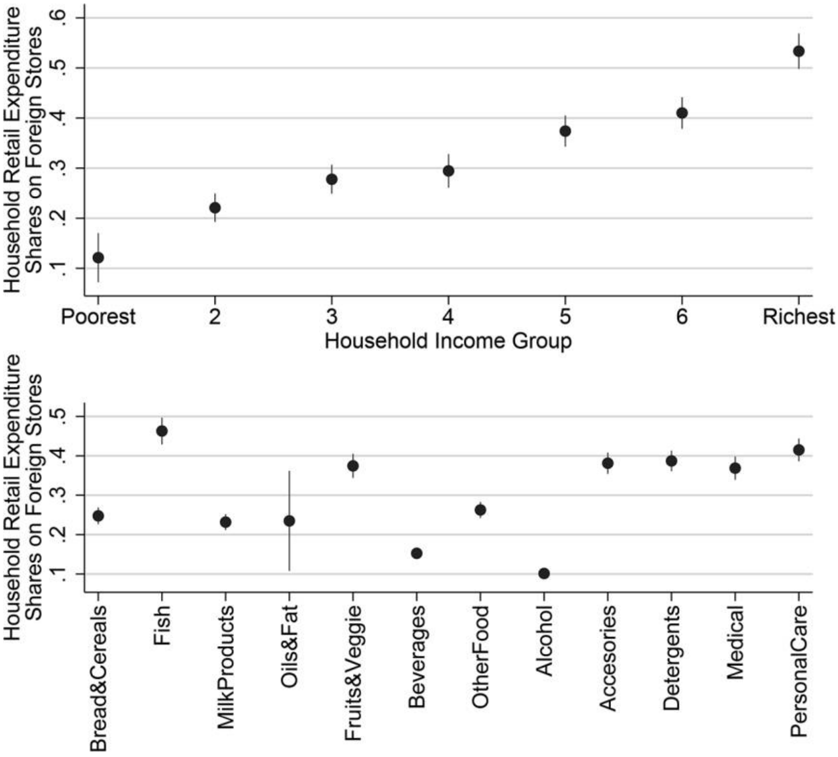

FIG. 3.—Foreign retail market shares after entry. The graphs plot the share of household retail expenditure spent at foreign stores. The data come from the Mexican consumer panel microdata for the years 2011–14. We restrict attention to municipalities where the first foreign store entered more than 2 years previously. Expenditure shares are weighted by household survey weights. Both graphs depict 95 percent confidence intervals based on standard errors clustered at the municipality level.

exercise, we find similar results when we restrict the estimation of equation (16) to municipalities where foreign retail first arrived between 2 and 3 years previously, where such concerns will be less pronounced (see online app. fig. A.1).[28]

The post-entry expenditure shares on foreign stores differ significantly across the income distribution. The upper panel of figure 3 shows that the wealthiest households spend more than 50 percent of their retail expenditure at foreign stores while the poorest spend just over 10 percent. These substantial differences in the extent to which local households sub-

[28] Note that since the consumer panel starts in 2011 and covers large urban municipalities that almost all had foreign stores by that time, we cannot carry out an event study looking at pre- and post-entry market shares of foreign stores.

stitute toward shopping at foreign stores suggest significant differences in how rich and poor households evaluate the amenities and product mix offered by foreign retailers (as captured by the taste shifters and store-income-group-specific price indices in our CES preference structure). Note that these differences across income groups come primarily from variation across income groups within locations rather than simply variation across poor and rich locations: we find similar differences when we include municipality-by-quarter fixed effects in figure A.2 of the online appendix.

The lower panel of figure 3 reports foreign store expenditure shares across 12 product groups. We find substantial differences across product groups; the foreign retail share in personal care products is above 40 percent but below 15 percent for beverages. For the quantification itself, we draw on the cross of the income and product group dimensions and allow foreign retail expenditure shares to be income-group-by-product-group-specific.

2. Effect on Domestic Store Exit

Empirical strategy.—To estimate the procompetitive exit effect in expression (7), we require data on the market shares of exiting stores. Since the consumer panel spans only a short time window and the CPI price data do not contain quantities, we rely on store counts from the microdata of the Mexican retail censuses collected in 2003 and 2008 (and implicitly assume that closing stores had average market shares).[29] Hence, we estimate the following specification (separately for traditional and modern establishments):

$$d \ln N\_\text{Establishments}_m^{03\text{–}08} = \beta d \text{Foreign Entry}_m^{03\text{–}08} + \gamma X_m + \epsilon_m, \quad (17)$$

where the dependent variable is the change in the log number of retail units across the two census rounds and the key independent variable is $d\text{Foreign Entry}_m^{03-08}$, the change in the foreign entry dummy between the two census rounds (i.e., whether the first foreign store opened between censuses). We also include a set of municipality controls X_m, the most important of which is a dummy for whether the municipality already had a foreign store at the time of the first census in 2003 (these municipalities are much larger, are more likely to be located in the center of the country, and presumably experienced different trends compared to municipalities that foreign stores entered in the major wave of expansions

[29] This assumption is likely conservative. If store closures were concentrated in stores with small market shares, then the welfare gains we report in our quantification would be even larger.

we study). In contrast to the household data sets, no survey weights are provided. We report the basic unweighted results as well as results weighting by municipality employment counts from the 2003 Economic Census for consistency with other estimates in the paper.

For β to yield an unbiased estimate of the effect of foreign entry on store counts, we require that foreign entry decisions between 2004 and 2008 are not correlated with other variables that drive changes in the number of local retail establishments. While the event studies in other subsections provide some support for this assumption, the lack of available high-frequency data on local store counts precludes a similar strategy here. To partially address such concerns, we present a number of robustness checks. As in the price regressions, we report estimation results that include additional sets of municipality-level controls: region fixed effects, municipality-size fixed effects, and contemporaneous changes in both log public expenditures and log GDP per capita.

Estimation results.—Table 5 presents the estimation results. Foreign entry has a negative and statistically significant effect on the number of traditional retailers. The size of the preferred point estimate in column 12 (weighted with controls) implies a 3.9 percent reduction. Reassuringly, the point estimate is similar across the various specifications. The coefficient estimate on store exit among modern domestic store formats is also negative and also equal to a 3.9 percent reduction, but it is not statistically significant at conventional levels. Hence, we find negative but moderate effects on domestic store exit over our 5-year time horizon.

C. *Effect on Nominal Labor Incomes, Business Incomes, and Employment*

Empirical strategy.—To calculate the income effect in expression (12), we require estimates of the causal impact of foreign retail entry on nominal incomes and employment in the location where retail entry occurred. To do so, we start by running an event study specification similar to our price event study above, but here using quarterly income and employment data from the employment and occupation survey microdata:

$$\ln \text{Income}_{kmt} = \sum_{j=-5}^{12} \beta_j I(\text{Quarters since Entry}_{mt} = j) + \gamma X_{kmt} + \delta_m + \eta_t + \epsilon_{kmt}, \tag{18}$$

where ln Income_{kmt} is log monthly nominal income for individual k residing in municipality m in quarter t; X_{kmt} are person controls including a gender dummy, dummies for completed degrees (below primary, primary, secondary, and higher), and third-order polynomials for age and years of

TABLE 5
STORE EXIT

	A. Unweighted Regressions: ΔLog(Number Stores) 2003–8							
	Traditional Store Formats				Modern Store Formats			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δforeign entry 2003–8	−.019	−.023	−.025*	−.024*	.0088	−.0065	−.036	−.035
	(.014)	(.014)	(.014)	(.014)	(.067)	(.068)	(.069)	(.069)
Foreign entry pre-2003	−.055**	−.057***	−.035**	−.032**	.20***	.16***	.17***	.17***
	(.013)	(.015)	(.015)	(.016)	(.053)	(.058)	(.060)	(.062)
Δlog(public expenditures)			.12***	.12***			.37***	.38***
			(.028)	(.028)			(.12)	(.12)
Δlog(GDP per capita)				−.020				−.012
				(.014)				(.066)
Geographical region fixed effects	X	✓	✓	✓	X	✓	✓	✓
Municipality size fixed effects	X	✓	✓	✓	X	✓	✓	✓
Observations	608	608	564	564	608	608	564	564
R^2	.022	.056	.107	.110	.015	.085	.107	.107
Median stores/municipality	2,088	2,088	2,088	2,088	33.5	33.5	33.5	33.5

	B. Population-Weighted Regressions: ΔLog(Number Stores) 2003–8							
	Traditional Store Formats				Modern Store Formats			
	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Δforeign entry 2003–8	−.025	−.026	−.039**	−.039**	.022	.0039	−.041	−.039
	(.022)	(.022)	(.018)	(.018)	(.074)	(.076)	(.074)	(.075)
Foreign entry pre-2003	−.091***	−.086***	−.050***	−.051***	.13**	.11*	.12*	.12*
	(.018)	(.019)	(.018)	(.019)	(.060)	(.063)	(.070)	(.069)
Δlog(public expenditures)			.14***	.14***			.17	.17
			(.035)	(.035)			(.14)	(.14)
Δlog(GDP per capita)				.0083				−.048
				(.025)				(.081)
Geographical region fixed effects	X	✓	✓	✓	X	✓	✓	✓
Municipality size fixed effects	X	✓	✓	✓	X	✓	✓	✓
Observations	608	608	564	564	608	608	564	564
R^2	.086	.117	.157	.158	.016	.180	.162	.163
Median stores/municipality	2,088	2,088	2,088	2,088	33.5	33.5	33.5	33.5

Note.—The table reports regressions of changes in log store counts between 2003 and 2008 on whether a foreign store first entered over that time period. The data come from the microdata of the Mexican retail census for the years 2003 and 2008 covering 608 urban municipalities. The dependent variable is the change in log municipality-level retail establishments between 2003 and 2008. Columns 1–4 and 9–12 report results for the traditional domestic retail segment, and cols. 5–8 and 13–16 report results for the modern domestic retail segment. Panel A is unweighted, and panel B weights using municipality employment counts from the 2003 Economic Census. Standard errors are clustered at the state level.

* Significant at the 10 percent level.

** Significant at the 5 percent level.

*** Significant at the 1 percent level.

schooling.[30] We also run an identical specification for employment that replaces ln Income_{kmt} with Employ_{kmt}, an indicator variable that takes the value of one if the person is employed. Hence, the β_j coefficients uncover pre and post foreign entry movements in income and employment.

As discussed above, the event study design allows us to explore preexisting trends in the run-up to the store opening event. Again, there are several potential scenarios. Foreign retailers could target municipalities with either higher or lower preexisting income growth rates. Alternatively, consistent with our price event study, foreign retailers may have expanded rapidly to establishing store presence across urban Mexico in a way that was uncorrelated with local shocks or preexisting trends in incomes or employment (at least for our urban estimation sample).

As in the price event study, we balance the estimation sample between 1 year before and 3 years after the store entry event.[31] In addition to these baseline specifications, we also estimate a number of additional robustness checks. As before, we replace the quarter fixed effects with region-by-quarter as well as municipality-size-by-quarter fixed effects and estimate specifications with both a 1-year pre period and an extended event study with a 2-year pre period.

To estimate the effect of foreign retail entry on business incomes among local store owners, we return to the retail census microdata and estimate specification (17) above after replacing the dependent variable $d\ln N\_\text{Establishments}_m^{03-08}$ by $d\ln(\overline{\text{Profits}_m})^{03-08}$, the change in log mean municipality profits for traditional retail establishments.[32] For the quantification exercise, we require an estimate for the effect on total retail profits accruing to local households (inclusive of the lost profits of exiting traditional stores). This total effect is given by the sum of the profit effect and the store exit effect already calculated in Section V.B.2.

Estimation results.—Figures 4 and 5 present the estimation results for the income and employment event study described above. In contrast

[30] Note that we do not include worker fixed effects as the data set is a rotating panel in which individuals are followed for a maximum of five quarters. Therefore, worker fixed effects sweep out any foreign entry–induced changes in income across cohorts entering the panel at different times. For completeness, online app. table A.5 also reports results with worker fixed effects.

[31] This restriction excludes 10 percent of our observations. The majority (6 percent) of these excluded observations are in urban municipalities that had not yet received a foreign store at the end of our sample in March 2014.

[32] We do not estimate the effect on modern retail profits for two reasons. The first reason is conceptual. Given that we are interested in the welfare effect of foreign entry on local households, the profits of retail chains that are repatriated to their headquarters in other locations should not enter the welfare expression in eq. (12) other than through the shareholdings of local residents, which are likely to be small. The second reason is data constraints. While in the exit regressions we were able to observe the total number of modern stores and subtract the number of foreign-owned stores in the ANTAD data, we cannot do the same for profits since ANTAD does not report them and the census data do not allow us to distinguish between domestic and foreign stores in the modern retail sector.

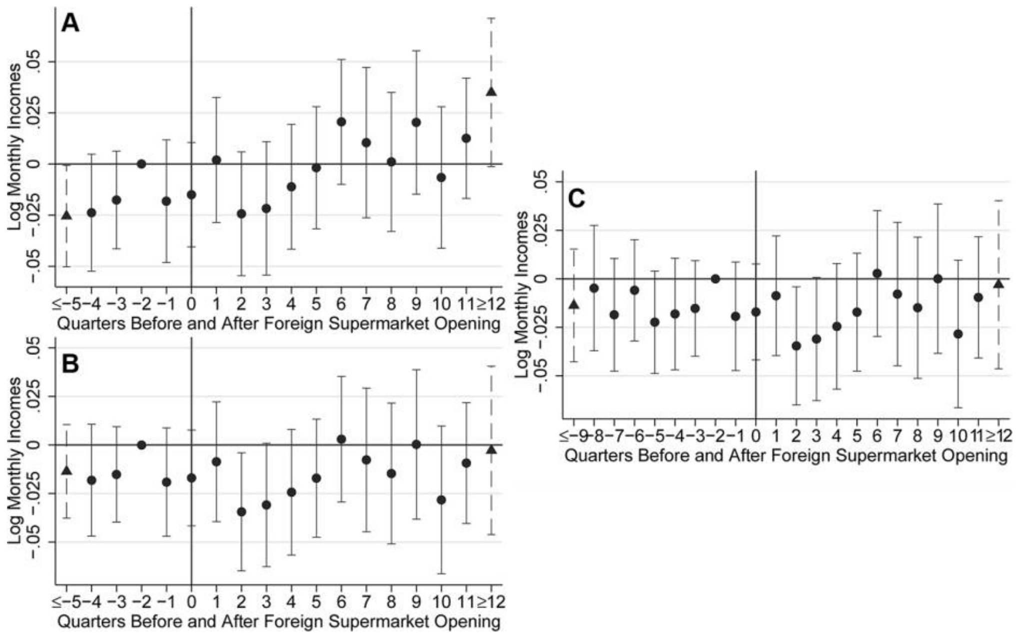

Fig. 4.—Effect on average municipality monthly incomes. Each circle corresponds to a treatment coefficient from a regression of log monthly incomes on 16 (or 20) quarterly treatment effects in addition to municipality and quarter fixed effects, as well as person controls for sex, education, and age. The triangles denote the coefficients on the dummy for the pre-event study period (≤5 or ≤9) and post-event study period (≥12) included in the same regression. Panel A presents the baseline specification, while panels B and C include additional region-by-quarter and municipality-size-by-quarter fixed effects. Panel C also extends the pre-event study period. The data come from the 273 urban municipalities in employment and occupation surveys over the period 2002–12. The reference category is incomes two quarters before foreign entry. Regressions are weighted by household survey weights. The graphs depict 95 percent confidence intervals based on standard errors clustered at the municipality level.

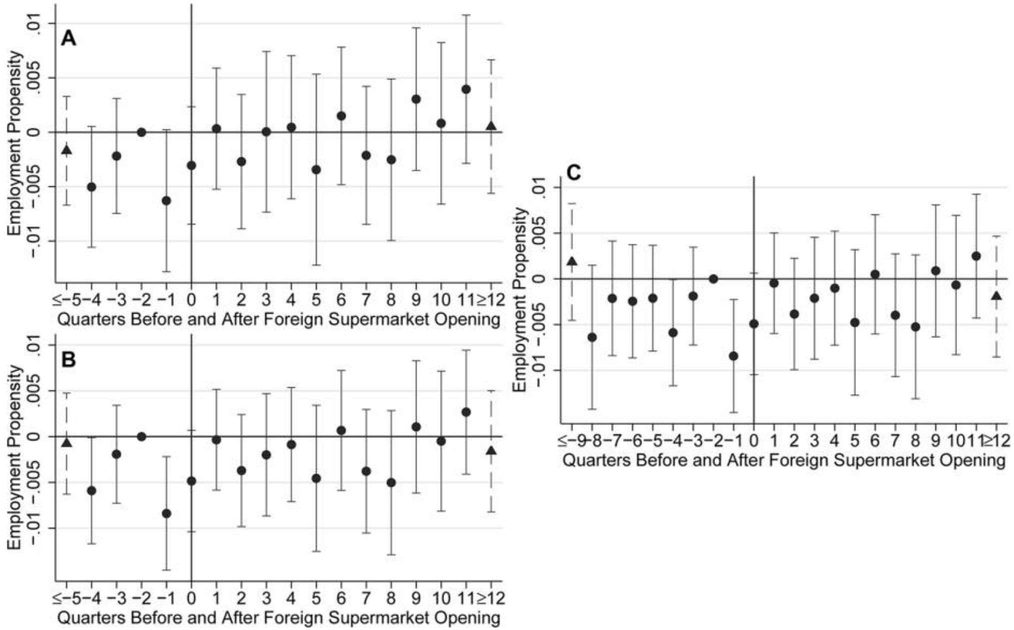

FIG. 5.—Effect on average municipality employment. Each circle corresponds to a treatment coefficient from a regression of individual employment indicators on 16 (or 20) quarterly treatment effects in addition to municipality and quarter fixed effects, as well as person controls for sex, education, and age. The triangles denote the coefficients on the dummy for the pre-event study period (≤5 or ≤9) and post-event study period (≥12) included in the same regression. Panel A presents the baseline specification, while panels B and C include additional region-by-quarter and municipality-size-by-quarter fixed effects. Panel C also extends the pre-event study period. The data come from the 273 urban municipalities in the employment and occupation surveys over the period 2002–12. The reference category is employment propensities two quarters before foreign entry. Regressions are weighted by household survey weights. The graphs depict 95 percent confidence intervals based on standard errors clustered at the municipality level.

to the price event study, we find no evidence of either jumps in levels or breaks in trends around the period of foreign retail entry (or evidence of preexisting trends). There appear to be no general equilibrium income or employment effects in the municipality, perhaps not surprisingly given that one store hires only a small number of employees. For completeness, table A.5 in the online appendix also provides the regression table. Table A.5 also shows that we find no evidence of changes in population, at least over the time horizon we study: we aggregate the worker weights up to the municipality level and regress log municipality population counts on quarters-since-entry dummies, municipality fixed effects, and quarter fixed effects and find statistically insignificant point estimates that are close to zero.

Table 6 presents the estimation results for domestic store profits using the retail census microdata. Foreign entry has a negative effect on retail profits for traditional store owners. The significance of the point estimate on profits depends on whether the specification is population weighted or not but ranges between −4.4 and −5.1 percent. Reassuringly, this point estimate varies little when we include regional fixed effects, include initial municipality size fixed effects, or control for contemporaneous changes in local government expenditures or GDP per capita.

Heterogeneity.—To quantify to what extent households may be affected differently depending on their primary source of income before foreign entry, we require not just the general equilibrium income and employment effects but also estimates broken down by sector. We now turn to exploring this heterogeneity. Since individuals may change sector or become unemployed over time, and the employment and occupation surveys follow individuals only over five consecutive quarters, we cannot directly assess longer-run outcomes for workers with different pre–foreign entry occupations. Instead, we calculate both the average income and employment changes across various sectors and apply the decomposition outlined in equation (12) of the theoretical framework.

To obtain average changes in nominal incomes across sectors, we regress individual log income on a foreign entry dummy that takes the value one when there is a foreign store in the municipality interacted with a sector dummy that takes the value one if a worker is employed in that sector:

$$\begin{aligned}\ln(\text{Income})_{kimt} = {} & \sum_i \beta_i(\text{Foreign Entry}_{mt} \times \text{Sector}_i) + \gamma X_{kimt} \\ & + \delta_{mt} + \eta_{im} + \theta_{it} + \epsilon_{kimt},\end{aligned} \tag{19}$$

where subscripts k, i, m, and t index individuals, sectors, municipalities, and quarters, respectively; δ_{mt} is a municipality-by-quarter fixed effect; η_{im} is a sector-by-municipality fixed effect; and θ_{it} is a sector-by-quarter fixed effect. Our sectoral categories consist of two retail categories, retail

TABLE 6
Effect on Store Profits

	Dependent Variable: ΔLog(Mean Profit) 2003–8							
	Unweighted Regressions				Population-Weighted Regressions			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δforeign entry 2003–8	−.049*	−.047	−.048	−.051*	−.039	−.041	−.043	−.044
	(.028)	(.029)	(.030)	(.030)	(.032)	(.033)	(.034)	(.034)
Foreign entry pre-2003	−.087***	−.082***	−.071***	−.081***	−.070**	−.074**	−.043	−.047
	(.024)	(.024)	(.026)	(.027)	(.031)	(.032)	(.032)	(.034)
Δlog(public expenditures)			.042	.038			.079*	.075*
			(.046)	(.045)			(.044)	(.045)
Δlog(GDP per capita)				.061*				.034
				(.035)				(.030)
Geographical region fixed effects	X	✓	✓	✓	X	✓	✓	✓
Municipality size fixed effects	X	✓	✓	✓	X	✓	✓	✓
Observations	608	608	564	564	608	608	564	564
R^2	.014	.061	.064	.071	.020	.071	.112	.115
Median number of stores per municipality in 2003 and 2008	2,088	2,088	2,088	2,088	2,088	2,088	2,088	2,088

Note.—The table reports regressions of changes in log profits between 2003 and 2008 on whether a foreign store first entered over that time period. The data come from the microdata of the Mexican retail census for the years 2003 and 2008 covering 608 urban municipalities. The dependent variable is the change in log mean municipality profits between 2003 and 2008 among traditional retail establishments. Columns 1–4 are unweighted and cols. 5–8 weight using municipality employment counts from the 2003 Economic Census. Standard errors are clustered at the state level.

* Significant at the 10 percent level.

** Significant at the 5 percent level.

*** Significant at the 1 percent level.

workers in either modern or traditional store formats, as well as two nonretail categories, individuals whose main income source is either agriculture or manufacturing. The omitted category contains individuals whose main income source is nonretail services (e.g., education, health, restaurant, or domestic services).[33] We assume that this large omitted category experiences no income changes with foreign entry based on the flat event study plots for average incomes above.

The coefficients β_i capture the differential effect of foreign store entry on the incomes of various sectors (conditional on flexible trends at the municipality-quarter level, initial earnings differences across sectors within the municipality, and national differences across sectors in that quarter). Finally, we remove observations covering the 24 months after the first foreign store entry to avoid capturing the short-run adjustment period we noted in the price regressions above.

Table 7 presents these results. In contrast to the average income regressions, we find a negative and significant effect on the incomes of traditional retail workers. This point estimate is robust to including income-group-by-quarter fixed effects as well as state-by-income-group-specific time trends, implying that this effect is not driven by preexisting trends that are specific to particular income groups. The point estimate corresponds to a reduction in the monthly incomes of traditional sector retail workers of 5.9 percent as a result of foreign retail entry. We find a smaller and insignificant 2.8 percent decline in the labor incomes of modern retail workers. In contrast, incomes in agriculture and manufacturing—individuals who may be supplying foreign retailers—rise but by small and insignificant amounts.

Finally, we turn to employment changes across sectors. We regress the log number of employed workers in municipality m, sector i, and quarter t on all the variables on the right-hand side of equation (19) except for the vector X_{kimt} of individual-level controls. The final three columns of table 7 report these coefficients. We find a significant and substantial reduction in traditional retail employment of 11.3 percentage points. This is partially compensated by an (insignificant) 3.9 percent increase in employment in the modern sector, potentially coming from employment at the foreign store itself. The employment changes in agriculture and manufacturing are small and insignificant (once again relative to the omitted category, nonretail services).

[33] Although these data also report business incomes in retail, we exclude the small fraction of retail business owners. The reason is that we already have an estimate for the effect on retail profits using the retail census data, where the profit data are more reliable for self-employed business owners and the number of observations is far larger. (The employment and occupation surveys contain a median of only nine store owners in each municipality-quarter cell.)

TABLE 7
EFFECT ON INCOMES: HETEROGENEITY

	Dependent Variable					
	Log(Monthly Income) (1)	Log(Monthly Income) (2)	Log(Monthly Income) (3)	Log(Employment) (4)	Log(Employment) (5)	Log(Employment) (6)
Foreign entry × modern retail workers	−.000278	−.0348*	−.0278	−.00396	.0369	.0392
	(.0192)	(.0204)	(.0212)	(.0653)	(.0714)	(.0561)
Foreign entry × traditional retail workers	−.0356*	−.0571***	−.0592**	−.104*	−.0942	−.113**
	(.0199)	(.0216)	(.0240)	(.0531)	(.0571)	(.0552)
Foreign entry × agriculture	.0265	.0218	.0202	−.0597	−.0285	−.00811
	(.0264)	(.0311)	(.0307)	(.0809)	(.101)	(.106)
Foreign entry × manufacturing	−.00513	−.00612	.0117	−.166***	.00572	−.0166
	(.0174)	(.0186)	(.0187)	(.0379)	(.0368)	(.0380)
Person controls	✓	✓	✓	X	X	X
Municipality-by-quarter fixed effects	✓	✓	✓	✓	✓	✓
Municipality-by-group fixed effects	✓	✓	✓	✓	✓	✓
Group-by-quarter fixed effects	X	✓	✓	X	✓	✓
State-by-group time trends	X	X	✓	X	X	✓
Observations	3,878,561	3,878,561	3,878,561	47,666	47,666	47,666
R^2	.340	.340	.341	.963	.965	.967
Number of individuals	1,455,911	1,455,911	1,455,911	1,455,911	1,455,911	1,455,911
Number of municipality-by-quarter cells	8,574	8,574	8,574	8,574	8,574	8,574
Number of state-by-group time trends	160	160	160	160	160	160
Number of municipality clusters	273	273	273	273	273	273

NOTE.—The table reports regressions of log incomes and log number of employees on an indicator for foreign entry interacted with sectoral dummies. Data come from the 273 urban municipalities in the employment and occupation surveys over the period 2002–12. The dependent variable in cols. 1–3 is individual log monthly incomes. The dependent variable in cols. 4–6 is the log number of employed individuals by sector. The regressions include different combinations of fixed effects and controls as indicated in the table. Columns 1–3 weight regressions by household survey weights, whereas cols. 4–6 use the survey weights to compute the sum of employment by sector and then weight the regressions by the sum of household survey weights at the municipality level. Standard errors are clustered at the municipality level and reported in parentheses.

* Significant at the 10 percent level.
** Significant at the 5 percent level.
*** Significant at the 1 percent level.

VI. Quantifying the Welfare Effect of Foreign Retail

In order to quantify the effects of foreign retail entry on local household welfare using equations (7) and (12), we require two more inputs beyond the causal effects on prices, quantities, and incomes we estimate above. First, we need an estimate of the elasticity of substitution across local retail outlets (the η_{gh} parameters in eq. [7]). Second, we require estimates of household budget shares across product groups (the α_{gh} parameters in eq. [7]), pre-entry store type expenditure shares (to calculate the ω_{gsh} weights in eq. [7]), and pre-entry income shares across sectors and occupations (the θ_i^0 terms in eq. [12]). Armed with all three types of moments, we then proceed to the quantification.

A. Estimation of the CES Elasticity Parameter

To estimate the elasticity of substitution across local retail outlets we use the consumer panel microdata. The strength of these data is that we observe how much each household purchases of a particular bar code equivalent product, at what price, and at which specific store. Thus, we can observe how store-level market shares vary with store-level prices across various locations. We exploit this cross-location variation (rather than the time series) for two reasons: this variation is more likely to provide estimates of the long-run elasticity relevant for estimating the gains from new foreign retail store openings, and the consumer panel has a relatively short duration (10 quarters between 2011 and 2014).

To derive our estimating equation, note that in the CES case the log expenditure share of store brand s (e.g., Walmart) within product group g (e.g., beverages) can be written as

$$\ln \phi_{gshmt} = (1 - \eta_{gh}) \ln r_{gshmt} - (1 - \eta_{gh}) \ln c_{ghmt} + \eta_{gh} \ln \beta_{gshmt}, \quad (20)$$

where

$$c_{ghmt} = \left(\sum_{s \in S_{gmt}} \beta_{gshmt}{}^{\eta_{gh}} r_{gshmt}{}^{1-\eta_{gh}} \right)^{1/(1-\eta_{gh})}$$

is the product group–specific CES price index, r_{gshmt} is the store-product-group-specific price index, and β_{gshmt} is a store-product-group-specific taste shifter for household h in municipality m in quarter t. Thus, if we regress log expenditure shares on local store-level price indices, use fixed effects to sweep out the CES price index, and deal with endogeneity due to the taste shifter, we can recover the elasticity of substitution across stores, η_{gh}.

To implement this procedure, we must place additional structure on the multitier preference structure we introduced in equation (3). First, since we will not be able to match individual households in the consumer panel to households in the income and expenditure surveys, we aggregate households into broad income groups (the seven income groups in the consumer panel). Second, we must impose some discipline on the store-specific taste shifters for them not to soak up all the variation in expenditure shares. Our basic specification allows tastes to differ both across time for each retail-chain-product-group pair (e.g., because of national advertising campaigns for a retail chain) and across municipalities for each retail chain (e.g., because a retail chain is located on the outskirts of some municipalities and at the center of others). We also report additional specifications that allow tastes to vary at the retail-chain-municipality-quarter level or retail-chain-municipality-product-group level.

Finally, in order to calculate the store-level price index we require a functional form for the lowest tier of consumer preferences—household preferences across products within a store product group—that we left unspecified up to now. In principle, we could use any demand system. For simplicity and transparency, we choose the widely used Stone price index, which is just a budget share–weighted sum of log prices: $\ln r_{gshmt} = \sum_{b \in B_{gshmt}} \phi_{gsbhmt} \ln p_{gsbhmt}$. As products differ across stores and some stores may sell higher-quality varieties, we use bar code fixed effects to ensure that we are comparing only identical products to extract these price index differences. To be precise, we recover $\ln r_{gshmt}$ from the store fixed effects in a regression of budget share–weighted log prices at the bar code level on both store and bar code fixed effects, where this regression is run separately for every product-group-income-group-municipality-time cell.[34] This procedure is similar to extracting firm fixed effects from matched employer-employee data, and we follow that literature by estimating price indices off the largest "mobility group" of connected stores and bar codes.

Given this additional structure, we can run the following regression:

$$\ln \phi_{gshmt} = b_{gh} \ln r_{gshmt} + \delta_{ghmt} + \gamma_{cgt} + \gamma_{cm} + u_{gshmt}, \tag{21}$$

where u_{gshmt} is the error term, δ_{ghmt} are product-group-by-income-group-by-municipality-by-time fixed effects that sweep out the CES price index $\ln c_{ghmt}$, and the γ terms are retail chain c fixed effects to capture the unobserved taste shifters. The object of interest is $1 - b_{gh}$, an estimate of the elasticity η_{gh} that governs the degree of substitutability between local retail outlets as a function of store price differences. To fix ideas via a simple example, we are essentially comparing one location where Soriana has

[34] Given the nature of the microdata, we first collapse the price data to average prices at the bar-code-store-municipality-quarter level. Alternatively, we collapse them to median prices, and we will report both specifications. For comparability across locations, the resulting store fixed effects are then demeaned within each product-group-income-group-municipality-time cell.

relatively high prices for beverages compared to Walmart with another location where the two stores have similar prices and inferring the elasticity of substitution from the difference in relative market shares across the two locations. As in previous sections, we allow this elasticity parameter to differ across food and nonfood retail product groups as well as across two broad income groups, "rich" and "poor" (defined as above and below median).[35]

A shortcoming of the consumer panel data is that unique store brand identifiers are recorded only for the modern retail sector. For this reason, we are restricted to calculating elasticities of substitution from the market shares of retailers in the modern sector. Given that similar store formats are presumably closer substitutes, this approach likely generates higher η_{gh} estimates (and hence lower direct price index effects). We will explore the sensitivity of our estimates to allowing the elasticity of substitution to differ both across and within modern and traditional store types.

As in any demand estimation, there are simultaneity concerns if demand shocks in the error term raise both store-level market shares and store-level price indices. To deal with this concern we follow Hausman (1997) and instrument the store-product-group price index with price indices in stores of the same retailer in other municipalities. In particular, we exploit the fact that there is both a local and national/regional component to prices at supermarket chains—with the local component potentially related to idiosyncratic local demand shocks and the national/regional component related to common suppliers and distribution networks as well as national/regional pricing rules (see, e.g., Beraja, Hurst, and Ospina 2014; DellaVigna and Gentzkow 2017).[36] This supply-side price variation allows us to identify the elasticity of substitution. Thus, we instrument using product group–specific price indices constructed either from national leave-out means for that retailer or from regional leave-out means (using five Mexican administrative regions). As recently shown by Beraja et al. (2014), these two instruments identify potentially different local average treatment effects. The national leave-out means estimate the elasticity of substitution based on retail chains that have com-

[35] Despite the richness of these data, the variation thins if we allowed for heterogeneity along more granular dimensions of h and g (we observe on average only 40 households in a given quarter-municipality cell).

[36] In support of this claim in the Mexican context, we find that in the consumer panel data the variance of log prices within a retailer-bar-code-quarter is 0.458 nationwide, 0.395 within a given region, and only marginally lower at 0.347 within a given state. When including cross-retailer variation (i.e., within bar-code-quarter), the variances of log prices are all higher (0.554 nationwide, 0.522 regionwide, and 0.476 statewide). More directly, prices are correlated across stores within a chain: a regression of log prices on retailer-bar-code-municipality fixed effects and the leave-out mean of prices for that bar code at that store in other locations nationwide results in a coefficient of 0.387 (and attenuated coefficients of 0.228 regionwide and 0.160 statewide).

mon national supply shocks or pricing rules, whereas the region-level leave-out means extend the complier group of the instrumental variable (IV) to regional supply shocks and pricing rules.

Table 8 presents the estimation results both for the average elasticity of substitution and allowing the elasticity to vary by income and product group. For both IV strategies and all three different taste shifter fixed-effect specifications, the elasticity estimates are consistently negative and significant and have substantial first-stage *F*-statistics. The average elasticities range between 2.28 and 4.36, well within the range of existing estimates that use scanner consumption microdata for the United States (e.g., Hausman and Leibtag 2007; Handbury and Weinstein, forthcoming). As is evident from the cost-of-living expression, equation (7), larger elasticity estimates result in smaller direct price index effects. Therefore, we base our welfare quantification on the specification that yields the most conservative (highest) average estimate of the elasticity of substitution across local retail outlets (those reported in cols. 9 and 23 of table 8). In addition to choosing the most conservative η_{gh} estimates, we also rerun our quantification both across a range of alternative average elasticities of substitution for which our preferred η_{gh} estimates form the midpoint and using the first-order approximation approach outlined in Section III that does not require estimates of this elasticity. Finally, online appendix C performs a variety of additional robustness checks (including controls for variety, using an unweighted price index, using only cross-sectional variation, restricting attention to a subset of common products, and restricting attention to stores with strong national pricing strategies), with the resulting elasticities all falling within this range.

B. *Combining the Estimated Moments for the Quantification*

In order to calculate the welfare expressions derived in Section III, we still require one final set of moments: estimates of pre-entry income and expenditure shares across the distribution of households. For these moments we draw on our sixth and final data set, the Mexican household income and expenditure surveys. The household-level shares from these data, combined with the previous moments estimated at the income-group-by-product-group level, will allow us to estimate the welfare gains from foreign retail entry separately for each household in the survey. As we require pre-entry income and expenditure shares, we restrict attention to the 12,293 households surveyed over the period 2006–12, who reside in the 240 urban municipalities without foreign stores at the time of the survey.[37]

[37] A potential concern is that the locations in our quantification sample may differ from the locations that did not have a foreign store at the start of 2002, the baseline sample we estimate various moments from. Therefore, online app. table A.6 presents quantification results that reweight the household sample so that the distribution of either municipality

On the income side, the surveys record the share of income from each sector and occupation (business owner or employee), which provides the pre-entry income shares (the θ_h^0 parameters in eq. [12]). Alongside the causal income, employment, and profit changes presented in Section V, we have all the moments necessary to quantify the three nominal income effects in equation (12).

To quantify the retail profit effect, we randomly assign a profit loss of 100 percent to a random 3.9 percent of traditional store owners in the income and expenditure surveys (based on the estimate of traditional store exit from Sec. V.B.2). We assign the remaining 96.1 percent of traditional store owners a profit loss of 4.4 percent (based on the traditional store profit decline estimated in Sec. V.C). To quantify the labor income effects, we use the estimates of the sectoral income and employment effects 2 years after foreign entry in columns 3 and 6 of table 7 to compute the expressions in equation (12) and footnote 17 (assigning an income loss of 100 percent to a randomly chosen fraction of households equal to the estimated employment decline and assigning the estimated wage decline to the remaining households).[38]

We now turn to the cost-of-living expression under the exact (CES) approach, expression (7). For each household, the income and expenditure surveys record household expenditures in each product group (the α_{gh} parameters). In Section V.B we obtained estimates of the share of retail expenditure at foreign stores after entry ($1 - \Sigma_{s \in S_g^{dc}} \phi_{gsh}^1$) by income group and product group (shown visually in fig. 3) and at exiting stores before entry ($1 - \Sigma_{s \in S_g^{dc}} \phi_{gsh}^0$) (cols. 12 and 16 of table 5). Coupled with the elasticity of substitution across stores η_{gh} estimated in Section VI.A above, we are in a position to calculate the key components of the direct price index effect and the procompetitive exit effect.

The remaining cost-of-living term is the procompetitive price effect. This is a function of causal changes in store price indices, r_{gsh}^1 / r_{gsh}^0, and ideal log change weights, ω_{gsh}. Recall that in Section V.A, we obtained estimates of causal price changes at domestic retail stores 2 years after foreign entry for food and nonfood groups in both traditional and modern stores (col. 4 of table 3). Accordingly, we assume that all product-level price changes

populations or supermarkets per capita matches the distribution in the baseline sample. The magnitudes of the welfare effects are almost identical.

[38] More precisely, households with members employed in a given sector (traditional retail, modern retail, agriculture, or manufacturing) receive their initial income share in that sector multiplied by the average income effect estimated in col. 3 of table 7. For sectors experiencing employment reductions of x percent in col. 6, we apply a 100 percent reduction in income to a random x percent of households working in that sector. For sectors experiencing gains of x percent, we increase income for a random $x\rho/(1 - \rho)$ percent of households not working in that sector by the average local earnings in that sector divided by total household income (where ρ is the proportion of local employment in that sector). We set the employment change to zero in agriculture given the point estimate of 0.00811 and the large standard errors due to the small number of agricultural employees in our urban sample.

TABLE 8
DEMAND PARAMETER ESTIMATES

Dependent Variable: Log Budget Shares (ϕ): A. Average Coefficient Estimates

	Average Prices OLS (1)	Median Prices OLS (2)	Average Prices National IV (3)	Median Prices National IV (4)	Average Prices Regional IV (5)	Median Prices Regional IV (6)	Average Prices National IV (7)	Median Prices National IV (8)	Average Prices Regional IV (9)	Median Prices Regional IV (10)	Average Prices National IV (11)	Median Prices National IV (12)	Average Prices Regional IV (13)	Median Prices Regional IV (14)
Log(store price index)	.214***	.205***	−1.341***	−1.281***	−1.856***	−1.578***	−2.648***	−2.421***	−3.362***	−2.736***	−1.913***	−1.821***	−2.367***	−2.052***
	(.006)	(.006)	(.145)	(.133)	(.608)	(.477)	(.338)	(.29)	(1.038)	(.747)	(.241)	(.219)	(.862)	(.689)
Product group-by-income group-by-municipality-by-quarter fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Retailer-by-product group-by-quarter fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Retailer-by-municipality fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Retailer-by-municipality-by-quarter fixed effects	X	X	X	X	X	X	✓	✓	✓	✓	X	X	X	X
Retailer-by-municipality-by-product group fixed effects	X	X	X	X	X	X	X	X	X	X	✓	✓	✓	✓
Observations	304,885	304,885	304,885	304,885	297,624	297,624	304,885	304,885	297,624	297,624	304,885	304,885	297,624	297,624
First-stage F-statistic			184.884	201.035	14.833	18.868	87.951	103.207	15.52	22.113	106.577	117.084	10.394	12.882

Dependent Variable: Log Budget Shares (ϕ): B. Heterogeneous Coefficient Estimates

	Average Prices OLS (15)	Median Prices OLS (16)	Average Prices National IV (17)	Median Prices National IV (18)	Average Prices Regional IV (19)	Median Prices Regional IV (20)	Average Prices National IV (21)	Median Prices National IV (22)	Average Prices Regional IV (23)	Median Prices Regional IV (24)	Average Prices National IV (25)	Median Prices National IV (26)	Average Prices Regional IV (27)	Median Prices Regional IV (28)
Log(store price index) × poor × food	.244***	.232***	−1.235***	−1.186***	−1.567***	−1.344***	−3.352***	−3.063***	−2.867***	−2.349***	−1.842***	−1.776***	−3.034***	−2.436**
	(.011)	(.011)	(.184)	(.174)	(.548)	(.439)	(.615)	(.528)	(.92)	(.679)	(.345)	(.324)	(1.748)	(1.225)

Log(store price index) × poor × nonfood	.21***	.206***	−.65***	−.597***	−1.602***	−1.352***	−1.688***	−1.496***	−2.854***	−2.317***	−.577***	−.511***	−.801***	−.784***
	(.008)	(.008)	(.1)	(.092)	(.571)	(.464)	(.308)	(.26)	(.89)	(.656)	(.124)	(.113)	(.304)	(.295)
Log(store price index) × rich × food	.21***	.201***	−1.621***	−1.581***	−1.814***	−1.58***	−4.008***	−3.697***	−3.267***	−2.717***	−2.274***	−2.219***	−3.341***	−2.716***
	(.009)	(.008)	(.205)	(.193)	(.573)	(.457)	(.692)	(.595)	(.988)	(.727)	(.384)	(.36)	(1.834)	(1.284)
Log(store price index) × rich × nonfood	.154***	.152***	−1.068***	−1.018***	−1.988***	−1.717***	−2.247***	−2.044***	−3.342***	−2.77***	−1.05***	−.983***	−1.165***	−1.146***
	(.008)	(.008)	(.11)	(.102)	(.623)	(.507)	(.344)	(.29)	(.968)	(.715)	(.141)	(.129)	(.337)	(.327)
Product group-by-income group-by-municipality-by-quarter fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Retailer-by-product group-by-quarter fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Retailer-by-municipality fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Retailer-by-municipality-by-quarter fixed effects	X	X	X	X	X	X	✓	✓	✓	✓	X	X	X	X
Retailer-by-municipality-by-product group fixed effects	X	X	X	X	X	X	X	X	X	X	✓	✓	✓	✓
Observations	304,885	304,885	304,885	304,885	297,624	297,624	304,885	304,885	297,624	297,624	304,885	304,885	297,624	297,624
First-stage F-statistic			31.287	33.191	3.915	4.894	10.783	12.694	4.14	5.81	13.586	14.534	.98	1.391

NOTE.—The table reports regressions of log budget shares on log store-specific price indices, with the coefficient corresponding to one minus the elasticity of substitution across stores. The data come from the Mexican consumer panel microdata over the period 2011–14. The dependent variable is log expenditure shares across stores within a municipality, quarter, product group, and income group cell. The independent variable is log store-specific price indices for that municipality, quarter, product group, and income group cell. The indices themselves are recovered from store fixed effects in a regression of budget share-weighted log prices at the bar code level on both store and bar code fixed effects as described in Sec. VI.A. All stages use household survey weights. Panel B reports results in which the coefficient on the price index is allowed to vary by income and product group. Standard errors are clustered at the municipality level.

* Significant at the 10 percent level.

** Significant at the 5 percent level.

*** Significant at the 1 percent level.

within these product-group-by-store types are the same, $p^1_{gsb}/p^0_{gsb} \equiv p_{gs}$, where the store s now indexes two types, modern domestic μ and traditional domestic τ; and product group g now indexes two groups, food retail and nonfood retail. We later relax this assumption and allow for different price changes across bar-coded and non-bar-coded products. Given that we found a precise zero on nonretail consumer prices (col. 5 of table 2), we set the price changes equal to zero across all nonretail product groups.

The ideal log change weights, ω_{gsh}, that weight these domestic store price changes in the CES price index can also be calculated. These are simple functions of pre- and post-entry expenditure shares for each store type within each product group. Fortunately, the design of the income and expenditure surveys are extremely helpful in this regard. As described in Section IV, the four survey rounds we use break down household expenditures in each product group into expenditures at different types of stores. This breakdown allows us to directly observe the pre-entry store type expenditure shares, $\tilde{\phi}^0_{gsh}$, for each household and product group. We then use the CES structure to calculate post-entry shares, $\tilde{\phi}^1_{gih}$, by combining these pre-entry shares with the causal price changes by store type discussed above and our elasticity of substitution estimates:

$$\tilde{\phi}^1_{gih} = \tilde{\phi}^0_{gih}\left[\frac{(p_{gi})^{1-\eta_{gh}}}{(p_{g\mu})^{1-\eta_{gh}}\tilde{\phi}^0_{g\mu h} + (p_{g\tau})^{1-\eta_{gh}}\tilde{\phi}^0_{g\tau h}}\right], \tag{22}$$

where i takes the value μ (modern domestic) or τ (traditional domestic).

Finally, the cost-of-living expression under the first-order approach is relatively straightforward to quantify. We require two sets of price changes: the pre-post price changes in domestic stores for the procompetitive effect (expression [9]) and the sum of these pre-post price changes and the post-entry price gaps between domestic and foreign stores for the direct price index effect (expression [10]). As above, we assume that these price changes and price gaps were common across all products within each product-group-by-store-type pair and draw on our estimates from table 3 (price changes) and table 4 (price gaps). For the post-entry expenditure weights, ϕ^1_{bsh}, that weight these price changes and price gaps in the Paasche indices, we start with the product group shares for each household from the income and expenditure surveys. These are then multiplied by the domestic (for the procompetitive price effect) or foreign (for the direct price index effect) market shares within each product group, as well as the store type shares within each product group, both obtained from households in the same income bin in the post-entry period consumer panel data.

Before performing the quantification, we need to confront the fact that the estimates of the causal effects of foreign retail entry on consumer prices, quantities, and incomes that enter the quantification are subject to sampling error. To obtain standard errors and confidence intervals that take this error into account, we bootstrap the entire quantification exercise 1,000 times. In each bootstrap, we both draw a random sample from our 12,293 households (sampling with replacement) and redraw each price, quantity, and income parameter from a normal distribution with a mean equal to the point estimate and a standard deviation equal to the standard error of the estimate.[39] Finally, as mentioned above, for robustness we recompute the quantification for a range of elasticity estimates (recall we chose the most conservative specification in our baseline quantification) and for a range of assumptions on price changes for non-bar-coded retail items (recall we assumed that these changes were equal to those of bar-coded items in our baseline quantification).

C. *Quantification Results*

1. The Average Welfare Gains from Foreign Retail Entry

We first present the results of the quantification under the exact (CES) approach. As described above, the income and expenditure surveys allow us to calculate the welfare gains for each household in our sample based on the occupations and sectors they work in, the products they spend their income on, and the types of stores they shop at. Column 1 of table 9 presents the mean of the total welfare gain across all these urban households (using the household survey weights) as well as the maximum, the minimum, and the proportion negative. The various subcomponents of the total welfare effect are reported in columns 2–7. On average, we find that foreign store entry leads to large and significant welfare gains for households in the municipality where the foreign-owned retailer enters. These gains are on the order of 6 percent of initial household income.

Turning to the subcomponents, our positive total welfare effect is driven by a significant reduction in the cost of living—a 6.4 percent welfare gain—that far outweighs the effects on the nominal income side. While the adverse effects on the incomes of traditional retail workers and local store owners are economically significant (recall that their income losses were around 6 percent and employment losses were even larger), these effects are muted for the municipality as a whole since only a fraction of households derive substantial shares of their total income from these sources. In contrast, retail constitutes a large part of household total consumption for every household, generating substantial cost-of-living reductions.

[39] This is a parametric bootstrap (e.g., Horowitz 2001) that implicitly assumes errors are uncorrelated across data sets.

TABLE 9

Household Welfare Effect: Decomposition

	A. Exact under CES Approach: Dependent Variable						
	Total Effect (1)	Direct Price Index Effect (2)	Procompetitive Price Effect (3)	Procompetitive Exit Effect (4)	Retail Labor Income Effect (5)	Retail Profit Effect (6)	Other Income Effect (7)
Average effect	.0601***	.0551***	.0158***	−.00705	−.00397**	−.00269**	.00289
	(.0104)	(.0006)	(.0050)	(.0053)	(.0020)	(.0013)	(.0078)
Maximum	.730	.177	.055	.000	.692	.000	.020
Minimum	−.986	.000	.000	−.014	−1.000	−1.000	−1.000
Proportion negative	.024	.000	.000	.999	.074	.058	.004
Observations (households)	12,293	12,293	12,293	12,293	12,293	12,293	12,293
Number of municipality clusters	240	240	240	240	240	240	240
	B. First-Order Approach: Dependent Variable						
	Total Effect (8)	Direct Price Index Effect (9)	Procompetitive Price Effect (10)	Procompetitive Exit Effect (11)	Retail Labor Income Effect (12)	Retail Profit Effect (13)	Other Income Effect (14)
Average effect	.0275***	.0204***	.0109***	0	−.00397**	−.00269**	.00289
	(.0093)	(.0014)	(.0037)	(.0000)	(.0020)	(.0013)	(.0078)
Maximum	.715	.060	.031	.000	.692	.000	.020
Minimum	−.995	.000	.000	.000	−1.000	−1.000	−1.000
Proportion negative	.057	.000	.000	.000	.074	.058	.004
Observations (households)	12,293	12,293	12,293	12,293	12,293	12,293	12,293
Number of municipality clusters	240	240	240	240	240	240	240

Note.—The table reports the welfare effects of foreign retail entry from the quantification exercise described in Sec. VI. The average effect is the survey-weighted mean across all 12,293 households in the income and expenditure surveys that reside in the 240 urban municipalities that had not yet experienced foreign retail entry at the time of the survey. Panel A uses the exact and complete CES approach while panel B uses the Paasche approximation that is purely based on observed price differences. Standard errors come from the bootstrap procedure described in Sec. VI.

* Significant at the 10 percent level.

** Significant at the 5 percent level.

*** Significant at the 1 percent level.

Homing in further, about one-quarter of the cost-of-living effect, or 1.6 percent, comes from the procompetitive price effect, the price reductions at domestic stores induced by the entry of foreign retailers (an effect size approximately equal to the following back-of-the-envelope calculation: the share of retail in expenditure [0.5 on average] multiplied by the average domestic market share before and after foreign entry [$(1 + 0.7)/2 = 0.85$ on average] multiplied by the price reduction in domestic stores [0.04 on average]). The procompetitive exit effect, the welfare losses due to shop closures, is small at -0.7 percent since the number of store closures was limited (at least over our 5-year window). The remainder of the cost-of-living effect, 5.5 percent, is due to the direct price index effect of foreign entry, the consumer gains from being able to shop at the foreign store itself (a finding that was foreshadowed in the raw data by the 12 percent lower prices charged by foreign stores and their 30 percent post-entry market shares).

On the income side, there is an average loss of 0.40 percent from declines in retail labor income, a loss of 0.27 percent from declines in retail profits, and a gain of 0.29 percent from increases in other incomes (with the intensive margin alone—i.e., wage and profit changes, not employment changes or store exit—accounting for losses of 0.19 and 0.14 and gains of 0.49 percent, respectively).

We report the results of the first-order approach in columns 8–14 of table 9. The total effect is smaller under this approach, with the average gains equal to 2.8 percent of initial household income. This is driven by a smaller direct price index effect of 2.0 percent (approximately the share of retail in expenditure [0.5 on average] multiplied by the foreign post-entry market share [0.3 on average] multiplied by the price reduction in domestic stores plus the post-entry price gap between foreign and domestic stores [0.16 on average]).

Recall that the first-order approach uses a Paasche approximation of the direct price index effect—post-entry expenditure shares multiplied by observable price changes—that is analogous to estimating the gains if foreign stores always existed but simply charged domestic prices in the pre-entry periods. Hence, it does not capture three potentially important welfare gains due to foreign store entry: new product variety on sale at foreign stores, different shopping amenities available at foreign stores, and the additional store variety that comes from having an extra shopping choice. In contrast, we prefer the exact (CES) approach above precisely because these three channels are captured in the CES direct price index effect alongside lower prices. Hence, in an approximate sense, the fact that the direct price index effect is 40 percent as large under the first-order approach (2.0 percent as opposed to 5.5 percent) suggests that these variety and amenity gains account for a substantial proportion of the cost-of-living effect.[40] Given the large differences between foreign and domestic stores that we highlighted in Section II, the sizable gains generated by these variety and amenity channels seem plausible (a discussion we will return to when assessing the distribution of the welfare gains).

Finally, we explore the sensitivity of our average total effect to alternative values of two parameters that are key drivers of the quantification results: the elasticity of substitution across local stores and the price changes for non-bar-coded products. Our baseline quantification above used the most conservative estimate of the elasticity of substitution, a specification that yielded point estimates around 4 for the product- and income group–specific elasticities (col. 23 of table 8). We take these η_{gh} estimates as the midpoints and reestimate the full quantification exercise using eight alternative sets of elasticities, either subtracting from or adding 0.5, 1, 1.5, and 2 to the midpoints. On the price side, our baseline quantification imposed the assumption that price changes were the same across bar-coded and non-bar-coded products within a product group and store type and then used the bar-coded estimates from Section V.A, where we could control for product characteristics.[41] Again, we take this assumption as our midpoint and reestimate the quantification assuming that there were no price changes for non-bar-coded items; the price changes were 50 percent as large, 150 percent as large, or 200 percent as large. Table 10 reports the total welfare effect for the 9×5 different combinations of elasticities and non-bar-coded price changes. Reassuringly, despite the wide parameter ranges, the total effects remain reasonable, varying between 3.3 and 14.7 percent.

Our multitiered preference structure imposes that the elasticity of substitution is the same across traditional and modern stores. We may expect the elasticity of substitution within modern store types to differ from that across modern and traditional stores. To allow for such heterogeneity while maintaining the CES middle tier, we rerun the quantification but placing traditional and modern stores in different CES nests within a product group—implicitly imposing a low elasticity of 1 across traditional and modern stores due to the Cobb-Douglas upper tier. As shown in online appendix table A.7, the average total effect falls in this scenario since we shut down the direct gains from foreign stores capturing market

[40] As noted in n. 16, the direct price index is biased upward using a Paasche approximation. However, this bias is small. A Sato-Vartia price index using our CES framework and the η_{gh} estimates from Sec. VI.A to estimate the initial market shares if foreign stores had charged domestic prices yields a direct price index effect of 1.9 percent rather than 2.0 percent. Conversely, the procompetitive price effect is biased downward, rising from 1.09 percent under the Paasche first-order approach to 1.15 percent when using the Sato-Vartia approach.

[41] This baseline assumption is consistent with that of Hausman and Leibtag (2007), who find similar price changes in branded and unbranded products upon the entry of big-box stores in the United States.

TABLE 10

Effect on Average Household Welfare: Robustness

$\boldsymbol{\eta} = \begin{bmatrix} \eta_{\text{poor,food}} & \eta_{\text{poor,nonfood}} \\ \eta_{\text{rich,food}} & \eta_{\text{rich,nonfood}} \end{bmatrix} = \begin{bmatrix} 3.87 & 3.85 \\ 4.27 & 4.34 \end{bmatrix}$	$\boldsymbol{\eta} - 2$	$\boldsymbol{\eta} - 1.5$	$\boldsymbol{\eta} - 1$	$\boldsymbol{\eta} - .5$	$\boldsymbol{\eta}$	$\boldsymbol{\eta} + .5$	$\boldsymbol{\eta} + 1$	$\boldsymbol{\eta} + 1.5$	$\boldsymbol{\eta} + 2$
0% of procompetitive effect for non-bar-coded products	.1355	.0964	.0716	.0626	.0520	.0468	.0394	.0374	.0329
50% of procompetitive effect for non-bar-coded products	.1392	.0986	.0785	.0649	.0554	.0485	.0453	.0389	.0388
100% of procompetitive effect for non-bar-coded products	.1423	.1004	.0821	.0667	.0601[a]	.0509	.0470	.0434	.0400
150% of procompetitive effect for non-bar-coded products	.1454	.1063	.0841	.0732	.0640	.0573	.0534	.0492	.0453
200% of procompetitive effect for non-bar-coded products	.1468	.1075	.0889	.0728	.0653	.0605	.0558	.0523	.0490

Note.—The table reports the welfare effects of foreign retail entry from the quantification exercise described in Sec. VI but using alternative estimates and assumptions. Each cell is the survey-weighted mean effect across all 12,293 households in the income and expenditure surveys that reside in the 240 urban municipalities that had not yet experienced foreign retail entry at the time of the survey. The estimate with the superscript "a" corresponds to the average effect in table 9, which applies our baseline assumption that the procompetitive effects on prices in domestic stores that we estimated using bar-coded items are identical for non-bar-coded items and that the elasticities are those from col. 23 of table 8. Other columns vary the value of these elasticities by adding or subtracting to the set of four elasticities in col. 23 of table 8 in increments of 0.5. Other rows vary the relative strength of the procompetitive effects on non-bar-coded retail items.

share from traditional stores. However, the magnitudes remain similar with average gains of 5.2 percent if we use the same elasticity of substitution within the traditional nests that we estimated in Section VI.A using data on consumer substitution across modern stores (rising to 5.4 percent if we assume the elasticity of substitution is instead 6 within traditional stores so that traditional store exit is less deleterious, and falling to 4.4 percent if we assume it is 2 instead).

2. The Distribution of the Gains from Foreign Retail Entry

In the previous section we reported average household effects. Since we have a separate estimate for each sample household in the income and expenditure surveys, it is straightforward to analyze the distribution of these gains. The upper panel of figure 6 plots the total welfare effect for each household against the initial position of the household in the income distribution using a nonparametric local polynomial regression. The lower panel decomposes these gains. For both, we focus on our preferred exact (CES) approach (online app. fig. A.3 plots the distribution using the first-order approach). While all income groups benefit substantially from foreign retail entry, richer households gain substantially more than poorer households (about 7.5 percent compared to 5 percent).

Where does this regressiveness come from? We present several counterfactual exercises that allow us to analyze the interplay of forces underlying this result. We focus on the role of several key differences in shopping and income patterns across the income distribution: the share of retail expenditure spent at foreign stores after entry, expenditure shares of retail relative to nonretail, expenditure shares of food relative to nonfood product groups within retail, and retail income shares relative to other income sources. Each row of figure 7 equalizes one of these differences across all households in the sample, setting it at its mean level (the left panel shows the distribution of the dimension before equalization). We then rerun the quantification and generate the counterfactual distribution of the welfare gains from foreign entry and compare it to the actual distribution (with the distribution of both the actual and counterfactual gains shown in the right panel).

As shown in the first row of figure 7, the richest households spend over 50 percent of their retail expenditure at foreign stores compared to just over 10 percent for the poorest households. These patterns suggest that household evaluations of store product variety and shopping amenities systematically differ across the income distribution (captured by the household- and store-specific taste shifters β_{gsh} and price indices r_{gsh} in Sec. III.A). As is evident from the figure, equalizing this moment alone is sufficient to eradicate the regressiveness. In fact, if poor households

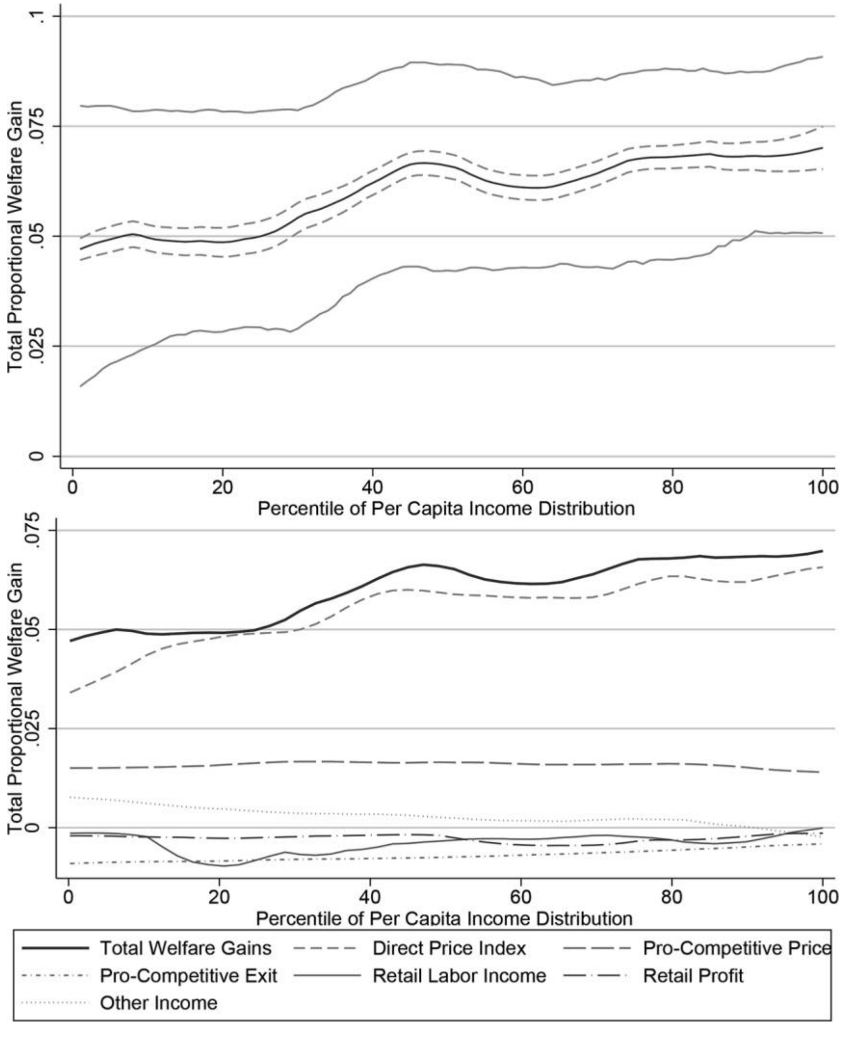

FIG. 6.—Gains from foreign retail entry across the household income distribution. The graphs are nonparametric plots of the household gains from foreign retail entry against the pre-entry location in the income distribution. Gains are calculated from the quantification exercise described in Section VI using the exact under CES approach. Pre-entry incomes as well as household-level income and expenditure shares come from the 12,293 households in the income and expenditure surveys that reside in the 240 urban municipalities that had not yet experienced foreign retail entry at the time of the survey. The upper panel depicts two sets of confidence intervals: The solid gray lines are the 95 percentile envelope of the nonparametric plots for each of the 1,000 bootstraps described in Section VI. The tighter dashed lines are the 95 percentile confidence interval that just takes account of sampling variation across households in the income and expenditure survey. The lower panel decomposes the total gains into its constituent parts. Plots are weighted by household survey weights.

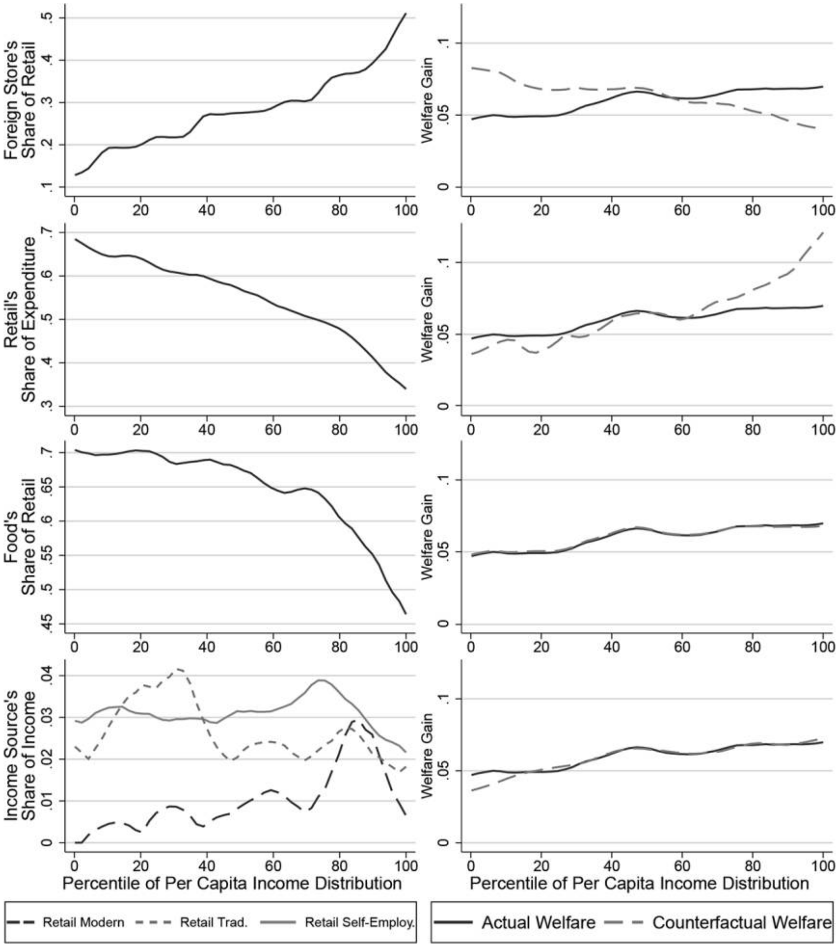

FIG. 7.—Counterfactual distributions of the gains from foreign retail entry. This figure explores the role of differences along several shopping and income dimensions in explaining the regressive total welfare effects we find using the exact under CES approach. Each row explores a different dimension (foreign retail shares, retail's share of total expenditure, the food share of retail, and retail income shares). The left panels plot the distribution of this dimension in the data. The right panels plot with a dashed line the counterfactual distribution of gains if this dimension was equalized across households at its mean (alongside the actual distribution displayed with a solid line). Pre-entry incomes as well as household-level income and expenditure shares come from the 12,293 households in the income and expenditure surveys that reside in the 240 urban municipalities that had not yet experienced foreign retail entry at the time of the survey. Plots are weighted by household survey weights.

valued the variety and amenity on offer at foreign stores as much as the rich did, the gains would actually be progressive. This finding is intuitive. The higher quality and greater product variety on offer at foreign stores, as well as their better hygiene, easier car accessibility, and parking amenities, are all benefits that are likely to be valued more by wealthier households. The fact that it is these households gaining most from variety and amenity gives credence to the large contribution these forces play in generating the substantial average gains we document above.

Turning to the second row of figure 7, there is a countervailing force at play. Wealthier households spend a significantly smaller share of their total expenditure on retail consumption compared to poorer households (35 percent vs. 70 percent). This force works in the opposite direction. In the absence of differences in retail expenditure shares, the gains from foreign retail entry would be much more regressive than we estimate.

The forces in the last two rows have much more muted effects on the distribution. The fact that poorer households spend more of their retail expenditure on food consumption contributes only very slightly to the regressiveness of the welfare gains because both the procompetitive price effect and the direct price index effect vary little across food and nonfood product groups. Finally, somewhat surprisingly, differences in income sources across the income distribution do not significantly contribute to the regressiveness we find. While there are clear distributional patterns in the sectors in which households obtain their income (e.g., poorer households derive a larger proportion of their income from working in the traditional retail sector), these differences have little effect on the distribution of total gains since only a fraction of households within any given income group derive the majority of their income from the retail sector.

D. *Discussion*

Before concluding the paper, we discuss several issues related to our findings above. First, we explore to what extent the results are specific to the entry of foreign retailers rather than the entry of modern store formats more generally. Second, we explore whether what we have labeled the procompetitive price effects are driven by reductions in markups or by reductions in marginal costs due to local spillovers from foreign entry. Finally, we relate our findings to those of the existing literature.

1. Foreign Entry or Modern Store Formats

To assess to what extent our findings are specific to foreign retail entry, we present two additional pieces of evidence. The first is to reestimate the two key moments that are driving our estimated welfare effects in table 9—price changes in preexisting stores and post-entry store market shares—using entry events involving domestic stores with store formats similar to those of the foreign entrants. The second is to verify to what extent these two effects are present among municipalities that already had domestic modern store formats at the start of the estimation period.

For the years 2002–6, we have data on store opening dates of all modern domestic stores.[42] We define comparable domestic entry events as store openings of Mexican national retail chains whose average floor space—observed in our ANTAD data set—is similar to that of foreign stores. Only four domestic chains—Soriana, Chedauri, Comercial Mexicana, and Gigante—both have a national presence and use the big-box store format used by foreign entrants.[43]

Figure 8 repeats the consumer price event study of Section V.A for the period 2002–6 and for both foreign store entry events (panel A) and domestic big-box store entry events (panel B). Reassuringly, the figure for foreign-retail entry is almost identical to the results in figure 2 for the full 2002–14 period. In contrast, when looking at comparable domestic entry events, we find no procompetitive effect on consumer prices in preexisting local stores.

Panels C and D of figure 8 carry out a similar exercise but on the post-entry retail market shares analyzed in Section V.B.1. We find that comparable domestic retailers, despite opening similar-sized stores, command much smaller local market shares after entry.[44] Conditional on entry, the average market share of the comparable domestic retailers is less than one-third of that of the foreign retailers (8.5 percent compared to 30 percent).

Finally, we explore heterogeneity in the effects of foreign entry with respect to differences in the preexisting level of local competition. In particular, if our results were driven by the entry of modern big-box store formats rather than specific to foreign entry, we would expect to find much weaker effects in municipalities that were already served by domestic modern big-box formats. To explore this hypothesis, we allow the procompetitive price effect of foreign entry and the post-entry foreign consumption shares to differ depending on the level of preexisting local competition (as measured by the number of domestic big-box stores per capita in 2002). As reported in online appendix tables A.8 and A.9, we find close to zero (and insignificant) heterogeneity in the procompetitive price effect. Post-entry foreign market shares are slightly higher (by 8.8 percent) in municipalities with low (i.e., below-mean) levels of preexisting local competition. However, this heterogeneity cannot explain the more than 20 percentage point difference in foreign and

[42] Recall that after 2006, ANTAD stopped reporting the date of store openings by municipality. We obtained subsequent opening dates only for foreign-owned supermarkets.

[43] These are the four largest domestic retail chains in terms of average floor space and are all within 25 percent of the average floor space of foreign supermarkets. In contrast, the average floor space of the fifth-largest is more than 85 percent smaller than the average floor space of foreign supermarkets.

[44] As before, we estimate average household retail expenditure shares using the consumer panel. We restrict attention to municipalities containing a foreign entrant (panel C) or a comparable modern domestic entrant (panel D).

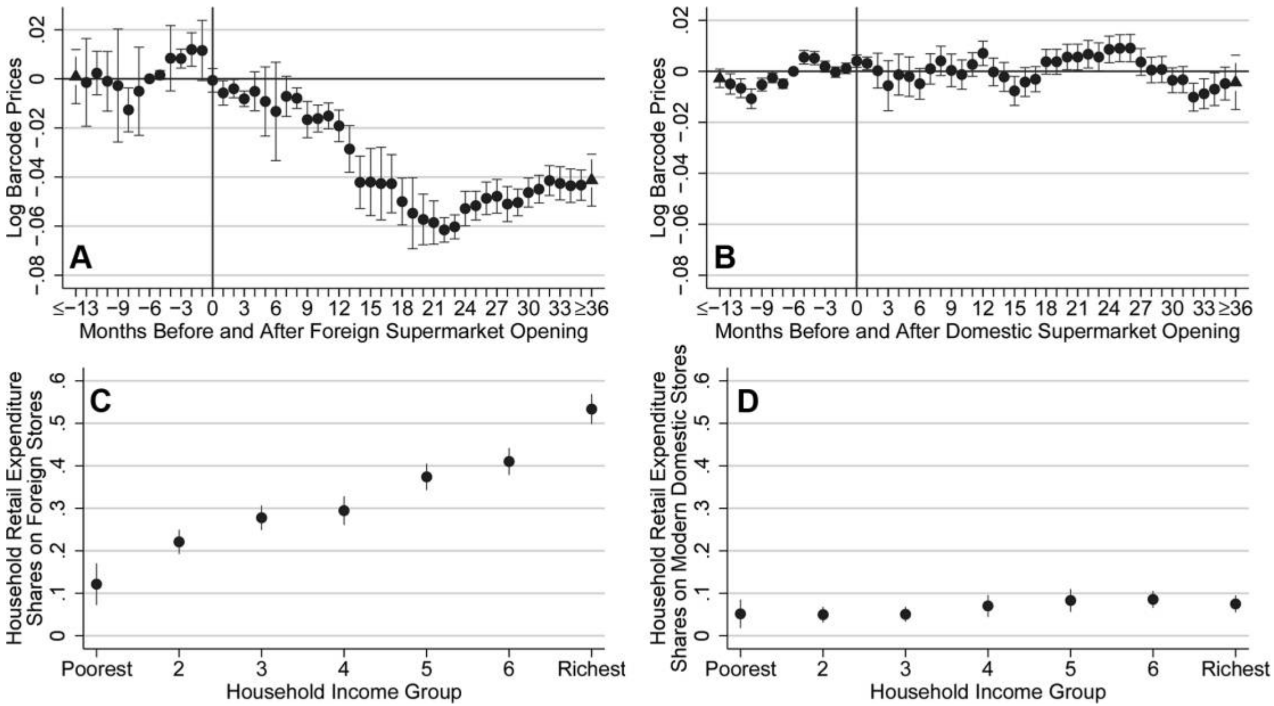

Fig. 8.—The procompetitive price effect and post-entry consumption shares of foreign versus domestic modern store entry. Panel A depicts the event study estimates for foreign entry over the period 2002–6 using the monthly CPI price microdata. Panel B depicts equivalent estimates for the entry of comparable domestic stores (i.e., domestic stores with similar average floor space). Panels C and D depict the average post-entry retail consumption shares of foreign and comparable domestic supermarkets, respectively, in the years 2011–14, using the consumer panel data. The comparable domestic retail chains in terms of average floor space are Soriana, Chedauri, Comercial Mexicana, and Gigante. See Section VI.D for further details.

domestic post-entry market shares we document in figure 8: even in the extreme case in which foreign stores chose uncompetitive locations and domestic big-box stores chose competitive ones, our heterogeneity estimates could explain only a little over a third of this gap.

Taken together, these additional exercises provide evidence that our findings are specific to the entry of foreign retailers rather than modern store formats more generally. As we discuss in Section II.B, the literature highlights multiple differences between global retailers and their domestic competitors in developing countries that underlie these results including the use of global supply chains, logistics centered around regional distribution centers, and modern information technology systems for real-time store monitoring.

2. Markups or Spillovers

An important part of our total welfare effect is the term we label the procompetitive price effect. These consumer price reductions in domestic stores can stem from two very different sources: markup reductions or reductions in marginal costs driven by spillovers from foreign entry. Examples of such spillovers include domestic stores adopting the better management practices and logistics used by foreign retailers and reductions in input prices at local suppliers due to scale effects or productivity spillovers from foreign retailers. As we note in Section III.A, our welfare estimation does not require this distinction since we separately capture the changes in both the local cost of living and local incomes (which include income from store profits along with other sources). Having said that, this distinction is an important and interesting one for thinking more broadly about the implications of retail FDI.

To assess whether the 2.4 percent price reduction in traditional stores reported in table 3 can be fully accounted for by markup reductions, we perform a back-of-the-envelope calculation using the retail census microdata. Table 6 showed that average store profits fell by 4.4 percent among traditional stores. We complement this finding by replacing the dependent variable in the exit and profits specification shown in equation (17) with either total revenues R or total costs C and find point estimates of -2 percent and -1.6 percent, respectively (see online app. tables A.10 and A.11). These results imply that the observed price reduction is driven by a 4.8 percent reduction in domestic store markups μ and a 2 percent reduction in marginal costs c, where price $p = \mu + c$.[45] Hence, these findings suggest a role for both markup reductions and marginal cost reductions through local spillovers from foreign entry.

[45] These numbers follow from the fact that $d \ln \mu = d \ln p - (d \ln R - d \ln \pi)$ and $d \ln c = d \ln p - (d \ln R - d \ln C)$. Note that they imply a realistic share of profits over revenues of 0.14.

This analysis is subject to one important caveat: these estimates provide a municipality-level perspective on spillovers from retail FDI as we do not capture potential spillovers on domestic stores or suppliers at the wider national level. In this sense, our estimates complement those of Javorcik, Keller, and Tybout (2008) and Iacovone et al. (2015), who find productivity gains among domestic suppliers in Mexico driven by foreign retail entry. In our empirical strategy, these national-level effects are absorbed by the time fixed effects. Such national-level effects should be added, along with any nominal income losses to Mexican shareholders of national retail chains, in order to calculate aggregate rather than local welfare effects.

3. Relationship to the Existing Literature

Given that Walmart is the most prominent of the foreign entrants in Mexico, we first relate the magnitude of our effects to studies from the United States that estimate the impacts of Walmart entry on some subset of employment, wages, prices, and store exit. In terms of employment and wages, this literature also finds negative impacts of Walmart's entry, but the magnitudes are substantially smaller than our estimates for Mexico in table 6.[46] In terms of price effects, Basker (2005b) documents citywide price reductions of 1.5–3 percent due to Walmart entry, compared to the 6.6 percent reduction we find when we take the weighted average of the price changes in domestic stores and the lower prices in foreign stores. In terms of store exit, Ellickson and Grieco (2013) show that 0.15 stores shut down for every Walmart entry, substantially smaller than our finding that 3.9 percent of retail units exited as a result of the foreign entry event (corresponding to 80 traditional stores closing and one modern retailer closing in the average municipality). Finally, Hausman and Leibtag (2007) calculate the welfare gains from the arrival of Walmart Supercenters. Despite ignoring the losses from store exit and nominal incomes that we include, the welfare gains they estimate—a direct price index effect of 2.4 percent and a procompetitive price effect of 0.06 percent—are more than 50 percent smaller than those we find.

The larger magnitudes we find are consistent with our claim in Section II.B that foreign retailers entering developing countries—whose retail landscapes, like Mexico's, are dominated by traditional retailers (Capizzani, Ramirez Huerta, and Rocha e Oliveira 2012; Bronnenberg and Ellickson 2015)—constitute a more dramatic shock than expan-

[46] Dube, Lester, and Eidlin (2007) find that county wage bills fall 1.4 percent; Neumark, Zhang, and Cicarella (2008) find that county retail employment falls 2.7 percent; and Ellickson and Grieco (2013) find 7 percent employment reductions for retailers within a 2-mile radius. In contrast, Basker (2005a) finds positive but small effects. These compare to a larger 7.8 percent reduction in local retail employment that we find for Mexico.

sions of companies such as Walmart within developed countries. It is also interesting to note that, for example, Broda, Leibtag, and Weinstein (2009) find that the price index implications of Walmart in the United States are pro-poor whereas we find the opposite distributional pattern for Mexico, consistent with foreign retailers' targeting upper-middle-class customers in the developing world.

Second, we turn to the question of how our results relate to an off-the-shelf quantification approach using the framework of Arkolakis et al. (2012) and the extension to multinational production by Ramondo and Rodríguez-Clare (2013). To calculate Ramondo and Rodríguez-Clare's expression for the gains from horizontal FDI (see their eq. [17]), we require three empirical moments: the share of retail in Mexican GDP (12 percent on average in Mexico's national accounts 2002–14), the foreign production share in Mexican retail, and the "trade elasticity" that governs how foreign production shares react to changes in FDI frictions. Assuming that foreign production shares—that is, foreign firms' share of retail value added—equal the post-entry market shares we observe in the consumer panel data (30 percent on average) and that the trade elasticity used in Ramondo and Rodríguez-Clare's study, $\theta = 4.45$, holds for Mexican retail FDI, the estimated gains from foreign retail entry are

$$\prod_{g\in G}\left[\left(\sum_{s\in S_g^{dc}}\phi_{gsh}^{1}\right)^{-1/\theta}\right]^{\alpha_{gh}} - 1 = [(0.7)^{-1/4.45}]^{0.12} - 1 = 0.01,$$

or about one-sixth of the gains we estimate.

There are several reasons why our numbers are higher. First, the Ramondo and Rodríguez-Clare approach assumes constant markups and productivity parameters and thus abstracts from lower prices in domestic stores due to procompetitive effects or productivity gains (which we estimate to yield a local gain of 1.6 percent). Second, while the formulas appear similar, our estimate of the direct price index effect uses a demand elasticity η rather than a trade elasticity θ, and our direct effect is scaled to account for the fact that 50 percent of expenditure is spent on retail rather than the fact that retail produces 12 percent of value added. Herrendorf, Rogerson, and Valentinyi (2013) explore a closely related issue: why do the effects of structural change differ when estimates come from the consumption side (as ours do) and the production side (as Ramondo and Rodríguez-Clare's do)? They conclude that it is crucial to use consumption-side elasticities when scaling by expenditure shares and production-side elasticities when scaling by value-added shares. On the consumption side, this is the elasticity of substitution across foreign and domestic stores we estimate. On the production side, this is the response of foreign production shares to changes in frictions restricting retail FDI. The elasticity used in Ramondo and Rodríguez-Clare's study is estimated off goods trade flows within a set of OECD countries (excluding Mexico) rather than production share changes in Mexican retail induced by changing FDI frictions. Our conjecture would be that the latter elasticity would be lower than $\theta = 4.45$, reducing the size of the discrepancy.

VII. Conclusion

The arrival of foreign retailers in developing countries is causing a radical transformation in the way in which households source their consumption. This paper sets out to evaluate the welfare consequences of retail globalization in a developing country context. To do so, we bring to bear newly available and uniquely rich microdata that allow us to estimate a general expression of the local welfare effect of retail FDI.

We find that foreign supermarket entry leads to large and significant welfare gains for the average household. The majority of these gains come from a significant reduction in the cost of living. About a quarter of this reduction is driven by reductions in prices at domestic stores, with the remainder coming from the consumer gains due to the lower prices, new product variety, and different shopping amenities offered by foreign retailers. In contrast, on the nominal income side we find no evidence of average income or employment effects. We do, however, find evidence of domestic store exit as well as employment, labor income, and profit declines in the traditional retail sector. Exploring the distribution of these gains from retail FDI, we find that while all income groups experience significant gains on average, these gains are 50 percent larger for the richest income group compared to the poorest, primarily because of the greater valuation wealthy households place on the product variety and shopping amenities on offer at foreign stores.

Our analysis provides a number of insights that relate to ongoing debates about developing country policies toward retail FDI. Our findings suggest that these debates may focus too little on the potential for reductions in the cost of living that benefit the vast majority of households, both those who end up shopping at the foreign retailer and those who enjoy price reductions at domestic retailers. Instead they commonly focus on the potentially adverse effects for an important but nevertheless select group of households working in the traditional retail sector. The empirical evidence suggests that while these adverse nominal income effects are present, they are swamped at the local level by reductions in the cost of living that give rise to real income gains across all household income groups.

References

Anderson, S. P., A. De Palma, and J. F. Thisse. 1992. *Discrete Choice Theory of Product Differentiation.* Cambridge, MA: MIT Press.

Arkolakis, C., A. Costinot, and A. Rodríguez-Clare. 2012. "New Trade Models, Same Old Gains?" *A.E.R.* 102 (1): 94–130.

Atkin, D. 2013. "Trade, Tastes and Nutrition in India." *A.E.R.* 103 (5): 1629–63.

Basker, E. 2005a. "Job Creation or Destruction? Labor Market Effects of Wal-Mart Expansion." *Rev. Econ. and Statis.* 87 (1): 174–83.

———. 2005b. "Selling a Cheaper Mousetrap: Wal-Mart's Effect on Retail Prices." *J. Urban Econ.* 58:203–29.

Beraja, M., E. Hurst, and J. Ospina. 2014. "The Regional Evolution of Prices and Wages during the Great Recession." Manuscript, Univ. Chicago.

Biles, J. J. 2008. "Wal-Mart and the Supermarket Revolution in Mexico." *Geographische Rundschau Internat. Ed.* 58 (2): 44–49.

Broda, C., E. Leibtag, and D. E. Weinstein. 2009. "The Role of Prices in Measuring the Poor's Living Standards." *J. Econ. Perspectives* 23 (2): 77–97.

Broda, C., and D. E. Weinstein. 2006. "Globalization and the Gains from Variety." *Q.J.E.* 121 (2): 541–85.

Bronnenberg, B. J., and P. B. Ellickson. 2015. "Adolescence and the Path to Maturity in Global Retail." *J. Econ. Perspectives* 29 (4): 113–34.

Capizzani, M., F. Ramirez Huerta, and P. Rocha e Oliveira. 2012. "Retail in Latin America: Trends, Challenges and Opportunities." Working Paper no. 170, IESE, Univ. Navarra.

Caron, J., T. Fally, and J. R. Markusen. 2014. "International Trade Puzzles: A Solution Linking Production and Preferences." *Q.J.E.* 129 (August): 1501–52.

Deaton, A. 1989. "Rice Prices and Income Distribution in Thailand: A Nonparametric Analysis." *Econ. J.* 99 (395): 1–37.

Deaton, A., and A. Tarozzi. 2000. *Prices and Poverty in India.* Princeton, NJ: Princeton Univ. Press.

DellaVigna, S., and M. Gentzkow. 2017. "Uniform Pricing in US Retail Chains." Working Paper no. 23996 (November), NBER, Cambridge, MA.

Deloitte Consulting. 2014. "Global Powers of Retailing 2014." Report, Deloitte, London.

Donaldson, D. Forthcoming. "Railroads of the Raj: Estimating the Impact of Transportation Infrastructure." *A.E.R.*

Dube, A., T. W. Lester, and B. Eidlin. 2007. "Firm Entry and Wages: Impact of Wal-Mart Growth on Earnings throughout the Retail Sector." Working Paper no. 126-05, Inst. Res. Labor and Employment, Univ. California Berkeley.

Dufey, A., M. Grieg-Gran, and H. Ward. 2008. *Responsible Enterprise, Foreign Direct Investment and Investment Promotion: Key Issues in Attracting Investment for Sustainable Development.* London: Internat. Inst. Environment and Development.

Einav, L., J. Levin, and P. Klenow. 2015. "The Importance of Customer Relationships for U.S. Retailer Size and Growth." Paper presented as the 2015 Allied Social Sciences Assoc. meetings, Boston.

Ellickson, P. B., and P. L. Grieco. 2013. "Wal-Mart and the Geography of Grocery Retailing." *J. Urban Econ.* 75:1–14.

Faber, B. 2014. "Trade Liberalization, the Price of Quality, and Inequality: Evidence from Mexican Store Prices." Working paper, Univ. California Berkeley.

Fajgelbaum, P. D., and A. K. Khandelwal. 2014. "Measuring the Unequal Gains from Trade." Working paper, Univ. California Los Angeles.

Feenstra, R. C. 1994. "New Product Varieties and the Measurement of International Prices." *A.E.R.* 84 (1): 155–57.

Feenstra, R. C., and D. E. Weinstein. 2017. "Globalization, Markups, and US Welfare." *J.P.E.* 125 (4): 1040–74.

Goldberg, P., and N. Pavcnik. 2007. “Distributional Effects of Globalization in Developing Countries.” *J. Econ. Literature* 45:39–82.

Handbury, J., and D. E. Weinstein. Forthcoming. “Goods Prices and Availability in Cities.” *Rev. Econ. Studies*.

Hausman, J. A. 1997. “Valuation of New Goods under Perfect and Imperfect Competition.” In *The Economics of New Goods*, edited by T. F. Bresnahan and R. J. Gordon, 207–48. Chicago: Univ. Chicago Press (for NBER).

Hausman, J. A., and E. Leibtag. 2007. “Consumer Benefits from Increased Competition in Shopping Outlets: Measuring the Effect of Wal-Mart.” *J. Appl. Econometrics* 22 (7): 1157–77.

Hausman, J. A., and G. K. Leonard. 2002. “The Competitive Effects of a New Product Introduction: A Case Study.” *J. Indus. Econ.* 50 (3): 237–63.

Herrendorf, B., R. Rogerson, and A. Valentinyi. 2013. “Two Perspectives on Preferences and Structural Transformation.” *A.E.R.* 103 (7): 2752–89.

Hicks, J. R. 1940. “The Valuation of the Social Income.” *Economica* 7:105–24.

Holmes, T. J. 2011. “The Diffusion of Wal-Mart and Economies of Density.” *Econometrica* 79 (1): 253–302.

Horowitz, J. L. 2001. “The Bootstrap.” In *Handbook of Econometrics*, vol. 5, edited by J. J. Heckman and E. Leamer, 3159–3228. Amsterdam: Elsevier.

Humphrey, J. 2007. “The Supermarket Revolution in Developing Countries: Tidal Wave or Tough Competitive Struggle?” *J. Econ. Geography* 7 (4): 433–50.

Iacovone, L., B. S. Javorcik, W. Keller, and J. R. Tybout. 2015. “Supplier Responses to Wal-Mart’s Invasion in Mexico.” *J. Internat. Econ.* 95 (1): 1–15.

Javorcik, B. S., W. Keller, and J. R. Tybout. 2008. “Openness and Industrial Response in a Wal-Mart World: A Case Study of Mexican Soaps, Detergents and Surfactant Producers.” *World Econ.* 31 (12): 1558–80.

Javorcik, B. S., and Y. Li. 2013. “Do the Biggest Aisles Serve a Brighter Future? Global Retail Chains and Their Implications for Romania.” *J. Internat. Econ.* 90 (2): 348–63.

Jia, P. 2008. “What Happens When Wal-Mart Comes to Town: An Empirical Analysis of the Discount Retailing Industry.” *Econometrica* 76 (6): 1263–1316.

Krugman, P. 1980. “Scale Economies, Product Differentiation, and the Pattern of Trade.” *A.E.R.* 70 (5): 950–59.

Lagakos, D. 2016. “Explaining Cross-Country Productivity Differences in Retail Trade.” *J.P.E.* 124 (2): 579–620.

Neumark, D., J. Zhang, and S. Cicarella. 2008. “The Effects of Wal-Mart on Local Labor Markets.” *J. Urban Econ.* 63:405–30.

Porto, G. 2006. “Using Survey Data to Assess the Distributional Effects of Trade Policy.” *J. Internat. Econ.* 70 (1): 140–60.

Ramondo, N., and A. Rodríguez-Clare. 2013. “Trade, Multinational Production, and the Gains from Openness.” *J.P.E.* 121 (2): 273–322.

Reardon, T., C. P. Timmer, C. B. Barrett, and J. Berdegué. 2003. “The Rise of Supermarkets in Africa, Asia, and Latin America.” *American J. Agricultural Econ.* 85 (5): 1140–46.

Topalova, P. 2010. “Factor Immobility and Regional Impacts of Trade Liberalization: Evidence on Poverty from India.” *American Econ. J.: Appl. Econ.* 2 (4): 1–41.

UNCTAD (UN Commission on Trade and Development). 2013. *World Investment Report: Global Value Chains: Investment and Trade for Development.* New York and Geneva: United Nations.

Varela, M. 2016. “The Costs of Growth: Accelerated Growth and Crowd-Out in the Mexican Supermarket Industry.” Working paper, Univ. Arizona.

Wrigley, N., and M. Lowe. 2010. “The Globalization of Trade in Retail Services.” OECD, Paris.