Prices and Immigration: Firm-Level Evidence*

Ryan Kim† Seoul National University

Justin H. Leung‡ Chinese University of Hong Kong

Ariel Weinberger§ George Washington University

This version: September 11, 2026

Abstract

Does immigration raise or lower consumer prices? The standard intuition points both ways: immigrants add labor supply, which should lower costs and prices where they work, and product demand, which should raise prices where they shop. Linking barcode-level scanner prices to the establishments that produce each good, we measure a firm's exposure to immigration separately in the counties where it sells and where it produces. Within the firm, prices fall with immigration in a firm's sales markets and rise modestly where it produces, the opposite of the standard intuition, with the demand effect three times larger. Since immigrant demand expands the market, the decline must reflect markup compression. Immigration shifts a firm's customer base toward more price-sensitive consumers, who search more, pay less for identical goods, and reallocate spending toward the firm's cheaper products. A model of heterogeneous price sensitivity rationalizes these patterns. At the county level, an inflow of 10,000 immigrants lowers four-year price growth by 0.58 percentage points.

JEL Codes: J61; E31; F22; L11.

Keywords: Immigration; Consumer Prices; Search; Demand Elasticity.

*We thank participants at Barcelona Summer Forum, CUHK, ESWC, INSEAD, Johns Hopkins SAIS, KATIS, the KEA Joint Conference, the Mid-Atlantic International Trade Workshop, Peking University, the Rocky Mountain Empirical Trade Conference, SEA, Seoul National University, Washington Area International Trade Symposium, the Workshop on Household Heterogeneity (NCSU), and Yonsei University for helpful comments and suggestions. A special thanks to Erhan Artuc, Jay Hyun, and Poorya Kabir for discussing the paper. Researcher(s)' own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researcher(s) and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

†Seoul National University, rkim59@snu.ac.kr

‡Chinese University of Hong Kong, jhleung@cuhk.edu.hk

§George Washington University, aweinberger@gwu.edu

1 Introduction

Immigration and inflation stand as two paramount concerns for policymakers and the public alike, frequently dominating political discourse and economic analyses, particularly in recent years.[1] Considering immigration, public sentiment has recently leaned towards a desire for reduced immigration, often driven by concerns about its perceived effects on native wages. Yet, to fully grasp its economic impact, particularly on consumers, it is crucial to examine how immigration affects the prices they pay. Concurrently, inflation has surged as a major concern—particularly for everyday consumer goods—with popular debates pointing to both labor shortages and increased consumer demand as culprits. While the economic connection between immigration and inflation is frequently debated in media[2], their interplay remains an underexplored area in academic research.

Against this backdrop of pressing policy debates and a sparsely studied academic landscape, this paper asks two fundamental questions: Does immigration affect consumer prices? If so, why?

The effect of immigration inflow on price is ambiguous, mainly arising from two complexities. First, immigration increases both labor supply and product demand simultaneously. Intuitions often raised in media highlight a dual effect: increased labor supply could lead to lower labor costs and thus cooling prices, while increased product demand may add to upward price pressures. Second, immigrants possess different characteristics than natives and may affect prices differently than a simple increase in total population. For instance, immigrant populations, often new to the market and requiring more information or possessing lower incomes relative to natives, tend to exhibit increased search effort and, consequently, higher price sensitivity compared to natives. This can compel firms to *lower* their prices even with the overall increased product demand from immigrants. Beyond these specific demand-side adjustments, immigration inflow can also alter the composition of demand, raise wages, spur the entry of new establishments, or lead to production-side adaptations like firm scaling, innovation, or the introduction of new product varieties, depending on the type of immigrants. A comprehensive understanding of firm pricing conditional on immigration inflows thus requires carefully considering these intricate and often countervailing mechanisms.

To investigate the causal effect of immigration and its underlying mechanisms, we construct a detailed micro-level dataset by combining product-destination-level scanner prices with producer production sites, alongside county-specific historical and current immigration inflows and their origins. This dataset is built by matching the NielsenIQ Retail Scanner and Consumer Panel database with the GS1 database and the NETs establishment database in 2006-2018; these combined datasets are further matched to historical and current immigration data from the Census and the American Community Survey based on county information.[3] The matched data allow us to separately observe firms' product

[1] See Gallup polls in Appendix Figure B.1. As of summer 2024, 55% of Americans responded that immigration should be decreased (highest since 2001), and 41% named inflation the most important financial problem facing their family (highest recorded since poll started in 2005). These figures highlight the widespread public concern over both issues.

[2] For example, see the related articles in Wall Street Journal and The Washington Post

[3] The matching of scanner price and firm production is done using firm names and address information via GS1

sales (destination) and their production (labor sites), which enables us to disentangle the product demand and labor supply channels through which immigration affects prices. The household panel characteristics—such as frequency of purchases and trips—and granular product information allow us to investigate mechanisms beyond these direct demand and supply channels, including consumer search behavior and product variety effects. To ensure external validity and conduct additional robustness exercises, we supplement these data with additional county-level information from various datasets, including the Quarterly Census of Employment and Wages (QCEW), Bureau of Labor Statistics (BLS) City Price Index, and CoreLogic and Zillow Housing Price Indexes.

Armed with detailed data, we first establish a robust reduced-form finding: immigration inflows *decrease* prices at the county level, with effects concentrated among low-income immigrants and disproportionately benefiting lower-income households. In this initial county-level analysis, we utilize geographic variation akin to the existing regional immigration literature. We consistently find that counties with more international immigrant inflows experience smaller price increases for consumer packaged goods (CPG), a result that holds whether using a chain-weighted county-level price index derived from barcode data or the official BLS price index for robustness.[4] To establish a causal impact and circumvent endogeneity concerns—such as immigrant self-selection into economically dynamic or more expensive regions—we employ a frontier Bartik-type instrument. This approach carefully follows and extends the method developed in Burchardi et al. (2019) and Terry et al. (2024) to recent years, leveraging more than 100 years of historical ancestry distributions and migration patterns to generate exogenous variation in current immigration, building on a broader literature using historical factors as instruments (e.g., Acemoglu et al. 2001; Nunn 2008; Dell 2010). The downward pressure is concentrated among low-income, less-educated immigrants—who are more price-sensitive and engage in more intensive price search—with the benefits accruing disproportionately to lower-income households, who consume similar products as newly arriving immigrants.

The effect is economically significant and extends broadly. An increase of ten thousand new immigrants over a four-year period lowers CPG prices by approximately 0.58 percentage points—about 8% of the mean price growth and 20% of the standard deviation of cross-county price differences, comparable to the effect on wages found in the literature. The local effect survives an array of identification checks. Shift-share-robust standard errors are smaller than our state-clustered baseline, and a randomization test that reassigns origin shifters across ancestry networks places our estimate in the far left tail of the placebo distribution. A Rotemberg decomposition shows that no single origin drives the result, which is unchanged when we drop Mexico as an origin, and past inflation does not respond to contemporaneous immigration. The effect is also not an artifact of measurement: alternative index formulas deliver the same answer, as does measuring immigration relative to initial county population rather than in levels. The disinflationary effect generalizes beyond CPG to

database. See Appendix for details.

[4] Our scanner price index, based on CES utility (Sato, 1976; Vartia, 1976), is similar to the BLS methodology and presents a similar trend to that of the official BLS price index; see the Appendix for details.

non-durable goods, with muted effects on durables and housing purchase prices, and positive effects on housing rents. The net effect on regional cost of living remains negative, as housing rent's small CPI share means that non-durable price reductions dominate the aggregate despite rental increases.

Quantities tell us more about what kind of adjustment accompanies the price decline. Within continuing barcodes, physical quantity rises only modestly while private-label quantity rises sharply. Spending is reallocating toward cheaper goods, a margin our price index deliberately excludes because it holds the product fixed. The price decline of continuing barcodes is itself uniform, negative in all sixteen samples we can construct. Thus, the price decline on identical barcodes runs alongside a shift in what people buy, and the firm data let us ask whether that shift also happens inside firms.

Having established a robust county-level reduced-form relationship, we next shift our unit of analysis to the firm to investigate the underlying mechanisms. Price changes are, to a first approximation, a firm-level phenomenon. The typical firm sells in about twenty counties, so its pricing reflects the immigration its customer base experiences across all of them, while a county regression relates a county's prices only to that county's own immigration. County estimates also bundle these responses with local general equilibrium forces such as establishment entry, consumer brand switching, and cross-county migration. Because a firm's customers are dispersed across its sales markets while its workers are concentrated where it produces, exposure to immigrant consumers and exposure to immigrant workers can be measured separately. Specifically, we measure the increased product demand from immigration by analyzing inflows in a firm's sales locations, and the increased labor supply by focusing on inflows in its production locations. For each of these distinct exposures, we analyze not only price responses but also other key outcomes such as purchaser characteristics, shopping behavior, and average prices paid, providing a comprehensive understanding of firm strategies. In these analyses, we control for industry fixed effects and other potential confounding variables that could affect firm pricing, such as retail wages and rental prices. Because prices reach consumers through retailers, we also bring the retailer dimension into the specification directly, and find that prices track the firm's immigration exposure rather than the retail chain's, confirming the manufacturer-pricing channel.

Our firm-level analyses reveal a key mechanism driven by the product demand channel: newly arriving immigrants engage in more extensive search behavior, which induces firms to lower their prices. To substantiate this finding, we first disentangle the effects of product demand and labor supply. We find that firms more exposed to immigration in their sales markets *lower* their prices, while their price response is weakly *positive* when immigrants arrive in their production sites. A one standard deviation increase in demand exposure lowers firm prices by 2.7 percentage points, three times the size of the supply-side response. The same decline appears when we run the regression at the barcode level, without aggregating prices to the firm. Since immigrant demand expands the market, which under a scale effect would raise prices, the decline must reflect markup compression as more price-sensitive consumers flatten the demand curve firms face. We also find that the demand-side effect is strongest for firms selling to low-income and immigrant-intensive consumers, mirroring the

county-level finding that price declines are concentrated among low-income immigrants and the products they buy.

We validate this search-theoretic mechanism by investigating detailed purchase behavior at the firm level, effectively testing whether firms facing greater immigration exposure in their sales (or production) markets are systematically linked to more search-intensive consumers. We find that firms facing a larger influx of immigrants see an increase in purchasers who make more shopping trips and spend more days searching. Crucially, we find that this heightened search effort is associated with consumers paying below-average prices for identical products they buy from other firms, providing direct evidence for the force driving the lower equilibrium prices. Lastly, in response to this demand shift, we observe an increase in the share of private-label products, which are known to be lower-priced options that are favored by these search-intensive, price-elastic consumers.

We corroborate the search mechanism by structurally estimating a nested demand system. This framework nests and extends our reduced-form indexes, yielding a utility-based price index that integrates considerations of product variety (Feenstra, 1994) and appeal (Hottman et al., 2016) at the firm and county levels. Three key empirical findings from this framework further support the product search mechanism. First, our estimation of location-specific demand elasticities across products and firms reveals that greater immigration inflows are associated with a higher demand elasticity, particularly within the product demand channel, consistent with the search mechanism. Second, the estimated product appeal term falls significantly in response to increased product demand, which is consistent with the narratives that search-intensive immigrants also purchase lower-quality products. Lastly, we find that the variety effect through the product demand channel is largely muted, as measured by our structural variety index and by the number of products purchased per household, so variety-based competition is not a primary channel through which immigrant demand lowers prices. On the supply side, prices, variety, and appeal all rise, consistent with firms upgrading their product lines when immigrants arrive where they produce (Kerr and Lincoln, 2010; Terry et al., 2024).

We rationalize these findings with a parsimonious model in which immigrants are more price-sensitive than natives, a single primitive that accounts for both structural results. Search is the natural source of that difference. It operates through two margins, which are the two tiers of the demand system we estimate. Substitution across firms identifies the markup, so as immigrants become a larger share of a firm's buyers the elasticity it faces rises and its markup compresses. That compression works against a scale effect that raises marginal cost as the market expands, so whether prices fall is a quantitative question. Evaluated at our estimated elasticities and returns to scale, markup compression dominates at every elasticity we observe, and the implied price decline matches the one we measure. Substitution within the firm identifies composition, so facing the same products at the same prices immigrants spend more on the firm's cheaper and lower-quality items, which lowers measured appeal without any product changing.

Literature Review. Research on the economic effects of immigration has long focused on the

labor market, with a large and diverse body of literature exploring its effects on nominal wages and employment.[5] The general equilibrium effects on consumer prices remain relatively underexplored, leaving an incomplete view of the overall economic impact on real wages and cost of living.

A few pioneering papers have partially filled this gap, finding that the direction and magnitude of price effects depend critically on whether immigration operates through labor supply or consumer demand channels. On the supply side, Cortes (2008) documents that low-skilled immigration reduces prices of immigrant-intensive nontradable services through lower wages.[6] On the demand side, Lach (2007) finds that Soviet immigrants reduced prices in Israel through greater search intensity, while McCully et al. (2025) find no effect of immigration on the prices of imported goods in the US: immigrants may be more familiar with goods from their origin country and thus exhibit less price search behavior, with immigrant demand instead expanding import variety through fixed trade cost reductions. Relatedly, Argente et al. (2026) document that migrant networks diffuse information about new varieties across consumers, inducing domestic retailers to expand product assortments in response to latent demand. These papers point to a variety/availability margin through which immigrant demand reshapes retail environments, distinct from the price and quality composition responses we emphasize on continuing products. In housing markets, immigrant demand has been shown to raise rents in supply-constrained cities (Saiz, 2007; Sanchis-Guarner, 2023), though supply-side adjustments and native out-migration can offset this in the long run (Saiz and Wachter, 2011; Monras, 2020; Sá, 2015). Across these settings, the direction of price effects varies with the dominant channel, with some studies finding small positive aggregate effects (Furlanetto and Robstad, 2019) and others negative (Zachariadis, 2012). More broadly, this points to the importance of jointly accounting for both channels: Cheremukhin et al. (2024) shows theoretically that the two effects may largely offset, Galaasen et al. (2025) find that demand exposure to immigration has large positive effects on native wages in Norway's labor market distinct from supply exposure, and Cravino et al. (2026) develops a quantitative GE framework tracing deportation effects on wages and prices through both channels simultaneously. Our paper complements this emerging literature by using firm-level data to decompose the two channels in a reduced-form sense, finding that the disinflationary effect of immigration on US consumer prices is driven by the demand side — markup compression through immigrants' greater price sensitivity — rather than labor costs.

Our analyses reveal the importance of consumer search behavior, and consumer heterogeneity more broadly, among various alternative mechanisms in which immigration affect prices, consistent

[5] For instance, while some studies find that immigration can harm the wages of low-skilled natives (Borjas, 2003; Borjas and Katz, 2007; Dustmann et al., 2017), other papers report positive or negligible impacts, citing factors like native worker specialization and productivity gains (Card, 2009; Peri and Sparber, 2009; Ottaviano and Peri, 2012; Mahajan et al., 2024). The meta-analyses suggest that the overall average wage effect is close to zero, though with substantial heterogeneity across different groups (Dustmann et al., 2016; Peri, 2016; Clemens and Hunt, 2019).

[6] Dustmann and Glitz (2015) finds a parallel pattern in labor markets: immigration depresses wages in nontradable sectors but not in tradables, where firms absorb labor supply shocks through output adjustment. Burstein et al. (2020) shows theoretically that the elasticity of output to local prices is endogenously higher in more tradable sectors, so that immigration-driven changes in production translate into quantities rather than prices.

with the work of Lach (2007). This finding is broadly consistent with theoretical models that link consumer price sensitivity to firm markups and macroeconomic phenomena (Kaplan and Menzio, 2016) and studies analyzing how consumer behavior changes with aggregate shocks, such as recessions (Aguiar et al., 2013; Dube et al., 2018; Nevo and Wong, 2019). In particular, our results are parallel to the empirical findings that local shocks (e.g. house prices, incomes of other households) can affect consumer spending and price sensitivity, which in turn leads to changes in markups (Stroebel and Vavra, 2019; Sangani, 2024). Relatedly, our work complements cross-sectional studies on how prices vary with consumer characteristics (Manova and Zhang, 2012; Jaravel, 2019; Handbury, 2021; Faber and Fally, 2022; Auer et al., 2024; Mongey and Waugh, 2025; Zhu et al., 2025). In this vein, we find that immigration-induced price changes are relatively larger for lower-income households—in contrast to the regressive distributional effects documented in the context of negative wealth shocks or innovation targeted at high-income consumers (Jaravel, 2019; Argente and Lee, 2021).

To credibly estimate the causal effect of immigration, we employ a cutting-edge method in the immigration literature by exploiting exogenous variation in immigration inflows. Our approach follows and extends the Bartik-type instruments developed by Card (2001) and more recently refined and applied by Burchardi et al. (2019) and Terry et al. (2024). This methodology leverages long-standing settlement patterns and country-of-origin networks to predict the destination of new immigrants, which allows us to isolate the immigration shock from other local economic factors that might simultaneously influence prices. In disentangling the product demand and labor supply effects, we utilize the NielsenIQ-GS1-NETS matched data to observe places where firms sell products as well as where they produce products.

To corroborate the mechanism, we estimate a nested demand system which allows for the time-varying product variety and consumer preferences at the micro-level, a technique widely used in macroeconomics, trade, and development (e.g., Feenstra and Romalis (2014); Hottman, Redding and Weinstein (2016); Atkin, Faber and Gonzalez-Navarro (2018); Lenzu, Rivers and Tielens (2022); Eslava, Haltiwanger and Urdaneta (2024)). This approach allows us to consistently aggregate barcode-level prices into firm- and county-level indices while accommodating crucial changes in product varieties, quality, and consumer substitution—important considerations given that immigration-induced price effects may operate through changes in available varieties and product quality, not just unit prices (Handbury and Weinstein, 2015). Our findings that immigration heterogeneously impacts how firms adjust prices across counties align with existing literature on price dispersion and discrimination (Simonovska, 2015; Fitzgerald et al., 2024).

The paper is organized as follows. Section 2 discusses the theoretical channels through which immigration can affect prices, distinguishing between labor supply and product demand mechanisms. Section 3 describes our extensive data construction and empirical strategy. Section 4 presents our main findings on how immigration affects prices at the county level, establishing the aggregate relationship and providing external validity across multiple sectors. Section 5 investigates the underlying mechanisms using firm-level variation in exposure to immigration through sales and

production locations. Section 6 develops a parsimonious theoretical framework with household heterogeneity to rationalize these findings. Section 7 concludes.

2 The Link Between Immigration and Prices

Immigration can influence consumer prices through two primary channels: labor supply effects and product demand effects. These channels can have opposing impacts on the equilibrium price level, making the net effect theoretically ambiguous. Table 1 summarizes the key mechanisms through which immigration affects prices based on our survey of the literature, acknowledging that other mechanisms may also be at play.

Table 1: Channels Through Which Immigration Affects Prices

Primary Channel		**Price**	**Description**	**Source**
Labor Supply	Lower Wages	Fall	Increased labor supply reduces wages and production costs.	Borjas 03, Cortes 08, Monras 20
	Productivity		Task specialization increases productivity, lowering unit costs.	Peri Sparber 09
	Entrepreneurship		Immigrant entrepreneurs increase firm entry and competition.	Olney 13, Mahajan 24
	Skill Upgrading	Rise	Natives shift to higher-skilled tasks, raising their wages.	Kerr et al. 15
	Innovation		Increased innovation and capital investment raise costs.	Kerr Lincoln 10, Burchardi et al. 19, Terry et al. 24
Product Demand	Demand Pull	Rise	Larger population increases overall demand.	Saiz 07
	Price Sensitivity	Fall	Search-intensive immigrants increase price competition.	Lach 07

2.1 Labor Supply Channel

Immigration directly expands the local labor force, which can affect prices through multiple supply-side mechanisms. The most direct effect operates through labor costs: an influx of workers increases labor supply and can reduce wages, particularly for occupations where immigrants concentrate (Borjas, 2003; Cortes, 2008; Monras, 2020). Beyond wages, immigration can enhance productivity through task specialization, as natives shift toward communication-intensive tasks (Peri and Sparber, 2009), or by selecting into the most productive firms (Brinatti and Morales, 2025). Immigrants also contribute to entrepreneurship and firm entry (Mahajan, 2024; Khanna and Lee, 2019; Olney, 2013), expanding competition and putting further downward pressure on prices.

On the other hand, several countervailing forces can lead to upward price pressure. As natives reallocate to more complex, higher-skilled occupations (Kerr et al., 2015), their wages increase, raising production costs. Immigration can also stimulate innovation, though evidence is mixed: while some studies find positive effects on patents and firm performance (Kerr and Lincoln, 2010; Clemens and Lewis, 2024; Amuedo-Dorantes et al., 2023), others document negative spillovers in

innovation-intensive fields (Borjas and Doran, 2012). Historical immigration networks may further facilitate innovation and entrepreneurship in ways that raise costs, such as through fixed R&D and overhead costs that are spread across a firm's products (Burchardi et al., 2019; Terry et al., 2024).

2.2 Product Demand Channel

On the demand side, immigration increases the local consumer base, creating upward pressure on prices through standard demand-pull mechanisms. This effect is most clearly documented in housing markets (Saiz, 2007).

However, immigrants may affect prices differently than a simple population increase would suggest. Lach (2007) documents that a large influx of Soviet immigrants to Israel actually reduced prices, attributing this to immigrants' greater price sensitivity and search intensity. New arrivals, often facing information frictions and lower initial incomes, have stronger incentives to search for better deals, increasing effective price competition and reducing markups even as total demand expands.

Compositional shifts also alter the aggregate willingness to pay for quality. Handbury (2021) and Argente and Lee (2021) document that demand for quality is highly income-dependent, creating spatial and demographic variation in price indices. Consequently, a shift in the consumer base can induce a "trading down" effect (Jaimovich et al., 2019), or prompt firms to adjust product appeal to target specific demographic segments (Goetz and Rodnyansky, 2023; Jaravel, 2019). In our setting, the arrival of immigrants—who may exhibit both higher price elasticity and a lower valuation for high-appeal goods—intensifies price competition and reweights consumption bundles.

3 Data and Empirical Strategy

This section describes our data construction and empirical approach. Our core contribution is the ability to analyze how firm prices respond to immigration exposure, decomposed by where firms sell their products (a product-demand exposure) and where they produce them (a labor-supply exposure). This firm-level analysis is enabled by matching retail scanner prices to producer locations, combined with an instrumental variable strategy that generates plausibly exogenous variation in immigration exposure.

We begin by describing our primary data sources: retail scanner prices and household shopping behavior for consumer packaged goods (CPG), supplemented by additional price measures across sectors. We then detail our key innovation, matching products to their producers' geographic locations, which allows us to construct firm-specific measures of immigration exposure through demand and supply channels. After describing our immigration data and instrumental variable strategy, we present summary statistics that document substantial spatial price variation even for identical products, motivating both our county-level baseline analysis and our main firm-level investigation.

3.1 Data Sources

3.1.1 Retail Scanner Prices and Household Shopping Behavior

Our primary data come from two NielsenIQ datasets that provide complementary information on prices and consumer behavior for CPG products. The Retail Scanner Dataset (RMS) contains weekly pricing, sales volume, and merchandising conditions from store point-of-sale systems across the United States from 2006 to 2019. The dataset includes approximately 35,000 participating stores in a balanced panel spanning drug stores, grocery stores, and mass merchandise retailers, covering roughly 2,500 counties, about 83% of all U.S. counties. Coverage is substantial: the data capture approximately 53-55% of national sales in food and drug stores and 32% of national sales in mass merchandise stores. Store locations are identified at the county level, which is the finest geographic unit available in the RMS data.

The dataset encompasses 2.6 million universal product codes (UPCs) organized into approximately 1,100 product modules, which aggregate further into 125 product groups covering consumer packaged goods (CPG), such as food, beverages, health and beauty products, and general merchandise. This rich product detail allows us to track prices of identical products, such as a specific cherry-flavored 500ml diet coke or a particular brand and size of laundry detergent, across different counties and time periods. This feature is critical for constructing quality-consistent price indices and documenting the spatial price variation that we exploit for identification.

To construct our analytical dataset, we perform several aggregation steps. First, we aggregate the raw store-UPC-week data to the store-UPC-year level, computing quantity-weighted average prices. We also construct alternative measures using modal prices within each year and obtain nearly identical results. Second, we aggregate to the county-UPC-year level by taking quantity-weighted averages across stores. Within each county we use only a balanced panel of stores, those present in every year of the 2006–2019 sample, so that store entry and exit do not mechanically drive our price measures. The resulting index measures price changes for continuing (incumbent) stores. This county-UPC-year panel is the foundation for all subsequent price index construction, described in Section 3.3.

We complement the retail scanner price data with the NielsenIQ Consumer Panel, which provides detailed shopping behavior and demographics for approximately 55,000 households annually from 2004 to 2019. Panel participants record every shopping trip using handheld scanners or mobile apps. For each household-year, we observe the number of shopping trips made, the number of days spent shopping, the number of visits to different retail stores, the store brands purchased, and the purchases made at each retailer. The panel also includes household income, education, household size, race, and ethnicity. To this we add a predicted immigrant household status based on a machine learning classifier trained on ACS microdata and domain-adapted to the panel using the 2008 Tell Us More About You (TUM) survey as a labeled bridge (see Appendix A.2).[7]

[7] An alternative approach is to directly use households surveyed in the 2008 TUM module, as in McCully et al.

We use the household panel in two ways. First, we measure how shopping behavior, in particular search intensity, varies with local immigration. Second, we use the demographic information to identify likely immigrant households and test whether they shop differently. These behavioral differences are central to the price-sensitivity mechanism we document in Section 5.

Supplementary Price Data Across Sectors We also assemble price measures for other expenditure categories to test whether immigration's effects extend beyond CPG. For housing purchase prices we use county home price indices from CoreLogic and Zillow. For rents we build county median gross rent indices from the ACS, which cover about 3,100 counties. For broader goods and services we use city-level BLS Consumer Price Index data for the 27 largest metro areas (261 counties). These measures have coarser geographic or product detail than the scanner data, but they let us gauge how general our findings are across the consumption basket.

3.1.2 Firm Production Locations

A key strength of this paper is matching consumer packaged products to their producers' geographic locations, which allows us to separately identify labor supply and product demand channels. We link firms in the scanner data to establishment-level information from the National Establishment Time-Series Database (NETS), which provides panel data on employment, sales, and locations for private and public businesses across the United States from 2006 to 2018.

The matching procedure uses business names and addresses from the GS1 Company Database, which maintains the registry of UPC codes and their associated manufacturers. GS1 is the global standards organization that administers the UPC system and assigns unique identifiers to products and companies. We use a multi-step algorithm to match GS1 companies to NETS establishments. First, we match entities with identical names, states, and cities, which yields high-confidence matches for firms with unique names in specific locations. For remaining firms with the same name and state but different cities, often large national companies with multiple facilities, we use the Google Places API to verify matches based on common URL appearances in search results. Specifically, if a search for the firm name, city, and state returns a common root domain in the top results for both the GS1 and NETS entries, we classify these as successful matches. We manually review a random sample of 500 matches to validate the algorithm and find a success rate of 94%, confirming that the procedure produces reliable linkages.

This procedure yields high-quality matches between products (UPCs) and their manufacturing establishments, which we categorize by industry using Standard Industrial Classification (SIC) codes from NETS. The matched data let us observe, for each product sold in the scanner data, the firm that produces it, the county or counties where the firm operates production establishments, the number of workers at each establishment, and the counties where the product is sold. This link of

(2025). Since our analysis focuses on new immigration inflows over 2006–2019, however, this would restrict the sample to households present in 2008 and miss immigrants who arrived in subsequent years.

production and sales locations is essential for our firm-level analysis in Section 5, where we construct firm-specific measures of exposure to immigration through labor supply at production counties and product demand at sales counties. Appendix A.1 provides additional detail on the matching algorithm, validation, match rates, and the representativeness of matched versus unmatched firms.

3.1.3 Immigration Data

We construct county-level immigration following Terry et al. (2024), using IPUMS microdata. For the historical periods behind the instrument we use the 1880, 1900–1930, and 1970–2000 census samples. For the main period we use the ACS, assigning each vintage to the immigration it best captures: the 2010 ACS for immigration through 2005, the 2014 ACS for 2005–2010, the 2018 ACS for 2010–2014, and the 2022 ACS for 2014–2018. We use later vintages because IPUMS pools five years of surveys after 2000, so the 2010 ACS alone would undercount recent arrivals. The 2022 ACS was the latest available at the data collection stage, which is why the sample ends in 2018.

Immigration is defined as the total number of foreign-born respondents by birthplace. We construct immigration flows (i.e. *new* immigrants) for multiple time periods using reported year of immigration. Ancestry is measured from respondents' reported ancestral origins, with country classifications harmonized across time using extended crosswalks based on Burchardi et al. (2019). We scale individual responses to county-level totals using IPUMS person weights and harmonize county codes to a consistent 1990 classification, yielding a balanced panel of 3,141 counties across all sample years.

The ACS data also provide rich information on immigrant characteristics at the county level, including educational attainment and skill, industry and occupation, income and earnings, and year of arrival. With these characteristics we examine heterogeneous effects by immigrant type, such as the price effects of high-skilled versus low-skilled immigration, explored in the mechanism analysis.

3.1.4 Supplementary County-Level Data

To control for confounders and examine additional outcomes, we incorporate several supplementary county-level datasets. The Quarterly Census of Employment and Wages (QCEW) from the Bureau of Labor Statistics provides county-level annual data on average wages (deflated by the PCE index), total employment, number of establishments, and total payrolls from 2006-2018. From the Bureau of Economic Analysis Regional Data we obtain county GDP, per capita income, and population for the same period, along with urban and rural classifications. The USDA Economic Research Service provides county population density and metropolitan status.

These data serve two roles. They provide controls for local economic conditions that could correlate with both immigration and prices, and they let us examine whether immigration affects local outcomes such as wages and employment that might mediate the price effects.

3.2 Baseline Empirical Specification

Our baseline specification examines how changes in county-level prices respond to immigration inflows:

$$\Delta \ln P_{dt} = \delta_t + \delta_s + \gamma I_{dt} + X'_{dt}\beta + \varepsilon_{dt}, \tag{3.1}$$

where $\Delta \ln P_{dt}$ is the log change in the price index for county d over period t, I_{dt} is the total number of new immigrants arriving in county d during period t, δ_t and δ_s are period and state fixed effects, and X_{dt} includes initial county-level controls such as population and GDP per capita. The coefficient γ captures the causal effect of immigration on local price inflation. We construct an instrument for I_{dt} to address endogeneity, discussed below.

We estimate this specification using a stacked long-difference approach covering three periods: 2006-2010, 2010-2014, and 2014-2018. We also report robustness checks using the full long-difference period (2006-2018). Pooling across periods focuses on medium-term adjustments, mitigates high-frequency noise in annual price changes, and allows time for firm-level adjustment. State and period fixed effects control for time-invariant state characteristics and common national shocks, isolating variation in immigration exposure across counties that deviates from national trends.

We estimate equation (3.1) separately for different price indices, using NielsenIQ CPG non-durables as our primary focus and supplementing with housing purchase prices, housing rents, durable goods, and services, to analyze immigration's overall price impact across the consumption basket as fully as the data permit.

3.3 Construction of the CPG Price Index

We build a price index from the barcode data that is consistent with CES preferences and accounts for consumer substitution across products and firms. The index follows the exact price index literature (Sato, 1976; Vartia, 1976; Diewert, 1976). In Section 5.4, we develop a full structural demand model that nests this reduced-form index and additionally incorporates product variety and quality adjustments. For now, we focus on a transparent, theory-consistent aggregation of observed prices for continuing varieties. We summarize the construction here and give the full detail in Appendix C.2.

Let P_{kdt} be the price of product (UPC) k in county d at time t, where k is made by firm f and belongs to product group g. We build the county index in three chained stages. Each stage uses only varieties present in both periods, so that entry and exit do not mechanically drive the measure. Stage 1 aggregates products within each firm, product group, and county using Sato-Vartia weights:

$$\frac{P_{fgdt}}{P_{fgd,t-1}} = \prod_{k \in \Omega_{fgdt,t-1}} \left(\frac{P_{kdt}}{P_{kd,t-1}} \right)^{w_{kdt}}, \tag{3.2}$$

where $\Omega_{fgdt,t-1}$ is the set of continuing products and the Sato-Vartia weights w_{kdt} are based on expenditure shares within the firm, product group, and county being aggregated (Appendix C.2).

Sato-Vartia weights are exact for CES and give more weight to products with larger expenditure shares. Stage 2 aggregates firms within each product group and county, again with Sato-Vartia weights:

$$\frac{P_{gdt}}{P_{gd,t-1}} = \prod_{f \in \Omega_{gdt,t-1}} \left(\frac{P_{fgdt}}{P_{fgd,t-1}} \right)^{w_{fgdt}}. \tag{3.3}$$

Stage 3 aggregates product groups within the county using Törnqvist weights, of which Cobb-Douglas is a special case:

$$\frac{P_{dt}}{P_{d,t-1}} = \prod_{g \in \Omega_{dt,t-1}} \left(\frac{P_{gdt}}{P_{gd,t-1}} \right)^{w_{gdt}}, \tag{3.4}$$

where $w_{gdt} = \frac{1}{2}(s_{gdt} + s_{gd,t-1})$ is the average expenditure share of group g across the two periods.

The log change $\Delta \ln P_{dt}$ is the outcome in equation (3.1). By construction it holds the set of products, firms, and groups fixed, so it captures pure price variation rather than changes in product mix or quality. This matters for interpretation. The price declines we document in Section 4 cannot come from immigrants buying different, cheaper products. That composition margin is measured separately by the appeal index of Section 5.4. Appendix Table B.4 shows the results are robust to Jevons, Laspeyres, Paasche, and Törnqvist alternatives, and Figure B.2 shows the national aggregate of our index tracks the BLS food price index closely.

3.4 Instrumental Variable Strategy

The immigration measure I_{dt} in equation (3.1) is likely endogenous. Immigrants select into counties on economic opportunities, wages, housing costs, and amenities, any of which also move local prices. If immigrants are drawn to booming counties with rising prices, OLS is biased upward. If they sort into cheaper, declining areas, OLS is biased downward. Either way, isolating the causal effect requires an instrument.

We use the predicted-ancestry shift-share instrument of Terry et al. (2024), who refine the canonical instrument of Card (2001). The idea is to predict where immigrants settle from a century of historical settlement patterns interacted with contemporary origin-country outflows, after purging the historical settlement measure of any component correlated with persistent local shocks. We summarize the instrument here and give the full construction, together with its motivation and the example that builds intuition, in Appendix B.2.

Our instrument for immigration to county d in period t interacts two components and sums across origin countries o:

$$Z_{dt} = \sum_{o} \hat{b}_t \left(\hat{A}_{od,t-1} \times \tilde{I}_{o,-r(d),t} \right). \tag{3.5}$$

The first component, predicted ancestry $\hat{A}_{od,t-1}$, is a forecast of how many residents of origin-o descent live in county d. It is built entirely from historical (1880 onward) push-pull variation, namely origin-country emigration pressures interacted with a destination's historical attractiveness

to an *unrelated* immigrant group (Europeans), and it is residualized against controls so that it carries no persistent local economic component (Appendix B.2, equations (B.2) and (B.3)). The second component, the shifter $\tilde{I}_{o,-r(d),t}$, is contemporary origin-o immigration to regions other than d's, scaled for region size, so it reflects global migration trends realized outside county d. The weights $\hat{b}_t$ come from the first stage that maps these interactions into actual immigration (Appendix equation (B.4)). Intuitively, $\hat{A}_{od,t-1}$ isolates settlement driven by century-old forces and $\tilde{I}_{o,-r(d),t}$ isolates origin-country conditions outside d, so their interaction predicts where new immigrants arrive without relying on d's current economy.

The exclusion restriction is that these historical settlement forces and out-of-region contemporary flows affect current county prices only through current immigration. This is plausible because the century-old economic, transportation, and institutional factors that shaped settlement are unlikely to bear directly on 2006–2018 CPG prices once we condition on broad geographic fixed effects, and because the contemporary shifter is realized outside county d. Terry et al. (2024) provide extensive evidence that this instrument satisfies the conditions for valid shift-share inference. We add checks specific to our price setting in Section 4, including randomization inference, a Rotemberg decomposition, and a specification that drops Mexico as an origin.

Implementation and first-stage validation. We construct the instrument for the three periods in our sample, 2006–2010, 2010–2014, and 2014–2018, using only non-European origin countries, since European settlement is used to measure destination attractiveness. Non-Europeans are over 80% of recent U.S. inflows, so this is the policy-relevant margin. Predicted ancestry draws on historical migration back to 1880, while the first stage that produces the $\hat{b}_t$ is estimated over nine periods spanning 1980–2018 (Appendix B.2). Appendix Table B.3 reports first-stage estimates that closely replicate Terry et al. (2024) for overlapping years and extend them through 2018. Our constructed instruments correlate at 0.997 with theirs over the common sample, and the excluded-instrument F-statistic exceeds 40 in every specification, well above weak-instrument thresholds.

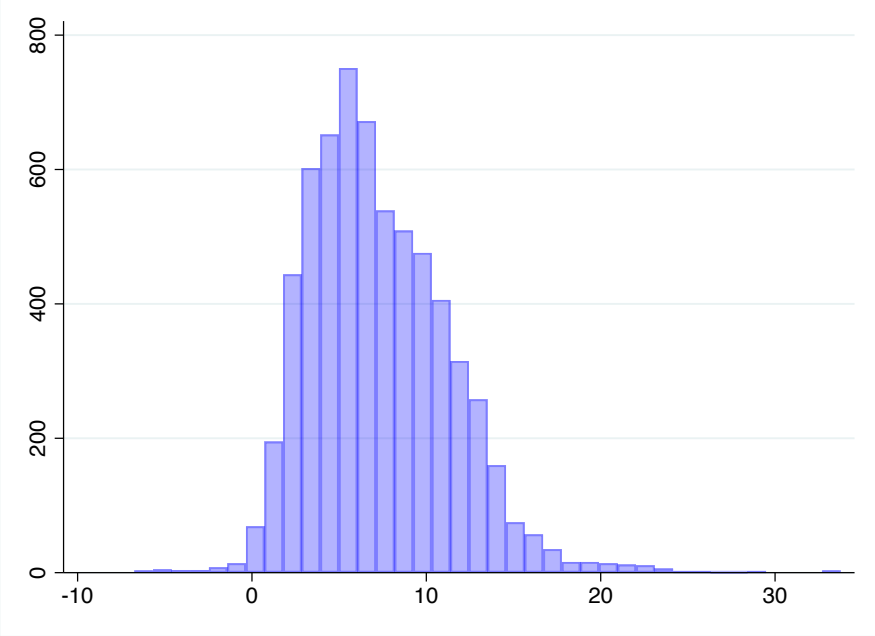
Finally, note that the predicted-ancestry approach produces immigration shocks in levels rather than population shares (Cortes, 2008; Peri and Sparber, 2009). This is consistent with the competing population-composition and scale channels in our theory (Section 6), and Table B.7 confirms robustness to various specifications that normalize by initial population size.

3.5 Summary Statistics and Spatial Variation in Prices

Before presenting our formal results, we document a key feature of our setting: substantial spatial variation in price changes across counties. This price dispersion is critical for understanding how local immigration shocks can affect local prices, and it motivates both our county-level baseline and our firm-level investigation of exposure through sales versus production locations.

Price Dispersion Across Counties Figure 1 presents the distribution of county-level price changes for CPG products across four-year periods from 2006-2018, using our Sato-Vartia price index. Even after holding the product set fixed to exclude entry and exit, while still allowing expenditure shares to change, we observe considerable geographic price dispersion. The median county experienced price increases of about 6.7% over four-year periods, and the interquartile range spans 5.6 percentage points, from 4.3% at the 25th percentile to 9.9% at the 75th. Counties at the 75th percentile experienced inflation more than twice as high as those at the 25th, indicating substantial variation in local price dynamics.

Figure 1: Distribution of County-Level Price Changes, Four-Year Periods 2006-2018

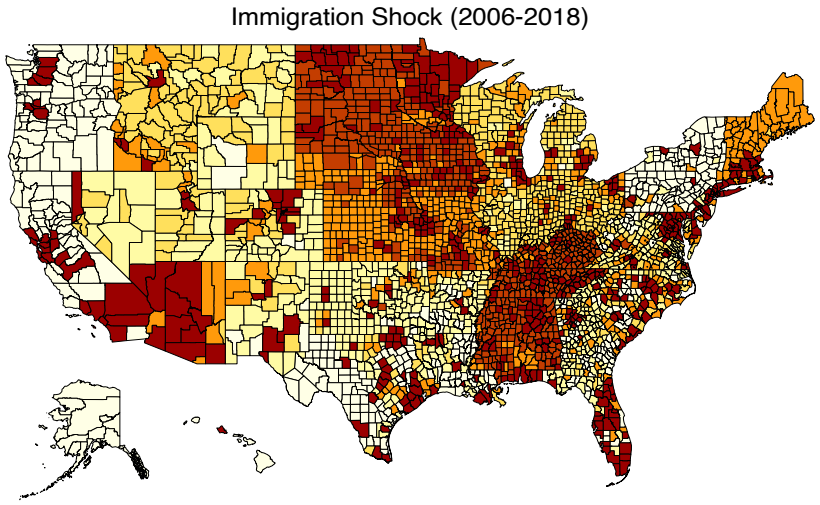


Notes: Histogram shows the unweighted distribution of price changes across counties over four-year periods from 2006-2018, measured using the Sato-Vartia price index for CPG products as described in Section 3.3. Each observation is a county-period (N=6,339). The price index is the baseline Sato-Vartia index over continuing products, holding the product set fixed while allowing expenditure shares to vary. Product entry and exit are handled separately via the Feenstra variety correction (Section 5.4). The distribution demonstrates substantial spatial price variation for consumer goods.

This spatial price variation is essential for our strategy. If purchasing power parity held for these goods, local immigration could not move local prices, because any gap would be arbitraged away through trade. In practice prices differ across locations for several reasons: manufacturers and retailers price-discriminate across markets (Fitzgerald et al., 2024), local promotions and competition vary, and transportation and supply-chain costs create wedges. Because our index aggregates across products and firms, it reflects these differences in pricing across manufacturers and retail formats.[8]

[8]Some retailers set near-uniform prices within their chains (DellaVigna and Gentzkow, 2019). Because our index aggregates across manufacturers and retail formats, it still captures the cross-market variation we exploit.

Figure 2: Geographic Distribution of Immigration Shocks, 2006-2018



Notes: Heat map shows predicted immigration shock Z_{dt} from equation (3.5) at the county level, cumulated over 2006-2018. Darker colors indicate larger predicted immigration inflows. The instrument uses only non-European immigration and is based on historical settlement patterns from 1880-1930 interacted with recent origin-country migration patterns. White areas indicate counties not included in the sample due to missing data.

Geographic Distribution of Immigration Shocks Figure 2 displays the spatial distribution of predicted immigration shocks across counties from 2006-2018. The map reveals substantial geographic variation in exposure that is plausibly orthogonal to local economic conditions, given our instrument construction (Section 3.4). The highest predicted inflows are concentrated in the Southwest, California, and Florida, but they also appear throughout the country, including the North-Central region, Appalachia, and the East Coast. This widespread variation, combined with the exogenous source of the instrument, provides strong identifying power.

Summary Statistics Across Price Measures and Sectors Table 2 presents summary statistics for our main variables at the county level over stacked four-year periods. For CPG products, the mean price increase of 7.35% with a standard deviation of 4.12 percentage points confirms the substantial dispersion in Figure 1.

Price variation across sectors follows patterns consistent with tradability. CPG products show considerable dispersion, while housing and rental prices show even more, reflecting housing's non-tradable nature. In the BLS data for 27 metropolitan areas, non-durables and durables show standard deviations of 5.40 and 3.35 percentage points, and services show more modest dispersion of 2.59 percentage points.[9]

Wages show a standard deviation comparable to CPG price dispersion, even though labor is less mobile than goods. We also report wages in levels (hundreds of 2012 dollars), and use this measure

[9]The modest service price dispersion within large metropolitan areas likely reflects greater price convergence in integrated urban markets.

Table 2: Summary Statistics: County Level

	count	mean	sd	p25	p50	p75	min	max
Δ ln $P_{d,t}$ (CPG)	6339	7.35	4.12	4.33	6.74	9.95	-6.70	35.4
Δ ln $P_{d,t}$ (Non-durables)	684	5.96	5.40	-0.18	8.14	10.1	-5.19	13.4
Δ ln $P_{d,t}$ (Durables)	684	-3.63	3.35	-6.42	-3.69	-1.57	-9.81	3.48
Δ ln $P_{d,t}$ (Services)	684	9.22	2.59	7.78	8.71	10.1	2.95	17.5
Δ ln $P_{d,t}$ (BLS, Rent)	684	11.2	4.77	8.16	10.1	14.1	2.09	25.8
Δ ln $P_{d,t}$ (ACS, Rent)	6336	10.7	6.45	7.31	10.9	14.3	-29.9	53.5
Δ ln $HP_{d,t}$(CL)	2446	-3.85	18.9	-11.8	-0.74	7.56	-87.9	44.6
Δ ln $HP_{d,t}$ (Z)	4781	6.03	16.2	-1.97	6.26	15.9	-108.6	57.2
Δ # $\text{Establ.}_{d,t}$	6339	-0.15	0.72	-0.45	-0.14	0.15	-8.44	22.4
Δ $\text{Employment}_{d,t}$	6339	1.77	15.0	-0.44	0.088	1.08	-305.0	297.9
Δ ln $\text{Wage}_{d,t}$	6339	3.27	5.30	0.36	3.17	6.00	-30.2	56.2
Δ ln $\text{Wage}_{d,t}$ (Retail)	6318	1.96	6.85	-1.85	2.04	5.87	-74.6	64.7
Δ $\text{Wage}_{d,t}$	6339	13.1	23.3	1.41	12.2	23.2	-163.8	288.0
Δ $\text{Wage}_{d,t}$ (Retail)	6339	4.79	22.6	-4.65	5.03	14.7	-288.1	452.6
$\text{Immigr}_{d,t}$	6339	1.68	4.95	0.066	0.17	0.62	0.0026	34.2
$\text{Immishock}_{d,t}$	6339	-0.035	0.48	-0.10	-0.042	-0.0013	-2.27	3.14

Notes: Summary statistics for stacked four-year periods (2006-2010, 2010-2014, 2014-2018). CPG price indices are constructed from barcode-level data following the procedure in equations (3.2)-(3.4). BLS non-durable, durable, service, and rental price indices are available for 27 metropolitan areas (261 counties). Housing purchase price indices are from CoreLogic (CL, available for 2006-2014 only) and Zillow (Z, available for all three periods). Housing rental prices are median gross rents from ACS data. Number of establishments (in thousands), employment, and wages are from QCEW data; wages are deflated by the national PCE index. Immigration ("Immigr") is measured for non-European origins from IPUMS data. The immigration shock ("Immishock") Z_{dt} is the predicted immigration instrument from equation (3.5), constructed using residualized predicted ancestry. Both Immigr and Immishock winsorize top and bottom 1% of the sample.

in regressions for comparability with Terry et al. (2024) wages results.

Immigration patterns confirm the geographic variation in Figure 2. The average county received 1,680 immigrants per four-year period, with a standard deviation of 4,950, three times the mean, reflecting extreme heterogeneity in settlement.[10] The predicted shock Z_{dt} shows substantial variation that, combined with its exogenous construction, provides strong identifying power.

4 County-Level Results

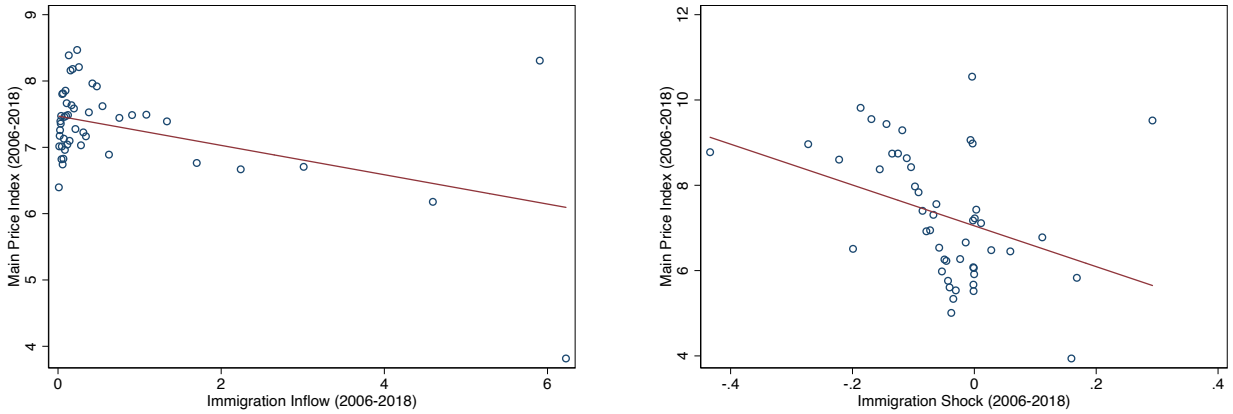
This section presents our baseline findings on how immigration affects consumer prices at the county level. We establish four main results. First, immigration reduces CPG prices at the county level. Second, this disinflationary effect is heterogeneous: it is concentrated among low-income, less-educated immigrants and is larger for products purchased by low-income households. Third, price reductions extend beyond CPG products to non-durable goods more broadly, and yield a net negative effect

[10]The reported dispersion is conservative, due to the 1% winsorization for county-level immigration. When applying the sales weights used in our regressions, the standard deviation increases to approximately 12,500 (comparable to Terry et al. (2024)). Note that we only apply this winsorization to our county-level analysis. For the firm-level specifications in Section 5, we use unwinsorized data because aggregating exposure across a firm's national sales footprint naturally smooths out extreme localized shocks.

in the regional cost-of-living despite modest rental increases. Fourth, immigration increases local economic activity—raising employment and the number of establishments—while we find no evidence of negative wage effects.[11] As we discuss below, the pattern of results across immigrants, product types, sectors, and local labor markets is consistent with a search-based mechanism whereby newly arriving immigrants engage in intensive price search, inducing firms to lower prices.

4.1 Immigration and CPG Prices

Figure 3: Immigration and Price Changes Across Counties, 2006-2018



(a) OLS: Actual Immigration

(b) Reduced Form: Immigration Shock

Notes: Scatter plots show the relationship between cumulative CPG price changes (y-axis) and immigration measures (x-axis) at the county level from 2006-2018. Panel (a) uses actual immigration inflows I_{dt}, while panel (b) uses predicted immigration shocks Z_{dt} from equation (3.5). Price changes are measured using the Sato-Vartia index described in Section 3.3. Both variables are winsorized at the 5% level to reduce the influence of outliers. County-level observations are grouped into 50 equal-sized bins based on their immigration measure. Each point represents the mean price change and mean immigration level within that bin, with point sizes proportional to the total initial population of the counties in the bin. Fitted lines are from bivariate regressions without controls or fixed effects.

Figure 3 provides initial visual evidence of the relationship between immigration and price changes. The left panel plots county-level price changes against actual immigration inflows, revealing a weak negative relationship. The right panel plots price changes against our predicted immigration shock Z_{dt}—the instrument based on historical settlement patterns—and shows a notably clearer negative relationship. This strong reduced-form relationship provides initial visual evidence that immigration causally reduces prices. Combined with the weak OLS pattern, this suggests that instrumental variables estimates may reveal larger causal effects than OLS, consistent with endogenous location choices attenuating OLS estimates. Both panels cumulate changes over 2006-2018 and show bivariate relationships without controls or fixed effects, motivating our formal regression analysis.

[11] Wage effects are positive or close to zero across specifications, with the magnitude sensitive to winsorization and the choice of log versus level specification. As specifications and sample periods approach those of Terry et al. (2024), the positive wage effect becomes more comparable to their findings (see Appendix Table B.13). Consistent with Terry et al. (2024) and the broader literature (Dustmann et al., 2013), we find that wages increase more in counties receiving

Table 3: Immigration and CPG Prices: Baseline Results

	Δ $P^{SV,Stacked}$ (2006-2018)					
	(1)	(2)	(3)	(4)	(5)	(6)
$Immigr_{d,t}$	-0.038***	-0.026***	-0.021***	-0.068***	-0.065***	-0.058***
	(0.009)	(0.004)	(0.004)	(0.013)	(0.012)	(0.013)
Obs.	6,339	6,339	6,249	6,339	6,339	6,249
First-Stage F-stat				254.4	94.4	85.5
Fixed Effects	Time	State, Time	State, Time	Time	State, Time	State, Time
Method	OLS	OLS	OLS	IV	IV	IV
Controls			Yes			Yes

Notes: Table reports regressions of log changes in county-level CPG price indices on immigration, estimated using equation (3.1). The price index uses Sato-Vartia weights. Columns (1)-(3) report OLS estimates; columns (4)-(6) report IV estimates instrumenting immigration with predicted shocks from equation (3.5). Immigration is winsorized at top and bottom 1%. All specifications use stacked four-year differences (2006-2010, 2010-2014, 2014-2018). Column (1) includes only period fixed effects. Column (2) adds state and period fixed effects. Column (3) adds interactions of initial county GDP per capita and urban indicator with period dummies. Columns (4)-(6) follow the same structure with IV estimation. All columns include sales weights based on the CPG data. Standard errors clustered by state. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 3 presents our baseline regression results, confirming that immigration reduces CPG prices. The OLS estimates in columns (1)-(3) show modest negative effects, and the IV specifications in columns (4)-(6) yield substantially larger negative coefficients, suggesting that immigrants positively select into counties with stronger economic conditions that would otherwise experience higher inflation. This pattern is consistent with immigrants seeking better economic opportunities, causing OLS to underestimate the causal price effects.

Our preferred specification in column (6) includes state and period fixed effects and controls for initial county characteristics interacted with period dummies. The estimated coefficient of -0.058 translates to an immigration inflow of 10,000 people (a shock comparable to that used by Terry et al. (2024)), reducing CPG prices by approximately 0.58 percentage points.[12] Relative to the mean cumulative CPG inflation rate of 7.27% over four years in our sample, this price reduction represents approximately an 8% ($\approx 0.58/7.27$) proportional reduction in the average inflation rate. This price response is economically meaningful: in proportional terms, it is comparable to the 8% effect on wage growth relative to the mean found by Terry et al. (2024), suggesting that immigration's effect on the cost of living may be as important as its much-debated effect on nominal wages—and that real wage effects are likely larger than nominal wage effects alone would imply.

The stability of IV coefficients across columns (4)-(6) as we add state fixed effects and county-level controls demonstrates robustness to controlling for regional trends and differential county characteristics. Results are also robust to alternative price index formulations: Appendix Table B.4 shows that Jevons, Laspeyres, Paasche, and Törnqvist indices all yield statistically significant

higher-educated immigrants.

[12] See discussion of Table B.7 in the Appendix for a comparison with existing estimates. Our estimates are within the range of existing studies finding negative price effects of immigration, though direct comparison is difficult given differences in settings, sectors, and specifications.

negative effects, with magnitudes broadly consistent with our baseline Sato-Vartia specification. Tables B.5–B.6 present further robustness checks—including long-difference specifications, controls for Mexican border proximity and local ICE enforcement, a test for anticipation/pre-trends through a one-period immmigration lead, state-by-period fixed effects, removal of winsorization, and county fixed effects—with remarkably stable results throughout. Finally, Table B.7 shows results are also robust to initial population controls, including an instrument normalized by initial county population, population-by-year fixed effects, and controls for native population growth (Lach, 2007). The positive coefficient on native population growth is consistent with population and immigration operating through distinct channels, as discussed in Section 6.

Inference and Placebo Terry et al. (2024) provide convincing evidence that the predicted-ancestry based instrument is robust to common empirical critiques of shift-share designs, satisfying the necessary conditions for valid causal inference. We add similar checks specific to our setting. First, we compute the standard error on the immigration shock using the Adao et al. (2019) method capturing the shared exposure to a small set of origin shocks, finding that it is 4 times smaller than our baseline of clustering by state and that the smaller AKM SE survives even when we additionally cluster by state.[13] Separately, we conduct a randomization-inference exercise that randomly reassigns the origin-country shifters across origins while holding each county's predicted-ancestry shares fixed, applying the same reassignment jointly across all three 4-year periods, and rebuild the IV as in the baseline specification. We then re-estimate the stacked reduced form for each of 1,000 permutation draws. The actual reduced-form t-statistic lies in the 1.9% left tail of the empirical placebo distribution (see Table B.8).

Second, following Goldsmith-Pinkham et al. (2020), we decompose the first-stage-fitted instrument by origin. Panel A of Table B.9 reports that Mexico receives a large Rotemberg weight, however the betas are negative for all of the top 15 origins and also very consistent around the weighted average. Panel B confirms that the quantitative effect is almost identical while dropping Mexico from the analysis altogether.[14] Finally, in Appendix Table B.10 we report a placebo test that confirms that contemporaneous migrations have no effect on *past inflation*.

Table 4: Heterogeneous Effects by Immigrant Characteristics

	Δ P$^{\text{Sato-Vartia,Stacked}}$ (2006-2018)			
	(1)	(2)	(3)	(4)
Panel A: Heterogeneity by Immigrant Characteristics				
$\text{Immigr}_{d,t}$ (over25)	-0.101***	-0.209***		
	(0.024)	(0.077)		
$\text{Immigr}_{d,t} \times \text{EducYears}_{d,t}$		0.110*		
		(0.063)		
$\text{Immigr}_{d,t}$ (No HS Degree)			-0.059**	
			(0.025)	
$\text{Immigr}_{d,t}$ (Some College)			0.019	
			(0.014)	
$\text{Immigr}_{d,t}$ (Low Inc. Occ.)				-0.045
				(0.028)
$\text{Immigr}_{d,t}$ (High Inc. Occ.)				0.055
				(0.053)
Obs.	6,249	6,245	6,249	6,249
First-Stage F-stat	85.5	8.2 / 7.0	59.3	19.3
Fixed Effects	State, Time	State, Time	State, Time	State, Time

Panel B: Heterogeneity in Product Basket (by Consumer Income and Predicted Immigration)				
	Low-Inc Goods	High-Inc Goods	High-Immigrant	Low Immigrant
	(1)	(2)	(3)	(4)
$\text{Immigr}_{d,t}$ (over25)	-0.127***	-0.077***	-0.133**	-0.103***
	(0.033)	(0.022)	(0.055)	(0.024)
Obs.	6,249	6,249	4,524	4,524
First-Stage F-stat	85.5	85.5	84.9	84.9
Fixed Effects	State, Time	State, Time	State, Time	State, Time

Notes: This table investigates heterogeneous effects of immigration on CPG prices across two dimensions: immigrant characteristics and product market segments. **Panel A** estimates effects by immigrant characteristics. Column (1) restricts the sample to immigrants aged 25+ since these are more relevant for educational attainment (though results are consistent throughout). Column (2) interacts immigration (over 25) with average years of education, following Terry et al. (2024)'s IV construction using predicted immigration from top-25 origin countries as instruments; for this column, the reported First-Stage F-statistics are the Sanderson-Windmeijer multivariate F-statistics for the two endogenous regressors (Immigration / Education), respectively. EducYears is demeaned annually across counties so that zero represents the average years of education among incoming immigrants. Columns (3)-(4) construct separate instruments for mutually exclusive immigrant subgroups. Column (3) distinguishes immigrants without high school degrees vs. those with some college education. Column (4) distinguishes low-income (occupation score < 26) vs. high-income immigrants. **Panel B** examines heterogeneity by product type, constructing separate Sato-Vartia price indices for subsets of products. Columns (1) and (2) estimate the effect of total immigration (age 25+) on separate price indices constructed for goods primarily purchased by low-income and high-income households, respectively. We note that the relative price difference between these two indices is statistically significant, with p-value equal to 0.019. In columns (3) and (4) we use price indices for goods primarily purchased by immigrants and those not primarily purchased by immigrants. We note that the difference between these two outcomes is not statistically significant. See Appendix A.2 for how to predict immigrants in the Household Panel. All specifications include state and period fixed effects and county controls, with immigration winsorized at 1%. They are also weighted by CPG sales. Standard errors clustered by state. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

4.2 Heterogeneous Effects by Immigrant and Consumer Characteristics

We next examine heterogeneity in these effects along two dimensions: the characteristics of the immigrants and the characteristics of the products consumed.

4.2.1 Immigrant Characteristics

The disinflationary effect is concentrated in places with lower-income, less-educated immigrants. We restrict the analysis to immigrants aged 25 and over to ensure accurate measurement of education and occupation, with results robust to this restriction (Table 4, Panel A, column 1). Column (2) examines this by interacting total immigration with the average years of education of incoming immigrants following Terry et al. (2024), finding that the disinflationary effect is attenuated among counties receiving more educated immigrants. Columns (3) and (4) pursue an alternative approach, constructing separate instruments for distinct immigrant subgroups. Column (3) distinguishes immigrants by educational attainment: less-educated immigrants generate substantial price reductions, while more-educated immigrants show no significant effect. Column (4) performs a similar decomposition by occupation-based income, again finding a more pronounced disinflationary effect among lower-income immigrants, while high-income immigrants show a positive but insignificant coefficient. These patterns reflect the fact that lower-income, less-educated immigrants are both more likely to search intensively due to budget constraints and to supply low-wage labor.

4.2.2 Product Characteristics

Panel (B) of Table 4 shows that the disinflationary effect is larger for *products* purchased by lower-income households—a group often perceived as most adversely affected by immigration through wage competition—and is pervasive across product types more broadly.

The first two columns report the immigration effect separately on low and high-income product price indices, classifying products by the average household income of their buyers in the NielsenIQ Consumer Panel.[15] Immigration lowers the price of both low and high-income goods, with a larger effect for low-income products (the difference is statistically significant with p-value = 0.019). This

[13]We apply AKM to the long-difference (2006–2018) specification, using 2005 predicted ancestry from the first stage. Following their method, we replicate the coefficient in column (5) of Table B.5 (long-diff specification), and find a SE = .00174, which is 27% of the size of the baseline with state clusters. To address the concern that AKM-origin clustering may miss within-state residual correlation that cluster-by-state captures, we also compute a two-way standard error combining the AKM-origin and state-cluster dimensions, obtaining SE = .00468 — still smaller than the cluster-by-state baseline, at 72% of its size. Finally, we also constructed a placebo IV using shares-normalized predicted ancestry times random shifters, and regressed actual non-EU immigration on the placebo IV with state fixed effects and state-clustered standard errors. Across simulations, the test rejects at the 5% level in 5.8% of draws.

[14]In Figure B.3 we also present the share of immigrants to the US across 9 major origins in our sample period. Mexico's share has decreased in the recent years, but note our IV is a multi-year weighted Bartik.

[15]Consistent with the inflation inequality literature (Jaravel, 2019; Faber and Fally, 2022; Argente and Lee, 2021), we find that inflation in our low-income price index is about 33% higher on average than the high-income index throughout the sample period.

differential effect likely reflects that low-income immigrants share similar consumption patterns with low-income native households, placing stronger pressure on prices in those market segments.

The last two columns differentiate products by predicted immigrant share, using a machine learning algorithm to predict whether a household is foreign born (see Appendix A.2). Price declines are substantial for both immigrant-intensive and non-immigrant-intensive products, slightly higher for the former, but with no statistically significant difference between the two. This suggests that immigrants and low-income natives share broadly similar consumption baskets, so that the disinflationary effect extends well beyond goods consumed exclusively by immigrants.

4.3 Effects Beyond the CPG Sector

Our analysis thus far has focused on CPG products, which constitute a major component of household expenditures but not the entirety of the consumption basket. Table 5 presents IV estimates of immigration effects on prices across several different sectors, revealing substantial heterogeneity with implications for both the mechanisms behind our CPG findings and the aggregate price effects of immigration.

Table 5: Immigration Effects on Prices Across Sectors

	$\Delta \ln P^{\text{Dur}}$	$\Delta \ln P^{\text{Non-dur}}$	$\Delta \ln P^{\text{Svc}}$	$\Delta \ln P^{\text{Rent,BLS}}$	$\Delta \ln P^{\text{Rent,ACS}}$		$\Delta \ln HP^{\text{CL}}$	$\Delta \ln HP^{\text{Z}}$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\text{Immigr}_{d,t}$	-0.017	-0.070***	-0.015	-0.015	0.075**	0.020	-0.103	-0.090
	(0.020)	(0.020)	(0.023)	(0.042)	(0.030)	(0.042)	(0.124)	(0.108)
Obs.	741	741	741	741	6,246	669	2,406	4,710
First-Stage F-stat	99.1	99.1	99.1	99.1	84.6	83.8	91.8	81.7
Fixed Effects	State, Time	State, Time	State, Time	State, Time	State, Time	State, Time	State, Time	State, Time
Method	IV	IV	IV	IV	IV	IV	IV	IV
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Sample						Cities		

Notes: Table reports IV estimates of immigration effects on prices across different sectors. Columns (1)-(4) use BLS city-level CPI data for 27 metropolitan areas (261 counties) covering durables, non-durables, services, and rent. Columns (5)-(6) use ACS median gross rent data: column (5) includes all counties, column (6) restricts to the 261 counties with BLS data. Columns (7)-(8) use housing purchase price indices: column (7) uses CoreLogic data (available 2006-2014 only), column (8) uses Zillow data (full period 2006-2018). All specifications instrument immigration with predicted shocks from equation (3.5) and include state and period fixed effects and county controls. All columns include population weights instead of sales weights, since weights on CPG products are not longer relevant. Standard errors clustered by state for ACS and housing price regressions; heteroskedasticity-robust standard errors for BLS regressions due to small number of states. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Tradable Goods and Services. Columns (1)-(2) examine BLS price indices for durables and non-durables. Non-durables show a statistically significant negative coefficient of -0.070, quantitatively similar to our baseline CPG results and providing important external validation using official BLS measures rather than scanner data. Durables show a negative but insignificant coefficient of -0.017. The weaker effect for durables is consistent with the infrequent nature of durable goods purchases: both immigrants and natives tend to search intensively when making high-stakes purchases like vehicles or appliances, so the arrival of new immigrants adds little additional search pressure. In contrast, for frequently purchased non-durables, established residents exhibit substantial inertia and brand loyalty (Dubé et al., 2010; Bronnenberg et al., 2012; Hitsch et al., 2019), reducing their search

intensity, while new immigrants—lacking established shopping routines—engage in more intensive search, amplifying competitive pressure. Column (3) examines services, showing small negative but insignificant effects. The muted effects for services reflect their local and less substitutable nature: consumers typically have limited ability to shop across providers, so the arrival of new immigrant consumers is less likely to alter pricing behavior.

Interpretation of the cross-sector effects highlights why separating demand and supply exposures is crucial. Burstein et al. (2020) and Dustmann and Glitz (2015) suggest that when supply-side adjustment can occur on the quantity margin, such as in tradable products, price effects could be dampened. Thus, the negative price effects of immigration on consumer packaged goods are more likely to reflect demand-side forces.

Housing Rent. Column (5), using ACS rental data across all counties, finds a positive and significant coefficient of 0.075, consistent with the notion that immigration increases housing demand. The BLS rental measure in column (4) and the city-restricted ACS sample in column (6) both yield smaller and insignificant coefficients of -0.015 and 0.020, respectively. Column (6) restricts to the 261 counties covered by BLS data to ensure comparability with column (4), suggesting the positive rental effect may be more pronounced outside major metropolitan areas.

Housing Purchase Prices. Columns (7)-(8) present effects on home purchase prices using CoreLogic (2006-2014) and Zillow (2006-2018) data, both showing small and statistically insignificant coefficients. This is consistent with the notion that newly arriving immigrants are more likely to rent than purchase homes, limiting their direct impact on housing purchase prices. The broader literature on immigration and housing prices yields mixed findings depending on factors such as local housing supply elasticity and native out-migration responses (Saiz, 2007; Sá, 2015; Albert and Monras, 2022; Sanchis-Guarner, 2023; Cabral and Steingress, 2026).

Aggregate (Regional) Cost-of-Living Effect. To translate these sector-specific findings into an overall assessment, we construct a weighted average using CPI expenditure weights (U.S. Bureau of Labor Statistics, 2018). Based on our baseline estimates, a 10,000 immigrant inflow reduces CPG prices by 0.58% (Table 3, column 6) and increases rental prices by 0.75% (Table 5, column 5). CPG products constitute approximately 18% of the CPI basket (Leung, 2021), while rent of primary residence accounts for approximately 7% of the CPI basket (U.S. Bureau of Labor Statistics, 2024). Weighting these effects by CPI shares and conservatively assuming null effects for all other categories yields a net aggregate disinflationary effect of approximately -0.052 ($= -0.58 \times 0.18 + 0.75 \times 0.07$) percentage points over four years. This aggregate calculation should be interpreted cautiously, as we lack precise estimates for all expenditure categories and our estimates reflect partial equilibrium effects. Nonetheless, specifications with aggregate CPI indices are broadly consistent (see Appendix).[16]

[16]Table B.12 reports results for available aggregate inflation indices: the "All Items" CPI using BLS city-level data and the state-level inflation rates provided by Hazell et al. (2022). These provide limited variation, and in most specifications the direction of aggregate prices in response to immigration is negative.

4.4 Effects on Establishments, Employment, and Wages

Table 6: Immigration Effects on Establishments, Employment, and Wages

	# Establ.	Emp.	Wages			Wages (Retail)
	(1)	(2)	(3)	(4)	(5)	(6)
$\text{Immigr}_{d,t}$	0.012**	0.746***	0.391**			0.001
	(0.005)	(0.247)	(0.147)			(0.141)
$\text{Immigr}_{d,t}$ (over25)				0.588**	-1.838*	
				(0.227)	(0.959)	
$\text{Immigr}_{d,t} \times \text{EducYears}_{d,t}$					2.828**	
					(1.156)	
Obs.	6,249	6,249	6,249	6,249	6,245	6,249
First-Stage F-stat	45.9	45.9	45.9	41.2	4.0	45.9
Fixed Effects	State, Time	State, Time	State, Time	State, Time	State, Time	State, Time
Controls	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Table reports IV estimates of immigration effects on establishments, employment, and wages. Column (1) reports the change in the number of establishments (in thousands) and column (2) the change in employment (in thousands). Columns (3)-(5) report wage effects: column (3) uses all immigrants as the endogenous variable; columns (4)-(5) restrict to immigrants aged 25 and older and additionally interact immigration with the (demeaned) average years of education among immigrants. Column (6) reports effects on retail wages. All wages are average annual real wages (in hundreds of 2012 dollars) deflated by the national PCE index. All specifications instrument immigration with predicted shocks from equation (3.5) and include state and period fixed effects and county controls. Standard errors clustered by state. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 6 examines immigration's effects on local economic outcomes beyond prices. Columns (1) and (2) show that immigration increases both the number of establishments and employment, consistent with an expansion of local economic activity (Olney, 2013; Mahajan, 2024).

Columns (3)-(6) examine wage effects. Column (3) shows a positive but specification-sensitive wage effect: immigration raises average wages in levels, though the effect becomes negligible under some specifications. Importantly, we find no evidence that wages decline in response to immigration across any specification—winsorized or not, level or log, restricting to workers aged 25 and older, or focusing on retail wages—suggesting that lower labor costs are unlikely to be the primary driver of our price results. Columns (4)-(5) examine education heterogeneity by restricting to immigrants aged 25 and older: wages increase more in counties receiving higher-educated immigrants, consistent with the large body of immigration literature and the notion that high-skilled immigration raises native wages through skill complementarities. Finally, column (6) shows that retail wages are essentially unaffected by immigration. Since retail labor costs constitute a significant share of retail pricing (Stroebel and Vavra, 2019), the absence of a retail wage effect suggests that supply-side cost reductions cannot account for the price declines we document. Appendix Table B.13 provides a comprehensive robustness analysis, showing that the positive wage effect is closer to Terry et al. (2024) as we get closer to their specification and sample, that the education heterogeneity result is robust across all specifications, and that the retail wage effect remains close to zero across most specifications.

4.5 Quantities and the Reallocation to Private Label

Our price results raise a natural question about quantities, and the answer speaks to what kind of adjustment accompanies the decline. If immigration simply expanded the local market, quantities would rise and, with decreasing returns to scale, prices would rise with them.[17] Since prices fall instead, whatever drives them down, such as more elastic consumers compressing markups or lower production costs, would ordinarily raise quantities as well. Quantities tell us whether the market is in fact bigger, or whether spending is instead reallocating toward cheaper goods, a margin our price index deliberately excludes. We measure within-barcode quantity as $\Delta \ln Q = \Delta \ln \text{Revenue} - \Delta \ln \text{Price}$ and study it in two complementary datasets. The retail-scanner census (RMS, our baseline) covers the most transactions but only larger retailers, at the county level. The Homescan panel (HMS) reaches the full retail landscape, including the discount, independent, and ethnic stores the scanner misses, but records far fewer purchases and supports only state-level aggregation. In each we separate branded (GS1) from private-label goods.[18] Across samples, and across both branded and private label, we compare price and quantity responses.

The quantity and prices responses diverge but point towards a coherent story. First, physical quantity in our baseline sample increases modestly and is not statistically different from zero (Appendix Table B.14). However, the increase is large and significant for private-label goods, suggesting a reallocation from branded products toward private label. The household panel tells a similar story with a smaller branded–private-label gap and, as expected given its smaller sample, noisier estimates.[19] Thus, it is possible that the arrival of new immigrants is associated with more consumption outside of the RMS sample (Albert and Monras, 2022), with the compositional effect clear in private label quantity responses. Second, whereas the quantity response mixes the demand shift with the compositional reallocation, *the price decline is uniform.* Prices fall in every sample we examine. Varying the data source, the product scope (branded and private label), the balance window, and the retailer restriction, the price coefficient is negative in all sixteen samples and significant in fifteen, ranging from -0.05 to -0.11 (Table B.15).[20] Both patterns are consistent with the search mechanism emphasized throughout—more price-sensitive immigrant consumers push prices down through a flatter demand curve, while also shifting spending toward cheaper alternatives.

[17] The immigration shock raises total county population (IV coefficient significant at 1%), while the native-population response is not statistically distinguishable from zero.

[18] The branded/private-label split is available only at the market level. Private-label goods carry no GS1 manufacturer prefix and therefore cannot be assigned to a producing firm, so they fall outside the firm-level analysis of Section 5; this is why we examine the quantity reallocation at the county and state level rather than the firm level.

[19] One reason the quantity effects are larger in HMS could be that it captures purchases at the smaller, untracked retailers that the scanner data do not observe. However, this is difficult to say definitively given the small sample.

[20] In this table, we extend the price results to also allow for different balanced samples (products present in at least k years in RMS), as well as several HMS samples. Price results are not an artifact of the scanner's retailer coverage.

5 Unpacking the Mechanism: Firm-Level Analysis

The county-level analysis establishes a robust reduced-form relationship between immigration and prices, but it cannot say why prices fall. In a county, the same immigrant inflow is simultaneously a shock to product demand and a shock to labor supply, while county-level estimates also conflate partial equilibrium firm responses with local general equilibrium effects such as establishment entry, consumer brand switching, and cross-county migration. Firm-level data make the separation possible and allow for sources of variation that local designs do not. Price changes are, to a first approximation, a firm-level phenomenon: examining the firm captures how a producer adjusts prices in response to the immigration its customer base experiences as a whole—a response that can extend across all of its markets, and that a county regression, which relates a county's prices only to that same county's immigration, cannot see.

Crucially, a firm's customers and its workers are located in different places, so its exposure to immigration through demand and through supply can be measured separately. We weight county immigration shocks by where a firm sells and, separately, by where it produces, yielding two firm-level exposures that isolate the channels the county estimates bundle.

5.1 Firm-Specific Immigration Shifters

We estimate the following firm-level specification, where a firm's exposure to immigration varies separately across its sales locations (demand) and production locations (supply):

$$\Delta \ln P_{ft} = \beta_1 \mathrm{I}_{ft}^{\text{demand}} + \beta_2 \mathrm{I}_{ft}^{\text{supply}} + \mathbf{X}'_{ft}\gamma + \alpha_{s(f)} + \delta_t + \varepsilon_{ft} \tag{5.1}$$

where $\Delta \ln P_{ft}$ is a change in the log price index of firm f, $\mathrm{I}_{ft}^{\text{demand}}$ is the immigration inflow in the areas where firms have sold their products, $\mathrm{I}_{ft}^{\text{supply}}$ is the immigration inflow in the areas where firms have their plants, $\mathbf{X}'_{ft}$ is the set of firm-specific control variables, and we control for industry (3-digit SIC) ($\alpha_{s(f)}$) and period (δ_t) fixed effects. Firm-specific immigration shifters are measured as follows:

$$\mathrm{I}_{ft}^{\text{demand}} = \sum_d \omega_{fd}^{\text{sales}} \, \mathrm{I}_{dt} \tag{5.2}$$

$$\mathrm{I}_{ft}^{\text{supply}} = \sum_d \omega_{fd}^{\text{empl}} \, \mathrm{I}_{dt} \tag{5.3}$$

where $\omega_{fd}^{\text{sales}}$ is the (initial) firm f's own sales share (of its total sales) in a market defined at the county level and $\omega_{fd}^{\text{empl}}$ is the firm f's own employment share (of its total employment) in a county market. Thus, each firm exposure is differentiated by where their sales and employment is situated. Note that I^{demand} and I^{supply} vary at the firm level. While it is possible to further differentiate shocks by product group using a crosswalk between 4-digit SIC industries and NielsenIQ product groups, we find that this introduces little additional variation given that most firms sell within a

single product group, and results are nearly identical.[21]

Three features of these exposures shape identification. First, identification comes from comparing firms in the same industry and period whose customer bases sit in differently-shocked counties. Second, sales footprints are dispersed while production footprints are concentrated—the typical firm sells in about twenty counties but produces in one—so the demand and supply exposures are far from collinear (their correlation across firms is roughly 0.4) and are separately identified. Third, both exposures inherit the exogeneity of the county design: we instrument them with the same predicted-ancestry shocks of Section 3.4, with sales weights, and the resulting first stage is strong.

5.2 Demand vs. Supply Channels

Table 7: Firm-Level Prices: Separate Demand and Supply Side Immigration Exposure

	Δ lnP^{SV,Stack} (2006-2018)					
	(1)	(2)	(3)	(4)	(5)	(6)
Immigration$_{f,d,t}$ (D)	-0.071***	-0.057***		-0.057***	-0.062***	-0.051***
	(0.020)	(0.022)		(0.020)	(0.020)	(0.019)
Immigration$_{f,d,t}$ (S)	0.016***		0.007	0.018***	0.019***	0.016***
	(0.005)		(0.005)	(0.004)	(0.005)	(0.005)
Δ lnW (Retail, D)					0.031	0.048
					(0.221)	(0.220)
Δ lnP (Rent, D)					0.224	0.181
					(0.185)	(0.177)
Δ lnW (Retail, S)					0.094	0.093
					(0.065)	(0.065)
Δ lnP (Rent, S)					-0.045	-0.034
					(0.065)	(0.067)
Obs.	11,037	11,037	11,037	11,037	11,037	11,037
First-Stage F-stat	3,743.4	4,896.3	4,908.1	3,733.5	3,977.9	
Fixed Effects	Industry, Time	Industry, Time	Industry, Time	Industry, Time	Industry, Time	Industry, Time
Method	IV	IV	IV	IV	IV	OLS
Controls				Yes	Yes	Yes

Notes: This table reports the effect of immigration on firm-level prices, separately identifying demand-side and supply-side immigration exposure as detailed in the main text. The price index refers to the Sato-Vartia index. All specifications use stacked 4-year differences (2006-2018) and include industry (SIC 3-digit) and time-period fixed effects, as well as controls for the sum of sales shares in firms' retail markets and the sum of employment shares in firms' producing establishments, to account for possible biases from missing county information. Columns (1)-(5) report IV estimates; column (6) reports OLS estimates. Column (4) additionally includes interactions of year dummies with the firm-level exposure to counties' urban indicator and GDP per capita. Columns (5) and (6) further add retail wages and rents at both the demand and supply side as controls. Standard errors clustered by SIC 3-digit industry. ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

Table 7 reports the effect of immigration on firm-level price changes, separately identifying demand and supply exposure. The price index is the Sato-Vartia index, capturing changes in prices of continuing products. All regressions include controls for the sum of sales shares in firms' retail markets and the sum of employment shares in firms' producing establishments, to account for possible biases from firms with incomplete observations in their sales or production counties. Columns (4)-(6)

[21] We construct a crosswalk between NETS 4-digit SIC industries and NielsenIQ product groups based on our NielsenIQ-NETS firm-level matching, and verify that results are robust to allowing shocks to vary at the firm-and-group level with group fixed effects.

additionally include interactions of year dummies with the firm-level demand and supply exposure to counties' urban indicator and GDP per capita.

Column (1) reveals the key finding: demand-side exposure significantly reduces prices, consistent with the county-level results, while supply-side exposure leads to a small but significant price increase.[22] Columns (2) and (3) show that omitting either channel biases the remaining coefficient: omitting supply exposure attenuates the demand-side effect, as the positive supply pressure partially offsets the negative demand effect, while omitting demand exposure attenuates the supply-side effect for the same reason. Column (4) adds firm-level controls and results remain stable. Column (5) further adds retail wages and rents as controls, to address concerns that supply-side cost factors in retail markets may confound the estimates.[23] These could be deemed "bad controls", but their inclusion as a robustness check confirms our results are not sensitive this. Column (6) shows that OLS estimates are attenuated toward zero relative to IV, consistent with attenuation bias from measurement error in firm-level exposure, but remain qualitatively similar.

These results confirm that the disinflationary effect of immigration operates through the demand channel, consistent with immigrants' greater price sensitivity inducing firms to lower markups. The positive supply-side effect is consistent with modest upward cost pressure when immigrants arrive in production locations, though the supply-side channel appears less important quantitatively than the demand channel. To gauge economic magnitude on a common scale, Appendix Table B.16 standardizes both exposures: a one-standard-deviation increase in demand exposure lowers firm prices by 2.7 percentage points, versus 0.9 for supply. Thus, the demand effect is three times the supply effect. The demand-side mechanism is examined more directly in the next section by asking how immigration affects consumer search behavior and demand elasticities at the firm level.

We again provide robustness tables in the Appendix. Because the firm-level demand and supply exposures I_{ft}^{demand} and I_{ft}^{supply} are shift-share aggregates of county immigration shocks (equations (5.2)–(5.3)), we report exposure-robust standard errors following Borusyak et al. (2022). We find that they are smaller than our SIC-3 clustered errors in all specifications (Appendix Table B.17), so industry clustering is the conservative choice for inference. The demand-side result is robust to how the firm price index is aggregated and holds at the individual barcode level. Aggregating UPC prices to the firm through alternative chain-weighted indices reproduces the baseline coefficient (Appendix Table B.18), and regressions estimated directly at the barcode level deliver comparable coefficients (Appendix Table B.19). Because these are price changes for identical barcodes, the decline cannot reflect shifts in product composition or quality, leaving markup compression on continuing goods as the operative channel.

Separately, Table B.20 confirms results hold with 12-year long-differences, and Table B.21 reports that the main coefficients are virtually unchanged when controlling for exporter and importer

[22] That the two exposures are positively correlated across firms yet produce opposite-signed effects is informative. The opposing signs point to two economically distinct channels rather than a common spurious factor.

[23] These are based on firm level exposure: wages and rents in the counties where firms sell and produce.

status, when excluding internationally active firms entirely, and when allowing the effects to differ by international trade status—with no significant difference between international and domestic firms.

Controlling for Retailers A natural concern is that the price response reflects retailer behavior rather than the manufacturer's: immigrants may concentrate at particular chains, large manufacturers may sell disproportionately through large retailers, or a retailer-level shock correlated with immigration may move shelf prices. We confirm that the effect reflects the manufacturer's exposure by bringing the retailer dimension into the specification directly.

To confirm that the effect reflects the manufacturer's exposure, we re-estimate the price regression at the firm-retailer level and reconstruct the demand shock in two alternative dimensions: a retailer's own counties (g_r) and the counties where a given firm sells through a given retailer (g_{fr}), each entered alongside the firm shock (g_f) and instrumented by its shift-share counterpart (Appendix Table B.22). The firm coefficient is stable throughout, while the retailer shock is an order of magnitude smaller and the firm-by-retailer interaction is at most marginal. Furthermore, absorbing retailer-by-period fixed effects, which eliminate any factor common to a chain in a period, leaves the firm effect essentially unchanged. Prices respond to the firm's immigration exposure, not the retailer's.[24]

5.3 Evidence for the Search Mechanism

Table 8 provides direct evidence that immigration shifts the composition of a firm's consumer base toward more search-intensive, price-sensitive shoppers. To construct firm-level measures of consumer search behavior, we link each firm to households that purchase from it using the NielsenIQ Consumer Panel, and measure each household's shopping behavior—excluding purchases from the focal firm to avoid mechanical correlation—then aggregate to the firm level weighted by purchase intensity.

Columns (1) and (2) show that firms with greater demand-side immigration exposure are increasingly linked to households that make more shopping trips and spend more days searching, consistent with a more search-intensive consumer base. Column (3) shows a positive but insignificant increase in the number of stores visited, consistent with cross-retailer search behavior. Column (5) shows an increase in the store brand share, reflecting a shift toward lower-priced alternatives among these consumers. Notably, the number of products purchased (column 4) does not increase, suggesting that the effect is driven by search intensity rather than variety-seeking behavior.

The most direct evidence comes from column (6), which examines the relative prices paid by households linked to each firm for purchases made elsewhere. Following Aguiar and Hurst (2007), we measure relative price as the price paid for a given barcode divided by the national average for that barcode. A negative coefficient indicates that firms in high-immigration counties are increasingly

[24] A concern we cannot completely account for is that the firm demand shock is actually a supply shock to the retailer (e.g. lower retail wages). Table 7 controls for the exposure to retail wages and rents at the firm level, while Table B.22 controls for firm-retail immigration exposure (g_{fr}). We have also checked that the firm demand coefficient is unchanged when the retail wage and rent exposure controls are measured at the *firm-retailer level* rather than the firm level. Immigration-driven changes in local retail costs do not account for the effect.

Table 8: Firm-Level Search Intensity

	Firm purchaser characteristics (stacked Δ, 2006-2018)					
	(1) N. Trips	(2) N. Days	(3) N. Stores	(4) N. Products	(5) % Store Brand	(6) Rel. Price
$\text{Immigration}_{f,d,t}$ (D)	0.012**	0.010**	0.008	-0.002	0.008*	-0.002***
	(0.005)	(0.004)	(0.005)	(0.005)	(0.004)	(0.001)
$\text{Immigration}_{f,d,t}$ (S)	-0.000	0.001	0.000	0.001	0.002	0.000
	(0.002)	(0.002)	(0.002)	(0.003)	(0.002)	(0.000)
Obs.	10,325	10,325	10,325	10,325	10,325	10,317
First-Stage F-stat	3,736.5	3,736.5	3,736.5	3,736.5	3,736.5	3,735.9
Fixed Effects	Industry, Time	Industry, Time	Industry, Time	Industry, Time	Industry, Time	Industry, Time
Method	IV	IV	IV	IV	IV	IV

Notes: Each observation is a firm × period, on stacked four-year differences (2006–2010, 2010–2014, 2014–2018); the outcomes measure the shopping behavior of the firm's customer base. For each firm we identify the NielsenIQ Consumer Panel households that purchase from it, measure each household's behavior on its purchases *excluding the focal firm's products*, and aggregate to the firm level weighted by purchase intensity. The columns are: (1) the number of shopping trips; (2) the number of distinct shopping days; (3) the number of distinct retail chains visited; (4) the number of distinct products purchased; (5) the store-brand (private-label) expenditure share; and (6) the relative price paid, defined following Aguiar and Hurst (2007) as the price paid for a barcode divided by its national average price. The regressors are the firm's demand-side (D) and supply-side (S) immigration exposure, entered jointly and instrumented by their shift-share predicted shocks as in Table 7. All specifications include industry (3-digit SIC) and period fixed effects and control for the sales- and employment-share of counties with missing shocks on the demand and supply sides. Observations are weighted by initial firm sales, and standard errors are clustered by 3-digit SIC industry. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

linked to households that pay below-average prices for identical products elsewhere—direct evidence that immigration shifts firms' consumer bases toward more price-sensitive shoppers who search more effectively. Supply-side exposure shows no significant effects across any outcome, further supporting the demand-side interpretation.

5.4 Structural Estimates: Elasticities, Varieties, and Appeal

Building on the search evidence above, this section formalizes the link from search to prices through a structural demand system. We estimate a nested CES demand system following Hottman et al. (2016), which decomposes the price index into a Sato-Vartia index for continuing varieties and a Feenstra (1994) variety adjustment for product entry and exit, and recovers firm-level appeal as a measure of quality. Unlike Hottman et al. (2016), who estimate national-level elasticities, we estimate elasticities separately by state, product group, and time period, and relate them to immigration inflows to test whether immigration raises demand elasticities, expands variety, or shifts quality. We sketch the main components here and refer the reader to Appendix C for full details.

Demand System. Consider a nested demand system in each location, where the upper nest is across product groups, the middle nest is across firms within each product group, and the lower nest is across products within each firm. Household utility is a Cobb-Douglas aggregate of real

consumption across product groups:

$$\ln U_{ct} = \int_{g \in \Omega_{ct}} \phi_{gct}^{C} \ln C_{gct} \, dg \tag{5.4}$$

where g is a product group, c is a county, t is time, and ϕ_{gct}^{C} is the Cobb-Douglas expenditure share on product group g. Within each product group, households have nested CES preferences across firms and products:

$$C_{gct} = \left[\sum_{f \in \Omega_{gct}} \left(\phi_{fgct}^{F} C_{fgct} \right)^{\frac{\sigma_\ell^F - 1}{\sigma_\ell^F}} \right]^{\frac{\sigma_\ell^F}{\sigma_\ell^F - 1}}, \quad C_{fgct} = \left[\sum_{u \in \Omega_{fgct}} \left(\phi_{ut}^{U} C_{ut} \right)^{\frac{\sigma_\ell^U - 1}{\sigma_\ell^U}} \right]^{\frac{\sigma_\ell^U}{\sigma_\ell^U - 1}} \tag{5.5}$$

where ϕ_{fgct}^{F} is the appeal of firm f, σ_ℓ^F is the elasticity of substitution across firms, ϕ_{ut}^{U} is the appeal of product u, and σ_ℓ^U is the elasticity of substitution across products within a firm. The subscript ℓ indexes a state, product group, and time period, so elasticities vary along all three.

Estimation. The key parameters to estimate are the demand elasticities σ^F and σ^U. We follow the approach in Feenstra (1994) and Hottman et al. (2016), which exploits the CES demand structure to identify elasticities from the covariance between price and market share changes. We estimate elasticities separately by state, product group, and time period, yielding σ_ℓ^F and σ_ℓ^U where ℓ indexes a state, product group, and period.[25] Identification requires a supply side as well as a demand side, and we impose the one in Online Appendix C.1. Firms set each price as a markup over marginal cost; marginal cost rises with the quantity produced at a rate δ_ℓ that we estimate jointly with the elasticities, subject to the nonnegativity we impose; and the markup is common across the products of a given firm, which is what lets double-differencing within a firm remove it. This pricing structure is consistent with monopolistic competition, Bertrand competition with differentiated products, and Cournot competition, and it is the supply side we carry into the model of Section 6.

Variety-adjusted Price Index and Firm Appeal. We construct three objects. The first, the Sato-Vartia price index, requires only observed prices and expenditure shares; the other two use the estimated elasticities and are not recoverable from reduced-form analysis.

First, the *Sato-Vartia price index* aggregates price changes of continuing varieties using expenditure-share weights:

$$\Phi_{fgct}^{SV} = \prod_{u \in \Omega_{fgct,t-1}} \left(\frac{P_{uct}}{P_{uc,t-1}} \right)^{w_{uct}} \tag{5.6}$$

where $\Omega_{fgct,t-1}$ is the set of products available in both t and $t-1$, and w_{uct} are Sato-Vartia weights based on expenditure shares. Abstracting from appeal and product entry and exit, aggregating

[25] Periods correspond to 2006–2010, 2011–2014, and 2015–2018.

barcode prices across products within each firm-group-county, then across product groups, and finally across counties (Sato-Vartia weighted) yields the firm-level price index used in Section 5.2; aggregating the same firm-group-county objects instead across firms within a group and then across groups (Cobb-Douglas) yields the county-level price index of Section 4.

Second, the *variety adjustment* captures the contribution of product entry and exit to the cost of living:

$$\Phi_{fgct}^{Variety} = \left(\frac{\lambda_{fgct}}{\lambda_{fgc,t-1}} \right)^{\frac{1}{\sigma_\ell^U - 1}} \tag{5.7}$$

where λ_{fgct} is the expenditure share on continuing products relative to all products. When entering products capture a larger share of period-t sales than exiting products held in period $t-1$, the ratio $\lambda_{fgct}/\lambda_{fgc,t-1}$ falls below one. Since $\sigma_l^U > 1$, this implies $\Phi_{fgct}^{Variety} < 1$, lowering the effective price index. The intuition is that consumers benefit from the availability of new varieties, even if prices of existing products are unchanged.

Combining these components, the firm-level utility-based price index decomposes as:

$$\frac{P_{fgct}}{P_{fgc,t-1}} = \Phi_{fgct}^{SV} \cdot \Phi_{fgct}^{Variety} \tag{5.8}$$

Third, firm-level *appeal* is recovered as a structural residual. Our reported price index is constructed holding quality fixed and appeal does not enter it directly. In the full CES cost-of-living dual, it appears as a utility shifter, so we analyze it separately as a quality-composition channel. From the CES demand structure, the market share of firm f depends on its price relative to competitors and its appeal. Rearranging, we obtain:

$$\phi_{fgct} = \frac{P_{fgct}}{P_{gct}} \left(S_{fgct} \right)^{\frac{1}{\sigma_\ell^F - 1}} \tag{5.9}$$

The key intuition is that appeal captures the component of market share that cannot be explained by relative prices, given the demand elasticity. If a firm has a higher market share than its price would predict, we attribute the residual to higher appeal. The growth rate of appeal is then:

$$\Phi_{fgct}^{Appeal} \equiv \frac{\phi_{fgct}}{\phi_{fgc,t-1}} = \frac{P_{fgct}/P_{fgc,t-1}}{P_{gct}/P_{gc,t-1}} \left(\frac{S_{fgct}}{S_{fgc,t-1}} \right)^{\frac{1}{\sigma_\ell^F - 1}} \tag{5.10}$$

We aggregate these firm-group-county-period level objects to the firm-period level using chain-weighted Sato-Vartia aggregation, first across product groups and then across counties, mirroring the price index construction in Section 3.3.

Results: Demand Elasticities. We first examine whether immigration is associated with higher demand elasticities, as implied by the search mechanism. Table 9 reports results from regressions of estimated elasticities on immigration exposure. We consider two specifications: (i) county-group-

Table 9: Demand Elasticities

	County-Level		State-Level	
	σ^F	σ^U	σ^F	σ^U
Immigration (100K)	0.040**	0.063**	0.019***	0.025***
	(0.016)	(0.025)	(0.006)	(0.007)
Obs.	465,736	465,736	11,033	11,033
Fixed Effects	Cty-Prdct, Yr-Prdct	Cty-Prdct, Yr-Prdct	St-Prdct, Yr-Prdct	St-Prdct, Yr-Prdct
Median Outcome	4.74	14.33	4.69	14.28

Notes: This table reports the effect of immigration on demand elasticities. Demand elasticities are estimated from barcode data at the state-product group-period level, with periods corresponding to 2006–2010, 2011–2014, and 2015–2018. There are 100 product groups, explaining the rise in observations relative to county-year results. Columns (1)-(2) use county-level aggregation, where counties within the same state share the same elasticity estimate; columns (3)-(4) use state-level aggregation, where immigration is aggregated across counties within a state and the instrument is reconstructed following the same procedure described in Section 3.4. Odd-numbered columns report results for the elasticity across firms (σ^F) as the outcome, while even-numbered column report results for the elasticity across products (σ^U). The first two columns include county-product and year-product fixed effects; the latter two columns include state-product and year-product fixed effects. Immigration is measured in 100,000s. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

period variation, where counties within the same state share the same elasticity estimate, and (ii) state-group-period variation, where immigration is aggregated to the state level. In the latter case, the instrument for state-wide immigration must be reconstructed; we do so following the same procedure described in Section 3.4 but using state-level immigration and ancestry data.

Both specifications show that immigration is associated with significantly higher demand elasticities, for both the elasticity across firms (σ^F) and across products within firms (σ^U). This is consistent with the search mechanism: markets with larger immigrant shares exhibit higher effective demand elasticities, which compresses firm markups.

Results: Price and Appeal Table 10 decomposes firm-level price responses into the utility-based price index and its components. Both the variety correction and appeal term are constructed using time- and state-varying demand elasticities, following the structural decomposition described above. *Price+Variety* is the utility-based (CES exact) price index combining the Sato-Vartia index with the Feenstra variety correction, *SV-Price* isolates the price index for continuing varieties, *Variety Corr.* is the Feenstra correction term, and *Appeal* reports changes in average firm-level appeal. We report the negative of the Feenstra factor so that a positive number intuitively indicates an increase in variety.

The demand-side exposure reveals two distinct effects. First, prices of continuing varieties fall (Column 2), indicating genuine disinflation—firms charge less for identical products, consistent with markup compression from more price-sensitive immigrant consumers. The variety correction is minimal (Column 3), minimally attenuating the disinflation in the utility-based index (Column 1), suggesting that product entry and exit do not significantly contribute to price changes in response to immigration. Second, the appeal index declines (Column 4), indicating that the composition of sales shifts toward lower-quality products. This composition effect operates *separately* from the Sato-Vartia price index—it captures changes in what consumers buy, not what they pay for identical

Table 10: Firm-Level Prices: Decomposition of the CES Price Index and Appeal

	Δ $\ln P^{Stack}$ (2006-2018), Decomposition			Quality
	(1) Price+Variety	(2) SV-Price	(3) Variety Corr.	(4) Appeal
$Immigration_{f,d,t}$ (D)	-0.051**	-0.057***	-0.005	-0.119***
	(0.020)	(0.020)	(0.005)	(0.019)
$Immigration_{f,d,t}$ (S)	0.014***	0.018***	0.004***	0.025***
	(0.004)	(0.004)	(0.001)	(0.009)
Obs.	10,333	11,037	10,333	10,333
First-Stage F-stat	3,779.3	3,733.5	3,779.3	3,779.3
Fixed Effects	Industry, Time	Industry, Time	Industry, Time	Industry, Time
Method	IV	IV	IV	IV
Controls	Yes	Yes	Yes	Yes

Notes: This table reports the effect of immigration shocks on the utility-based price index and its components, using the baseline IV specification with firm-level immigration exposure differentiated by demand- and supply-side as detailed in the main text. Column (1) reports the effect on the utility-based (CES exact) price index, which combines the Sato-Vartia price index for continuing varieties with the Feenstra variety correction. Column (2) isolates the Sato-Vartia price index, replicating Table 7 column (4). Column (3) reports the Feenstra variety correction (equation 5.7). To make interpretation intuitive, Variety Corr. (Col 3) is multiplied by -1 so that positive values reflect variety gains (not the case for column (1)). Column (4) reports the effect on firm-level appeal (equation 5.10), which captures shifts in quality composition. Demand elasticities used to construct the variety and appeal corrections vary over time and states as described above. All specifications use stacked 4-year differences (2006-2018) and include controls for the sum of sales shares in firms' retail markets and employment shares in producing establishments, interactions of year fixed effects with firm-average exposure to county GDP and unemployment rates, firm exposure to retail wages and rental rates, and industry (SIC 3-digit) and time-period fixed effects. Standard errors are clustered by SIC 3-digit industries. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

products. Broadly speaking, the demand-side effects dominate the supply-side ones in magnitude, consistent with the county-level reduced-form results.

The supply-side exposure also exhibits noteworthy findings. Prices of continuing products rise modestly (Column 2), and the variety correction (Column 3), while small in absolute terms, does increase and is relatively more important on the supply side than the demand side. The appeal index also increases (Column 4). The joint pattern of higher prices, greater variety, and higher appeal on the supply side is most consistent with a firm-level upgrading/innovation mechanism (Kerr and Lincoln, 2010; Terry et al., 2024): immigration to production locations enables firms to introduce higher-quality products while also raising overhead costs—such as research and development—that spread across existing varieties and raise their prices.

Table 11 extends the decomposition by examining whether these effects vary with the characteristics of each firm's consumer base, paralleling the county-level product heterogeneity in Table 4 Panel B. Using the NielsenIQ Consumer Panel, we compute the sales-weighted average household income and predicted immigrant status of each firm's buyers, residualizing within product group and year to remove compositional differences across product categories. Firms are classified as low-income-exposed if their average consumer income falls below the sample median, and as immigrant-intensive if a majority of their consumer base is predicted to be foreign-born using the same

Table 11: Firm-Level Prices: Decomposition of the CES Price Index and Appeal

	(1) Low-Inc Expos.	(2) High-Inc Expos.	(3) High-Immig Expos.	(4) Low-Immig Expos.
Panel A: SV-Price				
$\text{Immigration}_{f,d,t}$ (D)	-0.093***	-0.041*	-0.088***	-0.048**
	(0.033)	(0.023)	(0.033)	(0.019)
$\text{Immigration}_{f,d,t}$ (S)	0.014**	0.021**	-0.020	0.017***
	(0.006)	(0.010)	(0.014)	(0.005)
Obs.	4.569	4.674	708	7.890
Panel B: Appeal				
$\text{Immigration}_{f,d,t}$ (D)	-0.178***	-0.118***	-0.173***	-0.120***
	(0.039)	(0.038)	(0.055)	(0.031)
$\text{Immigration}_{f,d,t}$ (S)	0.018	0.038***	-0.007	0.022**
	(0.014)	(0.012)	(0.033)	(0.010)
Panel C: Variety Correction				
$\text{Immigration}_{f,d,t}$ (D)	-0.013	0.000	-0.022	-0.009*
	(0.008)	(0.008)	(0.015)	(0.005)
$\text{Immigration}_{f,d,t}$ (S)	0.004	0.007**	0.006	0.005***
	(0.002)	(0.003)	(0.006)	(0.002)

Notes: This table reports IV estimates of immigration effects on the SV-Price index (Panel A), firm appeal (Panel B), and the (negative) Feenstra variety correction (Panel C), splitting the sample by firm exposure type. Columns (1)–(2) split by whether the firm's consumers are predominantly low-income; columns (3)–(4) split by the (constructed) immigrant intensity of the firm's consumers. Observations are reported from Panel A; Panels B and C have slightly fewer observations due to elasticity availability requirements. All specifications use the baseline IV methodology with stacked 4-year differences (2006–2018), weighted by initial sales, with full controls and industry (SIC 3-digit) and period fixed effects. Standard errors clustered by SIC 4-digit industry. ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

machine learning algorithm described in Appendix A.2. The demand-side price effect on identical products is larger among firms selling to low-income consumers (column 1) compared to high-income consumers (column 2), and similarly stronger among firms with immigrant-intensive consumer bases (column 3) relative to those without (column 4).[26] The appeal effects follow a parallel pattern, with larger quality-composition declines for firms exposed to low-income and immigrant-intensive buyers. On the supply side, the positive price and appeal effects from Table 10 are concentrated among firms selling to *high*-income consumers, consistent with cost pressure from immigrant labor being passed through primarily by firms serving less price-sensitive markets. The asymmetry is informative about the mechanism: demand-side disinflation dominates in market segments most exposed to price-sensitive immigrant consumers, while supply-side cost pressures manifest for firms with customers that have a higher willingness to pay for quality.

6 Theoretical Framework

Section 5.4 documents that immigration lowers the price of continuing products, measured by the Sato-Vartia index Φ^{SV}, and lowers firm-level appeal, Φ^{Appeal}, and that both effects operate through

[26] In a specification where we interact the low-income/high-immigrant indicators with immigration, we find these differential effects are significant at the 10% level.

the demand channel. This section develops a model in which a single primitive, immigrants' greater price sensitivity, accounts for both facts. Given this heterogeneity, a larger immigrant population compresses markups on continuing products and reallocates the firm's sales toward its cheaper and lower-quality products. The two margins work through two different elasticities, substitution across firms in the first case and within the firm in the second, and these are the two tiers of the demand system we estimate in Section 5.4. We keep the model as parsimonious as this system allows, and all derivations are in Online Appendix D.

6.1 Environment

Consumers. Preferences are nested CES, as in (5.4) and (5.5) of Section 5.4, with one such system for each of two consumer types. There are natives (n) and immigrants (i), with populations L_n and L_i. Upper-tier preferences are Cobb-Douglas between an outside good and the differentiated bundle, so a type-h consumer spends a constant fraction α_h of income on differentiated products, an amount we write E_h. Consumers substitute across firms with elasticity σ_h^F and across products within a firm with elasticity σ_h^U, where $\sigma_h^U > \sigma_h^F > 1$. Firm f has appeal ϕ_f and its product u has quality ϕ_u and price p_u, with ϕ_f and ϕ_u common to the two types, so the types agree on quality and differ only in how they substitute across prices.[27]

We assume immigrants are more price-sensitive at both margins, $\sigma_i^U > \sigma_n^U$ and $\sigma_i^F > \sigma_n^F$. Our structural estimates support this, since both elasticities increase with immigration inflows (Table 9). One reason to expect the gap is income. Lower-income consumers are more price-elastic (Handbury, 2021; Faber and Fally, 2022; Auer et al., 2024), and the variation our instrument isolates is non-European by construction (Section 3.4), largely drawn from origins with lower income per capita than the United States. Search is the natural channel, since comparing prices costs such consumers less, so they shop around more and switch to a cheaper seller more readily. The reduced-form results support this channel. The disinflationary effect is concentrated among lower-income and less-educated immigrants and among low-income products (Table 4), and firms with greater demand-side exposure are increasingly linked to households that shop more and that pay below-average prices for the same barcodes elsewhere (Table 8). We take the gap as a primitive rather than deriving it. Preferences are homothetic within a type, so income fixes what a consumer spends and no elasticity depends on it. Mechanisms unrelated to income, such as weaker attachment to established brands, would rationalize the gap equally well, and nothing below depends on which one operates.

Note that each type's elasticities are fixed in the model. The estimation of Section 5.4 fits a single consumer for each product group, state and period, so what it recovers is an average of the two types' elasticities, weighted by their shares of the market. That average rises with immigration

[27] The appeal parameters here play the role of ϕ^F and ϕ^U of Section 5.4 but are not the same objects, since the relation between them involves the elasticity and so is not type-free. Online Appendix D sets out how the model relates to the estimated system.

even though neither type's price sensitivity has changed, which is what Table 9 reports.[28]

Firms. Each firm f produces a fixed set of products Ω_f with heterogeneous qualities. Product u has marginal cost $MC_u = a_u q_u^{\delta}$, where $\delta \geq 0$ governs returns to scale, with $\delta > 0$ decreasing returns and $\delta = 0$ constant marginal cost. Within a firm, the cost shifter is a function of quality, $a_u = a(\phi_u)$ with elasticity $\varepsilon_a \equiv d \ln a / d \ln \phi_u \geq 0$, so that higher quality carries a weakly higher unit cost. Each firm internalizes substitution among its own products but is small relative to the market, and it takes the aggregate indices P_h and aggregate expenditure as given, as in the standard multiproduct monopolistic competition setting.

We take product variety and quality as fixed. Two features of our setting motivate this assumption on the demand side. First, we find no variety response on that channel, since the Feenstra variety correction does not move with demand-side immigration exposure (Table 10). Second, a barcode is a fixed physical product, so its quality does not change. Everything that moves in a firm's average quality is therefore the composition of what it sells. Online Appendix D.5 integrates variety explicitly: markup compression and demand expansion push a product's operating profit in opposite directions, and the model does not sign the net effect, consistent with the muted variety response we find. Because quality is fixed, $a(\cdot)$ is not a technology for producing quality but a description of a firm's existing line: among the products it already sells, the better ones cost more per unit.

Demand. Type-h demand for product u of firm f is

$$c_{uh} = \phi_u \Big(\frac{p_u}{P_{fh}} \Big)^{-\sigma_h^U} \phi_f \Big(\frac{P_{fh}}{P_h} \Big)^{-\sigma_h^F} \frac{E_h}{P_h}, \qquad P_{fh} = \Big(\sum_{v \in \Omega_f} \phi_v p_v^{1-\sigma_h^U} \Big)^{1/(1-\sigma_h^U)}, \tag{6.1}$$

with total demand $q_u = \sum_h L_h c_{uh}$, and P_h the group-level CES price index aggregating the firm indices P_{fh} with weights ϕ_f, as in (5.4). Write $s_{uh|f}$ for the within-firm expenditure share of a type-h consumer, the share of what it spends at firm f that goes to product u. These shares sum to one over Ω_f:

$$s_{uh|f} = \frac{\phi_u \, p_u^{1-\sigma_h^U}}{\sum_{v \in \Omega_f} \phi_v \, p_v^{1-\sigma_h^U}}. \tag{6.2}$$

Both types face the same firms, the same products and the same prices, and they agree on quality, and each firm takes the group index P_h as given. They differ only in the elasticities with which they substitute, so each carries its own price index P_{fh} and its own within-firm shares. That is what lets $s_{un|f}$ and $s_{ui|f}$ differ over an identical product line at identical prices, and it is the whole of the composition mechanism below. No result rests on a difference in the taste for quality.

[28] A single-agent CES fitted to a two-type economy recovers a weighted average of the two type-specific elasticities, which is why the estimate should rise with the immigrant share of the market. Setting $\sigma_i^U = \sigma_n^U$ and $\sigma_i^F = \sigma_n^F$ collapses the model back to the estimated system exactly, as Online Appendix D verifies.

Equilibrium. Firms set prices to maximize profits. A firm owns the products that are closest substitutes for one another, so raising one price diverts sales to its own shelf rather than to a competitor. The firm keeps those sales, and within-firm substitution therefore puts no pressure on its markup; what does is substitution across firms. Online Appendix D.2 shows that the price of product u is then a markup over marginal cost governed by the across-firm elasticity:

$$p_u = \frac{\sigma^F_{f,\mathit{eff}}}{\sigma^F_{f,\mathit{eff}} - 1} \cdot MC_u, \qquad MC_u = a_u q_u^{\delta}, \tag{6.3}$$

where

$$\sigma^F_{f,\mathit{eff}} \equiv \sum_h \theta_{fh}\, \sigma^F_h, \qquad \theta_{fh} \equiv \frac{L_h \sum_{u \in \Omega_f} p_u c_{uh}}{\sum_{u \in \Omega_f} p_u q_u}. \tag{6.4}$$

The markup therefore depends on how price-sensitive the firm's customer base is when choosing among firms. Immigration raises θ_{fi}, and so raises $\sigma^F_{f,\mathit{eff}}$ and compresses the markup. Equation (6.3) sets a single markup across the firm's products rather than one tailored to the buyer mix of each barcode, which the firm does not observe. Online Appendix D.2 states the exact pricing rule this specializes, solves the firm's problem without restricting markups across products, and shows that the restriction does not change the direction of the response: in the unrestricted solution immigration lowers the markup on every product, as it does the common one. What the restriction suppresses is dispersion across products, and that dispersion is bounded by the across-firm elasticity gap. The single markup also leaves intact the reallocation of sales toward the products whose buyers are more immigrant, which enters through the sales weights in $\sigma^F_{f,\mathit{eff}}$.

6.2 Price Decomposition

We measure firm-level prices with the Sato-Vartia index, which aggregates product-level price changes with fixed expenditure weights. Taking logs of (6.3) and differentiating with respect to the immigrant population gives the following decomposition.

Proposition 1 (Direct Price Forces)**.** *Immigration exerts two direct forces on the firm's price:*

$$\left.\frac{d \ln P_f^{SV}}{dL_i}\right|_{\mathit{direct}} = \sum_{u \in \Omega_f} \bar{s}_u \left[\underbrace{-\frac{1}{\sigma^F_{f,\mathit{eff}}(\sigma^F_{f,\mathit{eff}} - 1)} \cdot \frac{\partial \sigma^F_{f,\mathit{eff}}}{\partial L_i}}_{\textit{Markup Effect } (<0)} + \underbrace{\delta \cdot \frac{\theta_{ui}}{L_i}}_{\textit{Scale Effect } (\geq 0)} \right] \tag{6.5}$$

where $\bar{s}_u$ are Sato-Vartia weights and

$$\frac{\partial \sigma^F_{f,\mathit{eff}}}{\partial L_i} > 0. \tag{6.6}$$

Note that we report the direct derivative, how immigration affects the firm's markup and marginal cost, holding the price fixed. In equilibrium the price moves too, creating a feedback effect

since the markup and quantity both depend on it. Online Appendix D.3 decomposes the total derivative exactly and, when the two types share the within-firm elasticity, solves it in closed form.

The markup effect comes from immigration raising the effective across-firm elasticity, and its size depends on the gap $\sigma_i^F - \sigma_n^F$. The scale effect goes the other way: immigration raises quantity demanded, which raises marginal cost whenever $\delta > 0$. Which force dominates is a restriction on parameters and a quantitative question. Ultimately, prices fall when the two types differ enough in price sensitivity relative to how steeply marginal cost rises.[29] Our finding that Φ^{SV} falls with demand-side exposure is consistent with the markup effect dominating in our CPG setting.

6.3 Quality Composition

Immigrants' price sensitivity operates at two margins. Across firms, it compresses markups. This is the price effect of Proposition 1. Within a firm, it shifts expenditure across the firm's own products. As the product basket is fixed, this shift leaves the Sato-Vartia index unchanged and instead is reflected in measured appeal; the object (5.10) recovers from a price index that holds quality fixed. The within-firm shift comes from $\sigma_i^U > \sigma_n^U$: facing the same products at the same prices, immigrants and natives split their spending across those products differently. Thus, a change in who buys from a firm changes the quality composition of what it sells, even though no price and no product quality has moved. This is the source of the decline in appeal, captured below by the object $\bar{\phi}_f$.[30]

Define firm-level quality as the within-firm expenditure-weighted average quality,

$$\bar{\phi}_f \equiv \sum_{u \in \Omega_f} s_{u|f}\, \phi_u, \tag{6.7}$$

where $s_{u|f}$ is product u's share of the firm's sales, averaging the two types,

$$s_{u|f} = \frac{L_n E_{fn}\, s_{un|f} + L_i E_{fi}\, s_{ui|f}}{L_n E_{fn} + L_i E_{fi}}, \qquad E_{fh} \equiv \phi_f \Big(\frac{P_{fh}}{P_h}\Big)^{1-\sigma_h^F} E_h, \tag{6.8}$$

with E_{fh} the expenditure of a type-h consumer on firm f.

Proposition 2 (Quality Composition Effect)**.** *Provided the firm's products are not all of the same quality and $\varepsilon_a + \delta > 0$, so that they do not all sell at the same price, immigration reduces firm-level quality, holding prices fixed:*

$$\frac{\partial \bar{\phi}_f}{\partial L_i} = \sum_{u \in \Omega_f} \frac{\partial s_{u|f}}{\partial L_i} \phi_u < 0. \tag{6.9}$$

[29] With a common within-firm elasticity, Online Appendix D.3 states the condition as $(\sigma_i^F - \sigma_n^F)\theta_{fn} > \delta\, \sigma_{f,\mathit{eff}}^F (\sigma_{f,\mathit{eff}}^F - 1)$.

[30] We also hold firm appeal ϕ_f and product qualities ϕ_u fixed. Composition is therefore the only channel through which the model can move measured appeal, and it attributes none of the decline in Φ^{Appeal} to the primitives. The estimation in Appendix C.1 lets relative appeal vary over time. The model isolates the part that is fixed to the product and reads the measured decline as a change in what the firm sells.

The effect operates through expenditure reallocation: immigrants' higher within-firm price sensitivity leads them to spend more on the firm's cheaper, and hence lower-quality, products.

The derivative holds prices fixed, which is what isolates the composition channel: at a given price vector, a change in the buyer mix alone reshuffles the firm's sales across its products. How prices themselves respond is the subject of Proposition 1. We prove the result in Online Appendix D.4. Both types see the same products at the same prices and agree on which ones are better, but immigrants substitute more readily on price, so a high price costs a product more of immigrants' spending than of natives'. Higher-quality products have higher prices in equilibrium, so as L_i rises the firm's sales tilt toward its cheaper and lower-quality products and $\bar{\phi}_f$ falls. This channel and the markup channel of Proposition 1 operate at two margins of the same heterogeneity, σ^U within the firm and σ^F across firms, so a composition shift does not enter Φ^{SV} and a price change does not enter Φ^{Appeal}.

7 Conclusion

This paper studies how immigration affects consumer prices through demand and supply channels. Using firm-level data linking product prices to both retail and production locations, we find that immigration's disinflationary effect operates primarily through the demand channel. Immigrants search more and are more price-sensitive than the consumers already in the market, so their arrival flattens the demand curve firms face rather than simply shifting it out, and firms respond by lowering their markups.

A natural expectation is that more consumers also means more quantity sold. We find only a modest increase in physical quantity, alongside a clear reallocation toward private-label goods. Our structural decomposition points the same way. The demand-side price effect operates through markup compression rather than through variety, which is largely muted, and firm-level appeal falls, indicating that immigration shifts the composition of sales toward cheaper and lower-quality products. On the supply side, immigration facilitates product variety expansion and raises the quality of the product composition, consistent with literature on immigrants and innovation. These findings point to heterogeneous demand elasticities, rather than lower labor costs, as the primary mechanism linking immigration to lower prices.

These results have direct implications for policy debates on immigration and inflation. The anti-inflationary effect of immigration operates through changing the composition of demand toward more price-sensitive consumers, suggesting that immigration restrictions aimed at fighting inflation may be counterproductive. Our theoretical framework is partial equilibrium and takes the market as given, so it does not speak to economy-wide inflation. Extending it so that firm-level pricing responses aggregate to aggregate price dynamics, and asking how the demand composition channel interacts with monetary policy and business cycles, is an important avenue for future research. More broadly, our findings suggest that demographic shifts in the consumer base—whether from immigration, aging,

or income redistribution—may be an underappreciated driver of price dynamics.

References

Acemoglu, Daron, Simon Johnson, and James A Robinson, "The Colonial Origins of Comparative Development: An Empirical Investigation," *American Economic Review*, 2001, *91* (5), 1369–1401.

Adao, Rodrigo, Michal Kolsear, and Eduardo Morales, "Shift-Share Designs: Theory and Inference," *The Quarterly Journal of Economics*, 2019, *1949*, 2010.

Aguiar, Mark and Erik Hurst, "Life-cycle prices and production," *American Economic Review*, 2007, *97* (5), 1533–1559.

_ , _ , and Loukas Karabarbounis, "Time Use during the Great Recession," *American Economic Review*, 2013, *103* (5), 1664–1696.

Albert, Christoph and Joan Monras, "Immigration and spatial equilibrium: the role of expenditures in the country of origin," *American Economic Review*, 2022, *112* (11), 3763–3802.

Amuedo-Dorantes, Catalina, Esther Arenas Arroyo, Parag Mahajan, and Bernhard Schmidpeter, "Low-wage jobs, foreign-born workers, and firm performance," 2023.

Argente, David and Munseob Lee, "Cost of living inequality during the great recession," *Journal of the European Economic Association*, 2021, *19* (2), 913–952.

_ , Esteban Méndez, and Diana Van Patten, "Cross-Border Product Adoption: Individual Imports, Migrant Networks, and Domestic Retailers," NBER Working Paper 35104, National Bureau of Economic Research 2026.

Atkeson, Andrew and Ariel Burstein, "Pricing-to-Market, Trade Costs, and International Relative Prices," *American Economic Review*, 2008, *98* (5), 1998–2031.

Atkin, David, Benjamin Faber, and Marco Gonzalez-Navarro, "Retail globalization and household welfare: Evidence from Mexico," *Journal of Political Economy*, 2018, *126* (1), 1–73.

Auer, Raphael, Ariel Burstein, Sarah Lein, and Jonathan Vogel, "Unequal expenditure switching: Evidence from Switzerland," *Review of Economic Studies*, 2024, *91* (5), 2572–2603.

Borjas, George J, "The labor demand curve is downward sloping: Reexamining the impact of immigration on the labor market," *The quarterly journal of economics*, 2003, *118* (4), 1335–1374.

_ and Kirk B Doran, "The Collapse of the Soviet Union and the Productivity of American Mathematicians," *The Quarterly Journal of Economics*, 2012, *127* (3), 1143–1203.

_ and Lawrence F Katz, "The evolution of the Mexican-born workforce in the United States," *Mexican Immigration to the United States*, 2007, *1*, 13–56.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "Quasi-Experimental Shift-Share Research Designs," *Review of Economic Studies*, 2022, *89* (1), 181–213.

Brinatti, Agostina and Nicolas Morales, "Firm heterogeneity and the impact of immigration: Evidence from german establishments," *Available at SSRN 3881995*, 2025.

Broda, Christian and David E. Weinstein, "Globalization and the Gains from Variety," *The Quarterly Journal of Economics*, May 2006, *121* (2), 541–585.

Bronnenberg, Bart J., Jean-Pierre H. Dubé, and Matthew Gentzkow, "The Evolution of Brand Preferences: Evidence from Consumer Migration," *American Economic Review*, 2012, *102* (6), 2472–2508.

Burchardi, Konrad B, Thomas Chaney, and Tarek A Hassan, "Migrants, ancestors, and foreign investments," *The Review of Economic Studies*, 2019, *86* (4), 1448–1486.

Burstein, Ariel, Gordon Hanson, Lin Tian, and Jonathan Vogel, "Tradability and the Labor-Market Impact of Immigration: Theory and Evidence From the United States," *Econometrica*, 2020, *88* (3), 1071–1112.

Cabral, James and Walter Steingress, "Immigration and US shelter prices: The role of geographical and immigrant heterogeneity," *European Economic Review*, 2026, *182*, 105214.

Card, David, "Immigrant inflows, native outflows, and the local labor market impacts of higher immigration," *Journal of Labor Economics*, 2001, *19* (1), 22–64.

_ , "Immigration and inequality," *American Economic Review*, 2009, *99* (2), 1–21.

Cheremukhin, Anton, Sewon Hur, Ronald Mau, Karel Mertens, Alexander Richter, and Xiaoqing Zhou, "The Postpandemic U.S. Immigration Surge: New Facts and Inflationary Implications," Technical Report 2407, Federal Reserve Bank of Dallas Working Paper October 2024.

Clemens, Michael A. and Ethan G. Lewis, "The Effect of Low-skill Immigration Restrictions on US Firms and Workers: Evidence From a Randomized Lottery," Technical Report No. 24-16, GMU Working Paper in Economics 2024.

Clemens, Michael A and Jennifer Hunt, "The Labor Market Effects of Refugee Waves: Reconciling Conflicting Results," *ILR Review*, 2019, *72* (4), 818–857.

Cortes, Patricia, "The Effect of Low-Skilled Immigration on U.S. Prices: Evidence from CPI Data," *Journal of Political Economy*, June 2008, *116* (3), 381–422. Publisher: The University of Chicago Press.

Cravino, Javier, Andrei A Levchenko, Francesc Ortega, and Nitya Pandalai-Nayar, "The Economic Impact of Mass Deportations," Working Paper 34790, National Bureau of Economic Research February 2026.

Dell, Melissa, "The Persistent Effects of Peru's Mining Mita," *Econometrica*, 2010, *78* (6), 1863–1903.

DellaVigna, Stefano and Matthew Gentzkow, "Uniform pricing in us retail chains," *The Quarterly Journal of Economics*, 2019, *134* (4), 2011–2084.

Diewert, W.E., "Exact and superlative index numbers," *Journal of Econometrics*, 1976, *4* (2), 115–145.

Dubé, Jean-Pierre, Günter J. Hitsch, and Peter E. Rossi, "State Dependence and Alternative Explanations for Consumer Inertia," *RAND Journal of Economics*, 2010, *41* (3), 417–445.

Dube, Jean-Pierre, Gunter J Hitsch, and Peter E Rossi, "Income and Wealth Effects on Private-Label Demand: Evidence from the Great Recession," *Marketing Science*, 2018, *37* (1), 22–53.

Dustmann, Christian and Albrecht Glitz, "How do industries and firms respond to changes in

local labor supply?," *Journal of Labor Economics*, 2015, *33* (3), 711–750.

_ , **Tommaso Frattini, and Ian P Preston**, "The effect of immigration along the distribution of wages," *Review of Economic Studies*, 2013, *80* (1), 145–173.

_ , **Uta Schönberg, and Jan Stuhler**, "Labor supply shocks, native wages, and the adjustment of local employment," *The Quarterly Journal of Economics*, 2017, *132* (1), 435–483.

_ , **Uta Schönberg, and Jan Stuhler**, "The impact of immigration: why do studies reach such different results?," *Journal of Economic Perspectives*, 2016, *30* (4), 31–56.

Eslava, Marcela, John Haltiwanger, and Nicolas Urdaneta, "The size and life-cycle growth of plants: The role of productivity, demand, and wedges," *Review of Economic Studies*, 2024, *91* (1), 259–300.

Faber, Benjamin and Thibault Fally, "Firm heterogeneity in consumption baskets: Evidence from home and store scanner data," *The Review of Economic Studies*, 2022, *89* (3), 1420–1459.

Feenstra, Robert C, "New Product Varieties and the Measurement of International Prices," *American Economic Review*, March 1994, *84* (1), 157–177.

Feenstra, Robert C. and John Romalis, "International Prices and Endogenous Quality," *Quarterly Journal of Economics*, 2014, *129* (2), 477–527.

Fitzgerald, Terry, Callum Jones, Mariano Kulish, and Juan Pablo Nicolini, "Is there a stable relationship between unemployment and future inflation?," *American Economic Journal: Macroeconomics*, 2024, *16* (4), 114–142.

Furlanetto, Francesco and Ørjan Robstad, "Immigration and the macroeconomy: Some new empirical evidence," *Review of Economic Dynamics*, 2019, *34*, 1–19.

Galaasen, Sigurd, Andreas Ravndal Kostøl, Joan Monras, and Jonathan Vogel, "The labor supply curve is upward sloping: The labor market effects of immigrant-induced demand shocks," Technical Report, IZA Discussion Papers 2025.

Goetz, Daniel and Alexander Rodnyansky, "Exchange rate shocks and quality adjustments," *Review of Economics and Statistics*, 2023, *105* (1), 86–100.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik Instruments: What, When, Why, and How," *American Economic Review*, 2020, *110* (8), 2586–2624.

Handbury, Jessie, "Are Poor Cities Cheap for Everyone? Non-Homotheticity and the Cost of Living Across US Cities," *Econometrica*, 2021, *89* (4), 1675–1724.

_ **and David E Weinstein**, "Goods prices and availability in cities," *The Review of Economic Studies*, 2015, *82* (1), 258–296.

Hazell, Jonathon, Juan Herreno, Emi Nakamura, and Jón Steinsson, "The slope of the Phillips Curve: evidence from US states," *The Quarterly Journal of Economics*, 2022, *137* (3), 1299–1344.

Hitsch, Günter J., Ali Hortaçsu, and Xiliang Lin, "Prices and Promotions in U.S. Retail Markets," *Quantitative Marketing and Economics*, 2019, *17*, 1–43.

Hottman, Colin J., Stephen J. Redding, and David E. Weinstein, "Quantifying the Sources

of Firm Heterogeneity *," *The Quarterly Journal of Economics*, August 2016, *131* (3), 1291–1364.

Jaimovich, Nir, Sergio Rebelo, and Arlene Wong, "Trading down and the business cycle," *Journal of Monetary Economics*, 2019, *102*, 96–121.

Jaravel, Xavier, "The Unequal Gains from Product Innovations: Evidence from the U.S. Retail Sector*," *The Quarterly Journal of Economics*, May 2019, *134* (2), 715–783.

Kaplan, Greg and Guido Menzio, "Shopping externalities and self-fulfilling unemployment fluctuations," *Journal of Political Economy*, 2016, *124* (3), 771–825.

Kerr, Sari Pekkala, William R. Kerr, and William F. Lincoln, "Skilled immigration and the employment structures of U.S. firms," *Journal of Labor Economics*, 2015, *33* (1), S147–S186.

Kerr, William R. and William F. Lincoln, "The Supply Side of Innovation: H-1B Visa Reforms and U.S. Ethnic Invention," *Journal of Labor Economics*, 2010, *28* (3), 473–508.

Khanna, Gaurav and Munseob Lee, "High-skill immigration, innovation, and creative destruction," in "The roles of immigrants and foreign students in US science, innovation, and entrepreneurship," University of Chicago Press, 2019, pp. 73–98.

Lach, Saul, "Immigration and Prices," *Journal of Political Economy*, August 2007, *115* (4), 548–587. Publisher: The University of Chicago Press.

Lenzu, Simone, David Rivers, and Joris Tielens, "Financial Shocks, Productivity, and Prices," Technical Report 2022.

Leung, Justin H., "Minimum Wage and Real Wage Inequality: Evidence from Pass-Through to Retail Prices," *The Review of Economics and Statistics*, October 2021, *103* (4), 754–769.

Mahajan, Parag, "Immigration and business dynamics: Evidence from US firms," *Journal of the European Economic Association*, 2024, p. jvae022.

_ , Nicolas Morales, Kevin Shih, Mingyu Chen, and Agostina Brinatti, "The impact of immigration on firms and workers: Insights from the h-1b lottery," 2024.

Manova, Kalina and Zhiwei Zhang, "Export prices across firms and destinations," *The Quarterly Journal of Economics*, 2012, *127* (1), 379–436.

McCully, Brett A., Torsten Jaccard, and Christoph Albert, "Immigrants, Imports, and Welfare: Evidence from Household Purchase Data," Working Paper 12278, CESifo 2025.

Mongey, Simon and Michael E Waugh, "Pricing Inequality," Technical Report, National Bureau of Economic Research 2025.

Monras, Joan, "Immigration and Wage Dynamics: Evidence from the Mexican Peso Crisis," *Journal of Political Economy*, 2020, *128* (8), 3017–3089.

Nevo, Aviv and Arlene Wong, "The Elasticity of Substitution between Time and Market Goods: Evidence from the Great Recession," *International Economic Review*, 2019, *60* (1), 25–51.

Nunn, Nathan, "The Long-Term Effects of Africa's Slave Trades," *The Quarterly Journal of Economics*, 2008, *123* (1), 139–176.

Olney, William W., "IMMIGRATION AND FIRM EXPANSION," *Journal of Regional Science*, 2013, *53* (1), 142–157.

Ottaviano, Gianmarco IP and Giovanni Peri, "Rethinking the effect of immigration on wages," *Journal of the European Economic Association*, 2012, *10* (1), 152–197.

Peri, Giovanni, "Immigrants, productivity, and labor markets," *Journal of Economic Perspectives*, 2016, *30* (4), 3–30.

_ **and Chad Sparber**, "Task Specialization, Immigration, and Wages," *American Economic Journal: Applied Economics*, June 2009, *1* (3), 135–169.

Prokhorenkova, Liudmila, Gleb Gusev, Aleksandr Vorobev, Anna Veronika Dorogush, and Andrey Gulin, "CatBoost: unbiased boosting with categorical features," *Advances in neural information processing systems*, 2018, *31.*

Sá, Filipa, "Immigration and house prices in the UK," *The Economic Journal*, 2015, *125* (587), 1393–1424.

Saiz, Albert, "Immigration and housing rents in American cities," *Journal of urban Economics*, 2007, *61* (2), 345–371.

_ **and Susan Wachter**, "Immigration and the neighborhood," *American Economic Journal: Economic Policy*, 2011, *3* (2), 169–188.

Sanchis-Guarner, Rosa, "Decomposing the impact of immigration on house prices," *Regional Science and Urban Economics*, 2023, *100*, 103893.

Sangani, Kunal, "Markups across the income distribution: Measurement and implications," *Available at SSRN 4092068*, 2024.

Sato, Kazuo, "The ideal log-change index number," *The Review of Economics and Statistics*, 1976, pp. 223–228.

Simonovska, Ina, "Income differences and prices of tradables: Insights from an online retailer," *The Review of Economic Studies*, 2015, *82* (4), 1612–1656.

Stroebel, Johannes and Joseph Vavra, "House Prices, Local Demand, and Retail Prices," *Journal of Political Economy*, 2019, *127* (3), 1391–1436.

Terry, Stephen J., Thomas Chaney, Konrad B. Burchardi, Lisa Tarquinio, and Tarek A. Hassan, "Immigration, Innovation, and Growth," *Mimeo*, 2024.

U.S. Bureau of Labor Statistics, "Relative Importance of Components in the Consumer Price Index: U.S. City Average, December 2018," CPI Detailed Report, U.S. Department of Labor 2018.

_ , "Measuring Price Change in the CPI: Rent and Rental Equivalence," CPI Factsheet 2024.

Vartia, Yrjö O, "Ideal log-change index numbers," *scandinavian Journal of statistics*, 1976, pp. 121–126.

Zachariadis, Marios, "Immigration and international prices," *Journal of International Economics*, 2012, *87* (2), 298–311.

Zhu, Yumin, Luca Macedoni, and Mingzhi Xu, "Inequality and Market Power: Evidence from the United States and China," 2025.

Appendix A Data Appendix

A.1 NielsenIQ GS1 + NETS Matching

To identify the production locations of the consumer goods in our price dataset, we link firms in the GS1 dataset (which provides the brand and manufacturer name via the UPC prefix) to the National Establishment Time-Series (NETS) database (which provides the headquarters address and employment counts).

We utilize the 2019 vintage of the NETS database and the 2017 vintage of the GS1 database. Because there is no common numeric identifier between these two administrative datasets, we develop a multi-step string matching algorithm. The procedure consists of three stages: Data Pre-processing, Hierarchical Matching, and Cardinality Resolution.

A.1.1 Data Pre-processing and Standardization

Before matching, we standardized company names and addresses in both datasets to maximize match rates.

1. **Lexical Standardization:** We utilized the Stata commands `stnd_compname` and `stnd_address` to harmonize common abbreviations (e.g., "Street" vs. "St", "Corporation" vs. "Corp") and remove non-identifying characters.

2. **Entity Type Cleaning:** To ensure that minor variations in legal registration did not prevent a match (e.g., matching "Smith Foods Inc" to "Smith Foods LLC"), we generated stripped versions of company names by systematically removing legal entity types (e.g., INC, LLC, LTD, CO) and their spacing.

3. **Geographic Harmonization:** City names were converted to uppercase to match the GS1 format. ZIP codes were split to isolate the primary 5-digit code from the 4-digit extension to allow for flexible matching levels.

4. **State Formatting:** We harmonized state codes across datasets, ensuring the inclusion of U.S. territories (Puerto Rico and Virgin Islands) which are present in our scanner data.

A.1.2 Hierarchical Matching Algorithm

We employed an iterative matching process that prioritizes precision over recall. We constructed a "match quality" variable based on four increasingly strict criteria.

First, we performed an exact join based on the standardized company name (`stn_name`). We then refined this pool using state and address data. For the fuzzy address matching step, we utilized the `reclink2` algorithm with a bigram string comparator. We assigned matching weights of

`wmatch(10 1 2 5 8)` corresponding to Name, State, City, Zip Code, and Street Address, respectively, requiring a mandatory exact match on Name and State.

We classified the resulting matches into the following quality tiers:

- **Tier 1 (Name Only):** Exact match on the standardized company name.
- **Tier 2 (Name + Entity):** Exact match on the standardized name AND an exact match on the stripped entity type (e.g., verifying that both files listed the firm as an "INC" or that the suffix was absent in both).
- **Tier 3 (Name + Entity + State):** Meets Tier 2 criteria AND the headquarters state in NETS matches the firm's state in GS1.
- **Tier 4 (Full Address Validation):** Meets Tier 3 criteria AND satisfies a high-confidence fuzzy match on the specific street address (calculated via `reclink2`).

As the baseline for the validation procedure described below, we first retain the strict **Tier 4** matches (labeled as `match >= 4` in our replication code), which match on name, legal entity structure, state, and specific street address simultaneously. The analysis sample is the augmented set defined at the end of this subsection, which combines these Tier 4 pairs with the same-city and API-validated different-city pairs.

A.1.3 Cardinality Resolution (Many-to-One vs. One-to-Many)

A structural challenge in linking these datasets is the cardinality of the relationship. A single firm often registers multiple GS1 prefixes (one for each brand or product line), while a single GS1 prefix should theoretically map to a single parent firm. However, name duplication can lead to "One-to-Many" (1:m) links where a single brand name matches multiple distinct NETS headquarters Duns Numbers.

To ensure proper identification of the firm's location, we applied the following resolution logic to the matched pairs:

- **Retained (1:1 and m:1):** We retained cases where a single GS1 prefix matched a single NETS ID (1:1), or where multiple GS1 prefixes matched the same NETS ID (m:1). The latter represents a multi-brand conglomerate (e.g., a parent company owning multiple distinct brand prefixes) and constitutes a valid link for our identification strategy, as these brands face the same firm-level shocks.
- **Dropped (1:m and m:m):** We dropped cases where a single GS1 prefix linked to multiple different NETS Duns Numbers. These represent ambiguous matches (e.g., a generic name like "Best Foods" appearing in multiple states) where we cannot definitively identify which headquarters owns the brand.

The resulting dataset serves as the baseline for the API-based validation described in the subsequent section.

A.1.4 Validation of Firm Linkages via Web Search API

A limitation of the strict address matching described above (specifically Tier 4) is that it may exclude valid firm-plant linkages. For example, a firm may list a manufacturing facility's address in the GS1 dataset while its corporate headquarters appears in the NETS dataset. These locations often share a state but differ in city and street address, causing them to fail the strict address fuzzy match. Conversely, matching solely on "Name + State" (Tier 3) risks false positives for common firm names (e.g., "Best Foods") that may exist independently in multiple cities within the same state.

To resolve this trade-off—expanding recall to include valid multi-establishment firms while maintaining high precision—we implemented a secondary validation procedure using a web search API. This process targets the subset of firm pairs that matched on *Name* and *State* but failed the strict *Address* match.

A.1.5 Sample Stratification and Pre-Screening

We first restricted the validation sample to firm pairs located in the same state. This decision followed a manual verification pilot (conducted in February 2024) which indicated that cross-state matches for this subset of firms were overwhelmingly false positives. We then stratified the same-state pairs into two categories based on city concordance:

1. **Same-City Matches (Auto-Accepted):** Pairs located in the same city (but with differing street addresses) were deemed high-probability matches. Given the low likelihood of two distinct firms with identical names operating in the same city, and considering the computational cost of API queries, these pairs were accepted without further validation. These observations are stored in the dataset `manual_match_all_small_samecity.dta`.

2. **Different-City Matches (API Validated):** Pairs located in the same state but different cities represent the ambiguous cases requiring digital verification. These pairs were subjected to the API matching algorithm described below. These observations are processed in the dataset `apimatch_diffct.dta`.

A.1.6 API Matching Algorithm

For the "Different-City" pairs, we automated a web search procedure to determine if the two physical locations belonged to the same corporate entity. The algorithm proceeds as follows:

1. **Query Generation:** For each firm in the pair (the GS1 entity and the NETS entity), we constructed a search query string combining the firm's name, city, and state (e.g., "[Firm Name]

[City] [State]").

2. **URL Extraction:** We utilized a custom search API to fetch the top search results for each query and extracted the root domains of the associated URLs (e.g., extracting `company-xyz.com` from `www.company-xyz.com/about`).

3. **Common Domain Verification:** We compared the set of URLs retrieved for the GS1 firm (U_{GS1}) with the set retrieved for the NETS firm (U_{NETS}). A pair is classified as a validated match (`MatchStatus = 1`) if the intersection of their domains is non-empty:

$$U_{GS1} \cap U_{NETS} \neq \emptyset$$

This criterion leverages the fact that even if a firm has distinct physical addresses for its plant and HQ, both locations will typically reference the same primary corporate website in search results.

This procedure allows us to distinguish between valid corporate networks (e.g., a "Traders Point Creamery" in Zionsville, IN matching a production facility in Indianapolis, IN via the common domain `traderspointcreamery.com`) and spurious name collisions. The final dataset combines the strictly matched address pairs (Tier 4), the same-city pairs, and the API-validated different-city pairs.

A.2 Identifying Immigrant Households in the NielsenIQ Panel

This appendix describes the machine learning framework used to predict immigrant households in the NielsenIQ Household Panel (HMS), enabling analysis of immigrant-specific shopping behavior and price sensitivity. Because the HMS does not directly identify immigration status, we develop a supervised learning approach that leverages labeled data from the American Community Survey (ACS) and a supplementary NielsenIQ survey module to predict immigrant households within the panel.

A.2.1 Data Sources and Domain Adaptation Challenge

We integrate three complementary data sources covering 2004–2019:

1. **American Community Survey (ACS):** Household-level microdata from IPUMS providing nationally representative information on demographics, education, occupation, and—crucially—immigration status and nativity. The ACS contains approximately 1.4 million household observations and serves as the primary labeled training data.

2. **NielsenIQ Household Panel (HMS):** Annual household demographic and purchasing information from participating panelists, including household composition, income, education, and shopping behavior. HMS serves as the prediction target but does not contain immigration labels.

3. **Tell Us More About You (TUM) 2008 Survey:** A supplementary questionnaire administered to approximately 40,000 HMS households in 2008, collecting place of birth and migration history. TUM provides ground-truth immigration labels within the HMS sampling frame.

The central methodological challenge is *domain adaptation*: while the ACS provides abundant labeled data on immigration status, it exhibits substantial covariate shift relative to the HMS target population. The TUM module bridges these domains by providing labeled observations drawn directly from the HMS sampling frame, though its sample size is limited. Table A.1 summarizes these three data sources.

Table A.1: Comparison of Data Sources for Immigrant Prediction

	ACS	**HMS**	**TUM 2008**
Years	2004–2019	2004–2019	2008
Sample size (households)	∼1.4M	∼200k	∼40k
Immigration label	Yes	No	Yes
Analytical role	Labeled source	Unlabeled target	Bridge domain

A.2.2 Data Harmonization

To ensure cross-source consistency, all datasets are transformed into a unified household-level schema. Person-level ACS records are aggregated to the household level, retaining only household heads. Occupational codes are mapped via crosswalks reconciling pre-2017 and post-2018 IPUMS schemes, and income and education categories are re-binned to match HMS codes. For HMS, households with both male and female heads are split into separate observations, variable names are aligned with ACS equivalents, and missing values are retained as categorical levels (since our algorithm handles them directly). TUM responses are merged into corresponding HMS household records using unique household identifiers, with careful cleaning to ensure one valid observation per household.

A.2.3 Machine Learning Algorithm

We employ CatBoost (Prokhorenkova et al., 2018), a gradient boosting decision tree (GBDT) algorithm specifically designed for categorical data. Several features make this approach well-suited to our prediction task:

- **Nonlinearity and interactions:** Decision trees automatically capture nonlinear effects and variable interactions (e.g., between age, income, and education) without manual specification.
- **Efficient categorical encoding:** CatBoost replaces each category with an ordered target statistic computed using only "past" observations, avoiding target leakage while efficiently handling high-cardinality categorical variables common in survey data.
- **Robustness to mixed data types:** Trees handle both continuous and categorical variables, as well as missing values, with minimal preprocessing.
- **Reduced overfitting:** Ordered boosting—building trees using different permutations of the training data—produces unbiased gradient estimates and improves generalization on small labeled samples such as TUM.

A.2.4 Domain Adaptation Strategy

To address covariate shift between the ACS (source domain) and HMS (target domain), we implement multiple training regimes and select the best-performing approach:

1. **ACS-only regimes:** Models trained on the full ACS sample, with validation either on an ACS hold-out fold or on TUM.
2. **IPS-weighted ACS regimes:** Models trained on ACS but reweighted using inverse propensity scores (IPS) to make ACS resemble TUM or HMS. IPS weights are constructed by training a

classifier to distinguish ACS from the target domain; observations more similar to the target receive higher weights:

$$w_i = \frac{1}{1 - \hat{p}(D = \text{TUM}|X_i)} \tag{A.1}$$

where $\hat{p}(D = \text{TUM}|X_i)$ is the predicted probability that household i originates from TUM given covariates X_i.

3. **ACS–TUM mixture regimes:** Models trained on mixtures of downsampled ACS and TUM training data, using either fixed ratios (1:1 or 2:1) or IPS-adjusted domain weights.

Across all models intended for HMS prediction, hyperparameter tuning, early stopping, and threshold selection rely exclusively on TUM validation data, since TUM is the only labeled dataset aligned with the HMS sampling frame. TUM samples are stratified by immigrant status and split 80:20 into training and validation sets at the household level.

A.2.5 Household-Level Aggregation

Model predictions are obtained at the individual-year level and then aggregated to the household level using one of four strategies: maximum, mean, median, or "noisy-or" probability aggregation:

$$\hat{p}_h^{\text{noisy-or}} = 1 - \prod_{i \in h} (1 - \hat{p}_i) \tag{A.2}$$

The optimal aggregation strategy is selected empirically using TUM validation data.

A.2.6 Model Performance

Table A.2 reports predictive performance on the held-out TUM validation split for each model family, evaluated using AUC (ranking performance) and F1 score (identification of the minority immigrant class).

Table A.2: Predictive Performance on TUM Validation by Model Family

Model Family	AUC	F1
ACS-IPS-TUM	0.770	**0.371**
ACS-IPS-HMS	0.771	0.365
ACS-transfer	0.758	0.346
ACS-only	0.753	0.334
TUM-only	0.754	0.318
Mix 2:1	0.761	0.230
Mix 1:1	0.766	0.224
Mix-IPS	0.732	0.182

Notes: AUC measures the probability that a randomly chosen immigrant household receives a higher predicted score than a randomly chosen non-immigrant household. F1 is the harmonic mean of precision and recall for identifying immigrant households. Best strategy refers to the optimal household-level aggregation method selected on TUM validation data.

Several patterns emerge from these results:

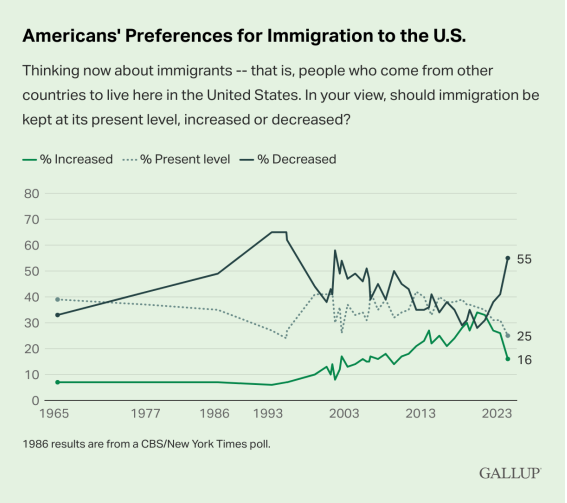
- **Covariate shift matters:** ACS-only models perform well on ACS but degrade substantially when evaluated on TUM, reflecting systematic differences between the nationally representative ACS and the NielsenIQ panel population.
- **IPS weighting is effective:** Reweighting ACS observations to match the target domain distribution substantially improves both AUC and F1, with the ACS-IPS-TUM model achieving the highest F1 score (0.371).
- **Ranking vs. identification trade-off:** Mixture models preserve high AUC close to the TUM-only benchmark, indicating strong ranking ability, but their F1 scores are lower due to more conservative identification of minority-class households.

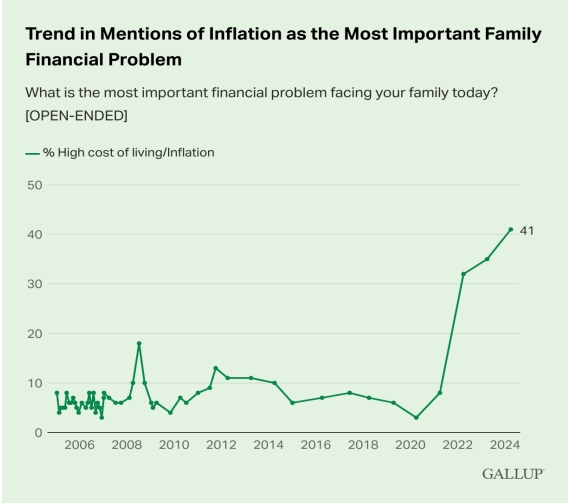
The best-performing model (ACS-IPS-TUM) achieves an AUC of 0.77 and F1 of 0.37, representing substantial improvement over rule-based proxies commonly used in economics (e.g., ethnicity-based or surname-based heuristics). These gains are important because the constructed immigrant measure feeds directly into our analysis of immigrant shopping behavior and price sensitivity in Sections 4 and 5.

A.2.7 Application to Shopping Behavior Analysis

The selected model is applied to all HMS households not included in TUM to generate predicted immigration probabilities. These predictions are then scaled using HMS projection factors to recover population-level immigrant shares by county. In our firm-level analysis, we use these predictions to characterize the consumer base of firms facing different levels of immigrant demand exposure, enabling tests of whether immigrants exhibit systematically different shopping patterns (e.g., more trips, more stores visited, greater coupon usage) that would be consistent with heightened price search intensity.

Figure B.1: Opinion Polls on Immigration and Inflation Concerns



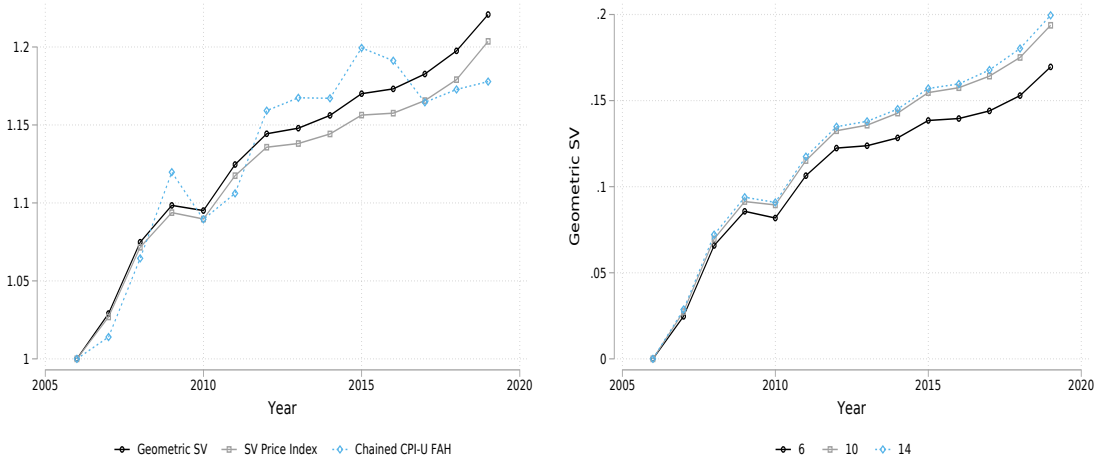


Notes: Figures are sourced from GALLUP polls. Figure on the left is sourced from an article on July 12, 2024, with the following link: Sharply More Americans Want to Curb Immigration to U.S.. Figure on the right is sourced from an article on May 2, 2024, with the following link: Americans Continue to Name Inflation as Top Financial Problem

Appendix B Additional Tables and Figures

B.1 Motivation, Data, and First Stage

Figure B.2: County Level Price Index: Comparison with CPI Food at Home and Varying Restrictions on the Set of Products



Notes: The left figure compares our constructed county-level price index using CPG data with the most closely related reported (national) BLS CPI. The CPI used is the "Food from Home" Chained Index. "Geometric SV" refers to just the normal Sato-Vartia price index, while the "SV Price Index" incorporates the Feenstra variety correction. We aggregate the county-level index to the national level. In the right figure, we compare our constructed price index under several restrictions on the set of products in our sample. The baseline in the paper is the most restrictive one, where products must be available for 14 years. We also construct the Sato-Vartia price index (geometric aggregation) for an expanded set of products, lowering the restriction to 10 and 6 years.

B.2 Instrument Construction

This appendix gives the full construction of the predicted-ancestry instrument summarized in Section 3.4, following Terry et al. (2024).

The standard shift-share instrument. The canonical shift-share instrument predicts current immigration to county d from origin o as the product of a "share" and a "shift":

$$\hat{I}_{odt} = A_{od,t-1} \times I_{o,-d,t}, \tag{B.1}$$

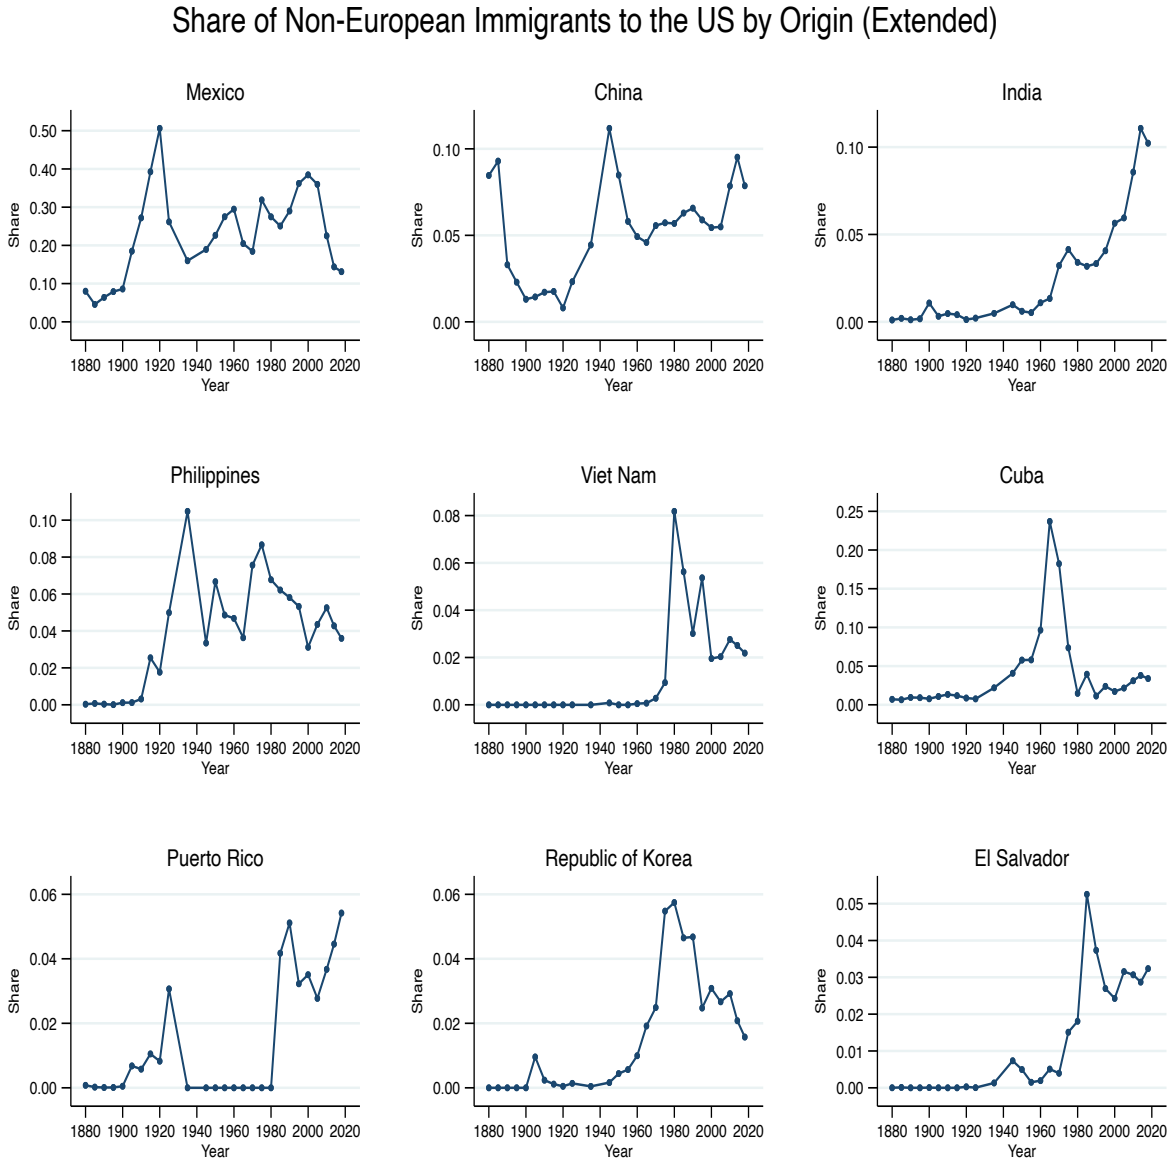
where $A_{od,t-1}$ is the ancestry network, the number of residents of origin-o ancestry living in county d in period $t-1$ (a raw network-exposure count, not a normalized share, whose scale is absorbed by the first-stage coefficient), and $I_{o,-d,t}$ is total immigration from o to all counties except d in period t. The identifying assumption is that historical ancestry networks predict where new immigrants settle through family ties and established communities, but are orthogonal to current local shocks. The concern is persistence. If the economic conditions that historically attracted origin-o immigrants to d persist, then $A_{od,t-1}$ is correlated with current unobserved shocks to d that also move prices, which violates the exclusion restriction.

Predicted ancestry. To break this persistence, we instrument for ancestry itself using exogenous historical push and pull factors starting from 1880, predicting county-level ancestry from the interaction of origin-country emigration and destination attractiveness to an unrelated immigrant group. The first stage predicts ancestry A_{odt} from historical migration patterns:

$$A_{odt} = \delta_{or(d)t} + \delta_{c(o)dt} + X'_{od}\zeta + \sum_{\tau \in \mathcal{T},\, \tau \leq t} a^{t}_{r(d)\tau} \left(I_{o,-r(d),\tau} \times \frac{I_{Europe,d\tau}}{I_{Europe,\tau}} \right) + \nu_{odt}. \tag{B.2}$$

Here $\mathcal{T}$ is the set of historical census years, and we estimate (B.2) separately for each ancestry census year t, using only push-pull terms from years $\tau \leq t$. The term $I_{o,-r(d),\tau}$ is total immigration from origin o in year τ, excluding migration to the census region $r(d)$ containing d (we use nine U.S. census divisions). This leave-out measure captures exogenous push factors, such as wars, famines, and crises, that drove emigration from o. The term $I_{Europe,d\tau}/I_{Europe,\tau}$ is the share of all European immigrants in year τ who settled in county d, capturing d's pull attractiveness to an unrelated group through transportation, port access, climate, and established communities. The fixed effects $\delta_{or(d)t}$ (origin by destination-region by time) and $\delta_{c(o)dt}$ (origin-continent by destination-county by time) absorb broad geographic settlement patterns, and X'_{od} collects time-invariant dyadic controls such as origin-destination distance and colonial ties. The interaction of origin-specific push and destination pull generates plausibly exogenous variation in where non-European immigrants' descendants live today.

Figure B.3: Share of Non-European Immigrants to the US by Origin



Notes: We plot the share of non-European immigration into the US from the 9 non-Euro origins with the largest immigration into the US. For the first 5, this is effectively an extension of Terry et al. (2024). We extend it to more countries to reflect the important share in this period from Central America and Korea.

Intuition through an example. Consider predicted ancestry of Chinese immigrants in Silver Bow County (home to the mining city of Butte), Montana. It is high when two conditions hold. First, many Chinese emigrants left China in 1880 under economic and political pressure and settled across the United States. Measured using settlement *outside* the Mountain West to avoid mechanical correlation with Montana, this captures strong origin-country push. Second, Silver Bow County drew a large share of European immigrants in 1880 through its copper-mining boom and railroad access, which made it attractive to immigrants generally, not just Europeans. The instrument interacts these two components, origin-country push with destination pull to an unrelated group, so the predicted ancestry arises entirely from historical factors plausibly unrelated to 2006–2018 CPG prices.

We obtain predicted ancestry as the fitted push-pull component, residualized against the controls:

$$\hat{A}_{odt} = \sum_{\tau \in \mathcal{T},\, \tau \leq t} \hat{a}^{\,t}_{r(d)\tau} \left(I_{o,-r(d),\tau} \times \frac{I_{Europe,d\tau}}{I_{Europe,\tau}} \right)^{\perp}, \tag{B.3}$$

where $\hat{a}^{\,t}_{r(d)\tau}$ are the coefficients from (B.2) and $\perp$ denotes residualization against all controls in (B.2). Predicted ancestry is time-varying for two reasons: the set of push-pull terms accumulates as t advances ($\tau \leq t$), and the coefficients are re-estimated for each ancestry census year. The $t-1$ subscript on $\hat{A}_{od,t-1}$ in (B.4) denotes ancestry evaluated at the census year preceding period t.

Final instrument. Returning to the shift-share logic, the final first stage predicts actual immigration:

$$I_{odt} = \delta_{or(d)} + \delta_{c(o)d} + \delta_t + X'_{od}\theta + b_t \left(\hat{A}_{od,t-1} \times \tilde{I}_{o,-r(d),t} \right) + u_{odt}, \tag{B.4}$$

where the shifter is scaled immigration from origin o to regions outside $r(d)$,

$$\tilde{I}_{o,-r(d),t} = I_{o,-r(d),t} \times \frac{I_{Europe,r(d),t}}{I_{Europe,-r(d),t}}. \tag{B.5}$$

The second term corrects for region size. Larger regions mechanically receive more immigrants, so we scale by the region's share of European immigration to make shifters comparable across regions. Because this scaling reflects European immigration and enters symmetrically across origins, it carries no origin-specific information, and the instrument's identifying variation comes only from its interaction with origin-specific predicted ancestry. Aggregating across origins gives the county instrument Z_{dt} in equation (3.5), whose variation comes from three components plausibly orthogonal to current local conditions: historical settlement embedded in predicted ancestry, contemporary origin-country outflows in the leave-out region, and historical destination attractiveness to Europeans. We estimate (B.4) over nine periods at five-year intervals through 2010 and four-year intervals thereafter (2014 and 2018), pairing each lagged predicted ancestry $\hat{A}_{od,t-1}$ (reaching back to 1975) with the contemporaneous shifter.

Table B.3: First Stage Results: Immigration on Push-Pull Instruments at Country-County Level

	$\text{Immigration}^t_{o,d}$				
	(1)	(2)	(3)	(4)	(5)
$\hat{A}_{o,d,1975} \times \tilde{I}_{o,-r(d),1980}$	0.0034***	0.0034***	0.0034***	0.0034***	0.0033***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
$\hat{A}_{o,d,1980} \times \tilde{I}_{o,-r(d),1985}$	0.0016***	0.0016***	0.0016***	0.0016***	0.0015***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
$\hat{A}_{o,d,1985} \times \tilde{I}_{o,-r(d),1990}$	0.0018***	0.0018***	0.0018***	0.0018***	0.0017***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
$\hat{A}_{o,d,1990} \times \tilde{I}_{o,-r(d),1995}$	0.0005***	0.0005***	0.0005***	0.0005***	0.0004***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
$\hat{A}_{o,d,1995} \times \tilde{I}_{o,-r(d),2000}$	0.0003***	0.0003***	0.0003***	0.0003***	0.0003***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
$\hat{A}_{o,d,2000} \times \tilde{I}_{o,-r(d),2005}$	0.0002***	0.0002***	0.0002***	0.0002***	0.0002***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
$\hat{A}_{o,d,2005} \times \tilde{I}_{o,-r(d),2010}$	0.0002***	0.0002***	0.0002***	0.0002***	0.0002***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
$\hat{A}_{o,d,2010} \times \tilde{I}_{o,-r(d),2014}$	0.0001***	0.0001***	0.0001***	0.0001***	0.0001***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
$\hat{A}_{o,d,2014} \times \tilde{I}_{o,-r(d),2018}$	0.0001***	0.0001***	0.0001***	0.0001***	0.0000***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
$I^t_{Euro,d}$				0.0255***	
				(0.0047)	
$\text{I}^t_{o,-r(d),t} \frac{I_{Europe,r(d)\tau}}{I_{Europe,-r(d)\tau}}$					0.2805**
					(0.1155)
Obs.	4,607,847	4,607,847	4,607,847	4,607,847	4,607,847
adj R-sq	0.602	0.602	0.631	0.637	0.649
Distance	No	Yes	Yes	Yes	Yes
Latitude Dis.	No	Yes	Yes	Yes	Yes
Region-Country FE	No	No	Yes	Yes	Yes
County-Continent FE	No	No	Yes	Yes	Yes
Year FE	No	No	Yes	Yes	Yes
Concurrent Euro Immigration	No	No	No	Yes	No
Contemporaneous Push-Pull	No	No	No	No	Yes

Notes: Specification follow Terry et al. (2024), to show we can replicate their instruments. In all cases, the push-pull instruments follow the specification in (B.4), with one instrument for every year t. We extend the data by adding 2014 and 2018. The first column represents a specification with no controls or fixed effects, the second column adds Distance controls, and column (3) adds the full set of fixed effects. The third column is our baseline specification. In column (4) we additionally control for concurrent European immigration, while in the last column we instead control for the contemporaneous interaction term from the ancestry prediction. Standard errors clustered are by country. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

B.3 County-Level Analysis

Table B.4: Robustness to Alternative Price Index Formulations

	Δ P$^{\text{Jevons}}$ (2006-2018)		Δ P$^{\text{Laspyres}}$ (2006-2018)		Δ P$^{\text{Passche}}$ (2006-2018)		Δ P$^{\text{Tornqvist}}$ (2006-2018)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Immigr$_{d,t}$	-0.021**	-0.014**	-0.069***	-0.067***	-0.067***	-0.063***	-0.068***	-0.065***
	(0.008)	(0.006)	(0.012)	(0.011)	(0.014)	(0.012)	(0.013)	(0.012)
Obs.	6,339	6,339	6,339	6,339	6,339	6,339	6,339	6,339
First-Stage F-stat	254.4	94.4	254.4	94.4	254.4	94.4	254.4	94.4
Fixed Effects	Time	State, Time	Time	State, Time	Time	State, Time	Time	State, Time
Method	IV	IV	IV	IV	IV	IV	IV	IV

Notes: Table reports IV regressions of log changes in county-level CPG price indices on immigration using alternative index formulations. Columns (1)-(2) use the Jevons (unweighted geometric mean) index; columns (3)-(4) use the Laspeyres index; columns (5)-(6) use the Paasche index; columns (7)-(8) use the Törnqvist index. Odd-numbered columns include only period fixed effects; even-numbered columns add state and period fixed effects. All specifications instrument immigration with predicted shocks from equation (3.5) and winsorize immigration at top and bottom 1%. All specifications use stacked four-year differences (2006-2010, 2010-2014, 2014-2018). Standard errors clustered by state. ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

Table B.5: Long-Difference (one period)

	Δ P$^{\text{SV,LongDiff}}$ (2006-2018)					
	(1)	(2)	(3)	(4)	(5)	(6)
Immigr$_{d,t}$	-0.041***	-0.029***	-0.024***	-0.054***	-0.036***	-0.031***
	(0.009)	(0.004)	(0.004)	(0.011)	(0.007)	(0.008)
Obs.	2,113	2,112	2,082	2,113	2,112	2,082
First-Stage F-stat				273.5	92.4	89.4
Fixed Effects	None	State	State	None	State	State
Method	OLS	OLS	OLS	IV	IV	IV
Controls			Yes			Yes

Notes: Specifications follow the order of Table 3, but for the 12-year long difference of prices and immigration. Standard errors are clustered by state. Both immigration shock and immigration are winsorized by top and bottom 1%. Controls: 2006 county GDP per capita and urban indicator, each interacted with the year indicators.

Table B.6: County Price Regressions: Robustness to Further Controls

	Δ $P^{\text{Sato-Vartia,Stacked}}$ (2006-2018)						
	(1) Baseline	(2) + Border	(3) + ICE Arrests	(4) + Lead	(5) StatexYear FEs	(6) Non-Winsorized	(7) + County FE
$\text{Immigr}_{d,t}$	-0.058***	-0.058***	-0.058***	-0.107**	-0.038***	-0.014***	-0.060**
	(0.013)	(0.013)	(0.012)	(0.049)	(0.009)	(0.004)	(0.027)
Mexico Border		-0.341***					
		(0.115)					
total ICE Arrests			0.047***				
			(0.012)				
Immigration (t+4)				0.049			
				(0.044)			
Obs.	6,249	6,249	6,249	4,166	6,246	6,249	6,339
First-Stage F-stat	85.5	86.4	101.5	21.7	77.7	43.8	266.5
Fixed Effects	State, Time	State, Time	State, Time	State, Time	State-Time	State, Time	County, Time

Notes: This table reports estimates of immigration effects on county-level CPG prices (Sato-Vartia index). All specifications use stacked four-year differences (2006–2010, 2010–2014, 2014–2018), include controls for initial county GDP per capita and urban indicator each interacted with period dummies, and cluster standard errors by state. All columns except the last two are weighted by initial county-level NielsenIQ sales. Column (1) replicates the baseline specification for reference. Column (2) adds a Mexican border county indicator; column (3) adds per-capita ICE arrests. To rule out anticipation effects and pre-existing trends, in column (4) we estimate a specification that includes both the contemporaneous immigration shock and its one-period lead, each instrumented by the corresponding predicted shock. Column (5) replaces separate state and period fixed effects with state-by-period fixed effects. Column (6) removes the top and bottom 1% winsorization applied to immigration and the immigration shock in all other columns. Column (7) replaces state fixed effects with county fixed effects. This last column does not include the county characteristic controls. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table B.7: County Price Regressions: Robustness to Alternative Population Specifications

	Δ $P^{\text{Sato-Vartia,Stacked}}$ (2006-2018)			
	(1) Imm / Pop'l Instr.	(2) Baseline	(3) + Init Pop × Year	(4) + Δ Native Pop
Immig/Pop	-1.120***			
	(0.252)			
Immigration		-0.058***	-0.052***	-0.055***
		(0.013)	(0.013)	(0.011)
dNative Pop				0.002***
				(0.000)
Obs.	6,339	6,249	6,249	6,247
First-Stage F-stat	9.3	85.5	75.7	83.2
Fixed Effects	State, Time	State, Time	State, Time	State, Time

Notes: This table reports estimates of immigration effects on county-level CPG prices (Sato-Vartia index). All specifications use stacked four-year differences (2006–2010, 2010–2014, 2014–2018), are weighted by initial county-level NielsenIQ sales, include state and period fixed effects, and cluster standard errors by state. Column (1) uses the immigrant-to-population ratio as the endogenous variable, with both the ratio and the instrument normalized by initial (2005) county population; controls include the urban indicator interacted with period dummies. Columns (2)–(4) use immigration levels (in thousands) as the endogenous variable, instrumented with baseline predicted shocks, and include controls for initial county GDP per capita and urban indicator, each interacted with period dummies. Column (3) additionally controls for initial (2005) county population interacted with period dummies, absorbing the concern that larger counties may exhibit differential price trajectories. Column (4) adds a control for the change in native population, constructed from the ACS as the change in total county population minus the total (European and non-European) immigration over the same period, using consistent survey weights throughout. This control separates the composition channel (markup compression from price-sensitive immigrants) from generic population scale effects; it is not instrumented. Immigration and the immigration shock are winsorized at the top and bottom 1% in all columns. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table B.8: Randomization Inference: Permuting Origin Labels in the Instrument

	Actual (anchor)		Placebo distribution				p-value	
	β	t	Mean	SD	5th	95th	1-sided	2-sided
Stacked reduced form	−0.080	−7.36	2.56	4.40	−5.88	8.10	0.019	0.115

Notes: Randomization-inference test of the stacked reduced form (1,000 draws). In each draw we permute the assignment of origin-country shifters to ancestry networks, holding the predicted-ancestry share matrix fixed, applied jointly across the three 4-year periods, and rebuild the instrument through the baseline pipeline (rolling ancestry vintages, per-period first-stage scaling, 1/99 winsorization, state and period fixed effects, state-clustered SEs, initial-sales weights); anchor and placebos differ only in the shifter-to-origin assignment. The anchor $t = -7.36$ reproduces the published stacked reduced form ($t = -7.23$) to within 2%, confirming the pipeline is calibrated to the estimator in the main text. The placebo distribution is centered right of zero because counties with concentrated ancestry exposure exhibit higher price growth regardless of which shifter is assigned—the same concentration documented in Table B.9—so the one-sided p-value is the relevant test given the directional (disinflationary) hypothesis.

Table B.9: Sensitivity to Origin Concentration: Rotemberg Weights and Dropping Mexico

Panel A. Goldsmith-Pinkham et al. (2020) decomposition (15 highest-weight origins)

Origin	α_o	$\hat{\beta}_o^{IV}$
Mexico	0.845	-0.036
China	0.057	-0.020
Philippines	0.029	-0.031
India	0.014	-0.042
Republic of Korea	0.012	-0.025
Puerto Rico	0.011	-0.012
Dominican Republic	0.007	-0.045
El Salvador	0.007	-0.021
Canada	0.005	-0.063
Japan	0.005	-0.050
Jamaica	0.005	-0.051
Guatemala	0.005	-0.026
Haiti	0.003	-0.040
Honduras	0.003	-0.028
Peru	0.002	-0.032
Sum (all origins)	1.000	
$\sum_o \alpha_o \hat{\beta}_o$		-0.035

Panel B. Dropping Mexico before the first stage

	Δ P$^{\text{SV,LongDiff}}$ (2006-2018)			
	(1)	(2)	(3)	(4)
Immigr$_{d,t}$	-0.029***	-0.035***	-0.036***	-0.044***
	(0.004)	(0.005)	(0.007)	(0.011)
Obs.	2,112	2,112	2,112	2,112
First-Stage F-stat			92.4	26.8
Sample	All	No Mex	All	No Mex
Fixed Effects	State	State	State	State
Method	OLS	OLS	IV	IV

Notes: Both panels analyze the long-difference IV estimator from Table B.5, column (5), based on 2005 predicted ancestry. **Panel A** reports the Goldsmith-Pinkham et al. (2020) decomposition: for each origin o, α_o is the Rotemberg weight and β_o^{IV} the just-identified estimator using only origin o's contribution; contributions are scaled within county so $\sum_o z_{(o,d)}$ reproduces the winsorized aggregate IV, and $\sum_o \alpha_o \beta_o$ equals the aggregate coefficient (up to rounding). Although Mexico carries a large weight, the β_o^{IV} are negative and tightly clustered across origins. **Panel B** drops Mexico before the first-stage prediction: columns (1)–(2) are OLS, (3)–(4) IV; (1) and (3) reproduce the published baseline and (2) and (4) the drop-Mexico case. The aggregate estimate is essentially unchanged. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.10: County Price Regressions: Does Past Inflation React to Contemporaneous Immigration?

	Δ P (BLS, Nondurables)		
	(1) Baseline	(2) Pre-Period Placebo	(3) Pre-Period Placebo
$\text{Immigr}_{d,t}$	-0.070***		
	(0.020)		
Immigration (t+5)			0.023
			(0.015)
Immigration (t+15)		0.022	
		(0.016)	
Obs.	741	693	1,189
Fixed Effects	State, Time	State, Time	State, Time
Method	IV	IV	IV

Notes: This table reports pre-trends and placebo specifications for the main price regressions. All columns use BLS regional nondurables price changes as the dependent variable, where pre-period prices are available. All specifications are IV regressions with state and time fixed effects, controls for county urbanization and 2006 GDP per capita interacted with year, and observations weighted by 2005 county population. Column 1 reports the baseline specification matching contemporaneous prices and immigration over 2006–2018 (Column 2 of Table 5 - which examines the sectoral BLS regional data). Column 2 is the central placebo: BLS nondurables price changes during 1990–2005 are regressed on immigration flows realized fifteen years later (2006–2018). Column 3 pools the pre- and post-periods (1990–2018), entering immigration just one period ahead. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Summary of Tables Table B.4 reports the baseline specification, but with price indices constructed with various aggregation methods. Using Laspeyres, Paasche, and Tornqvist weights results in remarkably similar results to Sato-Vartia, implying that effects are not due to major compositional changes in the basket of goods. The Jevons index is an unweighted average of prices and thus more likely to be driven by small outliers.

Table B.5 displays the results with a long-difference specification, an alternative to our baseline "stacked" specification. Prices reflect the 12-year difference in the chained price index by county, while immigration now accounts for all new immigrants by county from 2006 to 2018. Thus, the variation is strictly cross-sectional; relative to the stacked design it reduces reliance on period-by-period timing assumptions. The disinflationary effect of immigration persists, with similar quantitative implications.

Table B.6 includes miscellaneous specifications with further controls. Column (2) controls for proximity to the Mexican border, and Column (3) controls for ICE arrest activity; the immigration coefficient remains stable across both. In column (4), to test for anticipation and pre-existing trends, we estimate a specification that includes both the contemporaneous immigration shock and its one-period lead, each instrumented by the corresponding predicted shock. This isolates whether leads of the treatment predict prices conditional on the contemporaneous shock, which we can confirm is not the case. Only contemporaneous immigration is negative and statistically significant. With 2 instruments, the F-stat predictably drops. Column (5) replaces state and time fixed effects with state-by-year fixed effects, yielding a somewhat smaller but still significant coefficient. Column (6) reports results *without winsorization*, and Column (7) includes county fixed effects. The negative effect of immigration on prices is robust across all specifications.

Table B.7 presents robustness checks for the county-level price regressions. Column (1) uses

the immigrant-to-population ratio as the independent variable instead of the level of immigration. The instrument is re-constructed here following the baseline methodology covered in Section 3.4, using immigrants relative to the initial county population. The coefficient of -1.12 implies that a 1 percentage point increase in the immigrant share lowers CPG price growth by approximately 1.12 percentage points. Among existing studies that find negative price effects, Lach (2007) reports a 0.5 percentage-point decrease for a 1 percentage-point increase in the immigrant-to-native ratio in Israel; the smaller magnitude may reflect the short-run nature of the refugee shock he studies. Cortes (2008) uses a log-log specification with the share of low-skilled immigrants in the labor force, and finds that a 10 percent increase reduces prices of immigrant-intensive services by 4.8 to 6.3 percent; her estimates capture a supply-side wage channel in services where immigrants are concentrated, which differs from the broad CPG demand composition channel we study. Zachariadis (2012) also uses a log-log specification but with the immigrant-to-employment ratio across countries, and finds that a 10 percent increase reduces prices by about 2.9 percent. Despite these differences in specification, sector, and context, the magnitudes are broadly comparable.

We then also check whether the baseline specification is sensitive to population controls. In Column (3) we control for the initial (2005) population levels, and the interaction with time period indicators. This is to check whether the disinflationary effects of immigration are related to the size of the county (for example because larger counties will get more immigrants in levels). The baseline immigration coefficient remains quantitatively unchanged, suggesting our instrument is already handling this. The last column further shows that the price effect is unchanged when controlling for native population changes, a robustness check akin to Lach (2007). This is important for a few reasons. First, native displacement does not seem to be a major factor. Second, although only suggestive because it does not necessarily pick up the change in the total population, there is evidence that the disinflationary effect operates through the composition of population growth — the arrival of price-sensitive consumers — rather than scale. Consistent with the channels highlighted in the model in Section 6, more people through a larger native population raises the price level as evidenced by the coefficient on native population change.

Table B.8 reports a randomization-inference test that permutes the assignment of origin-country shifters to ancestry networks, holding the predicted-ancestry share matrix fixed, and rebuilds the instrument through the same pipeline as the baseline for each of 1,000 draws. The actual reduced-form estimate ($t = -7.36$, reproducing the published stacked reduced form of -7.23) falls in the 1.9% left tail of the placebo distribution, so only 19 of 1,000 random shifter assignments produce an effect as disinflationary as the one we estimate. Because concentrated ancestry exposure is associated with higher price growth regardless of which shifter is attached to it—the same origin concentration documented in Table B.9—the placebo distribution sits to the right of zero, and it is the specific pairing of shifters to ancestry networks that generates the strongly negative estimate.

Table B.9 reports the Goldsmith-Pinkham et al. (2020) decomposition of the long-difference IV estimator from Table B.5, column (5). For each origin o, α_o is the Rotemberg weight and β_o^{IV} is the

just-identified IV estimator using only origin o's contribution to the predicted immigration shock. By construction, $\sum_o \alpha_o \beta_o$ equals the aggregate IV coefficient (with the very slight difference due to rounding). Although Mexico receives a very large weight, we find that the betas are consistent across all origins. In fact, Panel B of the same table confirms that dropping Mexico from the analysis altogether (before the first stage prediction) does not affect the quantitative results at all (though the F-stat drops as expected). The fact that betas are very tightly distributed across origins is why the immigration shocks from all origins excluding Mexico still produce a very similar disinflationary effect.

Table B.10 conducts a robustness check to check for pre-trends. The goal is to confirm that past inflation does not respond to contemporaneous immigration, which we do in several ways. Since the baseline CPG prices are not available in the "pre-period" (before 2006), we instead rely mostly on the non-durable price index available for 27 metro areas from the BLS (as introduced in the main text). This price index responds to immigration in the most similar way from 2006-2018 (see Table 5 in the main text) – in fact we reproduce that result in Column (1), using the baseline specification with contemporaneous prices and immigration from 2006-2018. Column (2) produces the most natural robustness test: we lag prices by close to 15 years to produce regional inflation rates from 1990-1995, 1995-2000, and 2000-2005 (to match the ACS sample periods). Then we regress the past inflation rate to contemporaneous immigration, using our baseline sample of stacked immigration shocks from 2006-2018. Importantly, past inflation does not respond in any meaningful way to contemporaneous immigration. Columns (3) uses the full 1990-2018 sample for prices and immigration in order to show that if we regress inflation on a *one* period "lead" of immigration over this full sample, past inflation once again does not respond to future immigration. Note also that column (4) in Table B.6 controls for future immigration to test for pre-trends in a specification with CPG prices.

Table B.11: Alternative Rent Prices from ACS

	Δ P$^{\text{Rent,Stacked}}$ (2006-2018)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Immigr$_{d,t}$	0.070***	0.075***	0.083***	0.065***	0.078**	0.100***	0.014
	(0.019)	(0.020)	(0.017)	(0.023)	(0.033)	(0.036)	(0.035)
Obs.	6,336	6,336	6,246	6,336	6,336	6,246	669
First-Stage F-stat				23.7	49.8	45.8	157.7
Fixed Effects	Time	State, Time	State, Time	Time	State, Time	State, Time	State, Time
Method	OLS	OLS	OLS	IV	IV	IV	IV
Controls			Yes			Yes	Yes
Sample							Cities

Notes: Rent prices are reported median gross rents, variable NH063A in the NHGIS dataset, downloaded through IPUMS. Standard errors are clustered by state except the last column, where the robust standard errors are reported. Both immigration shock and immigration are winsorized by top and bottom 1%. Controls (columns 3, 6, and 7): 2006 county GDP per capita and urban indicator, each interacted with year indicators. Columns (1)-(3) are OLS and columns (4)-(7) are IV; columns (1) and (4) include time fixed effects only, while all other columns include state and time fixed effects. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table B.12: Aggregate Price Indices: BLS Cities and HNNS State

	Δ P (BLS City,All Items)	Δ P (HHNS State)		
	(1)	(2)	(3)	(4)
$\text{Immigr}_{d,t}$	-0.035*	0.007	-0.022**	
	(0.020)	(0.005)	(0.011)	
$\text{Immigr}_{s,t}$ (State)				-0.002
				(0.002)
Obs.	741	6,339	6,339	102
First-Stage F-stat	99.1	77.1	204.2	139.7
Fixed Effects	State, Time	State, Time	Time	Time
Controls	Yes	Yes	Yes	Yes

Notes: This table reports aggregate prices at the city-level and state-level, from the BLS and data shared by Hazell et al. (2022) respectively. The outcome in the first column is the "All Items" index from the BLS city-level CPI data (27 metro areas), which we merge to the county level. In the last 3 columns, the outcome is the state-level inflation rate for tradable industries produced by Hazell et al. (2022). Annual inflation data is transformed to 4-year rates. In columns (2) and (3) we maintain the aggregation at the county level, with prices only varying by the state. Column (2) includes state and time fixed effects, but in Column (3) we include only year fixed effects since prices vary by state. In the last column we aggregate immigration to the state level, where the shock is reconstructed so that the first-stage is aggregated by state-year. Controls: 2006 county GDP per capita and urban indicator, each interacted with the year indicators. At the state level we drop the urban control. Columns (2) and (3) report standard errors clustered by states, while columns (1) and (4) report heteroscedastic-robust SEs due to small numbers of states. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table B.11 turns from CPG prices to rents, using median gross rents from the ACS (NHGIS). In contrast to the disinflationary effect on retail goods, immigration *raises* rents. The IV estimate is positive and significant across specifications, from 0.065 to 0.100 in columns (4)–(6), consistent with Saiz (2007). This opposite sign confirms that our result is specific to searchable retail goods rather than a mechanical feature of high-immigration counties. The effect is smaller and insignificant in the smaller cities sample (column 7).

Table B.12 checks whether the disinflationary effect shows up in coarser, more aggregate price data. Column (1) uses the BLS city-level "All Items" CPI for 27 metro areas and yields a negative effect (-0.035). Columns (2)–(4) use state-level tradable inflation from Hazell et al. (2022). The estimate is negative when prices vary at the state level with year fixed effects (-0.022 in column 3), and close to zero once state fixed effects absorb most of the variation (column 2) or immigration itself is aggregated to the state level (column 4). These indices are far coarser than our scanner data, so power is limited, but the sign lines up with the CPG results and indicates the effect is not unique to the retail CPG index.

Table B.13: Immigration Effects on Wages: Robustness

Panel A: 2006-2018						
	Level Wage		Log Wage		Retail Wage	
	(1)	(2)	(3)	(4)	(5)	(6)
$\text{Immigr}_{d,t}$	0.391**	0.024	0.015	-0.003	0.001	-0.017
	(0.147)	(0.015)	(0.023)	(0.002)	(0.141)	(0.015)
Obs.	6,249	6,249	6,249	6,249	6,249	6,249
First-Stage F-stat	45.9	43.8	45.9	43.8	45.9	43.8
Fixed Effects	State, Time	State, Time	State, Time	State, Time	State, Time	State, Time
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Winsorized	Yes	No	Yes	No	Yes	No
Panel B: Education Heterogeneity (2006-2018)						
$\text{Immigr}_{d,t}$	-1.838*	-0.415***	-0.310*	-0.044***	-0.693	-0.485
	(0.959)	(0.050)	(0.183)	(0.006)	(0.979)	(0.393)
$\text{Immigr}_{d,t} \times$ EducYears	2.828**	0.675***	0.392*	0.060***	0.738	0.693
	(1.156)	(0.067)	(0.201)	(0.010)	(1.211)	(0.570)
Obs.	6,245	6,245	6,245	6,245	6,245	6,245
Fixed Effects	State, Time	State, Time	State, Time	State, Time	State, Time	State, Time
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Winsorized	Yes	No	Yes	No	Yes	No
Panel C: 1990-2005						
$\text{Immigr}_{d,t}$	1.209***	0.117***	0.158***	0.013***	0.106*	0.002
	(0.151)	(0.029)	(0.027)	(0.004)	(0.060)	(0.005)
Obs.	9,155	9,155	9,155	9,155	8,908	8,908
First-Stage F-stat	75.9	169.8	75.9	169.8	76.2	170.9
Fixed Effects	State, Time	State, Time	State, Time	State, Time	State, Time	State, Time
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Winsorized	Yes	No	Yes	No	Yes	No

Notes: Table reports IV estimates of immigration effects on county-level wages. All specifications instrument immigration with predicted shocks from equation (3.5) and include state and period fixed effects and county controls. Odd-numbered columns (1), (3), (5) winsorize the immigration regressor; even-numbered columns (2), (4), (6) do not. Panel A uses the 2006-2018 sample. Columns (1)-(2) report effects on average annual real wages in levels (in hundreds of 2012 dollars); columns (3)-(4) report effects on log wages; columns (5)-(6) report effects on retail sector wages. Panel B reports education heterogeneity for the 2006-2018 sample, interacting immigration with average years of education among immigrants following Terry et al. (2024) (immigrants restricted to those aged 25 and older); both immigration and its interaction with education are treated as endogenous and instrumented by the shift-share instrument and its interaction with immigrant education. Panel C replicates Panel A for the 1990-2005 sample for comparison with Terry et al. (2024). Standard errors clustered by state. ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

Appendix Table B.13 provides robustness checks for the wage results reported in Table 6. Panel A examines sensitivity to winsorization and the choice of log versus level specification for the 2006-2018 sample. The positive wage effect in levels is largely driven by winsorization, with non-winsorized estimates close to zero and more comparable to Terry et al. (2024). Log wage specifications yield estimates close to zero regardless of winsorization. Panel B shows that the education heterogeneity result—wages increasing more in counties receiving higher-educated immigrants—is robust across all specifications and winsorization choices. Panel C replicates the analysis for the 1990-2005 sample, showing substantially larger and more robust positive wage effects, consistent with Terry et al. (2024).

Table B.14: Price, Revenue, and Quantity: Branded versus Private Label

	Price	Revenue	Quantity
Panel A. RMS (store sales, county, weighted)			
Branded (GS1)	-0.065***	-0.001	0.064
	(0.012)	(0.056)	(0.063)
Private label	-0.050***	0.135**	0.185***
	(0.018)	(0.059)	(0.067)
Branded + priv. label	-0.066***	0.017	0.082
	(0.011)	(0.055)	(0.062)
Panel B. HMS (consumer panel, state, unweighted)			
Branded	-0.080***	0.166	0.246
	(0.029)	(0.191)	(0.208)
Private label	-0.091**	0.193	0.284
	(0.042)	(0.186)	(0.203)
Branded + priv. label	-0.079***	0.180	0.259
	(0.027)	(0.190)	(0.207)

Notes: Price is the Sato–Vartia (RMS) or projected (HMS) log price index. In RMS, quantity is the index-form residual $Q = \Delta \ln \text{Revenue} - \Delta \ln \text{Price}$, so revenue equals price plus quantity by construction. In HMS, quantity is the projected physical quantity recorded in the consumer panel, so price, revenue, and quantity are projected separately and need not satisfy this identity exactly. All variables are $\times 100$, with standard errors clustered by state in parentheses. Each row is a separate IV regression on stacked 4-year differences, with state and year fixed effects. Panel A is the NielsenIQ Retail Scanner census (RMS) at the county level ($N \approx 6{,}300$, base-sales weighted, the paper's baseline); Panel B is the NielsenIQ Consumer Panel (HMS, Homescan) at the state level ($N = 147$, unweighted, with the population projection factor applied in units). Within each panel, branded (GS1) products are manufacturer brands linked via GS1 prefix, private label are store brands, and branded plus private label pools both, all on the 14-year balanced product sample. $p < 0.10$, $p < 0.05$, $p < 0.01$.

Table B.14 reports price, revenue, and quantity responses in two datasets and for three product scopes each. Panel A uses the NielsenIQ Retail Scanner census (RMS) at the county level, weighted by base-year (2006) sales, with state and year fixed effects. Price is the Sato–Vartia index, revenue is log store sales, and quantity is the implied real quantity, revenue deflated by the price index ($\Delta \ln Q = \Delta \ln \text{Revenue} - \Delta \ln \text{Price}$). The three rows draw on separately constructed price files: branded goods are the products carrying a GS1 manufacturer prefix, private label are store brands with no GS1 prefix, and the pooled row combines the two. All three use the fully balanced product sample (present in all fourteen years, 2006–2019). Panel B uses the NielsenIQ Consumer Panel (Homescan, HMS) at the state level, the finest level at which its projection weights are valid, with year fixed effects. Here price, revenue, and physical quantity come directly from the projected panel aggregates, split into branded and private label the same way, with the pooled row a base-expenditure-weighted combination. In both panels immigration is instrumented by the shift-share prediction, winsorized at 1% over the full county sample, on stacked four-year differences, with standard errors clustered by state. For HMS, the county shock and its instrument are first aggregated to the state level using base-period expenditure weights.

Prices fall in every row. In RMS, real quantity for branded goods and for the pooled sample rises but is not significant (0.064 and 0.082), while the private-label quantity increase is large and significant (0.185, with revenue also up 0.135), pointing to a reallocation of spending from branded

Table B.15: Immigration Lowers Prices in Every Sample

	Price
RMS (county, weighted) — product scope	
Branded (GS1)	-0.065***
	(0.012)
Private label	-0.050***
	(0.018)
Branded + priv. label	-0.066***
	(0.011)
RMS (county, weighted) — balance window (GS1)	
RW2	-0.053***
	(0.015)
RW6	-0.065***
	(0.013)
RW10	-0.066***
	(0.013)
RW14	-0.065***
	(0.012)
HMS (state, unweighted) — product scope	
Branded	-0.080***
	(0.029)
Private label	-0.091**
	(0.042)
Branded + priv. label	-0.079***
	(0.027)
HMS (state, unweighted) — retailer/product restriction	
Whole market	-0.071**
	(0.031)
Balanced products	-0.111***
	(0.029)
RMS-tracked channels	-0.084***
	(0.032)
RMS firm sample	-0.057
	(0.037)
Big retailers	-0.077**
	(0.035)
RMS-tracked, large chains	-0.091***
	(0.032)

Notes: Each row is a separate IV regression of the price index on the immigration shock, instrumented by its shift-share prediction, winsorized at 1% over the full county sample, on 4-year stacked differences (2006–10, 2010–14, 2014–18) with state and year fixed effects; standard errors clustered by state in parentheses. RMS is the NielsenIQ Retail Scanner census (store-level), estimated at the county level and base-sales weighted (the paper's baseline weighting); HMS is the NielsenIQ Consumer Panel (Homescan household panel), aggregated to the state level, where its population projection factors are valid, and unweighted. Branded (GS1) products are manufacturer brands linked to firms via their GS1 prefix; private label are store brands; branded plus private label pools both. For the balance window, RW2-RW14 keeps products present in at least k consecutive years of 2006–2019: RW14 is fully balanced (present all 14 years; the paper's baseline product sample, \$44B of sales) and RW2 the loosest (\$139B). The HMS retailer and product restrictions tighten from the whole consumer market toward the RMS scanner scope: whole market uses all products and retailers in the panel; balanced products keeps continuing products only; RMS-tracked channels restricts to the retail channels the scanner census covers; RMS firm sample to the GS1–NETS matched manufacturers; big retailers to the largest tracked chains (the top 88, about 80% of panel spend); and RMS-tracked, large chains is the intersection of the last two, closest to the RMS scanner scope. The RMS branded estimate (-0.065) matches the county baseline in Table 3. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

toward private-label goods rather than an expansion of the overall market. HMS shows the same pattern with a smaller branded–private-label gap and, given its far smaller sample, noisier estimates that are individually insignificant.

Table B.15 collects the price effect across sixteen samples, grouped in four blocks that vary one construction choice at a time. The estimator is held fixed throughout: IV on the shift-share instrument, full-sample 1% winsorization, stacked four-year differences, standard errors clustered by state. The first block varies the RMS *product scope*, reusing the branded, private-label, and pooled price files from Table B.14, county level and sales-weighted. The second block varies the RMS *balance window* for branded goods: RWk keeps products present in at least k consecutive years of 2006–2019, so RW14 is the fully balanced baseline (about \$44B of sales) and RW2 the loosest (about \$139B), with RW6 and RW10 in between. Widening the window trades product comparability for coverage. The third block repeats the HMS product-scope split (branded, private label, pooled) at the state level. The fourth block is an HMS *restriction ladder* that progressively narrows the consumer panel toward the RMS scanner universe: "whole market" uses all products and retailers in the panel; "balanced products" keeps only continuing products; "RMS-tracked channels" restricts to the retail channels the scanner covers; RMS "firm sample" further restricts to the GS1–NETS matched manufacturers; "big retailers" keeps the largest tracked chains (the top 88, about 80% of panel spend); and RMS-tracked, "large chains" is the intersection of the last two, the closest HMS analogue to the RMS scanner scope.

The price coefficient is negative in all sixteen samples and significant in fifteen, ranging from -0.05 to -0.11. The single insignificant estimate (-0.057, the HMS GS1–NETS firm sample) is still negative and close to the baseline. The RMS branded estimate (-0.065) matches the county baseline in Table 3, and the HMS restriction ladder converges toward that RMS number as the panel is narrowed to the scanner's scope, so the two datasets agree where they overlap. That the decline also holds in the broad HMS samples, which include the discount, independent, and ethnic retailers the scanner census misses, shows the result is not an artifact of the scanner's retailer coverage.

B.4 Firm-Level Analysis

Table B.16: Economic Magnitude of the Demand and Supply Price Effects

	Demand	Supply
Effect of a one-standard-deviation increase	-2.664***	0.886***
	(0.748)	(0.265)
Effect of a p90–p10 move	-5.249	1.310
Raw coefficient (Table 7 units)	-0.071	0.016

Notes: Economic magnitudes based on the firm-level demand-and-supply price specification in Table 7, column (1). The first row reports the effect of a one-standard-deviation increase in each exposure (with the standard error from the corresponding standardized just-identified IV in parentheses); the second row the effect of moving a firm from the tenth to the ninetieth percentile of exposure; the last row the raw coefficient in the units of Table 7. Across-firm exposure standard deviations are 37.5 (demand) and 53.9 (supply), and tenth-to-ninetieth-percentile spreads are 73.8 and 79.7; all effects are in percentage points. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Shift-share robust inference. The firm-level demand and supply exposures are exposure-weighted aggregates $\mathrm{I}_f^{\text{demand}} = \sum_c \omega_{fc}^{\text{sales}} \mathrm{I}_c$ and $\mathrm{I}_f^{\text{supply}} = \sum_c \omega_{fc}^{\text{empl}} \mathrm{I}_c$ of the county immigration shock I_c. Following Borusyak et al. (2022), we compute exposure-robust ("shift-share") standard errors by residualizing the outcome and the endogenous variable on all controls and fixed effects, aggregating to the county level with the firm-county exposure shares, and running a county-level IV on the raw shock I_c. Table B.17 reports the result. The exposure-robust standard errors are smaller than or comparable to the SIC-3 clustered errors in every specification (ratios 0.36–0.87), and the point estimates reproduce the firm-level 2SLS exactly for the single-instrument specifications and to within 0.001 for the joint demand-and-supply specification—a just-identified one- versus two-instrument difference rather than an aggregation error. Consistent with the identification resting on historical ancestry shares rather than the contemporary shifters, industry clustering is the conservative choice.

Table B.17: Shift-Share Robust Standard Errors (BHJ) for the Firm-Level Specifications

	Coefficient	Industry-Clustered SE	BHJ Robust SE
Panel A. Single-instrument specifications			
Demand only (Col. 2)	-0.057	(0.022)	(0.010)
Supply only (Col. 3)	0.007	(0.005)	(0.004)
Panel B. Joint demand + supply (Col. 1)			
Demand	-0.071	(0.020)	(0.011)
Supply	0.016	(0.005)	(0.004)

Notes: The four Table 7 specifications (sales-weighted, SIC-3 and year fixed effects). The coefficient and the industry-clustered standard error are computed directly; the "BHJ Robust SE" is the exposure-robust standard error of Borusyak et al. (2022), obtained by residualizing the outcome and endogenous variable on all controls and fixed effects, aggregating to the county level with the firm-county exposure shares, and running county-level IV on the raw county shock. The BHJ coefficient reproduces the firm-level 2SLS estimate exactly in Panel A and to within 0.001 in Panel B. In every specification the BHJ standard error is no larger than the clustered one, so industry clustering is conservative.

Table B.18: UPC-to-Firm Aggregation: Sato–Vartia and Törnqvist vs. the Original Index

	(1) UPC→Firm SV	(2) UPC→Firm Tornqvist	(3) Original ($\Delta \ln \mathrm{SV}^{f,p}$)
$\text{Immigration}_{\text{f,d,t}}$ (D)	-0.0626***	-0.0596***	-0.0711***
	(0.0188)	(0.0188)	(0.0200)
$\text{Immigration}_{\text{f,d,t}}$ (S)	0.0129**	0.0133**	0.0164***
	(0.0050)	(0.0056)	(0.0049)
Obs.	11,037	11,037	11,037
First-Stage F	3,743.4	3,743.4	3,743.4

Notes: Firm-level regressions (N = 11,037). Columns (1)–(2) aggregate UPC-year prices to the firm via Sato–Vartia and Törnqvist chain indices over continuing UPCs; column (3) is the paper's original firm-level Sato–Vartia index. The three are statistically indistinguishable, so the firm-level baseline is robust to the UPC-to-firm aggregation. Firm-level demand (D) and supply (S) shift-share IV; SIC-3 and year fixed effects; sales-weighted; standard errors clustered by SIC-3 in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

UPC-level robustness. Our baseline firm-level price index is built through a multi-stage, county-nested procedure: within each county we aggregate barcode (UPC) prices into a firm price, and only then aggregate across the counties where the firm sells (Sato-Vartia weighted by its county sales). Two exercises confirm that the demand-side result is not an artifact of this construction and is present at the barcode level.

Table B.18 rebuilds the firm-level price change by a deliberately simpler and different route. Rather than nesting barcode prices through the county dimension, it collapses each barcode to a single price per year (pooled across stores) and aggregates those barcode prices *directly* to the firm using two off-the-shelf chain indices—Sato–Vartia and Törnqvist—over the firm's continuing barcodes, on the same firm sample as the baseline (N = 11,037). This discards the county nesting and the custom weighting of the baseline entirely, yet the demand-side coefficient is essentially unchanged: −0.063 (Sato–Vartia) and −0.060 (Törnqvist), against the baseline −0.071. The firm-level result is therefore not driven by how we construct the index.

Table B.19 instead estimates the price regression directly at the barcode level (N = 120,906 barcode-periods), with each barcode's price change measured across the markets in which it sells and the treatment still the firm's demand and supply exposure; the demand-side coefficient is −0.113 when barcodes are weighted by their firm's sales and −0.064 when weighted by their own revenue, bracketing the firm-level estimate. Because both regressions track prices of identical continuing barcodes, the decline cannot come from consumers reallocating toward cheaper products or from quality change, and instead reflects firms lowering markups on the same goods—the demand-side mechanism the paper emphasizes.

Table B.19: Within-UPC Prices: Chain-Weighted Across-County Aggregation

	Firm-Weighted		UPC-Weighted	
	(1) SV	(2) Tornqvist	(3) SV	(4) Tornqvist
$\text{Immigration}_{f,d,t}$ (D)	-0.1126***	-0.1134***	-0.0639***	-0.0642***
	(0.0343)	(0.0342)	(0.0190)	(0.0190)
$\text{Immigration}_{f,d,t}$ (S)	0.0287***	0.0288***	0.0141***	0.0141***
	(0.0108)	(0.0108)	(0.0044)	(0.0044)
Obs.	120,906	120,906	120,906	120,906
First-Stage F	779.4	779.4	3,723.8	3,723.8

Notes: UPC×period observations ($N = 120{,}906$). UPC-year prices are built from UPC×county×year prices by chain-weighted aggregation across counties (Sato–Vartia and Törnqvist log-mean county sales-share weights), preserving within-county price variation; firm-weighted (1)–(2) and UPC-weighted (3)–(4). The coefficients bracket the firm-level baseline of −0.071, confirming that UPC-level identification requires county-preserving aggregation. Firm-level shift-share IV; SIC-3 and year fixed effects; clustered by SIC-3. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.20: Firm-Level Prices: Demand and Supply Exposure, Long-Difference

	Δ $\ln P^{\text{SV,Long-Difference}}$ (2006-2018)			
	(1)	(2)	(3)	(4)
$\text{Immigration}_{f,d,t}$ (D)	-0.052***	-0.052***		-0.043**
	(0.018)	(0.019)		(0.018)
$\text{Immigration}_{f,d,t}$ (S)	0.014**		0.001	0.018***
	(0.006)		(0.006)	(0.006)
Δ lnW (Retail, D)				-0.129
				(0.264)
Δ lnP (Rent, D)				-0.335
				(0.297)
Δ lnW (Retail, S)				0.308**
				(0.148)
Δ lnP (Rent, S)				-0.228
				(0.231)
Obs.	3,629	3,629	3,629	3,629
First-Stage F-stat	3,152.4	3,494.1	6,156.1	1,664.2
Fixed Effects	Industry	Industry	Industry	Industry
Method	IV	IV	IV	IV
Controls				Yes

Notes: The regression is weighted by the initial sales. Missing shares are controlled. Standard errors are clustered by the industry; the industry is the initial 3-digit SIC code. Controls: 2006 county GDP per capita and urban indicator, each weighted by either initial sales share or employment share across counties within firm, and interacted with the year indicators. ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

Long-Difference Table B.20 re-estimates the firm-level demand-and-supply specification of Table 7 on a single 12-year long difference (2006–2018) rather than stacked four-year differences. Demand-side

exposure again lowers firm prices (−0.052) and supply-side exposure modestly raises them (+0.014), with the same joint-identification pattern as the baseline: the supply effect collapses to zero when the demand channel is omitted (column 3). The long-difference estimates are slightly smaller than the stacked baseline but qualitatively identical, so the result does not hinge on the period-by-period timing of the stacked design.

Table B.21: Firm-Level Prices: Robustness Using International Controls and Sub-Samples

	(1) Baseline	(2) +Exp/Imp Controls	(3) +Intl Control	(4) Domestic Only	(5) Intl Interactions
$\text{Immigration}_{f,d,t}$ (D)	-0.057***	-0.059***	-0.059***	-0.064**	-0.064**
	(0.020)	(0.019)	(0.019)	(0.030)	(0.029)
$\text{Immigration}_{f,d,t}$ (S)	0.018***	0.019***	0.019***	0.017***	0.017***
	(0.004)	(0.004)	(0.004)	(0.005)	(0.006)
Obs.	11,037	11,037	11,037	6,810	11,037
First-Stage F	3733.2	4016.6	4120.9	1696.4	1172.5
Fixed Effects	Industry, Time	Industry, Time	Industry, Time	Industry, Time	IndustryxIntl, TimexIntl

Notes: The dependent variable is the change in the firm-level Sato-Vartia price index (in percentage points). All columns report IV estimates using the shift-share instrument, with industry (SIC-3) and year fixed effects, baseline controls (urban share, GDP per capita, sales/employment missingness shares, each interacted with year where applicable), and sales weights. Column (1) reproduces the baseline specification. Column (2) adds separate exporter and importer indicator variables – their levels and interacted with year fixed effects; column (3) replaces these with a single international indicator (equal to one if the firm is an exporter or importer) interacted with year fixed effects. Column (4) drops all international firms from the sample. Column (5) interacts the immigration exposure variables (both demand and supply) with the international indicator, instrumenting the interactions with the corresponding instrument interactions; industry and year fixed effects are also fully interacted with international status. The reported coefficients in column (5) correspond to the base group (domestic-only firms); interaction terms are estimated but not shown. The interaction terms are both statistically insignificant (p-values of approximately 0.80, well above conventional thresholds), so there is no evidence that prices change differentially for international firms. Standard errors clustered by SIC-3 industry are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

International Controls Table B.21 checks whether internationally active firms drive the result, since exporters and importers may be exposed to immigration through global rather than local channels. The demand-side coefficient stays between −0.057 and −0.064 (supply between +0.017 and +0.019) whether we control for exporter and importer status (columns 2–3), drop international firms from the sample altogether (column 4, $N = 6{,}810$), or interact both exposures with international status (column 5). The interaction terms (not reported) are statistically insignificant ($p \approx 0.80$), so the price response does not differ between international and domestic firms.

Retailer robustness Table B.22 tests whether the firm-level demand effect is an artifact of the retailer dimension along which prices are observed. The outcome is the firm-retailer-period Sato–Vartia price change ($N = 76{,}437$). We build the demand shock over three alternative *footprints*, in each case averaging the county immigration shock I_c with fixed base-year sales weights but changing the set of counties over which the average is taken: g_f uses the counties where the firm sells (its overall footprint, as in the baseline); g_r uses the counties where the retail chain operates; and g_{fr} uses the counties where the particular firm sells *through the particular chain*. Each is instrumented by its shift-share counterpart. Column (1) is the firm-shock baseline ($g_f = -0.063$). Column (2) adds the retailer shock: it is small and an order of magnitude below the firm shock ($g_r = -0.006$)

Table B.22: Firm, Retailer, and Firm-by-Retailer Immigration Shocks

	Δ $\ln P_{r,f,t}^{SV}$ (stacked)			
	(1)	(2)	(3)	(4)
$Immigration_{f,t}$ (D, firm)	-0.0629***	-0.0558***	-0.0531**	-0.0486**
	(0.0209)	(0.0211)	(0.0214)	(0.0211)
$Immigration_{r,t}$ (D, retailer)		-0.0060**		
		(0.0026)		
$Immigration_{f,r,t}$ (D, firm×ret)				-0.0261*
				(0.0144)
$Immigration_{f,t}$ (S)	0.0196***	0.0195***	0.0191***	0.0192***
	(0.0055)	(0.0056)	(0.0057)	(0.0057)
Obs.	76,437	76,437	76,437	76,437
First-Stage F	3,591.6	2,547.1	4,566.9	797.6
Fixed Effects	Ind, Time	Ind, Time	Ind, RetxTime	Ind, RetxTime
Demand shocks	f	f, r	f	f, fxr

Notes: Each observation is a firm × retail chain × period, on stacked four-year differences (2006–2010, 2010–2014, 2014–2018); the outcome is the change in the firm-retailer Sato–Vartia log price index (in percentage points). All columns are IV. The endogenous demand-side exposures are built on the same shift-share logic as the firm shock, differing only in the footprint over which the county immigration shock I_c is averaged with fixed base-year sales weights: g_f over the firm's sales counties (varying by firm × period), g_r over the retail chain's counties (retailer × period), and g_{fr} over the counties where the firm sells through that chain (firm × retailer × period); each is instrumented by its shift-share counterpart. The firm's supply-side exposure (S), built over its production counties, is included and instrumented in every column. Columns (1)–(2) include industry (3-digit SIC) and period fixed effects; columns (3)–(4) replace the period effects with retail-chain-by-period fixed effects, which absorb g_r (constant within chain-period) and identify g_{fr} off residual within-firm-retailer variation. The "Demand shocks" row indicates which exposures enter each column. All columns use the same control vector as the firm-level baseline (Table 7): demand- and supply-side exposure to county urban share and 2006 GDP per capita, each interacted with period; the sales- and employment-share of counties with missing shocks on the demand and supply sides; and firm-level changes in retail wages and median rents on both sides. Columns including g_r or g_{fr} additionally control for the corresponding missing-county sales share. Observations are weighted by fixed 2006 firm-retailer sales; standard errors are clustered by firm. The sample contains 58 retail chains (parent companies). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

while g_f barely moves (-0.056), so prices track the firm's sales-market exposure, not the retailer's. Column (3) replaces the period effects with retailer-by-period fixed effects, absorbing cross-retailer composition and anything else common to a chain within a period; g_f is unchanged (-0.053), so the effect is within firm-retailer rather than a retailer-composition artifact. Column (4) addresses a subtler concern: a factor operating jointly at the firm-retailer level. Even granting that prices respond to the firm's exposure and not the retailer's, the effect could still run through specific firm-chain combinations—for instance, if a manufacturer's products at a given chain happen to be concentrated in immigrant-heavy counties, or if chains reallocate shelf space across competing manufacturers in response to local demand. The shock g_{fr} captures exactly this: it is the firm's immigration exposure measured over its counties *at that chain*, and it differs from the firm's overall exposure g_f because a firm's county footprint varies across the chains it supplies. Adding g_{fr} (with g_r now absorbed by the retailer-by-period effects) leaves the firm coefficient unchanged (-0.049), so no joint firm-retailer factor drives the result; g_{fr} itself is at most marginal (-0.026, $p < 0.10$) and in the same direction as the demand channel. The interaction is weak largely by construction: retailer-by-year effects alone absorb 97% of the variation in g_{fr}, because the large national chains have nearly identical footprints across the manufacturers they carry, leaving little independent firm-retailer variation. Throughout, the supply-side coefficient is stable and positive ($+0.019$ to $+0.020$). Bringing the retailer dimension in explicitly thus leaves the firm demand effect essentially unchanged.

Table B.23: Firm-Level Price Decomposition: Long-Difference

	Δ lnC$^{\text{Long-Difference}}$ (2006-2018), Decomposition			
	(1) SV-Price	(2) Variety	(3) Price+Variety Corr.	(4) Appeal
Immigration$_{f,d,t}$ (D)	-0.043**	-0.001	-0.041**	-0.087***
	(0.018)	(0.006)	(0.018)	(0.028)
Immigration$_{f,d,t}$ (S)	0.018***	0.005***	0.013*	0.022**
	(0.006)	(0.001)	(0.007)	(0.011)
Obs.	3,629	3,354	3,354	3,354
First-Stage F-stat	1,664.2	1,616.9	1,616.9	1,616.9
Fixed Effects	Industry	Industry	Industry	Industry
Method	IV	IV	IV	IV
Controls	Yes	Yes	Yes	Yes

Notes: This table is the long-difference (2006–2018) analog of the firm-level price decomposition in Table 10, estimated over a single 12-year change rather than stacked four-year differences. All columns report IV estimates of firm-level demand-side (D) and supply-side (S) immigration exposure. Column (1) is the Sato-Vartia price index for continuing varieties; column (2) is the Feenstra variety correction (equation 5.7), multiplied by -1 so that positive values reflect variety gains; column (3) is the utility-based (CES-exact) price index combining the two, with the variety term entering at its natural sign; and column (4) is firm-level appeal (equation 5.10), capturing shifts in quality composition. The variety and appeal terms use the state- and time-varying demand elasticities of Section 5.4. All specifications are weighted by initial firm sales and control for firm exposure to county GDP per capita and the urban indicator, to retail wages and median rents (all on both the demand and supply sides), and to the sales- and employment-share of counties with missing shocks; they include industry (3-digit SIC) fixed effects. Standard errors are clustered by 3-digit SIC industry. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Long-Difference Decomposition Table B.23 re-estimates the structural price decomposition of Table 10 on the 12-year long difference. The demand-side pattern is unchanged: prices of continuing varieties fall (-0.043), the variety correction is negligible (-0.001), and firm appeal declines sharply (-0.087)—so the demand channel operates through lower continuing-good prices and a shift in sales composition toward lower-quality products, not through variety. Supply-side exposure again raises prices ($+0.018$), variety ($+0.005$), and appeal ($+0.022$), consistent with the quality-upgrading interpretation in the main text.

Appendix C Structural Estimation: Details

This appendix provides details on the estimation and derivation of the structural decomposition in Section 5.4.

C.1 Elasticity Estimation

We estimate demand elasticities following the methodology of Feenstra (1994) and Hottman et al. (2016), adapted to our county-level setting. A key departure from prior work is that we allow elasticities to vary across states and time periods, enabling us to examine how immigration affects demand responsiveness. This section provides the complete estimation framework.

C.1.1 Notation and Setup

We define the following indexing convention used throughout this appendix:

- g: product group
- s: state
- $c \in s$: county (where $c \in s$ denotes that county c belongs to state s)
- $\tau \in \{1, 2, 3\}$: estimation period, corresponding to 2006–2010, 2011–2014, and 2015–2018
- $t \in \tau$: year within period τ
- f: firm
- u: product (UPC)

In the main text (Section 5.4) we use ℓ to index a product group, state, and period. Throughout this appendix we write $\ell \equiv (g, s, \tau)$ explicitly.

Our key structural parameters are the elasticities of substitution:

- $\sigma^U_{gs\tau}$: within-firm elasticity (substitution across products within a firm) for product group g, state s, period τ
- $\sigma^F_{gs\tau}$: across-firm elasticity (substitution across firms within a product group) for product group g, state s, period τ

Consider a firm f producing a set of products $\Omega_{fgc\tau} \subseteq \Omega_{gc\tau}$ in product group g, county $c \in s$, and period τ. From the nested CES demand system in Section 5.4, the expenditure share of product u within firm f is:

$$S^U_{ugct} = \frac{(\phi^U_{ut})^{\sigma^U_{gs\tau}-1}(P^U_{uct})^{1-\sigma^U_{gs\tau}}}{\sum_{k\in\Omega_{fgc\tau}}(\phi^U_{kt})^{\sigma^U_{gs\tau}-1}(P^U_{kct})^{1-\sigma^U_{gs\tau}}} \tag{C.1}$$

where ϕ^U_{ut} is product appeal, P^U_{uct} is the price of product u in county c at year t, and $\sigma^U_{gs\tau} > 1$ is the within-firm elasticity of substitution for product group g in state s during period τ.

Taking logs:

$$\ln S^U_{ugct} = (\sigma^U_{gs\tau} - 1)\ln \phi^U_{ut} + (1 - \sigma^U_{gs\tau})\ln P^U_{uct} - \ln \Phi_{fgct} \tag{C.2}$$

where $\Phi_{fgct} \equiv \sum_{k\in\Omega_{fgc\tau}}(\phi^U_{kt})^{\sigma^U_{gs\tau}-1}(P^U_{kct})^{1-\sigma^U_{gs\tau}}$ is a firm-product group-county-year specific term.

On the supply side, firms set prices as a markup over marginal cost:

$$P^U_{uct} = \mu^F_{fgct} \cdot MC_{ugct} \tag{C.3}$$

where μ^F_{fgct} is the firm-level markup and marginal cost takes the form:

$$MC_{ugct} = a_{ugct} \cdot (q_{ugct})^{\delta_{gs\tau}} \tag{C.4}$$

with a_{ugct} a product-specific cost shifter and $\delta_{gs\tau} \geq 0$ governing returns to scale. We maintain throughout that the markup is common across a firm's products, so that it cancels in the double-differencing procedure below. Online Appendix D.2 solves the firm's problem without that restriction and bounds the dispersion it suppresses by the across-firm elasticity gap (Corollary 1); it is a maintained conduct assumption rather than an implication of arbitrary differentiated-product conduct.

C.1.2 Within-Firm Elasticity $\sigma^U_{gs\tau}$

The within-firm elasticity governs substitution across products (UPCs) within a firm. We estimate $\sigma^U_{gs\tau}$ separately for each state s, product group g, and period τ, exploiting heteroskedasticity across products following Feenstra (1994).

Double-Differencing. We apply double-differencing to eliminate unobserved heterogeneity. First, take time differences (within period τ) to remove time-invariant product characteristics:

$$\Delta_t \ln S^U_{ugct} = (\sigma^U_{gs\tau} - 1)\Delta_t \ln \phi^U_{ut} + (1 - \sigma^U_{gs\tau})\Delta_t \ln P^U_{uct} - \Delta_t \ln \Phi_{fgct} \tag{C.5}$$

Second, difference relative to a reference product $\tilde{u}$ (the largest UPC within each firm-county) to remove firm-county-year specific terms:

$$\Delta_{u,t} \ln S^U_{ugct} \equiv \Delta_t \ln S^U_{ugct} - \Delta_t \ln S^U_{\tilde{u}gct} \tag{C.6}$$

This yields the double-differenced demand equation:

$$\Delta_{u,t} \ln S^U_{ugct} = (1 - \sigma^U_{gs\tau})\Delta_{u,t} \ln P^U_{uct} + \omega_{ugct} \tag{C.7}$$

where $\omega_{ugct} \equiv (\sigma^U_{gs\tau} - 1)(\Delta_t \ln \phi^U_{ut} - \Delta_t \ln \phi^U_{\tilde{u}t})$ is the demand residual capturing changes in relative appeal.

For the supply side, taking logs of the pricing equation (C.3):

$$\ln P^U_{uct} = \ln \mu^F_{fgct} + \ln a_{ugct} + \delta_{gs\tau} \ln q_{ugct} \tag{C.8}$$

Double-differencing eliminates the firm-level markup (which is common across products within a firm):

$$\Delta_{u,t} \ln P^U_{uct} = \Delta_{u,t} \ln a_{ugct} + \delta_{gs\tau}\Delta_{u,t} \ln q_{ugct} \tag{C.9}$$

Using the accounting identity $S^U_{ugct} = P^U_{uct} q_{ugct} / E_{fgct}$, equivalently $\ln q_{ugct} = \ln S^U_{ugct} + \ln E_{fgct} - \ln P^U_{uct}$ (where E_{fgct} is firm expenditure, which cancels in double-differencing), and collecting the

$\Delta_{u,t} \ln P^U_{uct}$ terms:

$$\Delta_{u,t} \ln P^U_{uct} = \frac{\delta_{gs\tau}}{1+\delta_{gs\tau}} \Delta_{u,t} \ln S^U_{ugct} + \kappa_{ugct} \tag{C.10}$$

where $\kappa_{ugct} \equiv \frac{1}{1+\delta_{gs\tau}}(\Delta_t \ln a_{ugct} - \Delta_t \ln a_{\tilde{u}ct})$ is the supply residual capturing changes in relative costs.

Identification via Heteroskedasticity. The system of equations (C.7) and (C.10) contains two endogenous variables ($\Delta_{u,t} \ln S^U_{ugct}$ and $\Delta_{u,t} \ln P^U_{uct}$) and two structural parameters ($\sigma^U_{gs\tau}$ and $\delta_{gs\tau}$). OLS estimation of either equation is biased due to simultaneity.

Identification comes from heteroskedasticity across products within each (g, s, τ). The key assumption is:

$$\text{Cov}(\omega_{ugct}, \kappa_{ugct}) = 0 \quad \text{for all } c \in s,\ t \in \tau \tag{C.11}$$

That is, demand shocks (changes in relative appeal) are uncorrelated with supply shocks (changes in relative costs) after double-differencing. This is plausible because double-differencing removes firm-level shocks that might affect both appeal and costs; the remaining variation is idiosyncratic to individual products.

Under this assumption, different products trace out different "hyperbolas" in $(\sigma^U_{gs\tau}, \delta_{gs\tau})$ parameter space, depending on their relative variances of demand and supply shocks. The true parameters must lie on all hyperbolas, so their intersection identifies both $\sigma^U_{gs\tau}$ and $\delta_{gs\tau}$.

GMM Implementation. We implement estimation via GMM, separately for each (g, s, τ). From equations (C.7) and (C.10), the residuals are:

$$\omega_{ugct} = \Delta_{u,t} \ln S^U_{ugct} - (1 - \sigma^U_{gs\tau}) \Delta_{u,t} \ln P^U_{uct} \tag{C.12}$$

$$\kappa_{ugct} = \Delta_{u,t} \ln P^U_{uct} - \frac{\delta_{gs\tau}}{1+\delta_{gs\tau}} \Delta_{u,t} \ln S^U_{ugct} \tag{C.13}$$

Identification exploits heteroskedasticity *across products*: each product u supplies its own moment, and the cross-product dispersion of these moments pins down $(\sigma^U_{gs\tau}, \delta_{gs\tau})$. For product u, define the time-series average of the residual cross-product,

$$G_{ugc}(\sigma^U, \delta) \equiv \frac{1}{|\tau|} \sum_{t \in \tau} \hat{\omega}_{ugct}(\sigma^U)\, \hat{\kappa}_{ugct}(\delta). \tag{C.14}$$

The orthogonality assumption (C.11) implies the population moment $\mathbb{E}[G_{ugc}] = 0$ *for each product*, so stacking over the $N_{gs\tau}$ products in (g, s, τ) yields a moment vector $\mathbf{G}(\sigma^U, \delta) = \big(G_{ugc}\big)_{u \in \Omega_{gs\tau}}$ of dimension far exceeding the two parameters. We estimate by GMM,

$$(\hat{\sigma}^U_{gs\tau}, \hat{\delta}_{gs\tau}) = \arg\min_{(\sigma,\delta)} \ \mathbf{G}(\sigma, \delta)'\, W_{gs\tau}\, \mathbf{G}(\sigma, \delta), \tag{C.15}$$

with weight matrix $W_{gs\tau}$. Following Feenstra (1994), Broda and Weinstein (2006), and Hottman et al. (2016), each product's moment is a function of its own variances and covariance of $(\Delta_{u,t} \ln P^U_{uct},\ \Delta_{u,t} \ln S^U_{ugct})$; products with different variance ratios trace out distinct hyperbolas in (σ^U, δ) space, and their intersection identifies both parameters. In practice we implement this as the nonlinear cross-product regression of these authors over the UPCs in each (g, s, τ).

C.1.3 Across-Firm Elasticity $\sigma^F_{gs\tau}$

The across-firm elasticity governs substitution across firms within a product group. We estimate $\sigma^F_{gs\tau}$ separately for each (g, s, τ) using an instrumental variables approach.

Double-Differenced Demand. The expenditure share of firm f within product group g in county $c \in s$ is:

$$S^F_{fgct} = \frac{(\phi^F_{fgct})^{\sigma^F_{gs\tau}-1}(P^F_{fgct})^{1-\sigma^F_{gs\tau}}}{\sum_{k\in\Omega_{gc\tau}}(\phi^F_{kgct})^{\sigma^F_{gs\tau}-1}(P^F_{kgct})^{1-\sigma^F_{gs\tau}}} \tag{C.16}$$

Double-differencing (over time within period τ, and relative to the largest firm $\tilde{f}$ in each product group-county) yields:

$$\Delta_{\tilde{f},t} \ln S^F_{fgct} = (1-\sigma^F_{gs\tau})\Delta_{\tilde{f},t} \ln P^F_{fgct} + \omega_{fgct} \tag{C.17}$$

where $\omega_{fgct} \equiv (\sigma^F_{gs\tau}-1)(\Delta_t \ln \phi^F_{fgct} - \Delta_t \ln \phi^F_{\tilde{f}gct})$ captures changes in relative firm appeal.

Instrument Construction. OLS estimation of (C.17) is biased because firm prices are correlated with firm appeal shocks. We construct an instrument from the CES price index structure.

The firm price index is:

$$P^F_{fgct} = \left[\sum_{u\in\Omega_{fgc\tau}} (\phi^U_{ut})^{\sigma^U_{gs\tau}-1}(P^U_{uct})^{1-\sigma^U_{gs\tau}}\right]^{\frac{1}{1-\sigma^U_{gs\tau}}} \tag{C.18}$$

Changes in this index admit a Sato-Vartia decomposition into UPC price changes and a within-firm share-dispersion term:

$$\Delta_{\tilde{f},t} \ln P^F_{fgct} = \sum_{u\in\Omega_{fgc\tau}} \tilde{w}_{ugct}\, \Delta_{\tilde{f},t} \ln P^U_{uct} + \frac{1}{1-\sigma^U_{gs\tau}} \Delta_{\tilde{f},t} \ln \left(\sum_{u\in\Omega_{fgc\tau}} \frac{S^U_{ugct}}{\tilde{S}^U_{fgct}}\right) + r_{fgct} \tag{C.19}$$

where $\tilde{w}_{ugct}$ are Sato-Vartia weights, the second term is a Theil-like index of within-firm share dispersion ($\tilde{S}^U_{fgct}$ the geometric mean of within-firm shares), and r_{fgct} collects higher-order UPC-appeal terms that vanish under the normalization (C.25). This is a decomposition of the price *change*, not an exact levels identity: UPC appeal enters the CES index and is absorbed into r_{fgct} under the maintained normalization.

The instrument is the double-differenced share dispersion term:

$$Z_{fgct} = \Delta_{\tilde{f},t}\left[\frac{1}{1-\sigma^U_{gs\tau}}\ln\left(\sum_{u\in\Omega_{fgc\tau}}\frac{S^U_{ugct}}{\tilde{S}^U_{fgct}}\right)\right] \tag{C.20}$$

Instrument Validity. *Relevance*: Share dispersion enters the firm price index by equation (C.19), so Z_{fgct} is mechanically correlated with $\Delta_{\tilde{f},t}\ln P^F_{fgct}$.

Exclusion (identifying assumption): we assume within-firm share dispersion is orthogonal to the firm-level appeal shock, $\mathbb{E}[Z_{fgct}\,\omega_{fgct} \mid g, s, \tau] = 0$. Economically, Z_{fgct} is driven by the cross-UPC distribution of prices and sales within the firm, whereas ω_{fgct} is a firm-level appeal innovation common to all of the firm's products; a firm-wide appeal shock shifts the level of within-firm shares but not their dispersion. This is an assumption on the data-generating process—which we assess in robustness checks—rather than an implication of defining appeal as a price residual.

We estimate $\sigma^F_{gs\tau}$ by 2SLS:

$$\hat{\sigma}^F_{gs\tau}: \quad \Delta_{\tilde{f},t}\ln S^F_{fgct} = (1-\sigma^F_{gs\tau})\Delta_{\tilde{f},t}\ln P^F_{fgct} + \omega_{fgct}, \quad \text{instrumented by } Z_{fgct} \tag{C.21}$$

separately for each (g, s, τ), pooling across all counties $c \in s$ and years $t \in \tau$.

C.1.4 Implementation Details

Estimation Units. We estimate both elasticities separately for each state s, product group g, and period τ, yielding parameter estimates $\hat{\sigma}^U_{gs\tau}$, $\hat{\delta}_{gs\tau}$, and $\hat{\sigma}^F_{gs\tau}$. The three periods are:

- $\tau = 1$: 2006–2010
- $\tau = 2$: 2011–2014
- $\tau = 3$: 2015–2018

This disaggregated estimation allows elasticities to vary across states and over time, which is essential for our analysis of how immigration affects demand responsiveness.

Data Requirements. Our retail scanner data provides sufficient variation for this disaggregated estimation. While Hottman et al. (2016) pool across locations to estimate national elasticities by product group, we observe prices and quantities at the store level across approximately 30,000–50,000 participating retailers nationwide. The county-level detail in our data provides the variation needed for identification at the state-period level.

Sample Restrictions. To ensure reliable estimates, we impose the following minimum sample size requirements within each (g, s, τ):

- For $\sigma^U_{gs\tau}$: minimum of 50 product-county-year observations with at least 5 products per firm
- For $\sigma^F_{gs\tau}$: minimum of 30 firm-county-year observations with at least 3 firms per product group-county

Cells failing these thresholds are dropped from the analysis. In robustness checks, we verify that our main results are not sensitive to these specific cutoffs.

Using Estimated Elasticities. The estimated elasticities $\hat{\sigma}^U_{gs\tau}$ and $\hat{\sigma}^F_{gs\tau}$ enter the construction of price indices and welfare decompositions in Section 5.4. For county c belonging to state s, we use the elasticities estimated for that state:

$$\text{For county } c \in s: \quad \sigma^U_{gc\tau} \equiv \hat{\sigma}^U_{gs\tau}, \quad \sigma^F_{gc\tau} \equiv \hat{\sigma}^F_{gs\tau} \tag{C.22}$$

This assignment assumes that consumer substitution patterns are similar across counties within the same state, while allowing for state-level heterogeneity that may reflect differences in market structure, demographics, or retail environments.

C.2 Sato-Vartia Price Index

The Sato-Vartia price index aggregates price changes using log-mean expenditure shares as weights. For continuing products in $\Omega_{fgc,t\cap t-1} \equiv \Omega_{fgct} \cap \Omega_{fgc,t-1}$ (where $t, t-1 \in \tau$):

$$\Phi^{SV}_{fgct} = \prod_{u \in \Omega_{fgc,t\cap t-1}} \left(\frac{P^U_{uct}}{P^U_{uc,t-1}} \right)^{w_{ugct}} \tag{C.23}$$

where the Sato-Vartia weights are:

$$w_{ugct} = \frac{L(S^{U*}_{ugct}, S^{U*}_{ugc,t-1})}{\sum_{k \in \Omega_{fgc,t\cap t-1}} L(S^{U*}_{kgct}, S^{U*}_{kgc,t-1})}, \quad L(a,b) = \frac{a-b}{\ln a - \ln b} \tag{C.24}$$

and S^{U*}_{ugct} is the expenditure share of product u among continuing products only.

The Sato-Vartia index is exact for CES preferences under the assumption that the geometric mean of appeal parameters is constant for continuing varieties:

$$\prod_{u \in \Omega_{fgc,t\cap t-1}} \left(\frac{\phi^U_{ut}}{\phi^U_{u,t-1}} \right)^{w_{ugct}} = 1 \tag{C.25}$$

This normalization is standard in the literature and implies that we measure price changes relative to a quality-constant basket.

Remark on Normalization. The Sato-Vartia index is exact under the normalization in equation (C.25), which holds UPC appeal constant for continuing varieties. Firm appeal is recovered as a *level* from observables in equation (C.29), $\phi^F_{fgct} = (P^F_{fgct}/P^G_{gct})(S^F_{fgct})^{1/(\sigma^F_{gs\tau}-1)}$, using the observed group price index P^G_{gct}; it therefore carries no free within-market normalization constant, and the growth rate $\Phi^{Appeal}_{fgct} \equiv \phi^F_{fgct}/\phi^F_{fgc,t-1}$ is well defined. (Had appeal instead been identified only up to a common market–year factor A_{gct}, the growth rate would be invariant only if $A_{gct} = A_{gc,t-1}$; here the factor is pinned down by P^G_{gct} and does not arise.) Our elasticity estimation (Section C.1) additionally relies on double-differencing, which eliminates appeal levels entirely.

C.3 Feenstra Variety Adjustment

To account for product entry and exit, we follow Feenstra (1994). Define the expenditure share on continuing products:

$$\lambda_{fgct} = \frac{\sum_{u \in \Omega_{fgc,t \cap t-1}} P^U_{uct} Q_{uct}}{\sum_{u \in \Omega_{fgct}} P^U_{uct} Q_{uct}} \tag{C.26}$$

The variety adjustment is:

$$\Phi^{Variety}_{fgct} = \left(\frac{\lambda_{fgct}}{\lambda_{fgc,t-1}} \right)^{\frac{1}{\sigma^U_{gs\tau}-1}} \tag{C.27}$$

where $\sigma^U_{gs\tau}$ is the within-firm elasticity for product group g, state s (with $c \in s$), and period τ (with $t \in \tau$).

When new products enter and capture market share, $\lambda_{fgct} < 1$, so $\lambda_{fgct} < \lambda_{fgc,t-1}$ (assuming exiting products had smaller shares than entering products). Since $\sigma^U_{gs\tau} > 1$, this implies $\Phi^{Variety}_{fgct} < 1$, lowering the effective price index. The intuition is that consumers benefit from the availability of new varieties, even if prices of existing products are unchanged.

C.4 Appeal Recovery

From the CES demand structure, the market share of firm f within product group g is:

$$S^F_{fgct} = \frac{(\phi^F_{fgct})^{\sigma^F_{gs\tau}-1} (P^F_{fgct})^{1-\sigma^F_{gs\tau}}}{\sum_{k \in \Omega_{gc\tau}} (\phi^F_{kgct})^{\sigma^F_{gs\tau}-1} (P^F_{kgct})^{1-\sigma^F_{gs\tau}}} \tag{C.28}$$

Rearranging and using the definition of the group-level price index P^G_{gct}:

$$\phi^F_{fgct} = \frac{P^F_{fgct}}{P^G_{gct}} \left(S^F_{fgct} \right)^{\frac{1}{\sigma^F_{gs\tau}-1}} \tag{C.29}$$

The growth rate of appeal is:

$$\frac{\phi^F_{fgct}}{\phi^F_{fgc,t-1}} = \frac{P^F_{fgct}/P^F_{fgc,t-1}}{P^G_{gct}/P^G_{gc,t-1}} \left(\frac{S^F_{fgct}}{S^F_{fgc,t-1}} \right)^{\frac{1}{\sigma^F_{gs\tau}-1}} \tag{C.30}$$

We denote this appeal growth rate as $\Phi^{Appeal}_{fgct} \equiv \phi^F_{fgct}/\phi^F_{fgc,t-1}$. Intuitively, if a firm's market share increases more than its relative price decline would predict (given $\sigma^F_{gs\tau}$), we attribute the residual to an increase in appeal.

Appendix D Theoretical Framework

This appendix develops the demand system and the firm problem behind the decompositions in Section 5.4, and proves Propositions 1 and 2 of the main text, together with the supporting results. We work throughout with the nested CES structure used in estimation, so that the two elasticities we estimate separately do separate work: the across-firm elasticity σ^F disciplines markups and therefore the Sato-Vartia price index, while the within-firm elasticity σ^U governs the reallocation of expenditure across a firm's products and therefore the appeal index.

D.1 Environment

D.1.1 Consumers

The economy has two consumer types, natives and immigrants, indexed by $h \in \{n, i\}$, with populations L_n and L_i. Preferences are nested. The upper tier is Cobb-Douglas between an outside good and a differentiated bundle, $U_h = C_{0h}^{1-\alpha_h} C_{Dh}^{\alpha_h}$, so a type-$h$ consumer with income y_h spends a constant fraction α_h of it on the differentiated bundle, an amount we write $E_h \equiv \alpha_h y_h$. The differentiated bundle is a two-level CES aggregate, across firms and across products within a firm,

$$C_{Dh} = \Big(\sum_f \phi_f^{1/\sigma^F_h} \, C_{fh}^{(\sigma^F_h-1)/\sigma^F_h} \Big)^{\sigma^F_h/(\sigma^F_h-1)}, \qquad C_{fh} = \Big(\sum_{u \in \Omega_f} \phi_u^{1/\sigma^U_h} \, c_{uh}^{(\sigma^U_h-1)/\sigma^U_h} \Big)^{\sigma^U_h/(\sigma^U_h-1)}, \tag{D.1}$$

with firm appeal $\phi_f > 0$, product quality $\phi_u > 0$, and elasticities $\sigma^U_h > \sigma^F_h > 1$, so that products of the same firm are closer substitutes than products of different firms. The exponent $1/\sigma^X_h$ on the taste weights makes the within-firm expenditure share proportional to $\phi_u \, p_u^{1-\sigma^U_h}$, with ϕ_u entering at a coefficient of one for both types, so that the two types value quality identically and differ only in σ^X_h. Immigrants are more price-elastic at both margins,

$$\sigma^U_i > \sigma^U_n > 1, \qquad \sigma^F_i > \sigma^F_n > 1. \tag{D.2}$$

Utility maximization gives the demand of a type-h consumer for product u of firm f,

$$c_{uh} = \phi_u \Big(\frac{p_u}{P_{fh}}\Big)^{-\sigma_h^U} \phi_f \Big(\frac{P_{fh}}{P_h}\Big)^{-\sigma_h^F} \frac{E_h}{P_h}, \tag{D.3}$$

with firm-level and aggregate price indices

$$P_{fh} = \Big(\sum_{u\in\Omega_f} \phi_u\, p_u^{1-\sigma_h^U}\Big)^{1/(1-\sigma_h^U)}, \qquad P_h = \Big(\sum_f \phi_f\, P_{fh}^{1-\sigma_h^F}\Big)^{1/(1-\sigma_h^F)}, \tag{D.4}$$

and total demand $q_u = \sum_h L_h c_{uh}$. Each tier is weakly separable and homothetic, so the consumer's problem decomposes tier by tier: the lower tier delivers a firm-level unit cost P_{fh} that summarizes the firm's entire product line, and the upper tier treats the firm as a single good priced at P_{fh}. This is a different model from the system estimated in Section 5.4, and it reuses that section's notation for the objects that play the same role. It is static and describes one product group in one market, so the location and time indices have no counterpart here and C_{0h} collects the remaining groups in (5.4). Dropping those indices, p_u, P_{fh}, P_h, Ω_f, and $(\sigma_h^U, \sigma_h^F, \delta)$ take the place of P_{uct}^U, P_{fgct}^F, P_{gct}^G, Ω_{fgct}, and $(\sigma_\ell^U, \sigma_\ell^F, \delta_\ell)$, with the consumer type h the only addition.

Two features of (D.1) matter for the results below. First, ϕ_u and ϕ_f are common to the two types, so the types agree on the ranking of products by quality and differ only in the elasticity with which they substitute across prices. No result below rests on a difference in the taste for quality. Second, demand is homothetic within a type: E_h enters (D.3) multiplicatively and appears in no elasticity, so income fixes what a consumer spends and nothing else. The gap $\sigma_i^X - \sigma_n^X$ for $X \in \{U, F\}$ is a separate primitive; the model does not derive it, and Section 6 motivates the gap.

Throughout, immigration means an increase in L_i holding L_n and the per-consumer expenditures E_n and E_i fixed, together with the firm's product set Ω_f, the appeal parameters, and the cost function; prices and quantities then adjust in equilibrium, holding the group index P_h fixed.

Remark 1 (Collapse to The Estimated System)**.** *If the two types coincide, $\sigma_i^U = \sigma_n^U \equiv \sigma^U$ and $\sigma_i^F = \sigma_n^F \equiv \sigma^F$, then $P_{fh} = P_f$ and $P_h = P$ no longer depend on h, and summing (D.3) over types gives*

$$q_u = \phi_u \Big(\frac{p_u}{P_f}\Big)^{-\sigma^U} \phi_f \Big(\frac{P_f}{P}\Big)^{-\sigma^F} \frac{\sum_h L_h E_h}{P}, \tag{D.5}$$

which is the single-consumer nested CES demand of Section 5.4 with aggregate expenditure $\sum_h L_h E_h$. The within-firm and across-firm expenditure shares reduce to (C.1) and (C.28) after the change of variable of Appendix D.6. The model has content only through the difference between the two types, and the results below record where that difference enters.

The remark fixes what the model and the estimating system have in common and where they diverge. They share the nesting and the role of the two elasticities; they differ in that the estimating system fits one representative consumer per product group, state, and period and takes prices as

data, whereas the model has two types facing prices that firms choose. The taste weights ϕ_u and ϕ_f play the role of the appeal parameters ϕ^U_{ut} and ϕ^F_{fgct} recovered there but are not the same objects, as Appendix D.6 makes precise. The exact link runs in the direction of the remark: the model collapses to the estimated system when the types coincide, not the other way around. What the model adds is a reading of the estimates, since a single-agent elasticity fitted to a two-type market averages the two types and should therefore rise with the immigrant share.

D.1.2 Firms

Each firm f produces a fixed set of products Ω_f with fixed qualities $\{\phi_u\}_{u\in\Omega_f}$. The cost of producing q_u units of product u is $C_u(q_u) = F_u + \frac{a_u}{1+\delta} q_u^{1+\delta}$, so marginal cost is

$$MC_u(q_u) = a_u\, q_u^{\delta}, \qquad \delta \geq 0, \tag{D.6}$$

with product-specific cost shifter $a_u > 0$ and returns to scale governed by δ. Firm profit is $\pi_f = \sum_{u\in\Omega_f}[\,p_u q_u - C_u(q_u)\,]$. Within a firm we let cost differences across products run through quality, $a_u = a(\phi_u)$ with $\varepsilon_a \equiv d\ln a / d\ln\phi_u \geq 0$, rather than through idiosyncratic cost draws. Qualities are fixed, so $a(\cdot)$ is not a technology for producing quality.[31] Proposition 6 shows that a firm's higher-quality products are more expensive when $\varepsilon_a + \delta > 0$, so either a costly quality or decreasing returns alone suffices.

D.1.3 Market Structure

Each firm is large within its own nest and atomistic across nests. In choosing $\{p_u\}_{u\in\Omega_f}$ it internalizes the substitution among its own products through P_{fh}, but takes the aggregate index P_h and aggregate expenditure $E_h L_h$ as given because its share of the group is negligible.

D.2 Equilibrium Pricing

D.2.1 Demand Derivatives

Raising p_u moves the firm's own index, $\partial \ln P_{fh} / \partial \ln p_u = s_{uh|f}$, where

$$s_{uh|f} \equiv \frac{\phi_u\, p_u^{1-\sigma^U_h}}{\sum_{v\in\Omega_f} \phi_v\, p_v^{1-\sigma^U_h}} \tag{D.7}$$

[31] The quality-price relation below is cross-sectional, while the supply residual of Appendix C.1 is a double-differenced *change* in relative cost. Since qualities are fixed here, $a(\phi_u)$ contributes nothing to that change and the Feenstra orthogonality condition is untouched.

is the within-firm expenditure share of a type-h consumer. Taking logs of (D.3) and differentiating, with P_h fixed, a type's demand responds as

$$\frac{\partial \ln c_{uh}}{\partial \ln p_v} = -\sigma_h^U \, \mathbb{1}\{u = v\} + \left(\sigma_h^U - \sigma_h^F\right) s_{vh|f}, \tag{D.8}$$

the direct effect plus the substitution through the firm's own index. The firm faces the sum of the two types' demands, $q_u = \sum_h L_h c_{uh}$, so aggregating (D.8) with the quantity weights $\theta_{uh} \equiv L_h c_{uh}/q_u$ gives

$$\frac{\partial \ln q_u}{\partial \ln p_v} = -\sigma_{u,\mathit{eff}}^U \mathbb{1}\{u = v\} + \sum_h \theta_{uh}\left(\sigma_h^U - \sigma_h^F\right) s_{vh|f}, \qquad \sigma_{u,\mathit{eff}}^U \equiv \sum_h \theta_{uh}\sigma_h^U. \tag{D.9}$$

Two shares appear in (D.9), and although both are indexed by a product and a type they divide different totals. The within-firm share $s_{vh|f}$ is a product share within a type: it divides what type h spends at the firm across the firm's products, and sums to one across products. The quantity share θ_{uh} is a type share within a product: it divides what product u sells across the two types, and sums to one across types.

Take the first term, the response of a product to its own price. The two types leave at different rates, and the firm loses each in proportion to what it was selling them, which is why the weight is θ_{uh}. The effective elasticity is product-specific even though the underlying σ_h^U are not: what differs across a firm's products is who buys them, so a product bought disproportionately by immigrants faces a more elastic demand than its neighbor on the same shelf. This is the object immigration moves.

The second term is what product u picks up when p_v rises. Each type moves spending in proportion to how much of it sat on v to begin with, which is $s_{vh|f}$, and the two types enter u's sales with the weights θ_{uh}. This term is what makes the firm's problem multidimensional, since a change in one price reaches every other product. It switches off in one direction: raising all of the firm's prices by the same proportion leaves the firm's internal relative prices, and so the shares $s_{vh|f}$, exactly where they were.

Lemma 1 (Uniform Price Scaling)**.** *Scale all of the firm's prices by a common factor $\xi > 0$. Then the within-firm shares $s_{uh|f}$ are unchanged, and each type's demand moves with the across-firm elasticity alone, $\partial \ln c_{uh}/\partial \ln \xi = -\sigma_h^F$, so that $\partial \ln q_u/\partial \ln \xi = -\sigma_{u,\mathit{eff}}^F$.*

Proof. In the firm's price vector, $s_{uh|f}$ is homogeneous of degree zero, which is the first claim, and P_{fh} is homogeneous of degree one; both by inspection of their formulas. Substituting ξp_v into (D.3), the ratio p_u/P_{fh} is unchanged and the only surviving factor is $(\xi P_{fh}/P_h)^{-\sigma_h^F}$, so $\partial \ln c_{uh}/\partial \ln \xi = -\sigma_h^F$. Aggregating over types with the quantity weights θ_{uh} gives $\partial \ln q_u/\partial \ln \xi = -\sigma_{u,\mathit{eff}}^F$. □

D.2.2 First-Order Condition and Markup

The firm chooses all of its prices jointly, so raising p_u moves the quantity of every product it sells, and the first-order condition retains the cross-product terms,

$$\frac{\partial \pi_f}{\partial p_u} = q_u + (p_u - MC_u)\frac{\partial q_u}{\partial p_u} + \sum_{v \neq u}(p_v - MC_v)\frac{\partial q_v}{\partial p_u} = 0. \tag{D.10}$$

Proposition 3 (Pricing). *Let $m_v \equiv (p_v - MC_v)/p_v$ be the Lerner index of product v and $\bar{m}_h \equiv \sum_{v \in \Omega_f} s_{vh|f}\, m_v$ the average a type-h consumer pays at the firm. At an interior optimum, for every $u \in \Omega_f$,*

$$1 = m_u\, \sigma^U_{u,eff} - \sum_h \theta_{uh}\big(\sigma^U_h - \sigma^F_h\big)\, \bar{m}_h, \tag{D.11}$$

so the firm's products are coupled only through $\bar{m}_n$ and $\bar{m}_i$.

Proof. Let $R_h \equiv L_h \sum_{v \in \Omega_f} p_v c_{vh}$ be what type h spends at the firm, so that $s_{vh|f} = p_v c_{vh} / \sum_{w \in \Omega_f} p_w c_{wh}$ and

$$p_v q_v\, \theta_{vh} = L_h\, p_v c_{vh} = R_h\, s_{vh|f} \qquad \text{for every } v \in \Omega_f. \tag{D.12}$$

Write (D.10) in logs as $p_u q_u + \sum_{v \in \Omega_f}(p_v - MC_v)\, \partial q_v / \partial \ln p_u = 0$, substitute (D.9) with the roles of u and v exchanged, and use $(p_v - MC_v)q_v = m_v p_v q_v$:

$$p_u q_u - m_u\, p_u q_u\, \sigma^U_{u,eff} + \sum_h \big(\sigma^U_h - \sigma^F_h\big)\, s_{uh|f} \sum_{v \in \Omega_f} m_v\, p_v q_v \theta_{vh} = 0.$$

By (D.12) the inner sum is $R_h \sum_v s_{vh|f} m_v = R_h \bar{m}_h$, and (D.12) applied to product u gives $R_h\, s_{uh|f} = \theta_{uh}\, p_u q_u$. Dividing by $p_u q_u$ yields (D.11). □

Equation (D.11) balances what a price increase earns against what it costs. Holding quantities fixed, a one percent rise in p_u raises revenue by one percent of $p_u q_u$, the left side. Against it the firm loses margin on the sales that leave, at the full within-firm elasticity, and recovers margin on the part that leaves product u but stays with the firm. A type-h buyer leaves at σ^U_h and reaches another firm at σ^F_h, so the difference is what stays within the firm's own line. What stays does not return at the margin of the product that was left but spreads across the line in proportion to $s_{vh|f}$, so the recovered sales are valued at $\bar{m}_h$ rather than at any single m_v. The two types are kept apart throughout because they divert differently, and each enters with the share θ_{uh} it accounts for in product u's sales.

With a single consumer type $\theta_{uh} = 1$, $\sigma^U_{u,eff} = \sigma^U$ and $\bar{m} = m$, so (D.11) becomes $1 = m\big(\sigma^U - (\sigma^U - \sigma^F)\big) = m\sigma^F$, the familiar constant markup over marginal cost. The within-firm elasticity drops out even though every cross-price derivative involves it, because a multiproduct firm owns the products its price increase diverts sales to. With two types it does not drop out product by product: $\sigma^U_{u,eff}$ remains in (D.11) and varies with the buyer mix. It does drop out of the firm's revenue-weighted average.

Proposition 4 (Aggregate Pricing Rule). *Weighting* (D.11) *by the revenue shares* $w_u \equiv p_u q_u / \sum_{v \in \Omega_f} p_v q_v$ *and summing over* $u \in \Omega_f$*, the within-firm elasticity cancels and*

$$\sum_h \theta_{fh}\, \sigma_h^F\, \bar{m}_h = 1, \tag{D.13}$$

with θ_{fh} *as in* (6.4).

Proof. Dividing (D.12) by the firm's revenue gives $w_u \theta_{uh} = \theta_{fh} s_{uh|f}$. Applying it to the first term of (D.11), $\sum_u w_u m_u \sigma^U_{u,\mathit{eff}} = \sum_h \sigma_h^U \sum_u (w_u \theta_{uh}) m_u = \sum_h \theta_{fh} \sigma_h^U \bar{m}_h$; applying it to the second, $\sum_u w_u \sum_h \theta_{uh} (\sigma_h^U - \sigma_h^F) \bar{m}_h = \sum_h \theta_{fh} (\sigma_h^U - \sigma_h^F) \bar{m}_h \sum_u s_{uh|f} = \sum_h \theta_{fh} (\sigma_h^U - \sigma_h^F) \bar{m}_h$. The left side weights to $\sum_u w_u = 1$, and the difference of the two right-hand terms is (D.13). □

Equation (D.13) holds whatever the dispersion of markups across the firm's products. The main text restricts them to a single markup, so that $\bar{m}_n = \bar{m}_i = m$ and (D.13) gives $m = 1/\sigma^F_{f,\mathit{eff}}$ with $\sigma^F_{f,\mathit{eff}} \equiv \sum_h \theta_{fh} \sigma_h^F$, which is (6.3). With one consumer type, or with two that share the within-firm elasticity, this is not a restriction: (D.11) then delivers a common markup on its own. It is one here, since we keep $\sigma_i^U > \sigma_n^U$ for the composition results of Appendix D.4, and the unrestricted solution then lowers m_u on the products whose buyers are more immigrant, where $\sigma^F_{u,\mathit{eff}}$ is higher. Both (D.11) and (D.13) are the two-type case of the multiproduct pricing conditions in Faber and Fally (2022), derived there with a continuum of consumer groups.[32]

Three things remain to be shown about the unrestricted system: that (D.11) has a solution at all, how far the markups in it can spread, and whether the restriction to a single markup changes the direction of the comparative statics. Define $A_h \equiv \sigma_h^U - \sigma_h^F \geq 0$ and solve (D.11) for m_u,

$$m_u = \frac{1 + \sum_h \theta_{uh} A_h \bar{m}_h}{\sum_h \theta_{uh} \sigma_h^U}. \tag{D.14}$$

This leaves the $\bar{m}_h$ on the right, and those are averages of the m_v. Writing $\bar{m}_h = \sum_v s_{vh|f} m_v$ and replacing each m_v in it by (D.14) eliminates the m_v, so the $|\Omega_f|$ equations become two, one for each type, with $\bar{m}_n$ and $\bar{m}_i$ the only unknowns. They are linear, and solving them returns every m_u through (D.14).

Proposition 5 (Existence and Uniqueness). *Write* $\ell_h \equiv (1 + A_h \bar{m}_h)/\sigma_h^U$ *and* $\lambda_u \equiv \theta_{ui} \sigma_i^U / \sigma^U_{u,\mathit{eff}}$. *Then* (D.14) *becomes*

$$m_u = \lambda_u\, \ell_i + (1 - \lambda_u)\, \ell_n, \tag{D.15}$$

[32] With finitely many firms the markup-relevant elasticity would instead interpolate between $\sigma^F_{f,\mathit{eff}}$ and the elasticity of demand faced by the group, according to the firm's group market share (Atkeson and Burstein, 2008). We work in the limit in which that share is negligible. Immigration raises the effective elasticities either way, so the direction of the markup response does not turn on this.

and, with $\bar{\lambda}_h \equiv \sum_{v\in\Omega_f} s_{vh|f}\lambda_v$, *the two constants solve*

$$\sigma_h^U \ell_h - A_h\big[\bar{\lambda}_h \ell_i + (1-\bar{\lambda}_h)\,\ell_n\big] = 1, \qquad h \in \{n, i\}. \tag{D.16}$$

Under $\sigma_h^U \geq \sigma_h^F > 1$ *the two equations in* (D.16) *have a unique solution, so* (D.11) *does too, and*

$$\frac{1}{\sigma_i^F} \;\leq\; \ell_h,\, m_u \;\leq\; \frac{1}{\sigma_n^F} \qquad \textit{for each } h \textit{ and every } u \in \Omega_f. \tag{D.17}$$

Every Lerner index therefore lies in $(0,1)$*, prices are positive and markups finite, and the two bounds coincide when the two types do.*

Proof. Use $\theta_{un} + \theta_{ui} = 1$ on the 1 in the numerator of (D.14): it becomes $\sum_h \theta_{uh}(1 + A_h\bar{m}_h) = \sum_h \theta_{uh}\sigma_h^U\ell_h$. The denominator is $\sum_h \theta_{uh}\sigma_h^U$, so m_u is a weighted average of ℓ_i and ℓ_n with weight λ_u on ℓ_i, which is (D.15). Substituting (D.15) into $\bar{m}_h = \sum_v s_{vh|f} m_v$ gives $\bar{m}_h = \bar{\lambda}_h\ell_i + (1-\bar{\lambda}_h)\ell_n$, and putting that into $\sigma_h^U\ell_h = 1 + A_h\bar{m}_h$ gives (D.16).

Write (D.16) as $B\ell = \mathbf{1}$ with $\ell = (\ell_i, \ell_n)'$ and

$$B = \begin{pmatrix} \sigma_i^U - A_i\bar{\lambda}_i & -A_i(1-\bar{\lambda}_i) \\ -A_n\bar{\lambda}_n & \sigma_n^U - A_n(1-\bar{\lambda}_n) \end{pmatrix}.$$

Two properties of B deliver the proposition: that it is invertible, which makes ℓ unique, and that $B^{-1} \geq 0$ entrywise, which lets ℓ be bounded without computing it, since $Bx \geq B\ell$ then implies $x \geq \ell$. Both follow from the sign pattern of B. Its off-diagonal entries are nonpositive, and each row sums to $\sigma_h^U - A_h = \sigma_h^F > 0$ because the $\bar{\lambda}_h$ terms cancel. A row sum is the diagonal entry less the size of the off-diagonal one, so the diagonal entries are positive and exceed the off-diagonal ones: B is strictly diagonally dominant, hence a nonsingular M-matrix.

For the bounds, apply B to a constant vector $c\mathbf{1}$. Both entries of $c\mathbf{1}$ being equal, this scales the row sums just computed and returns $(c\,\sigma_i^F,\, c\,\sigma_n^F)'$. At $c = 1/\sigma_n^F$ that vector is at least $\mathbf{1} = B\ell$, so $\ell \leq c\mathbf{1}$; at $c = 1/\sigma_i^F$ it is at most $\mathbf{1}$, so $\ell \geq c\mathbf{1}$. Each m_u is a convex combination of ℓ_i and ℓ_n by (D.15) and obeys the same bounds. □

By (D.15) every m_u lies between ℓ_i and ℓ_n, so the gap between those two is what a single markup for all of the firm's products gives up.

Corollary 1 (What a Common Markup Suppresses)**.** *Write* $\mu_u \equiv 1/(1-m_u)$ *for the markup on product* u*. Then*

$$\frac{1}{\sigma_i^F} < \ell_i < \ell_n < \frac{1}{\sigma_n^F}, \qquad \max_{u\in\Omega_f}\mu_u - \min_{u\in\Omega_f}\mu_u \;\leq\; \frac{\sigma_i^F - \sigma_n^F}{\big(\sigma_n^F - 1\big)\big(\sigma_i^F - 1\big)}. \tag{D.18}$$

Proof. Two claims: that $\ell_i < \ell_n$, and that both lie strictly inside $[1/\sigma_i^F, 1/\sigma_n^F]$. Both come from the difference $d \equiv \ell_i - \ell_n$, and the bound on markups then follows from (D.15).

Substituting $\ell_i = \ell_n + d$ into the two equations of (D.16) and subtracting them, the within-firm gap $\sigma_i^U - \sigma_n^U$ cancels and leaves

$$d = \frac{\sigma_i^F - \sigma_n^F}{A_i\bar{\lambda}_i - A_n\bar{\lambda}_n - \sigma_i^U}\,\ell_n.$$

Since $A_i\bar{\lambda}_i \le A_i$ and $A_n\bar{\lambda}_n \ge 0$, the denominator is at most $A_i - \sigma_i^U = -\sigma_i^F < 0$, while $\ell_n > 0$ by (D.17). So $d < 0$, which is the first claim.

For the second, return to (D.16) without subtracting, and in each equation replace the ℓ that is not being isolated. At $h = n$, substituting $\ell_i = \ell_n + d$,

$$\sigma_n^U\ell_n - A_n\big(\ell_n + \bar{\lambda}_n d\big) = 1 \qquad\Longrightarrow\qquad \sigma_n^F\,\ell_n = 1 + A_n\bar{\lambda}_n\,d,$$

using $\sigma_n^U - A_n = \sigma_n^F$; at $h = i$, substituting $\ell_n = \ell_i - d$,

$$\sigma_i^U\ell_i - A_i\big(\ell_i - (1-\bar{\lambda}_i)d\big) = 1 \qquad\Longrightarrow\qquad \sigma_i^F\,\ell_i = 1 - A_i(1-\bar{\lambda}_i)\,d.$$

With $d < 0$ the first right side is below one and the second above it, so $\ell_n < 1/\sigma_n^F$ and $\ell_i > 1/\sigma_i^F$.

Every m_u lies between ℓ_i and ℓ_n by (D.15) and so in $[1/\sigma_i^F, 1/\sigma_n^F]$. Since $\mu = 1/(1-m)$ is increasing, every μ_u lies in $\big[\sigma_i^F/(\sigma_i^F-1),\ \sigma_n^F/(\sigma_n^F-1)\big]$, an interval of width $\sigma_n^F/(\sigma_n^F-1) - \sigma_i^F/(\sigma_i^F-1) = \big(\sigma_i^F - \sigma_n^F\big)\big/\big[(\sigma_n^F-1)(\sigma_i^F-1)\big]$. □

A bound on the dispersion is not by itself a statement about the mechanism, since the two move together in $\sigma_i^F - \sigma_n^F$. What matters is that the restriction leaves the direction of the comparative statics unchanged.

Corollary 2 (The Restriction Does Not Affect the Direction)**.** *At fixed prices, an increase in L_i lowers every m_u in the solution of (D.11), as it lowers the common markup $1/\sigma_{f,\mathit{eff}}^F$ of (6.3).*

Proof. By (D.15), L_i affects m_u only through λ_u, ℓ_i and ℓ_n. We show that it raises λ_u and, through $\bar{\lambda}_h$, lowers both ℓ_i and ℓ_n; with $\ell_i < \ell_n$ the two movements push m_u the same way.

A rise in L_i raises every θ_{ui}. Writing $\lambda_u = \big[1 + \big((1-\theta_{ui})/\theta_{ui}\big)\big(\sigma_n^U/\sigma_i^U\big)\big]^{-1}$ shows that every λ_u rises with it, since $(1-\theta_{ui})/\theta_{ui}$ falls. The shares $s_{vh|f}$ are per-consumer and so do not depend on L_i at fixed prices, leaving $\bar{\lambda}_h = \sum_v s_{vh|f}\lambda_v$ higher for each h.

Solving $B\ell = \mathbf{1}$ by Cramer's rule,

$$\ell_i = \frac{Z+W}{\sigma_i^F Z + \sigma_n^F W}, \qquad \ell_n = \frac{X+Y}{\sigma_n^F X + \sigma_i^F Y},$$

where $X \equiv \sigma_i^U - A_i\bar{\lambda}_i$, $Y \equiv A_n\bar{\lambda}_n$, $Z \equiv \sigma_n^U - A_n(1-\bar{\lambda}_n)$ and $W \equiv A_i(1-\bar{\lambda}_i)$ are all nonnegative, with $X \ge \sigma_i^F$ and $Z \ge \sigma_n^F$. Each ℓ_h is therefore the reciprocal of a weighted average of σ_n^F and σ_i^F: the weight on σ_n^F is $X/(X+Y)$ in ℓ_n and $W/(Z+W)$ in ℓ_i. Raising $\bar{\lambda}_i$ lowers X and W, and raising $\bar{\lambda}_n$ raises Y and Z; either shifts weight from σ_n^F to the larger σ_i^F in both averages, so both ℓ_i and ℓ_n

fall. By (D.15), m_u is a weighted average of ℓ_i and ℓ_n, with weight λ_u on ℓ_i, which is the smaller by (D.18). A rise in L_i lowers both constants and raises λ_u, so both movements lower m_u. □

D.2.3 Quality and Price

The pricing results so far concern the level of markups, not which of a firm's products are expensive. Proposition 2 needs that ordering: immigration tilts spending toward a firm's cheaper products, and to turn a statement about prices into one about quality the two must be ranked the same way. That is what this subsection establishes.

Proposition 6 (Quality-Price Relationship)**.** *Fix a firm f and let the markup be common across its products as in (6.3). If $\varepsilon_a + \delta > 0$, then across the firm's products at any equilibrium*

$$\phi_u > \phi_v \implies p_u > p_v. \tag{D.19}$$

Holding the firm's price indices fixed, the local elasticity is

$$\frac{d\ln p_u}{d\ln \phi_u} = \frac{\varepsilon_a + \delta}{1 + \delta\, \sigma^U_{u,eff}}. \tag{D.20}$$

Proof. Two claims: the ranking (D.19), and the elasticity (D.20). Both come from writing the equilibrium condition with p_u on one side and ϕ_u on the other.

Holding P_{fh} and P_h at their equilibrium values, total demand for a product of firm f is $q_u = \phi_u\, G(p_u)$ with $G(p) = \sum_h D_{fh}\, p^{-\sigma^U_h}$ and $D_{fh} \equiv L_h\, \phi_f\, P_{fh}^{\sigma^U_h}\, (P_{fh}/P_h)^{-\sigma^F_h} E_h/P_h > 0$ common to all $u \in \Omega_f$; note $G' < 0$. Combining (6.3) and (D.6) gives $p_u = \mu_f a(\phi_u) q_u^\delta$ with $\mu_f \equiv \sigma^F_{f,eff}/(\sigma^F_{f,eff} - 1)$ common across the firm's products. Substituting $q_u = \phi_u G(p_u)$ and collecting the terms in p_u,

$$\frac{p_u}{G(p_u)^\delta} = \mu_f\, a(\phi_u)\, \phi_u^\delta. \tag{D.21}$$

The left side depends on p_u alone and is strictly increasing in it, since G is decreasing and $\delta \geq 0$; the right side depends on ϕ_u alone and is strictly increasing in it exactly when $\varepsilon_a + \delta > 0$. Comparing any two products of the firm then gives (D.19): the higher-quality one has the larger right side, hence the larger left side, hence the higher price. The argument needs only the two products themselves, not a continuum of qualities in between.

Taking logs of (D.21) and differentiating in $\ln \phi_u$,

$$\Big(1 - \delta\, \frac{\partial \ln G}{\partial \ln p_u}\Big) \frac{d\ln p_u}{d\ln \phi_u} = \varepsilon_a + \delta, \qquad \frac{\partial \ln G}{\partial \ln p_u} = -\sum_h \frac{D_{fh} p_u^{-\sigma^U_h}}{G}\, \sigma^U_h.$$

The weights in the second expression are the quantity shares of (D.9), since $q_u = \phi_u G$ and $L_h c_{uh} = \phi_u D_{fh} p_u^{-\sigma^U_h}$ give $\theta_{uh} = D_{fh} p_u^{-\sigma^U_h}/G$. So $\partial \ln G/\partial \ln p_u = -\sigma^U_{u,eff}$, which delivers (D.20). □

Higher quality shifts demand out, which under decreasing returns raises marginal cost and therefore price, and a costly quality reinforces it. The ranking compares two products within one equilibrium rather than tracing what happens when a price moves, so P_{fh} stays fixed; the demand derivatives (D.9) do the latter, and there P_{fh} moves with p_u.

D.3 Price Decomposition

This subsection proves Proposition 1, which isolates the two direct forces on price, and then gives the equilibrium response, first in general and then in the special case in which the two forces separate cleanly. All of it starts from one input. By (6.3) the firm's markup is set by $\sigma^F_{f,\mathit{eff}}$, so what immigration does to price begins with what it does to that elasticity, and the first result gives the two channels through which immigration raises $\sigma^F_{f,\mathit{eff}}$.

Lemma 2 (Immigration Raises the Effective Elasticities)**.** *At fixed prices,*

$$\frac{\partial \sigma^X_{u,\mathit{eff}}}{\partial L_i} = \frac{L_n c_{un} c_{ui}}{q_u^2}\left(\sigma_i^X - \sigma_n^X\right) > 0, \qquad X \in \{U, F\}, \tag{D.22}$$

and at the firm level

$$\frac{\partial \sigma^F_{f,\mathit{eff}}}{\partial L_i} = \underbrace{\sum_{u \in \Omega_f} w_u \frac{\partial \sigma^F_{u,\mathit{eff}}}{\partial L_i}}_{\text{buyer mix of each product}} + \underbrace{\frac{\sigma_i^F - \sigma_n^F}{L_i}\operatorname{Var}_w\left(\theta_{ui}\right)}_{\text{reallocation across products}} > 0, \tag{D.23}$$

where $\operatorname{Var}_w(\theta_{ui}) = \sum_{u \in \Omega_f} w_u(\theta_{ui} - \theta_{fi})^2$ *is the variance of the immigrant share across the firm's products under the sales weights. The second term vanishes when every product has the same buyer mix.*

Proof. The per-consumer demands c_{uh} depend on L_i only through prices, so at fixed prices $q_u = L_n c_{un} + L_i c_{ui}$ gives $\partial \theta_{ui}/\partial L_i = c_{ui} L_n c_{un}/q_u^2 > 0$. Since $\sigma^X_{u,\mathit{eff}} = \sigma_n^X + \theta_{ui}(\sigma_i^X - \sigma_n^X)$ is linear in θ_{ui}, this gives (D.22).

Differentiating $\sigma^F_{f,\mathit{eff}} = \sum_u w_u \sigma^F_{u,\mathit{eff}}$ gives $\sum_u w_u\, \partial \sigma^F_{u,\mathit{eff}}/\partial L_i + \sum_u \sigma^F_{u,\mathit{eff}}\, \partial w_u/\partial L_i$, the first term of which is positive by (D.22). For the second term we need $\partial w_u/\partial L_i = w_u\, \partial \ln w_u/\partial L_i$. With prices fixed, $w_u = p_u q_u / \sum_v p_v q_v$ gives

$$\frac{\partial \ln w_u}{\partial L_i} = \frac{\partial \ln q_u}{\partial L_i} - \frac{\partial \ln \sum_v p_v q_v}{\partial L_i}.$$

Here $\partial \ln q_u/\partial L_i = c_{ui}/q_u = \theta_{ui}/L_i$. The second derivative is the sales-weighted average of the first across the firm's products, since the firm's revenue is their sum, and equals θ_{fi}/L_i by $\sum_v w_v \theta_{vi} = \theta_{fi}$. So $\partial \ln w_u/\partial L_i = (\theta_{ui} - \theta_{fi})/L_i$: a product gains sales weight exactly when its own immigrant share exceeds the firm's.

The second term of (D.23) is therefore $L_i^{-1}\sum_u w_u(\theta_{ui}-\theta_{fi})\sigma^F_{u,eff}$, a weighted sum of the deviations $\theta_{ui}-\theta_{fi}$. Those deviations have weighted sum zero, since $\sum_u w_u\theta_{ui}=\theta_{fi}$ and $\sum_u w_u=1$, so a term that does not vary across products contributes nothing to the sum. This removes σ^F_n from $\sigma^F_{u,eff}=\sigma^F_n+\theta_{ui}(\sigma^F_i-\sigma^F_n)$, and then θ_{fi} from $\theta_{ui}=(\theta_{ui}-\theta_{fi})+\theta_{fi}$:

$$\sum_u w_u\big(\theta_{ui}-\theta_{fi}\big)\sigma^F_{u,eff}=\big(\sigma^F_i-\sigma^F_n\big)\sum_u w_u\big(\theta_{ui}-\theta_{fi}\big)\theta_{ui}=\big(\sigma^F_i-\sigma^F_n\big)\,\mathrm{Var}_w\big(\theta_{ui}\big).$$

□

The two terms of (D.23) are distinct channels. The first holds the firm's product composition fixed and asks how the buyers of each product change; the second holds each product's buyers fixed and asks which products sell more. Immigration favors the products already bought disproportionately by immigrants, and those are the ones facing the more elastic demand, so the reallocation reinforces the first channel rather than offsetting it.

Proof of Proposition 1. Take logs of the pricing rule (6.3) to get $\ln p_u=\ln\mu_f(\sigma^F_{f,eff})+\ln a_u+\delta\ln q_u$ and differentiate in L_i, holding the price fixed wherever it appears on the right. The markup term then responds through Lemma 2 and the scale term through the quantity:

$$\Big(\frac{d\ln p_u}{dL_i}\Big)^{\text{direct}}=\underbrace{\frac{d\ln\mu_f}{d\sigma^F_{f,eff}}\frac{\partial\sigma^F_{f,eff}}{\partial L_i}}_{\text{markup effect }(<0)}+\underbrace{\delta\,\frac{1}{q_u}\frac{\partial q_u}{\partial L_i}}_{\text{scale effect }(>0)},\tag{D.24}$$

with $d\ln\mu_f/d\sigma^F_{f,eff}=-1/[\sigma^F_{f,eff}(\sigma^F_{f,eff}-1)]<0$, $\partial\sigma^F_{f,eff}/\partial L_i>0$ by (D.23), and $\partial q_u/\partial L_i=c_{ui}$ with $c_{ui}/q_u=\theta_{ui}/L_i$, so the scale term is $\delta\,\theta_{ui}/L_i$. Aggregating across the firm's products with the log-mean expenditure weights $\bar s_u$ gives the display of the proposition.[33] □

Remark 2 (Equilibrium Response: The General Case)**.** *In equilibrium the price responds as well, and that response feeds back into both the markup and marginal cost. Write $\zeta_h\equiv d\ln P_{fh}/dL_i$ for the response of the type-h firm price index and $\zeta_f\equiv d\ln P^{SV}_f/dL_i$ for that of the firm's Sato-Vartia index, which to first order aggregates the product-level responses with the revenue weights w_u. Without restricting markups across products, ζ_f satisfies two exact identities, one decomposing it by consumer type and one by the force at work:*

$$\zeta_f=\sum_h\theta_{fh}\zeta_h,\qquad \zeta_f=\sum_{u\in\Omega_f}w_u\frac{d\ln\mu_u}{dL_i}+\delta\Big[\frac{\theta_{fi}}{L_i}-\sum_h\theta_{fh}\sigma^F_h\zeta_h\Big].\tag{D.25}$$

[33] Those weights respond to the sales composition, but they enter the index only multiplicatively against price changes, so to first order their response does not contribute to the price effect. To that order the index captures the two price forces and not the reallocation of expenditure across products, which is the composition channel of Appendix D.4.

Both follow from differentiating $\ln p_u = \ln \mu_u + \ln a_u + \delta \ln q_u$ product by product and aggregating with w_u. For the first, $\zeta_h = \sum_u s_{uh|f}\, d\ln p_u/dL_i$ by (D.7), so $w_u\theta_{uh} = \theta_{fh}s_{uh|f}$ and $\sum_h \theta_{uh} = 1$ turn $\sum_u w_u\, d\ln p_u/dL_i$ into $\sum_h \theta_{fh}\zeta_h$. For the second, letting prices respond in (D.8) gives

$$\frac{d\ln q_u}{dL_i} = \frac{\theta_{ui}}{L_i} + \sum_h \theta_{uh}\Big[-\sigma_h^U \frac{d\ln p_u}{dL_i} + A_h\zeta_h\Big],$$

the population term plus the own-price and firm-index responses. Aggregating with w_u and using the same identity, the second term becomes $-\sum_h \theta_{fh}\sigma_h^U\zeta_h$ and the third $\sum_h \theta_{fh}A_h\zeta_h$; their sum is $-\sum_h \theta_{fh}\sigma_h^F\zeta_h$ because $\sigma_h^U - A_h = \sigma_h^F$.

The first identity connects the index we measure to the two the model has: what we observe is the revenue-weighted average of what the two types experience. The second separates the markup response from the cost response, and in doing so shows that the within-firm elasticity cancels from the comparative static exactly as it cancels from the pricing rule (D.13). Only σ^F appears in the bracket, so the mechanism operates through the across-firm margin whether or not markups are restricted. These are identities, not a solution: ζ_n and ζ_i remain on the right, and the markup response is itself endogenous. Closing the system needs a further restriction, which the next remark imposes.

Remark 3 (Special Case: A Common Within-Firm Elasticity)**.** *Suppose the two types share the within-firm elasticity, $\sigma_i^U = \sigma_n^U$. Then (D.25) closes in a formula,*

$$\zeta_f = \Lambda_f\Big[\,\textit{markup effect} + \delta\,\frac{\theta_{fi}}{L_i}\,\Big], \tag{D.26}$$

$$\textit{markup effect} = -\frac{(\sigma_i^F - \sigma_n^F)\theta_{fi}\theta_{fn}}{L_i\,\sigma_{f,\mathit{eff}}^F(\sigma_{f,\mathit{eff}}^F - 1)} < 0, \qquad \Lambda_f = \Big[1 + \delta\,\sigma_{f,\mathit{eff}}^F - \frac{\mathrm{Var}_{\theta_f}(\sigma_h^F)}{\sigma_{f,\mathit{eff}}^F(\sigma_{f,\mathit{eff}}^F - 1)}\Big]^{-1}, \tag{D.27}$$

with $\mathrm{Var}_{\theta_f}(\sigma_h^F) = \theta_{fn}\theta_{fi}\big(\sigma_i^F - \sigma_n^F\big)^2$. The bracket in (D.26) is the direct decomposition of Proposition 1, and Λ_f collects the two feedbacks the direct derivative held fixed.

The restriction does three things. By (D.7) the two types then hold the same within-firm shares, $s_{ui|f} = s_{un|f}$, so every product has the same buyer mix, $\theta_{uh} = \theta_{fh}$; every product therefore takes the same markup $\mu_f(\sigma_{f,\mathit{eff}}^F)$; and the two types' price indices respond alike, $\zeta_n = \zeta_i$, since ζ_h weights the same price changes by the same shares. Both of the endogenous objects in (D.25) are thereby resolved.

Take (D.25) term by term. Both feedbacks show up as terms in ζ_f on the right, and collecting them on the left is what produces Λ_f^{-1}.

Start with the cost term, $\sum_h \theta_{fh}\sigma_h^F\zeta_h$ inside the bracket. With $\zeta_n = \zeta_i = \zeta_f$ it is $\sigma_{f,\mathit{eff}}^F\zeta_f$, so (D.25) becomes

$$\zeta_f = \sum_u w_u \frac{d\ln\mu_u}{dL_i} + \delta\,\frac{\theta_{fi}}{L_i} - \delta\,\sigma_{f,\mathit{eff}}^F\zeta_f, \qquad \text{that is} \qquad \zeta_f\big(1 + \delta\,\sigma_{f,\mathit{eff}}^F\big) = \sum_u w_u \frac{d\ln\mu_u}{dL_i} + \delta\,\frac{\theta_{fi}}{L_i}.$$

That coefficient is the first piece of Λ_f^{-1}. A lower price raises quantity and with it marginal cost, which pushes the price back up, so this feedback dampens.

Now the markup term, $\sum_u w_u \, d\ln\mu_u/dL_i$. All markups equal $\mu_f(\sigma^F_{f,\mathit{eff}})$, so by the chain rule this term is

$$\frac{d\ln\mu_f}{d\sigma^F_{f,\mathit{eff}}}\cdot\frac{d\sigma^F_{f,\mathit{eff}}}{dL_i}, \qquad \frac{d\sigma^F_{f,\mathit{eff}}}{dL_i}=\underbrace{\frac{\partial\sigma^F_{f,\mathit{eff}}}{\partial L_i}}_{\text{direct}}+\underbrace{\frac{\partial\sigma^F_{f,\mathit{eff}}}{\partial\ln\xi}\,\zeta_f}_{\text{loop}}.$$

The loop part is the second feedback. Immigration lowers the price; a lower price raises what both types spend at the firm but raises it more for the more elastic type, so θ_{fi} rises and $\sigma^F_{f,\mathit{eff}}$ rises again. To size it, note that $\theta_{fh}=L_hE_{fh}/\sum_k L_kE_{fk}$ and $\partial\ln E_{fh}/\partial\ln\xi=1-\sigma^F_h$ by (D.30), so the share responds at its own rate less the firm total's, which is the share-weighted average of the same rates:

$$\frac{\partial\ln\theta_{fh}}{\partial\ln\xi}=\big(1-\sigma^F_h\big)-\big(1-\sigma^F_{f,\mathit{eff}}\big)=-\big(\sigma^F_h-\sigma^F_{f,\mathit{eff}}\big).$$

Differentiating $\sigma^F_{f,\mathit{eff}}=\sum_h\theta_{fh}\sigma^F_h$ and using $\sum_h\theta_{fh}\big(\sigma^F_h-\sigma^F_{f,\mathit{eff}}\big)=0$,

$$\frac{\partial\sigma^F_{f,\mathit{eff}}}{\partial\ln\xi}=-\sum_h\theta_{fh}\,\sigma^F_h\big(\sigma^F_h-\sigma^F_{f,\mathit{eff}}\big)=-\sum_h\theta_{fh}\big(\sigma^F_h-\sigma^F_{f,\mathit{eff}}\big)^2=-\mathrm{Var}_{\theta_f}\big(\sigma^F_h\big). \tag{D.28}$$

Since $d\ln\mu_f/d\sigma^F_{f,\mathit{eff}}=-1/[\sigma^F_{f,\mathit{eff}}(\sigma^F_{f,\mathit{eff}}-1)]$, multiplying the two factors gives

$$\frac{d\ln\mu_f}{d\sigma^F_{f,\mathit{eff}}}\cdot\frac{d\sigma^F_{f,\mathit{eff}}}{dL_i}=\underbrace{-\frac{1}{\sigma^F_{f,\mathit{eff}}\big(\sigma^F_{f,\mathit{eff}}-1\big)}\frac{\partial\sigma^F_{f,\mathit{eff}}}{\partial L_i}}_{\text{markup effect}}+\frac{\mathrm{Var}_{\theta_f}(\sigma^F_h)}{\sigma^F_{f,\mathit{eff}}\big(\sigma^F_{f,\mathit{eff}}-1\big)}\,\zeta_f,$$

the second piece a positive multiple of ζ_f because both factors were negative. Substituting into the equation left after the cost term and moving that piece across,

$$\zeta_f\Big(1+\delta\,\sigma^F_{f,\mathit{eff}}-\frac{\mathrm{Var}_{\theta_f}(\sigma^F_h)}{\sigma^F_{f,\mathit{eff}}\big(\sigma^F_{f,\mathit{eff}}-1\big)}\Big)=\text{markup effect}+\delta\,\frac{\theta_{fi}}{L_i},$$

which is (D.26). Unlike the cost side, this loop amplifies: lower prices raise the elasticity, which lowers the markup, which lowers prices again.

The two feedbacks therefore pull against each other in Λ_f^{-1}, and Λ_f is positive wherever the pricing rule (6.3) describes a profit maximum.[34] The feedback then changes the size of the response but not its sign, which the bracket alone settles. Dividing it by $\theta_{fi}/L_i>0$, the response is negative

[34] $\Lambda_f>0$ is the firm's second-order condition rather than an extra assumption. Multiplying Λ_f^{-1} in (D.27) by $\sigma^F_{f,\mathit{eff}}\big(\sigma^F_{f,\mathit{eff}}-1\big)>0$ shows that $\Lambda_f>0$ is exactly $\mathrm{Var}_{\theta_f}(\sigma^F_h)<\sigma^F_{f,\mathit{eff}}\big(\sigma^F_{f,\mathit{eff}}-1\big)\big(1+\delta\,\sigma^F_{f,\mathit{eff}}\big)$, which is what concavity of profit in $\ln\mu_f$ requires: with markups common across the firm's products, the firm's remaining choice is the level of μ_f, and raising it scales all of its prices in proportion. It holds whenever $\sigma_i^F\le 3\sigma_n^F-1$.

exactly when

$$\left(\sigma_i^F - \sigma_n^F\right)\theta_{fn} \;>\; \delta\,\sigma_{f,eff}^F\left(\sigma_{f,eff}^F - 1\right). \tag{D.29}$$

The immigrant population and the immigrant sales share cancel from the comparison; at $\delta = 0$ any positive gap delivers a decline; and the condition tightens as δ or the level of the elasticity rises.

What the restriction costs is the other channel. The tilt of Lemma 3 needs $\sigma_i^U > \sigma_n^U$, so under $\sigma_i^U = \sigma_n^U$ the quality composition of Proposition 2 is switched off and (D.26) isolates the price channel. The two identities in (D.25) hold with both channels active.

D.4 Quality Composition

Before stating the result we fix the three objects in play. The primitive ϕ_f and the product qualities ϕ_u are held fixed, so the model moves neither. What immigration moves is the composition of the firm's sales, summarized below by $\bar{\phi}_f$, the sales-weighted average quality of what the firm sells. The appeal residual of Appendix C.4 is a third object: with one consumer type it recovers ϕ_f exactly (Remark 1), and with two it responds to the sales composition as well, so the decline in $\bar{\phi}_f$ is the model's within-firm channel for the measured decline in Φ^{Appeal}; measured appeal is a relative object and can also fall through reallocation across firms, and how much each margin accounts for is an empirical question. The proposition delivers this decline from the elasticity gap alone: the two types agree on quality, and only $\sigma_i^U > \sigma_n^U$ tilts the sales mix toward the firm's cheaper, lower-quality products.

Firm-level quality is the within-firm expenditure-weighted average $\bar{\phi}_f \equiv \sum_{u\in\Omega_f} s_{u|f}\phi_u$, with type-$h$ within-firm share $s_{uh|f}$ from (D.7) and type-aggregated share

$$s_{u|f} = \frac{L_n E_{fn} s_{un|f} + L_i E_{fi} s_{ui|f}}{L_n E_{fn} + L_i E_{fi}}, \qquad E_{fh} \equiv \phi_f \left(\frac{P_{fh}}{P_h}\right)^{1-\sigma_h^F} E_h, \tag{D.30}$$

where E_{fh} is the expenditure of a type-h consumer on firm f. This share is identically the revenue weight w_u of Proposition 4, written in the form that shows where L_i enters. Both $s_{uh|f}$ and $s_{u|f}$ sum to one over Ω_f, so $\bar{\phi}_f$ is a within-firm quality average that isolates composition from firm size.

Lemma 3 (Immigrants Tilt Toward the Firm's Cheaper Products)**.** *For two products of the same firm with $p_u > p_v$,*

$$\frac{s_{ui|f}/s_{vi|f}}{s_{un|f}/s_{vn|f}} = \left(\frac{p_u}{p_v}\right)^{\sigma_n^U - \sigma_i^U} < 1. \tag{D.31}$$

Proof. By (D.7), $s_{uh|f}/s_{vh|f} = (\phi_u/\phi_v)(p_u/p_v)^{1-\sigma_h^U}$. Taking the ratio across types cancels the quality term and leaves $(p_u/p_v)^{(1-\sigma_i^U)-(1-\sigma_n^U)} = (p_u/p_v)^{\sigma_n^U-\sigma_i^U} < 1$ since $p_u > p_v$ and $\sigma_i^U > \sigma_n^U$. □

Quality enters both types' shares in the same way and therefore drops out of the comparison. The tilt is governed by relative prices and by the within-firm elasticities alone; the across-firm terms ϕ_f and $(P_{fh}/P_h)^{1-\sigma_h^F}$ are common to all of the firm's products and cancel as well.

Proof of Proposition 2. This is the composition channel, holding equilibrium prices fixed, and all sums are over $u \in \Omega_f$. Differentiating $\bar{\phi}_f = \sum_u s_{u|f}\phi_u$ with prices and qualities fixed, and writing $D_f \equiv L_n E_{fn} + L_i E_{fi}$,

$$\frac{\partial \bar{\phi}_f}{\partial L_i} = \sum_u \frac{\partial s_{u|f}}{\partial L_i}\phi_u = \frac{E_{fi}}{D_f}\sum_u \big(s_{ui|f} - s_{u|f}\big)\phi_u = \frac{E_{fi}}{D_f}\big(\bar{\phi}_{fi} - \bar{\phi}_f\big), \qquad \bar{\phi}_{fh} \equiv \sum_u s_{uh|f}\phi_u, \tag{D.32}$$

using $\partial s_{u|f}/\partial L_i = (E_{fi}/D_f)(s_{ui|f} - s_{u|f})$. Since $\bar{\phi}_f = (1-w_i)\bar{\phi}_{fn} + w_i\bar{\phi}_{fi}$ with immigrant expenditure weight $w_i = L_i E_{fi}/D_f$, we have $\bar{\phi}_{fi} - \bar{\phi}_f = (1-w_i)(\bar{\phi}_{fi} - \bar{\phi}_{fn})$, and

$$\bar{\phi}_{fi} - \bar{\phi}_{fn} = \sum_u \big(s_{ui|f} - s_{un|f}\big)\phi_u, \tag{D.33}$$

and we show this sum is negative. Write $d_u \equiv s_{ui|f} - s_{un|f}$ for the gap between the two types' spending shares on product u, and $r_u \equiv s_{ui|f}/s_{un|f}$ for their ratio, so that $d_u = s_{un|f}(r_u - 1)$: product u takes a larger share of immigrant spending than of native spending exactly when $r_u > 1$. Order the firm's products by ϕ_u. By Proposition 6 this is also the ordering by price, and Lemma 3 then makes r_u decreasing in ϕ_u. The first step is the sign pattern of d_u. Both share vectors sum to one, so $\sum_u d_u = 0$, which in terms of the ratio says $\sum_u s_{un|f}\, r_u = 1$: the average of r_u under the native shares is exactly one. A decreasing sequence with average one starts above one and ends below it. There is therefore a quality level ϕ^* that separates the two regimes: products with $\phi_u < \phi^*$ have $r_u \geq 1$ and so $d_u \geq 0$, and products with $\phi_u > \phi^*$ have $r_u \leq 1$ and so $d_u \leq 0$. Relative to natives, immigrants overspend on the firm's low-quality products and underspend on its high-quality ones. The second step turns this pattern into the sign of the sum. Because $\sum_u d_u = 0$, subtracting the constant ϕ^* from every quality changes nothing,

$$\sum_u d_u\, \phi_u = \sum_u d_u\big(\phi_u - \phi^*\big), \tag{D.34}$$

and now every term is nonpositive: a product below ϕ^* pairs $d_u \geq 0$ with $\phi_u - \phi^* < 0$, and a product above ϕ^* pairs $d_u \leq 0$ with $\phi_u - \phi^* > 0$. If the firm's qualities are not all equal, as the proposition assumes, then prices differ across its products, r_u strictly decreases, and some term is strictly negative. Hence $\bar{\phi}_{fi} < \bar{\phi}_{fn}$, and therefore

$$\frac{\partial \bar{\phi}_f}{\partial L_i} = \frac{E_{fi}}{D_f}\,(1-w_i)\big(\bar{\phi}_{fi} - \bar{\phi}_{fn}\big) < 0. \tag{D.35}$$

□

This is a fixed-price result. Letting prices respond adds two components. The component of the price response that is common across the firm's products leaves every within-firm share unchanged, by Lemma 1; it still shifts the type weight w_i, and since $\sigma_i^F > \sigma_n^F$ a common price decline raises w_i,

which reinforces the fall in $\bar{\phi}_f$. The non-uniform component, which shifts the firm's relative prices, is the one the model does not sign, and this is why we state the result at fixed prices. The composition channel is distinct from the Sato-Vartia price channel of Proposition 1, and it operates through a different primitive. Specifically, the tilt is governed by $\sigma_i^U - \sigma_n^U$ whereas markup compression is governed by $\sigma_i^F - \sigma_n^F$, so Φ^{SV} and Φ^{Appeal} contain separate information about the mechanism.

D.5 Variety Effects

We now let firms choose which products to offer: product u is offered if and only if operating profit covers its fixed cost F_u. Under the common markup, operating profit is $R_u/\sigma^F_{f,eff}$ with revenue $R_u = p_u q_u$, and immigration moves it in two directions: markup compression lowers it, by (D.23), while demand expansion raises it. The scale force here is demand-side and operates even at $\delta = 0$, unlike the cost-side scale effect of Proposition 1. The model does not sign the net effect, which is consistent with the finding that the Feenstra variety correction does not move with demand-side immigration exposure (Table 10). We also hold the quality of each product fixed: a barcode is a fixed physical product, so quality attaches to the product rather than adjusting within it, and a quality upgrade is the entry of a new product, the margin treated in this subsection.

D.6 How Quality Enters Preferences

Two conventions for quality coexist and it matters which one we adopt. The estimation of Section 5.4, following Hottman et al. (2016), is quality-augmenting: a unit of product u delivers ϕ_u^U efficiency units, and the within-firm expenditure share is $s_{u|f} \propto (\phi_u^U)^{\sigma^U - 1} p_u^{1-\sigma^U}$, as in (C.1). The model of this appendix instead enters quality as a taste weight in (D.1), so that the share is $s_{u|f} \propto \phi_u \, p_u^{1-\sigma^U}$. With a single consumer type the two are the same system written in different variables,

$$\phi_u = (\phi_u^U)^{\sigma^U - 1}, \tag{D.36}$$

and likewise $\phi_f = (\phi_f^F)^{\sigma^F - 1}$.

With two consumer types the two conventions are no longer equivalent, because the map (D.36) involves the elasticity, which differs across types. Holding ϕ_u common across types is not the same as holding ϕ_u^U common. Under the taste-weight convention, quality enters both types' shares identically, so comparing the two types product by product, only prices distinguish them: immigrants lean toward the firm's cheaper products, whatever their qualities. Under the quality-augmenting convention, ϕ_u^U enters the shares as $(\phi_u^U)^{\sigma_h^U - 1}$, so the same quality attracts the two types with different strength, and immigrants lean instead toward the products that are cheap per unit of quality, p_u/ϕ_u^U.

We adopt the taste-weight convention for two reasons. It is the one that implements the discipline of the model: the two types agree on quality and differ only in elasticities, so every composition result is forced through the elasticity gap alone, whereas under the alternative a taste difference

re-enters through the exponent. And it rests Proposition 2 on the weaker condition: with quality out of the cross-type comparison, the result needs only the price-quality ranking of Proposition 6, whereas under the alternative a high-quality product can attract immigrants rather than repel them, and the result would survive only if price rises more than one for one with quality, $d \ln p_u / d \ln \phi_u^U > 1$. The cost is that the model and the estimating system are written in different variables, with (D.36) the bridge between them; with a single type the bridge is exact, as Remark 1 records. The ordering itself is shared: because (D.36) is increasing, the two conventions rank a firm's products identically.

A separate caution applies to both conventions. Appeal is recovered as the component of a firm's market share that its relative price does not explain, so whether that residual reflects quality, in the sense of something consumers would pay for at given prices, or a shift in tastes is not identified. For that reason we keep appeal out of the reported price index, as we state in Section 5.4, and we read the appeal index as the composition of a firm's sales rather than as a welfare-relevant quantity.