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: February 25, 2026

Abstract

This paper investigates how immigration affects consumer prices. Using scanner data and instrumenting county-level immigration with historical ancestry patterns, we find that an inflow of 10,000 immigrants lowers four-year price growth by 0.58 percentage points. Leveraging variation in firm exposure through sales versus production locations, we show price declines stem entirely from the product demand channel: firms lower prices in response to immigrants in sales markets, not production locations. Evidence suggests that immigrants search more intensively, exhibit higher demand elasticity, pay lower prices for identical products, and shift expenditure toward lower-appeal products—consistent with a model of heterogeneous price sensitivity.

JEL Codes: J61; E31; F22; L11.

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

*We thank participants at 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 and Jay Hyun 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 immigrants 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 disinflationary effect generalizes beyond CPG to 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.

Having established a robust county-level reduced-form relationship, we next investigate the underlying firm-level mechanisms, shifting our unit of analysis from the county to the firm. The aggregate county-level effect represents the net outcome of diverse firm-level adjustments, encompass-

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.

ing their responses to the changes in product demand and labor supply, firm entry/exit, and other general equilibrium forces. To disentangle these channels, we leverage our detailed data to isolate a firm's specific exposure to both immigrant consumers and workers. 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. In these analyses, we control for state fixed effects and other potential confounding variables that could affect firm pricing, such as retail wages and rental prices. 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.

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 reduce prices in their major sales destinations in response to immigrant inflows, while their price response is weakly positive when immigrants arrive in their more labor-populated production sites. We also more directly examine the county-level result that the negative price effect is particularly pronounced for the inflow of low-skilled immigrants, who are typically more price-sensitive and search-intensive. 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 lower prices for the same products relative to the national average, 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.[5] 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 available per household and

[5]At the county level, this index primarily emphasizes variety due to the normalization of appeal.

trip. While supply-side immigration exposure is associated with greater product variety, immigrant consumers do not primarily drive entry of new products, suggesting variety-based competition is not a primary channel through which immigrant demand lowers prices.

We rationalize these findings with a parsimonious theoretical framework featuring household heterogeneity in price sensitivity. Immigrants, with lower average incomes, exhibit higher demand elasticity, which induces firms to lower prices. The model generates predictions consistent with our empirical findings: price declines concentrated among low-income immigrants and products, higher demand elasticity, and expenditure shifts toward lower-appeal products.

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.[6] However, 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 explored immigration's link to consumer prices, finding mixed effects depending on the sector, country, and specific mechanism studied. Some studies document price decreases for non-traded goods and services (Cortes, 2008), in specific contexts such as refugees in Israel (Lach, 2007) and Turkey (Tumen, 2016), or through cross-country analyses of international relative prices (Zachariadis, 2012). Others find small medium-run positive effects on general price levels (Furlanetto and Robstad, 2019). A theoretical exercise suggests that the disinflationary labor supply effect and the inflationary product demand effect of immigration may largely cancel each other out, resulting in a neutral overall impact on average (Cheremukhin et al., 2024). A related but separate strand of literature has examined immigration's impact on housing markets, revealing mixed evidence. While several studies document positive price effects (Saiz, 2007; Sanchis-Guarner, 2023), others find negative or neutral long-run effects attributable to native out-migration and flight from immigrant neighborhoods (Saiz and Wachter, 2011; Monras, 2020), or downward pressure on construction wages (Sá, 2015). Our paper builds on these foundations by empirically establishing the disinflationary effect of immigration on US consumer prices and highlighting the consumer search mechanism.

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

[6] 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). 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).

(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; Mongey and Waugh, 2025; Zhu et al., 2025). An important finding we highlight is that immigration tends to reduce the price level relatively more for lower income households.[7]

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. Integrating both the labor supply and consumer demand is in spirit of recent papers analyzing both effects of immigration inflow (Hong and McLaren, 2015; Albert and Monras, 2022; Mahajan, 2024; Galaasen et al., 2025).

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. This methodological rigor is necessary because price effects may be reflected in goods and services available to consumers (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 data construction and empirical strategy, including the instrumental variable approach and the construction of firm-specific immigration exposure measures. 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

[7] In contrast, negative wealth shocks and product innovation targeted at high-income consumers tend to generate higher inflation for low-income households (Jaravel, 2019; Argente and Lee, 2021).

production locations, revealing the critical role of consumer search behavior and quantifying the role of demand elasticity, product variety, and product appeal. 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
	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). Beyond wages, immigration can enhance productivity through task specialization, as natives shift toward communication-intensive tasks while immigrants concentrate in manual tasks (Peri and Sparber, 2009). Immigrants also contribute to entrepreneurship and firm entry (Mahajan, 2024; 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 (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 firms adjust prices when exposed to immigration through two distinct channels: where they sell their products (product demand) and where they produce them (labor supply). 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 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. Crucially, we use only a balanced panel of stores within each county to ensure that changes in store composition do not mechanically drive our price measures. This county-UPC-year panel forms the foundation for all subsequent price index construction, which we describe in detail 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 2006 to 2018. Panel participants record every shopping trip using handheld scanners or mobile apps, providing comprehensive information on their shopping patterns. 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 specific store brands purchased, and the purchases made at each retailer. Importantly, the panel also includes detailed demographic characteristics such as household income, education, household size, race and ethnicity, and in some cases proxies for immigration status or nativity.

This household-level data serves two critical purposes in our analysis. First, it allows us to examine demand-side mechanisms by observing how shopping behavior—particularly search intensity measured by store visits and shopping frequency—varies with local immigration inflows. If immigrants are more price-sensitive and engage in more intensive search, we should observe changes in the composition of shoppers and their behavior in high-immigration counties. Second, the demographic information helps us identify likely immigrant households or at least households with characteristics correlated with immigrant status, enabling us to test whether different demographic groups exhibit systematically different price sensitivity and shopping patterns. These behavioral differences are central to understanding the price sensitivity mechanism we document in Section 5.

Supplementary Price Data Across Sectors While our main analysis focuses on CPG products where we have the richest micro-data, we supplement our scanner-based price indices with additional measures covering other major expenditure categories to assess whether immigration's effects extend beyond CPG products. For housing purchase prices, we use county-level home price indices from CoreLogic and Zillow covering 2006-2018. For rental housing, we construct county-level median gross rent indices from American Community Survey (ACS) data, which provides broader geographic coverage than other rental price sources and covers approximately 3,100 counties. For durable goods and services, we use city-level Consumer Price Index data from the Bureau of Labor Statistics, available for the 27 largest metropolitan statistical areas (261 counties). While these supplementary price measures have more limited geographic coverage or product detail than our scanner data, they allow us to assess the generality of our findings across the full consumption basket and provide important context for interpreting the magnitude of effects we document for consumer packaged goods.

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 employ 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 Google search for the firm name and state returns the same URL 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's accuracy and find a success rate of 94%, confirming that our 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 allow us to 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 employed at each establishment, and the counties where the product is sold. This dual observation of production locations 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 versus product demand at sales counties. Additional details on the matching algorithm, validation procedures, match rates, and representativeness of matched versus unmatched firms are provided in Appendix A.1.

3.1.3 Immigration Data

We construct county-level immigration measures following the methodology of Terry et al. (2024), using microdata from the Integrated Public Use Microdata Series (IPUMS). For historical periods essential to our instrumental variable strategy, we use the 1880, 1900-1930, and 1970-2000 decennial census samples to measure immigration and ancestry patterns. For our main analysis period, we extend their approach using American Community Survey (ACS) data: the 2010 ACS for immigration through 2005, the 2014 ACS for 2005-2010 immigration, the 2018 ACS for 2010-2014 immigration, and the 2022 ACS for 2014-2018 immigration. We deviate slightly from Terry et al. (2024) in our use of later ACS samples because IPUMS reports five-year pooled samples after 2000, and using the 2010 ACS alone would undercount early-period immigrants. For example, the 2010 ACS pools surveys from 2006-2010, meaning 2010 arrivals appear only in the final survey year. Using the 2014 ACS, which includes surveys from 2010-2014, more accurately captures all 2005-2010 immigration. We apply this logic consistently for subsequent periods. The 2022 ACS is the most recent data available at the time of our analysis, which is why our sample ends in 2018.

Immigration is defined as the total number of foreign-born respondents by birthplace. We construct immigration flows for multiple time periods using reported year of immigration. Ancestry is measured from survey 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.

Beyond measuring total immigration flows, the ACS data also provide rich information on immigrant characteristics at the county level, including educational attainment and skill levels, industry and occupation of employment, income and earnings, and year of arrival and duration in the United States. These characteristics allow us to examine heterogeneous effects by immigrant type, such as comparing the price effects of high-skilled versus low-skilled immigration, which we explore in our mechanism analysis.

3.1.4 Supplementary County-Level Data

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

These supplementary data serve two main purposes throughout our analysis. They provide control variables for our main specifications, capturing local economic conditions that might correlate with both immigration and prices. Also, they enable us to examine whether immigration affects local economic conditions such as wages and employment in ways that might confound or mediate the price effects we document.

3.2 Baseline Empirical Specification

Our baseline empirical 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} measures the total number of new immigrants arriving in county d during period t, δ_t and δ_s are state and period fixed effects, and X_{dt} includes additional 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 concerns about endogeneity, which we discuss in detail below.

We estimate this specification using a stacked long-difference approach covering three periods: 2006-2010, 2010-2014, and 2014-2018; we also conduct a robustness check using the full long-difference period (2006-2018). The baseline specification pools the data across periods, focusing on medium-term adjustments while mitigating concerns about high-frequency noise in annual price changes and allowing sufficient time for firm-level adjustments to take effect. 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 CPG non-durables as our primary focus and supplementing this with housing purchase prices, housing rents, durable goods, and services, to conduct the most comprehensive analysis of immigration's overall price impact across the full consumption basket as is empirically feasible.

3.3 Construction of the CPG Price Index

To properly measure how immigration affects CPG prices, we construct a price index from our barcode-level scanner data that accounts for consumer substitution patterns across products and firms. Our approach builds on the exact price index literature (Sato, 1976; Vartia, 1976) and uses chain-weighted geometric means that are consistent with constant elasticity of substitution (CES) preferences. 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 denote P_{kdt} as the price of product (UPC) k in county d at time t. Product k is produced by firm f and belongs to product group g such as carbonated beverages, breakfast cereals, or cleaning products. Our county-level price index is constructed through a three-stage aggregation using chain-weighted geometric means.

Stage 1: Product to Firm-Group Level We first aggregate products within each firm-group-county combination, using only products available in both periods $t-1$ and t to focus on continuing varieties:

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

where $\Omega_{fgdt,t-1}$ denotes the set of continuing products and w_{kdt} are Sato-Vartia weights:

$$w_{kdt} = \frac{s_{kdt} - s_{kd,t-1}}{\ln s_{kdt} - \ln s_{kd,t-1}} \Bigg/ \sum_{k' \in \Omega_{fgdt,t-1}} \frac{s_{k'dt} - s_{k'd,t-1}}{\ln s_{k'dt} - \ln s_{k'd,t-1}} \tag{3.3}$$

where s_{kdt} is the expenditure share of product k within firm-group fg in county d at time t, computed over the set of continuing products $\Omega_{fgdt,t-1}$. Sato-Vartia weights are exact for CES utility and provide the theoretically appropriate aggregation of price relatives (Diewert, 1976). These weights account for the fact that consumers substitute toward relatively cheaper products, giving more weight to products with larger expenditure shares and appropriately capturing substitution patterns.

Stage 2: Firm-Group to Group Level We aggregate across firms within each product group and county, again using Sato-Vartia weights and restricting to continuing firms:

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

where w_{fgdt} are Sato-Vartia weights based on firm-group expenditure shares within each product group.

Stage 3: Group to County Level Finally, we aggregate across product groups using Törnqvist weights, consistent with a Cobb-Douglas upper nest that allows for different consumption intensities across product categories:

$$\frac{P_{dt}}{P_{d,t-1}} = \prod_{g \in \Omega_{dt,t-1}} \left(\frac{P_{gdt}}{P_{gd,t-1}} \right)^{w_{gdt}} \tag{3.5}$$

where $w_{gdt} = \frac{1}{2}(s_{gdt} + s_{gd,t-1})$ are Törnqvist weights based on the average expenditure share of group g in county d across the two periods.

This three-stage procedure yields our baseline CPG price index P_{dt}, which measures inflation in county d from period $t-1$ to t while holding the set of products, firms, and groups constant by focusing on continuing varieties only. The log change $\Delta \ln P_{dt} = \ln(P_{dt}/P_{d,t-1})$ enters equation (3.1) as our primary outcome variable. By construction, this index measures price changes for an identical basket of goods across locations and time, capturing pure price variation rather than changes in product mix or quality. This construction is important for interpreting our results: the price declines we document in Sections 4 and 5 cannot be explained by immigrants simply purchasing different, cheaper products. [8]

We undertake two complementary exercises related to the price index: first, for a sanity check, we aggregate the county-level index to the national level and compare it to the official BLS food price index, finding the two aggregate inflation measures are reasonably close; second, to demonstrate the robustness of our final results, we show they are qualitatively unchanged when using alternative index formulas (Törnqvist, Jevons, Laspeyres, and Paasche) in place of the Sato-Vartia weights (both sets of results are detailed in Appendix A).

3.4 Instrumental Variable Strategy

The immigration measure I_{dt} in equation (3.1) is likely endogenous. Immigrants may select into counties based on economic opportunities, wage levels, employment prospects, housing affordability, or amenities, generating spurious correlations between immigration and prices. For example, if

[8] If immigration shifted consumption toward lower-quality goods without changing what firms charge for identical items, this compositional effect would not appear in our SV index. The structural demand model in Section 5.4 *separately* estimates a product appeal index that captures this quality composition channel.

immigrants are attracted to counties experiencing positive economic shocks that also drive up local prices through increased demand, OLS estimates of γ will be biased upward. Conversely, if immigrants select into declining areas with falling prices due to lower housing costs, OLS will be biased downward. These selection concerns motivate our instrumental variable approach.

To address this endogeneity, we employ an instrumental variables strategy that leverages historical settlement patterns from more than a century ago and origin-country migration patterns to generate plausibly exogenous variation in contemporary immigration flows. Our approach closely follows Terry et al. (2024), who extend and refine the canonical shift-share instrument developed by Card (2001). We provide a detailed exposition of the methodology below, with care to explain the identifying assumptions and provide intuition for how the instrument generates exogenous variation.

3.4.1 The Standard Shift-Share Instruments

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

$$\hat{I}_{odt} = A_{od,t-1} \times I_{o,-d,t} \tag{3.6}$$

where $A_{od,t-1}$ is the ancestry network, measuring the number of residents with ancestry from country o living in county d in period $t-1$, and $I_{o,-d,t}$ is the shifter, measuring total immigration from country o to all other counties except d in period t. The identifying assumption is that historical ancestry networks predict where new immigrants settle due to family ties, established ethnic communities, and information networks, but are orthogonal to current local economic shocks affecting prices. One concern in using this instrument is that persistent local shocks can contaminate the use of past ancestry as an exposure measure. If the same economic conditions that attracted immigrants to county d in the past persist into the present, then past ancestry $A_{od,t-1}$ is correlated with current unobserved shocks to county d that also affect prices. This persistence violates the exclusion restriction, making the instrument fails to isolate truly exogenous variation in immigration flows.

3.4.2 Predicted Ancestry Instrument

To address this endogeneity problem, we follow Terry et al. (2024) and instrument for ancestry itself using exogenous historical push and pull factors starting from 1880. The key insight is to predict county-level ancestry stocks using variation that arises from the interaction of origin-country emigration patterns and destination-county attractiveness to unrelated immigrant groups. This two-way interaction breaks the link between current economic conditions in a county and the instrument by relying on historical factors that are plausibly orthogonal to present-day local shocks.

The first stage predicts ancestry A_{odt} in recent periods using historical migration patterns:

$$A_{odt} = \delta_{or(d)t} + \delta_{c(o)dt} + X'_{od}\zeta + \sum_{\tau \in \mathcal{T}} a_{r(d)\tau} \left(I_{o,-r(d),\tau} \times \frac{I_{Europe,d\tau}}{I_{Europe,\tau}} \right) + \nu_{odt} \tag{3.7}$$

where $\mathcal{T}$ denotes the set of historical census years (1880, 1900, 1910, 1920, 1930, 1970, 1980, 1990, 2000) used to construct the instrument. $I_{o,-r(d),\tau}$ is total immigration from origin o in year τ, excluding migration to the census region $r(d)$ where county d is located. We use nine U.S. census divisions to define regions. This leave-out measure captures exogenous push factors from the origin country, such as wars, famines, economic crises, or political upheavals that drove emigration more than a century ago. The term $\frac{I_{Europe,d\tau}}{I_{Europe,\tau}}$ is the share of all European immigrants in year τ who settled in county d. This captures the pull attractiveness of destination d to an unrelated immigrant group, reflecting factors like transportation networks, railroad connections, port access, economic opportunities, climate, and established communities that made certain destinations more attractive in the late 19th and early 20th centuries.

The specification also includes $\delta_{or(d)t}$, which are origin-country by destination-region by time fixed effects, and $\delta_{c(o)dt}$, which are continent-of-origin by destination-county by time fixed effects. These fixed effects control for broad geographic patterns in settlement. The term X'_{od} includes time-invariant origin-destination controls such as geographic distance between the origin country and destination county, whether the origin country had historical colonial ties to the region, and other dyadic characteristics.

The interaction between origin-specific push factors and destination attractiveness to Europeans generates plausibly exogenous variation in where non-European immigrants' descendants live today. The exclusion restriction requires that these historical factors do not directly affect current price dynamics in 2006-2018 except through their effect on current immigration patterns. This is plausible because the economic, transportation, and institutional factors that shaped settlement patterns more than a century ago are unlikely to have persistent direct effects on contemporary CPG prices after controlling for broad geographic patterns through fixed effects.

Intuition Through an Example Consider predicted ancestry of Chinese immigrants in Butte County, Montana. This ancestry measure is high when two conditions hold. First, many Chinese emigrants left China in 1880 due to economic hardship and political instability, settling throughout the United States—in California cities, East Coast ports, and elsewhere. This large-scale Chinese emigration, measured using settlement outside the Mountain West region to avoid mechanical correlation with Montana, indicates strong origin-country pressures driving people to leave China. Second, Butte County attracted a large share of European immigrants in 1880 due to its copper mining boom, railroad access, and established economic opportunities—factors that made it an attractive destination for immigrants generally, not just Europeans. The instrument interacts these two components: origin-country emigration pressures (the "push") with destination attractiveness to an unrelated immigrant group (the "pull"). Crucially, this predicted ancestry arises entirely from historical factors in the late 19th century—origin-country conditions and destination characteristics like mining activity and transportation networks—that are plausibly unrelated to contemporary economic conditions affecting CPG prices in 2006–2018.

We obtain predicted ancestry as:

$$\hat{A}_{odt} = \sum_{\tau \in \mathcal{T}} \hat{a}_{r(d)\tau} \left(I_{o,-r(d),\tau} \times \frac{I_{Europe,d\tau}}{I_{Europe,\tau}} \right)^{\perp} \tag{3.8}$$

where $\hat{a}_{r(d)\tau}$ are the estimated coefficients from equation (3.7) and the superscript $\perp$ indicates that the interaction term has been residualized with respect to all controls included in equation (3.7).

3.4.3 Final Instrument Construction

With predicted ancestry in hand, we return to the shift-share logic. The final first-stage equation predicts actual immigration I_{odt} as:

$$I_{odt} = \delta_{or(d)} + \delta_{c(o)d} + \delta_t + X'_{od}\theta + b_t \left(\hat{A}_{od,t-1} \times \tilde{I}_{o,-r(d),t} \right) + u_{odt} \tag{3.9}$$

where $\tilde{I}_{o,-r(d),t}$ is the shifter, measuring scaled immigration from origin o to regions outside $r(d)$ in period t. The shifter is defined as:

$$\tilde{I}_{o,-r(d),t} = I_{o,-r(d),t} \times \frac{I_{Europe,r(d),t}}{I_{Europe,-r(d),t}} \tag{3.10}$$

where the second term is a scale factor that corrects for regional size differences. Larger regions naturally receive more immigrants in absolute numbers, so we scale by the region's share of European immigration to make shifters comparable across regions of different sizes.

The instrument for county-level immigration is constructed by aggregating across origin countries:

$$Z_{dt} = \sum_{o} \hat{b}_t \left(\hat{A}_{od,t-1} \times \tilde{I}_{o,-r(d),t} \right) \tag{3.11}$$

where $\hat{b}_t$ are the estimated coefficients from equation (3.9).

This instrument generates variation in immigration to county d in period t through three components, all plausibly orthogonal to current local economic conditions: historical settlement patterns embedded in predicted ancestry, current origin-country emigration rates in the leave-out region reflecting global migration trends, and historical destination-region attractiveness to Europeans. The exclusion restriction requires that these factors affect current county prices only through their effect on current immigration to that county, not through any direct channel such as persistent effects of historical mining or agriculture on current price dynamics.

3.4.4 Implementation and First-Stage Validation

We implement this procedure for three periods in our sample: 2006-2010, 2010-2014, and 2014-2018, using only non-European origin countries consistent with the instrument construction. The instrument is constructed from non-European immigration because European migration patterns are used to measure destination attractiveness in equation (3.7). Our estimated effects, therefore, represent the impact of non-European immigration specifically. In practice, non-European immigrants constitute the vast majority of recent U.S. immigration—over 80% of total inflows during our sample period—so this is the policy-relevant margin.

For the first-stage estimation in equation (3.9), we include nine time periods spanning 1975-2018 in five-year intervals, with four-year intervals after 2010 to accommodate data availability as described in Section 3.1. Appendix Table B.3 reports our first-stage estimates for post-1975 periods, which closely replicate those in Terry et al. (2024) for overlapping years and extend the analysis through 2018. Our constructed instruments for overlapping periods correlate at 0.997 with those provided by Terry et al. (2024), validating our replication and extension of their methodology. The first stage is strong: the F-statistic on the excluded instrument exceeds 40 in all specifications, well above the conventional thresholds for concerns about weak instruments. This reflects the substantial predictive power of historical settlement patterns and origin-country push factors for current immigration flows.[9]

3.5 Summary Statistics and Spatial Variation in Prices

Before presenting our formal empirical results, we document a key feature of our empirical 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 analysis and our main firm-level investigation of differential 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, constructed using our Sato-Vartia price index described in Section 3.3. Even accounting for product entry, exit, and changing expenditure shares within counties, we observe considerable geographic price dispersion. The distribution shows substantial variation: while the median county experienced price increases of approximately 6.7% over four-year periods, the interquartile range spans 5.6 percentage points (from 4.3% at the 25th percentile to 9.9% at the 75th percentile). Counties at the 75th percentile experienced inflation more than twice as high as those at the 25th percentile, indicating substantial variation in local price

[9] Standard shift-share instruments may require corrected standard errors or additional conditions to ensure valid inference (Adão et al., 2019; Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022). Our instrument uses predicted ancestry rather than actual settlement shares as the exposure measure, following Terry et al. (2024), who conduct placebo simulations following Adão et al. (2019) and show that this approach yields rejection rates close to the nominal 5% level, unlike instruments based on realized immigration or ancestry shares.

dynamics.

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

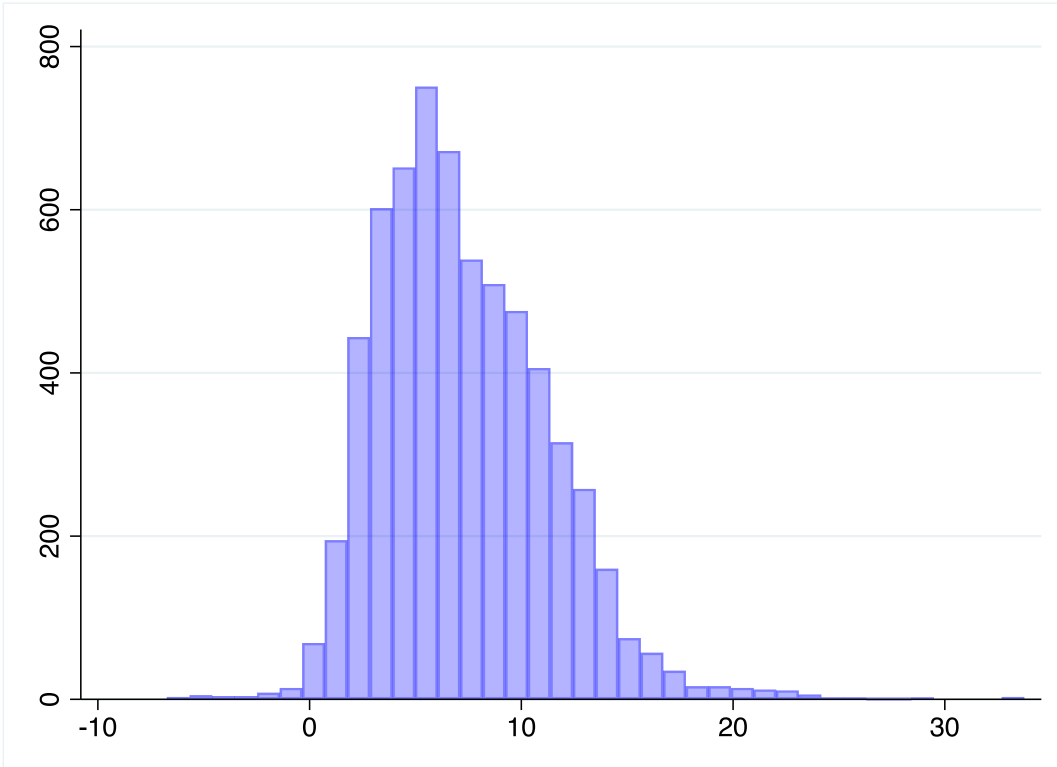

Notes: Histogram shows the unweighted distribution of price changes across counties over four-year periods from 2006-2018, measured using the Sato-Vartia price index for CPG products as described in Section 3.3. Each observation is a county-period (N=6,339). The price index uses a three-stage aggregation procedure that accounts for product entry and exit and allows expenditure shares to vary over time. The distribution demonstrates substantial spatial price variation for consumer goods.

This violation of purchasing power parity for tradable goods is essential for our empirical strategy. If purchasing power parity held perfectly, local immigration shocks could not affect local prices independently, as any price differentials would be immediately arbitraged away through trade. Several factors contribute to these persistent price differences. Manufacturers and retailers engage in spatial price discrimination, setting different prices in different markets based on local demand conditions and competitive environments (Fitzgerald et al., 2024). While some retailers adopt uniform pricing for identical products within their chains (DellaVigna and Gentzkow, 2019), our price index aggregates across both products and firms, capturing variation in pricing strategies across manufacturers and retail formats. Local promotions, discounts, and sales vary substantially across locations, reflecting differences in local competition and consumer composition. Local market structure, including the number and type of competing retailers, differs geographically and affects the effective prices consumers face. Transportation costs and local supply chain configurations introduce additional wedges between locations. This empirical reality—that consumer goods exhibit substantial spatial price variation—validates our approach of exploiting county-level variation in immigration exposure to study price effects. The observed variation in price changes across counties highlights that local conditions have a significant influence on price determination.

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

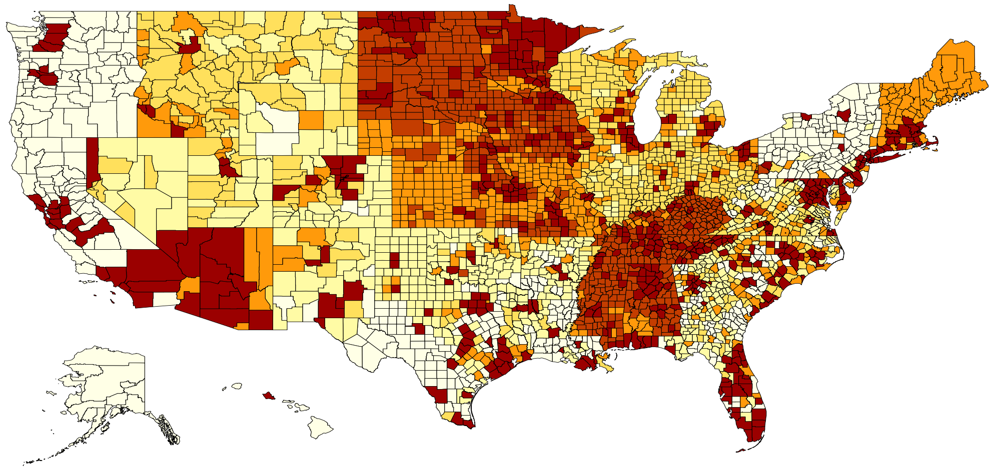

Notes: Heat map shows predicted immigration shock Z_{dt} from equation (3.11) at the county level, cumulated over 2006-2018. Darker colors indicate larger predicted immigration inflows. The instrument uses only non-European immigration and is based on historical settlement patterns from 1880-1930 interacted 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 heat map reveals substantial geographic variation in immigration exposure that is plausibly orthogonal to local economic conditions due to our instrumental variable construction (described in Section 3.4). Dark red areas indicate the highest predicted immigration inflows and are concentrated in the Southwest, California, and Florida, but also appear throughout the country including the North-Central region, Appalachia, and the East Coast. This widespread geographic variation, combined with the exogenous source of the instrument, provides strong identifying power for our analysis.

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 (2006-2010, 2010-2014, 2014-2018). For CPG products, the mean price increase of 7.28% with a standard deviation of 4.05 percentage points confirms the substantial spatial dispersion documented in Figure 1.

Examining price variation across sectors reveals patterns consistent with economic theory regarding tradability. While CPG products show considerable spatial dispersion, non-tradable housing exhibits even more dramatic variation: purchase prices have standard deviations of 18.9

Table 2: Summary Statistics: County Level

	count	mean	sd	p25	p50	p75	min	max
Δ ln $\mathrm{P_{d,t}}$ (CPG)	6339	7.35	4.12	4.33	6.74	9.95	-6.70	35.4
Δ ln $\mathrm{P_{d,t}}$ (Non-durables)	684	5.96	5.40	-0.18	8.14	10.1	-5.19	13.4
Δ ln $\mathrm{P_{d,t}}$ (Durables)	684	-3.63	3.35	-6.42	-3.69	-1.57	-9.81	3.48
Δ ln $\mathrm{P_{d,t}}$ (Services)	684	9.22	2.59	7.78	8.71	10.1	2.95	17.5
Δ ln $\mathrm{P_{d,t}}$ (BLS, Rent)	684	11.2	4.77	8.16	10.1	14.1	2.09	25.8
Δ ln $\mathrm{P_{d,t}}$ (ACS, Rent)	6336	10.7	6.45	7.31	10.9	14.3	-29.9	53.5
Δ ln $\mathrm{HP_{d,t}}$(CL)	2446	-3.85	18.9	-11.8	-0.74	7.56	-87.9	44.6
Δ ln $\mathrm{HP_{d,t}}$ (Z)	4781	6.03	16.2	-1.97	6.26	15.9	-108.6	57.2
Δ # $\mathrm{Establ._{d,t}}$	6339	-0.15	0.72	-0.45	-0.14	0.15	-8.44	22.4
Δ $\mathrm{Employment_{d,t}}$	6339	1.77	15.0	-0.44	0.088	1.08	-305.0	297.9
Δ ln $\mathrm{Wage_{d,t}}$	6339	3.27	5.30	0.36	3.17	6.00	-30.2	56.2
Δ ln $\mathrm{Wage_{d,t}}$ (Retail)	6318	1.96	6.85	-1.85	2.04	5.87	-74.6	64.7
Δ $\mathrm{Wage_{d,t}}$	6339	13.1	23.3	1.41	12.2	23.2	-163.8	288.0
Δ $\mathrm{Wage_{d,t}}$ (Retail)	6339	4.79	22.6	-4.65	5.03	14.7	-288.1	452.6
$\mathrm{Immigr_{d,t}}$	6339	1.68	4.95	0.066	0.17	0.62	0.0026	34.2
$\mathrm{Immishock_{d,t}}$	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)-(3.5). 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 (3.11), constructed using residualized predicted ancestry. Both Immigr and Immishock winsorize top and bottom 1% of the sample.

and 16.2 percentage points for CoreLogic and Zillow respectively, reflecting both housing's non-tradable nature and the boom-bust cycle during our sample period (which motivates robustness checks excluding the Great Recession years). Rental housing similarly shows large spatial variation with standard deviations of 4.77-6.45 percentage points. We also observe BLS price data for 27 metropolitan areas: non-durables and durables show standard deviations of 5.40 and 3.35 percentage points respectively, while services show relatively modest dispersion of 2.59 percentage points.[10]

Wages show a standard deviation of 5.30 percentage points—comparable to CPG price dispersion. Despite labor facing higher mobility costs than goods, overall wage dispersion is similar to price dispersion, likely reflecting institutional wage compression from minimum wage laws, union contracts, and standardized pay scales at large employers. Retail wages exhibit greater relative dispersion (in percent term), consistent with more localized labor markets. We also report wages in levels (in hundreds of 2012 dollars): the mean four-year wage change is 13.1 with a standard deviation of 23.3. Following Terry et al. (2024), we use this level measure in our regression analysis to facilitate direct comparison with their findings on immigration's wage effects.

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

Immigration patterns confirm the geographic variation shown in Figure 2. The average county received 2,190 immigrants per four-year period, but with a standard deviation of 9,930—more than four times the mean—reflecting extreme heterogeneity in settlement patterns. The predicted immigration shock Z_{dt} shows substantial variation that, combined with its exogenous construction from historical settlement patterns, provides strong identifying power for this analysis.

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 on the regional cost-of-living, despite modest increases in rental housing prices. 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 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.

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.

[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.10). Consistent with Terry et al. (2024) and the broader literature, we find that wages increase more in counties receiving higher-educated immigrants.

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

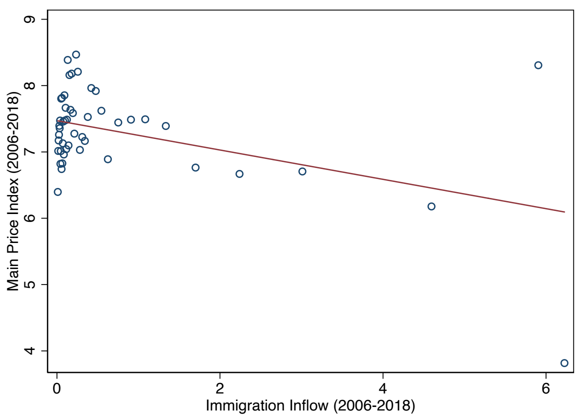

(a) OLS: Actual Immigration

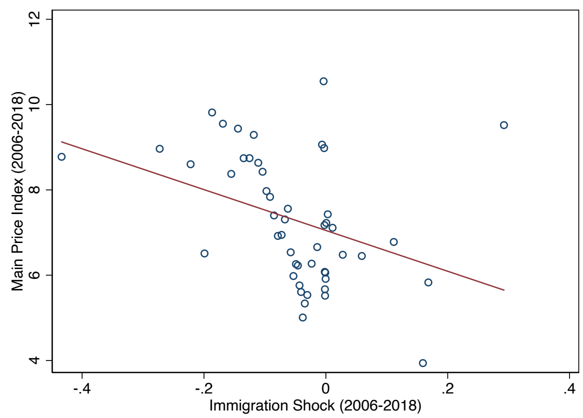

(b) Reduced Form: Immigration Shock

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 (3.11). 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. Each point represents a county, with point size proportional to initial county population. Fitted lines are from bivariate regressions without controls or fixed effects.

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 (3.11). 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. Standard errors clustered by state. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

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 negative effects, with magnitudes broadly consistent with our baseline Sato-Vartia specification. Tables B.5–B.7 present further robustness checks—including long-difference specifications, an expanded product sample (requiring only 6 rather than 12 years of presence), an alternative instrument that normalizes immigration shocks by initial county population, controls for Mexican border proximity and local ICE enforcement, state-by-period fixed effects, removal of winsorization, and county fixed effects—with remarkably stable results throughout. See appendix for further details.

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 among 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

[12] See 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.

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				
$\text{Immigr}_{\text{d,t}}$ (over25)	-0.101***	-0.209***		
	(0.024)	(0.077)		
$\text{Immigr}_{\text{d,t}} \times \text{EducYears}_{\text{d,t}}$		0.110*		
		(0.063)		
$\text{Immigr}_{\text{d,t}}$ (No HS Degree)			-0.059**	
			(0.025)	
$\text{Immigr}_{\text{d,t}}$ (Some College)			0.019	
			(0.014)	
$\text{Immigr}_{\text{d,t}}$ (Low Inc. Occ.)				-0.045
				(0.028)
$\text{Immigr}_{\text{d,t}}$ (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)
$\text{Immigr}_{\text{d,t}}$ (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. 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 estimate the effect of total immigration on the 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%. Standard errors clustered by state. ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

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 the disinflationary effect concentrated 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.[13] Immigration lowers the price of both low and high-income goods, with a somewhat larger effect for low-income products, with the difference statistically significant (p-value of 0.019). This differential effect likely reflects the fact 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, 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.

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

[13] 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.

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)
$\text{Immigr}_{\text{d,t}}$	-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 (3.11) 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. ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

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.

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).[14]

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 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 (3.11) 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 the expansion of local economic activity documented in the literature (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

[14] Table B.9 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, but in all specifications the direction of aggregate prices in response to immigration is negative.

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.10 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 Summary

The pattern of results documented in this section points toward a search-based demand mechanism. The disinflationary effects are concentrated among low-income, less-educated immigrants—who are typically more price-sensitive—and are larger for products disproportionately consumed by low-income households. County-level evidence on shopping behavior further supports this mechanism: in high-immigration counties, search intensity increases among likely immigrant households but not among native households (see Appendix Table B.11).[15] At the same time, wages do not decline and retail wages are essentially unaffected, making supply-side cost reductions an unlikely explanation. We test this mechanism directly in the next section.

5 Unpacking the Mechanism: Firm-Level Analysis

The county-level analysis establishes a robust reduced-form relationship between immigration and prices, but cannot separately identify the demand and supply channels through which immigration affects firm pricing decisions. This is because what is sold in a county is not necessarily produced there, and because county-level estimates conflate partial equilibrium firm responses with local general equilibrium effects such as establishment entry, consumer brand switching, and cross-county migration. To directly test the consumer search mechanism suggested by the county-level patterns, we turn to firm-level analysis, exploiting variation in firms' exposure to immigration separately through their sales locations (demand) and production locations (supply). This allows us to isolate the product demand channel and examine how firms respond to changes in the composition of their consumer base.

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):

$$\Delta \ln P_f = \beta_0 + \beta_1 \mathrm{I}_f^{\text{demand}} + \beta_2 \mathrm{I}_f^{\text{supply}} + \mathbf{X}'_f \gamma + \alpha_g + \varepsilon_f \tag{5.1}$$

[15] Since Hispanic individuals constitute a disproportionate share of recent immigrants, we use Hispanic households as a proxy for immigrant consumers. Hispanic households in high-immigration counties show significant increases in stores visited and coupon usage, while white households show no such changes.

where $\Delta \ln P_f$ is a change in the log price index of firm f, $\mathrm{I}_f^{\text{demand}}$ is the immigration inflow in the areas where firms have sold their products, $\mathrm{I}_f^{\text{supply}}$ is the immigration inflow in the areas where firms have their plants, $\mathbf{X}_f'$ is the set of firm-specific control variables, such as the initial firm revenue and employment, and we control for industry-specific fixed effects. $\mathrm{I}_f^{\text{demand}}$ and $\mathrm{I}_f^{\text{supply}}$ are measured as follows:

$$\mathrm{I}_f^{\text{demand}} = \sum_d \omega_{fd}^{\text{sales}} \mathrm{I}_d \tag{5.2}$$

$$\mathrm{I}_f^{\text{supply}} = \sum_d \omega_{fd}^{\text{empl}} \mathrm{I}_d \tag{5.3}$$

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.[16]

5.2 Demand vs. Supply Channels

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.

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. 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 when the two are conflated; symmetrically, omitting demand exposure attenuates the supply-side effect for the same reason. Column (4) adds firm-level controls including interactions of year dummies with firms' exposure to urban indicators and GDP per capita, 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; since these variables may themselves respond to immigration, we treat this as a robustness check, but the estimates are again similar. Column (6) shows that OLS estimates are attenuated toward zero relative to IV, consistent

[16] 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.

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. ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

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. We investigate the demand-side mechanism more directly in the next section by examining how immigration affects consumer search behavior and demand elasticities at the firm level. Appendix Table B.12 confirms results hold with 12-year long-differences.

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,

Table 8: Firm-Level Search Intensity

	Δ lnPurChar^{SV,Stack} (2006-2018)					
	(1) N. Trips	(2) N. Days	(3) N. Stores	(4) N. Products	(5) % Store Brand	(6) Rel. Price
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)
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: The regression is weighted by the initial sales. Missing shares are controlled. We also include industry (SIC 3-digit) and time-period fixed effects (each firm is assigned a main SIC industry in the NETS data). Standard errors are clustered by SIC 3-digit industries. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

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 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:

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

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

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

where ϕ^{F}_{fgct} is the appeal of firm f, σ^{F}_{ℓ} is the elasticity of substitution across firms, ϕ^{U}_{ut} is the appeal of product u, and σ^{U}_{ℓ} is the elasticity of substitution across products within a firm. The subscript ℓ indexes estimation cells; in our implementation, elasticities vary by state, product group, and time period.

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^{F}_{ℓ} and σ^{U}_{ℓ} where ℓ indexes state-group-period cells.[17]

Variety-adjusted Price Index and Firm Appeal. Given the estimated elasticities, we can construct three objects that are not recoverable from reduced-form analysis.

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

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

where $\Omega_{fgct,t-1}$ is the set of products available in both t and $t-1$, and w_{ut} are Sato-Vartia weights based on expenditure shares. Abstracting from appeal and product entry and exit, aggregating across products within a firm using Sato-Vartia weights yields the firm-level price index used in Section 5.2; further aggregating across firms within a group using Sato-Vartia weights, and then across groups using Cobb-Douglas shares, yields the county-level price index used in Section 4.

[17]Periods correspond to 2006–2010, 2010–2014, and 2014–2018.

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 new products enter and gain market share, $\lambda_t < \lambda_{t-1}$, which lowers the variety-adjusted price index since $\sigma_\ell^U > 1$. The intuition is that consumers benefit from substituting toward new, preferred varieties.

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. Unlike the price index, appeal does not directly affect the cost of living but operates as a separate quality shifter that we analyze independently. 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-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

Table 9: Demand Elasticities

	Demand Elast. (σ)			
	(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
Aggregation	County-Level	County-Level	State-Level	State-Level

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. 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. ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

demand elasticities, which compresses firm markups.

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
$\text{Immigration}_{f,d,t}$ (D)	-0.051**	-0.057***	-0.005	-0.119***
	(0.020)	(0.020)	(0.005)	(0.019)
$\text{Immigration}_{f,d,t}$ (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). 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. ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

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 where a positive number indicates an increase in variety, and *Appeal* reports changes in average firm-level appeal.

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 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, is relatively more important on the supply side than the demand side. The appeal index also increases (Column 4). These variety and appeal effects are consistent with evidence that firms innovate and upgrade quality in response to immigrant labor (Kerr and Lincoln, 2010; Terry et al., 2024). The joint pattern of higher prices, greater variety, and higher appeal on the supply side is most consistent with a firm-level "upgrading" mechanism: 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.

6 Theoretical Framework

The structural estimation in the previous section documents that immigration reduces prices of continuing products (the Sato-Vartia index Φ^{SV}) and lowers firm-level quality (Φ^{Appeal}) through the demand channel, which drives the bulk of the price changes. This section develops a theoretical framework with household heterogeneity to rationalize these findings. We interpret the appeal index as reflecting the quality composition of sales—the expenditure-weighted average quality of products sold. The key insight is that these two patterns emerge from a single primitive: immigrants are more price-sensitive than natives. For expositional simplicity, we present a flat CES demand structure with monopolistic competition and fixed product variety, abstracting from supply-side channels. Complete derivations and extensions—including nested CES with Bertrand competition and endogenous variety—appear in Online Appendix D.

6.1 Environment

Consumers. There are two consumer types: natives (n) and immigrants (i), with populations L_n and L_i. Consumers have CES preferences over differentiated products with type-specific elasticity of substitution $\sigma_h > 1$.

We assume immigrants are more price-sensitive than natives: $\sigma_i > \sigma_n$. Our structural estimates directly support this assumption, showing that demand elasticities are higher in markets with larger immigration inflows (Table 9). This likely reflects immigrants' lower incomes: the price-reducing effect of immigration concentrates among low-skilled immigrants and low-income products (Table 4), and these immigrants search more intensively, making more shopping trips and paying lower prices for identical products (Table 8). These patterns are consistent with Faber and Fally (2022) and Handbury (2021).[18]

Firms. Each firm f produces a set of products Ω_f with heterogeneous qualities. Product u has quality ϕ_u and price p_u. The firm faces marginal cost $MC_u = a_u q_u^{\delta}$, where $\delta > 0$ captures decreasing returns to scale. Each firm is a price-setter for its products but is small relative to the market, taking aggregate price indices and expenditures as given.

We take product variety and quality as fixed. Two features of our empirical setting motivate this assumption for the demand side: immigrants do not significantly expand product variety (the Feenstra variety correction does not respond to demand-side immigration exposure), and barcode-level product characteristics are fixed. Online Appendix D.4 provides theoretical foundations showing that variety is stable when markup compression and scale effects generate offsetting incentives for product variety.

Demand. Consumer h's demand for product u takes the standard CES form:

$$c_{uh} = \phi_u \left(\frac{p_u}{P_h}\right)^{-\sigma_h} \frac{E_h}{P_h}, \tag{6.1}$$

where E_h is expenditure on differentiated products and P_h is the CES price index. The expenditure share on product u is

$$s_{uh} = \frac{\phi_u p_u^{1-\sigma_h}}{\sum_k \phi_k p_k^{1-\sigma_h}}. \tag{6.2}$$

Equilibrium. Firms set prices to maximize profits. As shown in Online Appendix D, the optimal price for product u is a markup over marginal cost:

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

[18] Immigrants may also be more price-sensitive due to less familiarity with local markets, though we do not attempt to distinguish this channel from income effects.

where the effective elasticity is the quantity-weighted average across consumer types:

$$\sigma_{u,\mathit{eff}} \equiv \theta_{un}\sigma_n + \theta_{ui}\sigma_i, \qquad \theta_{uh} \equiv \frac{L_h c_{uh}}{q_u}. \tag{6.4}$$

The effective elasticity captures the price sensitivity of the firm's actual customer base. Immigration increases the immigrant share θ_{ui}, raising $\sigma_{u,\mathit{eff}}$ and lowering prices.

6.2 Price Decomposition

We measure firm-level prices using the Sato-Vartia index, which aggregates product-level price changes holding expenditure shares fixed. Taking logs of the pricing equation and differentiating with respect to immigrant population yields:

Proposition 1 (Price Decomposition)**.** *The effect of immigration on the Sato-Vartia price index decomposes as:*

$$\frac{d\ln P_f^{SV}}{dL_i} = \sum_{u\in\Omega_f} \bar{s}_u \left[\underbrace{-\frac{1}{\sigma_{u,\mathit{eff}}(\sigma_{u,\mathit{eff}}-1)} \cdot \frac{\partial \sigma_{u,\mathit{eff}}}{\partial L_i}}_{\text{M1: Markup Effect } (<0)} + \underbrace{\delta \cdot \frac{c_{ui}}{q_u}}_{\text{M2: Scale Effect } (>0)} \right] \tag{6.5}$$

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

$$\frac{\partial \sigma_{u,\mathit{eff}}}{\partial L_i} = \frac{L_n c_{un} c_{ui}}{q_u^2}(\sigma_i - \sigma_n) > 0. \tag{6.6}$$

The decomposition reveals two mechanisms. *M1 (Markup Effect)*: Immigration raises the effective elasticity, compressing markups and reducing prices. The magnitude depends on the elasticity gap $(\sigma_i - \sigma_n)$. *M2 (Scale Effect)*: Immigration increases quantity demanded, raising marginal cost when $\delta > 0$. Our finding that Φ^{SV} declines in response to demand-side exposure indicates M1 dominates M2.

6.3 Quality Composition

The Sato-Vartia index measures within-product price changes, holding the product basket fixed. It does not capture shifts in expenditure across products—precisely the margin through which immigration affects firm-level quality. We now characterize this composition effect.

Define firm-level quality as the expenditure-weighted average across products:

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

where s_u is the aggregate expenditure share on product u:

$$s_u = \frac{L_n E_n s_{un} + L_i E_i s_{ui}}{L_n E_n + L_i E_i}. \tag{6.8}$$

Proposition 2 (Quality Composition Effect)**.** *Immigration reduces firm-level quality:*

$$\frac{d\bar{\phi}_f}{dL_i} = \sum_{u \in \Omega_f} \frac{\partial s_u}{\partial L_i} \phi_u < 0. \tag{6.9}$$

The effect operates through expenditure reallocation: immigrants' higher price sensitivity leads them to spend more on lower-quality (cheaper) products, shifting aggregate expenditure shares toward low-ϕ goods.

The proof follows from CES demand (Online Appendix D.2 and D.3). With higher demand elasticity ($\sigma_i > \sigma_n$), immigrants' expenditure shares tilt toward lower-priced products relative to natives. Higher-quality products command higher prices in equilibrium. As L_i increases, aggregate shares therefore shift toward low-price, low-quality products, reducing the expenditure-weighted average $\bar{\phi}_f$.

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: firms reduce prices when exposed to more immigrant consumers, who engage in intensive price search and exhibit greater price sensitivity. Supply-side exposure—immigration in production locations—has small positive effects on prices, inconsistent with simple models where immigrant labor reduces costs.

Our structural decomposition reveals that the demand-side effect operates primarily through markup compression: firms lower markups in response to more price-elastic consumers, while the variety effect of immigration is largely muted on the demand side. On the supply side, immigration facilitates product variety expansion, 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. An important avenue for future research is developing a general equilibrium framework that aggregates these firm-level pricing responses to economy-wide inflation dynamics, and examining how the demand composition channel interacts with monetary policy and business cycles.

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.

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.

ao, Rodrigo Ad Michal Kolesár, and Eduardo Morales, "Shift-Share Designs: Theory and Inference," *Quarterly Journal of Economics*, 2019, *134* (4), 1949–2010.

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

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.

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.

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.

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.

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, 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 Kostøl, Joan Monras, and Jonathan Vogel, "The Labor Supply Curve is Upward Sloping: The Effects of Immigrant-Induced Demand Shocks," Working Paper 33930, National Bureau of Economic Research 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.

Hong, Gihoon and John McLaren, "Are Immigrants a Shot in the Arm for the Local Economy?," Working Paper 21123, National Bureau of Economic Research 2015. Accessed: October 17, 2023.

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.

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.

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

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.

Tumen, Semih, "The economic impact of Syrian refugees on host countries: Quasi-experimental evidence from Turkey," *American Economic Review*, 2016, *106* (5), 456–460.

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`).

For our final dataset, we imposed the most conservative specification, retaining only observations that satisfied **Tier 4** (labeled as `match >= 4` in our replication code). This ensures that we match on name, legal entity structure, state, and specific street address simultaneously.

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)	$\sim$1.4M	$\sim$200k	$\sim$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

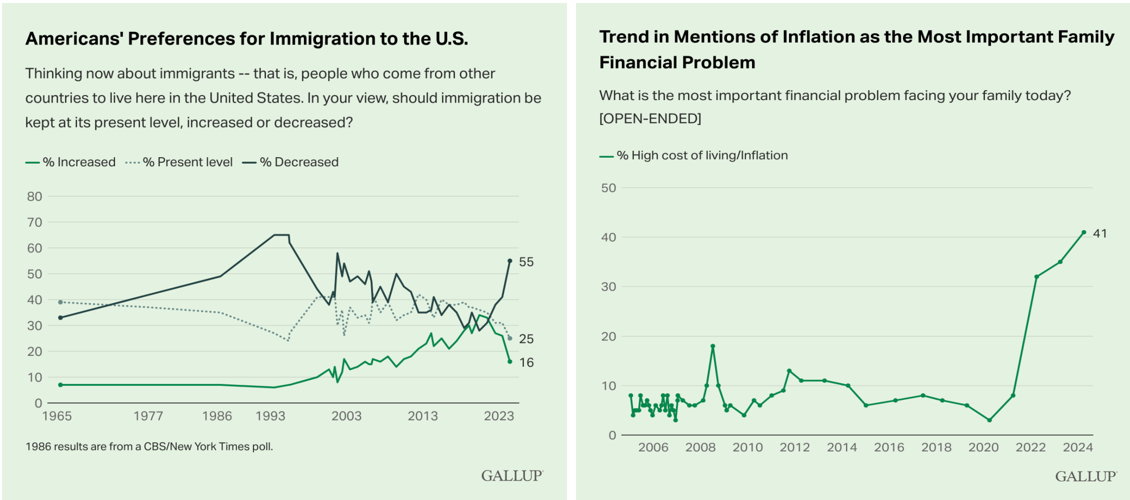

Notes: Figures are sourced from GALLUP polls. Figure on the left is sourced from an article on July 12, 2024, with the following link: Sharply More Americans Want to Curb Immigration to U.S.. Figure on the right is sourced from an article on May 2, 2024, with the following link: Americans Continue to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

Appendix B Additional Tables and Figures

B.1 Motivation and First Stage

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

	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)	
$\mathrm{I}^t_{o,-r(d),t}\frac{I_{Europe,r(d)\tau}}{I_{Europe,-r(d)\tau}}$					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 (3.9), 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.2 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)
$\text{Immigr}_{\text{d,t}}$	-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 (3.11) 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)
$\text{Immigr}_{\text{d,t}}$	-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: Price Regressions: Larger Sample of Products (6 years in sample)

	Δ P$^{\text{SV,Stacked}}$ (2006-2018)					
	(1)	(2)	(3)	(4)	(5)	(6)
$\text{Immigr}_{\text{d,t}}$	-0.034***	-0.028***	-0.023***	-0.062***	-0.064***	-0.056***
	(0.008)	(0.004)	(0.005)	(0.011)	(0.013)	(0.014)
Obs.	6,339	6,339	6,249	6,339	6,339	6,249
First-Stage F-stat				267.7	95.6	86.3
Fixed Effects	Time	State, Time	State, Time	Time	State, Time	State, Time
Method	OLS	OLS	OLS	IV	IV	IV
Controls			Yes			Yes

Notes: This table checks the effect of immigration shocks on changes in county-level average prices with the baseline IV specification. The price indices are constructed with products that are available for at least 6 years, instead of 14 years as in the baseline sample. Immigration shocks and endogenous immigration are winsorized by top and bottom 1%. We show specifications with stacked 4-year-differences (for 2006-2018), as well as just one 12-year long-difference (2006-2018). All regressions include state and time-period fixed effects, except columns (2) and (5), which replace state with county fixed effects. In column (4), we include controls that include retail wages, median rents, plus the interaction of GDP and unemployment rate at the county-level with year indicators. In columns (5) and (7) we continue to control for retail wages and median rents. Standard errors clustered are by state. ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

Table B.7: County Price Regressions: Robustness

	Δ $P^{\text{Sato-Vartia,Stacked}}$ (2006-2018)					
	(1)	(2)	(3)	(4)	(5)	(6)
$\text{Immigr}_{d,t}$ / $\text{Pop'l}_{d,t}$	-1.120***					
	(0.252)					
$\text{Immigr}_{d,t}$		-0.058***	-0.058***	-0.038***	-0.005***	-0.026***
		(0.013)	(0.012)	(0.009)	(0.001)	(0.008)
Mexico Border		-0.341***				
		(0.115)				
total ICE Arrests			0.047***			
			(0.012)			
Obs.	6,339	6,249	6,249	6,246	6,249	6,249
First-Stage F-stat	9.3	86.4	101.5	77.7	860.8	2,217.5
Fixed Effects	State, Time	State, Time	State, Time	State-Time	State, Time	County, Time
Baseline +	Imm / Popl	Border	ICE Arrests	State-Year FE	Non-Winsorized	County FE

Notes: Both immigration shock and immigration are winsorized by top and bottom 1%. All regressions include state and time-period fixed effects. Standard errors clustered are by state. ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

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 coefficient of -1.106 implies that a 1 percentage point increase in the immigrant share lowers CPG price growth by approximately 1.1 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.

Column (2) controls for proximity to the Mexican border, and Column (3) controls for ICE arrest activity; the immigration coefficient remains stable across both. Column (4) replaces state and time fixed effects with state-by-year fixed effects, yielding a somewhat smaller but still significant coefficient. Column (5) reports results without winsorization, and Column (6) includes county fixed effects. The negative effect of immigration on prices is robust across all specifications.

Table B.8: Alternative Rent Prices from ACS

	Δ $P^{SV,Stacked}$ (2006-2018)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\text{Immigr}_{d,t}$	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: 2006 county GDP per capita and urban indicator, each interacted with the year indicators. All regressions include state and time-period fixed effects. $***p < 0.01$, $**p < 0.05$, $*p < 0.1$.

Table B.9: 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.10: Immigration Effects on Wages: Robustness

Panel A: 2006-2018	Level Wage		Log Wage		Retail Wage	
	(1)	(2)	(3)	(4)	(5)	(6)
$\text{Immigr}_{\text{d,t}}$	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)						
$\text{Immigr}_{\text{d,t}}$	-1.838*	-0.415***	-0.310*	-0.044***	-0.693	-0.485
	(0.959)	(0.050)	(0.183)	(0.006)	(0.979)	(0.393)
$\text{Immigr}_{\text{d,t}} \times \text{EducYears}$	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						
$\text{Immigr}_{\text{d,t}}$	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 (3.11) 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 the immigration shock with average years of education among immigrants following Terry et al. (2024); immigrants are restricted to those aged 25 and older. 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.10 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.11: Immigration and Consumer Search Behavior by Demographics

	Hispanic			White		
	(1) (# days)	(2) (# stores)	(3) (coupon share)	(4) (# days)	(5) (# of stores)	(6) (coupon share)
$\text{Immigr}_{\text{d,t}}$	0.105	0.179**	0.020*	0.003	-0.047	0.000
	(0.090)	(0.079)	(0.010)	(0.023)	(0.035)	(0.004)
Obs.	1,343	1,343	1,343	1,343	1,343	1,343
First-Stage F-stat	55.1	55.1	55.1	55.1	55.1	55.1
Fixed Effects	State, Time	State, Time	State, Time	State, Time	State, Time	State, Time

Notes: Table examines effects of immigration on shopping behavior using NielsenIQ Consumer Panel data aggregated to county level. Analysis is conducted separately for Hispanic households (columns 1-3) and white households (columns 4-6). Outcomes are: number of shopping days per month, number of distinct stores visited, and share of transactions using coupons. All specifications use stacked four-year differences with state and period fixed effects. Standard errors clustered by state. ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

Appendix Table B.11 examines changes in consumer shopping behavior using NielsenIQ Consumer Panel data aggregated to the county level. Since Hispanic individuals constitute a disproportionate share of recent immigrants, we separately analyze Hispanic households (columns 1-3) as a proxy for immigrant consumers and white households (columns 4-6) as a proxy for natives. Outcomes include the number of shopping days per month, the number of distinct stores visited, and the share of transactions using coupons—all measures of price search intensity.

For Hispanic households in high-immigration counties, we find significant increases in the number of stores visited and coupon usage, with marginally significant increases in shopping days. In contrast, white households show no significant changes across any outcome. This demographic contrast suggests that newly arriving immigrants drive the increase in search intensity rather than a general change in local shopping norms, consistent with the search-based mechanism we hypothesize.

B.3 Firm-Level Analysis

Table B.12: Firm-Level Prices: Demand and Supply Exposure, Long-Difference

	Δ lnP$^{\text{SV,Long-Difference}}$ (2006-2018)			
	(1)	(2)	(3)	(4)
Immigration$_{f,d,t}$ (D)	-0.052***	-0.052***		-0.043**
	(0.018)	(0.019)		(0.018)
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. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table B.13: 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: 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. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

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, 2010–2014, and 2014–2018
- $t \in \tau$: year within period τ
- f: firm
- u: product (UPC)

In the main text (Section 5.4), we use ℓ to index estimation cells; throughout this appendix, we make this explicit by writing $\ell \equiv (g, s, \tau)$, i.e., the product group–state–period triplet.

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. This pricing structure is consistent with a broad class of competition models, including monopolistic competition, Bertrand competition with differentiated products, and Cournot competition. The key feature we exploit is that the markup is common across products within a firm, so it cancels in the double-differencing procedure below.

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 relationship $\ln q_{ugct} = \ln S^U_{ugct} + \ln E_{fgct}$ where E_{fgct} is firm expenditure (which cancels in double-differencing):

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

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 estimation cell (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) cell. From equations (C.7) and (C.10), the residuals are:

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

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

The moment condition is:

$$E[\omega_{ugct} \cdot \kappa_{ugct} \mid g, s, \tau] = 0 \tag{C.14}$$

We estimate $(\sigma^U_{gs\tau}, \delta_{gs\tau})$ by minimizing the sample analog of this moment condition:

$$(\hat{\sigma}^U_{gs\tau}, \hat{\delta}_{gs\tau}) = \arg\min_{(\sigma,\delta)} \left[\frac{1}{N_{gs\tau}} \sum_{c\in s} \sum_{t\in\tau} \sum_{u\in\Omega_{gs\tau}} \hat{\omega}_{ugct}(\sigma) \cdot \hat{\kappa}_{ugct}(\delta) \right]^2 \tag{C.15}$$

where $N_{gs\tau}$ is the number of observations in cell (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) cell 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:

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

This can be decomposed as:

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

where the first term is a weighted average of UPC prices (with $\tilde{w}_{ugct}$ being Sato-Vartia weights) and the second term is a Theil-like index of within-firm share dispersion, with $\tilde{S}^U_{fgct}$ being the geometric mean of within-firm shares.

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. The validity of this instrument follows directly from the CES structure:

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: Firm-level appeal ϕ^F_{fgct} is defined as the residual component of firm market share unexplained by firm prices. By construction, within-firm share dispersion—which reflects the distribution of sales across products within the firm—is orthogonal to this firm-level 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) cell, pooling across all counties $c \in s$ and years $t \in \tau$.

C.1.4 Implementation Details

Estimation Cells. We estimate both elasticities separately for each state s $\times$ product group g $\times$ period τ cell, 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
- $\tau = 2$: 2010–2014
- $\tau = 3$: 2014–2018

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 necessary within-cell variation 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) cell:

- 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 quality constant for continuing varieties. Separately, recovering the *level* of appeal ϕ^F_{fgct} (as in Section C.4) would require an additional normalization—for instance, Hottman et al. (2016) normalize the geometric mean of firm appeal to unity within each market. However, our regression analysis uses only appeal *growth rates* $\Phi^{Appeal}_{fgct} \equiv \phi^F_{fgct}/\phi^F_{fgc,t-1}$, which are invariant to level normalizations since both periods belong to the same estimation cell. Similarly, our elasticity estimation (Section C.1) relies on double-differencing, which eliminates appeal levels entirely. Thus, no additional normalization beyond equation (C.25) is required for our empirical analysis.

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 Model

D.1 Environment

D.1.1 Consumers

The economy has two consumer types $h \in \{n, i\}$—natives and immigrants—with populations L_n and L_i. Each consumer of type h has income y_h, with

$$y_i < y_n. \tag{D.1}$$

Consumers have two-tier preferences. The upper tier allocates expenditure between an outside good and differentiated products via Cobb-Douglas:

$$U_h = C_{0h}^{1-\alpha} C_{Dh}^{\alpha}, \tag{D.2}$$

implying each consumer spends $E_h = \alpha y_h$ on differentiated products.

The lower tier aggregates differentiated products via CES:

$$C_{Dh} = \left(\sum_f \sum_{u \in \Omega_f} \phi_u^{\frac{1}{\sigma_h}} c_{uh}^{\frac{\sigma_h - 1}{\sigma_h}} \right)^{\frac{\sigma_h}{\sigma_h - 1}}, \tag{D.3}$$

where $\phi_u > 0$ is product quality, c_{uh} is consumption, and $\sigma_h > 1$ is the elasticity of substitution. Recall from the main text: Immigrants have higher demand elasticity than natives ($\sigma_i > \sigma_n > 1$).

For expositional simplicity, the theoretical framework uses a flat CES structure with a single elasticity σ_h per consumer type. This corresponds to the special case $\sigma^F_h = \sigma^U_h = \sigma_h$ in the nested CES demand system used in our structural estimation (Section 5.4), where $\sigma^U_{gs\tau}$ governs within-firm substitution and $\sigma^F_{gs\tau}$ governs across-firm substitution. The qualitative predictions are unchanged under the more general nested structure; Section D.5 provides derivations for that case.

D.1.2 Firms

Each firm f produces a set of products Ω_f with heterogeneous qualities $\{\phi_u\}_{u \in \Omega_f}$. Production exhibits potentially decreasing returns to scale at the product level. The total cost of producing quantity q_u of product u is

$$C_u(q_u) = F_u + \frac{a_u}{1+\delta} q_u^{1+\delta}, \tag{D.4}$$

where $F_u \geq 0$ is a fixed cost, $a_u > 0$ is a product-specific cost shifter, and $\delta \geq 0$ governs the degree of decreasing returns. The marginal cost is

$$MC_u(q_u) = \frac{\partial C_u}{\partial q_u} = a_u \cdot q_u^{\delta}. \tag{D.5}$$

When $\delta = 0$, marginal cost is constant; when $\delta > 0$, marginal cost increases with quantity. Higher-quality products have higher costs: $\partial a_u / \partial \phi_u > 0$.

The firm's total profit from its product portfolio is

$$\pi_f = \sum_{u \in \Omega_f} \left[p_u q_u - C_u(q_u) \right]. \tag{D.6}$$

We assume that the set of products Ω_f and product-level qualities $\{\phi_u\}_{u \in \Omega_f}$ are fixed. This assumption reflects two empirical observations: Feenstra variety corrections are small (stable product sets), and barcode-level characteristics do not change (stable product qualities). Section D.4 provides theoretical foundations.

D.1.3 Market Structure

As stated in the main text, we consider monopolistic competition: each firm chooses prices $\{p_u\}_{u \in \Omega_f}$ for its products but takes aggregate price indices P_n, P_i and total expenditures $E_n L_n$, $E_i L_i$ as given. Since each product's weight in the aggregate price index is negligible, the firm's pricing decision for product u does not affect demand for its other products $v \neq u$. This yields a separability property: the firm's problem decomposes into independent single-product pricing decisions. Section D.5 relaxes this by considering Bertrand competition where firms internalize cannibalization across products.

D.2 Equilibrium Pricing

D.2.1 Firm's Problem

Taking aggregate variables $(P_n, P_i, E_n, E_i, L_n, L_i)$ as given, firm f chooses prices to maximize total profit:

$$\max_{\{p_u\}_{u \in \Omega_f}} \quad \pi_f = \sum_{u \in \Omega_f} \left[p_u q_u(p_u) - C_u(q_u(p_u)) \right], \tag{D.7}$$

where total demand for product u is

$$q_u(p_u) = L_n c_{un}(p_u) + L_i c_{ui}(p_u) = \sum_{h \in \{n,i\}} L_h \phi_u \left(\frac{p_u}{P_h} \right)^{-\sigma_h} \frac{E_h}{P_h}. \tag{D.8}$$

Under monopolistic competition, the firm treats P_h as fixed. Since q_v does not depend on p_u

for $v \neq u$, the first-order condition for each product is independent:

$$\frac{\partial \pi_f}{\partial p_u} = \frac{\partial}{\partial p_u}\left[p_u q_u - C_u(q_u)\right] = 0 \quad \text{for each } u \in \Omega_f. \tag{D.9}$$

D.2.2 First-Order Conditions

Differentiating profit with respect to p_u:

$$\frac{\partial \pi_u}{\partial p_u} = q_u + p_u \frac{\partial q_u}{\partial p_u} - MC_u(q_u)\frac{\partial q_u}{\partial p_u} = 0. \tag{D.10}$$

To evaluate $\partial q_u / \partial p_u$, we differentiate total demand (D.8):

$$\frac{\partial q_u}{\partial p_u} = \sum_h L_h \frac{\partial c_{uh}}{\partial p_u} = -\sum_h L_h \sigma_h \frac{c_{uh}}{p_u} = -\frac{q_u}{p_u} \cdot \sigma_{u,\mathit{eff}}, \tag{D.11}$$

where we define the effective demand elasticity:

Definition 1 (Effective Elasticity)**.** *The effective demand elasticity for product u is the quantity-weighted average of type-specific elasticities:*

$$\sigma_{u,\mathit{eff}} \equiv \theta_{un}\sigma_n + \theta_{ui}\sigma_i, \qquad \theta_{uh} \equiv \frac{L_h c_{uh}}{q_u}. \tag{D.12}$$

Substituting (D.11) into (D.10) and rearranging:

$$q_u - (p_u - MC_u)\frac{q_u}{p_u}\sigma_{u,\mathit{eff}} = 0 \quad \Longrightarrow \quad 1 = \frac{p_u - MC_u}{p_u} \cdot \sigma_{u,\mathit{eff}}. \tag{D.13}$$

Solving for the optimal price:

Proposition 3 (Markup Formula)**.** *The profit-maximizing price is*

$$p_u = \mu_u \cdot MC_u, \qquad \mu_u \equiv \frac{\sigma_{u,\mathit{eff}}}{\sigma_{u,\mathit{eff}} - 1}, \tag{D.14}$$

where $\mu_u > 1$ is the markup. The markup is decreasing in the effective elasticity:

$$\frac{\partial \mu_u}{\partial \sigma_{u,\mathit{eff}}} = -\frac{1}{(\sigma_{u,\mathit{eff}} - 1)^2} < 0. \tag{D.15}$$

The effective elasticity captures the price sensitivity of the firm's customer base. When immigrants—who have higher elasticity $\sigma_i > \sigma_n$—constitute a larger share of purchasers, $\sigma_{u,\mathit{eff}}$ rises and the markup μ_u falls.

D.2.3 Second-Order Condition

We establish that the profit-maximizing price is well-defined by verifying the second-order condition. The second derivative of profit with respect to price is:

$$\frac{\partial^2 \pi_u}{\partial p_u^2} = 2\frac{\partial q_u}{\partial p_u} + (p_u - MC_u)\frac{\partial^2 q_u}{\partial p_u^2} - \frac{\partial MC_u}{\partial q_u}\left(\frac{\partial q_u}{\partial p_u}\right)^2. \tag{D.16}$$

To evaluate this expression, we compute the second derivative of demand. From (D.8):

$$\frac{\partial^2 q_u}{\partial p_u^2} = \sum_h L_h \sigma_h(\sigma_h + 1)\phi_u \left(\frac{p_u}{P_h}\right)^{-\sigma_h} \frac{E_h}{P_h} \cdot \frac{1}{p_u^2} = \frac{q_u}{p_u^2}\tilde{\sigma}, \tag{D.17}$$

where we define:

$$\tilde{\sigma} \equiv \theta_{un}\sigma_n(\sigma_n + 1) + \theta_{ui}\sigma_i(\sigma_i + 1). \tag{D.18}$$

At the first-order condition, $p_u - MC_u = \frac{p_u}{\sigma_{u,\mathit{eff}}}$. Substituting into (D.16) along with (D.11):

$$\frac{\partial^2 \pi_u}{\partial p_u^2} = \frac{q_u}{p_u}\left[-2\sigma_{u,\mathit{eff}} + \frac{\tilde{\sigma}}{\sigma_{u,\mathit{eff}}} - \delta(\sigma_{u,\mathit{eff}} - 1)\sigma_{u,\mathit{eff}}\right]. \tag{D.19}$$

The second-order condition $\frac{\partial^2 \pi_u}{\partial p_u^2} < 0$ requires:

$$\tilde{\sigma} < \sigma_{u,\mathit{eff}}^2 \left(2 + \delta(\sigma_{u,\mathit{eff}} - 1)\right). \tag{D.20}$$

Verifying the Second-Order Condition. We show that condition (D.20) holds for all $\sigma_n, \sigma_i > 1$ and $\delta \geq 0$. Note that $\tilde{\sigma}$ and $\sigma_{u,\mathit{eff}}$ can be written as moments of the distribution of σ_h with weights θ_{uh}:

$$\sigma_{u,\mathit{eff}} = \mathbb{E}_\theta[\sigma_h], \tag{D.21}$$

$$\tilde{\sigma} = \mathbb{E}_\theta[\sigma_h(\sigma_h + 1)] = \mathbb{E}_\theta[\sigma_h^2] + \mathbb{E}_\theta[\sigma_h]. \tag{D.22}$$

Using the variance decomposition $\mathbb{E}_\theta[\sigma_h^2] = \mathrm{Var}_\theta(\sigma_h) + (\mathbb{E}_\theta[\sigma_h])^2$:

$$\tilde{\sigma} = \mathrm{Var}_\theta(\sigma_h) + \sigma_{u,\mathit{eff}}^2 + \sigma_{u,\mathit{eff}}. \tag{D.23}$$

Substituting into (D.20), the second-order condition becomes:

$$\mathrm{Var}_\theta(\sigma_h) < \sigma_{u,\mathit{eff}}(\sigma_{u,\mathit{eff}} - 1)(1 + \delta\sigma_{u,\mathit{eff}}). \tag{D.24}$$

In the two-type case, $\mathrm{Var}_\theta(\sigma_h) = \theta(1-\theta)(\sigma_i - \sigma_n)^2$, which is maximized at $\theta = \frac{1}{2}$. Consider the worst case: $\theta = \frac{1}{2}$ and $\delta = 0$. The condition (D.24) becomes:

$$\frac{(\sigma_i - \sigma_n)^2}{4} < \bar{\sigma}(\bar{\sigma} - 1), \tag{D.25}$$

where $\bar{\sigma} \equiv \frac{\sigma_n + \sigma_i}{2}$. Reparametrizing with $d \equiv \frac{\sigma_i - \sigma_n}{2}$, so that $\sigma_n = \bar{\sigma} - d$ and $\sigma_i = \bar{\sigma} + d$, the condition becomes $d^2 < \bar{\sigma}(\bar{\sigma} - 1)$. Since $\sigma_n > 1$ implies $d < \bar{\sigma} - 1$, we have:

$$d^2 < (\bar{\sigma} - 1)^2 < \bar{\sigma}(\bar{\sigma} - 1), \tag{D.26}$$

where the second inequality follows from $\bar{\sigma} - 1 < \bar{\sigma}$. Thus, the second-order condition holds for all $\sigma_n, \sigma_i > 1$ and $\delta \geq 0$.

D.2.4 Quality-Price Relationship

We now show that higher-quality products command higher prices in equilibrium. This result does not require assuming that production costs are increasing in quality; it emerges endogenously from the interaction of demand and decreasing returns to scale.

Setup. Define aggregate demand shifters for each consumer type:

$$D_h \equiv L_h \frac{E_h}{P_h^{1-\sigma_h}}. \tag{D.27}$$

Total demand for product u can then be written as:

$$q_u = \phi_u \left[D_n p_u^{-\sigma_n} + D_i p_u^{-\sigma_i} \right] \equiv \phi_u G(p_u), \tag{D.28}$$

where $G(p_u) \equiv D_n p_u^{-\sigma_n} + D_i p_u^{-\sigma_i}$.

Quantity Shares. The quantity share of consumer type h is:

$$\theta_{uh} = \frac{L_h c_{uh}}{q_u} = \frac{D_h p_u^{-\sigma_h}}{D_n p_u^{-\sigma_n} + D_i p_u^{-\sigma_i}}. \tag{D.29}$$

Note that θ_{uh} depends on p_u but not on ϕ_u (quality cancels in the ratio).

Equilibrium Condition. Substituting (D.28) into the pricing equation $p_u = \mu(p_u) \cdot a_u q_u^{\delta}$:

$$p_u = \mu(p_u) \cdot a_u \cdot \phi_u^{\delta} \cdot G(p_u)^{\delta}, \tag{D.30}$$

where $\mu(p_u) = \frac{\sigma_{u,\mathit{eff}}(p_u)}{\sigma_{u,\mathit{eff}}(p_u) - 1}$ and $\sigma_{u,\mathit{eff}}(p_u) = \theta_{un}(p_u)\sigma_n + \theta_{ui}(p_u)\sigma_i$.

Comparative Static. Having established that the equilibrium is well-defined, we apply the implicit function theorem to sign $\frac{\partial p_u}{\partial \phi_u}$. Write the equilibrium condition as:

$$F(p_u, \phi_u) \equiv p_u - \mu(p_u) \cdot a_u \cdot \phi_u^\delta \cdot G(p_u)^\delta = 0. \tag{D.31}$$

Derivative with respect to ϕ_u:

$$\frac{\partial F}{\partial \phi_u} = -\mu(p_u) \cdot a_u \cdot \delta \phi_u^{\delta-1} \cdot G(p_u)^\delta < 0. \tag{D.32}$$

Derivative with respect to p_u: To sign $\frac{\partial F}{\partial p_u}$, we establish its relationship to the second-order condition. The first-order condition can be written as $H(p_u) \equiv \frac{\partial \pi_u}{\partial p_u} = 0$. Direct calculation shows:

$$H = \frac{\partial q_u}{\partial p_u} \cdot \frac{\sigma_{u,\mathit{eff}} - 1}{\sigma_{u,\mathit{eff}}} \cdot F. \tag{D.33}$$

Differentiating with respect to p_u and evaluating at equilibrium where $F = 0$:

$$\frac{\partial^2 \pi_u}{\partial p_u^2} = \frac{\partial q_u}{\partial p_u} \cdot \frac{\sigma_{u,\mathit{eff}} - 1}{\sigma_{u,\mathit{eff}}} \cdot \frac{\partial F}{\partial p_u}. \tag{D.34}$$

Since $\frac{\partial q_u}{\partial p_u} < 0$ and $\sigma_{u,\mathit{eff}} > 1$, the second-order condition ($\frac{\partial^2 \pi_u}{\partial p_u^2} < 0$) holds if and only if $\frac{\partial F}{\partial p_u} > 0$.

By the implicit function theorem:

$$\frac{\partial p_u}{\partial \phi_u} = -\frac{\partial F / \partial \phi_u}{\partial F / \partial p_u}. \tag{D.35}$$

Since $\frac{\partial F}{\partial \phi_u} < 0$ and $\frac{\partial F}{\partial p_u} > 0$, we have $\frac{\partial p_u}{\partial \phi_u} > 0$.

Proposition 4 (Quality-Price Relationship)**.** *In equilibrium with decreasing returns to scale ($\delta > 0$) and $\sigma_n, \sigma_i > 1$, higher-quality products command higher prices:*

$$\frac{\partial p_u}{\partial \phi_u} > 0. \tag{D.36}$$

This result holds for arbitrary heterogeneity in demand elasticities ($\sigma_i \neq \sigma_n$).

The intuition is straightforward: higher quality ϕ_u shifts demand outward, increasing quantity q_u. With decreasing returns ($\delta > 0$), a higher quantity raises the marginal cost $MC_u = a_u q_u^\delta$, which in turn raises the equilibrium price.

D.2.5 Price Decomposition

Lemma 1 (Immigration Raises Effective Elasticity)**.**

$$\frac{\partial \sigma_{u,\mathit{eff}}}{\partial L_i} = \frac{L_n c_{un} c_{ui}}{q_u^2}(\sigma_i - \sigma_n) > 0. \tag{D.37}$$

Proof. Differentiating the immigrant share:

$$\frac{\partial \theta_{ui}}{\partial L_i} = \frac{c_{ui} \cdot L_n c_{un}}{q_u^2} > 0. \tag{D.38}$$

Then:

$$\frac{\partial \sigma_{u,\mathit{eff}}}{\partial L_i} = \frac{\partial \theta_{ui}}{\partial L_i}(\sigma_i - \sigma_n) = \frac{L_n c_{un} c_{ui}}{q_u^2}(\sigma_i - \sigma_n) > 0. \tag{D.39}$$

□

We now derive Proposition 1 from the main text.

Proof of Proposition 1. From $p_u = \mu_u \cdot MC_u$ with $MC_u = a_u q_u^\delta$:

$$\ln p_u = \ln \mu_u + \ln a_u + \delta \ln q_u. \tag{D.40}$$

Differentiating:

$$\frac{d \ln p_u}{dL_i} = \frac{d \ln \mu_u}{d\sigma_{u,\mathit{eff}}} \cdot \frac{\partial \sigma_{u,\mathit{eff}}}{\partial L_i} + \delta \cdot \frac{1}{q_u}\frac{\partial q_u}{\partial L_i}. \tag{D.41}$$

The markup elasticity is:

$$\frac{d \ln \mu_u}{d\sigma_{u,\mathit{eff}}} = \frac{1}{\sigma_{u,\mathit{eff}}} - \frac{1}{\sigma_{u,\mathit{eff}} - 1} = -\frac{1}{\sigma_{u,\mathit{eff}}(\sigma_{u,\mathit{eff}} - 1)} < 0. \tag{D.42}$$

The quantity effect is $\partial q_u / \partial L_i = c_{ui}$. □

Corollary 1 (Sato-Vartia Index)**.** *The firm-level Sato-Vartia index aggregates product-level effects:*

$$\frac{d \ln P_f^{SV}}{dL_i} = \sum_{u \in \Omega_f} \bar{s}_u \frac{d \ln p_u}{dL_i}, \tag{D.43}$$

where $\bar{s}_u$ are Sato-Vartia weights (log-mean shares). Since expenditure weights are held fixed, this index captures M1 and M2 but not composition effects.

D.3 Quality Composition

This section derives how immigration affects firm-level quality through expenditure reallocation across products.

D.3.1 Firm-Level Quality

Definition 2 (Firm-Level Quality). *Firm f's quality index is the expenditure-weighted average:*

$$\bar{\phi}_f \equiv \sum_{u \in \Omega_f} s_u \phi_u, \tag{D.44}$$

where the aggregate expenditure share is

$$s_u = \frac{L_n E_n s_{un} + L_i E_i s_{ui}}{L_n E_n + L_i E_i}. \tag{D.45}$$

D.3.2 Composition Effect

The expenditure share of product u for consumer type h is:

$$s_{uh} = \frac{\phi_u p_u^{1-\sigma_h}}{\sum_v \phi_v p_v^{1-\sigma_h}}. \tag{D.46}$$

Lemma 2 (Immigrants Tilt Toward Cheaper Products). *For any two products u and v with $p_u > p_v$:*

$$\frac{s_{ui}/s_{vi}}{s_{un}/s_{vn}} = \left(\frac{p_u}{p_v}\right)^{\sigma_n - \sigma_i} < 1. \tag{D.47}$$

Proof. From (D.46), the relative share is $\frac{s_{uh}}{s_{vh}} = \frac{\phi_u}{\phi_v}\left(\frac{p_u}{p_v}\right)^{1-\sigma_h}$. Taking the ratio across consumer types:

$$\frac{s_{ui}/s_{vi}}{s_{un}/s_{vn}} = \left(\frac{p_u}{p_v}\right)^{(1-\sigma_i)-(1-\sigma_n)} = \left(\frac{p_u}{p_v}\right)^{\sigma_n - \sigma_i} < 1, \tag{D.48}$$

since $p_u > p_v$ and $\sigma_i > \sigma_n$. □

We now prove Proposition 2 from the main text.

Proof of Proposition 2. Differentiating (D.44) with product-level qualities fixed:

$$\frac{d\bar{\phi}_f}{dL_i} = \sum_u \frac{\partial s_u}{\partial L_i} \phi_u. \tag{D.49}$$

Consider any two products u and v with $\phi_u > \phi_v$. By Proposition 4, $p_u > p_v$. By Lemma 2, immigrants tilt toward the cheaper product, so $\partial s_u / \partial L_i < \partial s_v / \partial L_i$.

Since shares sum to one, $\sum_u \frac{\partial s_u}{\partial L_i} = 0$. Order products by quality: $\phi_1 > \phi_2 > \cdots > \phi_N$, where $N = |\Omega_f|$. Then $\frac{\partial s_1}{\partial L_i} < \frac{\partial s_2}{\partial L_i} < \cdots < \frac{\partial s_N}{\partial L_i}$. Let $\phi_{\min} \equiv \phi_N$ denote the lowest quality. Then:

$$\sum_{u=1}^{N} \frac{\partial s_u}{\partial L_i} \phi_u < \sum_{u=1}^{N} \frac{\partial s_u}{\partial L_i} \phi_{\min} = \phi_{\min} \sum_{u=1}^{N} \frac{\partial s_u}{\partial L_i} = 0. \tag{D.50}$$

□

Relation to Chain-Weighted Indices. The Sato-Vartia index holds expenditure weights fixed:

$$\ln P_f^{SV,t} - \ln P_f^{SV,t-1} = \sum_u \bar{s}_u (\ln p_u^t - \ln p_u^{t-1}). \tag{D.51}$$

This captures within-product price changes (M1 and M2) but not between-product expenditure shifts. When immigrants shift spending toward low-quality products, this affects $\bar{\phi}_f$ but not P_f^{SV}. This is why our empirical decomposition shows effects on both Φ^{SV} (prices fall via M1) and Φ^{Appeal} (quality falls via composition), representing distinct channels.

D.4 Variety Effects

This section extends the baseline model to allow for endogenous product entry and exit. We show that immigration generates offsetting effects on firm profitability, rationalizing why variety effects are muted in our data.

D.4.1 Setup

The baseline model takes the product set Ω_f as fixed. We now allow firms to choose which products to offer. Each potential product u requires a fixed cost $F_u > 0$ to introduce. A firm offers product u if and only if operating profit covers this fixed cost.

D.4.2 Entry Condition

A product u is viable if operating profit covers fixed cost F_u. From the first-order condition (D.13), $p_u - MC_u = p_u/\sigma_{u,\mathit{eff}}$, so operating profit is:

$$\pi_u^{op} = (p_u - MC_u) q_u = \frac{p_u q_u}{\sigma_{u,\mathit{eff}}} \geq F_u. \tag{D.52}$$

D.4.3 Offsetting Effects

Proposition 5 (Ambiguous Variety Effect)**.** *Immigration has ambiguous effects on product entry because M1 and M2 generate offsetting effects on profitability:*

$$\frac{d\pi_u^{op}}{dL_i} = \underbrace{-\frac{R_u}{\sigma_{u,\mathit{eff}}^2} \frac{\partial \sigma_{u,\mathit{eff}}}{\partial L_i}}_{<0 \text{ (M1: margin falls)}} + \underbrace{\frac{1}{\sigma_{u,\mathit{eff}}} \frac{\partial R_u}{\partial L_i}}_{>0 \text{ (M2: revenue rises)}} . \tag{D.53}$$

Proof. Operating profit is $\pi_u^{op} = R_u/\sigma_{u,\mathit{eff}}$ where $R_u = p_u q_u$. Differentiating:

$$\frac{d\pi_u^{op}}{dL_i} = \frac{1}{\sigma_{u,\mathit{eff}}}\frac{\partial R_u}{\partial L_i} - \frac{R_u}{\sigma_{u,\mathit{eff}}^2}\frac{\partial \sigma_{u,\mathit{eff}}}{\partial L_i}. \tag{D.54}$$

The first term is positive (immigration raises demand and revenue). The second term is negative (immigration raises elasticity, compressing margins). The net effect depends on magnitudes. □

Corollary 2 (Stable Variety)**.** *Entry and exit occur at the margin where* $\pi_u^{op} = F_u$*. When M1 and M2 roughly offset for marginal products,* $d\pi_u^{op}/dL_i \approx 0$*, so immigration does not push these products above or below the entry threshold. Variety remains stable, consistent with our finding that Feenstra variety corrections are small.*

Fixed Costs and Quality Adjustment. We also require that product-level qualities are fixed. This is rationalized by high adjustment costs: reformulating products, changing suppliers, or repositioning brands involves substantial fixed investments. In the short run, firms respond to immigration through pricing and sales composition rather than product characteristics.

D.5 Extension: Bertrand Competition

This section extends the model to Bertrand competition where firms internalize substitution across their own products.

D.5.1 Setup

Each firm f produces products Ω_f and chooses prices simultaneously. Consumers have nested CES preferences:

$$C_{Dh} = \left(\sum_f \varphi_f^{\frac{1}{\sigma_h^F}} C_{fh}^{\frac{\sigma_h^F - 1}{\sigma_h^F}}\right)^{\frac{\sigma_h^F}{\sigma_h^F - 1}}, \quad C_{fh} = \left(\sum_{u \in \Omega_f} \phi_u^{\frac{1}{\sigma_h^U}} c_{uh}^{\frac{\sigma_h^U - 1}{\sigma_h^U}}\right)^{\frac{\sigma_h^U}{\sigma_h^U - 1}}, \tag{D.55}$$

with $\sigma_h^U > \sigma_h^F > 1$ (within-firm substitution exceeds across-firm substitution).

D.5.2 Demand

Consumer h's demand for product u from firm f is:

$$c_{uh} = \phi_u \left(\frac{p_u}{P_{fh}}\right)^{-\sigma_h^U} \cdot \varphi_f \left(\frac{P_{fh}}{P_h}\right)^{-\sigma_h^F} \cdot \frac{E_h}{P_h}, \tag{D.56}$$

where $P_{fh} = \left(\sum_{u \in \Omega_f} \phi_u p_u^{1-\sigma_h^U}\right)^{\frac{1}{1-\sigma_h^U}}$ is the firm-level price index for consumer h.

Total demand for product u is $q_u = \sum_h L_h c_{uh}$.

D.5.3 Firm's Problem

Firm f chooses prices $\{p_u\}_{u\in\Omega_f}$ to maximize total profit:

$$\max_{\{p_u\}} \sum_{u\in\Omega_f} \left[p_u q_u - C_u(q_u)\right]. \tag{D.57}$$

Unlike monopolistic competition, raising p_u affects demand for other products $v \in \Omega_f$ through the firm-level price index P_{fh}. The firm internalizes this cannibalization.

D.5.4 First-Order Condition

The first-order condition for product u is:

$$\frac{\partial \pi_f}{\partial p_u} = q_u + \sum_{v\in\Omega_f} (p_v - MC_v)\frac{\partial q_v}{\partial p_u} = 0. \tag{D.58}$$

To evaluate the cross-derivatives, note from (D.56):

$$\frac{\partial q_v}{\partial p_u} = \sum_h L_h \frac{\partial c_{vh}}{\partial p_u} = \sum_h L_h c_{vh} \cdot \frac{\sigma_h^U - \sigma_h^F}{p_u} \cdot s_{uh|f} \quad \text{for } v \neq u, \tag{D.59}$$

where $s_{uh|f} \equiv \phi_u p_u^{1-\sigma_h^U} / \sum_{k\in\Omega_f} \phi_k p_k^{1-\sigma_h^U}$ is product u's expenditure share within firm f for consumer h.

The own-derivative is:

$$\frac{\partial q_u}{\partial p_u} = -\frac{q_u}{p_u}\left[\sigma_{\mathit{eff}}^U - s_{u|f}(\sigma_{\mathit{eff}}^U - \sigma_{\mathit{eff}}^F)\right], \tag{D.60}$$

where we define:

$$\sigma_{\mathit{eff}}^U \equiv \theta_{un}\sigma_n^U + \theta_{ui}\sigma_i^U, \tag{D.61}$$

$$\sigma_{\mathit{eff}}^F \equiv \theta_{un}\sigma_n^F + \theta_{ui}\sigma_i^F, \tag{D.62}$$

$$s_{u|f} \equiv \theta_{un}s_{un|f} + \theta_{ui}s_{ui|f}, \tag{D.63}$$

and $\theta_{uh} \equiv L_h c_{uh}/q_u$ is the quantity share of consumer type h.

D.5.5 Markup Formula

Solving the first-order condition yields:

Proposition 6 (Bertrand Markup)**.** *Under Bertrand competition with symmetric markups across* *products within the firm, the markup for product u is:*

$$\mu_u = \frac{\tilde{\sigma}_{u,\mathit{eff}}}{\tilde{\sigma}_{u,\mathit{eff}} - 1}, \tag{D.64}$$

where the effective elasticity with cannibalization is:

$$\tilde{\sigma}_{u,\mathit{eff}} \equiv \sigma^U_{\mathit{eff}} - s_{u|f}(\sigma^U_{\mathit{eff}} - \sigma^F_{\mathit{eff}}) < \sigma^U_{\mathit{eff}}. \tag{D.65}$$

Proof. Substituting (D.60) into the first-order condition and assuming symmetric markups across products within the firm (so that cross-product terms simplify), we obtain:

$$q_u - (p_u - MC_u)\frac{q_u}{p_u}\tilde{\sigma}_{u,\mathit{eff}} = 0. \tag{D.66}$$

Solving for the markup: $\mu_u = p_u/MC_u = \tilde{\sigma}_{u,\mathit{eff}}/(\tilde{\sigma}_{u,\mathit{eff}} - 1)$. □

Interpretation. Cannibalization reduces the effective elasticity from σ^U_{eff} to $\tilde{\sigma}_{u,\mathit{eff}}$. When the firm raises p_u, some consumers switch to other products within the firm rather than to competitors. The firm captures fraction $s_{u|f}$ of these diverted sales, so the relevant elasticity is reduced by $s_{u|f}(\sigma^U_{\mathit{eff}} - \sigma^F_{\mathit{eff}})$.

D.5.6 Effect of Immigration

The qualitative effects of immigration are unchanged under Bertrand competition. Immigration raises the immigrant share θ_{ui}, which increases both σ^U_{eff} and σ^F_{eff}:

$$\frac{\partial \sigma^U_{\mathit{eff}}}{\partial L_i} = \frac{L_n c_{un} c_{ui}}{q_u^2}(\sigma^U_i - \sigma^U_n) > 0, \tag{D.67}$$

$$\frac{\partial \sigma^F_{\mathit{eff}}}{\partial L_i} = \frac{L_n c_{un} c_{ui}}{q_u^2}(\sigma^F_i - \sigma^F_n) > 0. \tag{D.68}$$

Since $\tilde{\sigma}_{u,\mathit{eff}}$ is increasing in both σ^U_{eff} and σ^F_{eff}, immigration raises the effective elasticity, compressing markups and reducing prices.

D.5.7 Quality Composition under Bertrand

The quality composition effect (Proposition 2) is unchanged under Bertrand competition. The expenditure share (6.2) depends only on consumer preferences, not on firm pricing behavior. Since $\sigma^U_i > \sigma^U_n$, immigrants still allocate more expenditure to lower-priced (and hence lower-quality) products, reducing firm-level quality $\bar{\phi}_f$. This demand-side mechanism is robust to the competitive structure.