Online Appendix for "Prices and Immigration: Firm-Level Evidence"

Ryan Kim Justin H. Leung Ariel Weinberger

Appendix Contents

Appendices **A-1**

Appendix A Data Appendix **A-2**
A.1 NielsenIQ GS1 + NETS Matching . . . A-2
A.2 Identifying Immigrant Households in the NielsenIQ Panel . . . A-7

Appendix B Additional Tables and Figures **A-12**
B.1 Motivation, Data, and First Stage . . . A-12
B.2 Instrument Construction . . . A-16
B.3 County-Level Analysis . . . A-20
B.4 Firm-Level Analysis . . . A-33

Appendix C Structural Estimation: Details **A-42**
C.1 Elasticity Estimation . . . A-42
C.2 Sato-Vartia Price Index . . . A-48
C.3 Feenstra Variety Adjustment . . . A-51
C.4 Appeal Recovery . . . A-51

Appendix D Theoretical Framework **A-52**
D.1 Environment . . . A-52
D.2 Equilibrium Pricing . . . A-54
D.3 Price Decomposition . . . A-60
D.4 Appeal Composition . . . A-64
D.5 Variety Effects . . . A-66
D.6 How Appeal Enters Preferences . . . A-66

Appendix E Calibrating the Price Response **A-67**

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 records every U.S. establishment with its county and employment, and links each establishment to its headquarters). We match GS1 firms to NETS headquarters records, and then measure each firm's production locations from all establishments linked to the matched headquarters (Section A.1.8).

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 Validation of Firm Linkages.

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 retail scanner data.

A.1.2 Hierarchical Matching Algorithm

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

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

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

We classified the resulting matches into the following quality tiers:

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

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

A.1.3 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 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.1.4 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.1.5 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 Google Places 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.1.6 Firms Linked to More Than One Headquarters Record

Matching on name and state does not always link a GS1 firm to a single NETS headquarters record. NETS sometimes records one company under several headquarters records in the same state, and a GS1 prefix that matches such a company is linked to all of them. We keep these links and measure

the firm's production locations from the establishments of all its records combined (Section A.1.8). The reverse case also occurs: a company often registers a GS1 prefix for each brand, so several prefixes can match the same headquarters record. Each prefix remains a separate firm in our analysis, and such firms share the same production locations.

Prefixes linked to more than one record account for 37 percent of the observations in Table 5 but 73 percent of its sales weight, since they include the largest firms. To check that combining records does not drive our results, we instead link each such prefix only to its record with the most employment, which accounts for 90 percent of the combined employment on average. Column (2) of Table A.2 reports the result, and column (1) reproduces the baseline of Table 5 for comparison. The results are essentially unchanged.

A.1.7 Coverage

Table A.1 reports how much of the NielsenIQ data the link covers. Of the 6,977 GS1 firms in the firm price index of Section 3.3, 4,727 are linked to NETS, and they account for 61 percent of these firms' sales in 2006. The estimation sample of Table 5 contains 3,679 of the linked firms, with 44 percent of these sales. The demand-side exposure is measured from NielsenIQ sales alone, so it is available for all 6,977 firms. Including the firms outside the link leaves the demand-side estimate essentially unchanged (columns (3) to (5) of Table A.2).

Table A.1: Coverage of the GS1–NETS Link

	Firms	Share of 2006 sales
GS1 firms in the firm price index	6,977	1.000
Linked to NETS	4,727	0.614
Estimation sample	3,679	0.441

Notes: Firms are GS1 prefixes in the firm price index of Section 3.3, which requires continuing products in every year from 2006 to 2019. Sales are these firms' NielsenIQ Retail Scanner sales in 2006, the year whose sales define the regression weights and the demand-side exposure. In every year from 2006 to 2019, the share lies between 0.61 and 0.63 for the linked firms and between 0.44 and 0.45 for the estimation sample. The estimation sample is that of Table 5, column (1). It excludes linked firms without an SIC code or without establishment employment in NETS.

A.1.8 From Matched Firms to Establishment Locations

NETS assigns each establishment (DUNS number) to the headquarters of the firm that owns it. For every matched GS1 firm we collect all establishments assigned to its matched headquarters records, with each establishment's county and employment in the initial year of each period (2006, 2010, and 2014). Summing employment by county and dividing by the firm's total gives the weights $\omega_{fd}^{\text{empl}}$ in equation (5.2). Employment in counties outside our immigration data enters the missing-share control included in all firm-level regressions.

Table A.3 compares the geography of a firm's sales with that of its production. Sales are dispersed and production is concentrated. The median firm sells in 18 counties in 3 states and earns a third of its sales in its largest county, while the median firm linked to NETS has establishments in

Table A.2: Firm-Level Prices and the GS1–NETS Link

	Δ $\ln P^{SV,Stack}$ (2006-2018)				
	(1)	(2)	(3)	(4)	(5)
$Immigration_{f,t}$ (D)	-0.072***	-0.069***	-0.071***	-0.076***	-0.074***
	(0.020)	(0.020)	(0.016)	(0.017)	(0.019)
$Immigration_{f,t}$ (S)	0.017***	0.013**			
	(0.005)	(0.005)			
Sample	Estimation	Estimation	All firms	Linked	Estimation
Firms	3,679	3,679	6,977	4,727	3,679
Obs.	11,037	11,037	20,931	14,181	11,037
First-Stage F-stat	3,089.1	3,083.2	6,405.9	5,841.3	4,158.9
Fixed Effects	Industry, Time	Industry, Time	Time	Time	Time

Notes: Each observation is a firm × period, on stacked four-year differences (2006–2010, 2010–2014, 2014–2018). The outcome is the firm price index of continuing products, as in Table 5. All columns report IV estimates, weighted by the firm's initial (2006) NielsenIQ sales, and control for the share of the firm's sales in counties without immigration data. Column (1) is column (1) of Table 5. Column (2) changes only the supply-side exposure: a GS1 prefix linked to more than one NETS headquarters record is assigned the exposure of its record with the most employment, instead of that of all its records combined. Columns (1) and (2) include industry (SIC 3-digit) and period fixed effects, control for the share of the firm's employment in counties without immigration data, and cluster standard errors by SIC 3-digit industry. Columns (3) to (5) include only the demand-side exposure, which is measured from NielsenIQ sales and is available for firms outside the link. Because these firms have no SIC code, the three columns include period fixed effects only and cluster standard errors by firm. Column (3) uses all 6,977 firms in the firm price index, column (4) the 4,727 firms linked to NETS, and column (5) the 3,679 firms in the sample of column (1). Immigration is measured in thousands. ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

a single county. The firms with most of the regression weight operate in many more locations, and the gap remains. Weighted by initial sales, the median number of counties is 1,482 for sales and 17 for establishments.[1]

Table A.3: Geographic Dispersion of Firm Sales and Production

		Percentiles					Sales-weighted
	Mean	p10	p25	p50	p75	p90	median
Panel A. Sales (NielsenIQ)							
Counties	195.5	1	4	18	119	741	1,482
States	12.4	1	1	3	17	46	49
Share of the largest county (%)	42.6	5.9	12.9	33.8	69.1	100.0	5.0
Panel B. Employment (NETS)							
Counties	15.2	1	1	1	2	12	17
States	3.7	1	1	1	2	7	11
Share of the largest county (%)	85.7	40.8	80.9	100.0	100.0	100.0	48.0

Notes: Distribution across the firms in the firm-level sample. Panel A covers all 6,977 firms and uses each firm's NielsenIQ Retail Scanner sales in 2006, the county shares that enter the demand-side exposure. Panel B covers the 3,774 firms that are linked to NETS and have establishment employment in 2006, and uses that employment, the county shares that enter the supply-side exposure. A firm's establishments are those of all headquarters records linked to its GS1 prefix. Percentiles and medians are computed separately for each row. The last column weights firms by their initial NielsenIQ sales, the regression weights.

[1]The two panels cover different firms. Restricting Panel A to the firms in Panel B gives very similar numbers (a median of 20 counties with sales, and a sales-weighted median of 1,314).

A.2 Identifying Immigrant Households in the NielsenIQ Panel

The 2008 Tell Us More About You (TUM) survey provides birthplace information for a subset of households in the NielsenIQ Consumer Panel. We use these data to estimate and evaluate an immigrant-household probability using demographics and purchase histories, then apply the model to households without observed status.

A.2.1 Labels, predictors, and estimation

We classify a household as immigrant if at least one household head was born outside the United States and its territories. Estimation and evaluation use TUM households for which all household heads report birthplace. The sample contains 19,522 households, including 1,056 immigrant households. Immigrants account for 5.4 percent of households without weights and 5.89 percent with Nielsen projection factors.

Demographic predictors include income, household size and composition, heads' age, education, occupation and employment, race, Hispanic origin, state, and reported birth years. We also include panel entry year and tenure. We use the earliest available response to each demographic question. For example, if education is missing when a household enters the panel, we use its first subsequent reported education.[2]

Purchase predictors describe both what households buy and where they shop. We pool purchases from 2004–2019 and construct expenditure shares by product module, department, brand within department, retailer, and retail channel. We also include annual expenditure, shopping trips, items per trip, years with reported purchases, and private-label expenditure shares.[3]

The full model also uses Magnet data on random-weight purchases, such as fresh produce, meat, and deli items that do not have standard UPCs. Nielsen collects these additional data from a subset of Consumer Panel households.

We combine regularized logistic regression with gradient boosting. Logistic regression penalizes large coefficients to limit overfitting with many predictors. Gradient boosting combines decision trees and allows nonlinear relationships and interactions between predictors (Ke et al., 2017). We recalibrate the boosting predictions using a logistic regression of observed status on their log odds. This step aligns predicted probabilities with immigrant frequencies in the training data. We average the logistic model's log odds and the recalibrated boosting model's log odds to form the household probability.[4]

We evaluate prediction using five household groups with similar immigrant shares. Each

[2]We represent categorical responses by indicator variables. Missing numeric values are replaced by medians among training households, and numeric variables are centered and scaled using their training-sample means and standard deviations. All these operations are repeated within each training sample.

[3]Years follow Nielsen's annual reporting files. Spending shares enter after a square-root transformation. This gives small shares more weight relative to dominant categories while preserving their ordering. Annual expenditure and trips enter in logs after adding one.

[4]Let p_h^L be the logistic prediction and $\widetilde{p}_h^G$ the recalibrated boosting prediction. The combined probability is $p_h = \text{logit}^{-1}\{[\text{logit}(p_h^L) + \text{logit}(\widetilde{p}_h^G)]/2\}$. Logistic regression uses an L2 penalty with $C = 0.1$. The base models are estimated without survey weights or class balancing. Recalibration and evaluation use projection factors.

household's prediction comes from models trained on the other four groups. Within each training sample, three additional groups provide predictions for recalibration without using a household's own label to generate its base prediction. For boosting, a separate part of the training sample determines how many trees to fit. We then refit using the selected number of trees. Neither preprocessing nor recalibration uses the outer evaluation households' labels. These comparisons use the historical TUM cohort on which we developed the method, rather than a newly collected validation sample.

A.2.2 Household prediction and the ACS comparison

Birthplace information in the American Community Survey (ACS) offers a much larger sample for learning how demographics predict immigrant status. We test whether relationships estimated in ACS improve prediction in the NielsenIQ Consumer Panel. We first estimate an immigrant-household probability in 2005–2007 ACS microdata using demographic variables available in both datasets. We then add this probability to the TUM model and recalibrate using only its training households.[5]

For the ACS comparison, we align the common demographic variables with Consumer Panel responses from 2008. These responses are available for 16,934 of the evaluation households. The remaining households retain their prediction without the ACS addition. To distinguish information learned in ACS from information in the additional demographic responses, we also compare each ACS specification with a model that has the same 2008 demographics but omits the ACS probability.

Table A.4 reports four specifications on the same households. The area under the receiver operating characteristic curve (AUC) measures how often an immigrant household receives a higher score than a native household in a randomly selected pair. An AUC of 0.5 corresponds to random ranking and 1 to perfect ranking. Log loss penalizes incorrect probabilities, especially confident errors. The Brier score is the average squared difference between a household's probability and its observed zero-or-one status. Smaller values of both measures indicate better prediction. We report the mean probability to compare the aggregate predicted share with the observed share.

Table A.4: Household prediction on a common TUM sample

Specification	AUC	Unweighted AUC	Log loss	Brier	Mean (%)
Demographics only	0.8058	0.7694	0.18075	0.04760	5.80
Demographics + ACS	0.8137	0.7781	0.17754	0.04668	5.93
Demographics + purchases	0.8772	0.8643	0.14951	0.03968	5.72
Demographics + purchases + ACS	0.8761	0.8645	0.15055	0.04009	6.22

Notes: All specifications use the same 19,522 complete-response households and five outer household folds. Metrics use Nielsen projection factors except the unweighted AUC. The observed weighted immigrant share is 5.89%. Evaluated probabilities do not substitute observed labels. ACS specifications add an independently estimated source probability and comparable 2008 demographics. Households without those demographics retain the corresponding prediction without ACS.

Adding purchases raises weighted AUC from about 0.81 to 0.88 and reduces log loss by about

[5]The ACS model uses a deterministic sample capped at one million households and survey-weighted gradient boosting. We identify household heads from the householder and an explicitly identified spouse or unmarried partner, requiring recognized birthplace information for both when present. The outcome measures geographic birthplace. Citizenship at birth is unavailable in the extract.

17 percent. The full model therefore ranks an immigrant household above a native household in about 88 percent of projection-weighted pairs. Its mean prediction is 5.72 percent, compared with 5.89 percent observed. We retain probabilities for aggregation rather than assigning every household a zero-or-one status using a 0.5 cutoff.

ACS information improves the demographic-only model. Relative to the model with the same 2008 demographic variables, it reduces log loss by about 2.08 percent. Once purchases enter the model, the corresponding improvement is about 0.08 percent and is not statistically distinguishable from zero. Adding ACS also gives slightly higher log loss than the preferred full specification in Table A.4. We therefore use demographics and purchases for the preferred model.[6]

A.2.3 Product immigrant intensity and customer income

A product's immigrant intensity is the share of its buyers or spending attributable to immigrant households. A household model that uses purchases of that product could mechanically associate its predicted status with the product. We therefore estimate ten additional household models, each excluding purchases from one Nielsen department. For a barcode in department d, we use $p_h^{(-d)}$, the probability estimated without purchases from d.

These models retain demographics and expenditure shares across modules, departments, and brands in the remaining departments. We renormalize shares by total expenditure in those departments. Brand shares are defined separately within each department, so a brand sold in several departments does not reintroduce the excluded purchases. We omit retailer and channel shares and all purchase-scale variables from these models. They use regularized logistic regression followed by logistic recalibration. Their weighted AUC ranges from 0.840 when excluding dry grocery to 0.867 when excluding general merchandise.[7]

Let H_b be the unique households buying barcode b, w_h their projection factors, and s_{hb}^+ their positive net spending on the barcode. We define buyer and expenditure intensities, respectively, as

$$I_b^{\text{buyer}} = \frac{\sum_{h \in H_b} w_h p_h^{(-d(b))}}{\sum_{h \in H_b} w_h}, \qquad I_b^{\text{spend}} = \frac{\sum_{h \in H_b} w_h s_{hb}^+ p_h^{(-d(b))}}{\sum_{h \in H_b} w_h s_{hb}^+}. \tag{A.1}$$

Each household enters the buyer measure once per barcode. For expenditure intensity, households that spend more on the barcode receive more weight. Annual measures use purchases and projection factors from the same year.[8]

We validate these measures using observed TUM status and predictions that exclude each

[6] The 95-percent intervals for the percentage reductions in log loss are $[1.18, 2.99]$ percent for demographics and $[-0.38, 0.55]$ percent for demographics and purchases. They use paired household resamples of the saved predictions, conditional on fitted models, and do not include survey-design uncertainty.

[7] Share denominators include all reported spending in classified departments other than the excluded department, including rare categories and purchases with missing module identifiers. The models exclude Magnet purchases because these purchases cannot be assigned unambiguously to a classified department. Differences from the full model also reflect the choice of learner and omission of retailer predictors.

[8] Barcode identity combines UPC and product version. Positive net spending sums $\max\{\text{paid}_r - \text{coupon}_r, 0\}$ across purchases before aggregation to a household-barcode pair.

evaluated household from training. The sample contains 16,936 buyer households in 2008. We construct observed immigrant shares on exactly the same purchases as the predicted shares. Table A.5 compares them among barcodes with at least 30 or 100 TUM buyers.[9]

Table A.5: Validation of product immigrant intensity

Minimum buyers	Measure	Barcodes	Predicted (%)	Observed (%)	Correlation	RMSE (pp)
30	Buyers	55,422	5.607	5.513	0.618	3.078
30	Spending	55,422	5.675	5.678	0.581	5.010
100	Buyers	17,923	5.479	5.365	0.669	2.058
100	Spending	17,923	5.493	5.526	0.605	3.806

Notes: The sample contains 2008 purchases by 16,936 TUM households. Eligibility counts unique buyers per UPC-version pair. Barcode moments are weighted by the denominator of the corresponding intensity measure. Predictions exclude purchases in the barcode's department. Observed shares use TUM status on the same purchases. RMSE is the root mean squared prediction error in percentage points.

Predicted and observed immigrant intensities have similar aggregate means. Among products with at least 100 buyers, their buyer-weighted correlation is 0.67, and the root mean squared prediction error is about 2.06 percentage points. Correlations are higher and errors lower than among products with at least 30 buyers, whose observed immigrant shares are based on fewer households.

For the immigrant-intensive groups in the price regressions, we use a two-step classification. First, we classify households using the department-excluded probabilities. For each department, we choose the probability cutoff so that the projection-weighted fraction of 2008 panel households meeting the cutoff is as close as possible to the weighted immigrant share in the TUM estimation sample. Equal probabilities receive the same classification, and the cutoff is held fixed across years. We then replace probabilities in the expenditure-intensity measure with these household indicators.

Second, we regress the annual barcode shares on barcode and product-group-by-year fixed effects, weighting by projected barcode expenditure. The barcode effects give scores that are fixed over time. Immigrant-intensive products are those above the 90th percentile across all scored UPC-version pairs, counting each product equally. Firm scores are averages of the continuous barcode scores, weighted by projected barcode expenditure pooled over 2004–2019. Immigrant-intensive firms are those above the separate 90th percentile across scored firms. Ties at either percentile belong to the other group. We classify firms using their continuous scores rather than averaging product classifications. The continuous buyer and expenditure measures evaluated in Table A.5 are distinct from these binary classifications used in the price regressions.

We also characterize products by the income of their buyers. We use reported household income-bracket midpoints and weight buyers by their positive net spending on the barcode and Nielsen projection factors. We regress these annual income averages on barcode and product-group-by-year fixed effects, weighting by projected barcode expenditure, and use the barcode effects as product income scores. Products are split at the median score. Firm scores average product scores with

[9] Classified departments account for about 79 percent of projection-weighted positive spending in this sample. The remaining spending is in Magnet purchases or purchases without a classified department. Barcode moments are pooled across classified departments rather than averaged from department-level statistics.

expenditure weights and use a separate median across firms. Both classifications are fixed over time.

A.2.4 Household prediction outside the labeled sample

We refit the preferred model on all 19,522 estimation households and estimate its recalibration using predictions that exclude each training household's own label. The resulting file covers 194,551 households. We use observed TUM status wherever it is determined, including incomplete responses that establish a foreign-born head. An incomplete response with no reported foreign-born head does not establish native status. In total, 20,294 households receive observed status, and the remaining households receive model probabilities.

For a TUM household, observed status can be used in each year that it appears in the panel. This identifies the household by its 2008 heads. The prediction model uses a pooled purchase history rather than a forecast made before purchases occur.

A.2.5 Aggregating household status

Household aggregates use the full model. Let q_h denote observed status when available and the full-model probability otherwise. With projection factor w_{hy}, annual CPG expenditure x_{hy}, and household size n_{hy}, the expenditure share and the share of household members associated with immigrant households are

$$\theta_y^E = \frac{\sum_h w_{hy} x_{hy} q_h}{\sum_h w_{hy} x_{hy}}, \qquad \theta_y^N = \frac{\sum_h w_{hy} n_{hy} q_h}{\sum_h w_{hy} n_{hy}}. \tag{A.2}$$

Both use the same eligible households. For the calibration years 2006, 2010, 2014, and 2018, pooled shares sum the numerators and denominators across years. The ratio θ^E/θ^N gives CPG expenditure per household member in immigrant households relative to the average across all households.[10]

These measures identify households with a foreign-born head. Household-size weighting therefore counts all members of such households, including members born in the United States, rather than identifying each individual's birthplace. They include immigrants of European and non-European origins, whereas the immigration shocks used in the regressions measure new non-European arrivals.

[10] Mean household size uses projection weights multiplied by q_h for immigrant households and $1 - q_h$ for the comparison group. Budget-share means first divide each household's CPG expenditure by its income and then average these ratios with the same weights, requiring positive income.

Appendix B Additional Tables and Figures

B.1 Motivation, Data, and First Stage

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

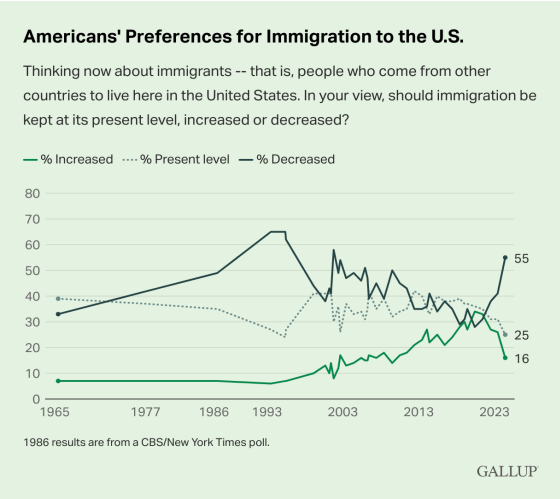

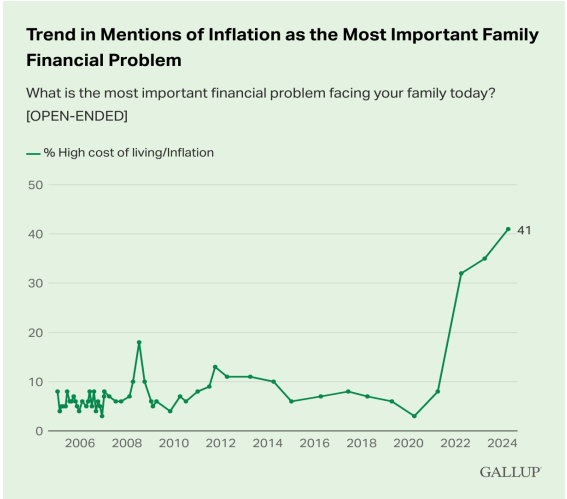

Notes: Gallup polls. The left panel is from "Sharply More Americans Want to Curb Immigration to U.S." (July 12, 2024), and the right panel from "Americans Continue to Name Inflation as Top Financial Problem" (May 2, 2024).

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

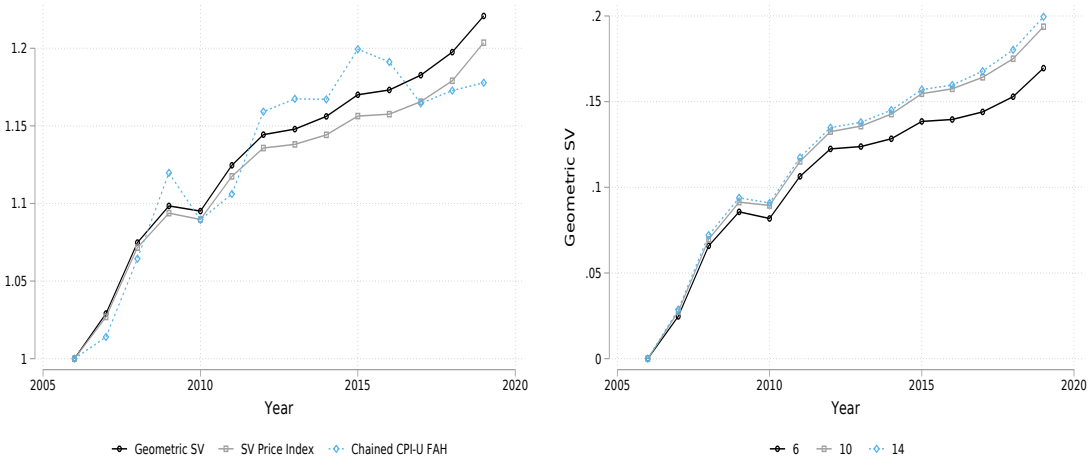

Notes: The left figure compares our constructed county-level price index using CPG data with the most closely related reported (national) BLS CPI. The CPI used is the "Food at Home" Chained Index. "Geometric SV" refers to the Sato-Vartia price index, while the "SV Price Index" incorporates the Feenstra (1994) variety correction. We aggregate the county-level index to the national level. In the right figure, we compare our constructed price index under several restrictions on the set of products in our sample. The baseline in the paper is the most restrictive one, where products must be available for 14 years. We also construct the Sato-Vartia price index (geometric aggregation) for an expanded set of products, lowering the restriction to 10 and 6 years.

Panel A of Table B.6 presents summary statistics for our main variables at the county level, and Panel B for the firm-level estimation sample over stacked four-year periods. Although CPG products are traded nationally, their price changes vary considerably across counties. Because our index holds each county's set of products fixed over time, this variation reflects price changes of continuing products rather than product turnover. We also report wages in levels (hundreds of 2012 dollars), and use this measure in regressions for comparability with Terry et al. (2026) wages results.

Table B.6: Summary Statistics

Panel A: County Level

	count	mean	sd	p25	p50	p75	min	max
Δ ln $P_{d,t}$ (CPG)	6405	7.34	4.11	4.33	6.73	9.96	-6.70	35.4
Δ ln $P_{d,t}$ (Non-durables)	699	5.95	5.40	-0.18	8.05	10.1	-5.19	13.4
Δ ln $P_{d,t}$ (Durables)	699	-3.65	3.34	-6.42	-3.69	-1.57	-9.81	3.48
Δ ln $P_{d,t}$ (Services)	699	9.24	2.58	7.83	8.74	10.1	2.95	17.5
Δ ln $P_{d,t}$ (BLS, Rent)	699	11.2	4.74	8.21	10.1	14.4	2.09	25.8
Δ ln $P_{d,t}$ (ACS, Rent)	6402	10.8	6.43	7.34	10.9	14.3	-29.9	53.5
Δ ln $HP_{d,t}$(CL)	2384	-2.12	15.4	-8.49	0.34	6.40	-87.1	51.3
Δ ln $HP_{d,t}$ (Z)	4831	6.00	16.2	-2.01	6.26	15.9	-108.6	57.2
Δ # $Establ._{d,t}$	6405	-0.15	0.72	-0.45	-0.14	0.15	-8.44	22.4
Δ $Employment_{d,t}$	6405	1.77	15.1	-0.45	0.085	1.08	-305.0	297.9
Δ ln $Wage_{d,t}$	6405	3.26	5.30	0.36	3.17	5.99	-30.2	56.2
Δ ln $Wage_{d,t}$ (Retail)	6382	1.95	6.88	-1.85	2.03	5.87	-74.6	64.7
Δ $Wage_{d,t}$	6405	13.1	23.3	1.38	12.1	23.2	-163.8	288.0
Δ $Wage_{d,t}$ (Retail)	6405	4.84	23.8	-4.65	5.01	14.7	-288.1	581.6
$Immigr_{d,t}$	6405	1.69	4.97	0.066	0.17	0.62	0.0026	34.2
$Immishock_{d,t}$	6405	-0.029	0.40	-0.087	-0.034	-0.0011	-1.99	2.60

Panel B: Firm Level

	count	mean	sd	p25	p50	p75	min	max
Δ ln $P_{f,t}$	11058	7.23	19.4	0	5.64	13.4	-315.6	358.4
$Immigr_{f,t}$ (Demand)	11058	33.7	38.5	6.83	22.2	45.3	0.016	299.3
$Immigr_{f,t}$ (Supply)	11058	31.9	55.8	1.65	10.7	30.7	0	299.3
$Immishock_{f,t}$ (Demand)	11058	4.41	9.85	0.11	1.03	4.37	-2.53	108.8
$Immishock_{f,t}$ (Supply)	11058	5.15	16.6	-0.0047	0.23	2.28	-2.57	108.8
Initial sales ($ million)	11058	4.14	25.5	0.011	0.092	0.74	0.00000080	649.5

Notes: Panel A: 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)-(3.3). 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, excluding distressed sales, 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 (3.4), constructed using residualized predicted ancestry. Both Immigr and Immishock winsorize top and bottom 1% of the sample. Panel B: Summary statistics for the firm-level estimation sample of Table 5: 3,679 firms over three stacked four-year periods (2006–2010, 2010–2014, 2014–2018). Δ ln P is the change in the firm's Sato-Vartia price index of continuing products, in percent. Immigr (Demand) and Immigr (Supply) are the firm's immigration exposures, built from county non-European immigrant inflows (in thousands) over the counties where the firm sells and where it produces (Section 5); Immishock are the corresponding instruments. Initial sales are the firm's 2006 NielsenIQ sales in millions of dollars, which serve as regression weights. Statistics are unweighted.

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

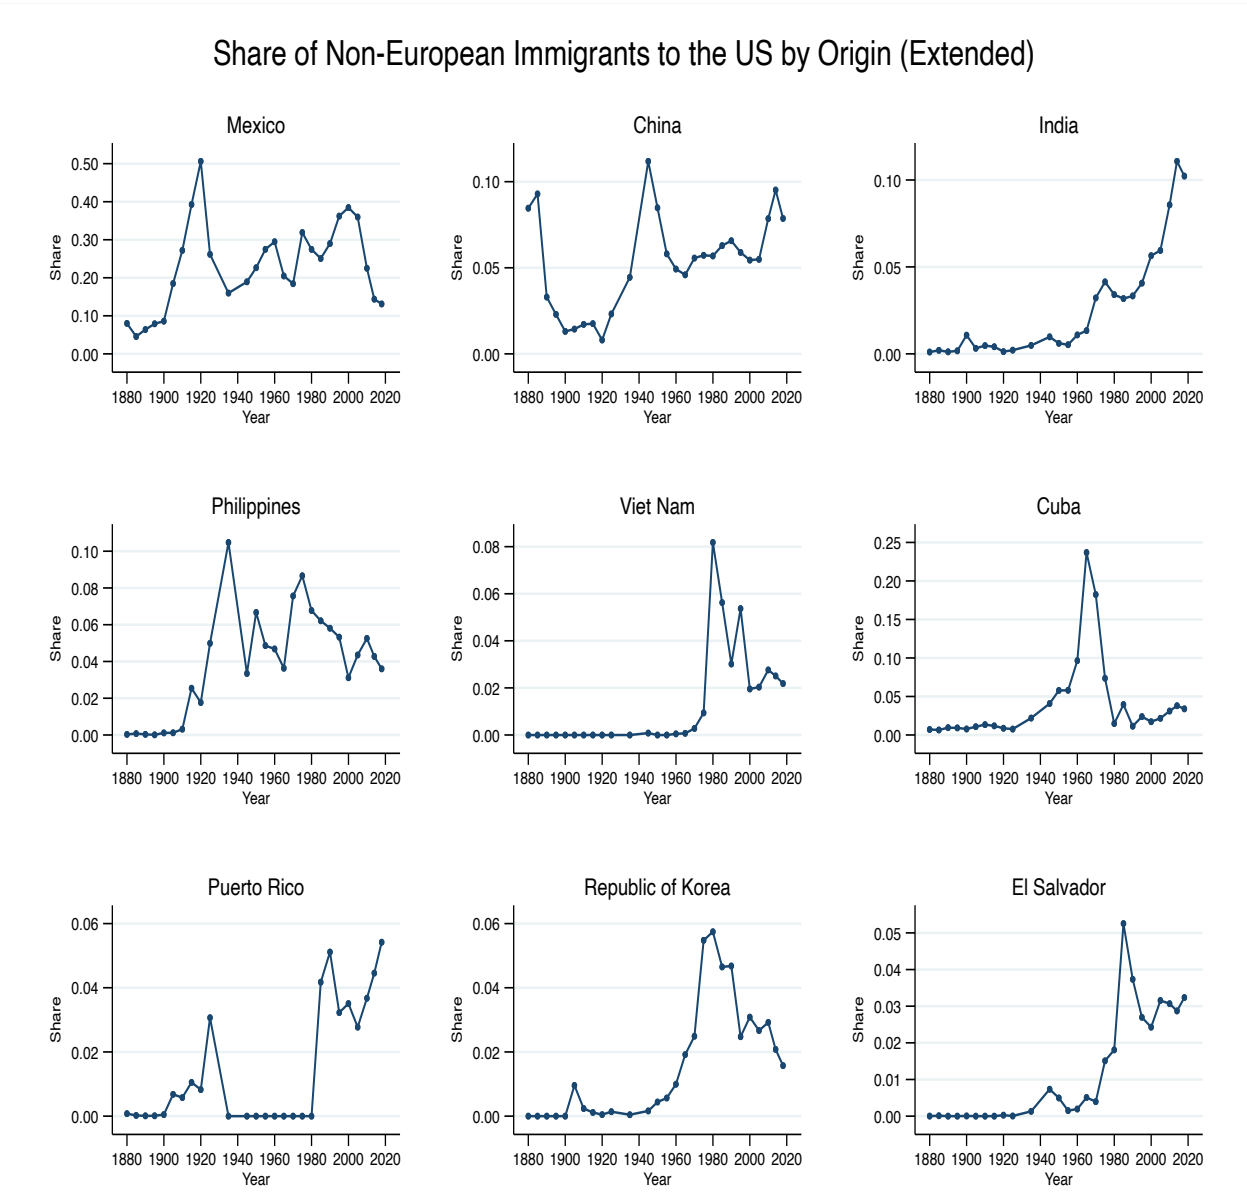

Notes: We plot the share of non-European immigration into the US from the 9 non-European origins with the largest immigration into the US. For the first 5, this is effectively an extension of Terry et al. (2026). We extend it to more countries to reflect the important share in this period from Central America and Korea.

B.2 Instrument Construction

This appendix gives the full construction of the predicted-ancestry instrument summarized in Section 3.4, following Terry et al. (2026).

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

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

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

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

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

Here $\mathcal{T}$ is the set of years with push-pull terms (every five years from 1880 to 1925, then 1935, 1945, every five years from 1950 to 2010, and 2014), and we estimate equation (B.2) separately for each ancestry year $t \in \{1975, 1980, \ldots, 2010, 2014\}$, using only push-pull terms from years $\tau \leq t$. The term $I_{o,-r(d),\tau}$ is total immigration from origin o in year τ, excluding migration to the census region $r(d)$ containing d (we use nine U.S. census divisions). This leave-out measure captures exogenous push factors, such as wars, famines, and crises, that drove emigration from o. The term $I_{Europe,d\tau}/I_{Europe,\tau}$ is the share of all European immigrants in year τ who settled in county d, capturing d's pull attractiveness to an unrelated group through transportation, port access, climate, and established communities. The fixed effects $\delta_{or(d)t}$ (origin by destination-region by time) and $\delta_{c(o)dt}$ (origin-continent by destination-county by time) absorb broad geographic settlement patterns, and X'_{od} collects two time-invariant dyadic controls, origin-destination distance and latitude difference. The interaction of origin-specific push and destination pull generates plausibly exogenous variation in where non-European immigrants' descendants live today.

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

We obtain predicted ancestry as the fitted push-pull component, with each push-pull term first residualized against the fixed effects and controls:

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

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

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

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

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

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

The second term corrects for region size. Larger regions mechanically receive more immigrants, so we scale by European immigration into the region relative to European immigration into all other regions, which makes shifters comparable across regions. Because this scaling reflects European immigration and enters symmetrically across origins, it contains no origin-specific information, and the instrument's identifying variation comes only from its interaction with origin-specific predicted ancestry. Aggregating across origins gives the county instrument Z_{dt} in equation (3.4), whose variation comes from three components plausibly orthogonal to current local conditions: historical settlement embedded in predicted ancestry, contemporary origin-country outflows in the leave-out region, and historical destination attractiveness to Europeans. We estimate equation (B.4) over nine periods at five-year intervals through 2010 and four-year intervals thereafter (2014 and 2018), pairing each lagged predicted ancestry $\hat{A}_{od,t-1}$ (reaching back to 1975) with the contemporaneous shifter.

State-Level Instrument. For the state-level specifications in Table 7, we rebuild the instrument at the state level with the same procedure. We sum immigration and ancestry over the counties in each state, by origin country and year, and average the distance controls over the state's counties. Predicted ancestry, the push and pull components, and the first stage then follow equations (B.2)–(B.5) with the state in place of the county: the push component leaves out the state's own census division, the pull component uses the state's share of European immigrants, and the fixed effects and controls are the same as for the county instrument. The state instrument is the sum over origins of the first-stage coefficient times the predicted ancestry interacted with the push-pull term, as in equation (3.4).

Table B.7: First Stage Results: Immigration on Push-Pull Instruments at Country-County Level

	$\text{Immigration}^t_{o,d}$				
	(1)	(2)	(3)	(4)	(5)
$\hat{A}_{o,d,1975} \times \tilde{I}_{o,-r(d),1980}$	0.0035***	0.0035***	0.0035***	0.0035***	0.0034***
	(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)
$\hat{A}_{o,d,1995} \times \tilde{I}_{o,-r(d),2000}$	0.0004***	0.0004***	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)
$I^t_{Euro,d}$				0.0081***	
				(0.0028)	
$\text{I}^t_{o,-r(d),t} \frac{I_{Europe,r(d)\tau}}{I_{Europe,-r(d)\tau}}$					0.2602***
					(0.0952)
Obs.	4,607,847	4,607,847	4,607,847	4,607,847	4,607,847
adj R-sq	0.600	0.600	0.655	0.655	0.670
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: The specification follows Terry et al. (2026), to show we can replicate their instruments. In all cases, the push-pull instruments follow the specification in (B.4), with one instrument for every year t. We extend the data by adding 2014 and 2018. The first column represents a specification with no controls or fixed effects, the second column adds distance controls, and column (3) adds the full set of fixed effects. The third column is our baseline specification. In column (4) we additionally control for concurrent European immigration, while in the last column we instead control for the contemporaneous interaction term from the ancestry prediction. Standard errors are clustered by country. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Figure B.4: Geographic Distribution of Predicted Immigration Shocks, 2006-2018

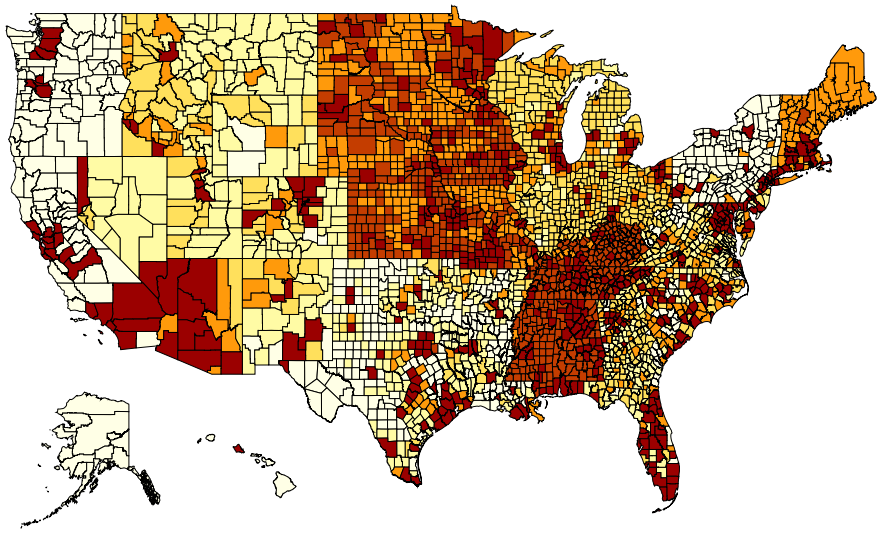

Notes: Heat map shows the predicted immigration shock Z_{dt} from equation (3.4) at the county level, cumulated over 2006-2018. Darker colors indicate higher values of the shock, shown in six quantile groups. Because predicted ancestry is residualized (Section 3.4), the shock averages zero across counties and is negative in 86 percent of them, so the map shows where predicted inflows are relatively large rather than their level. The instrument uses only non-European immigration and is based on historical settlement patterns interacted with recent origin-country migration patterns. White areas indicate counties not included in the sample due to missing data.

B.3 County-Level Analysis

Table B.8: Robustness to Alternative Price Index Formulations

	Δ P^{Jevons} (2006-2018)		Δ P^{Laspyres} (2006-2018)		Δ P^{Passche} (2006-2018)		Δ P^{Tornqvist} (2006-2018)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\text{Immigr}_{\text{d,t}}$	-0.022***	-0.013**	-0.071***	-0.069***	-0.069***	-0.064***	-0.070***	-0.067***
	(0.008)	(0.006)	(0.012)	(0.011)	(0.014)	(0.013)	(0.013)	(0.012)
Obs.	6,405	6,405	6,405	6,405	6,405	6,405	6,405	6,405
First-Stage F-stat	251.0	98.3	251.0	98.3	251.0	98.3	251.0	98.3
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 (3.4) 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.9: Long-Difference (one period)

	Δ $P^{SV,LongDiff}$ (2006-2018)					
	(1)	(2)	(3)	(4)	(5)	(6)
$Immigr_{d,t}$	-0.040***	-0.029***	-0.029***	-0.055***	-0.035***	-0.036***
	(0.009)	(0.004)	(0.006)	(0.011)	(0.007)	(0.007)
Obs.	2,135	2,134	2,134	2,135	2,134	2,134
First-Stage F-stat				263.3	92.8	76.8
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 2, 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, employment-to-population ratio, and urban indicator, each interacted with the year indicators.

Table B.10: County Price Regressions: Robustness to Further Controls

	Δ $P^{Sato\text{-}Vartia,Stacked}$ (2006-2018)						
	(1) Baseline	(2) + Border	(3) + ICE Arrests	(4) + Lead	(5) StatexYear FEs	(6) Non-Winsorized	(7) + County FE
$Immigr_{d,t}$	-0.066***	-0.066***	-0.066***	-0.149**	-0.044***	-0.013***	-0.059**
	(0.011)	(0.011)	(0.010)	(0.064)	(0.008)	(0.003)	(0.028)
Mexico Border		-0.288**					
		(0.133)					
total ICE Arrests			0.041***				
			(0.012)				
Immigration (t+4)				0.092			
				(0.057)			
Obs.	6,405	6,405	6,405	4,270	6,402	6,405	6,405
First-Stage F-stat	74.0	73.0	79.9	13.3	68.7	41.4	318.4
Fixed Effects	State, Time	State, Time	State, Time	State, Time	State-Time	State, Time	County, Time

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

Table B.11: County Price Regressions: Robustness to Alternative Population Specifications

	Δ $P^{\text{Sato-Vartia,Stacked}}$ (2006-2018)					
	(1)	(2)	(3)	(4)	(5)	(6)
Immigration	-0.066***		-0.061***	-0.063***	-0.064***	-0.059***
	(0.011)		(0.012)	(0.010)	(0.010)	(0.010)
Immig/Pop		-1.111***				
		(0.222)				
Δ Native Pop				0.002***		0.005***
				(0.001)		(0.002)
Δ Native Pop, Lag					0.002***	
					(0.001)	
Obs.	6,405	6,405	6,405	6,403	6,403	6,402
First-Stage F-stat	74.0	9.1	51.1	71.0	74.4	38.8
Initial Pop × Period	No	No	Yes	No	No	No
Δ Native Pop IV'd by Lag	No	No	No	No	No	Yes
Fixed Effects	State, Time	State, Time	State, Time	State, Time	State, Time	State, Time

Notes: This table reports estimates of immigration effects on county-level CPG prices (Sato-Vartia index). All specifications use stacked four-year differences (2006–2010, 2010–2014, 2014–2018), are weighted by initial county-level NielsenIQ sales, include state and period fixed effects, and cluster standard errors by state. Column (2) uses the immigrant-to-population ratio as the endogenous variable, with both the ratio and the instrument normalized by initial (2005) county population; controls include the urban indicator interacted with period dummies. Columns (1) and (3)–(6) use immigration levels (in thousands) as the endogenous variable, instrumented with baseline predicted shocks, and include controls for initial county GDP per capita, employment-to-population ratio, and urban indicator, each interacted with period dummies. Column (3) additionally controls for initial (2005) county population interacted with period dummies, absorbing the concern that larger counties may exhibit differential price trajectories. Column (4) adds a control for the change in native population, constructed from the ACS as the change in total county population minus the total (European and non-European) immigration over the same period, using consistent survey weights throughout. Column (5) instead controls for the change in native population over the previous period, constructed in the same way (2000–2005 for the first period), which is determined before the immigration inflow. Column (6) instruments the change in native population with its previous-period change. Immigration and the immigration shock are winsorized at the top and bottom 1% in all columns. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table B.12: Permutation Tests for the Baseline Specification

Instrument	(1) Coefficient (Mean)	(2) Coefficient (St. Dev.)	(3) Standard Error (Mean)	(4) Permutation p-value
Panel A: First Stage				
Actual	7.698		0.895	
Reassigned across all counties	-0.032	0.960	0.837	<0.001
Reassigned within period	-0.026	0.986	0.824	<0.001
Reassigned within period and division	-0.253	0.823	0.418	0.003
Panel B: Reduced Form				
Actual	-0.511		0.086	
Reassigned across all counties	-0.003	0.154	0.121	<0.001
Reassigned within period	-0.000	0.165	0.129	<0.001
Reassigned within period and division	0.154	0.185	0.095	0.051

Notes: This table reports permutation tests for our baseline county specification (column (6) of Table 2). All regressions use stacked four-year differences (2006–2010, 2010–2014, 2014–2018), include state and period fixed effects and controls for initial county GDP per capita, the employment-to-population ratio, and the urban indicator interacted with period dummies, are weighted by initial county-level NielsenIQ sales, and cluster standard errors by state. In each of 1,000 draws, we randomly reassign the predicted immigration shock across county-period observations, either across the full sample, within period, or within period and census division, and re-estimate the first stage (Panel A) and the reduced form (Panel B). The first row of each panel reports the estimate with the actual instrument. Columns (1) and (2) report the mean and standard deviation of the coefficient on the reassigned instrument across draws, and column (3) reports its mean standard error. Column (4) reports the permutation p-value, the share of draws in which the absolute t-statistic on the reassigned instrument is at least as large as with the actual instrument.

Table B.13: Sensitivity to Origin Concentration: Excluding or Including Only the Largest Origins

	Δ P$^{\text{Sato-Vartia,Stacked}}$ (2006-2018)				
	Mexico	China	India	Philippines	Guatemala
	(1)	(2)	(3)	(4)	(5)
Panel A: Excluding given origin					
$\text{Immigr}_{\text{d,t}}$	-0.068***	-0.071***	-0.066***	-0.068***	-0.068***
	(0.022)	(0.011)	(0.011)	(0.012)	(0.011)
Obs.	6,405	6,405	6,405	6,405	6,405
First-Stage F-stat	25.4	70.8	68.7	69.3	71.1
AR Wald test p-value	0.000	0.000	0.000	0.000	0.000
Panel B: Including only given origin					
$\text{Immigr}_{\text{d,t}}$	-0.071***	-0.059***	-0.066***	-0.052***	-0.077***
	(0.011)	(0.014)	(0.019)	(0.012)	(0.011)
Obs.	6,405	6,405	6,405	6,405	6,405
First-Stage F-stat	53.4	163.9	39.7	158.5	30.5
AR Wald test p-value	0.000	0.000	0.000	0.000	0.000

Notes: The table repeats the baseline specification of Table 2, column (6), changing only the origins used to build the instrument, following Appendix Table 11 of Terry et al. (2026). Each origin's contribution to the instrument in equation (3.4) is $\hat{b}_t \hat{A}_{od,t-1} \times \tilde{I}_{o,-r(d),t}$. Columns are the five origins with the largest Goldsmith-Pinkham et al. (2020) Rotemberg weights in the baseline estimate. Panel A drops the given origin's contribution from the instrument; Panel B uses only that origin's contribution. As in the baseline, each instrument is winsorized at 1% by period. All specifications use stacked four-year differences (2006–2010, 2010–2014, 2014–2018), include state and period fixed effects and county controls interacted with period, and are weighted by initial county CPG sales. We report the first-stage F-statistic on the excluded instrument and the p-value of the Anderson-Rubin Wald test. Standard errors clustered by state. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.14: County Price Regressions: Does Past Inflation React to Contemporaneous Immigration?

	Δ P (BLS, Nondurables)		
	(1) Baseline	(2) Pre-Period Placebo	(3) Pre-Period Placebo
Immigration	-0.069**		
	(0.034)		
Immigration (t+5)			0.029
			(0.021)
Immigration (t+15)		0.025	
		(0.018)	
Obs.	783	715	1,239
Wild cluster bootstrap p-value	0.019	0.139	0.238
Fixed Effects	State, Time	State, Time	State, Time
Method	IV	IV	IV

Notes: This table reports pre-trends and placebo specifications for the main price regressions. All columns use BLS regional nondurables price changes as the dependent variable, where pre-period prices are available. All specifications are IV regressions with state and time fixed effects, controls for county urbanization, 2006 GDP per capita, and the employment-to-population ratio interacted with year, and observations weighted by 2005 county population. Standard errors are clustered by state, and below each coefficient we report the wild cluster bootstrap p-value (state clusters, Webb weights, 9,999 draws), on which the significance stars are based. Column 1 reports the baseline specification matching contemporaneous prices and immigration over 2006–2018 (Column 2 of Table 4 - which examines the sectoral BLS regional data). Column 2 is the central placebo: BLS nondurables price changes during 1990–2005 are regressed on immigration flows realized fifteen years later (2006–2018). Column 3 pools the pre- and post-periods (1990–2018), entering immigration just one period ahead. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Summary of Tables Table B.8 reports the baseline specification, but with price indices constructed with various aggregation methods. Using Laspeyres, Paasche, and Tornqvist weights results in remarkably similar results to Sato-Vartia, implying that effects are not due to major compositional changes in the basket of goods. The Jevons index is an unweighted average of prices and thus more likely to be driven by small outliers.

Table B.9 displays the results with a long-difference specification, an alternative to our baseline "stacked" specification. Prices reflect the 12-year difference in the chained price index by county, while immigration now accounts for all new immigrants by county from 2006 to 2018. Thus, the variation is strictly cross-sectional; relative to the stacked design it reduces reliance on period-by-period timing assumptions. The disinflationary effect of immigration persists. With state fixed effects and controls, the long-difference estimate (-0.036) is close to the stacked estimate with state-by-period fixed effects (-0.044, Table B.10, column 5), which uses the same within-state variation, so each immigrant has a similar effect on the price level over twelve years as over four.

Table B.10 includes miscellaneous specifications with further controls. Column (2) controls for proximity to the Mexican border, and Column (3) controls for ICE arrest activity; the immigration coefficient remains stable across both. In column (4), to test for anticipation and pre-existing trends, we estimate a specification that includes both the contemporaneous immigration shock and its one-period lead, each instrumented by the corresponding predicted shock. This isolates whether leads of the treatment predict prices conditional on the contemporaneous shock, which we can confirm is

not the case. Only contemporaneous immigration is negative and statistically significant. With 2 instruments, the F-stat predictably drops. Column (5) replaces state and time fixed effects with state-by-year fixed effects, yielding a somewhat smaller but still significant coefficient. Column (6) reports results *without winsorization* (and, like column 7, without sales weights), while Column (7) is non-winsorized and also includes county fixed effects (as county FEs control for the outliers' potential distortion). The negative effect of immigration on prices is robust across all specifications.

Table B.11 presents robustness checks for the county-level price regressions. Column (2) uses the immigrant-to-population ratio as the independent variable instead of the level of immigration. The instrument is re-constructed here following the baseline methodology covered in Section 3.4, using immigrants relative to the initial county population. The coefficient of -1.11 implies that a 1 percentage point increase in the immigrant share lowers CPG price growth by approximately 1.11 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) finds that a 10 percent increase in the number of low-skilled immigrants reduces prices of immigrant-intensive services by close to 2 percent; her estimates capture a supply-side wage channel in services where immigrants are concentrated, which differs from the broad set of consumer packaged goods we study. Zachariadis (2012) also uses a log-log specification but with the immigrant-to-employment ratio across countries, and finds that a 10 percent increase reduces prices by about 2.9 percent. Despite these differences in specification, sector, and context, the magnitudes are broadly comparable.

We then also check whether the baseline specification is sensitive to population controls. In Column (3) we control for the initial (2005) population levels, and the interaction with time period indicators. This is to check whether the disinflationary effects of immigration are related to the size of the county (for example because larger counties will get more immigrants in levels). The baseline immigration coefficient remains quantitatively unchanged, suggesting our instrument is already handling this. Column (4) further shows that the price effect is unchanged when controlling for native population changes, a robustness check akin to Lach (2007). This is important for two reasons. First, native displacement does not seem to be a major factor. Second, the coefficient on native population change is positive: a larger population raises the price level, as the scale channel in the model of Section 7 predicts, while immigration still lowers it. Because the native population change may itself respond to immigration, column (5) instead controls for its change over the previous period, and column (6) instruments it with that change; both results hold, with a similar immigration coefficient and a positive coefficient on native population change.

Table B.12 reports permutation tests for our baseline county specification. In each of 1,000 draws, we randomly reassign the predicted immigration shock across county-period observations, either across the full sample, within period, or within period and census division, and re-estimate the first stage and the reduced form with the reassigned instrument. The coefficients on the reassigned instrument are far from the actual estimates in all three cases, and close to zero on average when the instrument is reassigned across the full sample or within period. The absolute t-statistic of the

actual reduced form exceeds that of more than 99.9% of draws in the first two tests (permutation p-value below 0.001) and of 94.9% of draws when the instrument is reassigned within period and census division). In the first stage, the permutation p-values are below 0.01 in all three tests.

Table B.13 examines how much the baseline estimate depends on individual origins, following Appendix Table 11 of Terry et al. (2026). The columns are the five origins with the largest Goldsmith-Pinkham et al. (2020) Rotemberg weights in the baseline estimate of Table 2, column (6). Panel A drops each origin's contribution from the instrument, and Panel B uses only that origin's contribution. The estimate is between -0.066 and -0.071 when one origin is dropped and between -0.052 and -0.077 when only one origin is used, all significant at the 1 percent level. Dropping Mexico lowers the first-stage F-statistic to 25 but leaves the estimate essentially unchanged (-0.068).

Table B.14 conducts a robustness check to check for pre-trends. The goal is to confirm that past inflation does not respond to contemporaneous immigration, which we do in several ways. Since the baseline CPG prices are not available in the "pre-period" (before 2006), we instead rely mostly on the non-durable price index available for 27 metro areas from the BLS (as introduced in the main text). This price index responds to immigration in the most similar way from 2006-2018 (see Table 4 in the main text) – in fact we reproduce that result in Column (1), using the baseline specification with contemporaneous prices and immigration from 2006-2018. Column (2) produces the most natural robustness test: we lag prices by close to 15 years to produce regional inflation rates from 1990-1995, 1995-2000, and 2000-2005 (to match the ACS sample periods). Then we regress the past inflation rate to contemporaneous immigration, using our baseline sample of stacked immigration shocks from 2006-2018. Importantly, past inflation does not respond in any meaningful way to contemporaneous immigration. Columns (3) uses the full 1990-2018 sample for prices and immigration in order to show that if we regress inflation on a *one* period "lead" of immigration over this full sample, past inflation once again does not respond to future immigration. Note also that column (4) in Table B.10 controls for future immigration to test for pre-trends in a specification with CPG prices.

Table B.15: Alternative Rent Prices from ACS

	Δ P^{SV,Stacked} (2006-2018)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Immigr$_{d,t}$	0.070***	0.076***	0.077***	0.064***	0.078**	0.101**	0.008
	(0.018)	(0.019)	(0.017)	(0.024)	(0.035)	(0.038)	(0.038)
Obs.	6,402	6,402	6,402	6,402	6,402	6,402	699
First-Stage F-stat				22.5	48.7	43.6	130.1
Fixed Effects	Time	State, Time	State, Time	Time	State, Time	State, Time	State, Time
Method	OLS	OLS	OLS	IV	IV	IV	IV
Controls			Yes			Yes	Yes
Sample							Cities

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

Table B.16: Aggregate Price Indices: BLS Cities and HHNS States

	Δ P (BLS City,All Items)	Δ P (HHNS State)
	(1)	(2)
$\text{Immigr}_{d,t}$	-0.039**	-0.018
	(0.019)	(0.011)
Obs.	783	5,172
First-Stage F-stat	122.7	155.0
Fixed Effects	State, Time	Time
Controls	Yes	Yes

Notes: This table reports aggregate prices at the city level and state level, from the BLS and from data shared by Hazell et al. (2022), respectively. The outcome in column (1) is the "All Items" index from the BLS city-level CPI data (27 metro areas), which we merge to the county level. In column (2), the outcome is the state-level inflation rate for tradable industries produced by Hazell et al. (2022). Annual inflation data are transformed to four-year rates. Column (2) keeps the county level of aggregation, with prices varying only by state, and includes only period fixed effects since prices vary by state. Controls: 2006 county GDP per capita, the employment-to-population ratio, and the urban indicator, each interacted with period indicators. Column (2) reports standard errors clustered by state, while column (1) reports heteroskedasticity-robust standard errors due to the small number of states. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table B.15 turns from CPG prices to rents, using median gross rents from the ACS (NHGIS). In contrast to the disinflationary effect on retail goods, immigration *raises* rents. The IV estimate is positive and significant across specifications, from 0.064 to 0.101 in columns (4)–(6), consistent with Saiz (2007). The opposite sign indicates that the price decline is not a mechanical feature of high-immigration counties.

Table B.16 checks whether the disinflationary effect shows up in coarser, more aggregate price data. Column (1) uses the BLS city-level "All Items" CPI for 27 metro areas and yields a negative effect (-0.039). Column (2) uses state-level tradable inflation from Hazell et al. (2022), with period fixed effects only since prices vary by state. The estimate is also negative (-0.018), though not significant. These indices are far coarser than our retail scanner data, so power is limited, but the sign lines up with the CPG results and indicates the effect is not unique to the retail CPG index.

Table B.17: Immigration Effects on Establishments, Employment, and Wages

	# Establ.	Emp.	Wages			Wages (Retail)
	(1)	(2)	(3)	(4)	(5)	(6)
$\text{Immigr}_{d,t}$	0.014***	0.703***	0.240			-0.025
	(0.005)	(0.258)	(0.152)			(0.147)
$\text{Immigr}_{d,t}$ (over25)				0.336	-1.176***	
				(0.237)	(0.409)	
$\text{Immigr}_{d,t} \times \text{EducYears}_{d,t}$					1.833***	
					(0.617)	
Obs.	6,405	6,405	6,405	6,405	6,405	6,405
First-Stage F-stat	43.6	43.6	43.6	38.1	4.8 / 2.3	43.6
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 across all industries per 1,000 residents and column (2) the change in employment (in thousands). Columns (3)-(5) report wage effects: column (3) uses all immigrants as the endogenous variable; columns (4)-(5) restrict to immigrants aged 25 and older and additionally interact immigration with the (demeaned) average years of education among immigrants. Column (6) reports effects on retail wages. All wages are average annual real wages (in hundreds of 2012 dollars) deflated by the national PCE index. In column (5), the first-stage F-statistics are the Montiel Olea–Pflueger effective F-statistics for the two endogenous regressors, as in Table 3. All specifications instrument immigration with predicted shocks from equation (3.4) 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 B.18: Immigration Effects on Wages: Robustness

Panel A: 2006-2018						
	Level Wage		Log Wage		Retail Wage	
	(1)	(2)	(3)	(4)	(5)	(6)
$\text{Immigr}_{d,t}$	0.240	-0.003	-0.005	-0.006**	-0.025	-0.024*
	(0.152)	(0.012)	(0.025)	(0.003)	(0.147)	(0.014)
Obs.	6,405	6,405	6,405	6,405	6,405	6,405
First-Stage F-stat	43.6	41.4	43.6	41.4	43.6	41.4
Fixed Effects	State, Time	State, Time	State, Time	State, Time	State, Time	State, Time
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Winsorized	Yes	No	Yes	No	Yes	No
Panel B: Education Heterogeneity (2006-2018)						
$\text{Immigr}_{d,t}$	-1.176***	-0.281***	-0.199***	-0.037***	-0.453*	-0.221
	(0.409)	(0.043)	(0.042)	(0.005)	(0.247)	(0.138)
$\text{Immigr}_{d,t} \times \text{EducYears}$	1.833***	0.603***	0.224***	0.057***	0.806**	0.397
	(0.617)	(0.083)	(0.058)	(0.005)	(0.315)	(0.308)
Obs.	6,405	6,405	6,405	6,405	6,405	6,405
Fixed Effects	State, Time	State, Time	State, Time	State, Time	State, Time	State, Time
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Winsorized	Yes	No	Yes	No	Yes	No
Panel C: 1990-2005						
$\text{Immigr}_{d,t}$	1.117***	0.106***	0.149***	0.012***	0.134*	0.004
	(0.177)	(0.025)	(0.030)	(0.003)	(0.077)	(0.005)
Obs.	9,398	9,398	9,398	9,398	9,145	9,145
First-Stage F-stat	78.5	173.5	78.5	173.5	78.9	173.9
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 (3.4) and include state and period fixed effects and county controls. Odd-numbered columns (1), (3), (5) winsorize the immigration regressor; even-numbered columns (2), (4), (6) do not. Panel A uses the 2006-2018 sample. Columns (1)-(2) report effects on average annual real wages in levels (in hundreds of 2012 dollars); columns (3)-(4) report effects on log wages; columns (5)-(6) report effects on retail sector wages. Panel B reports education heterogeneity for the 2006-2018 sample, interacting immigration with average years of education among immigrants following Terry et al. (2026) (immigrants restricted to those aged 25 and older); both immigration and its interaction with education are treated as endogenous and instrumented by the shift-share instrument and its interaction with immigrant education. Panel C replicates Panel A for the 1990-2005 sample for comparison with Terry et al. (2026). Standard errors clustered by state. ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

Table B.17 reports the local economic outcomes discussed in Section 4, all drawn from the QCEW. Column (1) is measured per 1,000 residents and column (2) in thousands, so the estimates imply that 10,000 immigrants raise the number of establishments across all industries by about 0.14 per 1,000 residents, and employment by about 7,000, over four years. Column (3) reports average annual real wages in hundreds of 2012 dollars, deflated by the national PCE index: the point estimate for the same inflow is about $240, against a mean four-year change of $1,310, though it is not statistically significant. Columns (4) and (5) restrict to immigrants aged 25 and older and interact immigration with the average years of education among them, following Terry et al. (2026), treating both as endogenous and instrumenting the interaction with the shift-share instrument interacted with immigrant education. The interaction is positive and significant, so wages rise more where arriving immigrants are more educated, though the effective first-stage F-statistics are only 4.8 and 2.3 once a second endogenous regressor enters, so the instruments are weak in this column. Column (6) repeats column (3) for the retail sector alone and returns essentially zero, about −$25 per 10,000 immigrants with a standard error of 147.

Table B.18 provides robustness checks for the wage results reported in Table B.17. Panel A examines sensitivity to winsorization and the choice of log versus level specification for the 2006-2018 sample. The positive but insignificant point estimate in levels is largely driven by winsorization, with non-winsorized estimates close to zero and more comparable to Terry et al. (2026). 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. (2026).

Table B.19: Price, Revenue, and Quantity: Branded versus Private Label

	Price	Revenue	Quantity
Panel A. RMS (store sales, county, weighted)			
Branded (GS1)	-0.067***	-0.011	0.055
	(0.012)	(0.057)	(0.064)
Private label	-0.055***	0.143**	0.198***
	(0.016)	(0.060)	(0.066)
Branded + priv. label	-0.067***	0.009	0.076
	(0.011)	(0.056)	(0.062)
Panel B. HMS (consumer panel, state, unweighted)			
Branded	-0.088***	0.155	0.243
	(0.030)	(0.198)	(0.217)
Private label	-0.099**	0.188	0.286
	(0.043)	(0.191)	(0.210)
Branded + priv. label	-0.087***	0.170	0.257
	(0.028)	(0.197)	(0.215)

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

Table B.19 reports price, revenue, and quantity responses in two datasets and for three product scopes each. Panel A uses the NielsenIQ Retail Scanner census (RMS) at the county level, weighted by base-year (2006) sales, with state and year fixed effects. Price is the Sato–Vartia index, revenue is log store sales, and quantity is the implied real quantity, revenue deflated by the price index ($\Delta \ln Q = \Delta \ln \text{Revenue} - \Delta \ln \text{Price}$). The three rows draw on separately constructed price files: branded goods are products with a GS1 manufacturer prefix, private label are store brands with no GS1 prefix, and the pooled row combines the two. Private-label goods cannot be linked to a producing firm, which is why they appear in this county- and state-level comparison but not in the firm-level analysis of Section 5. All three use the fully balanced product sample (present in all fourteen years, 2006–2019). Panel B uses the NielsenIQ Consumer Panel (Homescan, HMS) at the state level, the finest level at which its projection weights are valid, with year fixed effects. Here price, revenue, and physical quantity come directly from the projected panel aggregates, split into branded and private label the same way, with the pooled row a base-expenditure-weighted combination. In both panels immigration is instrumented by the shift-share prediction, winsorized at 1% over the full county sample, on stacked four-year differences, with standard errors clustered by state. For HMS, the county shock and its instrument are first aggregated to the state level using base-period expenditure weights.

Prices fall in every row. In RMS, real quantity for branded goods and for the pooled sample rises but is not significant (0.055 and 0.076), while the private-label quantity increase is large and

Table B.20: Immigration Lowers Prices in Every Sample

	Price
RMS (county, weighted) — product scope	
Branded (GS1)	-0.067***
	(0.012)
Private label	-0.055***
	(0.016)
Branded + priv. label	-0.067***
	(0.011)
RMS (county, weighted) — balance window (GS1)	
RW2	-0.056***
	(0.016)
RW6	-0.065***
	(0.014)
RW10	-0.067***
	(0.014)
RW14	-0.067***
	(0.012)
HMS (state, unweighted) — product scope	
Branded	-0.088***
	(0.030)
Private label	-0.099**
	(0.043)
Branded + priv. label	-0.087***
	(0.028)
HMS (state, unweighted) — retailer/product restriction	
Whole market	-0.079**
	(0.032)
Balanced products	-0.118***
	(0.030)
RMS-tracked channels	-0.090***
	(0.032)
RMS firm sample	-0.065*
	(0.037)
Big retailers	-0.082**
	(0.036)
RMS-tracked, large chains	-0.097***
	(0.032)

Notes: Each row is a separate IV regression of the price index on the immigration shock, instrumented by its shift-share prediction, winsorized at 1% over the full county sample, on 4-year stacked differences (2006–10, 2010–14, 2014–18) with state and year fixed effects; standard errors clustered by state in parentheses. RMS is the NielsenIQ Retail Scanner census (store-level), estimated at the county level and base-sales weighted (the paper's baseline weighting); HMS is the NielsenIQ Consumer Panel (Homescan household panel), aggregated to the state level, where its population projection factors are valid, and unweighted. Branded (GS1) products are manufacturer brands linked to firms via their GS1 prefix; private label are store brands; branded plus private label pools both. For the balance window, RW2-RW14 keeps products sold in at least k of the fourteen years of 2006–2019: RW14 is fully balanced (present all 14 years; the paper's baseline product sample, \$44B of sales) and RW2 the loosest (\$139B). The HMS retailer and product restrictions tighten from the whole consumer market toward the RMS scanner scope: whole market uses all products and retailers in the panel; balanced products keeps continuing products only; RMS-tracked channels restricts to the retail channels the retail scanner census covers; RMS firm sample to the GS1–NETS matched manufacturers; big retailers to the largest tracked chains (the top 88, about 80% of panel spend); and RMS-tracked, large chains is the intersection of the last two, closest to the RMS scanner scope. The RMS branded estimate (-0.067) matches the county baseline in Table 2. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

significant (0.198, with revenue also up 0.143), pointing to a reallocation of spending from branded toward private-label goods rather than an expansion of the overall market. HMS shows the same pattern with a smaller branded–private-label gap and, given its far smaller sample, noisier estimates that are individually insignificant.

Table B.20 collects the price effect across sixteen samples, grouped in four blocks that vary one construction choice at a time. The estimator is held fixed throughout: IV on the shift-share instrument, full-sample 1% winsorization, stacked four-year differences, standard errors clustered by state. The first block varies the RMS *product scope*, reusing the branded, private-label, and pooled price files from Table B.19, county level and sales-weighted. The second block varies the RMS *balance window* for branded goods: RWk keeps products sold in at least k of the fourteen years of 2006–2019, so RW14 is the fully balanced baseline (about $44B of sales) and RW2 the loosest (about $139B), with RW6 and RW10 in between. Widening the window trades product comparability for coverage. The third block repeats the HMS product-scope split (branded, private label, pooled) at the state level. The fourth block is an HMS *restriction ladder* that progressively narrows the consumer panel toward the RMS scanner universe: "whole market" uses all products and retailers in the panel; "balanced products" keeps only continuing products; "RMS-tracked channels" restricts to the retail channels the retail scanner data cover; RMS "firm sample" further restricts to the GS1–NETS matched manufacturers; "big retailers" keeps the largest tracked chains (the top 88, about 80% of panel spend); and RMS-tracked, "large chains" is the intersection of the last two, the closest HMS analogue to the RMS scanner scope.

The price coefficient is negative and significant in all sixteen samples, ranging from −0.055 to −0.118. The RMS branded estimate (−0.067) matches the county baseline in Table 2, and the HMS restriction ladder converges toward that RMS number as the panel is narrowed to the retail scanner's scope, so the two datasets agree where they overlap. That the decline also holds in the broad HMS samples, which include the discount, independent, and ethnic retailers the retail scanner census misses, shows the result is not an artifact of the retail scanner's retailer coverage.

B.4 Firm-Level Analysis

Table B.21: Economic Magnitude of the Demand and Supply Price Effects

	Demand	Supply
Effect of a one-standard-deviation increase	-2.800***	0.944***
	(0.778)	(0.287)
Effect of a p90–p10 move	-5.513	1.438
Raw coefficient (Table 7 units)	-0.073	0.017

Notes: Economic magnitudes based on the firm-level demand-and-supply price specification in Table 5, column (1). The first row reports the effect of a one-standard-deviation increase in each exposure (with the standard error from the corresponding standardized just-identified IV in parentheses); the second row the effect of moving a firm from the tenth to the ninetieth percentile of exposure; the last row the raw coefficient in the units of Table 5. Across-firm exposure standard deviations are 38.5 (demand) and 55.8 (supply), and tenth-to-ninetieth-percentile spreads are 75.8 and 85.0; all effects are in percentage points. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Shift-share robust inference. The firm-level demand and supply exposures are exposure-weighted aggregates $\mathrm{I}_f^{\text{demand}} = \sum_d \omega_{fd}^{\text{sales}} \mathrm{I}_d$ and $\mathrm{I}_f^{\text{supply}} = \sum_d \omega_{fd}^{\text{empl}} \mathrm{I}_d$ of the county immigration shock I_d. Following Borusyak et al. (2022), we compute exposure-robust ("shift-share") standard errors by residualizing the outcome and the endogenous variable on all controls and fixed effects, aggregating to the county level with the firm-county exposure shares, and running a county-level IV on the raw shock I_d. Table B.22 reports the result. The exposure-robust standard errors are smaller than or comparable to the SIC-3 clustered errors in every specification (ratios 0.42–0.87), and the point estimates reproduce the firm-level 2SLS exactly for the single-instrument specifications and to within 0.003 for the joint demand-and-supply specification—a just-identified one- versus two-instrument difference rather than an aggregation error. Consistent with the identification resting on historical ancestry shares rather than the contemporary shifters, industry clustering is the conservative choice.

Table B.22: Shift-Share Robust Standard Errors (BHJ) for the Firm-Level Specifications

	Coefficient	Industry-Clustered SE	BHJ Robust SE
Panel A. Single-instrument specifications			
Demand only (Col. 2)	-0.058	(0.022)	(0.009)
Supply only (Col. 3)	0.007	(0.005)	(0.004)
Panel B. Joint demand + supply (Col. 1)			
Demand	-0.073	(0.020)	(0.010)
Supply	0.017	(0.005)	(0.004)

Notes: The four Table 5 specifications (sales-weighted, SIC-3 and year fixed effects). The coefficient and the industry-clustered standard error are computed directly; the "BHJ Robust SE" is the exposure-robust standard error of Borusyak et al. (2022), obtained by residualizing the outcome and endogenous variable on all controls and fixed effects, aggregating to the county level with the firm-county exposure shares, and running county-level IV on the raw county shock. The BHJ coefficient reproduces the firm-level 2SLS estimate exactly in Panel A and to within 0.003 in Panel B. In every specification the BHJ standard error is no larger than the clustered one, so industry clustering is conservative.

Table B.23: UPC-to-Firm Aggregation: Sato–Vartia and Törnqvist vs. the Original Index

	(1) UPC→Firm SV	(2) UPC→Firm Tornqvist	(3) Original ($\Delta \ln \mathrm{SV}^{f,p}$)
$\text{Immigration}_{\text{f,t}}$ (D)	-0.0639***	-0.0606***	-0.0727***
	(0.0190)	(0.0190)	(0.0202)
$\text{Immigration}_{\text{f,t}}$ (S)	0.0132**	0.0135**	0.0169***
	(0.0052)	(0.0057)	(0.0051)
Obs.	11,058	11,058	11,058
First-Stage F	1,267.3	1,267.3	1,267.3

Notes: Firm-level regressions (N = 11,058). Columns (1)–(2) aggregate UPC-year prices to the firm via Sato–Vartia and Törnqvist chain indices over continuing UPCs; column (3) is the paper's original firm-level Sato–Vartia index. The three are statistically indistinguishable, so the firm-level baseline is robust to the UPC-to-firm aggregation. Firm-level demand (D) and supply (S) shift-share IV; SIC-3 and year fixed effects; sales-weighted; standard errors clustered by SIC-3 in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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

Table B.23 rebuilds the firm-level price change by a deliberately simpler and different route. Rather than nesting barcode prices through the county dimension, it collapses each barcode to a single price per year (pooled across stores) and aggregates those barcode prices *directly* to the firm using two off-the-shelf chain indices—Sato–Vartia and Törnqvist—over the firm's continuing barcodes, on the same firm sample as the baseline ($N = 11{,}058$). This discards the county nesting and the custom weighting of the baseline entirely, yet the demand-side coefficient is essentially unchanged: -0.064 (Sato–Vartia) and -0.061 (Törnqvist), against the baseline -0.073. The firm-level result is therefore not driven by how we construct the index.

Table B.24 instead estimates the price regression directly at the barcode level ($N = 121{,}368$ barcode-periods), with each barcode's price change measured across the markets in which it sells and the treatment still the firm's demand and supply exposure; the demand-side coefficient is -0.116 when barcodes are weighted by their firm's sales and -0.066 when weighted by their own revenue, bracketing the firm-level estimate. Because both regressions track prices of identical continuing barcodes, the decline cannot come from consumers reallocating toward cheaper products or from quality change, and instead reflects firms lowering markups on the same goods—the demand-side mechanism the paper emphasizes.

Table B.24: Within-UPC Prices: Chain-Weighted Across-County Aggregation

	Firm-Weighted		UPC-Weighted	
	(1)	(2)	(3)	(4)
	SV	Tornqvist	SV	Tornqvist
$\text{Immigration}_{f,t}$ (D)	-0.1164***	-0.1172***	-0.0657***	-0.0660***
	(0.0351)	(0.0350)	(0.0193)	(0.0193)
$\text{Immigration}_{f,t}$ (S)	0.0292***	0.0294***	0.0146***	0.0147***
	(0.0112)	(0.0111)	(0.0046)	(0.0046)
Obs.	121,368	121,368	121,368	121,368
First-Stage F	801.0	801.0	1,136.4	1,136.4

Notes: UPC×period observations ($N = 121{,}368$). UPC-year prices are built from UPC×county×year prices by chain-weighted aggregation across counties (Sato–Vartia and Törnqvist log-mean county sales-share weights), preserving within-county price variation; firm-weighted (1)–(2) and UPC-weighted (3)–(4). The coefficients bracket the firm-level baseline of -0.073, confirming that UPC-level identification requires county-preserving aggregation. Firm-level shift-share IV; SIC-3 and year fixed effects; clustered by SIC-3. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.25: Firm-Level Prices: Demand and Supply Exposure, Long-Difference

	Δ $\ln P^{SV,\text{Long-Difference}}$ (2006-2018)			
	(1)	(2)	(3)	(4)
$\text{Immigration}_{f,t}$ (D)	-0.056***	-0.052***		-0.048***
	(0.017)	(0.019)		(0.018)
$\text{Immigration}_{f,t}$ (S)	0.016**		0.001	0.019***
	(0.006)		(0.006)	(0.006)
Δ lnW (Retail, D)				-0.101
				(0.262)
Δ lnP (Rent, D)				-0.417
				(0.307)
Δ lnW (Retail, S)				0.274*
				(0.156)
Δ lnP (Rent, S)				-0.190
				(0.237)
Obs.	3,636	3,636	3,636	3,636
First-Stage F-stat	2,579.9	3,380.0	5,273.2	1,174.4
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, employment-to-population ratio, and urban indicator, each weighted by either initial sales share or employment share across counties within firm, and interacted with the year indicators. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Long-Difference Table B.25 re-estimates the firm-level demand-and-supply specification of Table 5 on a single 12-year long difference (2006–2018) rather than stacked four-year differences. Demand-side exposure again lowers firm prices (-0.056) and supply-side exposure modestly raises them ($+0.016$), with the same joint-identification pattern as the baseline: the supply effect collapses to zero when the demand channel is omitted (column 3). The long-difference estimates are slightly smaller than the stacked baseline but qualitatively identical, so the result does not hinge on the period-by-period timing of the stacked design.

Table B.26: Firm-Level Prices: Robustness Using International Controls and Sub-Samples

	(1) Baseline	(2) +Exp/Imp Controls	(3) +Intl Control	(4) Domestic Only	(5) Intl Interactions
$\text{Immigration}_{f,t}$ (D)	-0.062***	-0.065***	-0.064***	-0.071***	-0.071***
	(0.017)	(0.017)	(0.017)	(0.027)	(0.026)
$\text{Immigration}_{f,t}$ (S)	0.021***	0.022***	0.022***	0.022***	0.020***
	(0.005)	(0.004)	(0.004)	(0.006)	(0.005)
Obs.	11,058	11,058	11,058	6,828	11,058
First-Stage F	1911.3	2171.3	2103.3	1194.5	652.3
Fixed Effects	Industry, Time	Industry, Time	Industry, Time	Industry, Time	IndustryxIntl, TimexIntl

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

International Controls Table B.26 checks whether internationally active firms drive the result, since exporters and importers may be exposed to immigration through global rather than local channels. The demand-side coefficient stays between -0.062 and -0.071 (supply between $+0.020$ and $+0.022$) whether we control for exporter and importer status (columns 2–3), drop international firms from the sample altogether (column 4, $N = 6{,}828$), or interact both exposures with international status (column 5). The interaction terms (not reported) are statistically insignificant ($p \approx 0.80$), so the price response does not differ between international and domestic firms.

Retailer robustness Table B.27 tests whether the firm-level demand effect is an artifact of the retailer dimension along which prices are observed. The outcome is the firm-retailer-period Sato–Vartia price change ($N = 76{,}470$). We build the demand shock over three alternative *sets of counties*, in each case averaging the county immigration shock I_d with fixed base-year sales weights but changing the set of counties over which the average is taken: g_f uses the counties where the firm sells (all of its sales counties, as in the baseline); g_r uses the counties where the retail chain operates; and g_{fr} uses the counties where the particular firm sells *through the particular chain*. Each is instrumented by its shift-share counterpart. Column (1) is the firm-shock baseline ($g_f = -0.0703$). Column (2) adds the retailer shock: it is small and an order of magnitude below the firm shock ($g_r = -0.007$) while g_f barely moves (-0.064), so prices track the firm's sales-market exposure, not the retailer's. Column (3) replaces the period effects with retailer-by-period fixed effects, absorbing cross-retailer composition and anything else common to a chain within a period;

Table B.27: Firm, Retailer, and Firm-by-Retailer Immigration Shocks

	Δ $\ln P_{r,f,t}^{SV}$ (stacked)				
	(1)	(2)	(3)	(4)	(5)
$Immigration_{f,t}$ (D, firm)	-0.0703***	-0.0642***	-0.0585**	-0.0543**	-0.0596***
	(0.0229)	(0.0230)	(0.0230)	(0.0227)	(0.0223)
$Immigration_{r,t}$ (D, retailer)		-0.0073***			
		(0.0026)			
$Immigration_{f,r,t}$ (D, firm× ret)				-0.0257*	
				(0.0145)	
$Immigration_{f,t}$ (S)	0.0218***	0.0219***	0.0216***	0.0218***	0.0217***
	(0.0060)	(0.0060)	(0.0061)	(0.0061)	(0.0061)
Obs.	76,470	76,470	76,470	76,470	76,470
First-Stage F	727.9	491.6	746.2	775.2	694.1
Fixed Effects	Ind, Time	Ind, Time	Ind, RetxTime	Ind, RetxTime	Ind, RetxTime
Demand shocks	f	f, r	f	f, fxr	f
Retail cost controls	f	f	f	f	f, r

Notes: Each observation is a firm × retail chain × period, on stacked four-year differences (2006–2010, 2010–2014, 2014–2018); the outcome is the change in the firm-retailer Sato–Vartia log price index (in percentage points). All columns are IV. The endogenous demand-side exposures are built on the same shift-share logic as the firm shock, differing only in the set of counties over which the county immigration shock I_d is averaged with fixed base-year sales weights: g_f over the firm's sales counties (varying by firm × period), g_r over the retail chain's counties (retailer × period), and g_{fr} over the counties where the firm sells through that chain (firm × retailer × period); each is instrumented by its shift-share counterpart. The firm's supply-side exposure (S), built over its production counties, is included and instrumented in every column. Columns (1)–(2) include industry (3-digit SIC) and period fixed effects; columns (3)–(5) replace the period effects with retail-chain-by-period fixed effects, which absorb g_r (constant within chain-period) and identify g_{fr} off residual within-firm-retailer variation. The "Demand shocks" row indicates which exposures enter each column. Column (5) repeats column (3) with the demand-side retail wage and rent controls measured over the counties where the firm sells through that chain rather than over all of the firm's sales counties, as indicated in the "Retail cost controls" row. All columns use the same control vector as the firm-level baseline (Table 5): demand- and supply-side exposure to county urban share, 2006 GDP per capita, and the employment-to-population ratio, each interacted with period; the sales- and employment-share of counties with missing shocks on the demand and supply sides; and firm-level changes in retail wages and median rents on both sides. Columns including g_r or g_{fr} additionally control for the corresponding missing-county sales share. Observations are weighted by fixed 2006 firm-retailer sales; standard errors are clustered by firm. The sample contains 58 retail chains (parent companies). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

g_f is unchanged (-0.059), so the effect is within firm-retailer rather than a retailer-composition artifact. Column (4) addresses a subtler concern: a factor operating jointly at the firm-retailer level. Even granting that prices respond to the firm's exposure and not the retailer's, the effect could still run through specific firm-chain combinations—for instance, if a manufacturer's products at a given chain happen to be concentrated in immigrant-heavy counties, or if chains reallocate shelf space across competing manufacturers in response to local demand. The shock g_{fr} captures exactly this: it is the firm's immigration exposure measured over its counties *at that chain*, and it differs from the firm's overall exposure g_f because the set of counties in which a firm sells varies across the chains it supplies. Adding g_{fr} (with g_r now absorbed by the retailer-by-period effects) leaves the firm coefficient unchanged (-0.054), so no joint firm-retailer factor drives the result; g_{fr} itself is at most marginal (-0.026, $p < 0.10$) and in the same direction as the demand channel. Throughout, the supply-side coefficient is stable and positive ($+0.022$). Column (5) addresses retail costs. The baseline controls include changes in retail wages and rents averaged over all of the firm's sales counties; column (5) instead measures them over the counties where the firm sells through each chain, so that they match the retail costs that chain faces. The firm coefficient is unchanged (-0.060, against -0.059 in column 3). Bringing the retailer dimension in explicitly thus leaves the firm demand effect essentially unchanged.

Table B.28 examines store entry. Immigration may attract new stores, and incumbent stores may cut prices when they face more competition. We measure each firm's exposure to store entry as the change in the number of retail establishments per 100,000 residents from the BLS QCEW, averaged over the counties where the firm sells with the initial sales weights used for its demand-side immigration exposure. Column (1) regresses this exposure on the firm's demand- and supply-side immigration exposures, instrumented as in Table 5. The demand coefficient is small and not statistically significant. Columns (2) and (4) reproduce columns (1) and (5) of Table 5, and columns (3) and (5) add the store-entry exposure as a control. The demand coefficient is almost identical without the other controls and with them. The coefficient on store entry itself is positive, but because store entry may respond to the immigration inflow, we do not interpret it.

Table B.28: Store Entry and the Firm-Level Price Response

	Δ Retail Establ. (D)	Δ $\ln P^{SV,Stack}$ (2006-2018)			
	(1)	(2)	(3)	(4)	(5)
$Immigration_{f,t}$ (D)	0.005	-0.073***	-0.073***	-0.065***	-0.061***
	(0.010)	(0.020)	(0.020)	(0.018)	(0.018)
$Immigration_{f,t}$ (S)	0.003	0.017***	0.017***	0.021***	0.021***
	(0.002)	(0.005)	(0.005)	(0.005)	(0.005)
Δ Retail Establ. (D)			0.072**		0.094***
			(0.034)		(0.031)
Obs.	11,058	11,058	11,058	11,058	11,058
First-Stage F-stat	1,267.3	1,267.3	1,179.9	1,060.2	961.4
Fixed Effects	Industry, Time	Industry, Time	Industry, Time	Industry, Time	Industry, Time
Controls				Yes	Yes

Notes: Each observation is a firm × period, on stacked four-year differences (2006–2010, 2010–2014, 2014–2018). Δ Retail Establ. (D) is the change in the number of retail establishments per 100,000 residents (BLS QCEW), averaged over the counties where the firm sells with initial sales weights. Column (1) uses it as the outcome; columns (2)–(5) use the change in the firm's Sato-Vartia log price index. All columns are IV, with the demand-side (D) and supply-side (S) immigration exposures instrumented by their shift-share counterparts as in Table 5. Columns (2) and (4) reproduce columns (1) and (5) of Table 5. All columns include industry (3-digit SIC) and period fixed effects and control for the sales- and employment-share of counties with missing shocks; columns (4) and (5) add the controls of column (5) of Table 5. Observations are weighted by initial firm sales, and standard errors are clustered by 3-digit SIC industry. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.29: Firm-Level Price Decomposition: Long-Difference

	(1) CES price index	(2) Components of (1) Sato-Vartia	(3) Components of (1) Variety	(4) Appeal
Immigration_f (D)	-0.056***	-0.055***	-0.001	-0.096***
	(0.017)	(0.017)	(0.002)	(0.023)
Immigration_f (S)	0.016***	0.016***	0.000	0.041***
	(0.006)	(0.006)	(0.001)	(0.013)
Obs.	3,356	3,356	3,356	3,356
First-stage F-stat	2,564.2	2,564.2	2,564.2	2,564.2

Notes: This table is the long-difference (2006–2018) analog of Panel A of Table 8, estimated over a single 12-year change rather than stacked four-year differences. Column (1) is the CES price index, the sum of the Sato-Vartia price index of continuing varieties in column (2) and the variety term for product entry and exit in column (3) (equation 6.4), both in logs. Column (4) is firm appeal (equation 6.6). The variety and appeal terms use the state- and time-varying demand elasticities of Section 6. All columns report IV estimates, with demand-side (D) and supply-side (S) exposures instrumented jointly, and use the same firms, those for which all four outcomes can be constructed. All specifications include industry (3-digit SIC) fixed effects and control for the firm's exposure to the county urban indicator, GDP per capita, and employment-to-population ratio, and for the share of the firm's sales (demand side) and of its employment (supply side) in counties without immigration data. Observations are weighted by the firm's initial (2006) NielsenIQ sales. Immigration is measured in thousands. Standard errors are clustered by 3-digit SIC industry. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Long-Difference Decomposition Table B.29 re-estimates the structural price decomposition of Table 8 on the 12-year long difference. The demand-side pattern is unchanged: prices of continuing varieties fall, the variety correction is negligible, and firm appeal declines—so the demand channel operates through lower continuing-good prices and a shift in sales composition toward lower-appeal products, not through variety. Supply-side exposure again raises prices and appeal, consistent with the quality-upgrading interpretation in the main text.

Trademarks and Patents Supply-side exposure raises both prices and appeal in Table 8. Trademark and patent filings measure firms' investment in their brands and products directly. We match the GS1 company names of the firms in Table 8, and the NETS headquarters names linked to their prefixes, to applicants in the USPTO Trademark Case Files Dataset (Graham et al., 2013) and to assignees in the USPTO Patent Assignment Dataset (Graham et al., 2018). Names are matched exactly after removing legal suffixes, and for trademarks we also follow recorded changes of name. For patents, we count applications with an inventor-to-employer assignment, dated by filing. The outcome is 100 times the change in the inverse hyperbolic sine of the number of filings from years $t-7$ to $t-4$ to years $t-3$ to t, and the specification, weights, and clustering follow Table 8. We use the firms with at least one filing between 1995 and 2018, since a firm without a match may file under a name we do not match.

Supply-side exposure raises filings in all four columns of Table B.30. When immigrants arrive where a firm produces, the firm registers more brands and patents more inventions. This investment is consistent with the supply-side increases in appeal and in the prices of continuing products in Table 8,

since consumers value products sold under a stronger brand more highly.[11] Demand-side exposure lowers trademark filings, consistent with firms investing less in their brands as their customers shift toward cheaper products and store brands (Table 6), and has no significant effect on patents.

Table B.30: Immigration, Trademarks, and Patents

	(1)	(2)	(3)	(4)
	Trademark applications		Patent applications	
	All classes	Consumer goods	All	Granted
$\text{Immigration}_{f,t}$ (D)	-0.663**	-0.500*	-0.426	-0.191
	(0.268)	(0.259)	(0.406)	(0.582)
$\text{Immigration}_{f,t}$ (S)	0.252*	0.291**	0.362*	0.808***
	(0.138)	(0.142)	(0.189)	(0.184)
Obs.	8,042	8,042	2,478	2,478

Notes: IV estimates with the specification of Table 8 (instruments, controls, SIC 3-digit and period fixed effects, initial sales weights, standard errors clustered by SIC 3-digit industry). The sample is the Table 8 sample restricted to firms with at least one trademark (columns 1–2) or patent (columns 3–4) application in 1995–2018. Outcomes are 100× the inverse hyperbolic sine of the number of applications in years $t-3$ to t minus that in years $t-7$ to $t-4$. Consumer goods are Nice classes 3, 5, 16, 21, and 29–33. Column 4 counts applications that were later granted. ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

Appendix C Structural Estimation: Details

This appendix provides details on the estimation and derivation of the structural decomposition in Section 6.

C.1 Elasticity Estimation

We estimate demand elasticities following the methodology of Feenstra (1994) and Hottman et al. (2016), with each state as a market. 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
- $d \in s$: county (where $d \in s$ denotes that county d belongs to state s)

[11] Because many firm-periods have no filings, we also estimate a count model by Poisson pseudo-maximum likelihood, as Chen and Roth (2024) recommend, and the trademark results remain strong.

- $\tau \in \{1, 2, 3\}$: estimation period, corresponding to 2006–2010, 2011–2014, and 2015–2019
- $t \in \tau$: year within period τ
- f: firm
- u: product (UPC)

In the main text (Section 6) we use ℓ to index a product group, state, and period. Throughout this appendix we write $\ell \equiv (g, s, \tau)$ explicitly.

Our key structural parameters are the elasticities of substitution:

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

We aggregate the sales and quantities of each product across the counties of a state before estimation, so each state is one market in the equations below, and the state price of a product is its total sales divided by its total quantity.

Consider a firm f producing a set of products $\Omega_{fgs\tau} \subseteq \Omega_{gs\tau}$ in product group g, state s, and period τ. From the nested CES demand system in Section 6, the expenditure share of product u within firm f is:

$$S^U_{ugst} = \frac{(\phi^U_{ust})^{\sigma^U_{gs\tau}-1}(P^U_{ust})^{1-\sigma^U_{gs\tau}}}{\sum_{k\in\Omega_{fgs\tau}}(\phi^U_{kst})^{\sigma^U_{gs\tau}-1}(P^U_{kst})^{1-\sigma^U_{gs\tau}}} \tag{C.1}$$

where ϕ^U_{ust} is product appeal, P^U_{ust} is the price of product u in state s in 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:

$$\ln S^U_{ugst} = (\sigma^U_{gs\tau} - 1)\ln \phi^U_{ust} + (1 - \sigma^U_{gs\tau})\ln P^U_{ust} - \ln \Phi_{fgst} \tag{C.2}$$

where $\Phi_{fgst} \equiv \sum_{k\in\Omega_{fgs\tau}}(\phi^U_{kst})^{\sigma^U_{gs\tau}-1}(P^U_{kst})^{1-\sigma^U_{gs\tau}}$ is a firm-product group-state-year specific term.

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

$$P^U_{ust} = \mu^F_{fgst} \cdot MC_{ugst} \tag{C.3}$$

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

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

with a_{ugst} a product-specific cost shifter and $\delta_{gs\tau} \geq 0$ governing returns to scale. We maintain throughout that the markup is common across a firm's products, so that it cancels in the double-differencing procedure below. Online Appendix D.2 solves the firm's problem without that restriction

and bounds the dispersion it suppresses by the across-firm elasticity gap (Corollary 1); it is a maintained conduct assumption rather than an implication of arbitrary differentiated-product conduct.

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

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

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

$$\Delta_t \ln S^U_{ugst} = (\sigma^U_{gs\tau} - 1)\Delta_t \ln \phi^U_{ust} + (1 - \sigma^U_{gs\tau})\Delta_t \ln P^U_{ust} - \Delta_t \ln \Phi_{fgst} \tag{C.5}$$

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

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

This yields the double-differenced demand equation:

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

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

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

$$\ln P^U_{ust} = \ln \mu^F_{fgst} + \ln a_{ugst} + \delta_{gs\tau} \ln q_{ugst} \tag{C.8}$$

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

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

Using the accounting identity $S^U_{ugst} = P^U_{ust}q_{ugst}/E_{fgst}$, equivalently $\ln q_{ugst} = \ln S^U_{ugst} + \ln E_{fgst} - \ln P^U_{ust}$ (where E_{fgst} is firm expenditure, which cancels in double-differencing), and collecting the $\Delta_{u,t} \ln P^U_{ust}$ terms:

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

where $\kappa_{ugst} \equiv \frac{1}{1+\delta_{gs\tau}}(\Delta_t \ln a_{ugst} - \Delta_t \ln a_{\tilde{u}st})$ 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 ($\Delta_{u,t} \ln S^U_{ugst}$ and $\Delta_{u,t} \ln P^U_{ust}$) and two structural parameters ($\sigma^U_{gs\tau}$ and $\delta_{gs\tau}$). OLS estimation of either equation is biased due to simultaneity.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Double-Differenced Demand. The expenditure share of firm f within product group g in state s is:

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

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

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

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

Instrument Construction. OLS estimation of equation (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:

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

Following Hottman et al. (2016), we normalize product appeal so that its geometric mean over the firm's continuing products is constant:

$$\prod_{u\in\Omega_{fgs\tau}} \left(\frac{\phi^U_{ust}}{\phi^U_{us,t-1}}\right)^{1/N_{fgs\tau}} = 1, \tag{C.19}$$

where $N_{fgs\tau}$ is the number of continuing products of firm f. Under this normalization, changes in the firm price index decompose exactly into the change in the geometric mean of product prices (the Jevons index) and a within-firm share-dispersion term:

$$\Delta_{\tilde{f},t} \ln P^F_{fgst} = \frac{1}{N_{fgs\tau}} \sum_{u\in\Omega_{fgs\tau}} \Delta_{\tilde{f},t} \ln P^U_{ust} + \frac{1}{1-\sigma^U_{gs\tau}} \Delta_{\tilde{f},t} \ln \left(\sum_{u\in\Omega_{fgs\tau}} \frac{S^U_{ugst}}{\tilde{S}^U_{fgst}}\right) \tag{C.20}$$

where S^U_{ugst} is the share of product u in the firm's sales of continuing products and $\tilde{S}^U_{fgst}$ is the geometric mean of these shares. The decomposition follows from the CES demand for each product and has no remainder term.

The instrument is the double-differenced share dispersion term:

$$Z_{fgst} = \Delta_{\tilde{f},t}\left[\frac{1}{1-\sigma^U_{gs\tau}}\ln\left(\sum_{u\in\Omega_{fgs\tau}}\frac{S^U_{ugst}}{\tilde{S}^U_{fgst}}\right)\right] \tag{C.21}$$

Instrument Validity. *Relevance*: Share dispersion enters the firm price index by equation (C.20), so Z_{fgst} is mechanically correlated with $\Delta_{\tilde{f},t}\ln P^F_{fgst}$.

Exclusion (identifying assumption): we assume within-firm share dispersion is orthogonal to the firm-level appeal shock, $\mathbb{E}[Z_{fgst}\,\omega_{fgst} \mid g, s, \tau] = 0$. Economically, Z_{fgst} is driven by the cross-UPC distribution of prices and sales within the firm, whereas ω_{fgst} is a firm-level appeal innovation common to all of the firm's products; a firm-wide appeal shock shifts the level of within-firm shares but not their dispersion. This is an assumption on the data-generating process—which we assess in robustness checks—rather than an implication of defining appeal as a price residual. Because the instrument is built from the decomposition in equation (C.20), the assumption is stated under the normalization in equation (C.19).

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

$$\hat{\sigma}^F_{gs\tau}: \quad \Delta_{\tilde{f},t}\ln S^F_{fgst} = (1-\sigma^F_{gs\tau})\Delta_{\tilde{f},t}\ln P^F_{fgst} + \omega_{fgst}, \quad \text{instrumented by } Z_{fgst} \tag{C.22}$$

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

C.1.4 Implementation Details

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

- $\tau = 1$: 2006–2010
- $\tau = 2$: 2011–2014
- $\tau = 3$: 2015–2019

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. Aggregated to the state, these data still provide many products and firms in each state, product group, and period.

Sample Restrictions. Prices, market shares, and reference products are constructed from the state aggregates described above. We retain the state, product group, and period combinations with at least 30 non-reference products and 30 firm-year observations, a structural first-stage F-statistic above 10, and $1 < \hat{\sigma}^F_{gs\tau} < \hat{\sigma}^U_{gs\tau}$. This leaves 10,094 combinations.

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 6. For county d belonging to state s, we use the elasticities estimated for that state:

$$\text{For county } d \in s: \quad \sigma^U_{gd\tau} \equiv \hat{\sigma}^U_{gs\tau}, \quad \sigma^F_{gd\tau} \equiv \hat{\sigma}^F_{gs\tau} \tag{C.23}$$

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.

Table B.31: Estimated Elasticities and Marginal Cost Elasticity: Comparison with Hottman et al. (2016)

	p10	p25	Median	p75	p90
Panel A. This paper					
σ^U (products within firm)	8.4	11.0	15.6	22.1	30.9
σ^F (firms within product group)	2.1	3.0	4.3	5.8	7.5
δ (marginal cost elasticity)	0.04	0.06	0.11	0.19	0.28
Panel B. Hottman, Redding, and Weinstein (2016)					
σ^U (products within firm)	5.0	5.4	6.9	8.6	14.1
σ^F (firms within product group)	2.6	3.1	3.9	5.1	7.3
δ (marginal cost elasticity)	0.06	0.09	0.16	0.20	0.30

Notes: Panel A reports percentiles of our estimates across product groups, states, and periods (2006–2010, 2011–2014, 2015–2019). Panel B reproduces the distribution across product groups in Table V of Hottman et al. (2016), who estimate one set of parameters per product group for the country as a whole. σ^U is the elasticity of substitution across products within a firm, σ^F the elasticity of substitution across firms within a product group, and δ the elasticity of marginal cost with respect to output.

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 $\Omega_{fgd,t\cap t-1} \equiv \Omega_{fgdt} \cap \Omega_{fgd,t-1}$ (where $t, t-1 \in \tau$):

$$\Phi^{SV}_{fgdt} = \prod_{u \in \Omega_{fgd,t\cap t-1}} \left(\frac{P^U_{udt}}{P^U_{ud,t-1}} \right)^{w_{ugdt}} \tag{C.24}$$

where the Sato-Vartia weights are:

$$w_{ugdt} = \frac{L(S^{U*}_{ugdt}, S^{U*}_{ugd,t-1})}{\sum_{k \in \Omega_{fgd,t\cap t-1}} L(S^{U*}_{kgdt}, S^{U*}_{kgd,t-1})}, \quad L(a,b) = \frac{a-b}{\ln a - \ln b} \tag{C.25}$$

and S^{U*}_{ugdt} 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:

$$\prod_{u \in \Omega_{fgd,t \cap t-1}} \left(\frac{\phi^U_{udt}}{\phi^U_{ud,t-1}} \right)^{w_{ugdt}} = 1 \tag{C.26}$$

This normalization differs from the one used to estimate the elasticities, equation (C.19); the remark below explains why the two are consistent.

Remark on Normalization. The level of appeal is not identified. Scaling the appeal of every product in a nest by a common factor leaves expenditure shares unchanged, so prices and shares pin down only relative appeal, and each nest needs one normalization per period. A normalization changes neither the data nor the preferences. It determines how the change in the exact CES price index, which contains appeal, is split between a price index and an appeal term. We use two normalizations, that of Hottman et al. (2016) and the Sato-Vartia normalization, for three purposes: estimating the elasticities, constructing the price index, and measuring firm appeal.

First, the elasticities are estimated under the normalization of Hottman et al. (2016), equation (C.19). The estimates of σ^U and δ come from double differences relative to a reference product and do not depend on it. The parameter σ^F does not depend on it either, but its instrument, the share-dispersion term in equation (C.20), comes from the decomposition under this normalization, and so does the exclusion restriction. The normalization is used only to state this assumption and does not enter any reported index.

Second, the reported price index uses the Sato-Vartia normalization in equation (C.26), at the product level within firms and at the firm level within groups. Under this normalization the share-dispersion term is zero and the exact CES index is the Sato-Vartia index (Kurtzon, 2023). Feenstra and Reinsdorf (2007) show that the Sato-Vartia index is exact for CES preferences with changing tastes under this condition, and Martin (2022) shows that it measures the cost of living with tastes held at an intermediate point between the two periods. Martin (2020) finds that, when the CES specification is not exact, the Sato-Vartia index remains reliable while indices that attribute the CES residual to taste change do not. Estimating under one normalization and reporting under another is consistent: both describe the same preferences and data, written with a different normalization.

Third, firm appeal in equation (C.34) applies the normalization of Hottman et al. (2016) at the firm level. We compute the firm price P^F_{fgdt} as the geometric mean of the firm's product prices and normalize firm appeal so that its average change across firms in a county and product group is zero. Measured appeal therefore also reflects changes in the dispersion of sales across the firm's products. Because the normalizations differ only by a factor common to all firms in a county and product group, appeal under the Sato-Vartia normalization across firms is obtained by subtracting

the Sato-Vartia-weighted mean within each county and product group:

$$\Delta \ln \phi^{F,SV}_{fgdt} = \Delta \ln \phi^{F}_{fgdt} - \sum_{k \in \Omega_{gd,t \cap t-1}} w_{kgdt} \, \Delta \ln \phi^{F}_{kgdt}, \tag{C.27}$$

with w_{kgdt} the Sato-Vartia weights in equation (C.28). We report appeal separately and do not combine it with the price index.

Relation to Redding and Weinstein (2020). The CES unified price index of Redding and Weinstein (2020) holds the unweighted geometric mean of the tastes for common products constant, which adds a share-dispersion term to the Sato-Vartia index. The difference between the two indices comes from the choice of normalization, which the data cannot determine; Baqaee and Burstein (2023) contain a related discussion. Martin (2022) interprets the two indices as answers to different questions rather than as a biased and an unbiased measure. We use the Sato-Vartia index throughout, so the price index contains no taste term.

Aggregation to the County. Equation (C.24) is the first stage of the index in Section 3.3. The second stage aggregates firms within a product group and county, again over continuing firms and with Sato-Vartia weights:

$$\Phi^{SV}_{gdt} = \prod_{f \in \Omega_{gd,t \cap t-1}} \left(\Phi^{SV}_{fgdt}\right)^{w_{fgdt}}, \qquad w_{fgdt} = \frac{L(S^{F*}_{fgdt}, S^{F*}_{fgd,t-1})}{\sum_{j \in \Omega_{gd,t \cap t-1}} L(S^{F*}_{jgdt}, S^{F*}_{jgd,t-1})}, \tag{C.28}$$

where S^{F*}_{fgdt} is the expenditure share of firm f among continuing firms in group g and county d. The third stage aggregates product groups within the county with Törnqvist weights:

$$\Phi_{dt} = \prod_{g \in \Omega_{d,t \cap t-1}} \left(\Phi^{SV}_{gdt}\right)^{w_{gdt}}, \qquad w_{gdt} = \tfrac{1}{2}\left(s_{gdt} + s_{gd,t-1}\right), \tag{C.29}$$

where s_{gdt} is the share of group g in county expenditure. When these shares are constant, equation (C.29) is the price index of the Cobb-Douglas upper tier in equation (6.1). The county outcome in equation (3.1) is $\Delta \ln P_{dt} = \ln \Phi_{dt}$, which in logs is equation (3.3).

Chaining and Product Sample. Both the county index and the firm index below are computed for each pair of consecutive years and chained from 2006 to 2019, so the four-year outcome in each period is the sum of the annual log indices over the years in that period. The baseline keeps only products sold in all fourteen years of 2006–2019 in a county. Appendix Table B.20 shows that the results are similar when products sold in at least two, six, or ten of those years are included. The variety adjustment in Appendix C.3 is computed separately and uses all products, including those that enter or exit.

Firm-Level Index. The firm-level price index used in Section 5 starts from the same first stage, equation (C.24), and adds two stages with Törnqvist weights: product groups within a firm and county, and then counties within the firm,

$$\ln \Phi_{ft} = \sum_{d} \kappa_{fdt} \sum_{g} \kappa_{fgdt} \ln \Phi^{SV}_{fgdt}, \qquad \kappa_{fgdt} = \tfrac{1}{2}\left(s_{fgdt} + s_{fgd,t-1}\right), \quad \kappa_{fdt} = \tfrac{1}{2}\left(s_{fdt} + s_{fd,t-1}\right), \tag{C.30}$$

where s_{fgdt} is the share of group g in firm f's sales of continuing products in county d, and s_{fdt} is the share of county d in the firm's sales of continuing products. Each sum runs over the groups and counties in which the firm sells continuing products in both $t-1$ and t. The outcome in equation (5.1) is $\Delta \ln P_{ft} = \ln \Phi_{ft}$. Barcode prices are thus aggregated to the firm within each county first and only then across counties, so the index preserves within-county price variation and holds the firm's set of continuing products fixed. The county index and the firm index share the first stage and differ in the direction of the later aggregation: across firms and groups within a county for the former, across groups and counties within a firm for the latter. Appendix Table B.23 shows that aggregating barcode prices to the firm directly, with Sato-Vartia or Törnqvist chain weights, gives the same results.

C.3 Feenstra Variety Adjustment

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

$$\lambda_{fgdt} = \frac{\sum_{u \in \Omega_{fgd,t \cap t-1}} P^{U}_{udt} Q_{udt}}{\sum_{u \in \Omega_{fgdt}} P^{U}_{udt} Q_{udt}} \tag{C.31}$$

The variety adjustment is:

$$\Phi^{Variety}_{fgdt} = \left(\frac{\lambda_{fgdt}}{\lambda_{fgd,t-1}} \right)^{\frac{1}{\sigma^{U}_{gs\tau} - 1}} \tag{C.32}$$

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

When new products enter and capture market share, $\lambda_{fgdt} < 1$, so $\lambda_{fgdt} < \lambda_{fgd,t-1}$ (assuming exiting products had smaller shares than entering products). Since $\sigma^{U}_{gs\tau} > 1$, this implies $\Phi^{Variety}_{fgdt} < 1$, 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:

$$S^{F}_{fgdt} = \frac{(\phi^{F}_{fgdt})^{\sigma^{F}_{gs\tau}-1} (P^{F}_{fgdt})^{1-\sigma^{F}_{gs\tau}}}{\sum_{k \in \Omega_{gd\tau}} (\phi^{F}_{kgdt})^{\sigma^{F}_{gs\tau}-1} (P^{F}_{kgdt})^{1-\sigma^{F}_{gs\tau}}} \tag{C.33}$$

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

$$\phi^F_{fgdt} = \frac{P^F_{fgdt}}{P^G_{gdt}} \left(S^F_{fgdt}\right)^{\frac{1}{\sigma^F_{gs\tau}-1}} \tag{C.34}$$

The growth rate of appeal is:

$$\frac{\phi^F_{fgdt}}{\phi^F_{fgd,t-1}} = \frac{P^F_{fgdt}/P^F_{fgd,t-1}}{P^G_{gdt}/P^G_{gd,t-1}} \left(\frac{S^F_{fgdt}}{S^F_{fgd,t-1}}\right)^{\frac{1}{\sigma^F_{gs\tau}-1}} \tag{C.35}$$

We denote this appeal growth rate as $\Phi^{Appeal}_{fgdt} \equiv \phi^F_{fgdt}/\phi^F_{fgd,t-1}$. As in the main text, P^F_{fgdt} and S^F_{fgdt} are computed over the firm's continuing products, those sold in every year from 2006 to 2019. Intuitively, if a firm's market share increases more than its relative price decline would predict (given $\sigma^F_{gs\tau}$), we attribute the residual to an increase in appeal.

Appendix D Theoretical Framework

This appendix develops the demand system and the firm problem behind the decompositions in Section 6, and proves Propositions 1 and 2 of the main text, together with the supporting results. We work throughout with the nested CES structure used in estimation, so that the two elasticities we estimate separately do separate work: the across-firm elasticity σ^F disciplines markups and therefore the Sato-Vartia price index, while the within-firm elasticity σ^U governs the reallocation of expenditure across a firm's products and therefore the appeal index.

D.1 Environment

D.1.1 Consumers

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

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

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

σ_h^X. Immigrants are more price-elastic at both margins,

$$\sigma_i^U > \sigma_n^U > 1, \qquad \sigma_i^F > \sigma_n^F > 1. \tag{D.2}$$

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

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

with firm-level and aggregate price indices

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

and total demand $q_u = \sum_h L_h c_{uh}$. Each tier is weakly separable and homothetic, so the consumer's problem decomposes tier by tier: the lower tier delivers a firm-level unit cost P_{fh} that summarizes the firm's entire product line, and the upper tier treats the firm as a single good priced at P_{fh}. This is a different model from the system estimated in Section 6, and it reuses that section's notation for the objects that play the same role. It is static and describes one product group in one market, so the location and time indices have no counterpart here and C_{0h} collects the remaining groups in equation (6.1). Dropping those indices, p_u, P_{fh}, P_h, Ω_f, and $(\sigma_h^U, \sigma_h^F, \delta)$ take the place of P_{udt}^U, P_{fgdt}^F, P_{gdt}^G, Ω_{fgdt}, and $(\sigma_\ell^U, \sigma_\ell^F, \delta_\ell)$, with the consumer type h the only addition.

Two features of equation (D.1) matter for the results below. First, ϕ_u and ϕ_f are common to the two types, so the types agree on the ranking of products by appeal and differ only in the elasticity with which they substitute across prices. No result below rests on a difference in the taste for appeal. Second, demand is homothetic within a type: E_h enters equation (D.3) multiplicatively and appears in no elasticity, so income fixes what a consumer spends and nothing else. The gap $\sigma_i^X - \sigma_n^X$ for $X \in \{U, F\}$ is a separate primitive; the model does not derive it, and Section 7 motivates the gap.

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

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

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

which is the single-consumer nested CES demand of Section 6 with aggregate expenditure $\sum_h L_h E_h$. The within-firm and across-firm expenditure shares reduce to equations (C.1) and (C.33) after the change of variable of Appendix D.6. The model has content only through the difference between the two types, and the results below record where that difference enters.

The remark fixes what the model and the estimating system have in common and where they

diverge. They share the nesting and the role of the two elasticities; they differ in that the estimating system fits one representative consumer per product group, state, and period and takes prices as data, whereas the model has two types facing prices that firms choose. The taste weights ϕ_u and ϕ_f play the role of the appeal parameters ϕ^U_{udt} and ϕ^F_{fgdt} recovered there but are not the same objects, as Appendix D.6 makes precise. The exact link runs in the direction of the remark: the model collapses to the estimated system when the types coincide, not the other way around. What the model adds is an interpretation of the estimates. A single-agent elasticity fitted to a two-type market averages the two types, so it rises with the immigrant share even if neither type changes. This is one source of the rise in Table 7, and the model does not exclude others.

D.1.2 Firms

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

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

with product-specific cost shifter $a_u > 0$ and returns to scale governed by δ. Firm profit is $\pi_f = \sum_{u\in\Omega_f}[\,p_u q_u - C_u(q_u)\,]$. Within a firm we let cost differences across products run through appeal, $a_u = a(\phi_u)$ with $\varepsilon_a \equiv d\ln a/d\ln\phi_u \geq 0$, rather than through idiosyncratic cost draws. Appeal is fixed, so $a(\cdot)$ is not a technology for producing appeal.[12] Proposition 4 shows that a firm's products with higher appeal are more expensive when $\varepsilon_a + \delta > 0$, so either a cost that rises with appeal or decreasing returns alone suffices.

D.1.3 Market Structure

Each firm is large within its own nest and atomistic across nests. In choosing $\{p_u\}_{u\in\Omega_f}$ it internalizes the substitution among its own products through P_{fh}, but takes the aggregate index P_h and aggregate expenditure $E_h L_h$ as given because its share of the group is negligible. The comparative statics below hold P_h and $E_h L_h$ fixed for the same reason, so they are partial equilibrium.

D.2 Equilibrium Pricing

D.2.1 Demand Derivatives

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

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

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

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

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

the direct effect plus the substitution through the firm's own index. The firm faces the sum of the two types' demands, $q_u = \sum_h L_h c_{uh}$, so aggregating equation (D.8) with the quantity weights $\theta_{uh} \equiv L_h c_{uh}/q_u$, which equal type h's share of the product's revenue since both types pay p_u, gives

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

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

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

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

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

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

D.2.2 First-Order Condition and Markup

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

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

Proposition 1 (Pricing). *Let $m_v \equiv (p_v - MC_v)/p_v$ be the Lerner index of product v and $\bar{m}_h \equiv \sum_{v \in \Omega_f} s_{vh|f}\, m_v$ the average a type-h consumer pays at the firm. At an interior optimum, for every $u \in \Omega_f$,*

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

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

Proof. Let $L_h E_{fh} \equiv L_h \sum_{v \in \Omega_f} p_v c_{vh}$ be what type h spends at the firm, with E_{fh} as in equation (7.8), so that $s_{vh|f} = p_v c_{vh} / \sum_{w \in \Omega_f} p_w c_{wh}$ and

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

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

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

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

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

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

firm's revenue-weighted average.

Proposition 2 (Aggregate Pricing Rule)**.** *Weighting equation* (D.11) *by the revenue shares* $w_u \equiv p_u q_u / \sum_{v\in\Omega_f} p_v q_v$ *and summing over* $u \in \Omega_f$*, the within-firm elasticity cancels and*

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

with θ_{fh} *as in equation* (7.4)*.*

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

Equation (D.13) holds whatever the dispersion of markups across the firm's products. The main text restricts them to a single markup, so that $\bar{m}_n = \bar{m}_i = m$ and equation (D.13) gives $m = 1/\sigma^F_{f,\mathit{eff}}$ with $\sigma^F_{f,\mathit{eff}} \equiv \sum_h \theta_{fh}\sigma_h^F$, which is equation (7.3). With one consumer type, or with two that share the within-firm elasticity, this is not a restriction: equation (D.11) then delivers a common markup on its own. It is one here, since we keep $\sigma_i^U > \sigma_n^U$ for the composition results of Appendix D.4, and the unrestricted solution then lowers m_u on the products whose buyers are more immigrant, where $\sigma^F_{u,\mathit{eff}}$ is higher. Both equations (D.11) and (D.13) are the two-type case of the multiproduct pricing conditions in Faber and Fally (2022), derived there with a continuum of consumer groups.[13]

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

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

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

Proposition 3 (Existence and Uniqueness)**.** *Write* $\ell_h \equiv (1 + A_h\bar{m}_h)/\sigma_h^U$ *and* $\lambda_u \equiv \theta_{ui}\sigma_i^U/\sigma^U_{u,\mathit{eff}}$*. Then equation* (D.14) *becomes*

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

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

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

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

Under $\sigma_h^U \geq \sigma_h^F > 1$ the two equations in (D.16) have a unique solution, so, for given shares θ_{uh} and $s_{vh|f}$, equation (D.11) does too, and

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

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

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

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

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

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

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

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

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

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

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

Substituting $\ell_i = \ell_n + d$ into the two equations of (D.16) and subtracting them, then using

$\sigma_h^U - A_h = \sigma_h^F$, the within-firm elasticities drop out of the coefficient on ℓ_n and leave

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

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

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

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

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

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

With $d < 0$, and provided both types buy from the firm so that $0 < \bar{\lambda}_h < 1$, the first right side is below one and the second above it, so $\ell_n < 1/\sigma_n^F$ and $\ell_i > 1/\sigma_i^F$.

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

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

Corollary 2 (The Restriction Does Not Affect the Direction)**.** *At fixed prices, an increase in L_i lowers every m_u in the solution of equation (D.11), as it lowers the common markup $1/\sigma_{f,\mathit{eff}}^F$ of equation (7.3).*

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

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

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

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

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

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

D.2.3 Appeal and Price

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

Proposition 4 (Appeal-Price Relationship)**.** *Fix a firm f and let the markup be common across its products as in equation* (7.3)*. If $\varepsilon_a + \delta > 0$, then across the firm's products at any equilibrium*

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

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

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

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

Holding P_{fh} and P_h at their equilibrium values, total demand for a product of firm f is $q_u = \phi_u\, G(p_u)$ with $G(p) = \sum_h \Gamma_{fh}\, p^{-\sigma^U_h}$ and $\Gamma_{fh} \equiv L_h\, \phi_f\, P_{fh}^{\sigma^U_h}\, (P_{fh}/P_h)^{-\sigma^F_h} E_h/P_h > 0$ common to all $u \in \Omega_f$; note $G' < 0$. Combining equations (7.3) and (D.6) gives $p_u = \mu_f a(\phi_u) q_u^{\delta}$ with $\mu_f \equiv \sigma^F_{f,eff}/(\sigma^F_{f,eff} - 1)$ common across the firm's products. Substituting $q_u = \phi_u G(p_u)$ and collecting the terms in p_u,

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

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

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

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

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

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

D.3 Price Decomposition

This subsection proves Proposition 1, which isolates the two direct forces on price, and then gives the equilibrium response, first in general and then in the special case in which the two forces separate cleanly. All of it starts from one input. By equation (7.3) the firm's markup is set by $\sigma^F_{f,\mathit{eff}}$, so what immigration does to price begins with what it does to that elasticity, and the first result gives the two channels through which immigration raises $\sigma^F_{f,\mathit{eff}}$.

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

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

and at the firm level

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

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

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

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

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

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

The second term of equation (D.23) is therefore $L_i^{-1} \sum_u w_u(\theta_{ui} - \theta_{fi})\sigma^F_{u,\mathit{eff}}$, a weighted sum of the deviations $\theta_{ui} - \theta_{fi}$. Those deviations have weighted sum zero, since $\sum_u w_u \theta_{ui} = \theta_{fi}$ and

$\sum_u w_u = 1$, so a term that does not vary across products contributes nothing to the sum. This removes σ_n^F from $\sigma_{u,eff}^F = \sigma_n^F + \theta_{ui}(\sigma_i^F - \sigma_n^F)$, and then θ_{fi} from $\theta_{ui} = (\theta_{ui} - \theta_{fi}) + \theta_{fi}$:

$$\sum_u w_u(\theta_{ui} - \theta_{fi})\sigma_{u,eff}^F = (\sigma_i^F - \sigma_n^F)\sum_u w_u(\theta_{ui} - \theta_{fi})\theta_{ui} = (\sigma_i^F - \sigma_n^F)\,\mathrm{Var}_w(\theta_{ui}).$$

□

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

Proof of Proposition 1. Take logs of the pricing rule (7.3) to get $\ln p_u = \ln \mu_f(\sigma_{f,eff}^F) + \ln a_u + \delta \ln q_u$ and differentiate in L_i, holding the price fixed wherever it appears on the right. The markup term then responds through Lemma 2 and the scale term through the quantity:

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

with $d\ln\mu_f/d\sigma_{f,eff}^F = -1/[\sigma_{f,eff}^F(\sigma_{f,eff}^F - 1)] < 0$, $\partial\sigma_{f,eff}^F/\partial L_i > 0$ by equation (D.23), and $\partial q_u/\partial L_i = c_{ui}$ with $c_{ui}/q_u = \theta_{ui}/L_i$, so the scale term is $\delta\,\theta_{ui}/L_i$. Aggregating across the firm's products with the log-mean expenditure weights $\bar{s}_u$ gives the display of the proposition.[14] □

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

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

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

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

turn $\sum_u w_u \, d\ln p_u/dL_i$ into $\sum_h \theta_{fh}\zeta_h$. For the second, letting prices respond in equation (D.8) gives

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

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

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

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

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

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

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

The restriction does three things. By equation (D.7) the two types then hold the same within-firm shares, $s_{ui|f} = s_{un|f}$, so every product has the same buyer mix, $\theta_{uh} = \theta_{fh}$; every product therefore takes the same markup $\mu_f(\sigma_{f,\mathit{eff}}^F)$, by equation (D.15); and the two types' price indices respond alike, $\zeta_n = \zeta_i$, since ζ_h weights the same price changes by the same shares. Both of the endogenous objects in equation (D.25) are thereby resolved.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

The two feedbacks therefore pull against each other in Λ_f^{-1}, and Λ_f is positive wherever the pricing rule (7.3) describes a profit maximum.[15] The feedback then changes the size of the response but not its sign, which the bracket alone settles. Dividing it by $\theta_{fi}/L_i>0$, the response is negative exactly when

$$\big(\sigma^F_i-\sigma^F_n\big)\,\theta_{fn}\;>\;\delta\,\sigma^F_{f,\mathit{eff}}\big(\sigma^F_{f,\mathit{eff}}-1\big). \tag{D.29}$$

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

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

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

D.4 Appeal Composition

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

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

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

where E_{fh} is the expenditure of a type-h consumer on firm f. This share is identically the revenue weight w_u of Proposition 2, written in the form that shows where L_i enters. Both $s_{uh|f}$ and $s_{u|f}$ sum to one over Ω_f, so $\bar{\phi}_f$ is a within-firm appeal average that isolates composition from firm size.

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

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

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

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

Proof of Proposition 2. This is the composition channel, holding equilibrium prices fixed, and all sums are over $u \in \Omega_f$. Differentiating $\bar{\phi}_f = \sum_u s_{u|f} \phi_u$ with prices and appeal fixed, and writing

$D_f \equiv L_n E_{fn} + L_i E_{fi}$,

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

using $\partial s_{u|f}/\partial L_i = (E_{fi}/D_f)(s_{ui|f} - s_{u|f})$. Since $\bar{\phi}_f = \theta_{fn}\bar{\phi}_{fn} + \theta_{fi}\bar{\phi}_{fi}$ with $\theta_{fi} = L_i E_{fi}/D_f$ the immigrant share of the firm's sales, as in equation (7.4), we have $\bar{\phi}_{fi} - \bar{\phi}_f = \theta_{fn}(\bar{\phi}_{fi} - \bar{\phi}_{fn})$, and

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

and we show this sum is negative. Write $d_u \equiv s_{ui|f} - s_{un|f}$ for the gap between the two types' spending shares on product u, and $r_u \equiv s_{ui|f}/s_{un|f}$ for their ratio, so that $d_u = s_{un|f}(r_u - 1)$: product u takes a larger share of immigrant spending than of native spending exactly when $r_u > 1$. Order the firm's products by ϕ_u. By Proposition 4 this is also the ordering by price, and Lemma 3 then makes r_u decreasing in ϕ_u. The first step is the sign pattern of d_u. Both share vectors sum to one, so $\sum_u d_u = 0$, which in terms of the ratio says $\sum_u s_{un|f}\, r_u = 1$: the average of r_u under the native shares is exactly one. A decreasing sequence with average one starts above one and ends below it. There is therefore an appeal level ϕ^* that separates the two regimes: products with $\phi_u < \phi^*$ have $r_u \geq 1$ and so $d_u \geq 0$, and products with $\phi_u > \phi^*$ have $r_u \leq 1$ and so $d_u \leq 0$. Relative to natives, immigrants overspend on the firm's low-appeal products and underspend on its high-appeal ones. The second step turns this pattern into the sign of the sum. Because $\sum_u d_u = 0$, subtracting the constant ϕ^* from every appeal level changes nothing,

$$\sum_u d_u\, \phi_u = \sum_u d_u\left(\phi_u - \phi^*\right), \quad \text{(D.34)}$$

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

$$\frac{\partial \bar{\phi}_f}{\partial L_i} = \frac{E_{fi}}{D_f}\,\theta_{fn}\left(\bar{\phi}_{fi} - \bar{\phi}_{fn}\right) < 0. \quad \text{(D.35)}$$

□

This is a fixed-price result. Letting prices respond adds two components. The component of the price response that is common across the firm's products leaves every within-firm share unchanged, by Lemma 1; it still shifts the type weight θ_{fi}, and since $\sigma_i^F > \sigma_n^F$ a common price decline raises θ_{fi}, which reinforces the fall in $\bar{\phi}_f$. The non-uniform component, which shifts the firm's relative prices, is the one the model does not sign, and this is why we state the result at fixed prices. The composition channel is distinct from the Sato-Vartia price channel of Proposition 1, and it operates through a different primitive. Specifically, the tilt is governed by $\sigma_i^U - \sigma_n^U$ whereas markup compression is

governed by $\sigma_i^F - \sigma_n^F$, so Φ^{SV} and Φ^{Appeal} contain separate information about the mechanism.

D.5 Variety Effects

We now let firms choose which products to offer: product u is offered if and only if operating profit covers its fixed cost F_u. Under the common markup, operating profit is proportional to revenue $R_u = p_u q_u$, with factor $(1 + \delta\sigma_{f,eff}^F)/[\sigma_{f,eff}^F(1+\delta)]$, which is $1/\sigma_{f,eff}^F$ at $\delta = 0$, and immigration moves it in two directions: markup compression lowers it, by equation (D.23), while demand expansion raises it. The scale force here is demand-side and operates even at $\delta = 0$, unlike the cost-side scale effect of Proposition 1. The model does not sign the net effect, which is consistent with the finding that the Feenstra variety correction does not move with demand-side immigration exposure (Table 8). We also hold the appeal of each product fixed: a barcode is a fixed physical product, so appeal attaches to the product rather than adjusting within it, and an improvement in appeal is the entry of a new product, the margin treated in this subsection.

D.6 How Appeal Enters Preferences

Two conventions for appeal coexist and it matters which one we adopt. The estimation of Section 6, following Hottman et al. (2016), is appeal-augmenting: a unit of product u delivers ϕ_u^U efficiency units, and the within-firm expenditure share is $s_{u|f} \propto (\phi_u^U)^{\sigma^U - 1} p_u^{1-\sigma^U}$, as in equation (C.1). The model of this appendix instead enters appeal as a taste weight in equation (D.1), so that the share is $s_{u|f} \propto \phi_u \, p_u^{1-\sigma^U}$. With a single consumer type the two are the same system written in different variables,

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

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

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

We adopt the taste-weight convention for two reasons. It is the one that implements the discipline of the model: the two types agree on appeal and differ only in elasticities, so every composition result is forced through the elasticity gap alone, whereas under the alternative a taste difference re-enters through the exponent. And it rests Proposition 2 on the weaker condition: with appeal out of the cross-type comparison, the result needs only the price-appeal ranking of Proposition 4, whereas under the alternative a high-appeal product can attract immigrants rather than repel them, and the result would survive only if price rises more than one for one with appeal, $d\ln p_u / d\ln \phi_u^U > 1$. The cost is that the model and the estimating system are written in different

variables, with equation (D.36) the bridge between them; with a single type the bridge is exact, as Remark 1 records. The ordering itself is shared: because equation (D.36) is increasing, the two conventions rank a firm's products identically.

A separate caution applies to both conventions. Appeal is recovered as the component of a firm's market share that its relative price does not explain, so whether that residual reflects quality, in the sense of something consumers would pay for at given prices, or a shift in tastes is not identified. For that reason we keep appeal out of the reported price index, as we state in Section 6, and we read the appeal index as the composition of a firm's sales rather than as a welfare-relevant quantity.

Appendix E Calibrating the Price Response

Proposition 1 signs the price response only under a restriction on parameters and does not determine its magnitude. We evaluate both at our estimates, taking the estimated rise in the elasticity as given and asking whether, passed through the pricing rule (7.3), it is large enough to account for the measured price decline. This rise need not come only from the changing composition of buyers, since any force that raises the elasticity firms face lowers their markups through the same rule.

Let s denote immigration relative to the initial population, the immigrant inflow over a period divided by the market's population in 2005.[16] Taking logs of the pricing rule (7.3) and differentiating with respect to s,

$$\frac{d\ln P_f^{SV}}{ds} = \underbrace{\frac{1}{1+\delta\,\sigma_{f,\mathit{eff}}^F}}_{\textit{Feedback}\ (<1)} \left[\underbrace{-\frac{1}{\sigma_{f,\mathit{eff}}^F - 1}\frac{d\ln\sigma_{f,\mathit{eff}}^F}{ds}}_{\textit{Markup Effect}\ (<0)} + \underbrace{\delta\,\theta_{fi}\frac{d\ln L_i}{ds}}_{\textit{Scale Effect}\ (\geq 0)} \right]. \tag{E.1}$$

The bracket holds the two direct forces. The markup effect is the compression of the markup as the elasticity the firm faces rises. The scale effect is the cost of serving more demand: demand for the firm's products shifts out by $\theta_{fi}\,d\ln L_i/ds$, since immigrants buy a share θ_{fi} of its sales and their population rises by $d\ln L_i/ds$, and marginal cost rises with quantity at rate δ. The factor in front is a feedback that dampens both: the price fall itself raises the quantity the firm sells, and with it marginal cost, which offsets part of the fall.

Evaluating equation (E.1) requires the across-firm elasticity $\sigma_{f,\mathit{eff}}^F$, the returns to scale δ, the elasticity response $d\ln\sigma_{f,\mathit{eff}}^F/ds$, the immigrant share of the firm's sales θ_{fi}, and the proportional rise in the immigrant population $d\ln L_i/ds$ (Table E.1).

[16] The main regressions measure immigration by the number of arrivals, following Terry et al. (2026). We calibrate in shares for two reasons. First, the pricing rule depends on immigration only through the immigrant share of a firm's buyers, so equation (E.1) is a derivative with respect to s. Second, the elasticity response is estimated across states and the price response across counties, so a coefficient on the number of arrivals refers to markets of very different size in the two regressions, while a coefficient on the share does not. We therefore re-estimate both responses in shares; column (2) of Appendix Table B.11 reports the county-level price regression in this unit. We evaluate equation (E.1) per percentage point of s and convert the implied response to a count in the last step, so we can compare it with column (1) of Table 5.

Table E.1: Calibration of the Price Response

	Value	Description	Source
Panel A. Inputs			
$\sigma^F_{f,eff}$	4.31	Across-firm elasticity	Median estimate (Appendix C.1)
δ	0.106	Returns to scale	Median estimate (Appendix C.1)
$d\ln\sigma^F_{f,eff}/ds$	7.1%	Elasticity response	Table 7 column (3), in s; in counts, the published 0.021
θ_{fi}	0.098	Immigrant share of the firm's sales	Consumer panel
$d\ln L_i/ds$	8.96	Rise in the immigrant population	Consumer panel, 1/0.112
Pop	1.47m	Market population, for the count conversion	Sales-weighted mean over counties
Panel B. Price response			
Model, per pp of s	−1.410	Equation (E.1) at Panel A	
Model, per thousand immigrants	−0.096	Converted at *Pop*	
Firm demand side	−0.073	Measured, demand channel	Table 5 column (1)

The elasticity and the returns to scale are estimated jointly, product group by state by period, in Appendix C.1, one pair for each of the 9,141 in column (3) of Table 7. We take each at its median, 4.31 and 0.106. Hottman et al. (2016) report medians of 3.9 for the across-firm elasticity and 0.16 for the returns to scale. Our across-firm elasticity is a little higher, as expected when it is estimated within region rather than nationally. Our returns to scale are comparable to theirs and consistent with the flat marginal cost in Stroebel and Vavra (2019).

The elasticity response is measured against the number of immigrants in column (3) of Table 7, which needs to be converted per unit of s to map into equation (E.1). We re-estimate that specification in those units and obtain 7.1 percent per percentage point.[17] The scale effect requires the product $\theta_{fi}\, d\ln L_i/ds$, the shift in demand for the firm's products. The first term, θ_{fi}, is the share of the firm's sales bought by immigrants. Averaged across firms with sales as weights, it equals the immigrant share of CPG spending. We measure this share in the consumer panel, weighting each household by its immigrant probability from Appendix A.2 (equation (A.2)). Immigrant households account for 9.8 percent of CPG spending, so $\theta_{fi} = 0.098$. The second term, $d\ln L_i/ds$, follows from the definition of s. Since s is the inflow divided by the initial population N, the immigrant population rises by $dL_i = N \times ds$, so

$$\frac{d\ln L_i}{ds} = \frac{1}{L_i}\frac{dL_i}{ds} = \frac{N}{L_i}.$$

With the same weights, immigrant households account for 11.2 percent of household members, so $L_i/N = 0.112$ and $d\ln L_i/ds = 1/0.112 = 8.96$. The product, $0.098/0.112 = 0.88$, is CPG spending per member of an immigrant household relative to the average, pooling 2006, 2010, 2014, and 2018;

[17] The sample, weights and fixed effects are those of column (3) of Table 7. We replace the number of immigrants with immigration relative to the initial population and divide the shift-share instrument by the same population. Estimated against the number of immigrants, it reproduces the published coefficient of 0.021.

it lies between 0.87 and 0.91 in each year.[18] An inflow equal to one percent of the initial population therefore shifts demand for the firm's products by 0.88 percent, since each new immigrant spends 88 percent of what the average person spends.

Substituting these values into equation (E.1) gives

$$\frac{d \ln P_f^{SV}}{ds} = \frac{1}{1 + 0.106 \times 4.31} \left[-\frac{1}{4.31 - 1} \times 7.1 + 0.106 \times 0.098 \times 8.96 \right]$$
$$= \frac{1}{1.46} [-2.15 + 0.09] = -1.41 \text{ percentage points}$$

per percentage point of s.[19] The markup effect lowers the price by 2.15 percentage points and the scale effect raises it by 0.09, and the feedback, which divides their sum by 1.46, offsets about a third of the fall. Per thousand immigrants this is -0.096, the units in which Table 5 reports the demand-side coefficient of -0.073.[20]

The calibration asks whether the measured rise in the elasticity, passed through the pricing rule, is large enough to account for the measured price decline. We emphasize it is a consistency check instead of a structural estimate, where the mapping requires slight deviations from the model itself. First, the model has a single market, while the price response is identified across counties and the elasticities across states. Second, it divides consumers by the household probability of Appendix A.2, which refers to households with a foreign-born head of any origin and counts all of their members, while the regressions measure non-European immigration. Third, it has a single elasticity, while the estimates vary across product groups, states, and periods (evaluating the model at each pair of estimates and weighting by sales, instead of at the medians, leaves the model close to the regression). Finally, as mentioned above, the rise in the effective demand elasticity (7.1%) need not come only from the changing composition of buyers since the estimate is a reduced form result (Table 7). Any force that raises the elasticity firms face lowers their markups. Nevertheless, at the estimated elasticity and returns to scale, the measured rise in the elasticity implies a price decline close to the one we measure.

[18] This is spending per household member, not per household. Immigrant households spend slightly more on consumer goods than the rest, but they are larger (3.11 versus 2.49 members), so spending per member is lower. The ratio of 0.88 is computed over all immigrant households, while the scale effect depends on the spending of new arrivals. Among the 2008 TUM households, the 62 whose most recent foreign-born adult arrived within the previous four years spend 0.92 of the average per member, close to 0.88.

[19] Computing each side weighted by initial sales instead, as in the paper's other specifications, leaves the two close: the model gives -3.95 and the regression -3.75 (standard error 2.33). The model is then evaluated at each pair of estimates and averaged using sales as weights; the regression is already weighted that way.

[20] The conversion divides by a market population of 1.47 million, the sales-weighted mean across the counties in which firms sell. Markets differ in population by orders of magnitude, so this is one representative value. Expressed instead in immigration relative to population, where the two regressions and the model share one unit, the model is -1.41 and the same specification re-estimated in that unit is -1.03.

References

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

Baqaee, David R. and Ariel Burstein, "Welfare and Output with Income Effects and Taste Shocks," *Quarterly Journal of Economics*, 2023, *138* (2), 769–834.

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

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

Chen, Jiafeng and Jonathan Roth, "Logs with Zeros? Some Problems and Solutions," *Quarterly Journal of Economics*, 2024, *139* (2), 891–936.

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.

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 Marshall B. Reinsdorf, "Should Exact Index Numbers Have Standard Errors? Theory and Application to Asian Growth," in Ernst R. Berndt and Charles R. Hulten, eds., *Hard-to-Measure Goods and Services: Essays in Honor of Zvi Griliches*, University of Chicago Press, 2007, pp. 483–513.

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

Graham, Stuart J. H., Alan C. Marco, and Amanda F. Myers, "Patent Transactions in the Marketplace: Lessons from the USPTO Patent Assignment Dataset," *Journal of Economics & Management Strategy*, 2018, *27* (3), 343–371.

_ , Galen Hancock, Alan C. Marco, and Amanda F. Myers, "The USPTO Trademark Case Files Dataset: Descriptions, Lessons, and Insights," *Journal of Economics & Management Strategy*, 2013, *22* (4), 669–705.

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.

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.

Ke, Guolin, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu, "LightGBM: A Highly Efficient Gradient Boosting Decision Tree," in "Advances in Neural Information Processing Systems," Vol. 30 2017, pp. 3146–3154.

Kurtzon, Gregory, "The Problem with Normalizing Preferences that Change in a Cost-of-Living Index," BLS Working Paper 534, U.S. Bureau of Labor Statistics 2023.

Lach, Saul, "Immigration and Prices," *Journal of Political Economy*, August 2007, *115* (4), 548–587.

Martin, Robert S., "Changing Tastes Versus Specification Error in Cost-of-Living Measurement," BLS Working Paper 531, U.S. Bureau of Labor Statistics 2020.

_ , "Revisiting Taste Change in Cost-of-Living Measurement," *Journal of Economic and Social Measurement*, 2022, *46* (2), 109–147.

Redding, Stephen J and David E Weinstein, "Measuring Aggregate Price Indices with Taste Shocks: Theory and Evidence for CES Preferences*," *The Quarterly Journal of Economics*, February 2020, *135* (1), 503–560.

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

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," *American Economic Review*, 2026, *116* (3), 828–861.

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