The Economic Architecture of Retail Price Systems

Focal Prices, Complete Transaction-Gradient Paths, and Category-Dependent Price Sensitivity

Iván Franco | Triplethree International | Mexico

Research manuscript | October 2026

Abstract

What can a retail price distribution tell us when transaction quantities are unavailable? We study prices for the same UPC during the same week, rather than pooling brands or package sizes. Each distribution has a modal price, a complete sequence of frequency slopes between neighboring prices, and potentially distinct regular-price and promotional regions. We develop an exact accounting decomposition linking price-frequency gradients to realized unit-share gradients and test when price-only information approximates transaction activity. Using five Dominick's Finer Foods categories, we construct 128,817 UPC-week systems and 596,713 consecutive-price gradients. On a stricter sample of 66,567 off-focal comparisons, frequency and transaction-gradient signs agree in 86.9–91.3% of comparisons for four categories, but only 59.4% for shampoo. Placebo reallocations of per-price sales intensity show that most sign agreement is mechanical; the incremental advantage of the observed price-sales association is small and category-dependent. Price-frequency shape also contains ordinal information about independently estimated price–quantity sensitivity for cereals and shampoo, but not consistently elsewhere. Promotions are disproportionately observed below a historical price anchor. The evidence supports a measurable price-system architecture, not a universal demand elasticity, optimal pricing rule, or causal proof of retailers' motives.

Keywords: price-frequency gradient; focal prices; transaction shares; retail price dispersion; promotions; category-specific sensitivity.

JEL: C81, D40, D43, L11, L81

1. The question and the contribution

A price distribution for one identical product, observed in one period, reveals more than a mean or a range. It shows where prices concentrate, how frequencies change between neighboring prices, and whether lower- and higher-price regions differ in promotional and sales outcomes. These patterns can inform economic analysis even when unit quantities are missing. The challenge is to establish exactly what the price distribution reveals—and what it does not.

We investigate five propositions. First, modal prices form reproducible focal references. Second, the shape around a focal price can help classify relative price sensitivity. Third, the sequence of adjacent price-frequency slopes contains information about the allocation of transactions across prices. Fourth, promotional and non-promotional prices exhibit distinct volume and margin patterns around an established reference. Fifth, price-only analysis is informative under testable conditions, rather than a universal substitute for sales data.

Our contribution is the integrated measurement and validation of these five propositions at the UPC-week level. We construct complete local gradients from prices alone, compare them with gradients based on realized unit shares, and quantify the performance of category-specific sensitivity signals on held-out UPCs. Because the observed price distribution cannot by itself identify a retailer's objectives, economic interpretation is distinguished from causal claims throughout the analysis.

The existence of frequent or reference prices is not novel: Eichenbaum, Jaimovich, and Rebelo (2011) analyze persistent reference prices. A related but distinct methodological antecedent is Antoniades, Feenstra, and Xu (2022), who approximate expenditure shares from seller distributions. Our narrower empirical question concerns consecutive-price gradients for an identical UPC-week and the category-specific information contained in focal decay and regime asymmetry. See also Varian (1980), Kaplan and Menzio (2015), and Campbell and Eden (2014).

2. Five testable economic propositions

Table 1. Five economic propositions and empirical evidence.

Proposition	Economic content	Evidence / status
P1. Focal organization	Observed price modes recur across observations and time.	Descriptive regularity; not novel by itself.
P2. Shape and sensitivity	Steeper frequency decay near the mode can identify more price-sensitive UPCs.	Validated in cereals and shampoo only.
P3. Gradient path	All neighboring frequency slopes trace relative changes in price mass.	Fully computed; sales correspondence varies.
P4. Commercial regimes	Below-anchor discounts and above-anchor prices show different outcomes.	Promotional/volume/margin associations, not motives.
P5. Price-only inference	Frequency data contain limited information about unseen unit distribution.	Useful in selected contexts; not quantity recovery.

3. Data, observation unit, and validation rules

The central dataset is the Dominick's Finer Foods store-level movement panel provided by the Kilts Center for Marketing, University of Chicago Booth School of Business. Five original category CSV archives are used: cereals, laundry detergents, paper towels, shampoo and canned tuna. Each valid observation is a store–UPC–week record. The empirical system is one UPC in one week: all its valid store-level observed prices are grouped into exact unit-price points. This avoids conflating different brands and package sizes within a category.

We set unit price to PRICE/QTY, rounded to cents, and read MOVE as reported units sold. The principal sample retains OK=1, MOVE>0, PRICE>0, QTY>0. For the complete-gradient test, an eligible system has at least 10 valid price records, at least three distinct price points, and one uniquely most frequent price. These eligibility rules yield 128,817 systems across the five categories. Earlier modal-split and regression tests use different eligibility rules, and their denominators must not be treated as identical.

Crucially, the movement data do not give a valid posted price when MOVE=0: those records show PRICE=0. We therefore observe frequencies conditional on positive sales, not the universe of shelf offers. This selection limits the extension of our results to datasets containing posted prices regardless of transactions. The Kilts Center describes the movement files and requires source acknowledgment in publications (Kilts Center for Marketing, n.d.).

We distinguish three forms of evaluation: a within-product-week comparison of observed price-frequency and unit-share distributions; a held-out UPC validation of sensitivity classifiers; and a placebo that reassigns average sales intensity across the price points while preserving the observed frequencies. We cluster uncertainty by UPC when multiple weeks or price intervals from the same product enter a calculation.

4. Proposition 1 — Focal prices are reproducible but partly mechanical

For any UPC-week, the focal price pF is the observed price with the highest record count. A mode is guaranteed for any finite nonempty distribution; what matters empirically is its stability, reproducibility, and concentration. In an independent partition of stores, focal prices agree in 92.82% of 130,614 eligible

UPC-week comparisons. Restricting to systems with at least three distinct prices in both halves yields 90.07% agreement. When the mode has less than half of the observations in each half, agreement falls to 57.92%.

Table 2. Split-store agreement in modal prices for identical UPC-week systems.

Split-store test	Eligible UPC-week systems	Exact mode agreement
Unique mode in both halves	130,614	92.82%
At least three prices per half	89,064	90.07%
Each half's mode below 50%	18,255	57.92%

The sharp decline in the last row is important: strong modal concentration itself increases split-sample agreement, even without any strategic signal. A separate weekly persistence check finds an 86.79% match with last week's mode in 113,249 eligible comparisons. As supporting external evidence, a Profeco Mexico analysis of identical products finds much higher focal agreement within a chain than across chains; the within- and between-chain sets have different filters, so they cannot be interpreted as a causal competition contrast.

Economic interpretation. The focal price is an empirically reproducible reference around which observations concentrate. It is consistent with an operational pricing anchor; its existence alone does not establish an intended market signal or a profit-maximizing decision.

5. Proposition 2 — Does the distribution reveal relative price sensitivity?

A central hypothesis from the research program initiated in 2014 is that a narrowly organized price distribution may be associated with greater price sensitivity: prices deviate less from the focal reference when sensitivity is high. The appropriate test is not simply whether the mode is tall, but how frequency declines as distance from that mode increases. We therefore compare focal decay with modal concentration and conventional price dispersion.

Let $d(p)=|\ln(p/pF)|$ be the log distance from the unique focal price, and $N(p)$ the number of records at price p. We define the decay statistic as the fitted rate at which log frequency declines with distance from the focal point:

$$\kappa N = - \sum_{p \neq pF} d(p) \cdot \ln[N(p)/N(pF)] \,/\, \sum_{p \neq pF} d(p)^2 \qquad (1)$$

The reference for sensitivity comes from separate within-UPC price–quantity regressions on regular-price observations with store and week fixed effects. These coefficients are observational, not causally identified elasticities. We test the classification on UPCs held out of calibration. A high-sensitivity UPC belongs to the highest training-defined group of absolute price–quantity coefficients; a high price-only signal belongs to the high focal-decay group. The hit-rate baseline is the prevalence of high-sensitivity UPCs in the held-out category.

Table 3. Held-out UPC classification using the focal-decay statistic.

Category	Held-out UPCs	Focal-decay hit rate	Category baseline	Lift
Cereals	283	33.3%	20.5%	1.63×
Shampoo	368	40.5%	20.4%	1.99×
Laundry detergents	357	18.7%	20.2%	0.93×
Canned tuna	132	23.1%	22.0%	1.05×
Paper towels	64	7.1%	20.3%	0.35×

This supports a category-specific ordinal signal in cereals and shampoo. It does not support a universal inverse relation between dispersion and sensitivity. A second held-out evaluation uses area under the receiver-operating curve (AUC) over all products rather than the hit rate of the selected top group. In cereals a model combining several geometric features reaches an AUC of 0.684 versus 0.590 for simple dispersion (paired bootstrap difference +0.094, 95% interval +0.015 to +0.178). In shampoo, simple

dispersion (AUC 0.731) performs better than focal decay alone (0.672). Thus the more elaborate measure is not uniformly the best.

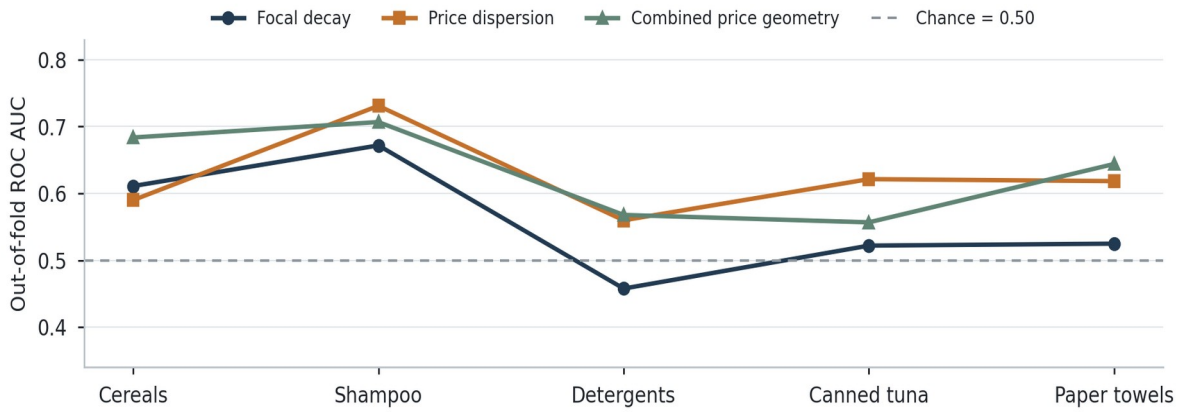


Figure 1. Out-of-fold identification of high-sensitivity UPCs: ROC AUC for three price-only measures. The dashed line indicates chance-level ranking (AUC = 0.50).

Economic interpretation. The shape of prices can help rank products by relative sensitivity, but the informative feature varies by category. Focal decay performs well for identifying the upper-sensitivity group in cereals and shampoo; broader dispersion sometimes offers a better overall ranking. Universal shape-to-elasticity rules are not supported.

6. Proposition 3 — The complete transaction-gradient path

Sort the k distinct prices of the same UPC-week as $p1 < p2 < \ldots < pk$. Let $N(p)$ be the number of valid observed store-weeks at price p and $U(p)$ the total units sold there. Their shares are $sN(p)=N(p)/N$ and $sU(p)=U(p)/U$. For each consecutive price pair we calculate the two level slopes:

$$G_{N,j} = [sN(p_{j+1}) - sN(p_j)] / [p_{j+1} - p_j] ; \quad G_{U,j} = [sU(p_{j+1}) - sU(p_j)] / [p_{j+1} - p_j] \qquad (2)$$

The full path $\Gamma N=(G_{N,1},\ldots,G_{N,k-1})$ runs from the minimum price through the focal price to the maximum. The two endpoints of each slope are prices, never two stores. No monotonicity or bell shape is imposed: the path can increase again to the right of the modal price. An analogous ΓU uses unit shares.

A useful accounting relationship explains why these objects can be close. If $\bar{q}(p)=U(p)/N(p)$ is the average number of units sold per valid price record and $\bar{q}=U/N$ is the overall average, then $sU(p)=sN(p)\bar{q}(p)/\bar{q}$. For strictly positive masses, the log-price slopes obey an exact decomposition:

$$\Psi U(p_1,p_2) = \Psi N(p_1,p_2) + \Psi \bar{q}(p_1,p_2) \qquad (3)$$

The additive identity provides an empirical bridge, not an input requirement for price-only applications. When average units per observation are approximately stable between prices, the frequency-based and transaction-based log slopes coincide. When sales intensity differs sharply across price points, the difference is measurable and may be economically substantial. This identity does not make the observed gradients causal demand responses.

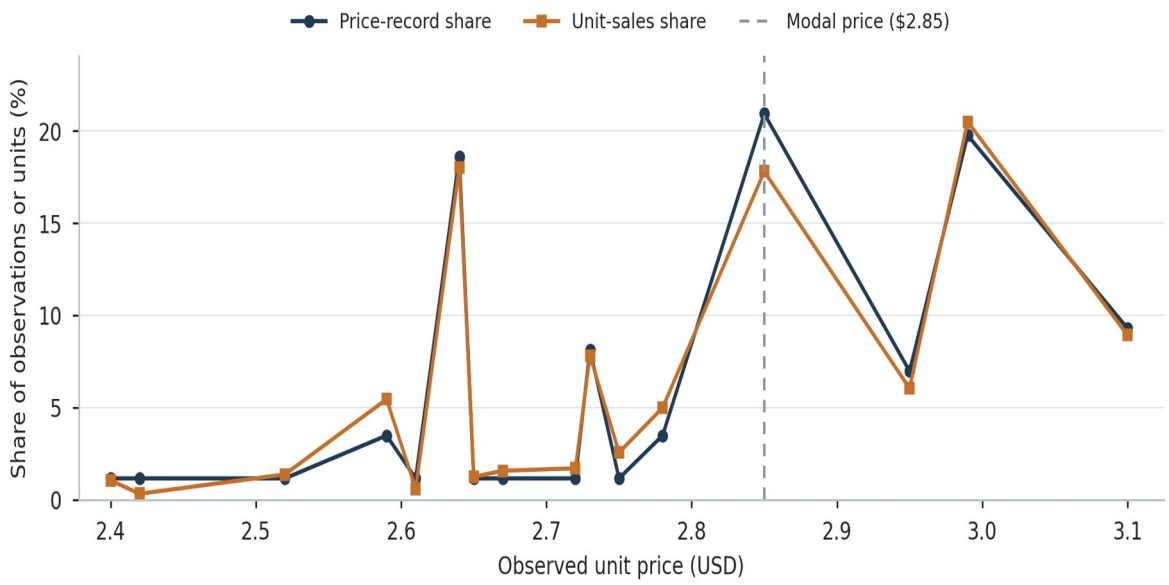


Figure 2. Observed price-record shares and unit-sales shares for the same cereal UPC in a single week (Dominick's scanner data). The vertical line marks the modal price.

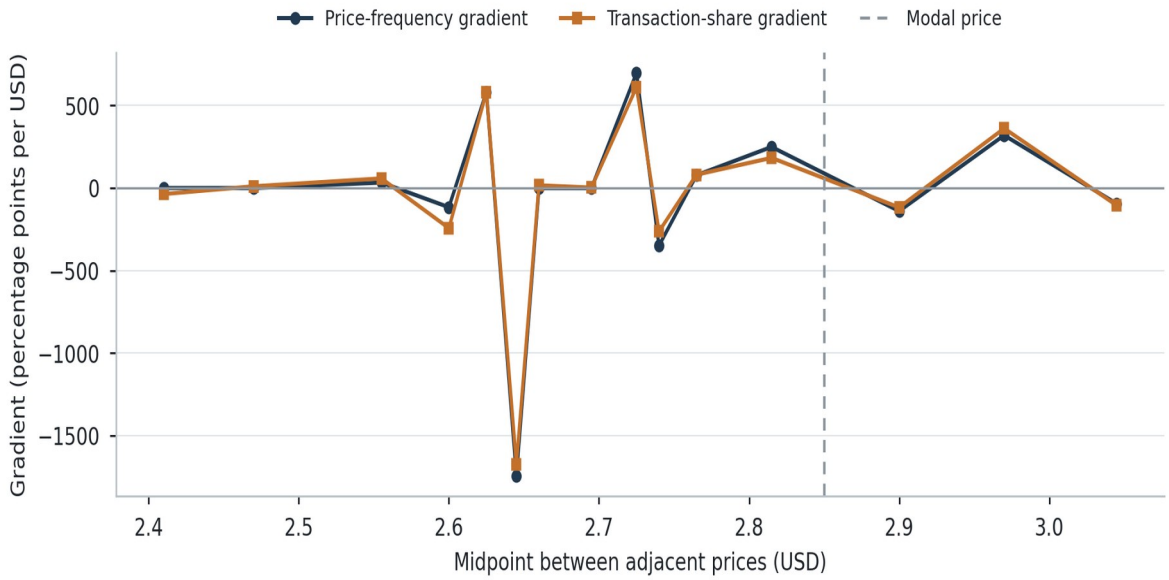


Figure 3. Complete adjacent-price gradient paths for the UPC-week in Figure 2. Each point is the slope between two consecutive prices; no store-to-store comparison is implied.

We calculate 596,713 neighboring-price pairs in 128,817 systems. A stricter evaluation retains price pairs with at least three records at each endpoint and a gap of at least 1% and excludes intervals touching the focal price. This yields 66,567 pairs. The table evaluates exact agreement of the three possible signs (negative, zero, positive); it does not imply that unit shares or quantities are accurately recovered.

Table 4. Adjacent-price gradient validation on the strict evaluation sample.

Category	Strict intervals	Exact sign match	Majority-sign baseline	Median TV distance
Cereals	28,638	91.3%	59.0%	5.15%
Laundry detergents	13,506	87.8%	63.9%	5.83%
Paper towels	5,387	86.9%	57.7%	5.67%
Shampoo	5,501	59.4%	53.4%	8.19%
Canned tuna	13,535	87.5%	54.2%	7.27%

Shampoo is an important exception. When intervals with a zero frequency or zero unit-share gradient are excluded, directional agreement rises to 75.9% across the remaining 4,203 strict comparisons. That

conditional number cannot replace the 59.4% full-sample result: zeros mean the indicator frequently does not supply a usable directional forecast.

More demanding null tests show why even large percentages should not be interpreted as new predictive information. Within each UPC-week, we reassign the existing per-price average sales intensities randomly across price points while keeping the prices and their observed record counts fixed. This preserves the arithmetic connection between high frequency and high sales opportunities without preserving their empirical pairing. The real minus shuffled sign-match advantage is +1.4 percentage points in cereals, +0.6 in detergents, +5.4 in paper towels, −2.7 in shampoo and +3.2 in tuna (40 exploratory reshuffles per category). These are descriptive placebo comparisons rather than a final causal randomization test.

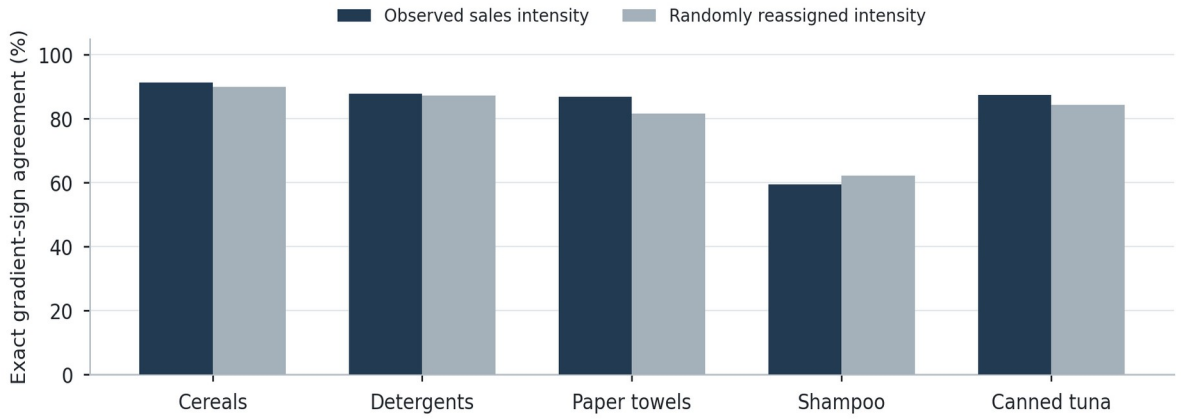


Figure 4. Exact gradient-sign agreement in observed data versus random reassignment of average sales intensity across price points within each UPC-week. The placebo preserves observed price frequencies.

Economic interpretation. Adjacent price-frequency gradients trace the observed distribution of price records, and frequently point in the same direction as adjacent transaction-share gradients. However, randomized sales-intensity reallocations generate similarly high agreement in several categories. The incremental correspondence is therefore category-dependent; the frequency gradient cannot be treated as a universal predictor of transaction quantities.

7. Proposition 4 — Prices around the focal anchor have different commercial roles

Retailers may pursue several commercial objectives within one pricing system. Discounts can stimulate purchases, accelerate stock turnover or attract store traffic, while higher prices can produce different margin outcomes. The data record prices, transactions and promotion indicators—not the retailer's intentions—so we use these mechanisms to interpret observed patterns rather than to assign an identified motive to any specific price.

Using an independently calculated historical regular-price anchor rather than the current-week mode, the five-category Dominick's analysis finds a promotion flag in 31.34% of eligible observations below the anchor, compared with 1.17% at the anchor and 2.79% above it. The lower-price side is therefore far more likely to appear in a promotional regime. The historical anchor is distinct from the contemporaneous UPC-week mode and should never be conflated with it.

Table 5. Promotional observations relative to a historical regular-price anchor.

Position vs. historical anchor	Store–UPC–week observations	Promotion flags	Flag rate
Below	2,060,619	645,819	31.34%
At anchor	4,154,740	48,414	1.17%
Above	2,368,381	65,972	2.79%

Among matched non-promotional UPC-week comparisons, lower-price observations often show higher average units per record, whereas higher-price observations often show higher gross profit per record.

These patterns range by category and are not within-store price experiments. A different analysis reports that median price–transaction slopes over regular-price points are positive, while the promotional subsets have negative medians in all five categories. Promotions change both the price regime and the allocation of observed quantities (Gupta, 1988; Hendel and Nevo, 2006). Experimental research shows that broader Hi-Lo and EDLP policies can yield different profit outcomes (Hoch, Drèze, and Purk, 1994).

Economic interpretation. Observed price dispersion is consistent with distinct commercial regimes around an established anchor. Prices below that anchor are disproportionately promotional, while other comparisons reveal differences in units sold and gross profit per record. This pattern helps connect distributional shape with retail pricing practice without assuming that inventory clearance or margin extraction has been causally established.

8. Proposition 5 — What prices can and cannot reveal without sales

The price-only use case is limited but real. One can calculate the mode, relative frequency shares, dispersion, focal decay and the complete sequence of frequency slopes without ever observing quantities. In categories calibrated against independent scanner data, these statistics can help classify relative sensitivity or approximate the allocation of sold units across already observed price points. They cannot reconstruct the number of sales at each price unless per-price sales intensity is sufficiently stable or externally calibrated.

An independent-store test from earlier replication work asks how much additional information sales, margins and promotions provide beyond the price-frequency pattern of a separate store group. On held-out UPCs, the distributional total-variation error falls from 10.55% to 10.39% for A→B and from 11.05% to 10.85% for B→A when economic covariates are added. Those modest gains do not identify economic motives, and this model conditions on a price support already observed in the target group. It predicts mass over given price points, not which new prices will appear.

Coverage also matters. When many UPC–store–weeks lack a valid positive-sale price, the empirical distributions describe a selected set of selling stores. For cereals, exact sign agreement is 59.5% on a small low-coverage group (<40%, n=190 strict pairs) versus 92.2% in a high-coverage group (≥80%, n=26,755 strict pairs). These are distinct systems; the difference is not a treatment effect. In shampoo, high coverage still leaves weak correspondence. A posted-price dataset independent of realized sales would be the strongest external test of the price-only promise.

9. Discussion — What is the economic architecture?

The evidence reveals an economic structure in the distribution of prices for identical UPCs. Focal prices recur, while the full distribution can rise and fall on both sides of the mode. Promotions are concentrated below independently estimated regular-price anchors. Gradients describe how frequency and realized transaction mass change between consecutive price points; the relationship is usually informative but is not equally accurate across categories. The shape of the distribution also helps identify relatively price-sensitive products in specific categories.

The economic architecture is thus an observable organization of prices and market activity—not an assertion that every retailer optimizes the same objective. Volume, margin, price positioning, product availability, and the timing of promotions are plausible drivers of that organization. Establishing the causal contribution of each driver remains a separate empirical task.

10. Conclusion

A price distribution contains economically useful information beyond its mean. For each identical UPC and week, it reveals a focal reference, a local shape and a complete path of frequency changes between neighboring prices. Comparing those gradients with transaction data shows a substantial but partly mechanical correspondence. Price-only signals must therefore be assessed against appropriate null benchmarks and filtered for data coverage. Out-of-sample tests also show that price geometry can identify relatively sensitive UPCs in cereals and shampoo, although the result does not generalize uniformly to other categories.

The practical contribution is a conditional measurement framework. Analysts can use focal prices to locate concentration, adjacent gradients to map local changes in price frequency, and promotion markers to distinguish commercial regimes. Where category-level calibration supports it, those features can help rank relative sensitivity and describe how observed transactions are distributed across prices. They do not, on their own, recover demand elasticity, unseen transaction totals or retailers' underlying objectives.

Acknowledgment and data availability

The Dominick's Finer Foods scanner dataset is provided by the Kilts Center for Marketing at the University of Chicago Booth School of Business. We acknowledge the Kilts Center as the source. All processing, tests, errors and interpretations are the author's responsibility. The derived UPC-week gradient tables, scripts and data dictionaries can accompany the manuscript; the original third-party movement files should be obtained from the provider or supplied under the data provider's terms.

References

Antoniades, A., Feenstra, R. C., & Xu, M. (2022). Using the retail distribution of sellers to impute expenditure shares. American Economic Review, 112(7), 2213–2236. https://doi.org/10.1257/aer.20190896

Campbell, J. R., & Eden, B. (2014). Rigid prices: Evidence from U.S. scanner data. International Economic Review, 55(2), 423–442. https://doi.org/10.1111/iere.12055

Eichenbaum, M., Jaimovich, N., & Rebelo, S. (2011). Reference prices, costs, and nominal rigidities. American Economic Review, 101(1), 234–262. https://doi.org/10.1257/aer.101.1.234

Franco, I. (2026). The hidden transaction structure of price distributions: Recovering realized market activity from prices alone. Working paper, Triplethree International. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=7466081

Gupta, S. (1988). Impact of sales promotions on when, what, and how much to buy. Journal of Marketing Research, 25(4), 342–355. https://doi.org/10.1177/002224378802500402

Hendel, I., & Nevo, A. (2006). Measuring the implications of sales and consumer inventory behavior. Econometrica, 74(6), 1637–1673. https://doi.org/10.1111/j.1468-0262.2006.00721.x

Hoch, S. J., Drèze, X., & Purk, M. E. (1994). EDLP, Hi-Lo, and margin arithmetic. Journal of Marketing, 58(4), 16–27. https://doi.org/10.1177/002224299405800402

Kaplan, G., & Menzio, G. (2015). The morphology of price dispersion. International Economic Review, 56(4), 1165–1206. https://doi.org/10.1111/iere.12134

Kehoe, P. J., & Midrigan, V. (2015). Prices are sticky after all. Journal of Monetary Economics, 75, 35–53. https://doi.org/10.1016/j.jmoneco.2014.12.004

Kilts Center for Marketing. (n.d.). Dominick's Finer Foods data. University of Chicago Booth School of Business. https://www.chicagobooth.edu/research/kilts/research-data/dominicks

Varian, H. R. (1980). A model of sales. American Economic Review, 70(4), 651–659.

Appendix A. Definitions and identification

A.1. Exact identities and slope definitions

For each UPC-week with observed positive-sale price records $N(p)>0$ and recorded units $U(p)>0$, write $N=\Sigma_p N(p)$, $U=\Sigma_p U(p)$, and $\bar{q}(p)=U(p)/N(p)$. Then $sU(p)=sN(p)\bar{q}(p)/\bar{q}$ is a counting identity. For $p2>p1>0$, define $\Psi X=[\ln sX(p2)-\ln sX(p1)]/[\ln p2-\ln p1]$. It follows exactly that $\Psi U=\Psi N+\Psi\bar{q}$. This is a decomposition of simultaneously observed price-point masses. It is not a counterfactual response or a supply or demand elasticity.

A.2. Full-path empirical sample

All distinct observed positive unit prices for the same UPC-week are ordered. The 596,713 adjacent pairs cover the full price domain. The strict subset requires $N(pj)\geq 3$ and $N(pj+1)\geq 3$, $pj+1/pj\geq 1.01$, and neither endpoint equal to the modal price. Sign concordance treats zero as a distinct outcome. Bootstrap confidence intervals in the replication package resample UPC identifiers, thereby retaining all eligible weeks for a sampled product.

A.3. Sales-intensity null and limitations

For each system the null shuffles $\bar{q}(p)$ across price points while preserving $N(p)$ and the multiset of intensities. Synthetic mass is $U^*(p)=N(p)\bar{q}^*(p)$, renormalized to shares. The published exploratory null uses 40 shuffles per category. Because pair differences share UPC-weeks and the placebo is conditional on observed intensities, strong inferential claims require a fully nested grouped-permutation design, preferably with more draws and a second dataset.

A.4. Held-out sensitivity test

The sensitivity benchmark is the magnitude of a regular-price within-UPC log-price/log-MOVE coefficient with store and week fixed effects, estimated on products satisfying minimum sample requirements. Its coefficient can be endogenous. UPC-level folds reserve complete UPCs, preventing cross-week sharing of the same product between training and evaluation. Thresholds are set using training UPCs; reported hit rates are the share of high-benchmark-sensitivity UPCs among those flagged by high focal decay. AUC evaluates the entire score ranking rather than only one selected upper group. Category differences prevent a universal threshold.

Appendix B. Evidence classification for referees

Table A1. Summary of evidentiary scope and remaining identification limits.

Statement	Evidence strength	What is not inferred
Observed focal modes persist and reproduce.	Strong descriptive, partly construction-driven.	Not necessarily strategic signals.
Steep focal decay classifies relative sensitivity.	Out-of-sample in cereals and shampoo.	Not a universal or causal elasticity.
Price gradient signs align with unit-share signs.	Full-path scanner validation, heterogeneous.	Not exact quantities or sales prediction.
Promotions disproportionately occur below an anchor.	Strong descriptive association.	Not proof of inventory liquidation motive.
Prices alone sometimes convey economic information.	Calibrated, coverage-dependent application.	Not complete sales reconstruction.