Article

Not peer-reviewed version

Diminishing Returns to Promotion Depth in Grocery: Evidence from Dominick's Cereal (1989–1997)

Milad Javadi *

Posted Date: 25 March 2026

doi: 10.20944/preprints202603.1966.v1

Keywords: retail promotions; scanner data; cereal; price cuts; stockpiling; high-dimensional fixed effects; event study

Article

Diminishing Returns to Promotion Depth in Grocery: Evidence from Dominick's Cereal (1989–1997)

Milad Javadi

Department of Economics, The University of Texas at Dallas, Richardson, TX 75080, USA; milad.javadi@utdallas.edu

Abstract

This study examines how the depth of temporary price cuts is related to weekly unit sales in the ready-to-eat cereal category of the Dominick's Finer Foods scanner data for 1989–1997. Promotion depth is modeled as a dose–response treatment in a store $\times$ UPC $\times$ week panel. The baseline specification regresses $\ln(1 + \text{MOVE})$ on discount-depth bins while absorbing store $\times$ UPC, store $\times$ week, and UPC $\times$ week fixed effects, with two-way clustered standard errors by store and UPC. The cleaned and trimmed panel contains 4.64 million observations, and the preferred estimating sample uses the top 100 UPCs ranked by cumulative sales. Deeper discounts are consistently associated with larger same-week lift. Relative to the 0–5% bin, the preferred top-100 estimates imply approximately 4.7% higher weekly sales for 5–10% discounts, 10.4% for 10–20% discounts, and 28.0% for discounts above 20%. Additional episode evidence indicates that most promotions are short, although a long right tail motivates treating very long runs as price regimes. A corrected event study around promotion starts shows a large contemporaneous spike but a significantly positive near lead, so those dynamics are best read as descriptive. By contrast, a promotion-end event study documents a persistent post-promotion dip of roughly 2–5% for at least eight weeks after promotions end, consistent with inventory drawdown. The evidence therefore supports strong contemporaneous lift and dynamic payback, although the coarse-bin static design does not, by itself, establish concavity in the depth–sales relationship.

Keywords: retail promotions; scanner data; cereal; price cuts; stockpiling; high-dimensional fixed effects; event study

---

1. Introduction

Temporary price reductions are central to grocery retailing, but the relationship between promotion depth and sales is more nuanced than the simple claim that deeper discounts always raise demand. In storable categories such as ready-to-eat cereal, a sale may increase purchases because consumers are price sensitive in the current week, because they stockpile against future high prices, or because retailers time promotions when demand is already strengthening. These distinctions matter for retailers deciding how aggressively to discount, for manufacturers allocating trade-promotion support, and for researchers interpreting scanner-based evidence on consumer demand.

This paper studies a focused question: within Dominick's Chicago-area stores, how does the depth of a temporary price cut affect weekly unit sales for ready-to-eat cereal, and do deeper discounts exhibit diminishing marginal returns? The project also evaluates dynamic adjustment around promotion episodes as an initial test of stockpiling-related payback. The empirical design is a reduced-form, high-dimensional fixed-effects model that compares the same UPC across stores in the same week and the same store across UPCs in the same week, while absorbing store $\times$ UPC, store $\times$ week, and UPC $\times$ week heterogeneity.

The paper is organized around three hypotheses. H1 states that deeper discounts increase same-week unit sales. H2 states that the depth–sales relationship is concave. H3 states that sales fall below baseline after a promotion ends because households draw down inventory. The empirical setting is the publicly available Dominick's scanner data, with cereal identified by commodity code 311 and

 doi:10.20944/preprints202603.1966.v1

promotion depth coded through mutually exclusive discount bins. The resulting design emphasizes transparent replication and interpretable reduced-form estimates rather than proprietary retailer analytics or a fully structural inventory model.

The main empirical result is clear. Deeper temporary price cuts are associated with larger contemporaneous weekly lift. In the preferred top-100 specification, discounts of 5–10%, 10–20%, and more than 20% are associated with approximately 4.7%, 10.4%, and 28.0% higher weekly unit sales, respectively, relative to the 0–5% bin. These effects remain positive and ordered across alternative regular-price definitions, alternative bin schemes, exclusions of chain-wide promotions, stable-UPC restrictions, and trimming rules designed to reduce likely stockout contamination.

The strongest interpretation, however, is narrower than the title alone might suggest. The static bin specifications establish monotone lift, but they do not provide a clean direct test of concavity. Likewise, the corrected promotion-start event study displays a very large week-0 spike and positive short-run post-start effects, yet the immediate lead at $rel = -1$ is also positive and statistically significant. That pattern weakens a causal interpretation of the full dynamic path and implies that the start-based event-study evidence should be read as descriptive rather than definitive evidence on stockpiling.

The paper contributes in four ways. First, it operationalizes promotion depth as a dose–response treatment rather than reducing promotions to a single sale indicator. Second, it estimates a transparent fixed-effects design that fits the structure of public store-level scanner data. Third, it combines static depth-bin estimation with dynamic event-study evidence and a broad robustness program. Fourth, it draws a sharper line between results that are strongly supported and claims that remain provisional.

The remainder of the paper proceeds as follows. Section 2 reviews the relevant literature. Section 3 describes the data, variable construction, and econometric design. Section 4 reports the main static and dynamic findings. Section 5 discusses interpretation, managerial implications, and limits. Section 6 concludes. Appendix A provides a master table-and-figure guide.

2. Related Literature

The literature on grocery promotions makes clear that the empirical statement "larger discounts raise sales" is only a starting point. Temporary price reductions generally increase contemporaneous unit sales, but the mechanisms behind that response are heterogeneous. Promotions are strategically chosen, and many grocery goods are storable. As a result, observed promotion spikes can reflect both equilibrium pricing behavior and household inventory management.

A first strand of work emphasizes equilibrium price dispersion under consumer information frictions. Varian [1] models sales as temporal price dispersion supported by mixed-strategy pricing, with firms randomizing prices across informed and uninformed consumers. The central implication for empirical work is not a specific functional form for the depth–sales curve, but an identification warning: promotions are equilibrium outcomes and are therefore unlikely to be exogenous.

A second strand highlights history-dependent promotion timing. Pesendorfer [2] shows that, in storable categories, the profitability of a sale rises with the time elapsed since the previous sale because effective demand accumulates. Using supermarket scanner data, he documents that both sale probability and sales-period demand increase with time since the last promotion. That logic directly motivates the present paper's attention to event timing and to the possibility that simple sale-versus-no-sale comparisons obscure economically meaningful nonlinearities across discount depths.

A third strand places household inventory behavior at the center of demand dynamics. Hendel and Nevo [3] model consumers as forward-looking inventory managers who trade off current prices, future opportunities, and holding costs. Their framework implies that temporary promotions can induce stockpiling and that static demand models may mismeasure substitution when inventory behavior matters. In the context of the present study, this mechanism provides the theoretical basis for the diminishing-returns hypothesis: once households accumulate sufficient inventory, incremental reductions in price may generate progressively smaller incremental purchases.

 doi:10.20944/preprints202603.1966.v1

Neslin and Schneider Stone [4] add an important measurement insight. They argue that post-promotion dips may appear weak or diffuse in aggregate weekly scanner data even when household-level stockpiling is real. The displacement can be spread across several weeks rather than concentrated in a single sharp trough. This point is directly relevant here because the available dynamic evidence is organized at the weekly store–UPC level.

Finally, Hastings and Washington [5] reinforce the broader identification problem by showing that predictable demand cycles and retailer responses can move together. In scanner-data environments, this implies that causal interpretation requires rich controls for store, product, and calendar shocks rather than naive before-and-after comparisons.

Taken together, the literature leaves an applied gap. Theory and mechanism are well developed, but there is less direct evidence from a transparent reduced-form design that estimates the depth–sales relationship in a high-dimensional store-level panel using public grocery data. This paper addresses that gap by estimating same-week sales responses across promotion-depth bins and by pairing those static estimates with event-study evidence on promotion timing and post-promotion payback.

3. Materials and Methods

3.1. Data and Sample Construction

The empirical setting is the Dominick's Finer Foods public scanner database, using the cereal movement file and the cereal UPC dictionary [6]. The raw weekly movement file (`wcer`) contains 6,602,582 store–UPC–week observations, and the UPC dictionary contains 490 product records prior to cleaning. After merging, each observation includes store, UPC, week, units sold (MOVE), quantity (QTY), price (PRICE), sale status, a data-quality flag, and product descriptors. The analysis focuses on ready-to-eat cereal, identified by commodity code 311.

The Stage 1–4 cleaning pipeline standardizes variable names, preserves UPC identifiers, merges the movement and UPC files, and constructs per-unit price. Because the Dominick's price field can reflect multi-unit deals, the cleaned pipeline defines per-unit item price as

$$p_{sit}^{\text{item}} = \frac{PRICE_{sit}}{QTY_{sit}}, \tag{1}$$

when multi-buy variation is not separately identified. This choice matters even though non-unit quantities are rare, because ignoring them would mechanically overstate effective unit price in bundled promotions.

The main sample is restricted to observations with positive price and nonnegative unit sales. The cleaning workflow also retains standard product observations and excludes merchandise-like placeholders that cannot be interpreted as meaningful cereal package sizes. The merged notebook reports a 100% UPC match rate and then constructs the core variables used throughout the analysis.

Regular price is defined at the store $\times$ UPC level as the median observed per-unit price. Discount depth is therefore measured as

$$depth_{sit} = \frac{p_{si}^{\text{reg}} - p_{sit}^{\text{item}}}{p_{si}^{\text{reg}}}, \tag{2}$$

which is then clipped at zero to obtain $depth\_pos_{sit}$. The dependent variable in the static analysis is

$$y_{sit} = \ln(1 + MOVE_{sit}). \tag{3}$$

The treatment is coded through mutually exclusive discount-depth bins: 0–5% (omitted baseline), 5–10%, 10–20%, and 20%+.

After basic cleaning, the constructed panel contains 4,751,202 observations. Trimming removes the 1% tails of p_{sit}^{item} and p_{si}^{reg} and caps $depth\_pos$ at 0.80, leaving a final trimmed panel of 4,639,362 observations. The identifying variation is substantial. Among 71,245 UPC-weeks, 97.1% include at least two stores, 69.4% exhibit positive within-UPC-week depth variation, and 54.1% display cross-store

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 25 March 2026 doi:10.20944/preprints202603.1966.v1

depth-bin variation within a UPC-week. The preferred static sample uses the top 100 UPCs ranked by cumulative sales, producing 2,685,320 observations across 93 stores, 100 UPCs, and 366 weeks. The corrected balanced event-study sample contains 293,335 observations.

Table 1. Sample construction and estimation samples.

Stage / sample	Observations	Notes
Raw weekly movement file (`wcer`)	6,602,582	Store × UPC × week cereal movement records
UPC dictionary (`upccer`)	490	Product records before cleaning
Constructed clean panel	4,751,202	After merge, positive price, nonnegative sales, and unit-price construction
Trimmed analysis panel	4,639,362	1% tails of p_{sit}^{item} and p_{si}^{reg} removed; *depth_pos* capped at 0.80
Main static sample (top 100 UPCs)	2,685,320	Preferred baseline fixed-effects estimating sample
Event-study working sample (top 50 UPCs)	1,431,668	Pre-event panel before episode construction
Balanced corrected event sample	293,335	Top 50 UPCs; full 17-week windows around promotion start

Notes: Counts are taken from the Stage 1–4 cleaning output and the Stage 5–8 estimation files. The top-100 and top-50 samples are sales-ranked subsets used for static and dynamic estimation, respectively.

Table 2. Summary statistics for key variables in the clean panel.

Variable	N	Mean	SD	Min	Median	95th pct.	99th pct.
price	4,751,202	3.114	0.764	0.05	3.14	4.37	4.85
qty	4,751,202	1.002	0.058	1.00	1.00	1.00	1.00
p_{sit}^{item}	4,751,202	3.113	0.765	0.05	3.14	4.36	4.84
p_{si}^{reg}	4,751,202	3.118	0.708	0.25	3.15	4.25	4.75
depth_pos	4,751,202	0.037	0.073	0.00	0.00	0.169	0.389
MOVE	4,751,202	19.565	58.725	1.00	13.00	46.00	135.00

Notes: Statistics come from Stage 3 before trimming. MOVE is weekly unit sales. *depth_pos* is the positive-clipped discount depth relative to the store × UPC median regular price.

Table 3. Promotion-depth bins in the trimmed sample.

Depth bin	Count	Share (%)
0–5%	3,474,890	74.90
5–10%	537,027	11.57
10–20%	469,218	10.11
20%+	158,227	3.41

Notes: Shares are computed from the 4,639,362-observation trimmed sample. The dominance of the lowest bin motivates using 0–5% as the omitted baseline category.

 doi:10.20944/preprints202603.1966.v1

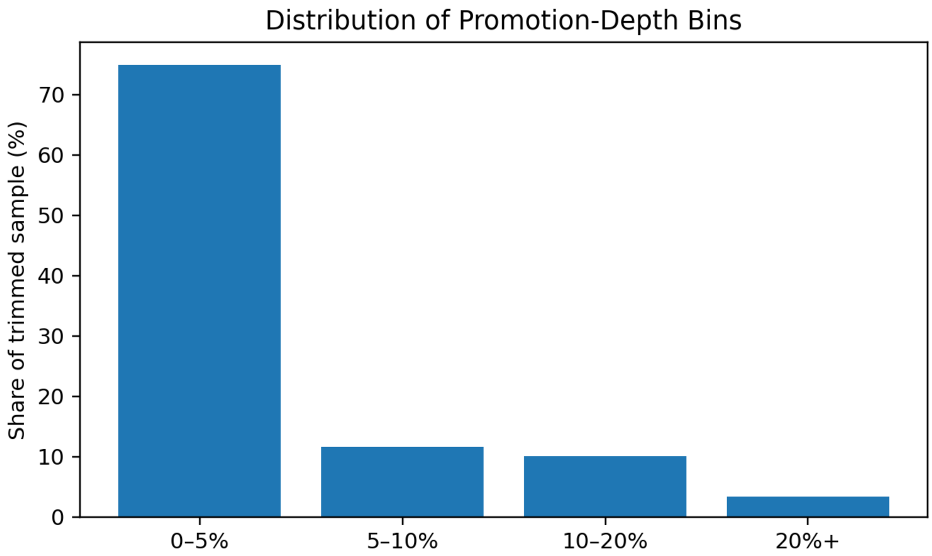

Figure 1. Distribution of promotion-depth bins. The figure shows that most observations are regular-price or shallow-discount weeks, while very deep discounts are comparatively rare.

To describe promotion dynamics beyond depth intensity, promotion episodes are defined as consecutive weeks with $depth\_pos \geq 0.10$ for a given store–UPC pair. Table 4 reports the duration distribution, and Table 5 summarizes episode length and depth.

Table 4. Promotion episode duration distribution.

Duration (weeks)	Episodes	Share (%)
1	99,398	66.26
2	24,538	16.36
3	7,088	4.72
4+	18,989	12.66

Notes: Episodes are defined as consecutive weeks with $depth\_pos \geq 0.10$ for each store–UPC pair.

Table 5. Promotion episode length and depth summary statistics.

Statistic	Episode length (weeks)	Mean depth in episode	Max depth in episode
Mean	4.183	0.272	0.285
Std.	11.744	0.130	0.135
Min	1	0.100	0.100
P25	1	0.155	0.167
Median	1	0.249	0.257
P75	2	0.377	0.401
Max	147	0.784	0.784

Notes: Episodes are defined as consecutive weeks with $depth\_pos \geq 0.10$ for each store–UPC pair. Mean episode length exceeds the median because the distribution has a long right tail.

3.2. Econometric Design

The baseline static design is a high-dimensional fixed-effects model estimated on the top-50 and top-100 UPC subsamples. The identifying variation comes from within-week, within-product, and within-store differences in discount depth while flexibly absorbing confounding demand and supply conditions. The static specification is

$$y_{sit} = \beta_{5\text{–}10} D_{sit}^{5\text{–}10} + \beta_{10\text{–}20} D_{sit}^{10\text{–}20} + \beta_{20+} D_{sit}^{20+} + \alpha_{si} + \gamma_{st} + \delta_{it} + \varepsilon_{sit}, \tag{4}$$

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 25 March 2026 doi:10.20944/preprints202603.1966.v1

where α_{si} denotes store × UPC fixed effects, γ_{st} denotes store × week fixed effects, and δ_{it} denotes UPC × week fixed effects. Estimation uses `AbsorbingLS`, and standard errors are clustered two ways, by store and by UPC.

Under this specification, the coefficients measure the incremental association between discount depth and weekly unit sales relative to the 0–5% reference category. Intuitively, the model compares the same UPC across stores in the same week and the same store across UPCs in the same week, net of persistent store-specific product heterogeneity.

The dynamic design builds promotion episodes from a threshold indicator equal to one when $depth\_pos \geq 0.10$. A clean promotion start is defined as a promotional week preceded by a non-promotional week and by an eight-week pre-window with no earlier promotion for the same store–UPC pair. Non-overlapping starts are retained, event time is indexed from -8 to $+8$, and the corrected event-study sample is restricted to balanced 17-week windows. The promotion-start event-study specification is

$$y_{et} = \sum_{\tau=-8,\,\tau\neq-2}^{8} \theta_\tau \mathbf{1}\{rel_{et} = \tau\} + \lambda_e + \mu_t + u_{et}, \tag{5}$$

where λ_e are event fixed effects and μ_t are calendar-week fixed effects.

The paper evaluates three empirical hypotheses. H1 predicts monotone positive sales lift across depth bins. H2 predicts diminishing marginal returns, which in the present coarse-bin design would imply that the incremental gain from 10–20% to 20%+ is smaller than the gain from 5–10% to 10–20%. H3 predicts a post-promotion dip consistent with inventory drawdown. The static model provides strong evidence on H1, mixed evidence on H2, and the promotion-end event study provides direct evidence on H3.

Stage 9 and Stage 10 robustness checks target three concerns: regular-price mismeasurement, treatment-coding choices, and sample contamination. The first replaces the store × UPC median regular price with the 75th percentile. The second redefines bins more coarsely. The third excludes nearly chain-wide promotions, restricts attention to stable UPCs, and trims likely stockout-risk observations.

4. Results

The static results strongly support H1. Table 6 reports the baseline top-50 and top-100 estimates side by side. Relative to the 0–5% bin, the preferred top-100 model implies approximately 4.7% higher weekly sales for 5–10% discounts, 10.4% for 10–20% discounts, and 28.0% for discounts above 20%. The top-50 sample displays the same ordering with somewhat larger magnitudes.

These effect sizes are economically meaningful. Even modest discounts are associated with sizable weekly lift, while very deep discounts generate much larger contemporaneous volume expansions. The broad conclusion that deeper discounts are associated with greater same-week sales is therefore very strong in these data, and the ranking survives the move from the top-50 exploratory sample to the top-100 preferred specification.

The evidence is weaker for the stronger claim of diminishing marginal returns. In the preferred top-100 baseline, the increase from the 5–10% coefficient to the 10–20% coefficient is approximately 0.053 log points, whereas the increase from 10–20% to 20%+ is approximately 0.147 log points. That pattern is not concave. If anything, the largest marginal gain appears in the deepest discount range. The most defensible interpretation is therefore robust monotonic lift rather than robust concavity in the coarse-bin static models.

The corrected event study around promotion starts provides useful dynamic context, but not a clean causal path. The week-0 coefficient is extraordinarily large, which is plausible for scanner-data promotion starts. Post-start coefficients remain positive and then decay. However, the immediate lead at $rel = -1$ is also positive and highly significant, indicating short-run anticipation, timing contamination, or related marketing activity that begins before the recorded price cut. Far leads from -8 to -3 are much flatter and are not jointly rejected at the 5% level, suggesting that the pretrend issue is concentrated near the start of the promotion window.

Because that event study is centered on promotion start rather than promotion end, it is not by itself a direct test of H3. To examine post-promotion payback more directly, Table 8 reports promotion-end coefficients relative to the pre-promotion baseline. Sales fall below baseline immediately after the promotion ends and remain negative over the full eight-week post window, with implied effects ranging from about −2.45% to roughly −5%. This pattern is consistent with stockpiling followed by inventory drawdown.

Table 6. Main static estimates by promotion depth.

Sample	5–10% coef.	Approx. lift (%)	10–20% coef.	Approx. lift (%)	20%+ coef.	Approx. lift (%)
Top 50 baseline	0.0572	5.88	0.1258	13.40	0.3338	39.63
Top 100 baseline	0.0460	4.71	0.0991	10.42	0.2465	27.95

Notes: The omitted category is the 0–5% discount bin. All specifications absorb store × UPC, store × week, and UPC × week fixed effects and cluster standard errors by store and UPC. The top-100 model is the preferred baseline.

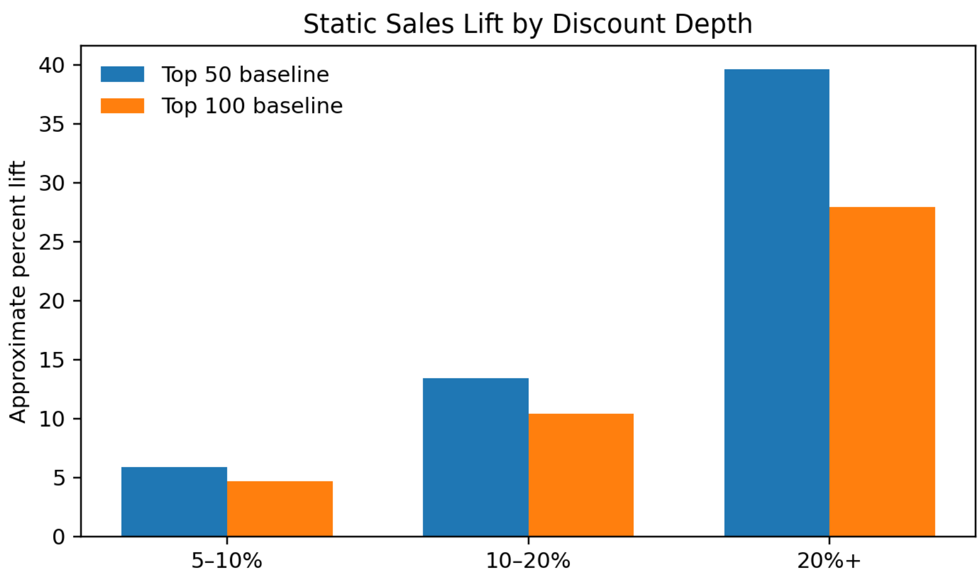

Figure 2. Static sales lift by discount depth in the baseline models. The figure translates log-point coefficients into approximate percentage effects, $100 \times \left(e^{\beta} - 1\right)$, for the top-50 and top-100 baseline specifications.

Table 7. Corrected event-study coefficients around promotion start (selected rows).

Event time	Normalized coef.	SE	Approx. effect (%)	Interpretation
−8 to −3 mean	0.000 (normalized)	—	—	Far-pre-period average normalized to zero
−1	0.0658	0.0182	6.80	Immediate anticipation / timing contamination
0	1.6619	0.0633	426.95	Large contemporaneous promotion-start spike
+1	0.1787	0.0335	19.56	Positive immediate decay
+2	0.1080	0.0291	11.40	Still positive
+3	0.1099	0.0246	11.61	Still positive
+8	0.0307	0.0339	3.12	Small and statistically weak

Notes: Balanced 17-week event windows around non-overlapping promotion starts in the top-50 UPC sample. The omitted event time is $rel = -2$. Far leads (−8 to −3) are jointly borderline flat at conventional levels ($p = 0.0679$), but the immediate near lead at $rel = -1$ is highly significant ($p < 0.001$).

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 25 March 2026 doi:10.20944/preprints202603.1966.v1

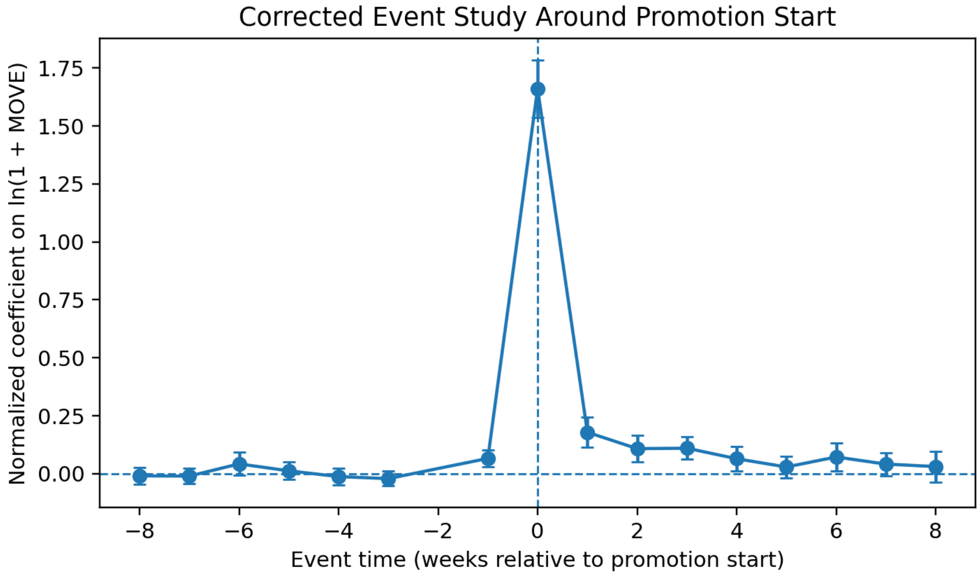

Figure 3. Corrected event study around promotion start. Points plot normalized event-time coefficients from the corrected Stage 8 specification, with 95% confidence intervals. The positive near lead implies that the full path should be interpreted as descriptive rather than fully causal.

Table 8. Promotion-end event-study coefficients (post-promotion dip).

Weeks after end (*rel_end*)	β	SE	Approx. % vs pre
+1	-0.0248	0.0105	-2.45
+2	-0.0367	0.0124	-3.60
+3	-0.0436	0.0096	-4.27
+4	-0.0496	0.0094	-4.84
+5	-0.0527	0.0096	-5.14
+6	-0.0441	0.0103	-4.31
+7	-0.0436	0.0112	-4.27
+8	-0.0478	0.0115	-4.67

Notes: Event time is indexed relative to the end week of a promotion episode ($depth\_pos \geq 0.10$). The dependent variable is $\ln(1 + MOVE)$ minus the store–UPC mean in weeks $start - 8$ to $start - 1$. The model absorbs event fixed effects and calendar-week fixed effects and clusters standard errors by store and UPC. Approximate percentage effects are computed as $100 \times (\exp(\beta) - 1)$.

Placeholder for Figure 4.

Export the promotion-end event-study plot and save it as `figures/figure4_promotion_end_event_study.png`.

Figure 4. Promotion-end event study around promotion end. Coefficients are negative throughout the post period, indicating a persistent post-promotion dip.

 doi:10.20944/preprints202603.1966.v1

Table 9. Robustness comparison across alternative specifications.

Specification	5–10%	10–20%	20%+	Comment
Baseline top 100	0.0460	0.0991	0.2465	Preferred static benchmark
Alternative regular price (Q75)	0.0543	0.0985	0.1758	Reference price = store × UPC 75th percentile
Alternative bins (0–10% omitted)	—	0.0587	0.2073	Different omitted category; directionally informative
Exclude chain-wide promotions	0.0393	0.0786	0.4385	Drops UPC-weeks with >90% of stores on promotion
Stable UPCs only	0.0396	0.0862	0.2240	Keeps UPCs observed in ≥90% of weeks
Stockout-risk trim	0.0479	0.1316	0.2865	Drops bottom-5% store × UPC sales observations

Notes: Entries are coefficient estimates from specifications parallel to the baseline top-100 model unless otherwise noted. The alternative-bins specification uses 0–10% as the omitted category, so those coefficients are not numerically comparable one-for-one with the baseline estimates.

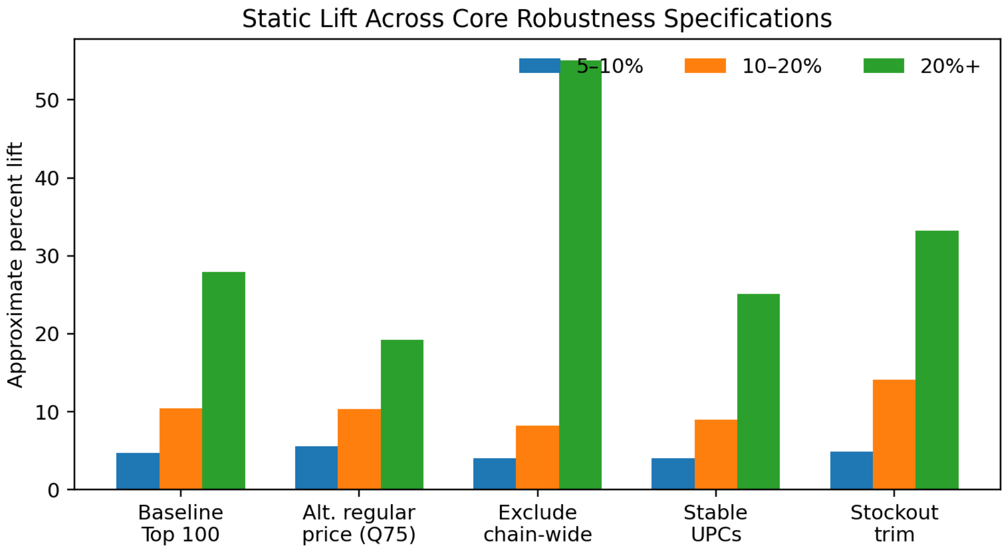

Figure 5. Static lift across core robustness specifications. The ordered depth–lift pattern survives every static perturbation shown.

5. Discussion

The paper's most credible contribution is narrower than the headline claim of diminishing returns. In a rich store–UPC–week fixed-effects design, temporary price cuts in cereal are strongly associated with higher same-week unit sales, and that pattern survives a long list of coding and sample perturbations. For retailers and manufacturers, the immediate implication is straightforward: promotion depth matters materially, and shallow discounts are not operationally equivalent to deep ones.

What the results do not yet identify is an interior optimum or a convincing saturation point within the observed range of discount depths. In other words, the evidence supports the proposition that deeper discounts still sell more, but it does not yet support the stronger managerial conclusion that discounts beyond a particular threshold are inefficient. That distinction matters because a retailer choosing discount depth ultimately cares about profits, category expansion, intra-category substitution, and post-promotion payback, not only UPC-level same-week lift.

Relative to the literature, the static results are consistent with the view that promotions are strategic and history dependent in storable categories. The corrected promotion-start event study also illustrates why identification is difficult. The near-lead coefficient is too large to dismiss and is exactly the kind of pattern one would expect if displays, advertising, or other promotional activity begin before the recorded price cut, or if promotions are timed to coincide with strengthening demand. This does not undermine the project; it simply narrows the claims that can be made from the current dynamic specification.

The safest managerial implication is therefore tactical rather than structural. First, shallow discounts should not be treated as substitutes for genuinely deep promotions if the objective is

 doi:10.20944/preprints202603.1966.v1

immediate volume. Second, any attempt to optimize discount depth should pair contemporaneous lift estimates with direct measurement of post-promotion payback. Third, chains should be cautious in extrapolating item-level lift from highly synchronized chain-wide campaigns, because the robustness results suggest that the behavior of the deepest discounts may differ when promotions are coordinated across stores.

For manufacturers, the findings imply that trade-promotion funding tied to deeper temporary price cuts can generate substantial short-run volume, but the economic meaning of that volume remains unsettled. Without category-level outcome measures, the additional units could reflect category expansion, business stealing within cereal, or intertemporal shifting. That decomposition remains one of the highest-value next steps for the paper.

6. Conclusions

Using Dominick's Finer Foods store–UPC–week cereal scanner data from 1989–1997, this paper estimates how weekly unit sales vary with the depth of temporary price cuts. Across the preferred top-100 baseline and a broad set of robustness checks, deeper discounts are consistently associated with larger same-week sales.

The strongest conclusion is therefore that promotion depth matters and that the contemporaneous depth–lift relationship is robustly monotone across the observed bins. The weaker, but more accurate, conclusion is that the current evidence does not yet establish diminishing marginal returns. In the preferred static model, the incremental gain from moving to the deepest discount range is not smaller than the gain from moving from 5–10% to 10–20%. Likewise, the promotion-start event study does not deliver a clean causal test of post-promotion dip because it is centered on promotion start and contains significant near-lead contamination.

These limits are productive rather than fatal. The cleaned panel, the fixed-effects design, and the robustness program provide a strong reduced-form foundation. The promotion-end event-study evidence strengthens the mechanism interpretation by documenting post-promotion payback. The most valuable next extensions are to estimate curvature with a continuous-depth or spline specification and to study heterogeneity by episode length and promotion intensity.

Author Contributions: Conceptualization, M.J.; methodology, M.J.; software, M.J.; validation, M.J.; formal analysis, M.J.; investigation, M.J.; resources, M.J.; data curation, M.J.; writing—original draft preparation, M.J.; writing—review and editing, M.J.; visualization, M.J.; supervision, M.J.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The study uses the Dominick's Finer Foods Scanner Data distributed through the Kilts Center for Marketing, University of Chicago Booth School of Business. Replication materials can be supplied by the corresponding author upon request.

Acknowledgments: The author thanks the independent-study advisor and readers who provided comments on earlier versions of the project.

Conflicts of Interest: The author declares no conflicts of interest.

Appendix A. Master Table and Figure Guide

Table A1 maps each exhibit to its source stage, recommended location in the paper, and intended analytical role.

doi:10.20944/preprints202603.1966.v1

Table A1. Exhibit map for the full paper.

Exhibit	Title	Source stage / file	Paper section	Purpose
Table 1	Sample construction and cleaning flow	Stages 1–4 PDF / notebook	Materials and Methods	Raw file counts, merge, cleaned and trimmed samples
Table 2	Summary statistics for key variables	Stage 3 output	Materials and Methods	Means, dispersion, medians, and tails for core variables
Table 3	Promotion-depth bins in trimmed sample	Stage 3 / new summary	Materials and Methods	Counts and shares by 0–5%, 5–10%, 10–20%, and 20%+
Table 3A	Promotion episode duration distribution	Step 11 output	Materials and Methods	Duration profile of promotion episodes
Table 3B	Promotion episode length and depth summary statistics	Step 11 output	Materials and Methods	Typical episode duration and discount depth
Figure 1	Distribution of promotion-depth bins	Stage 3 figure	Materials and Methods	Visual treatment distribution
Table 4	Main static estimates (top 50 and top 100)	Stages 5 and 6	Results	Core reduced-form evidence on same-week lift
Figure 2	Static sales lift by discount depth	Stages 5 and 6	Results	Visual comparison of lift across depth bins
Table 5	Corrected event-study coefficients around promotion start	Stage 8 corrected	Results	Balanced-window dynamic coefficients around promotion start
Figure 3	Corrected event study around promotion start	Stage 8 corrected plot	Results	Visual dynamic path around promotion start
Table 6	Promotion-end event-study coefficients	Step 12 output	Results	Direct test of post-promotion dip
Figure 4	Promotion-end event study around promotion end	Step 12 plot	Results	Visual dynamic payback after promotions end
Table 7	Robustness comparison across alternative specifications	Stages 9 and 10	Results	Ordered lift under alternative definitions and restrictions
Figure 5	Static lift across core robustness specifications	Stages 9 and 10	Results	Compressed visual summary of static robustness

References

1. Varian, H.R. A Model of Sales. *Am. Econ. Rev.* **1980**, *70*, 651–659.
2. Pesendorfer, M. Retail Sales: A Study of Pricing Behavior in Supermarkets. *J. Bus.* **2002**, *75*, 33–66.
3. Hendel, I.; Nevo, A. Measuring the Implications of Sales and Consumer Inventory Behavior. *Econometrica* **2006**, *74*, 1637–1673. https://doi.org/10.1111/j.1468-0262.2006.00721.x.
4. Neslin, S.A.; Schneider Stone, L.G. Consumer Inventory Sensitivity and the Postpromotion Dip. *Marketing Lett.* **1996**, *7*, 77–94.
5. Hastings, J.S.; Washington, E.L. The First-of-the-Month Effect: Consumer Behavior and Store Responses. *Am. Econ. J. Econ. Policy* **2010**, *2*, 142–162. https://doi.org/10.1257/pol.2.2.142.
6. Kilts Center for Marketing, University of Chicago Booth School of Business. Dominick's Finer Foods Scanner Data. Data resource, n.d.