Reference-Dependent Pricing

Ariel Xu*

March 26, 2026

Click here for the latest version

Abstract

I provide the first causal evidence that consumers' reference prices, shaped by past price exposures, affect current demand. In a large-scale two-stage field experiment across over 800 vending machines, I first randomly assign machines to different base price levels for 3 to 6 months, then implement common discounted prices across all machines. Arrivals dip immediately after the price increase, then recover over time, consistent with consumers forming new reference prices. In the second stage, consumers with a 10% higher reference price generate 1 to 3% more net revenue when facing identical current prices. I estimate a structural demand model that separates price sensitivity from reference dependence, enabling counterfactual policy analysis. In my setting, the focal firm earns about 6% higher profits by accounting for reference dependence, while promoting 38 percentage points less frequently than a myopic benchmark. Myopic firms overlook promotions' long-run cost to pricing power and misattribute reference dependence to price sensitivity. Over-promoting costs firms more than they gain.

*The University of Chicago, Booth School of Business. Email: arielnxu@chicagobooth.edu. I am deeply indebted to Jean-Pierre Dubé, Avner Strulov-Shlain, Sanjog Misra, and Bradley Shapiro for their advice, encouragement, and support. Thanks go to Xinyao Kong, Walter Zhang, Tesary Lin, Malika Korganbekova, Joshua Dean, Alex Imas, Pradeep Chintagunta, Carl Mela, Puneet Manchanda, Sanjay Jain, Gary Russell, and participants at Booth Marketing Brown Bag, Behavioral Economics Lab, AMA-Sheth Foundation Doctoral Consortium for comments and suggestions. I am grateful to the anonymous company whose collaboration made this research possible, particularly their general manager who facilitated the experimental implementation. All errors are my own.

1 Introduction

Firms spend billions on promotions each year to boost short-term sales. Far less is known about what these promotions cost in the long run. This paper provides the first causal evidence that frequent promotions erode consumers' reference prices, reducing their willingness to pay at regular prices. Using a large-scale field experiment across over 800 vending machines with digital menus, I show that past prices shape consumers' reference points, and these points have real consequences for demand and firm profitability.

That past prices shape current demand has long been recognized (Winer, 1986; Kalwani et al., 1990; Briesch et al., 1997; Bell and Lattin, 2000). The concern for firms is dynamic. A promotion that lifts sales today could lower consumers' future reference prices, making them more price-sensitive tomorrow (Greenleaf, 1995; Mela et al., 1997; Dekimpe et al., 1998). Any short-term gain must be weighed against this long-term erosion. As digital technology enables ever more frequent price changes, these dynamic effects become increasingly important.

Credible evidence on the dynamic effect has been scarce. Prior empirical work relies almost entirely on observational data, where reference price effects are confounded with consumer heterogeneity (Winer, 1986, 1989; Kalwani et al., 1990; Kalwani and Yim, 1992; Krishnamurthi et al., 1992; Hardie et al., 1993; Briesch et al., 1997). Price-sensitive customers self-select into lower prices, creating lower price anchors mechanically (Bell and Lattin, 2000). Separating genuine reference dependence from this selection requires exogenous variation in reference prices. That is what my experiment provides.

I ground the analysis in the prospect-theoretic framework of Kőszegi and Rabin (2006). The theory predicts a simple pattern: consumers exposed to higher past prices should buy more at any given current price. Formally, current demand is monotonically increasing in the reference price. A price cut today lowers tomorrow's reference price, reducing future demand and making it more elastic. This prediction is directly testable.

The experiment has two stages. The first stage creates exogenous variation in reference prices: machines are randomly assigned to different base price levels (the original price, a 3% increase, or a 6% increase) for 3 to 6 months. The second stage tests how these different past prices affect demand. I periodically set discounts so that all consumers face the same current price after discounts,

regardless of their first-stage assignment. Any difference in demand must therefore reflect the reference price. I further split machines into "No Discount," "Discount," and "Promotion" groups (9 groups total), where "Promotion" machines display a sign reading "Up to 10% off." This variation provides additional moments for identifying the structural model.

The Stage 1 price increase produces a pattern consistent with reference price formation: arrivals dip immediately after the increase, then recover as consumers adapt over the following months. The Stage 2 results confirm that these past prices affect demand. Consumers with a 10% higher past price generate 1 to 3% more net revenue and quantity sales when facing identical discounted prices. This is a lower bound of the nonparametric reference price effect. The effect decays over time, consistent with consumers updating their expectations toward the promotional price.

To set prices optimally, a firm needs to separate price sensitivity from reference dependence. Prior gain-loss models allow a flexible functional form that can absorb both effects, making counterfactual analysis unreliable (Putler, 1992). My model, grounded in Kőszegi and Rabin (2006), disciplines the gain-loss structure through consumption utility, achieving this separation.

Estimating the model requires distinguishing two types of consumers. Repeat customers have reference prices shaped by the firm's pricing history. New customers do not. Identifying the share of each type is the key estimation challenge. The promotional sign variation resolves it: the sign shifts only repeat customers' arrival behavior, serving as an exclusion restriction.

The counterfactual analysis compares two scenarios: the focal firm accounting for reference dependence versus myopically optimizing with A/B testing and static demand estimates. The reference-aware scenario earns about 6% higher profits (roughly \$700k to \$1.2 million per year) while discounting 38 percentage points less frequently. Myopic firms maximize current-period profit, ignoring how today's promotions erode tomorrow's reference prices. They also misattribute reference dependence to price sensitivity, overestimating the returns to discounting. As firms increasingly adopt dynamic pricing technologies, myopic demand models might lead them to systematic over-promotion.

These findings contribute to the literature on the long-run effects of price promotions (Gupta, 1988; Mela et al., 1997; Dekimpe et al., 1998; Pauwels et al., 2002). Prior observational work has found that promotions can increase price sensitivity over time, but has not been able to establish the causal mechanism. Anderson and Simester (2004) use field experiments to show that promotion

depth has different long-run effects on new versus established customers, but cannot distinguish among the mechanisms. My experimental design isolates reference dependence from competing explanations such as selection, stockpiling, and income effects. Greenleaf (1995) and Heidhues and Kőszegi (2014) show that reference dependence itself rationalizes the use of promotions, the former using observational scanner data and the latter theoretically. The micro-founded model adds what these papers lack: a clean separation of price sensitivity from reference dependence, making the counterfactual comparison between an RD-aware and a myopic firm credible. In practice, most firms already promote for operational reasons. The question is how they should adjust when consumers are reference-dependent.

The rest of the paper proceeds as follows. Section 2 develops the theoretical framework, derives testable predictions, and lays the foundation for the structural model. Section 3 describes the experimental design and setting. Section 4 presents the experimental results. Section 5 develops and estimates the structural model. Section 6 explores counterfactual pricing strategies. Section 7 concludes.

2 Theoretical Framework

The experimental test requires a prediction, and the counterfactual analysis requires a model that disciplines the gain-loss structure. I model reference-dependent preferences following Kőszegi and Rabin (2006) and derive the central prediction: holding actual price fixed, demand is increasing in the reference price.

2.1 General Model of Reference-Dependent Preferences

A consumer with budget y chooses x units of product 1 at price p and z units of a composite good representing everything else to maximize utility:

$$\max_{x,z} U(x,z) \quad \text{s.t.} \quad p \cdot x + z \leq y$$

A reference-dependent consumer does not maximize consumption utility alone. They compare their outcome to a reference level: doing better feels like a gain, doing worse feels like a loss. This

comparison can be defined over price or over consumption utility.

The prior literature chooses price. Following Putler (1992), gains and losses are defined over the reference price directly through a flexible function $E(p, p^r)$. Putler's framework is practical and has guided much of the empirical work on reference pricing. The problem is that $E(\cdot)$ has no structural constraints. It can absorb both price sensitivity and reference dependence into a single functional form. To set prices optimally, a firm needs to separate the two. The flexibility that makes $E(\cdot)$ useful for fitting data makes it unreliable for prediction.

To separate the two, I define the reference point over consumption utility, not price, following Kőszegi and Rabin (2006). Price enters only through the budget constraint. Total utility becomes

$$\max_{x,z} V(x,z) \equiv U(x,z) + \text{Gain-Loss}(U, U^r) \quad \text{s.t.} \quad p \cdot x + z \leq y$$

where U^r is the reference consumption utility level.[1] I use superscript r to denote reference levels throughout.

Utility depends on two goods. Gains or losses can arise in either one, so the Gain-Loss term must evaluate each good's contribution separately.

The natural decomposition holds one good at its reference level and asks how much changing the other affected utility, then reverses the roles. Kőszegi and Rabin (2006) achieve this by assuming U is additively separable across goods. The structural model in Section 5 uses discrete choice, where U is not separable. I instead decompose the change in U at the reference point into each good's marginal contribution, a finite-difference approach that preserves dimension-by-dimension evaluation without requiring separability.

Assumption 1. $\text{Gain-Loss}(x,z) = \mu(U(x,z^r) - U(x^r,z^r)) + \mu(U(x^r,z) - U(x^r,z^r))$.

The function μ must capture how consumers value gains and losses. Three properties are natural: no gain or loss at the reference point, utility increasing in gains and decreasing in losses, and loss aversion, meaning losses loom larger than equivalent gains.

Assumption 2. μ satisfies: (i) $\mu(0) = 0$, (ii) $\mu'(x) > 0$, (iii) loss aversion: $y > x > 0 \Rightarrow \mu(y) + \mu(-y) < \mu(x) + \mu(-x)$, and (iv) piecewise linearity: $\mu(x) = \eta x$ for $x > 0$ and $\mu(x) = \eta\lambda x$ for $x \leq 0$, where $\eta > 0$ and $\lambda > 1$.

[1] For now, I assume a single reference point, as if the consumer has a degenerate prior.

Piecewise linearity rules out diminishing sensitivity; Kőszegi and Rabin (2006) make the same choice. In my setting, consumption utility differences are small and there is no risk, so diminishing sensitivity adds complexity without empirical bite.

Including both assumptions, the consumer's problem becomes: given (x^r, z^r, p, y),

$$\max_{x,z} U(x,z) + \mu(U(x,z^r) - U(x^r,z^r)) + \mu(U(x^r,z) - U(x^r,z^r)) \quad \text{s.t.} \quad p \cdot x + z \leq y$$

This is the object from which I next derive the comparative statics and testable prediction.

2.2 Comparative Statics and Main Hypothesis

When the actual price equals the reference price, the reference-dependent consumer behaves identically to a classical consumer.

Proposition 1 (Benchmark Case)**.** If (x^r, z^r) solves the classical utility maximization problem at price p^r, and the actual price equals p^r, and product 1 is a normal good, then (x^r, z^r) also solves the reference-dependent problem.

Sketch of Proof. See Appendix, Proof 1.

What happens when the actual price differs from the reference price? Demand increases in the reference price: consumers exposed to higher past prices buy more at any given current price.

Proposition 2 (Reference Price Effects)**.** Consider a consumer with reference consumption (x^r, z^r) formed at reference price p^r. When facing a new price $p \neq p^r$, the consumer's total utility function is approximately, by first-order Taylor expansion:

$$V(x,z) \approx U(x,z) + p^r \cdot \beta\mu(x - x^r) + \beta\mu(z - z^r) \tag{2.3}$$

where $\beta = \partial U(x^r, z^r)/\partial z$. For continuous demand, $dx/dp^r > 0$.

Under logit demand ($x \in \{0,1\}$), the effect of reference price on purchase probability is:

$$\frac{\partial}{\partial p^r} \Pr\{\text{purchase}\} \geq 0$$

with strict inequality when the reference price is not too far from the actual price, and equality (zero effect) at the extremes.

Sketch of Proof. See Appendix, Proofs 2–4. The discrete case requires additive separability of $U(x,z)$; see Proof 3 for why this is needed. The logit case is revisited in detail in the structural model (Section 5).

Fix the actual price p and compare two consumers with reference prices $p^{r2} > p^{r1}$:

1. **Both loss** ($p > p^{r2} > p^{r1}$): Both consumers face a price above their reference. The higher reference price p^{r2} means a smaller loss, so that consumer is more likely to buy.

2. **Gain-loss trade-off** ($p^{r2} > p > p^{r1}$): The consumer with p^{r2} experiences a gain (price below reference); the consumer with p^{r1} experiences a loss. The gain consumer is more likely to buy.

3. **Both gain** ($p^{r2} > p^{r1} > p$): Both consumers see a price below their reference. The higher reference price means a larger gain, so that consumer is again more likely to buy.

At the extremes, reference prices stop mattering. If the actual price far exceeds both reference prices, both consumers face such large losses that the difference between p^{r1} and p^{r2} is irrelevant. Symmetrically, if the actual price is far below both reference prices, the gains are so large that marginal differences in the reference price do not affect behavior.[2]

These comparative statics yield the prediction I take to the field experiment.

Hypothesis 1 (Main Hypothesis)**.** Holding actual price constant at $\bar{p}$, demand D (total sales or purchase probabilities) increases with reference price:

$$\frac{\partial D(p^r, p = \bar{p})}{\partial p^r} > 0 \tag{2.5}$$

[2]This is exact under piecewise linear μ. A consumer who did not plan to buy has a reference point of (no product, full budget for numeraire). This reference does not depend on p^r: if you never planned to spend, the expected price is irrelevant to your baseline. When the actual price is far below the reference price, this type of purchase decision dominates, and p^r drops out.

3 Experimental Design

The theoretical framework did not specify how reference prices form. I focus on one type: price expectations shaped by past prices, consistent with adaptation-level theory (Helson, 1964; Mazumdar et al., 2005; Briesch et al., 1997; Kalyanaram and Winer, 1995; Monroe, 1973; Bolton et al., 2003; Campbell, 1999). This is the reference price that responds to firms' pricing and promotional decisions, making the dynamic cost of promotions visible. The key assumption is that extended exposure to a price leads consumers to adapt their reference point to that price. The experiment identifies whether reference prices causally affect demand. How consumers form these reference prices is a separate question that I do not address here.

3.1 Institutional Background

I partner with a food retail company that operates over 1,000 vending machines across the United States. The machines offer fresh, healthy food in offices, hospitals, schools, and transportation hubs, with little direct competition at each location.

Each machine has a digital menu display that can be updated remotely, allowing the firm to implement price changes flexibly and track them precisely. Consumers see products on shelves when they approach, but prices appear only when they activate the digital screen. This separation lets me distinguish arrival decisions from purchase decisions conditional on arrival.

The system records every session when a consumer activates the screen, whether or not a purchase occurs. Sessions without a purchase are anonymous. When a purchase is made, payment information identifies whether the customer is new or returning, together with their full history.

Prices had been stable for over 8 months before the experiment. Consumers likely had well-established reference prices at baseline.

Data and Key Metrics. The company provides: (1) price menu data (all shelf prices and the exact time of each change); (2) inventory-level data at each machine down to the session level; (3) fulfillment schedules (when each machine was last restocked); and (4) marginal cost data, needed for profit maximization.

On the demand side, the data include: (1) foot traffic data from a third-party provider using

phone location at daily frequency for each location; (2) session counts from the digital screen, recording each interaction even when no purchase occurs; and (3) transaction data with customer identity, from which I observe whether each buyer is a repeat or new customer.

Motivating Evidence. Before the main experiment, the company observed a pattern: after promotional discounts, customers who had experienced the promotion were less likely to complete purchases when returning at regular prices, even after approaching the machine and activating the screen. These were customers who had demonstrated purchase intent. The pattern is consistent with reference dependence, but historical data alone could not rule out selection or other confounds. The experiment described below tests the causal mechanism directly.

3.2 Experimental Design

Testing Hypothesis 1 requires exogenous variation in reference prices while holding current prices fixed. I implement a two-stage randomized experiment. The first stage randomly assigns locations to different price levels for 3 to 6 months, creating variation in reference prices. The second stage equalizes current prices through temporary discounts, so that any demand difference reflects the reference price alone. Figure 1 summarizes the design.

3.2.1 Stage 1: Creating Reference Price Variation (3–6 Months)

The first stage creates the reference price variation. I assigned 818 locations to three price levels: a control group (367 locations) maintaining baseline prices (p_0), a 3% increase group (286 locations), and a 6% increase group (165 locations). To ensure balance across groups, the randomization is stratified by location site (e.g., multiple locations in the same university), type (office, hospital, school, etc.), and geographic market.

These price changes are applied to the base prices of all products[3] and maintained for 3 to 6 months, depending on the date of the price increase. The earliest increase was implemented in mid-August 2024, with two subsequent waves in late October and November 2024. This extended period allows consumers' reference prices to adjust to the new price levels.

[3]There are about 20 products at each location, with little to no menu variation across locations. The products are almost all in the same category. I therefore treat them as one product.

For simplicity, I refer to price groups using price indices. For each location ℓ, the price index at time t is:

$$\text{Price Index}_{\ell,t} = \frac{\sum_s p_{t,s} \cdot r_{0,s}}{\sum_s p_{0,s} \cdot r_{0,s}}$$

where s indexes product SKUs, $p_{t,s}$ is the price of SKU s at time t, and $r_{0,s}$ is the annual pre-experiment revenue share of SKU s.

3.2.2 Stage 2: Identifying Reference Price Effects (10 weeks)

The second stage tests whether these different reference prices affect demand. From March 3 to April 27, 2025, I implement temporary discounts that equalize actual prices across groups. Each first-stage group is further divided into three cells: No Discount (maintaining Stage 1 prices), Discount (price brought to 93% of the pre-experiment baseline), and Promotion (same discount plus promotional messaging visible on the screen before purchase). All price changes are implemented simultaneously on the physical fridge screen and the mobile app.

The three Disc. groups provide the identifying variation. All face the same actual price (93% of p_0) but carry different reference prices from Stage 1:

Group	Actual price p	Reference price p^r
Control	0.93	≈ 1.00
3% increase	0.93	≈ 1.03
6% increase	0.93	≈ 1.06

Same current price, different reference prices. Any difference in demand must reflect the reference price.

One complication: consumers might adapt to the Stage 2 price ($p^r \to p = 0.93$), erasing the variation from Stage 1. Two precautionary measures address this:

1. The discounts are brief and sporadic, with no prior announcement to consumers.

2. Even during the discount window, no explicit discount message appears on the screen; consumers learn about the discount only after activating the screen.

Crossed-out price. The discounted price in the Disc. and Promo groups is shown with the Stage 1 price *crossed out* next to it. Since I am concerned with firms' promotional strategies, which are almost always implemented with crossed-out prices, this design aligns with real-world practice. This does raise the question of whether consumers react to the crossed-out price itself (an "external reference price," Mayhew and Winer, 1992) rather than to memory-based reference prices. I discuss this in detail in Section 4.

Supplemental groups in Stage 2. The No Disc. group serves as the baseline for experiment validity checks (see Section 3.3). The Promo group serves a different purpose: its additional messaging ("Up to 10% off!") creates variation in arrival rates, providing identifying moments for the structural model in Section 5. For a subset of Promo locations in the latter half of Stage 2, registered app users also receive push notifications about the discount. All promotional messages are shown only during HH (defined below).

3.3 Experimental Validity

3.3.1 Could Consumers Have Learned the Discount Pattern?

If consumers learn the discount schedule, they may form a new reference price based on the discounted price, erasing the Stage 1 variation the experiment relies on. The Stage 2 discounts occur during fixed 2-hour windows on Mondays and Tuesdays, 3–5 pm. I refer to this window as Happy Hour (HH). Ideally, the window would vary week to week, but the company's technology does not support this. If consumers learned the pattern, they should shift a higher fraction of their purchases into the HH window.

To test this, I compare the share of purchase sessions during HH over all weekday purchase sessions across No Disc., Disc., and Promo groups, filtering on sessions after consumers have seen the discounted price once.[4] At the aggregate level over the full Stage 2 period, there is no systematic difference: the Disc./Promo groups do not shift purchases to the HH window more than the No Disc. group. The result holds when restricting the denominator to sessions during 3–5 pm on all weekdays. This does not rule out that some individual consumers may have learned the discount, or

[4]For the No Disc. group, this means they have bought once during HH hours, but without a discounted price.

that learning increased toward the end of Stage 2.

3.3.2 Is Selection an Issue?

A natural concern is that higher demand in the high-reference-price group reflects selection rather than reference dependence. The Stage 1 price increase may drive away price-sensitive consumers, leaving only less price-sensitive ones. My data can identify consumers only when they make purchases, not when they start a session without buying, so I cannot fully separate the two mechanisms at the individual level.

The aggregate data can, however, provide a bound. If selection were the only explanation, the high-reference-price group should have **lower** total sales and revenue in Stage 2, since fewer consumers would remain active. The reference price effect predicts **higher** total sales and revenue. My results therefore provide a lower bound on the true reference price effect when the outcome is sales or revenue.

4 Experimental Results

4.1 Randomization and Balance Check

The nine experimental groups are balanced on pre-experiment outcomes. Figure 2 shows residual variation in net revenue, the main dependent variable, after controlling for location type and market fixed effects, before Stage 1 begins. I cannot reject that the nine groups have equal means. Table 2 reports balance on additional variables.

I use location type and market fixed effects in the main analyses because some market-type cells contain fewer sites, making within-cell randomization imperfect. These fixed effects absorb any remaining imbalance at the randomization level.

A second concern is spillovers: consumers who visit multiple locations may be exposed to different Stage 1 prices. To assess this, I track each customer's exposure to distinct price groups across all locations in both Stage 1 and Stage 2.[5] Overall, 97.95% of customers were exposed to

[5] I include Stage 2 because a customer exposed to the regular price in Week 1 of Stage 2 and returning during the HH window in Week 8 is still in the experiment population. This makes the check conservative: it includes customers whose cross-price exposure occurred after the Stage 2 discount, which would not affect the reference price effect I study.

only one price level. Within individual markets, the share is even higher: in most markets, more than 99% of consumers saw a single price. Even in the markets with the highest spillover rates, only 3–4% of consumers encountered multiple prices. Cross-contamination is unlikely to bias the estimates.

4.2 Stage 1 Results: Treatment Effect of the Price Increase

Stage 1 provides preliminary evidence that reference prices adapt to price changes. Figure 4 shows the event study of the 3% price increase for locations treated in mid-August 2024, estimated using the following specification based on Callaway and Sant'Anna (2021):

$$Y_{\ell,t} = \alpha_1^{g,t} + \alpha_2^{g,t} \cdot \mathbb{I}\{G_\ell = g\} + \alpha_3^{g,t} \cdot \mathbb{I}\{T = t\} + \beta^{g,t} \cdot \mathbb{I}\{G_\ell = g\}\mathbb{I}\{T = t\} + \gamma X_{\ell,t} + \varepsilon_{\ell,t} \quad (4.1)$$

where t indexes two-week periods, G is the time period when a location first becomes treated, and g is a specific treatment cohort (e.g., mid-August 2024). Controls $X_{\ell,t}$ include location type and market fixed effects (the randomization strata), an indicator for fridge closure (e.g., school vacations), and $\log(\text{foot traffic} + 1)$. Both fridge closures and foot traffic are measured at the building level and are independent of treatment assignment.

Arrivals dip sharply after the price increase, then recover over the following months. This recovery cannot be explained by an influx of new customers, since the control group would attract new arrivals at a similar rate. The pattern is consistent with consumers initially treating the higher price as a loss, then adapting their reference price upward.

The conversion rate (purchases conditional on arrival) also dips initially, but recovers less. By the end of Stage 1, it remains significantly negative: the law of demand holds even as arrivals recover.

Total quantity sales and net revenue appear roughly flat because these two margins offset: traffic recovers while conversion stays slightly depressed. By the end of Stage 1, I cannot reject that the treatment effects on quantity sales and net revenue are zero.

4.3 Main Results: Effect of Reference Prices

A 10% higher reference price generates 1 to 3% more net revenue when consumers face identical current prices.

The test compares the three Disc. groups from Section 3.2, which share the same actual price ($p = 0.93$) but differ in reference prices from Stage 1. Tables 3 and 5 report regressions of demand outcomes on the Stage 1 price, controlling for location type and market fixed effects, fridge closure, and building-level foot traffic:

$$Y_{\ell,t} = \alpha_1 + \beta_p P^{S1} + \gamma X_{\ell,t} + \varepsilon_{\ell,t} \tag{4.2}$$

where t indexes HH windows in four-week periods. A \$1 higher reference price (base price is about \$7.75) increases the conversion rate by 9.4 percentage points. A 1% increase in reference price leads to about 0.105% increase in net revenue within HH windows. These are lower bounds of the true effect, given the selection argument in Section 3.3.2.

Extrapolating the HH-window revenue effect to approximately 16 operating hours per day, the implied annual revenue gain across the company's 1,000 machines is roughly \$390,000 (95% CI: \$322,000 to \$454,000). This estimate does not account for the revenue decrease during Stage 1; the full dynamic tradeoff is the subject of the counterfactual analysis in Section 6.

These results are robust to alternative specifications. The hypothesis is directional, so the continuous price variable serves as a proxy. Using discrete treatment indicators (3% and 6% dummies) yields similar patterns.

The effects decay over time. They are strongest in Weeks 1–4 (coefficient of 2.296) and moderate to insignificant by Weeks 5–8. The pattern persists across outcomes. As consumers receive more exposure to the promotional price, their reference prices shift toward it, narrowing the gap between treatment and control.

4.3.1 Alternative Explanations

Stockpiling and income effects. The products are perishable (fresh food) with a low price point (about \$7.75 per item). Stockpiling during the discount is not feasible, and spending during the promotional window is too small to meaningfully affect budgets.

Crossed-out prices. As noted in Section 3.2, the discounted price is displayed with the Stage 1 price crossed out. Consumers might react to this "external reference price" (Mayhew and Winer, 1992) rather than to memory-based beliefs. Two patterns cut against this. The Stage 1 results involve no crossed-out prices, yet still show adaptation: arrivals dip and recover over time. And crossed-out prices are static across Stage 2, so if they were the primary mechanism, the treatment effect should not decay. Instead, it weakens steadily over the Stage 2 period, consistent with consumers updating memory-based reference prices.

Selection. As discussed in Section 3.3.2, the Stage 1 price increase may screen out price-sensitive consumers, leaving only less price-sensitive ones in the treatment group. Selection works against my finding. Fewer active consumers in the high-reference-price group should reduce total sales, yet I observe higher total sales. The positive effect I measure is net of this downward bias, making it a lower bound on the true reference price effect.

5 Structural Estimation

The reduced-form results confirm the effect's direction. The counterfactual analysis requires estimates of the underlying parameters: price sensitivity, loss aversion, and the share of repeat customers.

5.1 Consumer Decision Process

The model captures two margins: arrival (whether to activate the screen) and purchase conditional on arrival. Reference prices may affect these margins differently. I consider two types of consumers: repeat customers (type R), who have formed reference prices from past price exposure, and new customers (type N), who have not. Let ρ_R denote the share of potential customers who are repeat customers, treated as fixed during the experimental window.

5.1.1 Repeat Customer Decisions

Arrival. Before activating the screen, a repeat customer's information set includes their reference price p^r and whether a promotional message is displayed. I assume a degenerate distribution

concentrated at a single reference price p^r for each individual repeat consumer as their belief on p^r. The consumer draws an arrival cost $c_R \sim$ Type I Extreme Value and chooses arrival option $j_A \in \{0,1\}$ based on:

$$\tilde{u}_1 = \mathbb{E}_{p,\varepsilon}\left[U(\text{purchase decision}) \mid p^r,\ \text{promo}\right] + c_R$$
$$\tilde{u}_0 = 0$$

where $\mathbb{E}_{p,\varepsilon}\left[U(\text{purchase decision}) \mid p^r,\ \text{promo}\right]$ is the expected utility from starting a session, integrated over purchase utility shocks and the realized price distribution. This assumes rational expectations about the price distribution conditional on p^r and promo. Since $c_R \sim$ Type I Extreme Value, the arrival probability is:

$$\mathscr{A}_R \equiv \Pr\{\text{Arrival} \mid \text{Repeat}\} = \Phi(\mathbb{E}_{p,\varepsilon}\left[U(\text{purchase decision}) \mid p^r,\ \text{promo}\right])$$

where $\Phi(x) = \frac{\exp(x)}{1+\exp(x)}$.

When a promotional message ("up to 10% off") is displayed, repeat customers interpret it relative to their reference price, updating their expectation to $0.9 \cdot p^r$.

Purchase conditional on arrival. After activating the screen, the repeat customer observes the actual price p and draws random utility $\varepsilon \sim$ Type I Extreme Value. They maximize total utility as defined in Section 2. Under logit demand, the purchase probability takes four cases depending on p^r/p (for proof, see Appendix 4):

$$\Pr\{\text{purchase}\} = \begin{cases} \Phi\left(\bar{\psi}_1 - \frac{1+\eta\lambda}{1+\eta}\alpha p\right) & \frac{p^r}{p} \leq \frac{1+\eta\lambda}{\lambda+\eta\lambda} \\ \Phi\left(\bar{\psi}_1 - \alpha p + \alpha\eta\lambda\left(p^r - p\right)\right) & \frac{1+\eta\lambda}{\lambda+\eta\lambda} \leq \frac{p^r}{p} < 1 \\ \Phi\left(\bar{\psi}_1 - \alpha p + \alpha\eta\left(p^r - p\right)\right) & 1 \leq \frac{p^r}{p} < 1 + \frac{\lambda-1}{1+\eta\lambda} \\ \Phi\left(\bar{\psi}_1 - \frac{1+\eta}{1+\eta\lambda}\alpha p\right) & 1 + \frac{\lambda-1}{1+\eta\lambda} \leq \frac{p^r}{p} \end{cases}$$

The four cases correspond to how far the actual price deviates from the reference price. When the actual price is far above the reference (case 1), losses on both the product and numeraire dimensions are so large that the reference price drops out of the purchase probability. When the actual price

is moderately above (case 2), the consumer experiences a loss, weighted by $\eta\lambda$. When the actual price is moderately below (case 3), the consumer experiences a gain, weighted by η. When the actual price is far below (case 4), gains are so large that the reference price again drops out. The middle two cases are where reference dependence has bite.

Adding the promotional message as a constant shift $\Delta_{R,\text{promo}}$ in utility (capturing any informative or prestige effect of the advertisement, Ackerberg, 2003) and subscripting for repeat customers, the structural purchase probability becomes $\mathscr{P}_R = \Phi(\psi_R + \mathbb{I}\{\text{promo}\}\Delta_{R,\text{promo}} + g(p^r, p; \alpha_R, \eta, \lambda))$, where $g(\cdot)$ is the gain-loss component from the corresponding case above.

The purchase probabilities above are derived conditional on different values of (x^r, z^r) (see Appendix 3). To obtain the unconditional expected utility for the arrival decision, note that for purchase decision $j_P \in \{0,1\}$:

$$U_1 = f(p, p^r, \text{promo}; \alpha_R, \psi_R, \eta, \lambda) + \varepsilon_1$$
$$U_0 = \varepsilon_0$$

Given $\mathscr{P}_R = \Phi(U_{\text{purch}})$ where U_{purch} is the deterministic component (the term inside $\Phi(\cdot)$ in $\mathscr{P}_R$), the arrival probability simplifies to:

$$\mathscr{A}_R = \Phi(\mathbb{E}_p\left[\ln\left(1 + \exp(U_{\text{purch}})\right) \mid p^r,\ \text{promo}\right])$$

U_{purch} is a function $f(p^r, p, \text{promo}; \psi_R, \alpha_R, \eta, \lambda)$.

5.1.2 New Customer Decisions

New customers follow the same two-stage process, but their reference prices are exogenous to the firm's pricing decisions. They form price expectations from general knowledge about food prices and nearby competitors, not from past purchases at this vendor.

Arrival. New customers observe the promotional message, but since their reference price is exogenous, the promotion does not shift their belief about the firm's prices in the same way. The

arrival probability is:

$$\mathscr{A}_N \equiv \Pr\{\text{Arrival} \mid \text{New}\} = \Phi(\mathbb{E}_{p,p^{\text{crossed}},\varepsilon}\left[U_N(\text{purchase decision}) \mid \text{ promo}\right])$$

which simplifies to $\mathscr{A}_N = \Phi(\mathbb{E}_{p,p^{\text{crossed}}}\left[\ln\left(1+\exp(U_{N,\text{purch}})\right) \mid \text{ promo}\right])$.

Purchase conditional on arrival. New customers' purchase decisions depend on the current price and, in the Disc. and Promo groups, on the crossed-out Stage 1 price p^{crossed}, which serves as an external reference point:

$$\mathscr{P}_N = \begin{cases} \Phi(\psi_N + g(p;\alpha_N)) & \text{no disc.} \\ \Phi(\psi_N + g(p^{\text{crossed}}, p;\alpha_N,\eta,\lambda)) & \text{disc.} \\ \Phi(\psi_N + \Delta_{N,\text{promo}} + g(p^{\text{crossed}}, p;\alpha_N,\eta,\lambda)) & \text{promo} \end{cases}$$

The promotional effect $\Delta_{N,\text{promo}}$ is allowed to differ from that of repeat customers.

5.2 Data Moments and Identification

The identification challenge centers on ρ_R, the share of repeat customers. At each location ℓ and time t, I observe four quantities: customer flow M_{flow} (potential customers passing by, type unknown), total sessions M_{sess} (screen activations, type unknown), purchases by repeat customers $M_{\text{R,purch}}$, and purchases by new customers $M_{\text{N,purch}}$. These yield three moment conditions:

$$\mathbb{E}[M_{\text{sess}} - M_{\text{flow}} \cdot (\rho_R \mathscr{A}_R + (1-\rho_R)\mathscr{A}_N)] = 0 \tag{5.1}$$

$$\mathbb{E}[M_{\text{R,purch}} - M_{\text{flow}} \cdot \rho_R \mathscr{A}_R \mathscr{P}_R] = 0 \tag{5.2}$$

$$\mathbb{E}[M_{\text{N,purch}} - M_{\text{flow}} \cdot (1-\rho_R)\mathscr{A}_N \mathscr{P}_N] = 0 \tag{5.3}$$

The first condition matches observed sessions to predicted arrivals from both types. The second and third match observed purchases to predicted purchases by repeat and new customers, respectively. These three moments are insufficient to identify all nine parameters. Sessions pool both types: I observe total screen activations but not which type initiated them, so ρ_R is confounded

with the arrival probabilities.

5.2.1 Leveraging Experimental Variation

The experiment resolves this. Stage 1 generates random variation in prices and reference prices; Stage 2 adds variation in gains and losses; the Promo treatment adds variation in promotional messaging. The key exclusion restriction is that promotion shifts repeat customers' arrival behavior (they update their reference price expectation to $0.9p^r$) but does not differentially affect how new customers respond to Stage 1 price variation (their reference prices are exogenous). The resulting difference-in-differences isolates ρ_R:

$$\begin{aligned}
&\left(\mathbb{E}\left[\frac{M_{\text{session}}}{M_{\text{flow}}} \mid p^r = p', \text{promo}\right] - \mathbb{E}\left[\frac{M_{\text{session}}}{M_{\text{flow}}} \mid p^r = p, \text{promo}\right]\right) \\
&- \left(\mathbb{E}\left[\frac{M_{\text{session}}}{M_{\text{flow}}} \mid p^r = p', \text{disc.}\right] - \mathbb{E}\left[\frac{M_{\text{session}}}{M_{\text{flow}}} \mid p^r = p, \text{disc.}\right]\right) \\
&= \rho_R(\mathscr{A}_R(0.9p', \Delta_{R,\text{promo}}) - \mathscr{A}_R(0.9p, \Delta_{R,\text{promo}})) - \rho_R(\mathscr{A}_R(p') - \mathscr{A}_R(p))
\end{aligned} \tag{5.4}$$

The new customer terms cancel because their reference prices do not depend on the Stage 1 price. This moment, combined with standard moments identifying utility and gain-loss parameters, provides sufficient variation to identify all nine parameters. I have one overidentifying restriction from assuming common η and λ across consumer types.

5.3 Estimation Approach

5.3.1 Reference Price Formation

The structural model requires a reference price p_t^r for each consumer at each point in time. I calibrate the formation process using Stage 1 data, comparing three parametric specifications: exponential smoothing ($p_t^r = \theta p_{t-1} + (1-\theta)p_{t-1}^r$), simple averaging over the past N days, and frequency-weighted averaging over the past N days. I fit reduced-form demand regressions on the actual price and the gap $p - p^r$ under each specification. The frequency-weighted average from the past 30 days fits best, consistent with the firm's qualitative observations about customer behavior.

5.3.2 Likelihood

I estimate the model by maximum likelihood. The likelihood has two components.

Arrival likelihood. Session counts follow a Poisson process. For large customer flows, the session rate $M_{\text{sess}}/M_{\text{flow}}$ is approximately Normal:

$$\frac{M_{\text{sess},\ell,t}}{M_{\text{flow},\ell,t}} \sim \mathcal{N}\left(\gamma \cdot [\rho_R \mathscr{A}_R + (1-\rho_R)\mathscr{A}_N],\ \frac{\gamma \cdot [\rho_R \mathscr{A}_R + (1-\rho_R)\mathscr{A}_N]}{M_{\text{flow},\ell,t}}\right)$$

where γ is a scaling parameter accounting for measurement differences between mobile phone ping data and actual potential customers.

Purchase likelihood. For purchases conditional on arrival, I fit the model to observed purchase rates for each customer type. Non-purchasers' types are unobserved, so I derive the posterior probability that a non-purchaser is a repeat customer:

$$\phi_{\ell,t} \equiv \Pr\{\text{Repeat} \mid \text{No Purchase, Arrival}\} = \frac{\rho_R \mathscr{A}_R (1-\mathscr{P}_R)}{\rho_R \mathscr{A}_R (1-\mathscr{P}_R) + (1-\rho_R)\mathscr{A}_N (1-\mathscr{P}_N)}$$

Under large-sample conditions, the observed purchase rates are approximately Normal with binomial sampling variances, yielding a quasi-log-likelihood in closed form.

5.3.3 Handling Experimental Variation

I aggregate data to two-week periods, computing daily-level metrics before aggregating to preserve economic interpretation. For Stage 2, I focus on the discount window with appropriate scaling. To control for the stratified randomization, I use a hierarchical model with market and location-type random effects (integrated out via Laplace approximation) and unrestricted time fixed effects across the 26 two-week periods.

5.4 Estimation Results

The customer base consists of 53% repeat customers ($\rho_R = 0.534$). The two segments differ markedly in price sensitivity: new customers are nearly 70% more responsive to price changes than

repeat customers ($\alpha_N = -0.261$ vs. $\alpha_R = -0.157$).

The loss aversion parameter $\lambda = 3.54$ and gain sensitivity $\eta = 0.49$ yield a loss-gain ratio of approximately 7.2:1. Losses loom far larger than equivalent gains, implying that price increases relative to the reference price have a disproportionate negative effect on demand.

Repeat customers have higher baseline utility ($\psi_R - \psi_N = 0.064$), reflecting their established preference for the product. New customers respond more strongly to promotions ($\Delta_N - \Delta_R = 1.258$), consistent with promotions serving as a trial mechanism for consumers unfamiliar with the product.

Table 6 reports the full parameter estimates for the homogeneous model.

6 Counterfactual

How should firms promote when consumers are reference-dependent? Current practice ranges from heuristic approaches (seasonal sales, competitor matching) to data-driven methods (A/B testing, demand-based optimization). Even the most sophisticated of these maximize current-period profit only. If today's prices affect future demand through reference prices, the firm needs to think dynamically.

I compare two scenarios using the structural estimates to simulate consumer behavior and company data to parametrize the firm's environment.

6.1 Setup

The firm sells a product at regular price $p^H = \$10$ or promotional price $p^L = \$7$, with marginal cost $MC = \$5$. The 30% discount reflects typical promotional depths in grocery retail and the firm's planned discount. Each day, the firm decides whether to promote ($a_t \in \{0,1\}$).

The state space includes three variables: traffic $M_t \in \{M^H = 1000, M^L = 100\}$, traffic regime $z_t \in \{\text{peak}, \text{regular}\}$ capturing seasonal patterns, and reference price p_t^r. Promotional cost reflects coordination required to change the UI and adjust prices: \$0.05 for high-traffic periods and \$0.50 for low-traffic periods.

Tracking the full 30-day price history for reference price formation would create an intractable state space. I approximate the reference price using a traffic-weighted exponentially weighted moving average (EWMA), which captures the key feature that prices during high-traffic periods

have greater influence on reference price formation. The approximation achieves a correlation above 0.88 with the 30-day session-weighted reference price.

6.2 Simulation Scenarios

Scenario 1: Myopic firm. The firm does not know consumers are reference-dependent. It runs a promotional A/B test: one group of machines receives the promotion, the other does not. I simulate consumer responses to this test using the structural estimates, so consumers behave according to their true reference-dependent preferences even though the firm ignores this. The firm estimates the effect of promotion on profit across traffic levels and uses these coefficients to decide when to promote, solving a static maximization in each period:

$$\max_{a \in \{0,1\}} \pi^{belief}(x, z, a; \hat{\theta})$$

where $\pi^{belief}(x, z, a; \hat{\theta}) = (p(a) - c) \cdot Q^{belief}(x, z, p(a); \hat{\theta})$. This excludes the continuation value because the firm does not model the effect of current pricing on future states.

Scenario 2: Reference-aware firm. The firm fully understands reference-dependent preferences and solves the dynamic programming problem:

$$\max_{a \in \{0,1\}} \pi(s, a) + \beta \mathbb{E}[V(s') | s, a]$$

where s collects the current traffic state, season, and reference price. The demand function $Q^{true}(s, a)$ reflects the full reference-dependent preferences. I solve this using modified policy iteration.

6.3 Results

The reference-aware firm earns 6.15% higher daily profits per machine (empirical 95% interval: 4.22% to 8.50%). At the company's current scale, this amounts to \$700k to \$1.2 million per year.

The profit gain comes from promoting less, not more. The reference-aware firm promotes 38 percentage points less frequently than the myopic firm. Figure 6 shows daily purchase sessions over time for both customer types under each scenario. The myopic firm, running frequent promotions,

attracts more sessions from new customers. The reference-aware firm secures more sessions from repeat customers. Because the reference-aware firm charges higher prices for these sessions, and repeat customers are less price-sensitive (as the structural estimates show), the tradeoff favors long-run profit over short-term volume.

The source of the myopic firm's mistake is misattribution. When the model is misspecified with reference price as an omitted variable, the firm conflates loss aversion with price sensitivity. It observes strong negative responses when prices exceed unobserved reference prices and concludes that consumers are highly price-elastic. This pushes the firm toward more frequent discounting, which further erodes reference prices, compounding the error.

7 Conclusion

Reference prices causally affect demand: a 10% higher reference price generates 1 to 3% more net revenue at identical current prices. The effect decays as consumers update their expectations, confirming that reference prices are not fixed anchors but evolving beliefs shaped by price exposure.

The structural model reveals that reference prices operate on two margins. They affect not only whether consumers purchase conditional on arrival, but also whether they show up in the first place. Standard models that focus only on the purchase decision miss this first margin entirely.

The counterfactual analysis quantifies what this means for pricing strategy. In my setting, a firm that accounts for reference dependence promotes 38 percentage points less frequently and earns about 6% higher profits than a myopic benchmark. The myopic firm over-promotes because it misattributes reference dependence to price sensitivity, overestimating the returns to discounting while ignoring the long-run erosion of pricing power.

This dynamic is not unique to vending machines. Any setting where consumers form price expectations from repeated purchases faces the same tradeoff. Identifying it requires two ingredients: variation in past prices that creates different reference prices, and a period where current prices align across consumer groups. These need not come from an experiment; price zone borders, staggered rollouts, or common promotional events can provide both naturally. The structural model, grounded in consumer utility, generalizes across these settings.

As digital pricing technology makes promotions easier and cheaper to implement, the temptation

to promote more frequently grows. The evidence here suggests the opposite prescription: restraint in promotion frequency can be more profitable than aggression, precisely because consumers remember.

References

Daniel A. Ackerberg. Advertising, learning, and consumer choice in experience good markets: an empirical examination. *International Economic Review*, 44(3):1007–1040, 2003. ISSN 1468-2354. doi: 10.1111/1468-2354.t01-2-00098. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/1468-2354.t01-2-00098. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/1468-2354.t01-2-00098.

Eric T. Anderson and Duncan I. Simester. Long-Run Effects of Promotion Depth on New Versus Established Customers: Three Field Studies. *Marketing Science*, 23(1):4–20, February 2004. ISSN 0732-2399. doi: 10.1287/mksc.1030.0040. URL https://pubsonline.informs.org/doi/10.1287/mksc.1030.0040. Publisher: INFORMS.

David R. Bell and James M. Lattin. Looking for Loss Aversion in Scanner Panel Data: The Confounding Effect of Price Response Heterogeneity. *Marketing Science*, 19(2):185–200, May 2000. ISSN 0732-2399. doi: 10.1287/mksc.19.2.185.11802. URL https://pubsonline.informs.org/doi/abs/10.1287/mksc.19.2.185.11802. Publisher: INFORMS.

Lisa E Bolton, Luk Warlop, and Joseph W Alba. Consumer perceptions of price (un) fairness. *Journal of consumer research*, 29(4):474–491, 2003.

Richard A. Briesch, Lakshman Krishnamurthi, Tridib Mazumdar, and S. P. Raj. A Comparative Analysis of Reference Price Models. *Journal of Consumer Research*, 24(2):202–214, September 1997. ISSN 0093-5301, 1537-5277. doi: 10.1086/209505. URL https://academic.oup.com/jcr/article-lookup/doi/10.1086/209505.

Brantly Callaway and Pedro HC Sant'Anna. Difference-in-differences with multiple time periods. *Journal of econometrics*, 225(2):200–230, 2021.

Margaret C Campbell. Perceptions of price unfairness: antecedents and consequences. *Journal of marketing research*, 36(2):187–199, 1999.

Marnik G Dekimpe, Dominique M Hanssens, and Jorge M Silva-Risso. Long-run effects of price promotions in scanner markets. *Journal of econometrics*, 89(1-2):269–291, 1998.

Eric A. Greenleaf. The Impact of Reference Price Effects on the Profitability of Price Promotions. *Marketing Science*, 14(1):82–104, 1995. ISSN 0732-2399. URL https://www.jstor.org/stable/184209. Publisher: INFORMS.

Sunil Gupta. Impact of sales promotions on when, what, and how much to buy. *Journal of Marketing research*, 25(4):342–355, 1988.

Bruce G. S. Hardie, Eric J. Johnson, and Peter S. Fader. Modeling Loss Aversion and Reference Dependence Effects on Brand Choice. *Marketing Science*, 12(4):378–394, November 1993. ISSN 0732-2399, 1526-548X. doi: 10.1287/mksc.12.4.378. URL https://pubsonline.informs.org/doi/10.1287/mksc.12.4.378.

Paul Heidhues and Botond Kőszegi. Regular prices and sales. *Theoretical Economics*, 9(1):217–251, 2014. ISSN 1555-7561. doi: 10.3982/TE1274. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/TE1274. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/TE1274.

Harry Helson. *Adaptation-level theory*. Adaptation-level theory. Harper & Row, Oxford, England, 1964. Pages: xvii, 732.

Manohar U. Kalwani and Chi Kin Yim. Consumer Price and Promotion Expectations: An Experimental Study. *Journal of Marketing Research*, 29(1):90–100, February 1992. ISSN 0022-2437. doi: 10.1177/002224379202900108. URL https://doi.org/10.1177/002224379202900108. Publisher: SAGE Publications Inc.

Manohar U Kalwani, Chi Kin Yim, Heikki J Rinne, and Yoshi Sugita. A Price Expectations Model of Customer Brand Choice. *JOURNAL OF MARKETING RESEARCH*, 1990.

Gurumurthy Kalyanaram and Russell S. Winer. Empirical Generalizations from Reference Price Research. *Marketing Science*, 14(3):G161–G169, 1995. ISSN 0732-2399. URL https://www.jstor.org/stable/184158. Publisher: INFORMS.

Lakshman Krishnamurthi, Tridib Mazumdar, and S. P. Raj. Asymmetric Response to Price in Consumer Brand Choice and Purchase Quantity Decisions. *Journal of Consumer Research*, 19 (3):387–400, December 1992. ISSN 0093-5301. doi: 10.1086/209309. URL https://doi.org/10.1086/209309.

Botond Kőszegi and Matthew Rabin. A MODEL OF REFERENCE-DEPENDENT PREFERENCES. *Quarterly Journal of Economics*, 121(4):1133–1165, November 2006. ISSN 0033-5533. doi: 10.1093/qje/121.4.1133. URL https://research.ebsco.com/linkprocessor/plink?id=92c76f26-fdee-3847-8bb0-9e5c31e3c480. Publisher: Oxford University Press / USA.

Glenn E. Mayhew and Russell S. Winer. An Empirical Analysis of Internal and External Reference Prices Using Scanner Data. *Journal of Consumer Research*, 19(1):62–70, 1992. ISSN 0093-5301. URL https://www.jstor.org/stable/2489188. Publisher: Oxford University Press.

Tridib Mazumdar, S.P. Raj, and Indrajit Sinha. Reference Price Research: Review and Propositions. *Journal of Marketing*, 69(4):84–102, October 2005. ISSN 0022-2429. doi: 10.1509/jmkg.2005.69.4.84. URL https://doi.org/10.1509/jmkg.2005.69.4.84. Publisher: SAGE Publications Inc.

Carl F. Mela, Sunil Gupta, and Donald R. Lehmann. The long-term impact of promotion and advertising on consumer brand choice. *JMR, Journal of Marketing Research*, 34(2):248–261, May 1997. ISSN 00222437. URL https://www.proquest.com/docview/235236369/abstract/4AEF11F0C23439EPQ/1. Num Pages: 14 Place: Chicago, United States Publisher: SAGE PUBLICATIONS, INC.

Kent B Monroe. Buyers' subjective perceptions of price. *Journal of marketing research*, 10(1): 70–80, 1973.

Koen Pauwels, Dominique M Hanssens, and Sivaramakrishnan Siddarth. The long-term effects of price promotions on category incidence, brand choice, and purchase quantity. *Journal of marketing research*, 39(4):421–439, 2002.

Daniel S. Putler. Incorporating Reference Price Effects into a Theory of Consumer Choice. *Marketing Science*, 11(3):287–309, August 1992. ISSN 0732-2399. doi: 10.1287/mksc.11.3.287. URL https://pubsonline.informs.org/doi/abs/10.1287/mksc.11.3.287. Publisher: INFORMS.

Russell S. Winer. A Reference Price Model of Brand Choice for Frequently Purchased Products. *Journal of Consumer Research*, 13(2):250–256, 1986. ISSN 0093-5301. URL https://www.jstor.org/stable/2489230. Publisher: Oxford University Press.

Russell S. Winer. A Multi-Stage Model of Choice Incorporating Reference Prices. *Marketing Letters*, 1(1):27–36, 1989. ISSN 0923-0645. URL https://www.jstor.org/stable/40216131. Publisher: Springer.

Table 1: Average Difference in Percentage of Sessions happened during HH across Consumers

$\bar{Y}_{\text{Control}}$	$\bar{Y}_{\text{Control}} - \bar{Y}_{\text{Disc.}}$	$\bar{Y}_{\text{Control}} - \bar{Y}_{\text{Promo}}$	$\bar{Y}_{\text{Disc.}} - \bar{Y}_{\text{Promo}}$
0.413	0.015	0.001	-0.014
(0.0135)	(0.0187)	(0.0192)	(0.0188)

This table reports the difference between the average fraction of purchase sessions during HH over all sessions ($\bar{Y}$) in Disc./Promo groups compared to the Control group during stage 2. The fraction is calculated as $\bar{Y}_T = \frac{1}{N_T}\sum_{i\in T} M_{i,HH}/M_{i,\text{all}}$ where $T \in \{\text{Control, Disc., Promo}\}$ are the stage 2 groups and M_i is the number of purchase sessions consumer i made after they've made one purchase in HH, i.e. seen the discounted price. Reported standard errors are based on one-sided t-tests under the alternative hypothesis that $\bar{Y}_{\text{Control}} - \bar{Y}_{\text{Disc. or Promo}} < 0$, and $\bar{Y}_{\text{Disc.}} - \bar{Y}_{\text{Promo}} < 0$

Table 2: Balance Check: Residual Variation after controlling for Market and Type FE

	$rp = 1$			$rp = 1.03$			$rp = 1.06$			
	No Disc.	Disc.	Promo	No Disc.	Disc.	Promo	No Disc.	Disc.	Promo	p-value
Number of Sessions	0.646	-2.438	11.454	-5.667	-0.044	2.759	-4.258	1.145	-12.795	0.9030356
	9.729	9.001	8.461	9.102	8.013	9.697	12.370	13.198	10.838	
Number of Purchase Sessions	0.979	-1.954	0.767	-0.231	1.918	2.439	-3.177	1.101	-4.282	0.9972712
	5.668	4.347	4.027	5.861	4.722	6.575	6.073	7.037	6.576	
Net Revenue	7.920	-16.970	5.248	2.154	22.558	20.071	-32.909	12.835	-45.547	0.9924455
	48.489	37.193	33.767	51.425	42.190	54.416	51.636	60.771	55.839	
Gross Revenue	8.098	-15.177	10.257	-2.453	21.811	20.131	-34.478	9.958	-47.232	0.9927326
	49.544	38.328	35.230	51.748	43.211	55.615	53.625	61.524	56.823	
Purchase Amount	3.454	-2.850	2.517	-0.855	2.675	2.375	-4.886	-1.147	-7.465	0.9854883
	7.694	5.616	5.285	7.901	6.180	8.337	7.753	8.732	8.126	
Number of Visits	34.922	-109.844	140.581	-112.468	37.303	-5.563	62.629	134.342	-193.954	0.7008912
	82.903	75.821	79.566	131.196	113.138	160.651	182.697	226.537	136.438	
Conversion Rate	-0.068	-0.035	-0.112	-0.022	0.016	-0.075	-0.031	-0.069	0.017	0.6575014
	0.040	0.045	0.042	0.044	0.055	0.049	0.061	0.063	0.070	

Table 3: Effect of Stage 1 Price on Continuous Outcomes

	log(Net Revenue)			log(Purchase Amt)		
	Pooled	Week 1–4	Week 5–8	Pooled	Week 1–4	Week 5–8
log(Stage 1 Price)	0.105***	0.279***	-0.085	0.103***	0.331***	-0.124
	(0.000)	(0.000)	(0.165)	(0.000)	(0.000)	(0.180)
Stage 2 Control Mean		59.168			9.682	
Stage 2 Control SEM		3.500			0.549	
Type FE		X			X	
Market FE		X			X	

This table reports the reference price effect estimates. The estimates are obtained using cross-sectional regression, controlling for market, location-type fixed effects, building level foot traffic (log-transformed), and fridge closure. Standard errors are two-way clustered at the market level and the location type level, in line with the randomization.

Table 4: Effect of Stage 1 Price on Stage 2 Continuous Outcomes ($)

	Net Revenue			Purchase Amt		
	Pooled	Weeks 1–4	Weeks 5–8	Pooled	Weeks 1–4	Weeks 5–8
Stage 1 Price	1.063***	2.296**	-0.196	0.174*	0.362*	-0.017
	(0.092)	(0.549)	(0.169)	(0.064)	(0.131)	(0.160)
Stage 2 Control Mean		59.168			9.682	
Stage 2 Control SEM		3.500			0.549	
Type FE		X			X	
Market FE		X			X	

This table reports the reference price effect estimates. The estimates are obtained using cross-sectional regression, controlling for market, location-type fixed effects, building level foot traffic (log-transformed), and fridge closure. Standard errors are two-way clustered at the market level and the location type level, in line with the randomization.

Table 5: Average Marginal Effect of Stage 1 Price on Stage 2 Conversion Rate

	Conversion Rate		
	Pooled	Weeks 1 - 4	Weeks 5 - 8
Stage 1 Price	0.094** (0.031)	0.175*** (0.044)	0.007 (0.043)
Stage 2 Control Mean		0.590	
Stage 2 Control SEM		0.009	
Type FE	X	X	X
Market FE	X	X	X

This table reports the reference price effect estimates. The estimates are obtained using binomial GLM controlling for market and vertical fixed effects, building level foot traffic and fridge closure. Standard errors are two-way clustered at the market level and the location type level, in line with the randomization.

Table 6: Structural Model Estimates

Parameter	Estimate	Interpretation
Customer Composition		
Base Repeat Share (ρ_R)	0.5345	53% of potential customers are repeat
Price Sensitivities		
α_R	-0.1567	Repeat customers less price sensitive
α_N	-0.2611	New customers most price sensitive
Baseline Utility Difference		
$\psi_R - \psi_N$	0.0642	Repeat customers higher baseline utility
Difference in Promotion Response		
$\Delta_N - \Delta_R$	1.2583	New customers respond to promotions more
Reference Dependence		
Gain Slope (η)	0.49	Mild gain slope
Loss Aversion (λ)	3.54	Strong loss aversion
Model Fit		
Log-likelihood	68203.910	
Observations	15,682	Location-biweek level

Table 7: Differences between Dynamic Policy and Myopic Policy Outcomes: 1 machine

Metric	Dynamic Policy Mean	Myopic Policy Mean	Difference							
			Mean	Median	Min	2.5%	25%	75%	97.5%	Max
Profit	37414.649	35269.564	2145.086 [2089, 2201]	2124.670	1497.735	1639.053	1910.474	2344.931	2699.686	2857.717
Relative Profit, %	–	–	6.150 [6, 6]	5.968	3.805	4.221	5.275	6.912	8.508	9.122
Discount share	25.592	63.757	-38.165 [-39, -37]	-37.979	-47.877	-45.584	-41.284	-35.154	-30.830	-26.438

Note:
Relative row reports (Dynamic Myopic) / Myopic; values shown in percent.

Figure 1: Simplified Experiment Timeline and Price Implementations

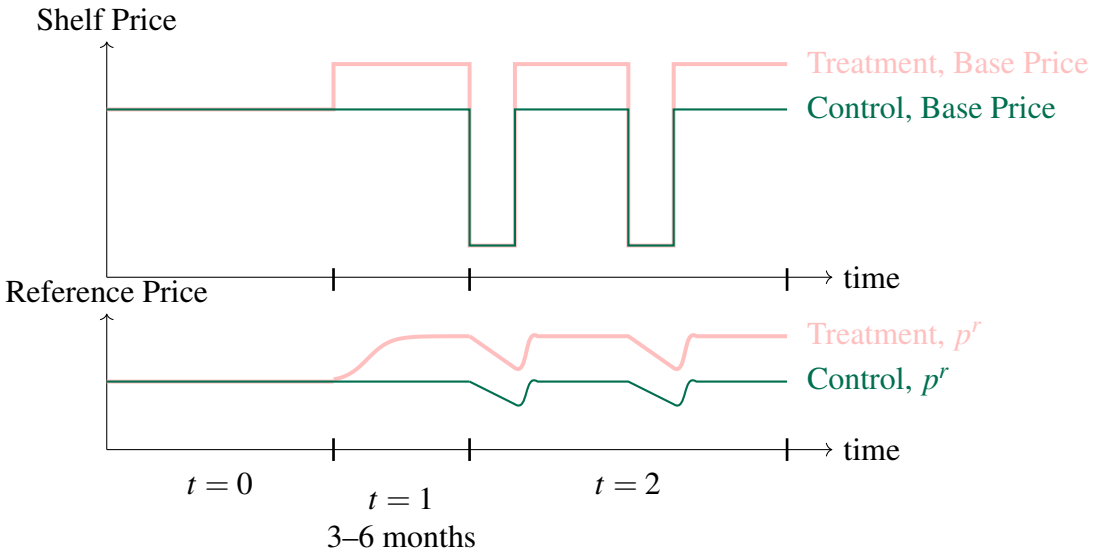
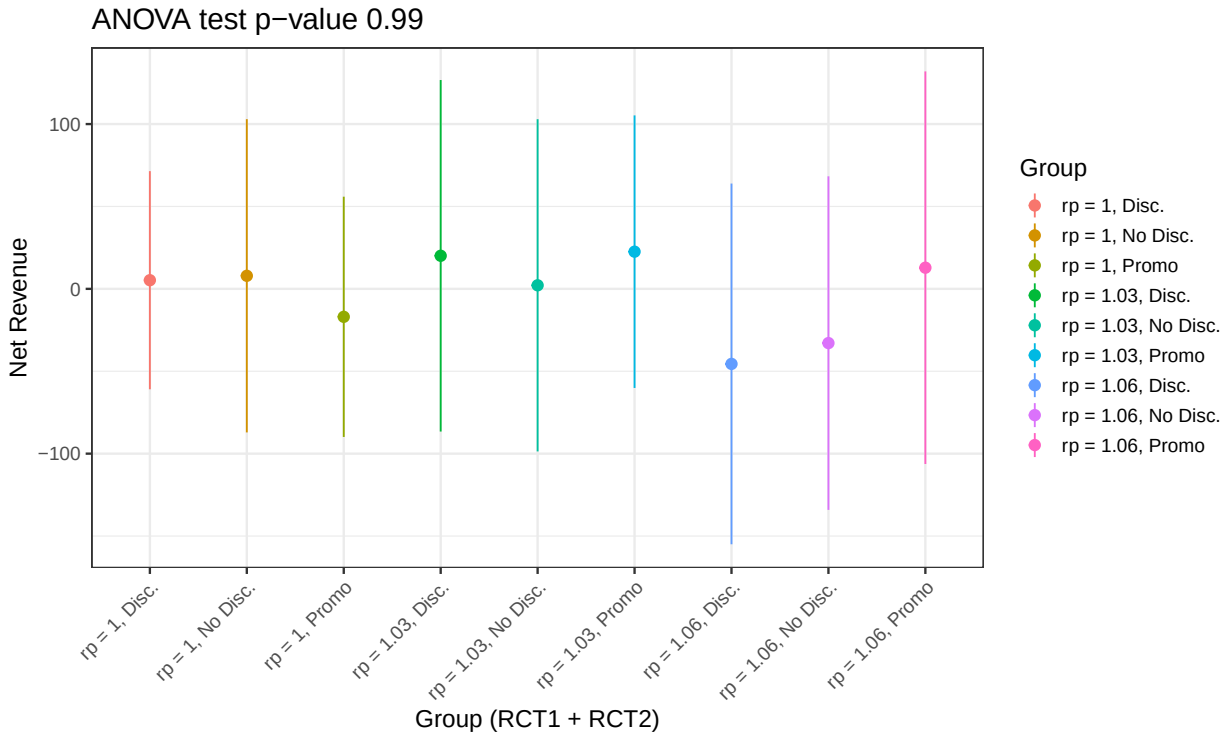


Figure 2: Residual Variation in Net Revenue conditional on Market-Type FE



Note: This figure presents residual variation in net revenue after controlling for Market and Type FE before RCT1 starts. The error bar shows 95% confidence interval. The *rp* values are price indices to denote the treatment assignments. The p-value 0.99 is based on an ANOVA test. I also ran pair-wise t-tests, and we cannot reject that any pairs of groups are equal.

Figure 3: SUTVA Check: Market Level

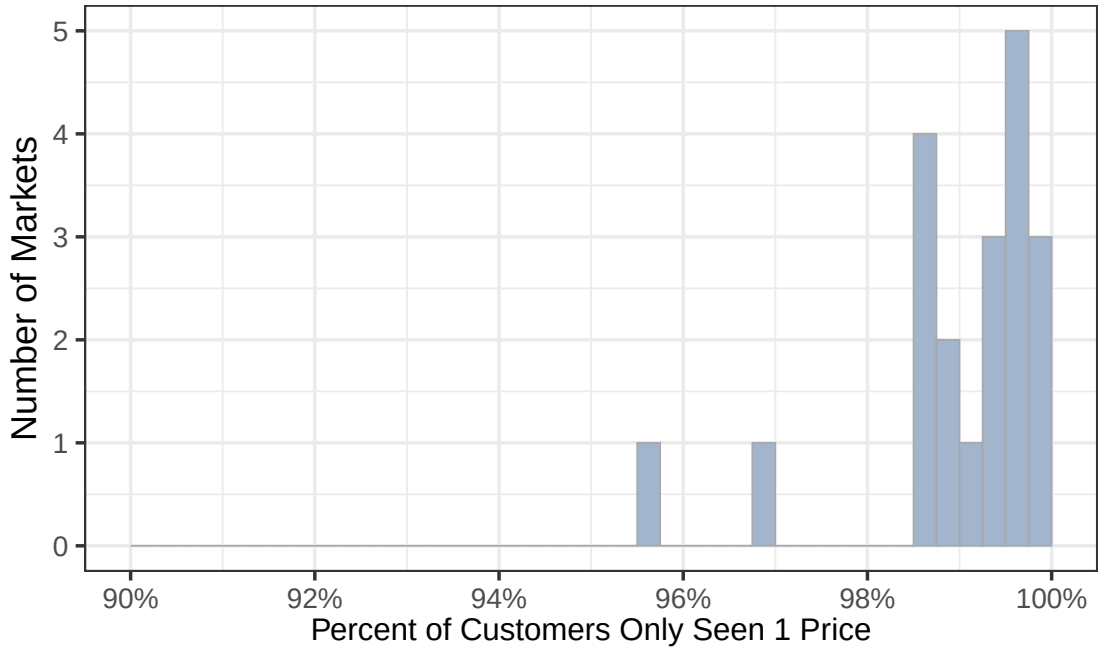
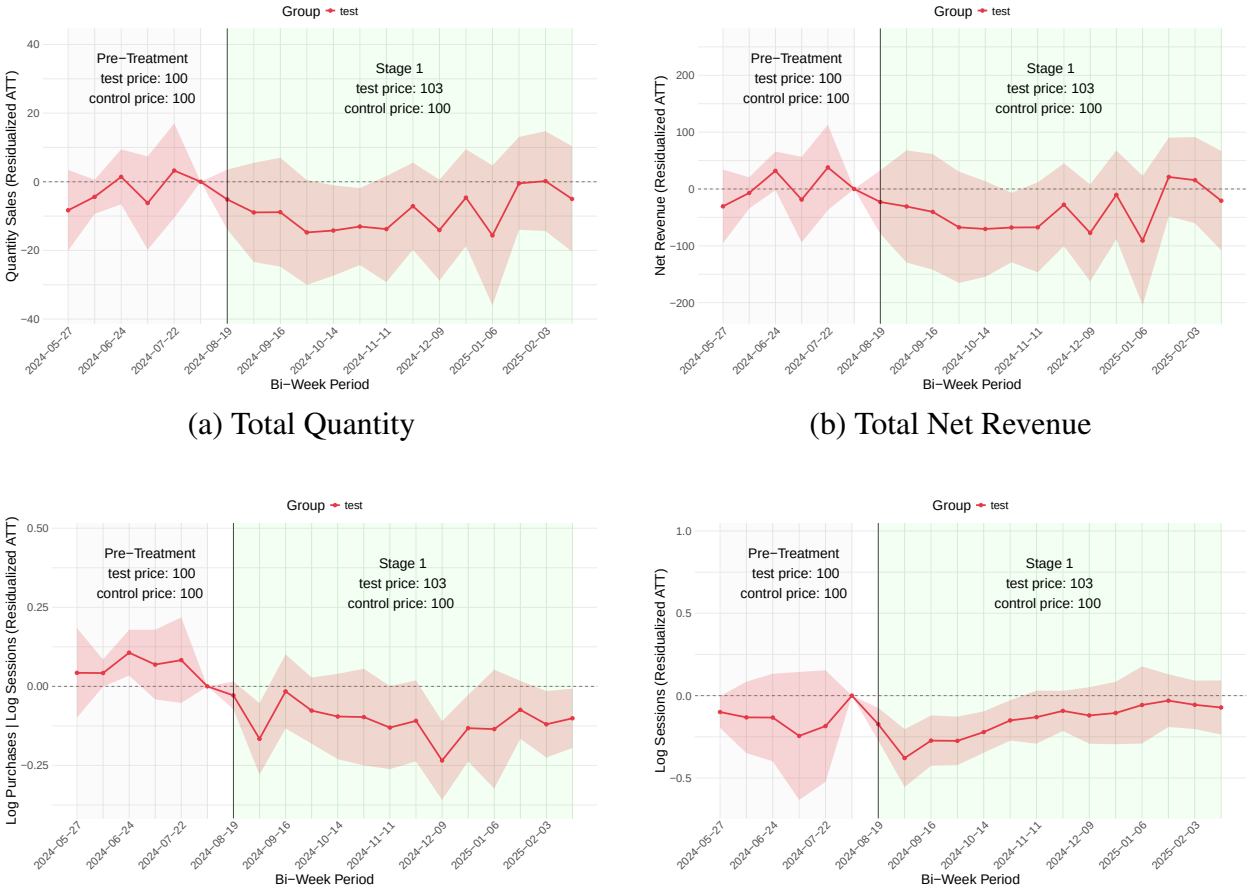


Figure 4: Event Study: ATT of Price Increases in Stage 1



(a) Total Quantity

(b) Total Net Revenue

(c) Purchase Sessions conditional on Sessions

(d) Sessions

Note: This figure presents event study results for locations that adopted a 3% price increase in mid-August 2024, examining four outcomes: (a) quantity sold, (b) net revenue, (c) number of purchase sessions, and (d) number of sessions. The specification includes market and location-type fixed effects, as well as fridge closure (due to school or building closures) and phone-tracked building foot traffic. Outcomes (c) and (d) are log-transformed. For outcome (c), log-tranformed number of sessions is included as an additional control to identify the effect on purchases conditional on sessions. Standard errors are clustered at the market-vertical level (the randomization unit) and calculated using 5,000 bootstrap iterations following Callaway and Sant'Anna (2021). Shaded areas represent 95% confidence intervals.

Figure 5: Aggregate Customer Journey

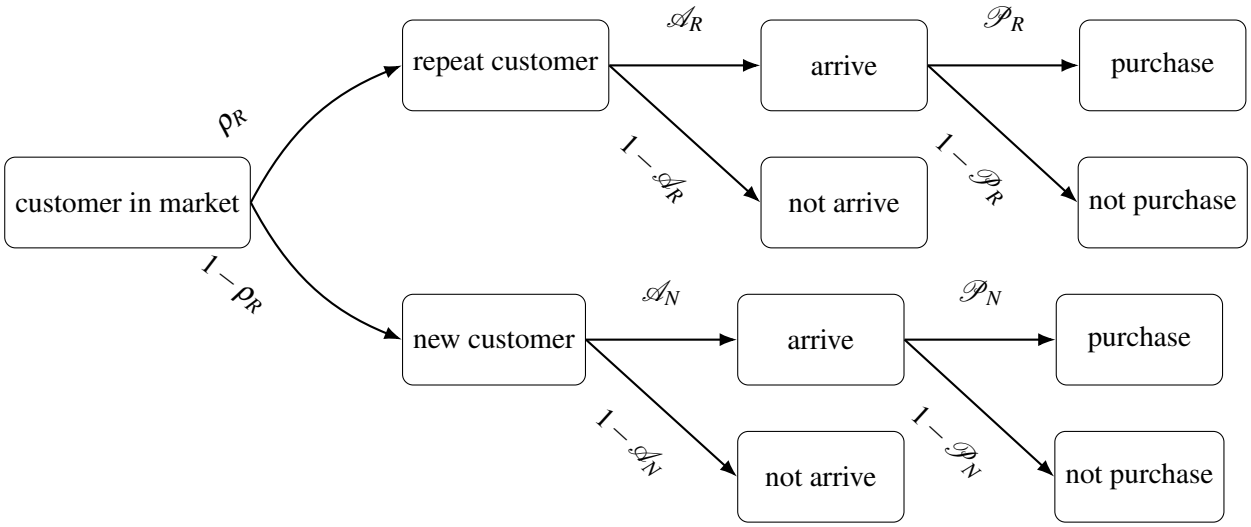


Figure 6: Number of Purchase Sessions from Different Types of Customer for 1 Machine



A Appendix

A.1 Proof of Propositions

Sketch of Proof 1 (for Proposition 1)**.** Suppose

$$(x_1, z_1) = \arg \max_{p^{[r]} \cdot x + z \leq y} \quad U(x,z) + \mu\left(U(x,z^r) - U(x^r,z^r)\right) + \mu\left(U(x^r,z) - U(x^r,z^r)\right)$$

and $(x_1, z_1) \neq (x^{[r]}, z^{[r]})$. This implies

$$\begin{aligned}
& p^{[r]} \cdot x_1 + z_1 = y \\
& U(x_1,z_1) + \mu\left(U(x_1,z^{[r]}) - U(x^{[r]},z^{[r]})\right) + \mu\left(U(x^{[r]},z_1) - U(x^{[r]},z^{[r]})\right. \\
>& U(x^{[r]},z^{[r]}) + \mu\left(U(x^{[r]},z^{[r]}) - U(x^{[r]},z^{[r]})\right) + \mu\left(U(x^{[r]},z^{[r]}) - U(x^{[r]},z^{[r]})\right. \\
\Rightarrow & U(x_1,z_1) - U(x^{[r]},z^{[r]}) + \mu(U(x_1,z^{[r]}) - U(x^{[r]},z^{[r]})) + \mu(U(x^{[r]},z_1) - U(x^{[r]},z^{[r]})) \\
>& 0
\end{aligned}$$

We know that

$$(x^{[r]}, z^{[r]}) = \arg \max_{p^{[r]} \cdot x + z \leq y} \quad U(x,z)$$

$p^{[r]} \cdot x^{[r]} + z^{[r]} = p^{[r]} \cdot x_1 + z_1 = y$, Therefore

$$\begin{aligned}
& \underbrace{U(x_1,z_1) - U(x^{[r]},z^{[r]})}_{\equiv D_u < 0} + \mu(\underbrace{U(x_1,z^{[r]}) - U(x^{[r]},z^{[r]})}_{\equiv D_x}) + \mu(\underbrace{U(x^{[r]},z_1) - U(x^{[r]},z^{[r]})}_{D_z}) \\
\Rightarrow & D_x \cdot D_z < 0
\end{aligned}$$

WLOG, assume $D_x > 0$, Then we need $\mu(D_x) > -(D_u + \mu(Dz))$ to hold for the inequality to hold. Use assumption 2, $\eta D_x > -D_u - \eta\lambda D_z$. This requires:

$$\eta(U(x_1,z^{[r]}) - U(x^{[r]},z^{[r]})) > -\left[U(x_1,z_1) - U(x^{[r]},z^{[r]}) + \eta\lambda(U(x^{[r]},z_1) - U(x^{[r]},z^{[r]}))\right]$$

Use first-order Taylor expansion to approximate, this is approximately

$$\eta(x_1 - x^{[r]})\frac{\partial U(x^{[r]},z^{[r]})}{\partial x} > -\left[U(x_1,z_1) - U(x^{[r]},z^{[r]}) + \eta\lambda(z_1 - z^{[r]})\frac{\partial U(x^{[r]},z^{[r]})}{\partial z}\right]$$

$$\frac{\partial U(x^{[r]},z^{[r]})}{\partial x} \Big/ \frac{\partial U(x^{[r]},z^{[r]})}{\partial z} = p_x^{[r]}, \quad \text{Let } \beta \equiv \frac{\partial U(x^{[r]},z^{[r]})}{\partial z}$$

$$\eta(x_1 - x^{[r]})\beta p_x^{[r]} > -\left[U(x_1,z_1) - U(x^{[r]},z^{[r]}) + \eta\lambda(z_1 - z^{[r]})\beta\right]$$

$$\eta\beta\left[(x_1 - x^{[r]})p_x^{[r]} + \lambda(z_1 - z^{[r]})\right] > -(U(x_1,z_1) - U(x^{[r]},z^{[r]}))$$

$$\eta\beta\left[y - y + (\lambda - 1)(z_1 - z^{[r]})\right] > -(U(x_1,z_1) - U(x^{[r]},z^{[r]}))$$

$$\underbrace{\eta\beta(\lambda - 1)}_{>0}(z_1 - z^{[r]}) > -(\underbrace{U(x_1,z_1) - U(x^{[r]},z^{[r]})}_{<0})$$

Since I assumed $D_x > 0$, i.e. $x_1 > x^{[r]}$, then given (x_1,z_1), $(x^{[r]},z^{[r]})$ are on the same budget line, $z_1 < z^{[r]}$ has to be true, so the LHS of the inequality above is negative. However, the RHS is positive. We arrive at a contradiction. Therefore,

$$(x^{[r]},z^{[r]}) = \arg\max_{p^{[r]}\cdot x + z \le y} \quad U(x,z) + \mu\left(U(x,z^r) - U(x^r,z^r)\right) + \mu\left(U(x^r,z) - U(x^r,z^r)\right)$$

Sketch of Proof 2 (for Proposition 2)**.** What we started with:

$$\begin{aligned} \max_{x,z} \quad & U(x,z) + \mu\left(U\left(x,z^{[r]}\right) - U\left(x^{[r]},z^{[r]}\right)\right) + \mu\left(U\left(x^{[r]},z\right) - U\left(x^{[r]},z^{[r]}\right)\right) \\ \text{s.t.} \quad & p\cdot x + z \le y \end{aligned}$$

By first-order Taylor expansion, this is

$$\approx U(x,z) + \mu\left(\frac{\partial U\left(x^{[r]},z^{[r]}\right)}{\partial x}\left(x - x^{[r]}\right)\right) + \mu\left(\frac{\partial U\left(x^{[r]},z^{[r]}\right)}{\partial z}\left(z - z^{[r]}\right)\right)$$

By our assumption 2,

$$\approx U(x,z) + \frac{\partial U\left(x^{[r]},z^{[r]}\right)}{\partial x}\mu\left(x - x^{[r]}\right) + \frac{\partial U\left(x^{[r]},z^{[r]}\right)}{\partial z}\mu\left(z - z^{[r]}\right)$$

Our proposition stated $\left(x^{[r]}, z^{[r]}\right)$ optimizes $U(x,z)$ when price is $p_x^{[r]}$, so we have $\frac{\partial U\left(x^{[r]},z^{[r]}\right)}{\partial x} / \frac{\partial U\left(x^{[r]},z^{[r]}\right)}{\partial z} = p_x^{[r]}$.

Then by Implicit Function Theorem:

$$\begin{aligned}
\frac{dx}{dp_x^{[r]}} &= -\frac{\frac{\partial^2 V}{\partial x \partial p_x^{[r]}}}{\frac{\partial^2 V}{\partial x^2}} \\
&= -\frac{\beta\mu'\left(x - x^{[r]}\right) - \beta p_2 x^{[r]} \mu''\left(p_x^{[r]} x^{[r]} - p_2 x\right)}{\frac{d}{d_x}\left(U_x - p_2 U_z\right) + p_x^{[r]} \beta\mu''\left(x - x^{[r]}\right) + \beta p_2^2 \mu''\left(p_x^{[r]} x^{[r]} - p_2 x\right)} \\
\text{By our assumption 2, } &= -\frac{\beta\mu'\left(x - x^{[r]}\right)}{\frac{d}{d_x}\left(U_x - p_2 U_z\right)} \\
&> 0
\end{aligned}$$

6

Sketch of Proof 3 (for Proposition 2)**.** For indivisible good $x \in \{0,1\}$, I want to focus on how $\Pr\{x=1\}$ changes with p^r, therefore I introduce stochastic error terms ε's into the utility.

For $\max_{x\in\{0,1\},z} U(x,z,\varepsilon_x)$, s.t. $p^r x + z \le y$, where $U(x,z,\varepsilon_x) \equiv u(x,z) + \varepsilon_x$, the optimal solutions would be at:

$$(x^{[r]}, z^{[r]}) \equiv (x^*, z^*) = \begin{cases} (1, y - p^r), & U(1, y - p^r, \varepsilon_1) > U(0, y, \varepsilon_0) \\ (0, y), & U(1, y - p^r, \varepsilon_1) \le U(0, y, \varepsilon_0) \end{cases}$$

$$\frac{\partial U(x,z,\varepsilon)}{\partial z}\Big|_{(x^*,z^*)} = \begin{cases} \frac{\partial u(1,z)}{\partial z}, & U(1, y - p^r, \varepsilon_1) > U(0, y, \varepsilon_0) \\ \frac{\partial u(0,z)}{\partial z}, & U(1, y - p^r, \varepsilon_1) \le U(0, y, \varepsilon_0) \end{cases}$$

Rewrite the choice-specific total utility function in reduced-form (replacing z with $y - px$):

$$\begin{aligned}
\max_{x\in\{0,1\}} \Big\{ & v_0 \equiv U(0,y,\varepsilon_0) + \mu(U(0,z^{[r]},\varepsilon_0) - U(x^{[r]},z^{[r]},\varepsilon_{x^{[r]}})) + \mu(U(x^{[r]},z,\varepsilon_{x^{[r]}}) - U(x^{[r]},z^{[r]},\varepsilon_{x^{[r]}})), \\
& v_1 \equiv U(1,y-p,\varepsilon_1) + \mu(U(1,z^{[r]},\varepsilon_1) - U(x^{[r]},z^{[r]},\varepsilon_{x^{[r]}})) + \mu(U(x^{[r]},z,\varepsilon_{x^{[r]}}) - U(x^{[r]},z^{[r]},\varepsilon_{x^{[r]}})) \Big\}
\end{aligned}$$

[6] I didn't impose the non-negativity constraint on consumption unit explicitly, otherwise there should be a bound

Use first difference on x instead of taylor expansion, I have

$$
\begin{aligned}
v_0 \approx & U(0,y,\varepsilon_0) + \mu\left((U(1,z^{[r]},\varepsilon_1) - U(0,z^{[r]},\varepsilon_0))\left(0 - x^{[r]}\right)\right) + \mu\left(\frac{\partial u(x,z)}{\partial z}\mid_{x^{[r]},z^{[r]}}\left(y - z^{[r]}\right)\right) \\
\approx & U(0,y,\varepsilon_0) + (U(1,z^{[r]},\varepsilon_1) - U(0,z^{[r]},\varepsilon_0))\mu\left(-x^{[r]}\right) + \frac{\partial u(x,z)}{\partial z}\mid_{x^{[r]},z^{[r]}}\mu\left(y - z^{[r]}\right) \\
= & \begin{cases} u(0,y) + \varepsilon_0 + (u(1,y-p^r) + \varepsilon_1 - u(0,y-p^r) - \varepsilon_0)\mu(-1) + \frac{\partial u(1,z)}{\partial z}\mu(p^r), & (x^{[r]},z^{[r]}) = (1,y-p^r) \\ u(0,y) + \varepsilon_0, & (x^{[r]},z^{[r]}) = (0,y) \end{cases} \\
\equiv & \begin{cases} v(0,1) \\ v(0,0) \end{cases}
\end{aligned}
$$

$$
\begin{aligned}
v_1 \approx & U(1,y-p,\varepsilon_1) + \mu\left((U(1,z^{[r]},\varepsilon_1) - U(0,z^{[r]},\varepsilon_0))\left(1 - x^{[r]}\right)\right) + \mu\left(\frac{\partial u(x,z)}{\partial z}\mid_{x^{[r]},z^{[r]}}\left(y - p - z^{[r]}\right)\right) \\
= & \begin{cases} u(1,y-p) + \varepsilon_1 + \frac{\partial u(1,z)}{\partial z}\mu(p^r - p), & (x^{[r]},z^{[r]}) = (1,y-p^r) \\ u(1,y-p) + \varepsilon_1 + (u(1,y) + \varepsilon_1 - u(0,y) - \varepsilon_0)\mu(1) + \frac{\partial u(0,z)}{\partial z}\mu(-p), & (x^{[r]},z^{[r]}) = (0,y) \end{cases} \\
\equiv & \begin{cases} v(1,1) \\ v(1,0) \end{cases}
\end{aligned}
$$

$v(0,0)$ and $v(1,0)$ doesn't depend on p^r, yet the choice in Stage 1 does. In addition, conditional on $(x^{[r]},z^{[r]}) = (1,y-p^r)$, the probability of purchase is:

$$
\begin{aligned}
\Pr\Big\{ & u(1,y-p) + \varepsilon_1 + \frac{\partial u(1,z)}{\partial z}\mu(p^r - p) > u(0,y) + \varepsilon_0 + (u(1,y-p^r) + \varepsilon_1 - u(0,y-p^r) - \varepsilon_0)\mu(-1) \\
& + \frac{\partial u(1,z)}{\partial z}\mu(p^r)\Big\} \\
\Pr\Big\{ & \eta\lambda(u(1,y-p^r) + \varepsilon_1 - u(0,y-p^r) - \varepsilon_0) + C > \frac{\partial u(1,z)}{\partial z}(\mu(p^r) - \mu(p^r - p))\Big\}
\end{aligned}
$$

I put the terms that are unrelated to p^r in C. Holding p fixed, higher p^r lowers $y - p^r$, i.e. z^r. Because $\partial^2 U/\partial x \partial z > 0$ is generally true when z is numeraire[7], the term $\eta\lambda(u(1,y-p^r) + \varepsilon_1 - u(0,y-p^r) - \varepsilon_0)$ becomes smaller with higher p^r. For the term $\mu(p^r) - \mu(p^r - p)$, If $p^r \geq p$, it's

[7] having more z makes the marginal utility of x higher

not affected by p^r increasing. If $p^r < p$,

$$\begin{aligned}\mu(p^r) - \mu(p^r - p) &= \eta p^r - \eta\lambda(p^r - p) \\ &= \eta \underbrace{(1-\lambda)}_{<1} p^r - \eta\lambda p\end{aligned}$$

With higher p^r, this term gets more negative.

Therefore, $\Pr\{v(1,1) > v(0,1) \mid (x^{[r]}, z^{[r]}) = (1, y - p^r)\}$ doesn't necessarily increase with p^r. We would need to inspect

$$\Pr\{v_1 \geq v_0\} = \Pr\left\{ \left[v(1,1) \geq v(0,1) \bigcap v(1,1) \geq v(0,0) \right] \bigcup \left[v(1,0) \geq v(0,1) \bigcap v(1,1) \geq v(0,0) \right] \right\}$$

It's not trivial to show that this probability increases with p^r.

Therefore, for simplicity, I assume $\max_{x\in\{0,1\},z} U(x,z,\varepsilon_x)$, s.t. $p^r x + z \leq y$, where $U(x,z,\varepsilon_x) \equiv \tau(x) + \omega(z) + \varepsilon_x$, the optimal solutions would be at:

$$(x^{[r]}, z^{[r]}) \equiv (x^*, z^*) = \begin{cases} \tau(1) + \omega(y - p^r), & \tau(1) + \omega(y - p^r) + \varepsilon_1 > \tau(0) + \omega(y) + \varepsilon_0 \\ \tau(0) + \omega(y), & \tau(1) + \omega(y - p^r) + \varepsilon_1 \leq \tau(0) + \omega(y) + \varepsilon_0 \end{cases}$$

$$\frac{\partial U(x,z,\varepsilon)}{\partial z} \mid_{(x^*,z^*)} = \frac{\partial \omega(z)}{\partial z} \mid_{(z^*)}$$

We now have:

$$v(0,1) = \tau(0) + \omega(y) + \varepsilon_0 + (\tau(1) + \varepsilon_1 - \tau(0) - \varepsilon_0)\mu(-1) + \frac{\partial \omega}{\partial z} \mid_{y-p^r} \mu(p^r)$$

$$v(0,0) = \tau(0) + \omega(y) + \varepsilon_0$$

$$v(1,1) = \tau(1) + \omega(y-p) + \varepsilon_1 + \frac{\partial \omega}{\partial z} \mid_{y-p^r} \mu(p^r - p)$$

$$v(1,0) = \tau(1) + \omega(y-p) + \varepsilon_1 + (\tau(1) + \omega(y-p) + \varepsilon_1 - \tau(0) - \omega(y) - \varepsilon_0)\mu(1) + \frac{\partial \omega}{\partial z} \mid_{y} \mu(-p)$$

To derive $\Pr\{v_1 \geq v_0\} = \Pr\left\{ \left[v(1,1) \geq v(0,1) \bigcap v(1,1) \geq v(0,0) \right] \bigcup \left[v(1,0) \geq v(0,1) \bigcap v(1,1) \geq v(0,0) \right] \right\}$, we need to compare term by term which inequality would dominate. This would rely on some assumptions on τ and ω. I omit the proof here since it can get very long, but proof 4 derived the

final form with the assumptions for traditional logit-type discrete choice demand.

Sketch of Proof 4 (for Proposition 2)**.** For the discrete variable x, there are several adjustments we made. First, the gain-loss utility compares the logarithm of consumption utility, and the separable additivity is imposed on this level as well. This is intuitive as we assume the choice is made based on $\ln(U)$. However, we cannot consider $\ln(U)$ as a consumption utility function itself, which would generate misleading results. This is because αz and the budget constraint implicitly set the scale of our utility. Second, since x is discrete, we use the first difference instead of the first-order Taylor expansion to approximate:

$$\ln U\left(x^{[r]},z^{[r]}\right)=\varepsilon_0\left(1-x^{[r]}\right)+(\bar{\psi}_1+\varepsilon_1)x^{[r]}+\alpha z^{[r]}$$

$$\begin{aligned}\text{Assume: } \ln V(x,z) &= \ln U(x,z)+\mu\left(\ln U\left(x,z^{[r]}\right)-\ln U\left(x^{[r]},z^{[r]}\right)\right)+\mu\left(\ln U\left(x^{[r]},z\right)-\ln U\left(x^{[r]},z^{[r]}\right)\right)\\ &\approx \ln U+\mu\left(\Delta\ln U\left(z^{[r]}\right)\left(x-x^{[r]}\right)\right)+\mu\left(\frac{\partial\ln U\left(x^{[r]},z^{[r]}\right)}{\partial z}\left(z-z^{[r]}\right)\right)\\ &\approx \ln U+\Delta\ln U\left(z^{[r]}\right)\mu\left(x-x^{[r]}\right)+\frac{\partial\ln U\left(x^{[r]},z^{[r]}\right)}{\partial z}\mu\left(z-z^{[r]}\right)\end{aligned}$$

$$\text{where } \Delta\ln U\left(z^{[r]}\right)\equiv\ln U\left(1,z^{[r]}\right)-\ln U\left(0,z^{[r]}\right)$$

$$\frac{\partial\ln U\left(x^{[r]},z^{[r]}\right)}{\partial z}\equiv\alpha$$

Assume $\mu(x)=\begin{cases}\eta x & x>0\\ \eta\lambda x & x\le 0\end{cases}$ where $\eta>0$ and $\lambda>1$. Now, the choice-specific gain-loss utility is determined by the reference utility. The expressions for different choice-reference scenarios are:

$$n(j=1\mid r=1)=(\ln V(1,z^r)-\ln V(0,z^r))\,\mu(0)+\alpha\mu(y-p_1-(y-p_1^r))=\alpha\mu(p_1^r-p_1)$$

$$n(j=0\mid r=1)=(\ln V(1,z^r)-\ln V(0,z^r))\,\mu(-1)+\alpha\mu(p_1^r)=(\bar{\psi}_1+\varepsilon_1-\varepsilon_0)\mu(-1)+\alpha\mu(p_1^r)$$

$$n(j=1\mid r=0)=(\ln V(1,z^r)-\ln V(0,z^r))\,\mu(1)+\alpha\mu(-p_1)=(\bar{\psi}_1+\varepsilon_1-\varepsilon_0)\mu(1)+\alpha\mu(-p_1)$$

$$n(j=0\mid r=0)=0$$

The total choice-specific utility for each occasion is:

$$u(j=1 \mid r=1) = v_1 + n(j=1 \mid r=1) = \bar{\psi}_1 + \alpha(y-p_1) + \varepsilon_1 + \alpha\mu(p_1^r - p_1)$$

$$u(j=0 \mid r=1) = \alpha y + \varepsilon_0 + (\bar{\psi}_1 + \varepsilon_1 - \varepsilon_0)\mu(-1) + \alpha\mu(p_1^r)$$

$$u(j=1 \mid r=0) = \bar{\psi}_1 + \alpha(y-p_1) + \varepsilon_1 + (\bar{\psi}_1 + \varepsilon_1 - \varepsilon_0)\mu(1) + \alpha\mu(-p_1)$$

$$u(j=0 \mid r=0) = \alpha y + \varepsilon_0.$$

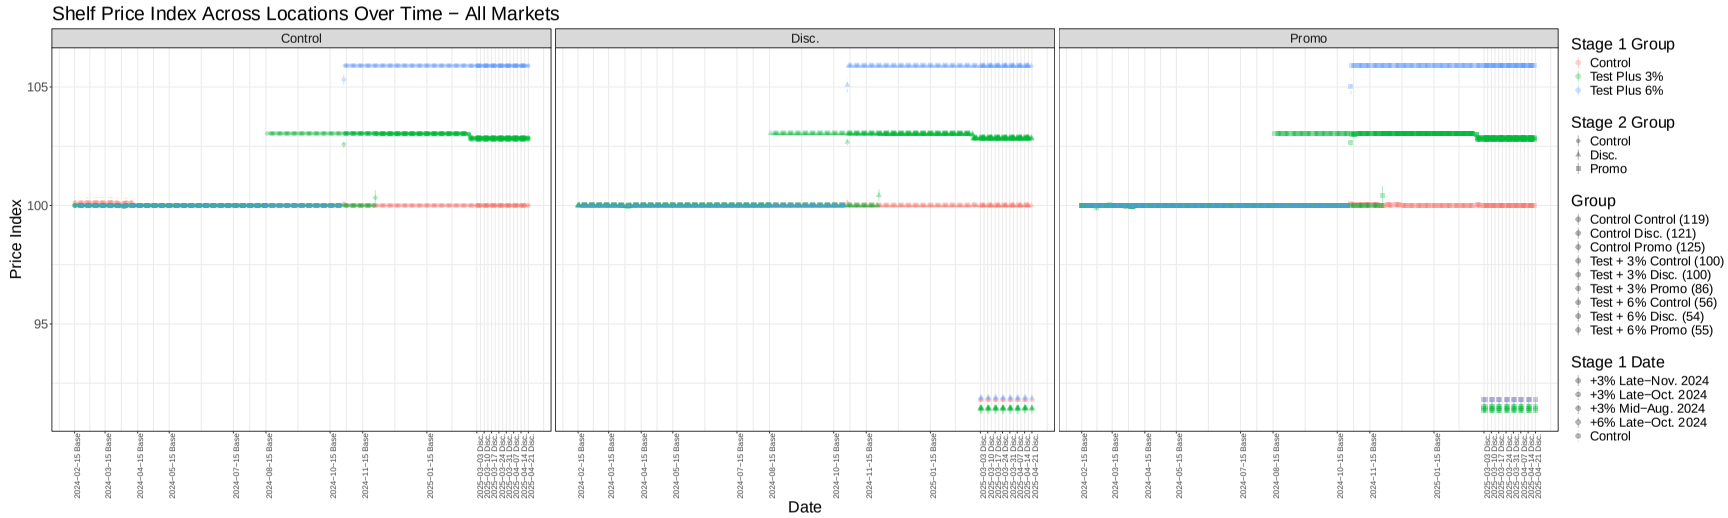
One can solve:

$$\begin{aligned}\Pr\{\text{purchase}\} = \Pr\Big\{ &\Big[u(1,1) \geq u(0,1) \bigcap u(1,1) \geq u(0,0)\Big] \bigcup \\ &\Big[u(1,0) \geq u(0,1) \bigcap u(1,0) \geq u(0,0)\Big]\Big\}\end{aligned}$$

$$\Pr\{\text{purchase}\} = \begin{cases} \frac{1}{1+\exp\left(-\bar{\psi}_1 + \frac{1+\eta\lambda}{1+\eta}\alpha p_x^{\text{new}}\right)} & \frac{p_x^{[r]}}{p_x^{\text{new}}} \leq \frac{1+\eta\lambda}{\lambda+\eta\lambda} \\ \frac{1}{1+\exp\left(-\bar{\psi}_1 + \alpha p_x^{\text{new}} - \alpha\eta\lambda\left(p_x^{[r]} - p_x^{\text{new}}\right)\right)} & \frac{1+\eta\lambda}{\lambda+\eta\lambda} \leq \frac{p_x^{[r]}}{p_x^{\text{new}}} < 1 \\ \frac{1}{1+\exp\left(-\bar{\psi}_1 + \alpha p_x^{\text{new}} - \alpha\eta\left(p_x^{[r]} - p_x^{\text{new}}\right)\right)} & 1 \leq \frac{p_x^{[r]}}{p_x^{\text{new}}} < 1 + \frac{\lambda-1}{1+\eta\lambda} = \frac{\lambda+\eta\lambda}{1+\eta\lambda} \\ \frac{1}{1+\exp\left(-\bar{\psi}_1 + \frac{1+\eta}{1+\eta\lambda}\alpha p_x^{\text{new}}\right)} & \frac{\lambda+\eta\lambda}{1+\eta\lambda} \leq \frac{p_x^{[r]}}{p_x^{\text{new}}} \end{cases}$$

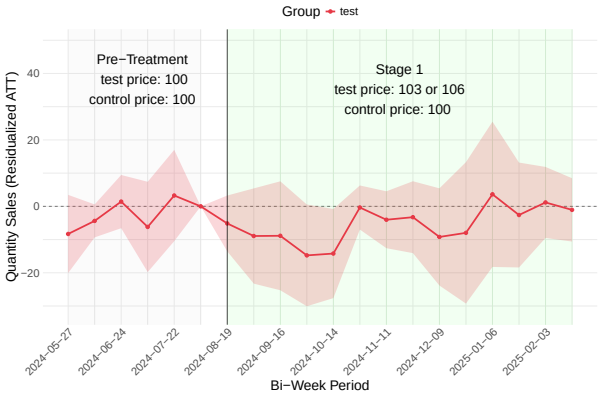
The comparative statics then follows.

Figure 7: Time Series of Price Indices for Each Location

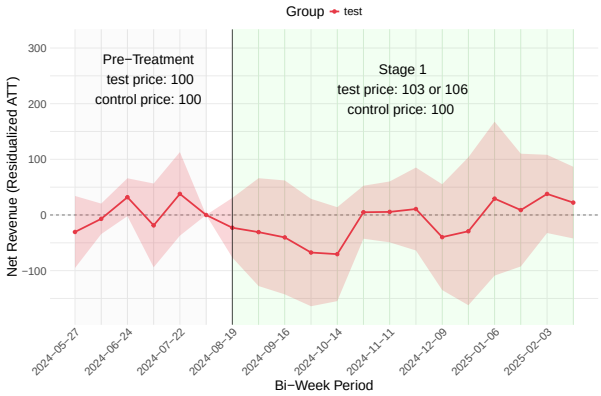


Note: This figure presents time series plot of price indices for each location in my experiment. The price index is calculated using each location's own SKU-level annual revenue prior to the Stage 1 experiment as weight. In February 2025, the firm decided to lower two low-priced items (under $3) by 50 cents in their locations with 3% price increase without informing me. Therefore there is a small dip in the price indices. However, these two items only take up less than 2.5% net revenue at any given time/duration (year, month, the same season as my Stage 2, during the same HH window, etc.) before the experiment starts. During Stage 2's HH windows, these two items even had higher net revenue both in absolute value and in relative value (to the total net revenue of all SKUs) in the group without price increase. Therefore it's reasonable to believe it's not the lowered prices of these two items that drove the Stage 2 results.

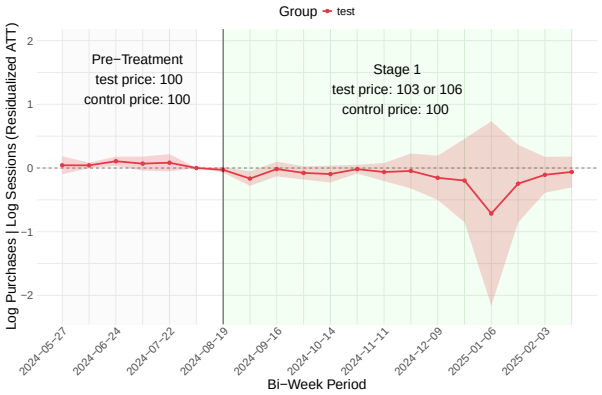
Figure 8: Event Study: ATT of Price Increases in Stage 1



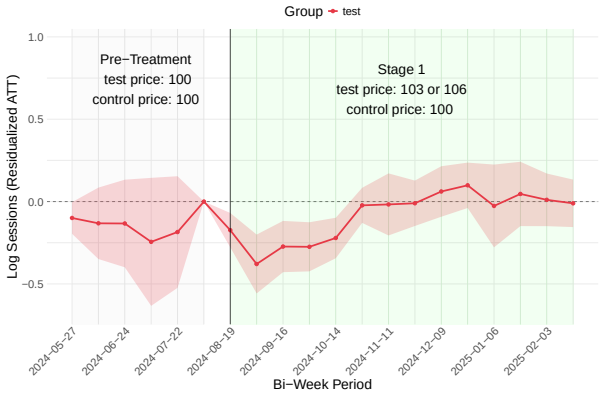
(a) Total Quantity



(b) Total Net Revenue



(c) Purchase Sessions conditional on Sessions



(d) Sessions

Note: This figure presents aggregated event study results for all locations, including both 3% and 6% groups, examining four outcomes: (a) quantity sold, (b) net revenue, (c) number of purchase sessions, and (d) number of sessions. Due to the aggregation, the effect is for a mix of treated and not-yet-treated group (i.e. not pure ATT) before the end of November, 2024. The specification includes market and location-type fixed effects, as well as fridge closure (due to school or building closures) and phone-tracked building foot traffic. Outcomes (c) and (d) are log-transformed. For outcome (c), log-tranformed number of sessions is included as an additional control to identify the effect on purchases conditional on sessions. Standard errors are clustered at the market-vertical level (the randomization unit) and calculated using 5,000 bootstrap iterations following Callaway and Sant'Anna (2021). Shaded areas represent 95% confidence intervals.