Reference-Dependent Demand and Reference Price Persistence:

Pre-Analysis Plan for an Additional Survey Wave

Sharath Sonti

This Version: July 30, 2026

JEL classification: D12, D84, E31, E52

Keywords: inflation, reference prices, reference-dependent preferences, survey experiment

AEA Registry Summary

Trial title. Reference-Dependent Demand and Reference Price Persistence: Pre-Analysis Plan for an Additional Survey Wave.

Country and sampling frame. United States; Prolific representative U.S. sample, ages 18–65.

Status and timing. Additional online survey-experiment wave to be fielded after registration.

Unit of observation and randomization. Individual survey respondent. Randomization is at the respondent level, with no clustered randomization and no stratification.

Planned observations. The new wave targets 1,000 collected responses, with approximately 850 expected to remain after baseline exclusions. Pooled with prior waves, the combined analysis sample is expected to contain approximately 2,100 respondents eligible for analysis and approximately 1,700 respondents in the balanced asymmetric estimation sample.

Randomized components. Respondents are randomized to a reference-price shift size $s \in \{0.12, 0.20, 0.28\}$, to the order in which the loss and gain scenarios are shown, and to one of two reference-price persistence framings: expected-price arm (E) or feels-normal arm (N).

Primary outcomes. Scenario purchase quantities q_L and q_G; elicited scenario reference prices R_L and R_G; and the post-shock updated reference price $\hat{R}$ in the E or N arm. The primary derived estimands are Γ_L, Γ_G, $\Delta\Gamma = \Gamma_L - \Gamma_G$, and per-trip reference-price persistence ρ_{trip} by arm.

Confirmatory hypotheses. H1: $\Gamma_L > 0$; H2: $\Gamma_G > 0$; H3: $\Delta\Gamma > 0$; H4: per-trip ρ is positive in both updating arms.

1 Introduction

Why is inflation so unpopular? Households dislike inflation far more than standard models can rationalize (Shiller, 1997; Stantcheva, 2024; Binetti et al., 2024; Georgarakos et al., 2025). The classical welfare costs are quantitatively small: the inflation tax on real money balances (Fischer, 1981; Lucas, 1981, 2000) and the price-dispersion cost emphasized by the New Keynesian literature (Gali, 2015) both deliver modest welfare losses, and the latter shrinks further once realistic menu-cost frictions are imposed (Nakamura et al., 2018). A recent literature locates larger costs in the labor market, where nominal wages lag prices (Afrouzi et al., 2024; Guerreiro et al., 2024), and in the revaluation of nominal asset positions (Del Canto et al., 2023; Pallotti et al., 2024).

This project studies a complementary reason, rooted in consumer psychology, for why people dislike inflation: they dislike paying a price that is high *relative to what they are used to paying*. Consumers carry an internal *reference price* for the goods they buy, formed from past prices (Mazumdar et al., 2005), and a posted price above that reference is experienced as a loss (Tversky and Kahneman, 1991; Thaler, 1985). Embedding such a reference price into an otherwise standard demand system makes inflation costly through a new channel: by repeatedly pushing posted prices above slowly-updating references, inflation imposes a transaction-utility loss that the standard, second-order price-dispersion accounting misses. As we show below, this loss is *first-order* in inflation and is governed by two structural parameters — a reference-dependence elasticity Γ and reference-price persistence ρ — both of which this survey is designed to estimate. The remainder of this section lays out the reference-dependent demand system, the welfare cost it implies, and how the survey identifies Γ and ρ.

1.1 Individual demand

Consider consumer h on a purchase occasion, choosing among N inside varieties. The utility of inside variety j is

$$u_{hj} = \bar{v}_j + \Gamma \log\left(\frac{R_j^h}{P_j}\right) - \varepsilon \log P_j + \xi_{hj}, \tag{1}$$

Here P_j is the posted price, R_j^h is consumer h's reference price for variety j, $\bar{v}_j$ is a quality term, and $\varepsilon > 1$. The coefficient of interest $\Gamma \geq 0$ is the slope on the log reference-price gap $\log(R_j^h/P_j)$: a good priced below its reference ($R_j^h > P_j$) earns positive *transaction utility*, in the mental-accounting tradition (Thaler, 1985), while a good priced above its reference is penalized. With i.i.d. Type-I extreme-value (Gumbel) taste shocks ξ, choice probabilities are multinomial logit, and if consumer h has M_h independent purchase opportunities in the period the expected quantity of variety j is $\mathbb{E}[q_{hj}] = M_h\,\mathbb{P}_h(j)$, where

$$\mathbb{P}_h(j) = \frac{e^{\bar{v}_j}\,(P_j/R_j^h)^{-\Gamma}\,P_j^{-\varepsilon}}{\sum_{k=1}^{N} e^{\bar{v}_k}\,(P_k/R_k^h)^{-\Gamma}\,P_k^{-\varepsilon}}. \tag{2}$$

1.2 Aggregation to reference-dependent CES demand

When the transaction-utility slope Γ is common across consumers, reference prices are common within a market ($R_j^h = R_j$), and inside-good quality is symmetric, aggregating expected quantities across consumers delivers *exactly* a reference-dependent CES (RD-CES) demand system. Demand is generated by the reference-dependent CES quantity aggregator and its dual price index,

$$C = \left[\frac{1}{N} \sum_{j=1}^{N} s_j^{1/\varepsilon} \, C_j^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}}, \qquad P^{\mathrm{RD}} = \left[\frac{1}{N} \sum_{j=1}^{N} s_j \, P_j^{1-\varepsilon} \right]^{\frac{1}{1-\varepsilon}}, \qquad s_j = \left(\frac{P_j}{R_j} \right)^{-\Gamma}, \tag{3}$$

where the taste shifter s_j tilts the effective CES weight toward varieties priced below their reference ($R_j > P_j$). Minimizing expenditure $\sum_j P_j C_j$ subject to attaining a given C yields variety demand

$$C_j = s_j \left(\frac{P_j}{P^{\mathrm{RD}}} \right)^{-\varepsilon} C, \tag{4}$$

where C is the CES composite and P^{RD} the corresponding reference-dependent price index. Heterogeneity in shopping intensity M_h enters only through a common multiplier, so it shifts the *level* of aggregate demand but not its *composition* across varieties. The RD-CES shifter $s_j = (P_j/R_j)^{-\Gamma}$ is therefore just the exponential of the common transaction-utility term at the individual level — the logit–CES equivalence familiar from the discrete-choice foundations of CES demand (Anderson et al., 1989). The same Γ that governs an individual's purchase probability in (2) governs reference dependence in aggregate demand.

For the macroeconomic embedding it is convenient to index varieties by a continuum $i \in [0, 1]$ in place of the discrete set $\{1, \ldots, N\}$. The reference-dependent CES aggregator, its dual price index, and the implied variety demand are then

$$C = \left[\int_0^1 s(i)^{1/\varepsilon} \, C(i)^{\frac{\varepsilon-1}{\varepsilon}} \, di \right]^{\frac{\varepsilon}{\varepsilon-1}}, \qquad P^{\mathrm{RD}} = \left[\int_0^1 s(i) \, P(i)^{1-\varepsilon} \, di \right]^{\frac{1}{1-\varepsilon}}, \qquad C(i) = s(i) \left(\frac{P(i)}{P^{\mathrm{RD}}} \right)^{-\varepsilon} C, \tag{5}$$

with taste shifter $s(i) = (P(i)/R(i))^{-\Gamma}$. This is the form used in the welfare analysis below, where each variety's reference price $R(i)$ is a state variable.

1.3 The welfare cost of inflation

Embedding the RD-CES shifter (4) in an otherwise standard New Keynesian economy makes the reference price a state variable. A price-setting firm therefore faces a forward-looking problem: in choosing its current price it internalizes not only current demand and the cost of price adjustment, but also the effect of today's price on future demand, which operates through the reference price its customers carry forward via the updating law below. The reference price updates geometrically toward the posted price,

$$R_{t+1}(i) = R_t(i)^{\rho} \, P_t(i)^{1-\rho} e^{\bar{\pi}}, \qquad \rho \in [0, 1), \tag{6}$$

so that the log gap between the posted price and the reference, $d_t \equiv \log(P_t/R_t)$, evolves as $d_t = (\pi_t - \bar{\pi}) + \rho \, d_{t-1}$. The factor $e^{\bar{\pi}}$ makes the reference price grow at the steady-state inflation rate, so the reference-dependence friction is not active in the long run. Sustained inflation above trend keeps posted prices above their slowly-adjusting references, and ρ governs how persistently. Reference dependence enters the household's effective (reference-dependent) price index as a wedge, $p_t^{\mathrm{RD}} = p_t + \chi_t$, with

$$\chi_t = \frac{\Gamma}{\varepsilon - 1}\, d_t. \tag{7}$$

A second-order welfare expansion around the steady state after normalizing by trend inflation gives the per-period welfare loss

$$\ell_t = \underbrace{\chi_t}_{\text{reference-dependence cost}} + \underbrace{\frac{\psi_p}{2}\,\pi_t^2 + \frac{1+\varphi}{2(1-\alpha)}\,\tilde{y}_t^2}_{\text{standard New Keynesian terms}}, \qquad \mathcal{L}_0 = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \ell_t, \tag{8}$$

where $\tilde{y}_t$ is the flexible-price output gap, ψ_p the price-adjustment cost, and φ, α, β are standard preference and technology parameters. The novelty is the leading term: unlike the quadratic price-dispersion and output-gap terms, χ_t enters *linearly*, so reference dependence generates a *first-order* welfare cost of inflation. By (7)–(8), that cost is governed by exactly the two parameters this survey estimates: the reference-dependence elasticity Γ (with separate loss- and gain-side slopes Γ_L, Γ_G, see §1.5) scales the wedge, and reference-price persistence ρ determines, through d_t, how long inflation keeps posted prices above their references. Reference dependence in the price domain thus delivers a welfare loss from inflation that standard models do not capture, consistent with a growing macro literature that builds reference dependence and loss aversion into pricing and consumption (Heidhues and Kőszegi, 2008; Pagel, 2017; Santoro et al., 2014; Angelis, 2025).

1.4 Survey identification of Γ and ρ

The survey identifies both parameters directly, without taking a stand on the full macro model. For Γ, it manipulates the reference price of a single focal product for a respondent while holding its posted price fixed, and elicits the resulting quantity. Let $x_h = \log(R_h/P_h)$ be the respondent-specific log reference-price gap induced by the manipulation, and let ω_h be the focal product's baseline share in respondent h's choice set. Taking the ratio of expected scenario quantity to expected baseline quantity *for the same respondent*, the logit denominator and all fixed respondent-product taste terms cancel, leaving

$$\frac{\mathbb{E}[q_{h1}]}{\mathbb{E}[q_{h0}]} = \frac{e^{\Gamma x_h}}{1 + \omega_h\,(e^{\Gamma x_h} - 1)}. \tag{9}$$

For frequently purchased grocery or household items, the focal product is expected to be a small share of the respondent's overall consumption basket, so $\omega_h \approx 0$, the denominator correction is negligible, and the within-respondent relation is log-linear:

$$\log \mathbb{E}[q_{hb} \mid h, x_{hb}] = \alpha_h + \Gamma\, x_{hb} + O(\omega_h), \tag{10}$$

where b indexes survey scenarios and $\alpha_h = \log \mathbb{E}[q_{h0}]$ is a respondent baseline. Equation (10) is the identification result the design exploits: a within-respondent regression of log quantity on the log reference-price gap recovers the same structural Γ as the aggregate RD-CES system in (4), while permitting respondent-specific reference prices and fixed respondent-by-product tastes, which both cancel in the within-respondent comparison. It also motivates the estimator — a Poisson pseudo-maximum-likelihood (PPML) regression with $\log q_0$ as a respondent baseline (offset), which recovers Γ as a conditional-mean slope without requiring the quantity data to be literally Poisson. Reference-price persistence ρ in (6) is identified separately from a single-trip updating question: after one exposure to a price shock, the respondent's reported updated reference $\hat{R}$ identifies the per-trip ρ as the weight on the old reference relative to the posted price. We do not calibrate this per-trip persistence directly; the object the model requires is persistence at its quarterly frequency. We therefore compute each respondent's per-trip $\hat{\rho}$, compound it to a quarterly horizon using their self-reported shopping frequency (the shopping-frequency question is reproduced in Appendix Figure A8), and average the resulting quarterly $\hat{\rho}$ across respondents.

1.5 Gain and loss sides, and the additional wave

The reference-price manipulation comes in two framings. In the loss scenario the posted price sits *above* the manipulated reference, identifying a loss-side elasticity Γ_L; in the gain scenario the posted price sits *below* the reference, identifying a gain-side elasticity Γ_G. For inflationary episodes in which posted prices rise above reference prices, the welfare-relevant slope in (8) is primarily Γ_L. The gain-side slope Γ_G is used to document asymmetry and is relevant for price decreases, disinflation, or deflation episodes. The updating elicitation separately measures ρ_{trip} in a cognitive (expected-price) and an affective (feels-normal) arm.

This pre-analysis plan covers an additional Prolific wave (approximately 1,000 collected responses) that extends an existing dataset (1,323 respondents collected April 24–27, 2026; 1,262 after baseline exclusions) on which Γ and ρ_{trip} have already been estimated. The additional wave **pre-registers** the primary estimation procedure going forward, eliminating analyst degrees of freedom in the confirmatory tests. Prior-wave evidence used for design and precision planning is reported in §8, after the specifications and exclusion restrictions are fixed.

2 New-Wave Design

2.1 Sampling frame

The new wave will be fielded on Prolific with the representative U.S. sample, ages 18–65, providing a defensible sampling frame for the final paper. Pooling implications relative to the existing data are discussed in §3.

2.2 Sample size and stopping rule

Target sample size: $N = 1{,}000$ **collected responses.** The expected retention rate after the baseline exclusions is approximately 85%, giving roughly $N_{\text{new}} \approx 850$ respondents eligible for analysis. Pooled with prior waves, the combined sample eligible for analysis is expected to be $N \approx 2{,}100$, with a balanced asymmetric estimation sample of approximately 1,700 respondents.

Stopping rule: data collection stops at exactly 1,000 collected responses regardless of interim results. No interim analyses will be performed and no data-dependent (interim-look) stopping is permitted. If Prolific over-recruits by a few responses (as commonly occurs), the final N at the moment data collection is paused is taken as the analysis sample; we do not subsample down.

2.3 Randomization

Respondents are randomized independently at the individual level by the survey software. First, each respondent is assigned a reference-price shift size $s \in \{0.12, 0.20, 0.28\}$ with equal probability. Second, the order of the loss and gain scenarios is randomized with equal probability. Third, respondents are randomized with equal probability to the expected-price arm (E) or the feels-normal arm (N) in the reference-price persistence block. Randomization is not clustered and is not stratified. Before exclusions, expected new-wave cell sizes are approximately 333 respondents per shift size, 500 respondents per order, 500 respondents per updating arm, and $1{,}000/(3 \times 2 \times 2) \approx 83$ respondents per full factorial cell.

2.4 Survey content

The core survey instrument carries over from the prior waves without modification; this subsection describes it in full. Screenshots of every screen are reproduced in Appendix A. Throughout, the respondent's own product, package size, and store are inserted into the question text, so that every question refers to a concrete, familiar good rather than an abstract price.

2.4.1 Opening and baseline elicitation

After consent and instructions, each respondent first fixes the object of measurement:

- **Product and store.** The respondent names a grocery or household product they buy regularly, the package size in which they buy it (so that all subsequent prices and quantities refer to a fixed unit), and the store where they usually buy it.
- **Reference price R_0.** The respondent reports the per-unit price they *typically expect to pay* for that product at that store. This elicited "typically expected" price is the empirical counterpart of the reference price in §1 and serves as the posted price in the scenarios that follow.
- **Baseline quantity q_0.** The respondent reports how many units (of the stated package size) they would buy over the next month at R_0. This is the offset $q_{0,i}$ in the estimating equations of §7.

Two short follow-ups validate the reference price: whether R_0 feels "about right," "too high," or "too low" for the item, and how often the respondent has that expected price in mind while shopping. The respondent also reports shopping frequency for the item (used to compound per-trip ρ to the model's quarterly horizon) and whether the item is perishable. Two attention checks are embedded in this block. The exact setup, reference-price, and baseline-quantity screens are reproduced in Appendix Figures A1–A3.

2.4.2 Reference-dependence elicitation (Γ)

Each respondent then faces two hypothetical scenarios that hold the next-month posted price fixed at R_0 while exogenously shifting the *usual* (long-run) price away from R_0; whether the loss or the gain frame is shown first is randomized across respondents and recorded, and enters the order robustness check of §7. A single shift size of 12, 20, or 28 percent ($s \in \{0.12, 0.20, 0.28\}$) is randomly assigned to the respondent and held constant across both scenarios.

- **Loss scenario.** The usual price is set s *below* R_0, i.e. $(1-s)R_0$. Next month only, the price temporarily *rises* to R_0; thereafter it returns to $(1-s)R_0$. The respondent thus contemplates paying a price (R_0) that is high relative to the scenario's usual price — the loss domain.
- **Gain scenario.** The usual price is set s *above* R_0, i.e. $(1+s)R_0$. Next month only, the price temporarily *falls* to R_0; thereafter it returns to $(1+s)R_0$. The respondent contemplates paying a price below the scenario's usual price — the gain domain.

Within each scenario the respondent answers two questions: (i) the per-unit price they would *typically expect to pay going forward* once the usual price resumes (the elicited scenario reference, R_L or R_G), and (ii) the number of units they would buy next month at the posted price R_0. The quantity question is worded to isolate within-month demand and to fix the unit; the loss-scenario version reads:

> *Reminder: the item is [product], one unit is [package size], and your usual store is [store]. Assume that at your usual store the price of this item is usually $(1-s)\,R_0$ per unit. Next month only, the price will temporarily increase to R_0 per unit. After next month, the price will return to $(1-s)\,R_0$ per unit. Assume your income and the prices of other goods stay the same, and answer based only on how much you would want to buy for next month's use — ignore buying extra to store, or delaying purchases to a later month. How many units would you buy next month at R_0 per unit at your usual store?*

The gain-scenario wording is identical, with $(1+s)R_0$ as the usual price and a temporary *decrease* to R_0. Comparing the stated quantity at the fixed posted price R_0 across the exogenously shifted references identifies the reference-dependence elasticity: respondents whose usual price is below R_0 (loss) trace out Γ_L, and those whose usual price is above R_0 (gain) trace out Γ_G, exactly as in (10). The assigned shift s supplies the instrument for the self-reported reference in the control-function specifications of §7. The exact loss- and gain-scenario screens that elicit this information — both the quantity and the going-forward reference-price questions — are reproduced in Appendix Figures A4–A7.

2.4.3 Reference-price persistence elicitation (ρ)

A separate block measures reference-price persistence after a price change. The respondent is told that the price on their previous trips was R_0 and that, on their *next* shopping trip, they observe a price 20% higher, $1.2\,R_0$. They are then asked what price they would carry forward to subsequent trips. Respondents are randomly assigned to one of two framings, which separate the respondent's articulated price expectation from the internal reference price:

> *Suppose that on your previous shopping trips, the price of [product] was R_0 per unit. Assume that on your next shopping trip to [usual store], the price of [product] is $1.2\,R_0$ per unit.*
>
> ***Expected arm (E):*** *After seeing the higher price on that one trip, what price per unit would you typically expect to pay for this item going forward?*
>
> ***Normal arm (N):*** *After seeing the higher price on that one trip, what price would feel like the normal price for this item going forward — that is, a price you have gotten used to paying that feels neither too high nor too low?*

The reported updated reference $\hat{R}$ implies per-trip reference-price persistence $\hat{\rho}_{\text{trip}}$ in (17): $\hat{\rho}_{\text{trip}} = 0$ means the reference jumps fully to the new price after a single exposure, while $\hat{\rho}_{\text{trip}} = 1$ means it does not move at all. As the marketing literature has found, a consumer's articulated price expectation may diverge from the internal reference price against which deals are evaluated (Thomas and Menon, 2007). The E arm captures the former — the price a respondent rationally *expects* to pay (a belief) — while the N arm captures the latter — the price that *feels* normal, the evaluative benchmark against which posted prices are judged; the distinction matters because the welfare-relevant reference in (7) is the internal benchmark measured by the N arm. Per-trip persistence estimates are compounded to the quarterly horizon using the respondent's reported shopping frequency. The exact survey screens that elicit this information in the two arms are reproduced in Appendix Figures A10 (expected arm) and A11 (normal arm).

2.4.4 Reference-price diagnostics and emotional response

After the scenario and updating questions, the survey closes with a battery of short descriptive items that characterize how the reference price is held and how respondents react to price changes. None of these items enters the confirmatory family (H1–H4); they are collected for sample description and descriptive correlates. These variables do not enter the confirmatory specifications; any analysis using them is exploratory unless explicitly listed in §9. The reference-price updating questions are described in §2 above and are not repeated here.

Item characteristics. Two single-item questions record properties of the good. A *perishability* indicator asks whether the item spoils within a short time (yes/no). A *shopping-frequency* question asks how often the respondent makes a trip to the usual store to buy the item, on a five-point scale from "about twice a week" to "less than once a month"; this frequency is the input that converts the per-trip persistence to a common horizon. The exact shopping-frequency screen is reproduced in Appendix Figure A8.

Reference-price diagnostics. Three items probe how the reference price is formed and used.

First, a *feeling* question asks whether a posted price equal to the elicited reference R_0 feels "too low," "about right," or "too high" for the item; a parallel feeling question is asked once the one-trip higher price in the updating scenario has persisted for a full month, recording whether the new, higher price has come to feel "about right." Second, a *salience* question asks how often the respondent has the expected price in mind while shopping ("never" to "always," five points). Third, a *formation* question asks the respondent to select all factors that influence the price they expect to pay — past shopping experiences, the store's regular shelf price, pre-trip information such as an app or website announcement, prices of similar items at the same store, prices of the same item at other stores, or no specific expected price in mind. Together these characterize the empirical inputs to the reference price R in (6) and how actively it is used. The exact survey screens that elicit this information — the reference-price feeling, formation, and salience questions — are reproduced in Appendix Figures A9, A13, and A14, and the post-shock feeling follow-up in Appendix Figure A12.

Emotional response. Finally, two check-all-that-apply items record the emotional reaction to a surprise price change presented with no explanation, with income and other prices held fixed. For a price *increase* of 20% above R_0, the options are: angry, annoyed, or frustrated; the price feels unfair (Kahneman et al., 1986; Rotemberg, 2005); no strong feeling / feels normal; an understanding that the store likely faces higher costs; or being happy to pay more because the respondent values the product. For a symmetric 20% *decrease*, the options are: happy, pleased, or excited; the price feels like a great deal; no strong feeling / feels normal; worry that the lower price will not last; or suspicion that the quality or package size may have changed. The surprise magnitude is fixed at 20% for every respondent, independent of the assigned scenario shift s. Because both items are asked *after* the scenario quantity questions, the emotional response is post-treatment relative to the quantity choices and cannot mechanically drive them. The exact survey screens that elicit feelings in response to the price increase and the price decrease are reproduced in Appendix Figures A15 and A16.

3 Pooling and the Combined Analysis Sample

The primary specification pools the new wave with the prior April 24–27 data into a single combined sample, with no wave indicator in the primary regressions. Pooling is justified ex ante on three grounds:

1. Identification is within-respondent (Equation (10)), so cross-wave composition differences should not bias $\hat{\Gamma}$ unless they correlate with the structural parameter.

2. The prior-wave analysis (Appendix in the companion report) showed demographic balance across the three pilot/wave-2 batches on age, sex, ethnicity, employment, student status, and lifetime Prolific approvals.

3. Any level difference between the samples shifts a respondent's baseline q_0, which the offset $\log q_0$ absorbs respondent by respondent, so it does not bias the within-respondent slope $\hat{\Gamma}$.

As a robustness check reported in an appendix — and not as part of the primary estimation strategy — we augment the pooled specification with a wave indicator $W_i \in \{0, 1\}$ (0 = prior waves, 1 =

new wave) and its interaction with the reference-price gap x, to test whether the two waves differ in level or in the estimated Γ. A significant wave term will be reported and discussed but will not, by itself, trigger a deviation from the pooled-primary framing.

4 Outcomes

Primary survey outcomes. The primary survey outcomes are the loss-scenario quantity q_L, the gain-scenario quantity q_G, the elicited loss-scenario reference price R_L, the elicited gain-scenario reference price R_G, and the post-shock updated reference price $\hat{R}$ elicited in either the expected-price arm (E) or the feels-normal arm (N).

Primary derived estimands. The primary derived estimands are the loss-side reference-dependence elasticity Γ_L, the gain-side reference-dependence elasticity Γ_G, the asymmetry $\Delta\Gamma = \Gamma_L - \Gamma_G$, and per-trip reference-price persistence ρ_{trip} by updating arm. The quarterly persistence parameters used in the calibration are derived by compounding respondent-level ρ_{trip} values to a quarterly horizon using the respondent's reported shopping frequency and averaging within arm.

Secondary and descriptive outcomes. Secondary and descriptive outcomes include the reference-price feeling response, the post-shock feeling response, reference-price salience, reported reference-price formation channels, perishability, shopping frequency, and emotional reactions to surprise price increases and decreases. These outcomes do not enter the confirmatory tests unless explicitly listed in the hypotheses below.

5 Hypotheses

Four confirmatory hypotheses are pre-registered. All other analyses (pooled-Γ and the parallel OLS/2SLS specification) are sensitivity or descriptive, not confirmatory.

H1 (Loss-side reference dependence exists). On the combined sample with the primary exclusion criteria, $\Gamma_L > 0$.

- Null: $\Gamma_L = 0$. Alternative: $\Gamma_L > 0$ (one-sided).
- Test: bootstrap-SE Wald test on Γ_L from the single-scenario loss specification (Equation (14) below), $\alpha = 0.05$. Prior-wave evidence used for design and precision planning is reported in §8.

H2 (Gain-side reference dependence exists). On the combined sample with the primary exclusion criteria, $\Gamma_G > 0$.

- Null: $\Gamma_G = 0$. Alternative: $\Gamma_G > 0$ (one-sided).
- Test: bootstrap-SE Wald test on Γ_G from the single-scenario gain specification (Equation (15) below), $\alpha = 0.05$. Prior-wave evidence used for design and precision planning is reported in §8.

H3 (Γ asymmetry). On the combined sample, the balanced asymmetric specification produces $\Delta\Gamma = \Gamma_L - \Gamma_G > 0$.

- Null: $\Delta\Gamma = 0$. Alternative: $\Delta\Gamma > 0$ (one-sided; prior-wave evidence used for design and precision planning is reported in §8).
- Test: bootstrap-SE Wald test on $\Delta\Gamma$ from the asymmetric balanced specification (Equation (16)), $\alpha = 0.05$.

H4 (Per-trip ρ is positive in both arms). Excluding respondents whose implied $\hat{\rho}$ lies outside $[0, 1]$, per-trip reference-price persistence is strictly positive in both arms.

- Null (per arm): $\bar{\rho}_{\text{trip}} = 0$. Alternative: $\bar{\rho}_{\text{trip}} > 0$.
- Test: one-sample t-test on per-respondent $\hat{\rho}_i$, separately for E and N arms, $\alpha = 0.05$ each. No asymmetry test (ρ^N vs. ρ^E) is pre-registered.

No formal multiple-testing adjustment is applied across H1–H4: each is a test on a conceptually distinct parameter. The pre-registered family is exactly these four tests; no additional hypotheses will be added to the confirmatory family after this PAP is filed.

6 Exclusion Restrictions

The confirmatory tests use the primary exclusion criteria from the main analysis, fixed in advance for this PAP. They combine a respondent-level exclusion with a row-level reference-price window.

Baseline respondent-level exclusions. A respondent is dropped from every analysis if they failed both attention checks, reported a baseline quantity above 50 units, a posted price outside $[0.50, 100]$ dollars, any scenario quantity above three times the baseline, or any reference-price response at or below one cent.

Reference-price window. On top of the respondent-level exclusions, the Γ estimation restricts elicited reference prices to a window anchored at the posted price P:

$$\text{loss scenario: } R_L \in [P\,(1 - 1.5s),\ 1.03\,P], \tag{11}$$

$$\text{gain scenario: } R_G \in [0.97\,P,\ P\,(1 + 1.5s)], \tag{12}$$

where $s \in \{0.12, 0.20, 0.28\}$ is the assigned shift. The outer edge is widened to $1.5\times$ the announced shock past P (admitting respondents who overshoot the window on the theory-consistent side); the inner edge carries a 3% tolerance to retain responses indicating essentially no change in the reference price on the wrong side of P without separately classifying them. Scenario domain is assigned by the randomized scenario, not by the sign of the respondent's elicited x_{ij}. Near-zero wrong-side responses within the 3% tolerance are retained with their observed x_{ij} values. This window is *fixed in advance* for this PAP as the primary criterion; in an appendix we report robustness to a less aggressive trim of theory-consistent overshooters, widening the outer edge up to $2.5\times$ the announced shock past P (in place of $1.5\times$) on the theory-consistent side, while holding the inner edge fixed at the 3% tolerance.

Reference-price persistence sample. For the ρ_{trip} analysis, the baseline exclusions are augmented by dropping respondents whose stated post-shock reference price lies below the original price (an internally inconsistent response), and respondents whose implied $\hat{\rho}_i$ lies outside $[0, 1]$ are excluded.

7 Primary Specifications

This section pre-registers the exact estimating equations used for H1–H4. All four specifications operate on the exclusion criteria of §6, pooled across prior waves and the new wave (§3).

7.1 First stage and control function

For each scenario $j \in \{L, G\}$, the endogenous regressor is the log reference–price gap $x_{ij} = \log(R_{ij}/P_i)$. The first-stage regression is OLS of x_{ij} on the announced log gap z_{ij} and a constant:

$$x_{ij} = \pi_0^j + \pi_1^j z_{ij} + \xi_{ij}, \qquad \hat{v}_{ij} = x_{ij} - \hat{x}_{ij}, \tag{13}$$

where $z_{L,i} = \log(p_{\text{down},i}/P_i)$ and $z_{G,i} = \log(p_{\text{up},i}/P_i)$. The residual $\hat{v}_{ij}$ enters the second-stage Poisson PML as a control function. First-stage uncertainty is propagated to second-stage standard errors via respondent-clustered bootstrap (1,000 replications).

The order indicator is $O_i = \mathbf{1}\{\text{loss scenario shown first}\}$. The primary (calibration) specifications below include neither a wave indicator nor an order interaction: the parameter used for calibration is the one *without* the order term, which is the theory-consistent object. We retain $\log q_{0,i}$ as a fixed offset (unit coefficient, no free intercept) throughout, because the structural model requires that at $R_i = P_i$ the scenario quantity equal the baseline $q_{0,i}$, which pins the intercept at $\log q_{0,i}$ rather than leaving it free. As robustness checks we report, in the main text, a version that adds the order interaction $\delta_O\, O_i\, \hat{v}_{ij}$, and, in an appendix, a version that adds a wave indicator $W_i \in \{0, 1\}$ (0 = prior waves, 1 = new wave) together with its interaction with x to test for differences between the two waves (§3); neither is part of the primary estimation strategy.

7.2 Single-scenario Γ_L (for H1)

Sample: respondents who pass the baseline exclusions and whose loss-scenario reference price falls within the window ($R_L \in [P(1 - 1.5s),\, 1.03P]$). One loss-scenario observation per respondent.

Estimator: Poisson PML, no intercept, $\log q_{0,i}$ as offset.

Specification (control function; primary/calibration):

$$\log \mathbb{E}[q_{L,i} \,|\, x_{L,i},\, \hat{v}_{L,i}] = \log q_{0,i} + \Gamma_L \cdot x_{L,i} + \delta_L\, \hat{v}_{L,i}. \tag{14}$$

The H1 test is a one-sided bootstrap-SE Wald test on Γ_L against the alternative $\Gamma_L > 0$.

7.3 Single-scenario Γ_G (for H2)

Sample: respondents who pass the baseline exclusions and whose gain-scenario reference price falls within the window ($R_G \in [0.97P,\, P(1 + 1.5s)]$). One gain-scenario observation per respondent.

Estimator: Poisson PML, no intercept, $\log q_{0,i}$ as offset.

Specification (control function; primary/calibration):

$$\log \mathbb{E}[q_{G,i} \,|\, x_{G,i},\, \hat{v}_{G,i}] = \log q_{0,i} + \Gamma_G \cdot x_{G,i} + \delta_G\, \hat{v}_{G,i}. \tag{15}$$

The H2 test is a one-sided bootstrap-SE Wald test on Γ_G against $\Gamma_G > 0$.

7.4 Asymmetric balanced specification (for H3)

Sample: respondents who pass the baseline exclusions and fall within the reference-price window on *both* scenarios. Two observations per respondent (one loss, one gain), stacked.

Estimator: Poisson PML, no intercept, $\log q_{0,i}$ as offset, respondent-clustered.

Specification (control function; primary/calibration): let i index respondents and $j \in \{L, G\}$ index scenarios; $D_{L,ij} = \mathbf{1}\{j = L\}$.

$$\log \mathbb{E}[q_{ij} \,|\, x_{ij},\, \hat{v}_{ij}] = \log q_{0,i} + \Gamma_G \cdot x_{ij} + \Delta\Gamma \cdot D_{L,ij}\, x_{ij} + \delta\, \hat{v}_{ij} + \delta_L\, D_{L,ij}\, \hat{v}_{ij}. \tag{16}$$

This is the balanced analog of the two single-scenario specifications above. The H3 test is a one-sided bootstrap-SE Wald test on $\Delta\Gamma$ against $\Delta\Gamma > 0$. Derived $\Gamma_L = \Gamma_G + \Delta\Gamma$ is reported with delta-method SE.

7.5 Per-trip ρ (for H4)

Sample: the reference-price persistence sample, excluding respondents whose implied $\hat{\rho}_i$ lies outside $[0, 1]$ (see §6).

Estimator: respondent-level closed-form,

$$\hat{\rho}_{\mathrm{trip},i} = \frac{\log \hat{R}_i - \log p_{\mathrm{shock},i}}{\log P_i - \log p_{\mathrm{shock},i}}. \tag{17}$$

The H4 test is a one-sample t-test of $\mathbb{E}[\hat{\rho}_i] > 0$, run separately within each framing arm (E and N), at $\alpha = 0.05$ each. H4 concerns per-trip persistence; the object used for calibration is instead quarterly persistence, formed by compounding each respondent's $\hat{\rho}_i$ to a quarterly horizon using their reported shopping frequency and averaging across respondents.

Estimation. All Γ specifications are estimated by Poisson pseudo-maximum-likelihood, with standard errors from a respondent-cluster bootstrap (1,000 replications).

8 Prior-Wave Evidence Used for Design and Precision Planning

The prior-wave point estimates below are reported as evidence used for design and precision planning. They are shown here—after the specifications (§7) and exclusion restrictions (§6) are committed—so that the pre-registered analysis choices are stated before the prior-wave estimates are reviewed. All numbers use the exclusion criteria of §6.

Table 1: Prior-wave estimates (April 2026 data), reported for design and precision planning. The Γ rows use the no-order control-function (CF) specification; ρ_{trip} rows report per-trip means restricted to $\hat{\rho} \in [0, 1]$, and ρ_{quarter} the calibrated quarterly persistence (per-trip $\hat{\rho}$ compounded to the quarterly horizon and averaged across respondents).

Parameter	Estimate	SE	Sample
Γ_L (single-scenario)	0.891^{***}	0.114	Loss scenario, $N = 1{,}085$
Γ_G (single-scenario)	0.285^{**}	0.121	Gain scenario, $N = 1{,}073$
$\Delta\Gamma$ (balanced asymmetric)	0.557^{***}	0.200	Balanced, $N_{\text{resp}} = 995$
$\bar{\rho}^E_{\text{trip}}$ (E arm, $\hat{\rho} \in [0, 1]$)	0.328	—	$N = 571$
$\bar{\rho}^N_{\text{trip}}$ (N arm, $\hat{\rho} \in [0, 1]$)	0.745	—	$N = 554$
ρ^E_{quarter} (E arm, calibrated)	0.22	—	—
ρ^N_{quarter} (N arm, calibrated)	0.55	—	—

Loss-side reference dependence is roughly three times the gain-side magnitude, and the asymmetry $\Delta\Gamma$ is significant at the 1% level; the expected-price arm (E) has lower persistence than the normal-price arm (N), at both the per-trip and the calibrated quarterly horizon. Full numerical detail and the complete set of tables are in the companion analysis report; this PAP does not re-derive that material.

9 Secondary Specifications

A standard set of sensitivity and descriptive tables, all already implemented in the analysis workflow, will be reported alongside the primary specifications: a pooled-Γ regression (single coefficient on x with no L/G asymmetry); an OLS/2SLS counterpart to the full set of tables (with the $q > 0$ restriction required by $\log(q/q_0)$); and a demographics-by-wave balance table extended to include the new wave. All regressions follow the structural model in retaining no intercept, consistent with the prediction that at $R = P$ both scenarios yield the same baseline. None of these are part of the confirmatory family (H1–H4); they are reported for transparency. Analyses using the diagnostic, item-characteristic, or emotional-response variables are exploratory unless explicitly listed in this section.

10 Planned Use in the Welfare Calibration

The welfare simulations are not part of the confirmatory survey-experiment hypothesis family. They will be updated mechanically using the pre-specified survey estimates. For inflationary price increases, the baseline calibration uses the primary PPML-CF estimate of Γ_L, because posted prices are above reference prices in the loss domain. The gain-side estimate Γ_G is reported to document asymmetry and is the relevant slope for price decreases, disinflation, or deflation episodes, but it is not the baseline slope for the inflationary reference-dependence wedge. The calibration will report results using the quarterly persistence estimates derived from the expected-price and feels-normal arms; the feels-normal arm is the main measure of the internal evaluative benchmark used in the welfare-relevant reference price.

11 Deviations

Any deviation from the specifications in this PAP will be reported transparently in a Deviations section of the final paper, including both the PAP specification and the deviated specification side-by-side. Deviations are anticipated only for issues that could not be foreseen at the time of pre-registration (e.g., a software bug discovered during estimation, a Prolific recruitment failure that requires re-targeting).

A Survey Instrument

This appendix reproduces each survey screen in the order respondents encounter them. Throughout, the respondent's own product, package size, store, and elicited reference price are piped into the question text; the screens below show these fields filled in for one example respondent (Clover Organic Whole Milk, half-gallon, at Target, with $R_0 = \$4.99$).

Berkeley
UNIVERSITY OF CALIFORNIA

Please think of one everyday item that you purchase frequently. If possible, provide the most specific description you can (e.g., Kirkland Signature Organic Whole Milk). If you do not remember exact details, a general description that clearly identifies the item is fine (e.g., "Organic Whole Milk").

For the package/unit size, your best estimate is sufficient (e.g., "1 gallon," "a standard dozen," "a 1-liter bottle"). If you purchase this item from multiple stores, please enter the store where you buy it most often.

Please keep the same item and unit size in mind for the rest of the survey.

Product / Item Description	Clover Organic Whole Milk
Package size / Unit size (the size you typically purchase)	Half-Gallon
Store (e.g., Target, Walmart, Local Grocery Store)	Target

Figure A1: Setup: selection of the product, package/unit size, and usual store. All later questions pipe these values into the question text.

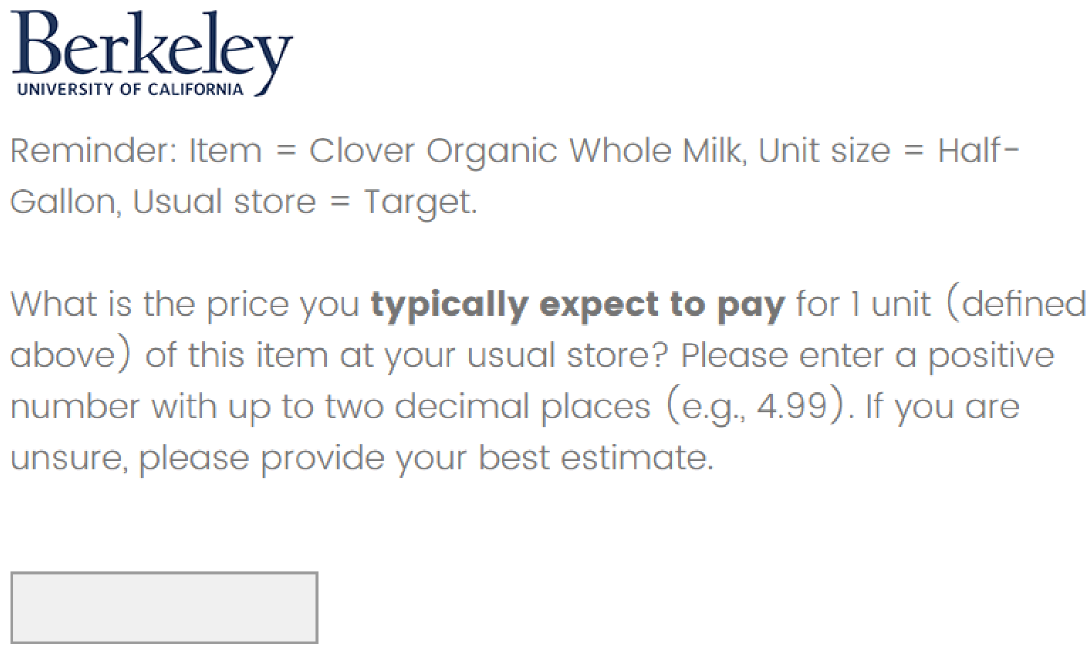

Figure A2: Reference-price elicitation: the per-unit price the respondent typically expects to pay for the item at the usual store (R_0).

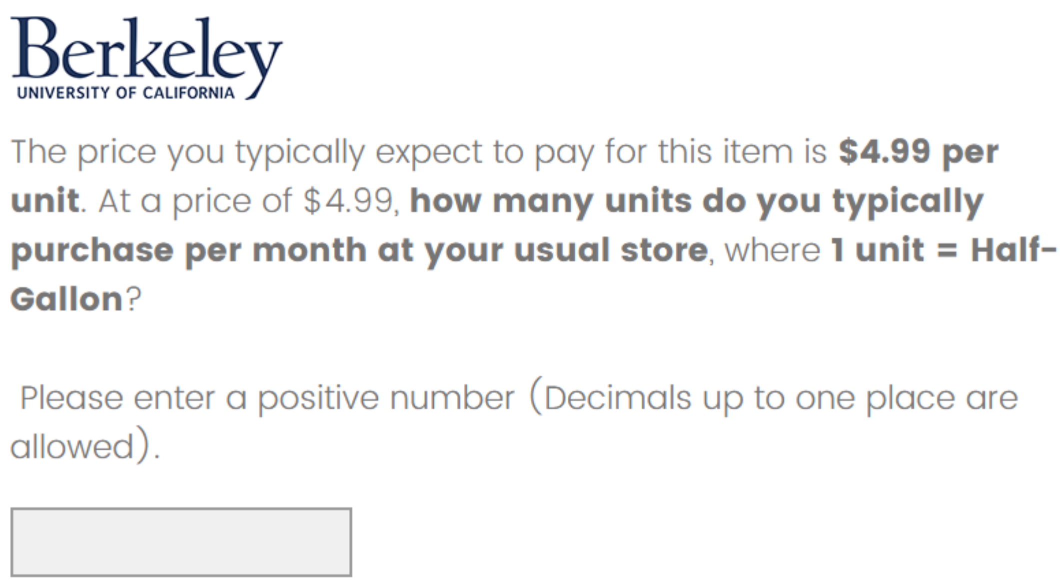

Figure A3: Baseline quantity: units purchased per month at the reference price R_0 (q_0, the offset in the estimating equations).

Berkeley
UNIVERSITY OF CALIFORNIA

Reminder: Item = Clover Organic Whole Milk, Unit size = Half-Gallon, Usual store = Target.

Assume that at your usual store (Target), the price of this item **is usually $3.99 per unit.**

Next month only, the price will temporarily increase to **$4.99** per unit.

After next month, the price will return to $3.99 per unit. Assume your income and the prices of other goods stay the same. Please answer based only on how much you would want to buy for next month's use at that price. Ignore any possibility of buying extra to use later or delaying purchases to a later month.

Recall that one unit is: Half-Gallon.

How many units would you buy next month at **$4.99 per unit at your usual store**? (Enter a number up to one decimal place; 0 is allowed.)

Figure A4: Loss scenario, quantity question: a usual price of $(1-s)R_0$, a temporary one-month increase to the posted price R_0, and the units the respondent would buy that month.

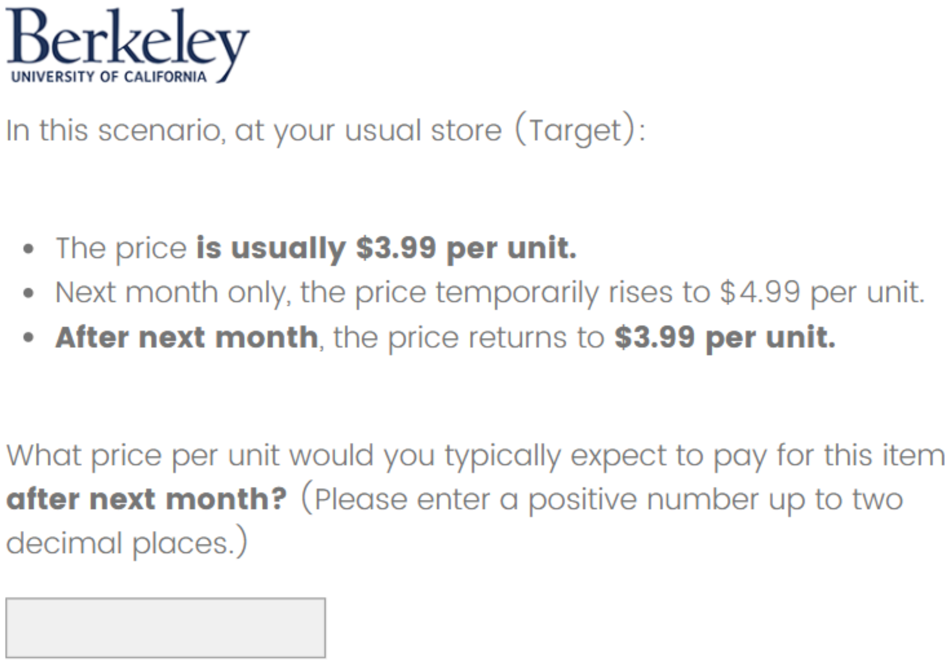
Berkeley
UNIVERSITY OF CALIFORNIA

In this scenario, at your usual store (Target):

- The price **is usually $3.99 per unit.**
- Next month only, the price temporarily rises to $4.99 per unit.
- **After next month**, the price returns to **$3.99 per unit.**

What price per unit would you typically expect to pay for this item **after next month?** (Please enter a positive number up to two decimal places.)

Figure A5: Loss scenario, elicited reference: the per-unit price the respondent would expect to pay going forward once the usual price resumes (R_L).

Berkeley
UNIVERSITY OF CALIFORNIA

Reminder: Item = Clover Organic Whole Milk, Unit size = Half-Gallon, Usual store = Target.

Assume that at your usual store (Target), the price of this item **is usually $5.99 per unit.**

Next month only, the item will be on sale for **$4.99** per unit.

After next month, the price will return to $5.99 per unit. Assume your income and the prices of other goods stay the same. Please answer based only on how much you would want to buy for next month's use at that price. Ignore any possibility of buying extra to use later or delaying purchases to a later month.

Recall that one unit is: Half-Gallon.

How many units would you buy next month at **$4.99 per unit at your usual store**? (Enter a number up to one decimal place; 0 is allowed.)

Figure A6: Gain scenario, quantity question: a usual price of $(1+s)R_0$, a temporary one-month sale to the posted price R_0, and the units the respondent would buy that month.

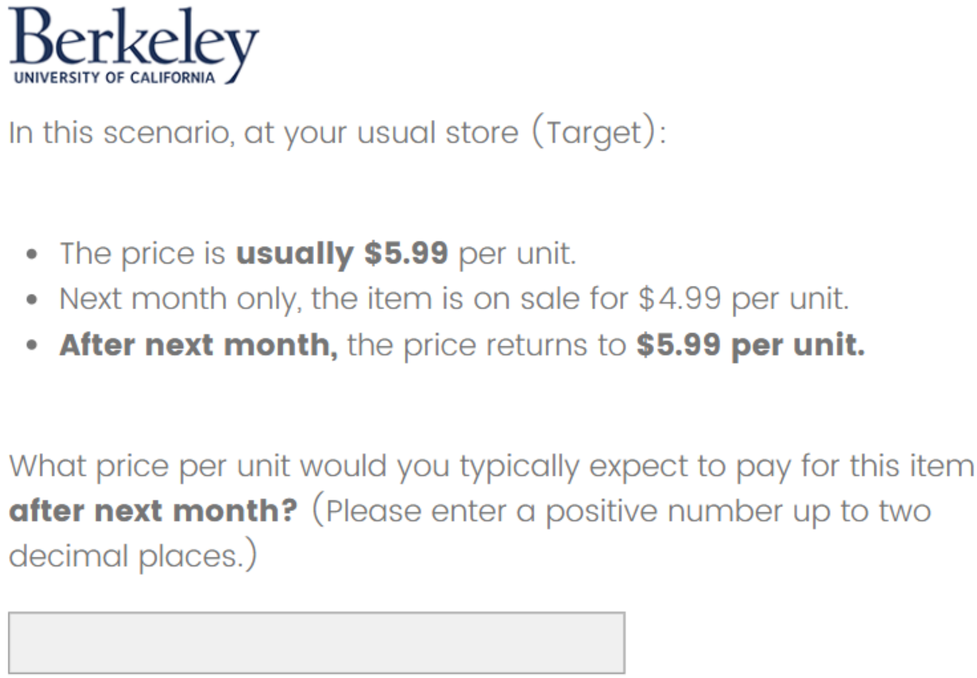

Berkeley
UNIVERSITY OF CALIFORNIA

In this scenario, at your usual store (Target):

- The price is **usually $5.99** per unit.
- Next month only, the item is on sale for $4.99 per unit.
- **After next month,** the price returns to **$5.99 per unit.**

What price per unit would you typically expect to pay for this item **after next month?** (Please enter a positive number up to two decimal places.)

Figure A7: Gain scenario, elicited reference: the per-unit price the respondent would expect to pay going forward once the usual price resumes (R_G).

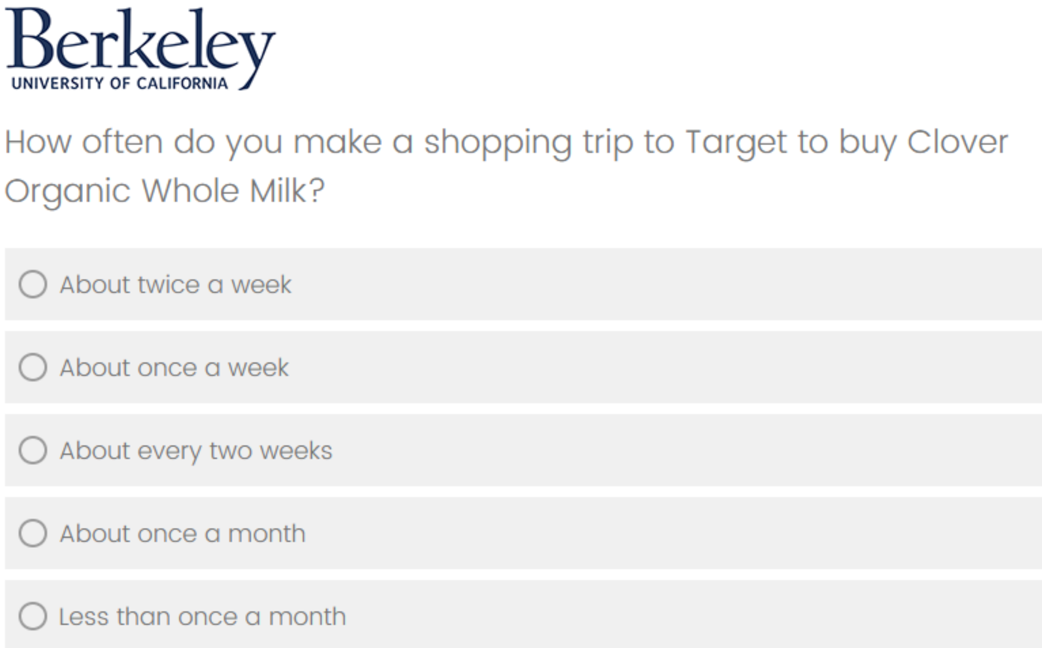

Berkeley
UNIVERSITY OF CALIFORNIA

How often do you make a shopping trip to Target to buy Clover Organic Whole Milk?

- About twice a week
- About once a week
- About every two weeks
- About once a month
- Less than once a month

Figure A8: Shopping frequency: how often the respondent makes a trip to the usual store to buy the item.

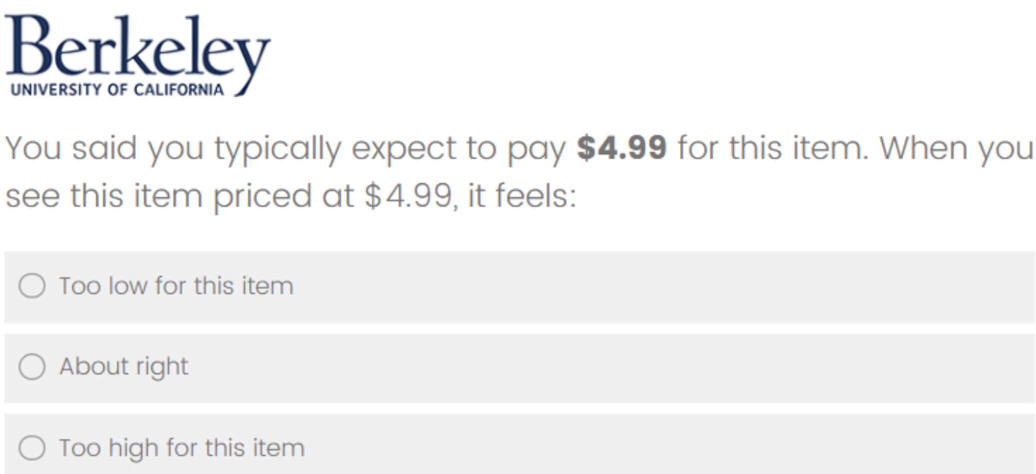

Berkeley
UNIVERSITY OF CALIFORNIA

You said you typically expect to pay **$4.99** for this item. When you see this item priced at $4.99, it feels:

- Too low for this item
- About right
- Too high for this item

Figure A9: Reference-price feeling: whether a posted price equal to R_0 feels too low, about right, or too high for the item.

Berkeley
UNIVERSITY OF CALIFORNIA

Suppose that on your previous shopping trips, the price of Clover Organic Whole Milk was **$4.99** per unit.

Assume that on your **next shopping trip** to Target, the price of Clover Organic Whole Milk is **$5.99** per unit.

After seeing the higher price on that one trip**, what price per unit would you typically expect to pay for this item going forward** at Target? (Please enter a positive number with up to two decimal places.)

Figure A10: Reference-price persistence elicitation, expected arm (E): the price the respondent would expect to pay going forward after one trip at a price 20% above R_0.

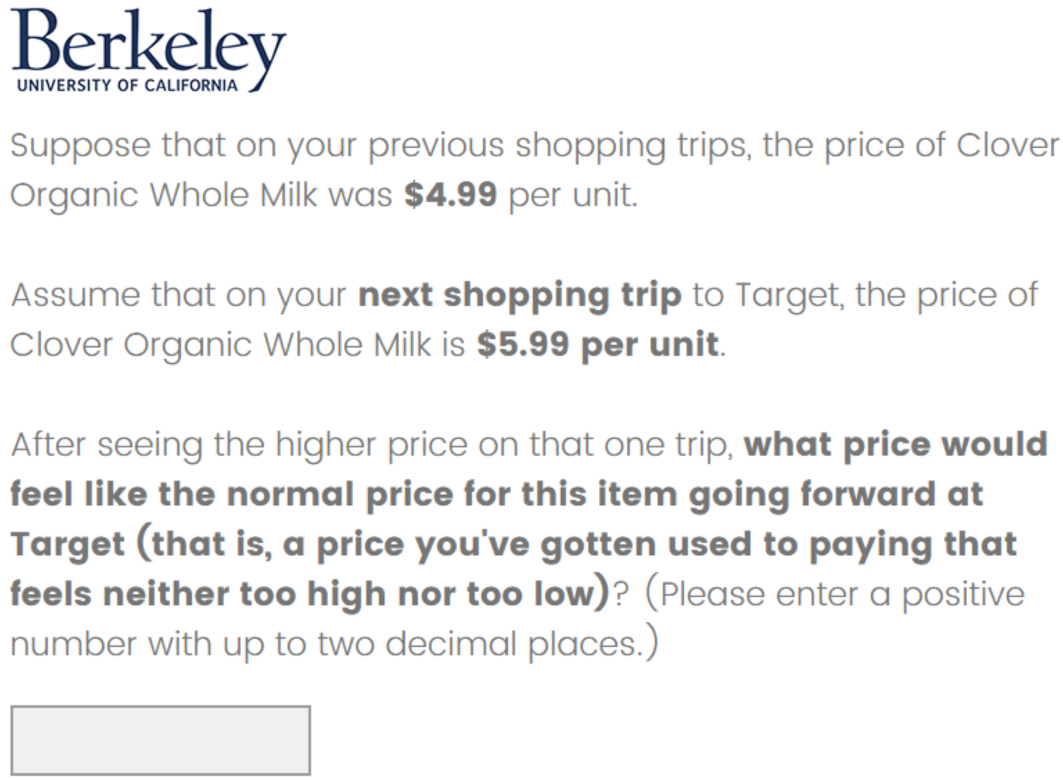
Berkeley
UNIVERSITY OF CALIFORNIA

Suppose that on your previous shopping trips, the price of Clover Organic Whole Milk was **$4.99** per unit.

Assume that on your **next shopping trip** to Target, the price of Clover Organic Whole Milk is **$5.99 per unit**.

After seeing the higher price on that one trip, **what price would feel like the normal price for this item going forward at Target (that is, a price you've gotten used to paying that feels neither too high nor too low)**? (Please enter a positive number with up to two decimal places.)

Figure A11: Reference-price persistence elicitation, normal arm (N): the price that would feel normal going forward after one trip at a price 20% above R_0.

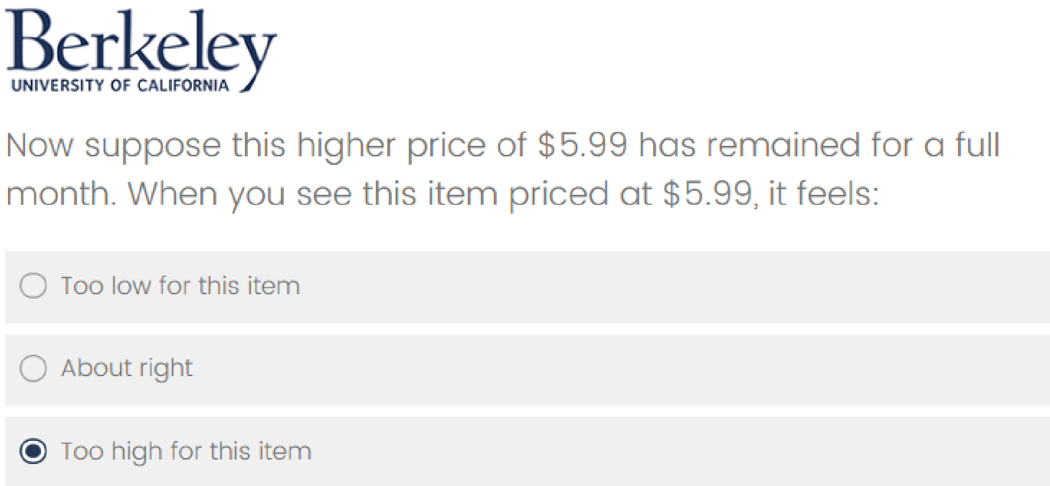

Berkeley
UNIVERSITY OF CALIFORNIA

Now suppose this higher price of $5.99 has remained for a full month. When you see this item priced at $5.99, it feels:

- ○ Too low for this item
- ○ About right
- ◉ Too high for this item

Figure A12: Post-shock feeling follow-up: whether, once the higher price has persisted for a full month, it feels too low, about right, or too high.

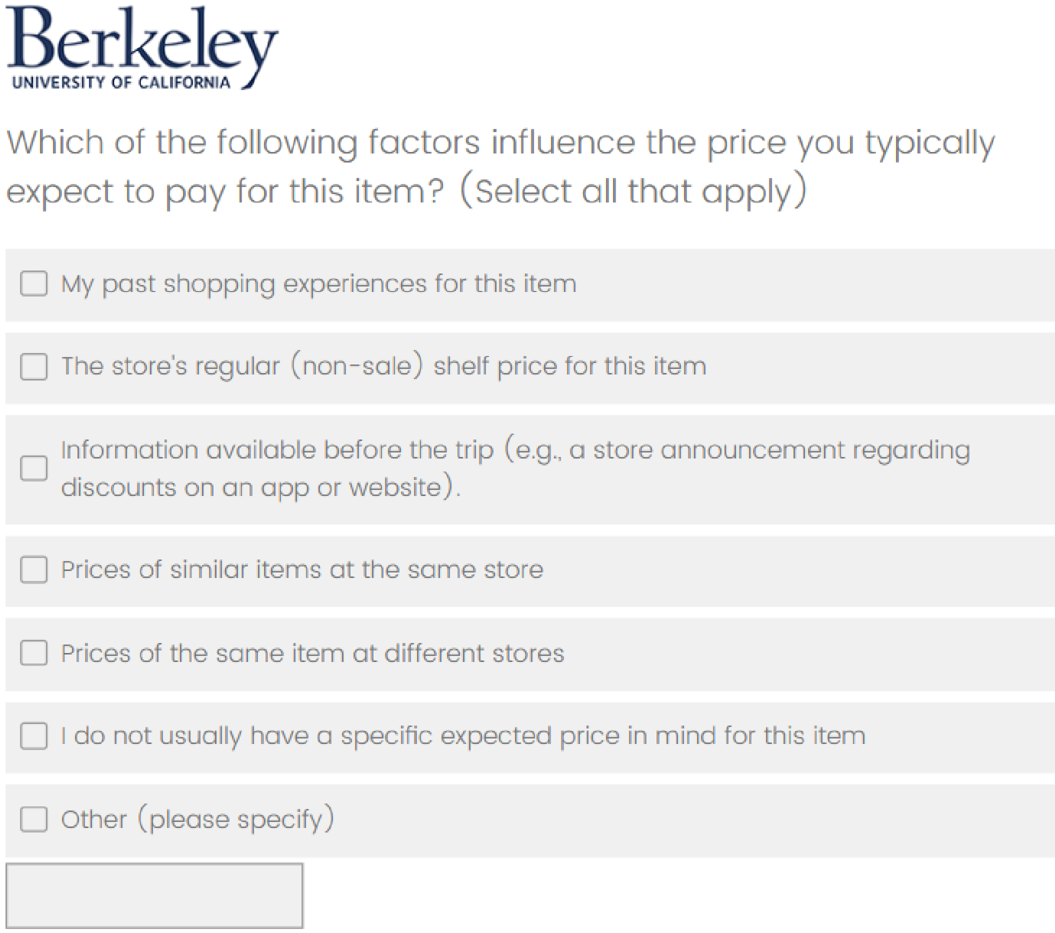

Berkeley
UNIVERSITY OF CALIFORNIA

Which of the following factors influence the price you typically expect to pay for this item? (Select all that apply)

- ☐ My past shopping experiences for this item
- ☐ The store's regular (non-sale) shelf price for this item
- ☐ Information available before the trip (e.g., a store announcement regarding discounts on an app or website).
- ☐ Prices of similar items at the same store
- ☐ Prices of the same item at different stores
- ☐ I do not usually have a specific expected price in mind for this item
- ☐ Other (please specify)

Figure A13: Reference-price formation: the factors the respondent reports as influencing the price they expect to pay (select all that apply).

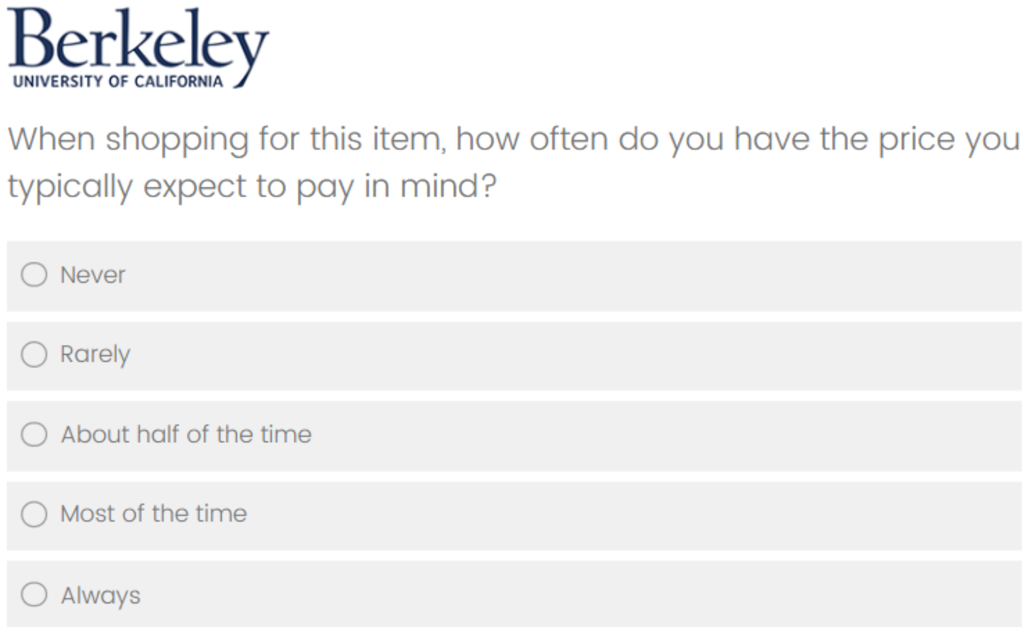
Berkeley
UNIVERSITY OF CALIFORNIA

When shopping for this item, how often do you have the price you typically expect to pay in mind?

- Never
- Rarely
- About half of the time
- Most of the time
- Always

Figure A14: Reference-price salience: how often the respondent has the expected price in mind while shopping.

Berkeley
UNIVERSITY OF CALIFORNIA

Reminder: Item = Clover Organic Whole Milk, Unit size = Half-Gallon, Usual store = Target.

The price per unit you typically expect to pay for this item is **$4.99.**

Now imagine you go to your usual store to purchase this item, and **the price today is: $5.99,** with no explanation provided. Assume your income and the prices of other goods are unchanged.

Which of the following reactions come to mind? (Check all that apply)

- [ ] Angry, annoyed or frustrated
- [ ] This price feels unfair
- [ ] No strong feeling/this feels normal
- [ ] I understand, the store likely has higher costs
- [ ] I am happy to pay more since I value this product
- [ ] Other (please specify)

Figure A15: Emotional response to a surprise price increase: reactions to paying 20% above R_0 with no explanation (check all that apply).

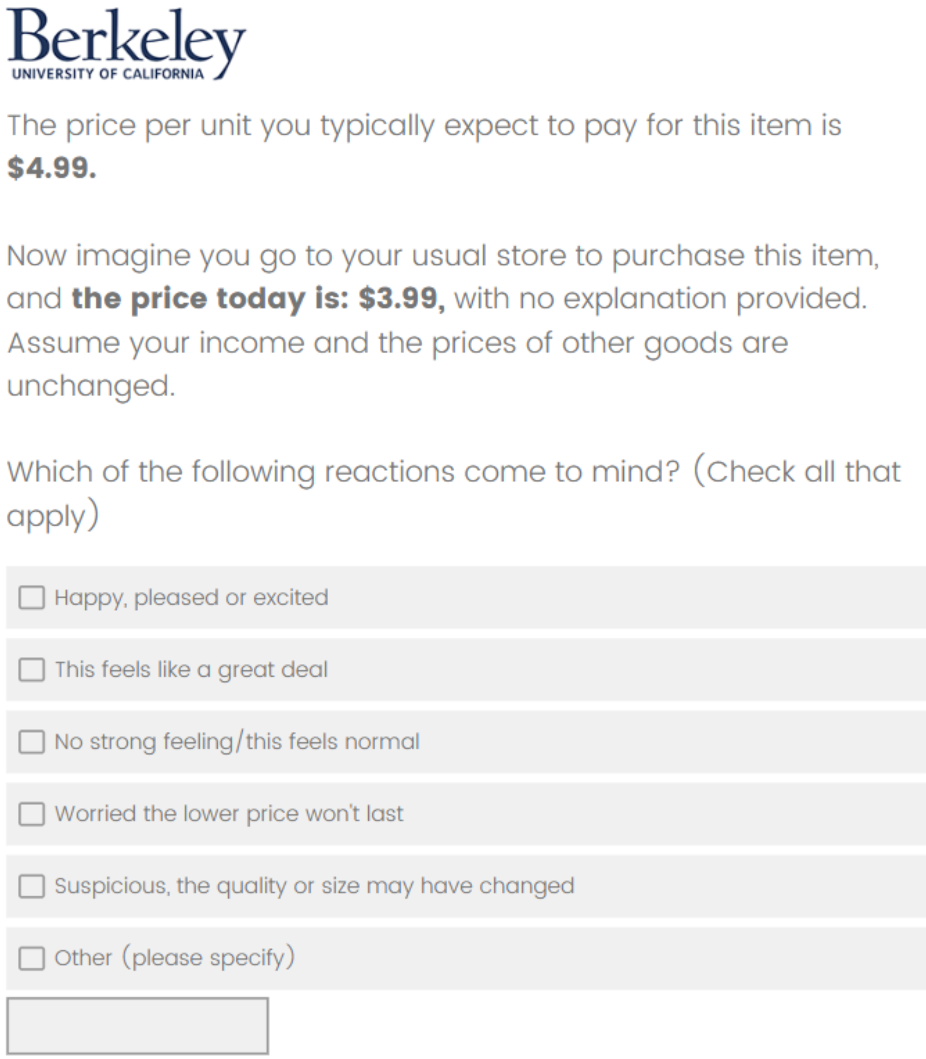

Berkeley
UNIVERSITY OF CALIFORNIA

The price per unit you typically expect to pay for this item is **$4.99.**

Now imagine you go to your usual store to purchase this item, and **the price today is: $3.99,** with no explanation provided. Assume your income and the prices of other goods are unchanged.

Which of the following reactions come to mind? (Check all that apply)

- [ ] Happy, pleased or excited
- [ ] This feels like a great deal
- [ ] No strong feeling/this feels normal
- [ ] Worried the lower price won't last
- [ ] Suspicious, the quality or size may have changed
- [ ] Other (please specify)

Figure A16: Emotional response to a surprise price decrease: reactions to paying 20% below R_0 with no explanation (check all that apply).

References

Hassan Afrouzi, Andrés Blanco, Andrés Drenik, and Erik Hurst. A theory of how workers keep up with inflation. Working Paper 33233, National Bureau of Economic Research, December 2024. URL http://www.nber.org/papers/w33233.

Simon P Anderson, André De Palma, and Jacques-François Thisse. Demand for differentiated products, discrete choice models, and the characteristics approach. *The Review of Economic Studies*, 56(1):21–35, 1989.

Georgios Angelis. Price setting and price stickiness: A behavioral foundation of inaction bands. *Journal of the European Economic Association*, 23(3):1095–1132, 2025.

Alberto Binetti, Francesco Nuzzi, and Stefanie Stantcheva. People's understanding of inflation. *Journal of Monetary Economics*, 148:103652, 2024. ISSN 0304-3932. doi: https://doi.org/10.1016/j.jmoneco.2024.103652. URL https://www.sciencedirect.com/science/article/pii/S0304393224001053. Inflation in the COVID Era and Beyond.

Felipe N Del Canto, John R Grigsby, Eric Qian, and Conor Walsh. Are inflationary shocks regressive? a feasible set approach. Working Paper 31124, National Bureau of Economic Research, April 2023. URL http://www.nber.org/papers/w31124.

Stanley Fischer. Towards an understanding of the costs of inflation: Ii. *Carnegie-Rochester Conference Series on Public Policy*, 15:5–41, 1981. ISSN 0167-2231. doi: https://doi.org/10.1016/0167-2231(81)90016-6. URL https://www.sciencedirect.com/science/article/pii/0167223181900166.

Jordi Gali. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications Second edition*, volume None of *Economics Books*. Princeton University Press, 2 edition, None 2015. doi: None. URL https://ideas.repec.org/b/pup/pbooks/10495.html.

Dimitris Georgarakos, Kwang Hwan Kim, Olivier Coibion, Myungkyu Shim, Myunghwan Andrew Lee, Yuriy Gorodnichenko, Geoff Kenny, Seowoo Han, and Michael Weber. How costly are business cycle volatility and inflation? a vox populi approach. Technical report, National Bureau of Economic Research, 2025.

Joao Guerreiro, Jonathon Hazell, Chen Lian, and Christina Patterson. Why do workers dislike inflation?: Wage erosion and conflict costs. 2024.

Paul Heidhues and Botond Kőszegi. Competition and price variation when consumers are loss averse. *American Economic Review*, 98(4):1245–1268, 2008.

Daniel Kahneman, Jack L Knetsch, and Richard Thaler. Fairness as a constraint on profit seeking: Entitlements in the market. *The American economic review*, pages 728–741, 1986.

Robert E. Lucas. Discussion of : Stanley fischer, "towards an understanding of the costs of inflation: Ii". *Carnegie-Rochester Conference Series on Public Policy*, 15:43–52, 1981. ISSN 0167-2231. doi: https://doi.org/10.1016/0167-2231(81)90017-8. URL https://www.sciencedirect.com/science/article/pii/0167223181900178.

Robert E. Lucas. Inflation and welfare. *Econometrica*, 68(2):247–274, 2000. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/2999427.

Tridib Mazumdar, Sevilimedu P Raj, and Indrajit Sinha. Reference price research: Review and propositions. *Journal of marketing*, 69(4):84–102, 2005.

Emi Nakamura, Jón Steinsson, Patrick Sun, and Daniel Villar. The elusive costs of inflation: Price dispersion during the u.s. great inflation*. *The Quarterly Journal of Economics*, 133(4):1933–1980, 08 2018. ISSN 0033-5533. doi: 10.1093/qje/qjy017. URL https://doi.org/10.1093/qje/qjy017.

Michaela Pagel. Expectations-based reference-dependent life-cycle consumption. *The Review of Economic Studies*, 84(2):885–934, 2017.

Filippo Pallotti, Gonzalo Paz-Pardo, Jiri Slacalek, Oreste Tristani, and Giovanni L Violante. Who bears the costs of inflation? euro area households and the 2021–2023 shock. *Journal of Monetary Economics*, 148:103671, 2024.

Julio J Rotemberg. Customer anger at price increases, changes in the frequency of price adjustment and monetary policy. *Journal of monetary economics*, 52(4):829–852, 2005.

Emiliano Santoro, Ivan Petrella, Damjan Pfajfar, and Edoardo Gaffeo. Loss aversion and the asymmetric transmission of monetary policy. *Journal of Monetary Economics*, 68:19–36, 2014.

Robert J. Shiller. Why Do People Dislike Inflation? In *Reducing Inflation: Motivation and Strategy*, NBER Chapters, pages 13–70. National Bureau of Economic Research, Inc, April 1997. URL https://ideas.repec.org/h/nbr/nberch/8881.html.

Stefanie Stantcheva. Why Do We Dislike Inflation? NBER Working Papers 32300, National Bureau of Economic Research, Inc, April 2024. URL https://ideas.repec.org/p/nbr/nberwo/32300.html.

Richard Thaler. Mental accounting and consumer choice. *Marketing science*, 4(3):199–214, 1985.

Manoj Thomas and Geeta Menon. When internal reference prices and price expectations diverge: The role of confidence. *Journal of Marketing Research*, 44(3):401–409, 2007.

Amos Tversky and Daniel Kahneman. Loss aversion in riskless choice: A reference-dependent model. *The Quarterly Journal of Economics*, 106(4):1039–1061, 1991. ISSN 00335533, 15314650. URL http://www.jstor.org/stable/2937956.