Finance and Economics Discussion Series

Federal Reserve Board, Washington, D.C.
ISSN 1936-2854 (Print)
ISSN 2767-3898 (Online)

Paying More and Buying Less: 2025 Tariffs and U.S. Household Spending

Sinem Hacioglu-Hoke, Leo Feler

2026-035

Please cite this paper as:
Hacioglu-Hoke, Sinem, and Leo Feler (2026). "Paying More and Buying Less: 2025 Tariffs and U.S. Household Spending," Finance and Economics Discussion Series 2026-035. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2026.035.

Paying More and Buying Less: 2025 Tariffs and U.S. Household Spending

Sinem Hacıoğlu-Hoke *
Federal Reserve Board
CEPR

Leo Feler†
Visiting Economist
Federal Reserve Bank of Chicago

May 21, 2026

Abstract

This paper estimates the effects of the 2025 U.S. tariffs on household spending using transaction-level data linked to tariff exposure and a tariff sentiment survey. Comparing high versus low tariff-exposed categories, we find 15 to 20 percent price pass-through. At the mean increase in tariff exposure, prices rise by 1 to 2 percent while spending falls by roughly 4 percent. Survey evidence linking stated intentions to revealed behavior identifies a mechanism for the large spending response: reallocation toward essentials and trade-down within categories, concentrated among middle-income households with discretionary slack who express tariff concerns. Low-income households bear a disproportionate welfare burden through regressive pass-through.

Keywords: Tariffs, Pass through, Prices, Household Spending, Welfare, Uncertainty

JEL Classification: F13, D12, E31

*2001 C St NW, Washington, DC 20037. Email: sinem.haciogluhoke@frb.gov Web: www.sinemhaciogluhoke.com

†Email: feler.leo@gmail.com Web: https://leofeler.com/

We are grateful to Fariha Kamal for the guidance on the trade data and to Samantha Mitchell, Jack Chylak, and Amanda Schoenbauer for help with the survey. Mason Thieu provided excellent research assistance. We also thank Aaron Flaaen, Colin Hottman, Chris Kurz, Robbie Minton, Ekaterina Peneva, Erick Sager, Daniel Villar Vallenas, as well as the participants of the SNDE 2026 conference and Federal Reserve Board seminar for their comments and feedback. Some of this work was done while Leo Feler was the chief economist of Numerator. Artificial intelligence was used in text editing and coding, but all AI-generated text and code were thoroughly verified by the authors. The views expressed in this paper are the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System.

1 Introduction

In early 2025, the United States imposed broad-based tariffs covering virtually all of its trading partners, leading to one of the most substantial shifts in U.S. trade policy in decades. How much of this shock households face or feel and how they respond has become a central question for assessing its implications. A growing literature has made progress on the price side, documenting that pass-through to retail prices has so far been partial and gradual (Cavallo et al., 2025; Minton and Somale, 2025; Dvorkin et al., 2025; Barbiero and Stein, 2025; Hacıoğlu-Hoke et al., 2026). Considerably less is known about the subsequent margin of adjustment as to whether households continue to purchase the same goods at higher prices, whether they reduce expenditure, and whether the response varies systematically across the income distribution.

This paper takes up these questions using transaction-level data on over 125,000 U.S. households observed continuously from January 2024 through December 2025. The data come from the Numerator household panel and record the price, quantity, and product attributes for household purchases from brick-and-mortar and online retailers and also contain a rich set of household characteristics including income. The consumption categories that dominate the panel, including groceries, health and beauty, apparel, home goods, and electronics, are precisely those that are most directly affected by tariffs on traded goods. They are also the categories households weight most heavily when forming inflation expectations and adjusting spending decisions (D'Acunto et al., 2021; Weber et al., 2022), which makes them particularly salient for the questions we explore. The central empirical advantage of the data is the joint observation of prices and quantities at the household-product level. This allows us to estimate the price, spending, and quantity responses to tariffs from a single source. For a subset of approximately 21,000 panel households, we further link the transaction records to a tariff sentiment survey fielded by Numerator, which records tariff awareness, concern about price increases, and stated behavioral intentions. The survey allows us to connect household beliefs to spending behavior and to distinguish price-driven responses from precautionary responses to tariffs.

We link each purchase to a measure of tariff exposure constructed at the 4-digit Harmonized System level from U.S. Census monthly merchandise trade data, combining realized effective tariff rates with product-category import penetration. Our empirical strategy is a difference-in-differences design that compares categories with higher tariff exposure to categories with lower exposure before and after tariffs were implemented in February 2025. We estimate the same specification on three outcomes: a Fisher price index constructed from the household data, real spending, and quantities purchased.

We document three main findings. The first is that tariffs translate into substantially larger movements in spending than in prices. We report pass-through under two complementary treatment definitions that answer related but distinct questions. Using the realized tariff rate directly, which weights all product categories equally regardless of their import dependence, a one percentage point increase in the tariff rate raises retail prices by approximately 0.15 percentage points, implying retail pass-through of around 15 percent. Our benchmark specification instead scales the tariff rate by import penetration to isolate categories where the tariff is actually borne by domestic consumers, and yields a larger pass-through estimate of roughly 20 percent. Both estimates are broadly consistent with the partial price pass-through documented for the 2025 episode (Cavallo et al., 2025; Minton and Somale, 2025; Dvorkin et al., 2025; Hacıoğlu-Hoke et al., 2026). On the other hand, a one percentage point increase in tariff exposure reduces real household spending on affected categories by approximately 0.6 percentage points. Evaluated at the mean increase in tariff exposure over the sample period, these estimates imply that average prices in an affected consumption category rise by approximately 1 to 2 percent, depending on the specification, while the corresponding spending decline is roughly 4 percent, a contraction three to four times larger than the implied price increase.

Mechanically, the wedge between prices and spending is difficult to explain by households simply cutting back in response to higher prices. This observation is consistent with households curtailing purchases of tariff-exposed goods beyond what the price change would predict, a channel recently emphasized as operating through higher tariff-rate uncertainty in Candia et al. (2026). Moreover, den Besten et al. (2026) document that tariffs depress real output and our household-level evidence on spending contractions offers a micro channel consistent with that aggregate finding. We also find no evidence of stockpiling as an explanation. The spending contraction is as large in the steady state as in the first post-tariff months, and stated stockpilers reduce rather than accelerate their purchases, consistent with anxiety about tariffs dominating any anticipatory buying motive.

The second finding concerns the distribution of the burden across income groups. Low-income households bear a disproportionate price burden. Tariffs operate through two distinct channels with different income profiles. The first is a price-incidence channel that dominates at the bottom of the distribution. The price response to tariff exposure is monotonically decreasing in income. When these differences are combined into welfare terms, low-income households bear losses of approximately 0.19 percent of income compared with 0.02 percent for high-income households, creating a gap disproportionally larger for low-income households. We compute this cost using two complementary methods: a standard consumer surplus formula (Harberger deadweight loss) and an exact cost-of-living index (Sato-Vartia) that uses observed spending reallocation rather than an assumed demand elasticity.

The third finding uses survey evidence to ask why spending falls by so much more than prices rise. While the aggregate results document the fact, the survey identifies who drives it and how. All income groups contract spending on tariff-exposed goods by similar magnitudes, but the margin of adjustment differs. Within categories that households flag as tariff-affected, trade-down to cheaper varieties is common across the entire income distribution: spending falls while quantities hold, a pattern that does not vary with income. The margins that do vary are across categories. Middle-income households additionally cut the level of non-essential spending, driving the basket-composition shift toward essentials, a reallocation identified from the differential response across tariff exposure categories and confirmed by linking stated intentions to revealed behavior. Low-income households absorb the tariff primarily through higher prices on continuing purchases, with limited room to restructure the composition of their spending. The two channels thus differ not in the total magnitude of the spending response but in its composition. The price channel is regressive in welfare terms, the precautionary channel is a reallocation story concentrated where households have discretionary room to reweight their baskets. We use the term ‘precautionary’ throughout to refer to this defensive reallocation of the basket under uncertainty, rather than to the aggregate savings response that it usually stands for in the macro literature.

Related literature and contribution. Our paper relates to several strands of work on the effects of trade policy.

A first strand studies tariff pass-through. Work on the 2018–19 trade war established near-complete pass-through to import prices (Amiti et al., 2019, 2020; Fajgelbaum et al., 2020) and examined transmission to retail prices (Cavallo et al., 2021a). For the 2025 episode, a growing body of work using online prices, PCE data, firm surveys, and scanner data finds retail pass-through that is partial and gradual (Cavallo et al., 2025; Minton and Somale, 2025; Dvorkin et al., 2025; Barbiero and Stein, 2025; Hacıoğlu-Hoke et al., 2026) and our estimates sit within this range. The modest percentage pass-through we find partly reflects data composition toward everyday consumer goods and two conceptual points in the literature. First, Sangani (2026) shows that incomplete pass-through in percentages can disguise complete pass-through in levels, and second, Ganapati and Hottman (2026) show that accounting for endogenous quantity responses reduces measured pass-through to around 60 percent. Our estimates should therefore be interpreted as quantity-attenuated percentage pass-through.

A second strand examines household consumption responses to trade shocks. Waugh (2019) documents that U.S. counties exposed to Chinese retaliatory tariffs experienced large consumption declines, operating through the labor market channel as lost export access reduced local employment and spending. Ma et al. (2025) and DeDad and Ghosh (2026) use NielsenIQ scanner data for the same episode, with Ma et al. (2025) showing that wealthier households' greater substitution capacity drives distributional differences in welfare losses. We provide among the first household-level estimates of price, spending, and quantity responses jointly identified from the same data for the 2025 episode, and link the panel to a survey that measures tariff awareness, sentiment, and behavioral intentions. This allows us to separate price incidence from precautionary contraction in a way not possible with scanner data alone. In contrast to the substitution-capacity story in Ma et al. (2025), our evidence points toward differential access to discretionary adjustment such that low-income households are distinguished not by substitution technology but by lack of slack.

A third strand studies how trade affects household welfare through the consumption basket. Fajgelbaum and Khandelwal (2016) show that trade liberalization disproportionately benefits lower-income households, Acosta and Cox (2024) provide evidence that U.S. tariffs are regressive, and Porto (2006) develops the methodological framework for assessing distributional effects through household consumption. Our contribution is to construct welfare measures from two identification margins, each anchored to income-group-specific behavioral responses.[1]

A fourth strand takes a macroeconomic perspective. den Besten et al. (2026) show that tariff increases are contractionary for output and, post-WWII, disinflationary, pointing toward aggregate demand channels. Candia et al. (2026) argues that tariff-induced policy uncertainty links small price effects to large spending declines. Our household-level estimates complement these findings by showing that households reduce purchases even when retail prices move modestly, and we provide micro evidence on the mechanism: spending contractions on goods facing negligible tariffs, isolated using survey measures of pessimism, identify a precautionary channel concentrated in the middle of the income distribution.

Finally, we contribute to work linking a household survey on tariff expectations to spending behavior. Consumer sentiment surveys have long been used to study macroeconomic expectations (Coibion et al., 2024), but direct tariff-sentiment measures linked to transaction data are rare. We show that stated reduction intentions predict revealed spending cuts for higher-income households, that stated substitution intentions do not predict revealed substitution, and that stated stockpiling is associated with lower post-tariff spending. These patterns support an interpretation in which the behavioral response is dominated by contraction rather than reallocation, concentrated among households with slack.

The remainder of the paper proceeds as follows. Section 2 describes the data and the construction of key measures. Section 3 presents the empirical strategy and the main results on prices, spending, and quantities, documents heterogeneity across income groups and product essentiality, and provides evidence against stockpiling as an explanation for the spending contraction. Section 4 quantifies the welfare losses from the price channel.

[1] Related work explores broader labor market and welfare effects of trade shocks (Topalova, 2010; Autor et al., 2013; Pierce and Schott, 2016; Autor et al., 2021; Caliendo and Parro, 2023; Borusyak and Jaravel, 2021). Jaravel and Sager (2025) show that U.S.-China trade historically reduced consumer prices, especially for lower-income households. Feng et al. (2023) document price effects of Chinese imports and retaliatory tariffs.

Section 5 uses the linked household survey to identify the two channels through which tariffs operate and their distinct incidence across income groups. Section 6 concludes.

2 Data and the Construction of Key Measures

This section presents the data and key measures used throughout the paper. We begin with the transaction-level household data and trade data. These two datasets are merged using a crosswalk, which we describe next. Following the data overview, we outline the key measures employed in our analysis: the tariff exposure measure and the Fisher Price Index.

2.1 Transaction-Level Household Data

Our primary data source is the Numerator household panel, which provides comprehensive transaction-level records of consumer purchases combined with detailed demographic information. Panelists submit purchase records through the Receipt Hog mobile app via three channels: uploading photographs of paper receipts, connecting email accounts for automatic digital receipt collection, and linking retailer apps including Amazon, Walmart, and other major retailers. Numerator processes these receipts using automated tools supplemented by manual human verification. The dataset includes rich demographic attributes obtained through periodic surveys, such as income bracket, education levels, age groups, 5-digit zip code, household composition, and race/ethnicity. The dataset provides 16 categorical income brackets that are self-reported by panelists. Appendix A provides more information on data collection and representativeness.

The panel by the end of 2025 comprises 200,000 households selected from a broader pool exceeding one million registered users to ensure demographic and geographic representativeness. The Numerator panel operates as a rolling panel with panelists joining and exiting each month. Numerator designates households as "static panelists" if they maintain continuous and complete transaction records for a minimum of 12 months, enabling reliable longitudinal analysis. For our analysis, we restrict the sample to the 126,448 households who remained active throughout the entire 24-month period from January 2024 through December 2025, minimizing compositional changes.

The household data contain products that are identified with their unique item identification numbers (item ID). The data structure links three core identifiers: user ID (household), basket ID (shopping trip/transaction), and item ID (product), allowing us to track individual purchase decisions at the household-product-trip level. All products are mapped to a hierarchy of classifications. To ensure data quality, we omit items with prices below $1, as some zero or very low prices reflect buy-one-get-one-free deals or transcription errors, and we remove any product where unit price exceeds $5,000, or $500 for grocery items, which constitute a negligible share of total spending. The final data covers approximately $3.8 billion in expenditure on over 29 million unique items, over 2024 and 2025.

Table 1 lists the broader sectoral classification with each sector's expenditure shares in the final data. The dataset primarily captures expenditures on physical goods purchased through both brick-and-mortar and online retailers. However, several spending categories have limited or no coverage. Among services, only restaurants are well-represented while other services are mostly missing. Healthcare spending is partially captured through over-the-counter pharmaceutical products, co-pays, and some medical services. Motor vehicle purchases and leases are not captured, though automotive products such as oil and parts are included. Certain major spending categories are entirely absent, including housing expenditures (rent, mortgage, utilities) and financial services. Given these features of the data, we systematically exclude restaurant spending from our analysis, due to its minimal direct exposure to tariffs, and medical spending since we cannot capture its full extent. The resulting coverage represents approximately 20 percent of the consumption basket captured in the PCE and CPI. However, this subset captures the high-frequency purchases most directly affected by tariffs on traded goods. Moreover, recent research indicates that households form inflation expectations primarily based on prices of frequently purchased items such as groceries rather than expenditure-weighted basket averages (D'Acunto et al., 2021; Weber et al., 2022), suggesting that the categories in our data are particularly salient for understanding spending responses to tariff-policy driven shocks.

Table 1: Expenditure Shares by Sector

Sector Description	Expenditure Share (%)
Grocery	46.36
Health & Beauty	9.57
Home & Garden	9.52
Apparel, Footwear, and Accessories	6.57
Household	5.13
Tools & Home Improvement	4.23
Pet	4.14
Electronics	3.83
Party & Occasions	2.80
Toys	2.06
Automotive	1.40
Baby	1.19
Sports	1.00
Tobacco Products and Accessories	0.87
Office	0.83
Books	0.40
Entertainment	0.11

Note: Values in percent. Expenditure shares are calculated by aggregating total expenditure across all households in the Numerator panel for each sector for data spanning January 2024 through December 2025; total spending over the 24-month period is $3,802,841,523 on 29,362,394 items. Sectors are sorted in descending order by expenditure share.

2.2 Trade Data

In trade, products are classified into Harmonized System (HS) codes, a standardized international classification system used by customs authorities to identify and categorize traded goods. Every product has a 10-digit HS code, providing detailed product classification. Trade data can be aggregated to different levels, such as 6- or 4-digit HS codes. In this paper, we work with 4-digit HS codes (HS4), which provide broad product categories for reasons we discuss below in the household-to-trade data crosswalk section.

We assemble three components from public U.S. Census sources. First, monthly tariff rates at the HS4 level for 2024 and 2025 are constructed from the Census Bureau's monthly import merchandise microdata.[2] For each HS10 record, we observe consumption import value and calculated duties by source country and month. We aggregate to HS4

[2]Source: Merchandise data, Trade Imports. `https://www.census.gov/foreign-trade/data/dataproducts.html`.

by summing duties and import values across all 10-digit codes within each HS4 and across all source countries separately for each month. The realized effective tariff rate for HS4 code h in month t is then computed as in Equation (1) below. This yields a panel of HS4 × month tariff rates spanning January 2024 through December 2025.

Second, import penetration ratios are constructed from 2022 trade and production data and held fixed throughout the sample period. We obtain 2022 monthly import and export merchandise microdata from the Census Bureau and aggregate to the HS4 level by summing across all source (destination) countries and across all twelve months of 2022, yielding annual HS4-level imports M_h and exports X_h.[3] Domestic shipments S_h come from the 2022 Economic Census, covering manufacturing and agriculture, restricted to U.S. national totals.[4] We map shipments from North American Industry Classification System (NAICS) codes to HS4 codes using the Census Bureau's 2022 import and export concordances, with hierarchical matching at the 6-, 5-, 4-, 3-, and 2-digit NAICS levels to maximize coverage.[5] The 2022 Economic Census is the most recent comprehensive source of domestic production data by detailed industry, since the Economic Census is conducted every five years. Consequently, we fix import shares, import and export values, shipments, and import penetration ratios at their 2022 values throughout our analysis. The time variation in our tariff exposure measure comes exclusively from the monthly HS4 tariff rates over 2024–2025.

We now define the tariff rate and tariff exposure measure formally. For each HS4 code h in month t, the realized effective tariff rate is:

$$\tau_{h,t} = \frac{\sum_{k \in h} \text{Duties}_{k,t}}{\sum_{k \in h} \text{ImportValue}_{k,t}} \tag{1}$$

where the sum runs over all 10-digit HS codes k within HS4 code h and across all source countries, and $\text{Duties}_{k,t}$ and $\text{ImportValue}_{k,t}$ are total calculated consumption duties and

[3]Exports are matched to HS4 via the 10-digit Schedule B code, which we truncate to 4 digits in the same way as imports.

[4]Manufacturing and agriculture shipments cover all the HS4 codes in Numerator data. Manufacturing is covered under NAICS 31–33 and agriculture is under NAICS 11 where NAICS stands for North American Industry Classification System. These data are from the EC2231 Basic file for manufacturing and EC2211 Basic file for agriculture, downloaded from the 2022 Economic Census: https://www.census.gov/programs-surveys/economic-census/data/tables.html.

[5]Concordances available at https://www.census.gov/foreign-trade/reference/index.html.

total consumption import value for HS10 code k in month t. This yields effective applied tariff rates that reflect the actual tariff burden on imports, rather than announced statutory rates at the Federal Register. Due to delays in implementation and exemptions at the country or product level, realized rates can differ substantially from announced rates (Azzimonti, 2025; Cavallo et al., 2025; Gopinath and Neiman, 2026).

Import penetration is calculated using 2022 annual values:

$$\text{ImportPenetration}_h = \frac{M_h^{2022}}{M_h^{2022} + S_h^{2022} - X_h^{2022}} \tag{2}$$

where M_h^{2022}, X_h^{2022}, and S_h^{2022} denote 2022 annual imports, exports, and domestic shipments at the HS4 level constructed as described above. This ratio captures the extent to which domestic consumption relies on foreign sources. A value near 1 indicates nearly complete import dependence, while a value near 0 indicates that consumption is primarily met by domestic production.

The tariff exposure measure combines these two components. For each HS4 code h in month t:

$$\text{TariffExposure}_{h,t} = \tau_{h,t} \times \text{ImportPenetration}_h \tag{3}$$

Since import penetration is fixed at its 2022 value, all time variation in $\text{TariffExposure}_{h,t}$ comes from $\tau_{h,t}$. This measure weights the tariff rate by import penetration, capturing the effective tariff burden faced by consumers. A given tariff rate has substantially larger effects on categories with high import penetration, where domestic production alternatives are limited, than on categories with low import penetration. For example, a 10 percent tariff on electronics (with import penetration near 0.9) yields a tariff exposure of 9 percent, while the same tariff on dairy products (with import penetration near 0.1) yields an exposure of only 1 percent. Since our panel is constructed at the product-hierarchy × HS4 code level, each panel unit maps directly to a single HS4 code, and no within-category aggregation across HS codes is required. Identification comes from cross-sectional variation in import penetration combined with time-series variation in monthly tariff rates around the February 2025 tariff implementation.

2.3 Household to Trade Data Crosswalk

A central challenge in linking retail transaction data to trade-based tariff measures involves mapping retail product classifications in the household data to standardized HS codes used in international trade. To address this challenge, we create a crosswalk that maps retail product categories in the household data to HS4 codes preserving the full retail hierarchy available in the household data: sector → department → major category → category → HS4 code.

Our approach maps product categories to HS4 codes rather than matching individual Universal Product Codes (UPCs) to 10-digit HS codes as in Cavallo et al. (2025). This methodological choice reflects data constraints, where UPCs are only consistently available across a small share of products in the household data but the full retail hierarchy is universally available.[6] In the crosswalk, products are first matched at the most granular category level, if it exists, and unmatched items are successively matched at increasingly aggregate levels (major category, department, then sector).[7] Table 2 presents an example of the hierarchical classification structure, which we label as panel unit, ranging from the broadest (sector) to the most detailed (category) level, showing how household expenditure categories correspond to HS4 codes. The final data has 4,918 panel units.

Table 2: Example Crosswalk: Sector Hierarchies to HS Codes

Sector	Department	Major Category	Category	HS Code
Apparel	Apparel Accessories	Belts	Leather Belts	4203
Apparel	Baby Apparel	Baby Pajamas	Baby Girl Pajamas	6111
Electronics	Video Games & Acc.	Video Games	PlayStation Games	9504
Entertainment	Movies & TV	Blu-ray	Comedy Blu-ray Movies	8523
Grocery	Beverages	Tea	Earl Grey Tea	0902
Grocery	Dairy	Cheese	Soft Cheese	0406
Grocery	Meat	Sausage	Sausage-Pork	1601
Health & Beauty	Makeup	Eye Makeup	Mascara	3304

Notes: The table demonstrates the four-level hierarchical structure (Sector → Department → Major Category → Category) in the Numerator data mapping to standardized HS codes.

[6]Restricting our sample to only UPC-identified products would significantly reduce coverage and compromise the representativeness of household consumption baskets.

[7]Artificial intelligence was used in the construction of the category to HS4 crosswalk. The accuracy of the crosswalk was manually verified by the authors and was then cross-checked using artificial intelligence. This hierarchy match often returns 6-digit HS codes for some categories. But to ensure consistency across the wide range of products, we aggregate them to the 4-digit level.

With a household-to-trade data crosswalk in place, we can visualize tariff rates and tariff exposure at the aggregate level for products in our household data. Figure 1 shows the evolution of both measures over the sample period. Average tariff rates are low and stable throughout 2024, then rise sharply after tariffs are imposed in February 2025. The mean post-period tariff rate across HS4 codes is approximately 10 percent, up from 2 percent pre-tariff period. The mean post-period tariff exposure is approximately 0.073 percent, up from 0.011 percent in the pre-period, reflecting the fact that many product categories have substantial domestic production that partially shields consumers from the full tariff rate.

FIGURE 1: Average Tariff Rate and Tariff Exposure

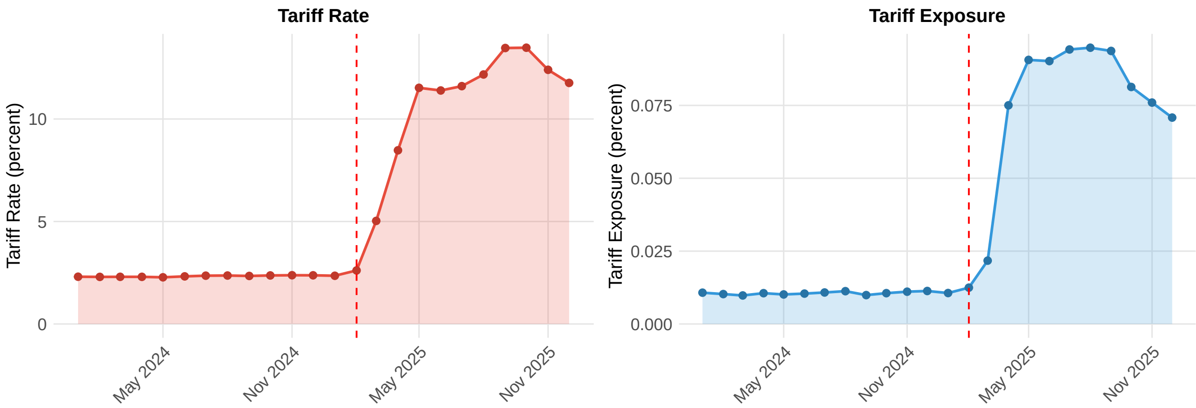

Notes: This figure shows the monthly average tariff rate (left panel) and tariff exposure (right panel) over the sample period for products in the Numerator panel. For each year-month, we calculate the mean across all observations. The tariff rate is the ratio of duties collected to import value (Equation 1). Tariff exposure multiplies the tariff rate by import penetration (Equation 3). Both measures are in percent. The red dashed vertical line marks February 2025, when tariffs were announced.

Appendix C shows tariff exposure by sector, country, and "essentiality," which we define in Section 3. In Appendix B, we construct a tariff measure for a smaller set of products for which we know the country of origin information, through another dataset matched with the household panel, to validate the accuracy of the tariff measures.

3 Price, Spending and Quantity Responses to Tariffs

This section presents our empirical strategy and the main results on price pass-through and spending responses to the 2025 tariffs. We first describe the difference-in-differences design used on the full household panel, then turn to aggregate results, heterogeneity across the income distribution, and heterogeneity across product essentiality.

3.1 Empirical Strategy

Our main specification is a difference-in-differences model that exploits variation in tariff rates across product categories:

$$\ln(Y_{h(i),t}) = \gamma_t + \delta_{i,m(t)} + \beta \left(\text{Tariff}_{h(i),t} \times \text{Post}_t\right) + \varepsilon_{i,t} \tag{4}$$

where i indexes the panel unit (sector × department × major category × category × HS code), $h(i)$ denotes the HS4 code to which panel unit i maps, t indexes the year-month, and $Y_{h(i),t}$ is one of three outcomes: the Fisher Price Index, real spending (deflated by the PCE goods price index) or quantity purchased.[8] The term γ_t is a year-month fixed effect absorbing common macroeconomic shocks. The binary indicator Post_t equals one for months from February 2025 onward, when tariffs were implemented.

Our benchmark specification augments Equation (4) by controlling for pre-existing linear trends in the price equation:

$$\ln(\text{Price}_{h(i),t}) = \gamma_t + \delta_{i,m(t)} + \theta_i \cdot \text{TimePre}_{i,t} + \beta \left(\text{Tariff}_{h(i),t} \times \text{Post}_t\right) + \varepsilon_{i,t} \tag{5}$$

where $\text{TimePre}_{i,t}$ is a linear time trend measured in the pre-tariff period and capped at zero in the post-period. This panel-specific detrending controls for the possibility that high-tariff product categories were already on different price trajectories before tariffs were imposed. The coefficient θ_i is estimated separately for each panel unit i. We apply detrending only to the price equation. This approach follows Cavallo et al. (2025) and removes product-specific linear pre-trends estimated over the pre-tariff period before esti-

[8] Data for the PCE goods index is downloaded from FRED. Personal consumption expenditures: Goods (chain-type price index) (DGDSRG3M086SBEA).

mating tariff pass-through. The spending and quantity equations are estimated without panel-specific pre-trends.

The key treatment variable $\text{Tariff}_{h(i),t}$ takes one of two forms. The first specification uses the realized monthly tariff rate directly:

$$\text{Tariff}_{h(i),t} = \tau_{h(i),t} \tag{6}$$

where $\tau_{h(i),t}$ is the realized effective tariff rate on HS4 code $h(i)$ in month t, defined in Equation (1). The second specification uses the tariff exposure measure:

$$\begin{aligned}\text{Tariff}_{h(i),t} =& \text{TariffExposure}_{h(i),t} \\ =& \tau_{h(i),t} \times \text{ImportPenetration}_{h(i)}\end{aligned} \tag{7}$$

where $\text{ImportPenetration}_{h(i)}$ is as defined in Equation (2). The interaction with Post_t in Equation (4) sets the regressor to zero in the pre-tariff period, so that the estimated coefficient β measures the response of the outcome to a one percentage point increase in the post-period tariff rate (or tariff exposure) for HS4 code $h(i)$, relative to the pre-period baseline.

The term $\delta_{i,m(t)}$ represents an interaction fixed effect structure that absorbs product-specific seasonality. Formally, $\delta_{i,m(t)} = \delta_{i \times m(t)}$ where $m(t) \in \{1, \ldots, 12\}$ maps each year-month to its calendar month. Identification of the tariff effect comes entirely from the within-cell pre-post difference, after absorbing aggregate time effects through γ_t. This structure absorbs product-specific seasonality without imposing parametric restrictions on how seasonal patterns vary across products. By allowing each product category to have its own seasonal baseline, the specification accounts for unit-specific cyclical demand variation that might otherwise confound the tariff effect.

The coefficient β identifies whether categories with higher tariff rates or higher tariff exposure experience differentially larger changes in prices or spending after tariffs are imposed, relative to the pre-tariff period and relative to lower-tariff categories. Because $\text{Tariff}_{h(i),t}$ is a rate measured in percent, β has a direct interpretation: in the tariff rate specification, β represents the percentage point change in outcome per one percentage point increase in the tariff rate; in the tariff exposure specification, β represents the percentage point change in outcome per one percentage point increase in tariff exposure. Note that only the tariff rate coefficient maps directly to pass-through in the traditional trade-literature sense (a coefficient of 0.15 implies approximately 15% pass-through); the tariff exposure coefficient instead measures the response to a composite of the tariff rate and import dependence. Standard errors are clustered at the HS4 code level throughout.

For estimating price pass-through, we leverage product-level prices and quantities from the household panel and compute a monthly Fisher Price Index (FPI) beginning in January 2024 at the HS4 level. The FPI accounts for both base- and current-period quantity weights, offering a key advantage over alternatives that rely solely on prices without quantity weighting. The fixed-base FPI holds the product basket fixed when computing price changes and is therefore immune to within-category composition effects.

For period t relative to base period 0, the FPI for an HS4 code h is:

$$F_{h,t} = \sqrt{L_{h,t} \cdot P_{h,t}} = \sqrt{\frac{\sum_{k\in\Omega_h} p_{hkt}q_{hk0}}{\sum_{k\in\Omega_h} p_{hk0}q_{hk0}} \cdot \frac{\sum_{k\in\Omega_h} p_{hkt}q_{hkt}}{\sum_{k\in\Omega_h} p_{hk0}q_{hkt}}} \tag{8}$$

where $L_{h,t}$ denotes the Laspeyres index using base-period quantities, $P_{h,t}$ the Paasche index using current-period quantities, and p_{hkt}, q_{hkt} are the price and quantity of product k within HS code h in period t. The set Ω_h denotes the products belonging to HS code h. The FPI is the geometric mean of the two indices. This construction parallels that of the Bureau of Economic Analysis's PCE Index, which similarly employs expenditure-weighted aggregation. Appendix A.5 provides a discussion of how the aggregate Fisher Price Index for a longer time period of January 2018 to December 2025, constructed from the Numerator household panel, compares with officially published price indices.

We estimate Equations (4) and (5) for the full sample of all product categories and for a restricted sample that excludes grocery products. Grocery constitutes 46% of household spending in our data as presented in Table 1, which might attenuate tariff effects. We present both specifications to assess whether aggregate results are driven by this dominant but distinctive category. To explore heterogeneity across the income distribution, we estimate income-group-specific versions of the main specification using Fisher Price Indices constructed separately for different income groups as outlined in the next section.

3.2 Aggregate Results

Table 3 presents our main estimates for the full sample of all product categories. Panel A shows results using the tariff rate as the treatment variable; Panel B uses the tariff exposure measure (tariff rate × import penetration). For each treatment definition, we report the results for prices (with and without panel-specific pre-trends), spending and quantities.

TABLE 3: Main Results: Price, Spending, and Quantity (All Sectors)

	Price		Spending	Quantity
	Detrended (1)	Not Detrended (2)	(3)	(4)
Panel A: Tariff Rate				
Tariff Rate × Post	0.15***	0.12***	−0.46***	−0.42***
	(0.05)	(0.04)	(0.10)	(0.09)
Observations	108,626	108,626	108,626	108,626
Adjusted R^2	0.63	0.55	0.97	0.98
Panel B: Tariff Exposure				
Tariff Exposure × Post	0.20***	0.13**	−0.63***	−0.57***
	(0.07)	(0.06)	(0.13)	(0.11)
Observations	108,626	108,626	108,626	108,626
Adjusted R^2	0.63	0.54	0.97	0.98
Year-month FE	Yes	Yes	Yes	Yes
Panel ID × Month FE	Yes	Yes	Yes	Yes
Panel-specific pre-trends	Yes	No	No	No

Notes: This table reports estimates from Equations (4) and (5) run on Numerator panel data. Dependent variables are log Fisher Price Index (columns 1–2), log real spending (column 3), and log quantity (column 4). Panel A uses the tariff rate as the treatment variable; Panel B uses the tariff exposure measure (tariff rate × import penetration). Price results are reported for both detrended (column 1) and non-detrended (column 2) specifications. Detrending includes panel-specific linear time trends in the price equation only. Spending and quantity equations do not include panel-specific pre-trends. All specifications include year-month and panel-unit-by-calendar-month interaction fixed effects. Standard errors clustered at HS4 code level in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

The two panels of Table 3 reflect different notions of treatment intensity. The tariff rate specification in Panel A treats all product categories as equally exposed to a given tariff, regardless of how much domestic consumption relies on imports, while the tariff exposure specification in Panel B scales the tariff rate by import penetration to capture how much the tariff actually affects consumers in each category. Based on Panel A, a one percentage point increase in the tariff rate raises retail prices by approximately 0.15 percentage points, implying retail pass-through of around 15 percent. Panel B yields systematically larger coefficients across all outcomes, indicating that the exposure measure better isolates categories where tariffs generate meaningful consumer impact. To provide a sense of magnitudes for the exposure specification, mean tariff exposure rose from 0.011 percent in the pre-tariff period to 0.073 percent post-tariff. Evaluated at the mean increase in tariff exposure over the sample period (from 0.011 to 0.073, a change of 0.062 percentage points), the benchmark Panel B coefficients imply that average prices in an affected category rise by approximately 1.2 percent while real spending falls by approximately 3.9 percent, a contraction more than three times the price increase. Under the tariff rate specification in Panel A, the mean tariff rate increase of roughly 8 percentage points implies a price increase of approximately 1.2 percent and a spending decline of 3.7 percent. Both treatment definitions thus yield economically similar magnitudes and confirm that the spending response substantially exceeds what price pass-through alone would predict. The not-detrended coefficients are slightly smaller with also slightly weaker model fit.[9]

The most noteworthy feature of Table 3 is that spending and quantity decline by magnitudes 3–4 times larger than the corresponding price increase, across every specification and treatment definition. The close alignment of the spending and quantity coefficients confirms that spending cuts on tariff-exposed goods are driven almost entirely by reductions in quantities purchased.[10] If households responded to tariffs solely through price-induced cutbacks in direct proportion to higher prices, the ratio of the spending coefficient to the price coefficient would imply a price elasticity of demand on the order of three, if interpreted mechanically. Note that, in our context, this ratio is a comparison across equations estimated on the same treatment variable but with different dependent variables and different levels of aggregation, and should not be read as a struc-

[9]Panel-specific pre-trend slopes $\hat{\theta}_i$ are weakly negatively correlated with post-period tariff exposure ($\rho = -0.09$), indicating that high-exposure categories experienced modest pre-existing price declines. Detrending removes this drift and yields a larger pass-through estimate than the non-detrended specification in Column 2). The non-detrended estimate thus provides a conservative lower bound on pass-through.

[10]Our spending measure is deflated by the aggregate PCE goods price index, so real spending and quantity differ only to the extent that category-specific prices differ from the aggregate deflator. The small residual gap between the two coefficients reflects this difference.

tural elasticity. It is instead informative as an upper bound on the demand response that would be needed to rationalize the spending contraction through price substitution alone and such a large price elasticity of demand as a response to tariffs would be supported by previous research such as Fajgelbaum et al. (2020); Cavallo et al. (2021b); Amiti et al. (2019). That the implied elasticity exceeds plausible values for e.g. groceries and household goods, however, suggests that quantity adjustments operate through channels beyond price responses, including precautionary behavior, uncertainty about future tariff policy, anticipatory cutbacks in discretionary spending, or broader demand contractions. Section 5 exploits survey-based measures of household beliefs to show that this wedge reflects precautionary reallocation of the household basket rather than intertemporal substitution or a uniform demand response. Appendix D provides the analyses to ensure the stability of our results when we exclude grocery and and also leaves one sector out at a time.

Figure 2 presents event study estimates that traces the dynamic path of price and spending coefficients over the sample period. The specification interacts $\text{Tariff}_{h(i)}$ with indicators for each year-month omitting January 2025 as the reference period, while retaining panel-unit-by-calendar-month fixed effects and panel-specific pre-trends in the price equation. The top row shows the tariff rate specification and the bottom row shows the tariff exposure specification.

The event study provides visual confirmation of the parallel trends assumption and the timing of tariff effects. In both specifications, price coefficients fluctuate around zero throughout the pre-tariff period with no discernible trend. Following the February 2025 implementation, price coefficients turn positive and rise gradually, consistent with lagged and incomplete pass-through to retail prices. The quantity and spending panels show a complementary pattern. Pre-period coefficients are centered near zero, while post-tariff coefficients shift sharply negative, with declines materializing starting June–July 2025, due to the gradual pass-through in prices, and these declines are persistent through the end of the sample.[11]

We assess the robustness of our event-study estimates to potential violations of the

[11]In the pre-period the tariff exposure is nearly zero for all HS4s, so the interaction coefficient is identified off essentially no cross-sectional variation and standard errors inflate accordingly, which dissipates once tariff changes are implemented in February 2025.

FIGURE 2: Event Study: Price, Spending, and Quantity Responses to Tariffs

Notes: The figure plots coefficients from event study versions of Equations (4) and (5) run on Numerator panel data, interacting the treatment variable with year-month indicators (January 2025 omitted as reference). The top row uses the tariff rate; the bottom row uses the tariff exposure measure. Left panels show price (log Fisher Price Index); right panels show quantity and spending (both in logs and spending is real spending). All specifications include year-month and panel-unit-by-calendar-month interaction fixed effects, with panel-specific linear pre-trends in the price equation. Shaded areas indicate 95% confidence intervals based on standard errors clustered at the HS4 code level. The vertical dashed line marks February 2025, the tariff implementation date.

parallel trends assumption using the methodology of Rambachan and Roth (2023), which constructs confidence sets for causal effects that remain valid even if parallel trends do not hold exactly, without requiring a specific assumption about the nature of any pre-existing differential trends. We implement two classes of restrictions: the relative magnitudes approach bounds the post-treatment trend violation as a fraction $\bar{M}$ of the largest pre-treatment violation, and the smoothness approach bounds the second difference of the trend (M), restricting how much the differential trend can change slope between consecutive periods, with $M = 0$ corresponding to linear extrapolation of any pre-existing trend. When we assume that parallel trends holds exactly ($\bar{M} = 0$ or $M = 0$), both the spending and quantity confidence intervals exclude zero, confirming statistically significant reductions. Under the smoothness approach at $M = 0$, the average post-treatment price effect also excludes zero. In other words, when we allow for the possibility that any pre-existing trend simply continued linearly into the post-period, the results for all three outcomes remain significant.

Based on these results, we adopt the specification with price detrending and the tariff exposure measure as our benchmark for the remainder of the paper. The detrended specification appropriately controls for pre-existing price trends that would otherwise attenuate the price coefficient. The tariff exposure measure is the conceptually appropriate treatment variable for welfare analysis. It captures the effective tariff burden on consumers by recognizing that a tariff on a product category with high import penetration directly affects a larger share of consumption than the same tariff on a category where domestic production can absorb the shock. We continue to report tariff rate results as a robustness check, wherever appropriate, and because the tariff rate coefficient has a direct pass-through interpretation useful for comparison with existing work.

3.3 Income-Specific Results

The aggregate results in the previous section provide a useful overview of tariff impacts but there might be significant heterogeneity on how tariffs affect households with different demographics. For instance, the aggregate FPI weights all products equally across the income distribution, and the pooled spending coefficient assumes that households at different income levels respond identically to tariff exposure. However, if consumption baskets and behavioral responses differ systematically by income, these aggregate estimates may obscure distributional consequences.

For this analysis, we define three income groups from households' reported income brackets: low-income households earning \$0–\$40,000, middle-income households earning \$40,000–\$125,000, and high-income households earning above \$125,000, all annual and gross.[12] We then estimate income-specific versions of Equations (4) and (5) where the price measure is the income-group specific FPIs. Table 4 presents price, spending, and quantity responses by income group, focusing on the benchmark specifications.[13]

Table 4: Price, Spending, and Quantity Responses by Income Group

	Price			Spending			Quantity		
	Low (1)	Middle (2)	High (3)	Low (4)	Middle (5)	High (6)	Low (7)	Middle (8)	High (9)
Panel A: All Sectors									
Tariff Rate × Post	0.18***	0.13**	0.18***	−0.33***	−0.40***	−0.42***	−0.29***	−0.36***	−0.41***
	(0.04)	(0.06)	(0.07)	(0.10)	(0.09)	(0.10)	(0.08)	(0.08)	(0.09)
Adjusted R^2	0.35	0.51	0.51	0.96	0.97	0.96	0.98	0.98	0.98
Tariff Exposure × Post	0.28***	0.20***	0.15**	−0.54***	−0.56***	−0.57***	−0.47***	−0.46***	−0.52***
	(0.06)	(0.07)	(0.07)	(0.11)	(0.13)	(0.13)	(0.09)	(0.10)	(0.11)
Adjusted R^2	0.36	0.52	0.50	0.96	0.97	0.96	0.98	0.98	0.98
Panel B: Excluding Grocery									
Tariff Rate × Post	0.19***	0.12	0.18**	−0.23**	−0.37***	−0.32***	−0.21**	−0.33***	−0.33***
	(0.05)	(0.07)	(0.08)	(0.09)	(0.10)	(0.09)	(0.08)	(0.09)	(0.09)
Adjusted R^2	0.30	0.45	0.46	0.94	0.96	0.95	0.97	0.97	0.97
Tariff Exposure × Post	0.27***	0.17*	0.10	−0.38***	−0.49***	−0.38**	−0.32***	−0.39***	−0.35***
	(0.08)	(0.09)	(0.10)	(0.13)	(0.15)	(0.14)	(0.11)	(0.14)	(0.14)
Adjusted R^2	0.30	0.45	0.44	0.94	0.96	0.95	0.97	0.97	0.97
Observations (Panel A)	103,511	107,190	105,771	103,511	107,190	105,771	103,511	107,190	105,771
Observations (Panel B)	76,617	80,053	78,756	76,617	80,053	78,756	76,617	80,053	78,756
Year-month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Panel ID × Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Panel-specific pre-trends	Yes (price)	Yes (price)	Yes (price)	No	No	No	No	No	No

Notes: This table reports income-group-specific estimates from Equation (5). Columns 1–3 use income-specific Fisher Price Indices as the dependent variable; columns 4–6 use log real spending; columns 7–9 use log quantity. Panel A uses the full sample; Panel B excludes grocery products. Income groups: Low (<\$40k), Middle (\$40k–\$125k), High (>\$125k). All specifications include year-month and panel-unit-by-calendar-month interaction fixed effects, with panel-specific pre-trends in the price equation only. Observation counts differ across income groups because income-specific Fisher Price Indices have missing values for some HS4 × month cells when an income group records no transactions in a given product category during a month. Standard errors clustered at HS4 code level in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Three patterns emerge from Table 4, from the benchmark tariff exposure specification in the second row of Panel A. First, price pass-through is monotonically decreasing in income. Low-income households face the largest price response at 0.28, middle-income households face 0.20, and high-income households face the smallest response at 0.15, landing support the cheapflation phenemenon documented in Cavallo and Kryvtsov (2024).

[12]Based on the 2024 Current Population Survey published in September 2025 by the U.S. Census Bureau, low-income households with \$40,000 in annual income would fall between the 20th and 30th percentile of household income, and high-income households with \$125,000 annual income would fall between the 60th and 70th percentile. See https://www.census.gov/library/publications/2025/demo/p60-286.html.

[13]Results without price detrending are qualitatively similar and available upon request.

This ordering is also consistent with evidence that higher-income households can substitute quality and change shopping behavior to mitigate inflation in line with the findings of Argente and Lee (2021), while low-income consumers shop at retailers with less competitive pricing and consume products with fewer close substitutes would need to absorb higher costs. Moreover, our aggregate income-specific Fisher Price Indices constructed from the household panel show that the basket purchased by low-income households experienced cumulatively higher price growth over the sample period than the baskets of middle- and high-income households (Appendix A.5, Figure A.4), consistent with the inflation-inequality patterns documented in Argente and Lee (2021). The fact that pass-through is highest at the bottom of the income distribution is a first indication that the welfare burden of tariffs is regressive even before accounting for differences in spending baskets and income.

Second, spending responses are large and similar across income groups. Low-income households reduce spending by 0.54 percent per percentage point of tariff exposure, middle-income households by 0.56 percent, and high-income households by 0.57 percent. The similarity of the spending response across the distribution reflects two offsetting forces. High-income households have the most discretionary spending and the most flexibility to defer or eliminate purchases. Low-income households face the highest pass-through and have the least slack in their budgets, so even modest price increases force substantial reallocation. These forces produce spending contractions of comparable magnitude at both ends of the distribution.

Third, the quantity responses show high-income households exhibiting the largest absolute adjustment, with low- and middle-income households responding by similar and smaller magnitudes. The combination of these three patterns has direct implications for welfare. Low-income households face the highest pass-through, which produces the largest welfare cost as a share of income, as we document in Section 4. Excluding grocery (Panel B) leaves the qualitative patterns unchanged, with slightly attenuated coefficients overall. The broadly similar spending and quantity responses across income groups indicate that all households contract by comparable total magnitudes. What differs across the distribution is the margin of adjustment: Section 5 shows that middle-income households execute the contraction through a compositional shift away from non-essentials, while low-income households absorb the shock primarily through higher prices on continuing purchases.

3.4 Heterogeneity by Product Essentiality

To understand which product categories drive the aggregate results, we decompose the sample by essentiality. Essential goods include grocery, household and health and beauty; non-essential goods include everything else (automotive, electronics, apparel, toys, and all others reported in table 1). Table 5 presents results for essential and non-essential goods separately, focusing on the benchmark specifications.

Table 5: Results by Product Essentiality

	Price		Spending		Quantity	
	Essential (1)	Non-Essential (2)	Essential (3)	Non-Essential (4)	Essential (5)	Non-Essential (6)
Panel A: Tariff Rate						
Tariff Rate × Post	0.00 (0.04)	0.09 (0.06)	0.09 (0.11)	−0.37*** (0.12)	0.08 (0.08)	−0.34** (0.13)
Observations	44,502	64,124	44,502	64,124	44,502	64,124
Adjusted R^2	0.52	0.56	0.99	0.96	0.99	0.98
Panel B: Tariff Exposure						
Tariff Exposure × Post	0.06 (0.05)	0.10 (0.08)	0.10 (0.14)	−0.46** (0.19)	0.03 (0.13)	−0.39** (0.18)
Observations	44,502	64,124	44,502	64,124	44,502	64,124
Adjusted R^2	0.52	0.55	0.99	0.96	0.99	0.98
Year-month FE	Yes	Yes	Yes	Yes	Yes	Yes
Panel ID × Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Panel-specific pre-trends	Yes (price only)	Yes (price only)	No	No	No	No

Notes: This table reports estimates from Equations (4) and (5) by product essentiality. Panel A uses the tariff rate as the treatment variable; Panel B uses the tariff exposure measure. Columns 1–2 report price effects; columns 3–4 report spending effects; columns 5–6 report quantity effects. Essential goods: Grocery, Household, Health & Beauty. Non-essential goods: all other sectors. All specifications include year-month and panel-unit-by-calendar-month interaction fixed effects, with panel-specific pre-trends in the price equation only. Standard errors clustered at HS4 code level in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

The decomposition reveals that the aggregate spending and quantity results are driven entirely by non-essential goods. Essential goods show no statistically significant price, spending, or quantity response under either treatment definition. The absence of significant price pass-through in either subsample, despite the significant aggregate effect, likely reflects attenuation from splitting the sample and the resulting loss of cross-category variation that identifies the pooled estimate. Non-essential goods exhibit significant spending and quantity declines of comparable magnitude in both specifications. The concentration of effects in discretionary categories aligns with the interpretation that households cut back where they have flexibility to do so, rather than responding mechanically to price changes across all goods. This pattern also helps explain the large wedge between price and spending responses documented in Table 3. Non-essential goods are precisely the categories where demand is most elastic and where precautionary or uncertainty-driven cutbacks are most likely to operate.

3.5 Evidence Against Stockpiling

Before moving from the aggregate spending and quantity responses documented above to the welfare and behavioral interpretation of those responses, we address a natural alternative explanation for the observed contractions: anticipatory stockpiling. If households pulled purchases forward in anticipation of tariff-driven price increases, the negative post-tariff coefficients documented in Table 3 might partly or entirely reflect the mechanical unwinding of that pull-forward rather than a genuine reduction in household consumption. The event study in Figure 2 shows larger coefficient movements immediately following the February 2025 tariff implementation under the exposure specification, which is exactly the pattern one would expect if stockpiling were dominant. We provide two independent pieces of evidence against this interpretation, one from the aggregate dynamics of the response and one from stated-vs-revealed survey evidence.

We first partition the post-tariff period into a three-month transition window (February–April 2025), during which any pull-forward behavior would plausibly unwind, and a steady-state window covering the remainder of the sample (May–December 2025). For each event-study specification, we compute the average monthly response within each window, constructing joint standard errors from the full variance-covariance matrix of the estimated coefficients. Figure 3 presents the results. Across all four outcome-by-treatment combinations, the transition-window averages are small and statistically indistinguishable from zero, consistent with modest anticipatory behavior but not with a large intertemporal reallocation that would be expected to swamp subsequent effects. In stark contrast, the steady-state averages are economically large, precisely estimated, and uniformly negative: spending contracts by roughly 0.45 to 0.60 percentage points per month and quantity by a similar magnitude, with confidence intervals that exclude zero at conventional levels in both the tariff rate and tariff exposure specifications. These steady-state averages also closely match the static difference-in-differences estimates reported in Table 3, confirming that the aggregate results capture sustained household behavior rather than transitional dynamics.

FIGURE 3: Average Spending and Quantity Responses: Stockpiling versus Sustained Response

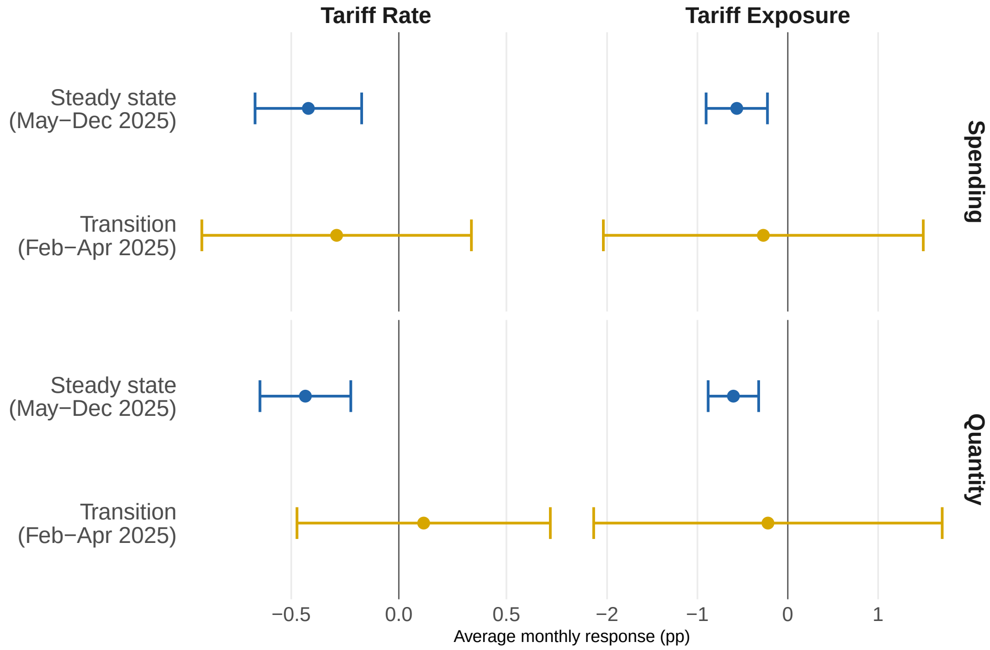

Notes: The figure reports average monthly responses of log real spending and log quantity to tariffs, separately for a three-month transition window (February–April 2025) and a steady-state window covering the remainder of the sample period (May–December 2025). Estimates are constructed as linear combinations of event-study coefficients from Equations (4) and (5), averaging the coefficients within each window. Standard errors are computed from the full variance-covariance matrix of the estimated event-study coefficients to account for correlation across event-time dummies, and are clustered at the HS4 code level. The transition window captures any anticipatory or stockpiling behavior that would unwind shortly after tariff implementation; the steady-state window captures the sustained response once such dynamics have dissipated. Horizontal bars report 95% confidence intervals. Rows correspond to outcomes (spending, quantity); columns correspond to treatment definitions (tariff rate, tariff exposure).

A complementary piece of evidence comes from the household survey. When we link stated stockpile intent to revealed behavior on tariff-exposed goods, the estimated quantity response among stated stockpilers is an order of magnitude smaller than the aggregate contraction and is accompanied by no commensurate spending response (Table 10). A further timing test confirms no front-loading signature in the early post-period

(Appendix E.6, Table E.7).

Taken together, these two pieces of evidence indicate that the negative spending and quantity responses documented in this section reflect genuine reductions in household consumption during the post-tariff period rather than the mechanical unwinding of anticipatory buying. The wedge between the modest price pass-through and the much larger spending contraction that motivates the survey and welfare analyses in Sections 5 and 4 therefore requires an explanation beyond intertemporal substitution.

4 Welfare Analysis

This section quantifies the welfare cost of the 2025 tariffs and documents that low-income households bear roughly five to nine times the welfare burden of high-income households when measured as a share of income. We present three complementary measures anchored to different identification strategies, all of which deliver the same regressive pattern. In this section, i indexes households and h indexes HS4 product codes.

Before turning to the measures, we clarify what they identify. Each is constructed from coefficients $\hat{\beta}^g_{\text{spend}}$, $\hat{\beta}^g_{\text{price}}$, and $\hat{\beta}^g_{\text{qty}}$ from Equations (4)–(5), in which the year-month fixed effect γ_t absorbs everything common across HS4 codes within each month. The coefficients therefore capture the differential response of high-exposure HS4 codes relative to low-exposure codes within the same month, not the response relative to a counterfactual without the 2025 tariffs. Any uniform component of the tariff shock, such as an across-the-board contraction from economy-wide tariff uncertainty, an income effect from the average tariff burden, or a precautionary saving response to policy news, is absorbed by γ_t. Our welfare totals therefore place a lower bound on the total household-level cost. Because γ_t is common across income groups within each month, any unidentified aggregate component cancels out of the ratio of welfare losses across income groups. The regressivity documented in this section is invariant to the unidentified aggregate level shift.

We define three welfare measures in this section.

Net Spending Effect. Our first measure captures the total observed change in household expenditure on tariff-affected goods:

$$\text{NetSpendingEffect}_{i,h} = \text{Spend}_{i,h}^{\text{pre}} \times \hat{\beta}_{\text{spend}}^{g(i)} \times \overline{\text{TariffExposure}}_h^{\text{post}} \tag{9}$$

where $\text{Spend}_{i,h}^{\text{pre}}$ is household i's annualized 2024 pre-tariff spending on HS4 code h, $\hat{\beta}_{\text{spend}}^{g(i)}$ is the income-group-specific spending coefficient from Table 4, and $\overline{\text{TariffExposure}}_h^{\text{post}}$ is the average post-period (February–December 2025) tariff exposure for HS4 code h, defined as $\overline{\text{TariffExposure}}_h^{\text{post}} = \frac{1}{|T^{\text{post}}|} \sum_{t \in T^{\text{post}}} \text{TariffExposure}_{h,t}$, with $\text{TariffExposure}_{h,t}$ as in Equation (7). This measure is useful for gauging the scale of cross-sectional consumption reallocation but is not itself a welfare measure. Instead, it nets the quantity reduction against the price increase, treating foregone consumption as a costless adjustment when it in fact represents a loss of consumer surplus.

Deadweight Loss. Our second measure is the Harberger consumer surplus loss, decomposed into a price burden on inframarginal units and a deadweight triangle on units no longer purchased:

$$\text{PriceBurden}_{i,h} = \text{Spend}_{i,h}^{\text{pre}} \times |\hat{\beta}_{\text{price}}^{g(i)}| \times \overline{\text{TariffExposure}}_h^{\text{post}} \tag{10}$$

$$\text{DWL}_{i,h} = \frac{1}{2} \times \text{Spend}_{i,h}^{\text{pre}} \times |\hat{\beta}_{\text{price}}^{g(i)} \times \overline{\text{TariffExposure}}_h^{\text{post}}| \times |\hat{\beta}_{\text{qty}}^{g(i)} \times \overline{\text{TariffExposure}}_h^{\text{post}}| \tag{11}$$

$$\text{WelfareLoss}_{i,h}^{\text{DWL}} = \text{PriceBurden}_{i,h} + \text{DWL}_{i,h} \tag{12}$$

The Harberger formulation requires only the local price and quantity responses and a linear demand approximation around the pre-tariff equilibrium. It is valid under the assumption that foreign supply is perfectly elastic, which is supported by recent evidence that 90 to 96 percent of tariff-induced price changes are absorbed domestically; see Amiti et al. (2026) and Hinz et al. (2026).[14]

Sato-Vartia Welfare Loss. Our third measure is the Sato-Vartia exact log-change cost-of-living index, following Sato (1976), Vartia (1976), Feenstra (1994), and Redding

[14]The DWL measure captures consumer welfare losses only and understates the total domestic welfare cost, which additionally includes reduced intermediary profits.

and Weinstein (2020). The change in the cost of living for income group g is:

$$\Delta \log P_g^{SV} = \sum_h w_{g,h}^{SV} \cdot \hat{\beta}_{\text{price}}^{g} \cdot \overline{\text{TariffExposure}}_h^{\text{post}} \tag{13}$$

with Sato-Vartia weights constructed from the logarithmic mean of pre- and post-tariff expenditure shares:

$$w_{g,h}^{SV} = \frac{L(\omega_{g,h}^{post}, \omega_{g,h}^{pre})}{\sum_{h'} L(\omega_{g,h'}^{post}, \omega_{g,h'}^{pre})}, \qquad L(a,b) = \frac{a-b}{\log a - \log b} \tag{14}$$

Pre-tariff shares $\omega_{g,h}^{pre}$ are computed from 2024 spending; post-tariff shares $\omega_{g,h}^{post}$ are computed symmetrically from observed post-period spending and renormalized within income group. No regression coefficient other than $\hat{\beta}_{\text{price}}^{g}$ enters the Sato-Vartia construction. Substitution is captured non-parametrically through observed expenditure reallocation rather than through an estimated elasticity. The household-level welfare loss is:

$$\text{Welfare}_i^{SV} = \text{TotalSpend}_i^{\text{pre}} \cdot \Delta \log P_{g(i)}^{SV} \tag{15}$$

The deadweight loss and the Sato-Vartia measure are complementary: the DWL uses the local quantity elasticity together with a linear demand approximation; the Sato-Vartia measure uses the global pattern of expenditure reallocation and takes no stand on demand curvature. Household-level welfare for all three measures is obtained by summing across HS4 codes: $\text{WelfareLoss}_i = \sum_h \text{WelfareLoss}_{i,h}$.

The welfare burden of the 2025 tariffs is sharply regressive when measured as a share of household income. Figure 4 presents all three welfare measures as a share of approximate household income, using bracket midpoints: \$20,000 for low-income, \$82,500 for middle-income, and \$187,500 for high-income.[15]

Under the deadweight loss measure, low-income households bear a welfare cost of 0.19 percent of income, compared with 0.05 percent for middle-income and 0.02 percent for high-income households. The Sato-Vartia measure yields 0.18, 0.05, and 0.02 percent

[15] The high-income midpoint of \$187,500 is a lower bound because this group includes all households above \$125,000, with the top bracket capturing \$250,000 and above with no upper bound. If actual mean incomes in this group exceed \$187,500, the regressive pattern is strengthened.

FIGURE 4: Welfare Losses as Share of Approximate Income

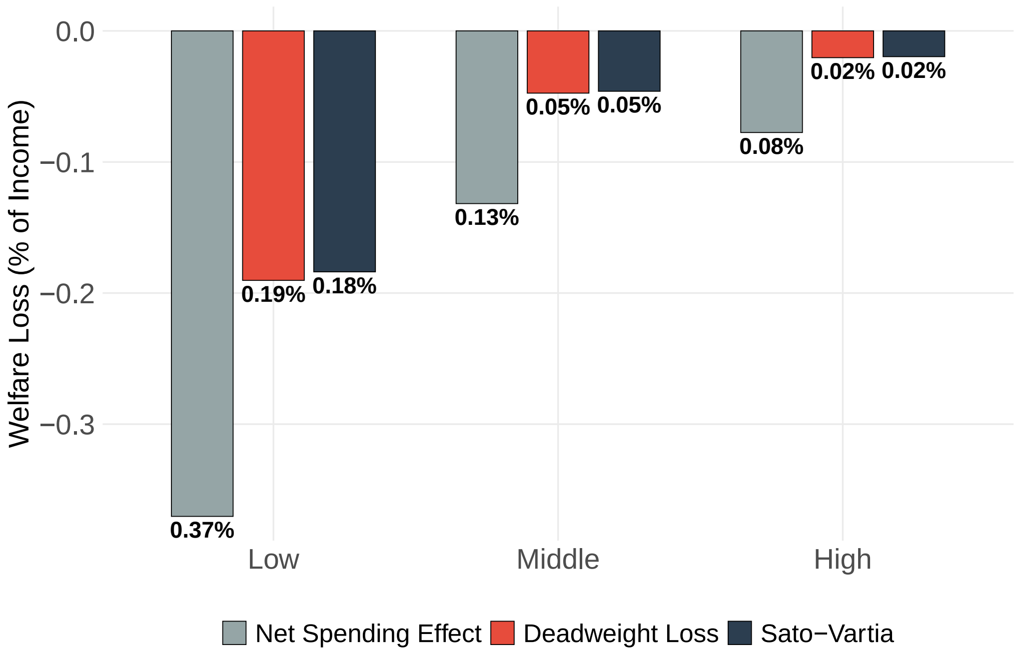

Notes: The figure presents three welfare measures: net spending effect, deadweight loss, and Sato-Vartia cost-of-living index, as a percentage of approximate household income. See Equations (9)–(15) for measure definitions. Income is approximated using bracket midpoints: \$20,000 (low), \$82,500 (middle), \$187,500 (high). The high-income midpoint is a lower bound on mean income in this group. The deadweight loss and Sato-Vartia bars are visually almost indistinguishable, reflecting the close agreement between the two price-anchored welfare measures.

respectively. Under both price-anchored welfare measures, low-income households bear roughly five and nine times the welfare cost of high-income households as a share of income. The net spending effect shows the same qualitative pattern, with low-income households facing 0.37 percent of income versus 0.13 and 0.08 percent for middle and high-income groups, respectively. High-income households reduce spending by more in dollar terms but can readily afford to do so; low-income households adjust by less but bear a proportionally much larger cost, in line with Acosta and Cox (2024).

The regressivity is driven by two forces documented in Table 4. First, price pass-through is monotonically decreasing in income: low-income households face pass-through of 0.28 per percentage point of tariff exposure, compared with 0.15 for high-income households. Second, low-income households have the smallest pre-tariff baseline income to absorb the burden. The combination produces a welfare cost that is modest in dollars but large in proportional terms.

Table 6 presents all three measures side by side, expressed both in dollars per month and as a share of spending on tariff-affected goods.

Table 6: Comparison of Welfare Measures by Income Group

		Monthly ($)			% of Spending		
Income Group	N HHs	Net Spend	DWL	Sato-Vartia	Net Spend	DWL	Sato-Vartia
Low	28,151	6.17	3.17	3.06	0.62	0.32	0.31
Middle	64,943	9.06	3.26	3.17	0.72	0.26	0.25
High	31,739	12.12	3.19	3.07	0.78	0.21	0.20
All	124,833	9.18	3.23	3.12	0.72	0.25	0.25

Notes: Net spending effect computed using Equation (9); DWL using Equations (10)–(12); Sato-Vartia using Equations (13)–(15). All measures use income-specific coefficients from Table 4. Mean annual spending of low, middle, high-income households respectively: $12,000, $15,044, $18,668.

Two features of Table 6 are worth emphasizing. First, the deadweight loss and Sato-Vartia measures agree closely. The Sato-Vartia measure equals approximately 95 percent of the deadweight loss across every income group, and the absolute gap never exceeds six cents per dollar of monthly welfare cost. This agreement is informative because the two measures share only the price pass-through coefficient in their construction. The DWL uses the estimated quantity coefficient with a linear demand approximation; the Sato-Vartia measure uses observed expenditure shares and imposes no functional form on demand. That two measures with overlapping but distinct identification margins land within five percent of each other in every income group provides convergent evidence that the price-anchored welfare cost is robust to the specific measurement assumption. Because the deadweight triangle is second-order in the price and quantity changes, the near-agreement is partly mechanical; nevertheless, the match confirms that the Sato-Vartia expenditure reallocation weights and the DWL quantity-elasticity approach produce consistent estimates of the first-order price burden.[16]

Second, the net spending effect substantially exceeds both price-anchored measures

[16]The welfare analysis applies income-specific but not essentiality-specific coefficients. Since non-essential goods exhibit larger spending and quantity responses than essentials (Table 5), using essentiality-specific coefficients would produce larger welfare numbers for non-essential categories, making the current approach conservative.

across all income groups. This gap is not a puzzle under the differential-identification framing. The price-anchored measures quantify the consumer surplus cost attributable to differential price pass-through across HS4 exposure, while the net spending effect captures the full differential behavioral response, including channels such as category-specific anticipation, uncertainty effects that operate more strongly on highly exposed categories, and precautionary reallocation across categories. Section 5 uses the household survey to show that the wedge between modest pass-through and large spending responses is driven by precautionary reallocation of the household basket toward essentials, not by intertemporal substitution or uniform demand contraction. The DWL and Sato-Vartia measures should therefore be read as a lower bound on the consumer welfare cost of the 2025 tariffs: they capture the component attributable to realized price changes, and the survey evidence in the next section documents the additional mechanism driving the broader consumption disruption.[17]

Overall, the welfare analysis yields three findings. First, the welfare burden of the 2025 tariffs is sharply regressive, with low-income households bearing approximately five to nine times the cost of high-income households as a share of income, depending on the measure. Second, this regressivity is robust across welfare measurement strategies, with the deadweight loss and Sato-Vartia measures agreeing to within five percent across every income group despite sharing only the price pass-through coefficient. Third, the gap between price-anchored welfare measures and the net spending effect points to a substantial behavioral response beyond realized pass-through, which Section 5 documents is driven by precautionary reallocation of household consumption.

5 Evidence from Household Survey Responses

The main results in Section 3 show that tariff-exposed categories experience large spending contractions alongside modest price pass-through. This section uses a survey run on a subset of households in the panel to explore what drives household behavior leading

[17] Note that a back-of-envelope upper bound on the uniform component absorbed by the year-month fixed effects, computed as the product of the mean tariff rate, mean import penetration, and the aggregate price coefficient applied to mean household spending, yields approximately \$14 per month, roughly four times the identified differential component of \$3.23 per month. The total welfare cost therefore likely lies between \$3 and \$17 per household per month, or roughly \$40 to \$200 per year.

to contractions in spending. The aggregate regressions estimate the total effect of tariff exposure on prices, quantities, and spending averaging across all households in an income group regardless of beliefs or intentions. The survey-based estimates ask a narrower question: among households facing the same tariff exposure in the same month, do those who are more worried or who report specific intent behave differently? By comparing households with different beliefs but the same exposure, we isolate the portion of the aggregate response that varies with sentiment from the portion common to all households. A group can therefore show a large total contraction with no sentiment differential if every household responds similarly, or a small total contraction with a large differential if only a subset of households drives the response.

Numerator fielded the survey in six biweekly waves between April 2025 and June 2025, asking households whether they were aware of the tariffs, how they expected the tariffs to affect their personal finances and the broader economy, how concerned they were about price increases, and what actions they planned to take in response. Throughout this section, we restrict the sample to households who both responded to the survey and remained in the panel throughout the analysis period. This merge yields 21,227 households observed monthly from January 2024 through December 2025. Appendix A compares the survey sample to the full sample and Appendix E.1 lists the full questionnaire.

Figure 5 shows that both stated concerns (Panel A) and anticipated behavioral reactions (Panel B) are stable across the six survey waves. Higher prices on everyday goods and general inflation dominate concerns in every wave. Cutting back on or delaying purchasing non-essentials and searching for discounts are the most common anticipated reactions. The stability in responses across the waves justifies pooling them and suggests that tariff-driven sentiment had crystallized by late April, before the bulk of the post-period price response materialized.

The headline finding of this section is that the wedge between modest price pass-through and large spending contraction reflects precautionary reallocation and trade-down within categories. Four pieces of evidence support this interpretation. First, pessimistic households shift their basket toward essentials by cutting non-essential spending while holding essential spending flat, with the cut concentrated in middle-income households who have slack to cut discretionary spending. Second, households who flag

Figure 5: Household Sentiment and Intentions from the Survey

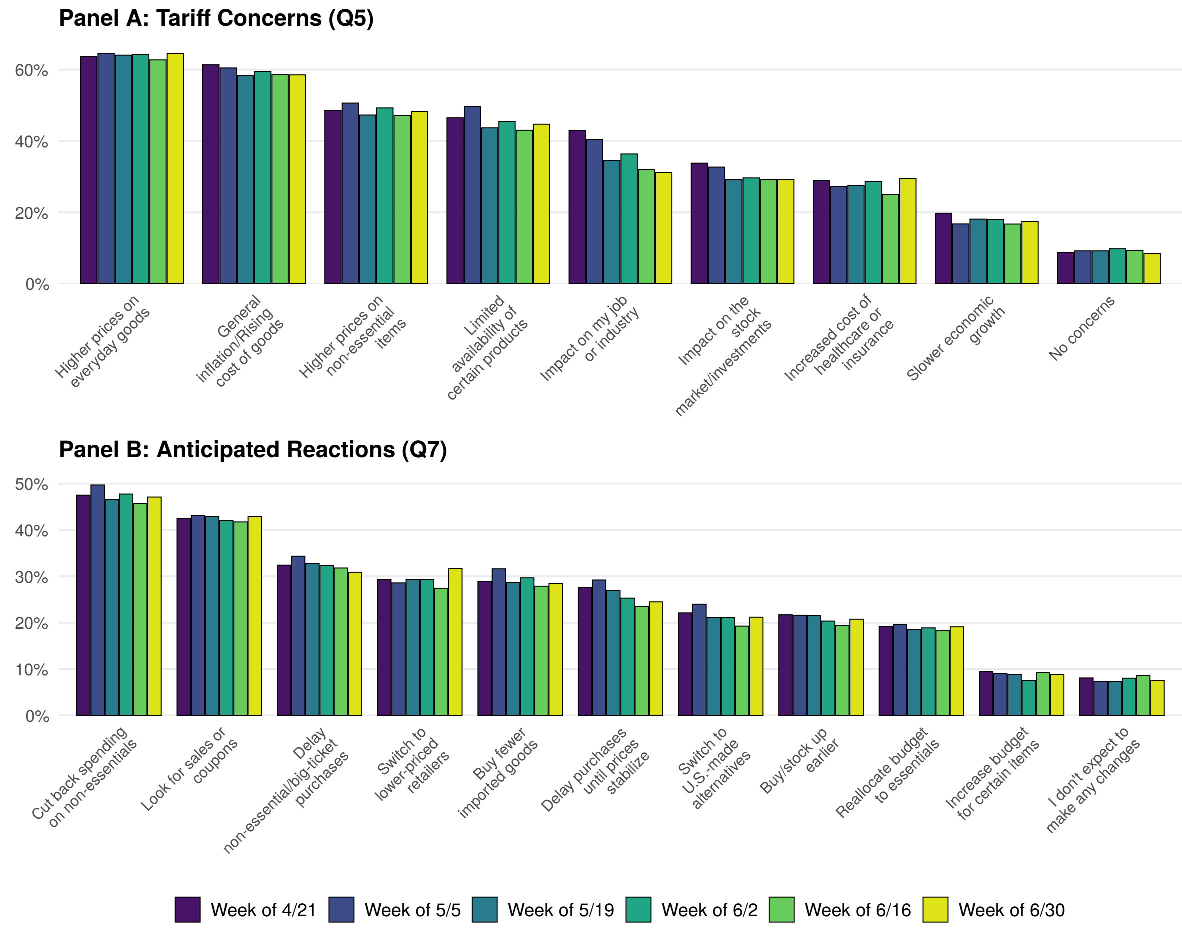

Notes: Panel A reports the share of respondents in the Numerator survey selecting each Q5 concern item in each of the six biweekly waves. Panel B reports the share selecting each Q7 anticipated-behavior item.

a specific category as tariff-affected cut spending on that category within HS4-month, with the cut concentrated on essential categories and with the quantity response indistinguishable from zero. Spending falls while units hold, so unit prices paid fall which indicates trading-down to cheaper varieties within the category of concern. Third, stated reduction intent predicts a differential cut in spending on tariff-exposed goods with no differential cut in quantity, again consistent with trade-down rather than quantity contraction. Stated substitution intent does not predict any differential response on the matched margin, and stated stockpile intent predicts a small, marginally significant positive quantity response with no spending response, an order of magnitude smaller than the aggregate quantity contraction and therefore inconsistent with a pull-forward story. Fourth, decomposing the stated reduction intent into its five components identifies stated reallocation toward essentials as the robust driver of the exposed-goods cut, with stated cutbacks as a secondary contributor and stated delay and delay of big-ticket purchases as null.

We do not find evidence that sentiment drives spending cuts specifically on tariff-exposed goods. Pessimism predicts significant reductions in spending on low-tariff goods, and the triple interaction between pessimism and tariff intensity is statistically indistinguishable from zero (Appendix E.3). This suggests that precautionary belt-tightening is broad-based rather than tied to realized tariff price increases, weighing against an interpretation in which the sentiment effect merely proxies for household-level pass-through.

The remainder of this section proceeds as follows: Section 5.1 documents the headline basket-reallocation result; Section 5.2 documents within-HS4 trade-down among worried households; and Section 5.3 maps stated intentions to revealed behavior. Section 5.4 summarizes the findings of this section.

Throughout this section, i indexes the household, h indexes the HS4 code, and t indexes the year-month. The treatment variable is $\text{TariffExposure}_{h,t}$ as defined in Equation (3); Post_t is the post-February-2025 indicator defined in Section 3.1. The survey panel is at the household-HS4 level rather than the panel-unit level used in Section 3, so the HS4 is the finest cross-sectional unit. A household-by-HS4 fixed effect $\alpha_{i,h}$ absorbs persistent household-product purchasing patterns, and an HS4-by-year-month fixed effect $\lambda_{h,t}$ absorbs all aggregate and category-level shocks at the HS4 level, including the tariff shock itself. Specifications that collapse to the household-month level use household fixed effects α_i and year-month fixed effects γ_t. Standard errors are clustered at the household level throughout. Finally, note throughout this section that the survey specifications include HS4-by-year-month fixed effects. So the aggregate spending and quantity declines documented in Section 3 are absorbed as a common response and the survey identifies the additional margin on which pessimistic households differ.

We measure household pessimism using four sentiment measures. Q6 asks about the perceived impact of tariffs on the U.S. economy over the next year (the only item explicitly anchored on tariffs); Q9 asks about expected macroeconomic conditions over the next year; Q10 asks about concern regarding recession in the coming year; Q5 is a count of tariff-related concerns selected by the household. Each item is standardized to have mean zero and unit variance across households. PC1 is the first principal component across the four items.[18] We use Sent_i as a generic notation for the household-level sentiment measure; the specific item is indicated in each table column. Essential_h is the HS4-level indicator defined in Section 3, covering spending on Grocery, Household, and Health & Beauty.

5.1 Precautionary Reallocation Toward Essentials

We first ask whether household pessimism predicts a shift in the composition of household spending toward essentials. Under uncertainty, households cut discretionary spending more sharply than essential spending, protecting baseline consumption while deferring non-essentials and durables (Coibion et al., 2024). In our set-up, this prediction operates on the composition of the basket rather than on the level of total spending. It does not require households to know the HS4-level tariff status of individual categories, only that they recognize tariffs as a source of aggregate uncertainty and reweight toward goods they regard as necessities.

We collapse the household-HS4-month panel to the household-month level and construct three outcomes for each household-month observation: total real spending $Y_{i,t}$, real spending on essential categories $Y_{i,t}^{\text{ess}}$, real spending on non-essential categories $Y_{i,t}^{\text{noness}}$, and the within-household share of spending on essentials, $\text{ShareEssential}_{i,t} = Y_{i,t}^{\text{ess}}/Y_{i,t}$, expressed in percentage points. We estimate

$$\text{ShareEssential}_{i,t} = \alpha_i + \gamma_t + \beta\,(\text{Sent}_i \cdot \text{Post}_t) + \varepsilon_{i,t}, \tag{16}$$

$$\ln Y_{i,t}^{k} = \alpha_i + \gamma_t + \beta\,(\text{Sent}_i \cdot \text{Post}_t) + \varepsilon_{i,t}, \quad k \in \{\text{ess, noness}\}, \tag{17}$$

[18] For each of the four sentiment items, we subtract the cross-household mean and divide by the cross-household standard deviation, so that each standardized item has mean zero and unit variance across households. PC1 is the first principal component of the resulting four-item matrix, computed by singular value decomposition on the standardized inputs. We sign the component so that it correlates positively with the Q6 tariff-impact item, and rescale it to have mean zero and unit variance across households; higher values therefore denote greater pessimism. PC1 explains 73.3 percent of the variance of the four inputs.

where α_i is a household fixed effect and γ_t is a year-month fixed effect. The sentiment main effect is absorbed by α_i and the post main effect is absorbed by γ_t, so β identifies whether pessimistic households shift their behavior in the post-period differentially relative to less-pessimistic households in the same month.

Table 7: Pessimism and Reallocation Toward Essentials

	Share essential (pp) (PC1)	Share essential (pp) (Q6 binary)	$\ln Y^{\text{ess}}$ (%) (PC1)	$\ln Y^{\text{noness}}$ (%) (PC1)
$\text{Sent}_i \cdot \text{Post}_t$ (PC1)	0.176***		0.001	−0.009**
	(0.065)		(0.002)	(0.004)
$\text{Sent}_i \cdot \text{Post}_t$ (Q6 binary)		0.324***		
		(0.124)		
Num. Obs.	382,866	457,863	377,845	377,845
R^2	0.460	0.461	0.635	0.551

Notes: The table reports estimates of Equations (16) (columns 1–2) and (17) (columns 3–4) on the household-month panel of survey respondents. Columns 1 and 2 use the within-household share of spending on essential categories in percentage points, $\text{ShareEssential}_{i,t} = 100 \cdot Y^{\text{ess}}_{i,t}/Y_{i,t}$. Columns 3 and 4 report log-level regression coefficients, so that entries multiplied by 100 are interpretable as approximate percentage changes. PC1 is the first principal component of standardized Q6, Q9, Q10, and Q5, signed so that higher values denote greater pessimism. Q6 binary equals one if the household reports a "somewhat negative" or "very negative" tariff impact in Q6. The level regressions restrict to households with positive essential and non-essential spending in every month. All specifications include household fixed effects α_i and year-month fixed effects γ_t. Standard errors clustered at the household level in parentheses. *, **, *** denote significance at the 10, 5, and 1 percent levels.

Table 7 reports the headline reallocation result. Pessimistic households increase the share of their spending devoted to essentials by 0.18 percentage points per standard deviation of PC1, and households reporting a negative tariff impact on Q6 increase their essential share by 0.32 percentage points. The level decomposition in columns 3 and 4 identifies where the share shift comes from. Non-essential spending falls by a significant 0.9 percent per standard deviation of PC1, while essential spending is statistically indistinguishable from zero. Pessimistic households defend essentials in absolute terms and cut non-essentials, and the share shift is the composition consequence of this asymmetric response. This pattern matches the heterogeneity-by-essentiality result documented in Section 3 that essentials exhibit no aggregate spending response to tariff exposure, precisely because households protect essential consumption when they contract.

Table 8 estimates the same specifications separately by income group. The non-essential level cut is concentrated in middle-income households (a significant 1.0 percent

Table 8: Reallocation Toward Essentials by Income Group

	Share essential (pp)			$\ln Y^{\text{ess}}$ (%)			$\ln Y^{\text{noness}}$ (%)		
	Low	Middle	High	Low	Middle	High	Low	Middle	High
$\text{Sent}_i \cdot \text{Post}_t$ (PC1)	0.216	0.157*	0.202	0.006	−0.001	−0.001	−0.010	−0.010**	−0.006
	(0.150)	(0.086)	(0.130)	(0.006)	(0.003)	(0.005)	(0.009)	(0.005)	(0.008)
Num. Obs.		378,834			373,839			373,839	
R^2		0.460			0.635			0.551	

Notes: The table reports estimates of Equations (16) and (17) estimated on the Numerator panel separately by income group, so that each cell reports the $\text{Sent}_i \cdot \text{Post}_t$ coefficient within the indicated income group. Columns 1–3 use $\text{ShareEssential}_{i,t}$ in percentage points; columns 4–9 report log-level regression coefficients, so that entries multiplied by 100 are interpretable as approximate percentage changes. Specifications, sample restrictions, and fixed effects as in Table 7. PC1 is the first principal component of standardized Q6, Q9, Q10, and Q5. Standard errors clustered at the household level in parentheses. *, **, *** denote significance at the 10, 5, and 1 percent levels.

decline per standard deviation of PC1) while low-income and high-income point estimates are smaller and not statistically significant. The magnitude of the middle-income share-essential shift matches the pooled estimate, consistent with this group driving the aggregate pattern. The pattern is consistent with the "slack, not budget" interpretation that runs through the rest of the paper. Low-income households face binding budget constraints that leave little room for additional discretionary contraction in response to pessimism. Middle-income households have the slack to execute the precautionary contraction on the discretionary margin.

5.2 Category-Specific Worry Predicts Category-Specific Trade-Down

We now turn to the within-HS4 belief-to-behavior result. Question Q8 asks households which categories of goods they are particularly concerned about being impacted by tariffs, with thirteen multi-select options spanning groceries, household goods, automobiles, electronics, apparel, and other major consumption categories. We map each Q8 category to the corresponding sectors in our HS classification (Appendix E.2) and construct, for each household-HS4-month observation, an indicator $\text{Worried}_{i,h}$ equal to one if household i flagged any Q8 category whose mapped sector contains HS4 code h. The specification is

$$Y_{i,h,t} = \beta_1 \left(\text{Worried}_{i,h} \cdot \text{Post}_t\right) + \beta_2 \left(\text{Worried}_{i,h} \cdot \text{Post}_t \cdot \text{Essential}_h\right) + \alpha_{i,h} + \lambda_{h,t} + \varepsilon_{i,h,t}, \quad (18)$$

where $Y_{i,h,t} \in \{\ln Y_{i,h,t}, \ln Q_{i,h,t}\}$ is log real spending or log quantity, $\alpha_{i,h}$ is a household-by-HS4 fixed effect, and $\lambda_{h,t}$ is an HS4-by-year-month fixed effect that absorbs all time-varying shocks at the category level, including the tariff shock itself. The pooled specification restricts $\beta_2 = 0$. In the essential split, β_1 captures the response in non-essential flagged categories and $\beta_1 + \beta_2$ captures the response in essential flagged categories.

TABLE 9: Category-Specific Worry (Q8) and Spending on Flagged Categories

	Pooled ($\beta_2 = 0$)		Essential split	
	Spend	Quantity	Spend	Quantity
$\text{Worried}_{i,h} \cdot \text{Post}_t$ (β_1)	-0.010^{***}	0.001	0.021^{***}	0.002
	(0.001)	(0.001)	(0.003)	(0.001)
$\text{Worried}_{i,h} \cdot \text{Post}_t \cdot \text{Essential}_h$ (β_2)			-0.035^{***}	-0.000
			(0.003)	(0.001)
Implied $\beta_1 + \beta_2$ (essentials)			-0.014^{***}	0.001
Num. Obs.	47,381,907	47,381,907	47,381,907	47,381,907
R^2	0.395	0.241	0.395	0.241

Notes: Estimates of Equation (18). The dependent variable is log real spending (columns 1, 3) or log quantity (columns 2, 4) on HS4 code h in month t. $\text{Worried}_{i,h} = 1$ if household i flagged any Q8 category mapped to a sector containing HS4 code h (mapping in Appendix E.2). All columns include household-by-HS4 ($\alpha_{i,h}$) and HS4-by-year-month ($\lambda_{h,t}$) fixed effects; category-level and aggregate post-period shocks are absorbed. In the split, β_1 gives the response in non-essential flagged categories and $\beta_1 + \beta_2$ the response in essential flagged categories, with sum significance computed from the joint variance. Standard errors clustered at the household level in parentheses. *, **, *** denote significance at the 10, 5, and 1 percent levels.

Table 9 reports the results. Pooled, worried households cut spending on flagged categories by a significant 1.0 percent relative to unworried households in the same HS4-month, with a quantity response indistinguishable from zero. The split reveals the mechanism. In essential flagged categories, the implied spending response is a significant 1.4 percent cut, and the quantity response remains statistically indistinguishable from zero. Spending falls, units hold, so unit prices paid fall, which is the signature of trade-down to cheaper varieties within the flagged category. In non-essential flagged categories the pattern inverts. Spending rises by a significant 2.1 percent and quantity is unchanged, indicating that worried households absorb higher unit prices rather than trading down, consistent with the narrower within-category substitution margin typically available in non-essential categories relative to essentials. The asymmetry across essentials and non-essentials within Table 9 therefore operates on the intensive margin conditional on continued category participation, while the aggregate non-essential spending cut in Table 7 operates on the extensive margin of category participation. So, pessimistic households exit or reduce engagement with non-essential categories in aggregate, but within the non-essential categories they continue to purchase, they absorb higher unit prices rather than trading down. The key observation is that in both cases the quantity does not change. Worried households do not pull forward purchases in flagged categories on either essential or non-essential margins, providing evidence against category-specific stockpiling as the channel.

The essential trade-down pattern here is the household-level counterpart of the aggregate heterogeneity-by-essentiality result in Section 3. Table 5 shows no aggregate spending or quantity response for essentials and the channel here shows why. The aggregate null reflects two offsetting forces: worried households trade down to cheaper varieties within essentials, reducing measured spending, while the quantity of units purchased holds steady. What appears as no response in the aggregate is, at the household level, a compositional shift within the category that leaves total units unchanged. The non-essential spending-increase-without-quantity-response result, in turn, provides a within-category analog to the aggregate pass-through result such that worried households continue purchasing non-essential items at higher unit prices rather than substituting along a quantity or variety margin.

Income-group breakdowns of the pooled specification are reported in Appendix E.5; the pooled within-HS4 coefficient is essentially uniform across the income distribution, consistent with trade-down as a universally available margin at the HS4 level.

5.3 Stated Intentions and Revealed Behavior

We now examine whether household-stated intentions in the post-tariff survey predict realized spending and quantity behavior on tariff-exposed goods. Question Q7 elicits which actions households plan to take in response to expected price increases (multi-select across reduce, substitute to U.S.-made, switch retailers, reallocate across categories, stockpile, and buy fewer imports). For each intent, we estimate

$$Y_{i,h,t} = \beta\left(\text{TariffExposure}_{h,t} \cdot \text{Post}_t \cdot \text{Intent}_i\right) + \alpha_{i,h} + \lambda_{h,t} + \varepsilon_{i,h,t} \tag{19}$$

at the household-HS4-month level, where $Y_{i,h,t} \in \{\ln Y_{i,h,t}, \ln Q_{i,h,t}\}$ and Intent_i is a household-level binary indicator for the stated intent in question. The two main effects of the triple interaction ($\text{TariffExposure}_{h,t} \cdot \text{Post}_t$ and $\text{Intent}_i \cdot \text{Post}_t$) are absorbed by $\lambda_{h,t}$ and $\alpha_{i,h}$ respectively, so the triple-interaction coefficient β identifies whether households who stated a given intent adjusted differentially on tariff-exposed categories in the post-period relative to households who did not state that intent. Estimating the specification on both spending and quantity isolates whether any differential response operates through basket composition (spending without quantity) or genuine consumption contraction (spending with quantity).

Table 10: Stated Intentions (Q7) and Revealed Behavior on Tariff-Exposed Goods

	Spend			Quantity		
	Reduce	Substitute	Stockup	Reduce	Substitute	Stockup
$\text{TariffExposure}_{h,t} \cdot \text{Post}_t \times \text{Reduce}_i$	-0.044^{**}			-0.003		
	(0.018)			(0.008)		
$\text{TariffExposure}_{h,t} \cdot \text{Post}_t \times \text{Substitute}_i$		0.009			-0.002	
		(0.017)			(0.008)	
$\text{TariffExposure}_{h,t} \cdot \text{Post}_t \times \text{Stockup}_i$			0.005			0.018^{*}
			(0.021)			(0.010)
Num. Obs.	47,381,907	47,381,907	47,381,907	47,381,907	47,381,907	47,381,907
R^2	0.395	0.395	0.395	0.241	0.241	0.241

Notes: Each column is a separate estimation of Equation (19) on the household-HS4-month panel of survey respondents. The dependent variable in columns 1–3 is $\ln Y_{i,h,t}$, the natural log of real (PCE-deflated) household spending on HS4 code h in month t; in columns 4–6 it is $\ln Q_{i,h,t}$, the natural log of real quantity purchased. Reduce_i is the union of the five Q7 reduction-related items; Substitute_i is the union of substitution to U.S.-made and retailer switching; Stockup_i is stated stockpile intent. All specifications include household-by-HS4 ($\alpha_{i,h}$) and HS4-by-year-month ($\lambda_{h,t}$) fixed effects. Standard errors clustered at the household level in parentheses. *, **, *** denote significance at the 10, 5, and 1 percent levels.

Table 10 reports the results. Only stated reduction intent translates into a differential cut on tariff-exposed goods, a significant 4.4 percent decline per unit exposure, while stated substitution and stockpile intent do not affect spending. The quantity columns are decisive for interpretation. The stated reducer's spending cut is accompanied by a quantity response of −0.003, indistinguishable from zero, so stated reducers cut measured spending without cutting units. This pattern indicates trading-down to cheaper varieties within exposed HS4 codes, parallel to the within-worried-category pattern documented in Section 5.2. Item-level data from the full household panel corroborate this interpretation by showing that within tariff-exposed HS4 codes, cheap items gain revenue share at the expense of expensive items, and the reallocation is monotonically decreasing in the within-category price percentile (Appendix E.8, Tables E.9–E.10). Stated substitution intent is null on both margins, an outcome informative in its own right: substitution-related responses dominate stated intentions in the survey (Figure 5, Panel B), yet produce no differential response on the exposed-good revealed margin. Stated stockpile intent is null on spending but delivers a marginally significant positive quantity response of 1.8 percent per unit exposure. The coefficient is small relative to the aggregate quantity contraction documented in Section 3 to explain the wedge so we conclude that modest pull-forward exists in the stated-stockpiler subsample but accounts for an order of magnitude less than the aggregate quantity decline.

The pooled Reduce_i indicator aggregates five Q7 sub-items: cut back spending on non-essentials, delay non-essential or big-ticket purchases, delay purchases until prices stabilize, buy fewer imported goods, and reallocate budget across categories to prioritize essentials. We construct an additional index Precaution_i that aggregates the four sub-items most directly consistent with a precautionary response (cut back, delay, delay big-ticket, and reallocate), excluding buy-fewer-imports, and re-estimate Equation (19) with each sub-item and the composite entered separately.

Table 11 reports the results. Three patterns stand out. First, stated reallocation is the single largest and most precisely estimated driver of the exposed-good spending cut at a significant 4.5 percent, with stated cutback on non-essentials providing a smaller, marginally significant secondary contribution of 3.0 percent. The precaution composite delivers a coefficient of a significant 4.3 percent that is close to the reallocate single-item estimate and within sampling variation of the pooled Reduce_i coefficient in Table 10, confirming that the pooled reduction effect is driven almost entirely by the precautionary-reallocation subset of the five sub-items. A horserace specification that

TABLE 11: Decomposition of Stated Reduction Intent: Spending and Quantity

	Cut back	Delay	Delay big-tkt	Fewer imports	Reallocate	Precaution
Panel A: Spend, $\ln Y_{i,h,t}$						
$\text{TariffExposure}_{h,t} \cdot \text{Post}_t \times \text{Intent}_i$	−0.030*	−0.018	−0.011	0.012	−0.045**	−0.043**
	(0.017)	(0.019)	(0.018)	(0.018)	(0.021)	(0.017)
Panel B: Quantity, $\ln Q_{i,h,t}$						
$\text{TariffExposure}_{h,t} \cdot \text{Post}_t \times \text{Intent}_i$	0.004	0.003	0.005	−0.016*	0.016	0.004
	(0.008)	(0.009)	(0.008)	(0.008)	(0.010)	(0.008)
Num. Obs.	47,381,907					
R^2 (spend / qty)	0.395 / 0.241					

Notes: Each column-panel cell reports a separate estimate of Equation (19) with Intent_i set to the indicated Q7 sub-item. "Cut back" is Q7 item 9 (cut back spending on non-essentials); "Delay" is item 4 (delay until prices stabilize); "Delay big-tkt" is item 11 (delay non-essential or big-ticket purchases); "Fewer imports" is item 5 (buy fewer imported goods); "Reallocate" is item 12 (reallocate budget toward essentials); "Precaution" is the household-level indicator equal to one if the household selected any of cut back, delay, delay big-ticket, or reallocate (i.e., the four reduction-related items excluding buy-fewer-imports). Panel A uses $\ln Y_{i,h,t}$ as the dependent variable; Panel B uses $\ln Q_{i,h,t}$. All specifications include household-by-HS4 ($\alpha_{i,h}$) and HS4-by-year-month ($\lambda_{h,t}$) fixed effects. A horserace specification with all five sub-items entered simultaneously is reported in Appendix E.4, Table E.5. Standard errors clustered at the household level in parentheses. *, **, *** denote significance at the 10, 5, and 1 percent levels.

enters all five sub-items simultaneously (Appendix E.4) confirms that reallocate is the only sub-item to retain marginal significance when the five compete for variance, reinforcing its status as the primary channel. Second, the two timing-related sub-items, delay and delay of big-ticket purchases, are insignificant on both margins. The null on the quantity margin is particularly informative because if households who reported intending to delay purchases were acting on that intent, we would expect a negative quantity response on tariff-exposed goods, and we find none. This provides further evidence against intertemporal substitution, complementing the analysis in Section 3.5. Third, the buy-fewer-imports sub-item exhibits a distinct pattern: an insignificant spending response of 1.2 percent alongside a marginally significant negative quantity response of 1.6 percent. Stated import-substituters reduce units of tariff-exposed goods while paying roughly the same aggregate amount, an outcome consistent with substitution toward U.S.-produced alternatives at higher unit prices. The pattern is distinct from the precautionary trade-down documented elsewhere in the section and identifies a separate behavioral margin.

Taken together, the belief-driven component of the response operates entirely through spending composition: on the margin that distinguishes pessimistic from non-pessimistic households, spending on tariff-exposed goods falls without corresponding quantity reductions, consistent with trading down within categories rather than buying less. The aggregate quantity contraction documented in Section 3 is a common response shared by all households regardless of stated beliefs, while the survey identifies the additional channel through which beliefs shape the composition of spending. Stated reducers, reallocators, cutbackers, and precautionary households all exhibit the trade-down pattern. Stockpilers show a small positive quantity effect without a spending response, while delay and big-ticket postponement channels are null throughout. These patterns identify precautionary reallocation combined with within-category trade-down as the primary mechanism behind the disproportionate spending contractions documented in the aggregate results, with stockpiling and intertemporal delay playing at most minor roles.

5.4 Summary and Discussion

Recent evidence by Candia et al. (2026) demonstrates that tariff uncertainty and expected tariff levels have opposing effects on durable goods. Higher expected tariff levels encourage intertemporal substitution toward current purchases, while greater tariff uncertainty discourages spending. In the 2025 tariff episode, where policy volatility was exceptionally high, the precautionary channel appears to have dominated. Consistent with this, the survey evidence in our analysis identifies precautionary reallocation toward essentials, combined with within-category trade-down to cheaper varieties, as the behavioral mechanism behind the wedge between modest price pass-through and large spending contractions. The two margins have different income profiles. Within-category trade-down is shared by all income groups. In categories households flag as tariff-affected, spending declines while quantity holds regardless of income, indicating that substitution toward cheaper varieties is universally available at the product level. The across-category margin is not. Pessimistic households cut non-essential spending while holding essential spending flat, shifting the composition of their basket rather than contracting total consumption uniformly, but this level cut on non-essentials is concentrated among middle-income households where discretionary slack permits the adjustment. Among stated behavioral intentions, only reduction intent maps onto revealed behavior on tariff-exposed goods, and it does so entirely through the spending margin with no quantity response. Stated substitution and stockpiling intent are null or economically negligible. Low-income households show no sentiment differential on either margin, consistent with budget constraints that leave limited room for compositional adjustment regardless of beliefs.

A natural question is why high-income households, who exhibit significant aggregate spending and quantity contractions in Section 3, show no differential response on any survey-based margin. Appendix Table E.11 estimates Equation (19) separately by income group for stated reduction intent, the precautionary composite, and stated reallocation. Middle-income stated reducers cut spending on tariff-exposed goods by a significant 6.5 percent per unit exposure with no quantity response, replicating the pooled trade-down pattern. High-income estimates are small and statistically insignificant across all three intent measures. One interpretation is that the high-income contraction documented in the aggregate results is a common response shared by high-income households regardless of their stated beliefs or intentions, leaving limited within-group variation for the survey interactions to detect. This is consistent with high-income households having sufficient financial capacity to adjust spending patterns in response to the tariff shock without the adjustment being contingent on the degree of pessimism or on specific stated plans.

A potential concern with the survey-based analysis is that sentiment is measured contemporaneously with the treatment because the survey was fielded between April and June 2025, after tariffs were implemented. Households who experienced larger price increases may report greater pessimism precisely because they were more exposed, raising a reverse-causality concern. Three features of the results mitigate this concern. First, if sentiment were simply proxying for realized exposure, generic macroeconomic pessimism (Q9, Q10) should also amplify the tariff response, since households experiencing larger price shocks would plausibly report worse macroeconomic outlooks as well. Yet, Q6 loads with the largest magnitude among the sentiment measures. Second, the within-household-across-HS4 specifications absorb the household's total realized experience through $\alpha_{i,h}$, so the interactions identify differential reallocation across categories conditional on the household's overall exposure. Third, in Appendix E.7 we re-estimate the Q6 sentiment specifications controlling for the household's basket-weighted realized price change between the pre- and post-periods. The sentiment coefficient remains unchanged in magnitude and significance, confirming that tariff-anchored pessimism does not proxy for realized price exposure (Table E.8). Finally, in addition to the survey-based evidence, although our post-period begins in February 2025, tariff implementation remained volatile during our survey window. Liberation Day tariffs were announced on April 2 and partially paused shortly thereafter, with sector-specific exemptions and reinstatements continuing through June. The survey waves therefore span a period of active policy uncertainty rather than a stable tariff implementation and realized-price environment, making it unlikely that respondents' sentiment simply reflected already-observed price changes for their baskets.

6 Conclusion

This paper uses transaction-level data from a large panel of U.S. households, linked to a subset with survey-reported tariff awareness, sentiment, and behavioral intentions, to study how the 2025 U.S. tariffs reshape household spending. The joint observation of prices, quantities, stated intentions, and revealed behavior from a single source allows us to document both the aggregate response and the behavioral channel behind it. Tariff exposure raises retail prices modestly, with pass-through monotonically decreasing in income, but reduces real spending by a far larger margin. At the aggregate level, both quantities and spending decline sharply, and the contraction is sustained rather than front-loaded, weighing against intertemporal substitution as the primary channel. The wedge between modest pass-through and large spending contractions cannot be rationalized by price pass-through alone. Survey evidence suggests an additional channel: conditional on the common quantity contraction that all households share, pessimistic households further reallocate their baskets toward essentials and trade down to cheaper varieties within tariff-exposed categories. This belief-driven margin operates through spending composition rather than through additional quantity adjustment. Stated stockpile and delay intentions do not predict corresponding revealed behavior.

The tariff operates on households through two channels with different income profiles. A price-incidence channel dominates at the bottom of the income distribution. Low-income households face the highest pass-through and allocate a larger share of their baskets to tariff-exposed goods, producing a welfare cost much greater as a share of income than for high-income households. A precautionary channel operates through the composition of spending rather than its total magnitude, and itself splits into two margins. At the product level, households across the income distribution switch to cheaper varieties within tariff-affected categories, a within-category trade-down that is universal. At the basket level, only middle-income households go further and cut the level of non-essential spending, generating the across-category shift toward essentials. Low-income households show no sentiment differential, consistent with budget constraints that leave limited room for compositional adjustment.

References

Acosta, M. and L. Cox (2024). The regressive nature of the US tariff code: Origins and implications. Working Paper.

Amiti, M., C. Flanagan, S. Heise, and D. E. Weinstein (2026, February). Who is paying for the 2025 u.s. tariffs? *Liberty Street Economics*. Federal Reserve Bank of New York blog post.

Amiti, M., S. J. Redding, and D. E. Weinstein (2019). The impact of the 2018 tariffs on prices and welfare. *Journal of Economic Perspectives 33*(4), 187–210.

Amiti, M., S. J. Redding, and D. E. Weinstein (2020). Who's paying for the US tariffs? a longer-term perspective. *AEA Papers and Proceedings 110*, 541–546.

Argente, D. and M. Lee (2021, 04). Cost of living inequality during the great recession. *Journal of the European Economic Association 19*(2), 913–952.

Autor, D., D. Dorn, and G. H. Hanson (2021). On the persistence of the China shock. *Brookings Papers on Economic Activity*.

Autor, D. H., D. Dorn, and G. H. Hanson (2013). The China syndrome: Local labor market effects of import competition in the United States. *American Economic Review 103*(6), 2121–2168.

Azzimonti, M. (2025, August). Why predicted and actual tariff rates diverged in may 2025. *Economic Brief* (25-29).

Barbiero, O. and H. Stein (2025, February). The impact of tariffs on inflation. Current Policy Perspectives 25-2, Federal Reserve Bank of Boston.

Borusyak, K. and X. Jaravel (2021). The distributional effects of trade: Theory and evidence from the United States. *NBER Working Paper No. 28957*.

Braun, R., A. Flaaen, and S. Hacıoğlu Hoke (2024, February). Supply vs demand factors influencing prices of manufactured goods. FEDS Notes. Washington: Board of Governors of the Federal Reserve System.

Caliendo, L. and F. Parro (2023). Lessons from US–China trade relations. *Annual Review of Economics 15*, 513–547.

Candia, B., J. Mitchell, and D. Pfajfar (2026, February). The causal effects of tariff uncertainty on consumers' macroeconomic expectations and spending plans. Working Paper 26-05, Federal Reserve Bank of Cleveland.

Cavallo, A., G. Gopinath, B. Neiman, and J. Tang (2021a). Tariff pass-through at the border and at the store: Evidence from US trade policy. *American Economic Review: Insights 3*(1), 19–34.

Cavallo, A., G. Gopinath, B. Neiman, and J. Tang (2021b, March). Tariff Pass-Through at the Border and at the Store: Evidence from US Trade Policy. *American Economic Review: Insights 3*(1), 19–34.

Cavallo, A. and O. Kryvtsov (2024). Price discounts and cheapflation during the post-pandemic inflation surge. *Journal of Monetary Economics 148*, 103644.

Cavallo, A., P. Llamas, and F. Vazquez (2025, November). Tracking the short-run price impact of US tariffs. Working Paper 34496, National Bureau of Economic Research.

Coibion, O., D. Georgarakos, Y. Gorodnichenko, G. Kenny, and M. Weber (2024). The effect of macroeconomic uncertainty on household spending. *American Economic Review 114*(3), 645–677.

D'Acunto, F., U. Malmendier, J. Ospina, and M. Weber (2021). Exposure to grocery prices and inflation expectations. *Journal of Political Economy 129*(5), 1615–1639.

DeDad, M. and S. Ghosh (2026). Us-china trade war: Heterogeneous effects on the us consumer. *Canadian Journal of Economics*. Forthcoming.

den Besten, T., R. Barnichon, D. R. Känzig, and A. Singh (2026, April). The macroeconomic effects of tariffs: Insights from 180 years of u.s. trade policy. Working Paper 35102, National Bureau of Economic Research.

Dvorkin, M. A., F. Leibovici, and A. M. Santacreu (2025, October). How tariffs are affecting prices in 2025. Federal Reserve Bank of St. Louis, On the Economy Blog.

Fajgelbaum, P. D., P. K. Goldberg, P. J. Kennedy, and A. K. Khandelwal (2020). The return to protectionism. *Quarterly Journal of Economics 135*(1), 1–55.

Fajgelbaum, P. D. and A. K. Khandelwal (2016). Measuring the unequal gains from trade. *The Quarterly Journal of Economics 131*(3), 1113–1180.

Feenstra, R. C. (1994). New product varieties and the measurement of international prices. *American Economic Review 84*(1), 157–177.

Feng, C., L. Han, and L. Li (2023). Who pays for the tariffs and why? a tale of two countries. Working Paper.

Ganapati, S. and C. J. Hottman (2026, February). Did foreigners pay america's tariffs? quantity discounts, scale economies and incomplete pass-through. Working Paper 34901, National Bureau of Economic Research.

Gopinath, G. and B. Neiman (2026, January). The incidence of tariffs: Rates and reality. Working Paper 34620, National Bureau of Economic Research.

Hacıoğlu-Hoke, S., S. Malladi, and L. Feler (2026, March). The slow climb: How tariffs gradually raised retail prices in 2025. FEDS notes, Board of Governors of the Federal Reserve System, Washington.

Hacıoğlu Hoke, S., L. Feler, and J. Chylak (2024, October). A Better Way of Understanding the US Consumer: Decomposing Retail Sales by Household Income. FEDS Notes. Washington: Board of Governors of the Federal Reserve System.

Hacıoğlu Hoke, S., L. Feler, and J. Chylak (2025). Lost in aggregation: Geographic mismeasurement of income and spending. CEPR Discussion Paper 20428, CEPR Press.

Hinz, J., A. Lohmann, H. Mahlkow, and A. Vorwig (2026, January). America's own goal: Who pays the tariffs? Technical Report 201, Kiel Institute for the World Economy.

Jaravel, X. and E. Sager (2025). What are the price effects of trade? Evidence from the U.S. and implications for quantitative trade models. *Review of Economics and Statistics*. Conditionally accepted.

Ma, H., L. Macedoni, J. Ning, and M. J. Xu (2025). Tariffs tax the poor more: Evidence from household consumption during the us-china trade war. CESifo Working Paper Series 11610, CESifo.

Minton, R. and M. Somale (2025, May). Detecting tariff effects on consumer prices in real time. FEDS Notes. Washington: Board of Governors of the Federal Reserve System.

Pierce, J. R. and P. K. Schott (2016). Surprisingly swift decline in US manufacturing employment. *American Economic Review 106*(7), 1632–1662.

Porto, G. G. (2006). Using survey data to assess the distributional effects of trade policy. *Journal of International Economics 70*(1), 140–160.

Rambachan, A. and J. Roth (2023). A more credible approach to parallel trends. *Review of Economic Studies 90*(5), 2555–2591.

Redding, S. J. and D. E. Weinstein (2020). Measuring aggregate price indices with taste shocks: Theory and evidence for CES preferences. *Quarterly Journal of Economics 135*(1), 503–560.

Sangani, K. (2026, May). Complete pass-through in levels. *The Quarterly Journal of Economics 141*(2), 1077–1135.

Sato, K. (1976). The ideal log-change index number. *Review of Economics and Statistics 58*(2), 223–228.

Shapiro, A. H. (2024). Decomposing supply- and demand-driven inflation. *Journal of Money, Credit and Banking*.

Topalova, P. (2010). Trade liberalization and firm productivity: The case of India. *The Review of Economics and Statistics 92*(3), 765–783.

Vartia, Y. O. (1976). Ideal log-change index numbers. *Scandinavian Journal of Statistics 3*(3), 121–126.

Waugh, M. E. (2019). The consumption response to trade shocks: Evidence from the US-China trade war. *NBER Working Paper No. 26353*.

Weber, M., F. D'Acunto, Y. Gorodnichenko, and O. Coibion (2022). The subjective inflation expectations of households and firms: Measurement, determinants, and implications. *Journal of Economic Perspectives 36*(3), 157–184.

Online Appendix

A More Detail on the Numerator Household Panel

This section provides technical details on the Numerator data collection, panel composition, weighting methodology, and data structure.

A.1 Data Collection and Panel Composition

Numerator collects purchase data through a mobile app called Receipt Hog, which enables consumers to upload photographs of paper receipts, connect email accounts for automatic digital receipt collection, and link retailer apps including Amazon, Walmart, and similar platforms. The only requirement from panelists is to own a mobile phone with a camera. Participants receive approximately $10 monthly in compensation through gift cards or direct payment. Receipt Hog has over a million registered users.

Our analysis uses Numerator's Static Panel households, a subset of the broader panel selected for consistent reporting behavior. Static Panelists must meet the following criteria: (1) maintain active status for a minimum of 12 consecutive months, (2) submit purchases consistently across the period with no extended gaps in reporting, and (3) meet Numerator's internal quality thresholds for transaction completeness, e.g. an average panelist records 32 transactions a month. The Static Panel designation is important because it enables accurate longitudinal tracking of household behavior without compositional changes that could confound difference-in-differences estimates. The static panel has grown to include up to 200,000 panelists as of 2025, up from approximately 150,000 in early 2024. Our analysis further restricts to the 126,448 households with complete data coverage from January 2024 through December 2025, ensuring balanced panel composition throughout the tariff implementation period.

A.2 Weighting and Calibration Methodology

Numerator applies a two-stage weighting procedure to ensure representativeness. The first stage applies demographic weights to align the sample with U.S. Census benchmarks using raking (iterative proportional fitting) across eight demographic dimensions simultaneously: gender of head of household, age, ethnicity, household income, household size, presence of children under 18 living at home, Census division, and urbanicity. Figure A.1 compares the Numerator panel to American Community Survey (ACS) targets across these eight dimensions as of March 2025. The weighted panel matches ACS benchmarks within 2 percentage points for all dimensions. The close alignment between the weighted panel and ACS targets validates our spending measures as representative of the broader population.

Figure A.1: Numerator Panel Demographics

Note: Numerator Total Commerce Static Panel weighted as of March 2025. Values shown as percentages. Blue bars show Numerator panelists before applying demographic weights; red bars show the panel after demographic weights are applied; yellow bars shows the American Community Survey (ACS) implied demographics in 2024.

The second stage applies calibration weights (national_factor) to scale aggregate spending to match external benchmarks. Numerator calibrates total dollar spend to the following sources: (1) Census Bureau Retail Sales (excluding motor vehicles and parts), (2) public company filings (10-K and 10-Q reports from major retailers), (3) direct retail relationships and partnerships, and (4) third-party data partnerships (e.g., location data providers). Additionally, for Static Panel households, Numerator applies a trend factor adjustment to account for missing buyers within product categories, ensuring that the panel represents not just frequent purchasers but also occasional buyers in each category. All weights are updated quarterly to reflect changes in panel composition and external benchmark data. This second stage weighting is not used in our main analyses but it is used in our analysis in Appendix A.5 that demonstrates Numerator's aggregated data performance against published statistics.

A.3 Data Structure and Core Identifiers

The Numerator disaggregated data feed is structured around three core identifiers: (1) user_id: anonymized household identifier linking all purchases by a given household across time, (2) basket_id: shopping trip/transaction identifier representing a single receipt or online order, and (3) item_id: product identifier at the most granular level, where different package sizes and formats receive distinct IDs (e.g., a single can versus a 12-pack of the same beverage). The data are delivered in six primary tables that can be joined using these identifiers: fact_table (transaction and trip details), item_table (product dimension with brand, category, department, sector hierarchy), static_table (Static household roster), people_table (core demographics), people_attributes_table (extended attributes such as psychographics), and banner_table (retailer information). This relational structure allows flexible aggregation from the transaction-item level to household-month or category-month panels while preserving the underlying granularity for detailed analysis.

A.4 Income Measurement

Numerator panelists self-report their annual household income in categorical brackets rather than exact dollar amounts. The categories are: less than \$20,000 through \$100,000 in \$10,000 increments, and \$100,000 through \$250,000 in \$25,000 increments, with a top-coded category for \$250,000 and above. This categorical approach minimizes measurement error from inaccurate reporting of exact earnings while retaining sufficient variation for distributional analysis. Household income information is updated through annual resurveys, with additional triggered updates when panelists indicate significant life events such as employment transitions during quarterly check-ins (approximately every 3 months). When resurveyed, panelists indicate whether they have moved to a different income bracket. The annual update frequency with event-triggered adjustments ensures that measurement error from income changes remains within acceptable bounds for our 24-month analysis window.

Finally, Table A.1 show some sample statistics from the full household panel and the survey panel.

A.5 Proof of Concept

To validate that the Numerator household transaction data reliably captures the spending and price dynamics faced by U.S. consumers, we construct an aggregate retail sales spending measure and monthly price index spanning 2018 to 2025, and compare them with the officially published statistics. What follows in this section borrows from our earlier analyses in Hacıoğlu Hoke et al. (2024) and Hacıoğlu Hoke et al. (2025) where the reader can refer to for the details.

Figure A.2 plots three lines. The solid red line indicates the Census Bureau's published seasonally-adjusted total retail sales excluding motor vehicles and parts. The blue dashed line is the retail spending measure we construct from the Numerator data, which we seasonally adjust using the Census Bureau's X13-ARIMA-SEATS package. The gray dash-dotted line is the non-seasonally adjusted Census series that we then seasonally adjust using the same X13-ARIMA-SEATS methodology we use for Numerator data. The purpose of the gray line is to show that our seasonal adjustment methodology aligns closely with the Census methodology, and therefore, the differences between the Numerator seasonally-adjusted series and the official Census seasonally-adjusted series is not due to differences in seasonal adjustment. The correlation between the Numerator measure and the official seasonally-adjusted retail sales measure is 0.993. This analysis allows us to verify that the bottom-up construction of the Numerator data, directly from household purchases, closely matches the top-down construction of the Census retail sales series, captured from surveys of retail and food service establishments.

Table A.1: Descriptive Statistics: Full Sample vs. Survey Sample

Variable		Full Sample		Survey Sample	
		%	N	%	N
Age	18-20	0.02	23	0.03	6
	21-24	1.73	917	1.31	278
	25-34	10.5	13271	13.49	2864
	35-44	18.39	23248	17.07	3623
	45-54	18.93	23940	14.43	3062
	55-64	21.23	26841	19.11	4056
	65+	30.22	38208	34.57	7338
Education	College	60.81	76896	61.15	12980
	Advanced degree	18.36	23214	18.54	3935
	High school	18.19	23005	18.23	3869
	Less than high school	2.64	3333	2.09	443
Employment Status	Employed full-time	49.06	62040	46.13	9791
	Retired	24.74	31286	28.24	5995
	Employed part-time	8.09	10227	8.45	1794
	Homemaker	5.12	6471	4.81	1020
	Self-employed	5.04	6371	4.81	1020
	Disabled	3.43	4332	3.1	659
	Unemployed	3.02	3815	2.89	614
	Student	1.04	1317	1.18	251
	Active military	0.47	589	0.39	83
Income	Less than $40,000	22.26	28151	23.2	4925
	$40,000-$124,999	51.36	64943	53.71	11400
	$125,000 or more	26.38	33354	23.09	4902
Race/Ethnicity	White/Caucasian	68.44	86538	71.72	15225
	Hispanic/Latino	12.99	16428	10.57	2243
	Black/African American	11.61	14675	10.99	2333
	Asian	5.53	6991	5.23	1110
	Other	1.44	1816	1.49	316
Marital Status	Married	58.06	73419	56.13	11914
	Never married	16.37	20702	17.89	3797
	Divorced	12.31	15560	12.6	2675
	Widowed	5.97	7543	6.13	1301
	Living with partner	5.71	7222	5.7	1211
	Separated	1.58	2002	1.55	329
Has Children	No	69.23	87541	72.98	15492
	Yes	30.77	38905	27.02	5735
Household Size	1	19.45	24591	20.37	4323
	2	38.95	49248	41.72	8855
	3	17.05	21563	16.57	3518
	4	14.19	17937	12.72	2701
	5	5.95	7519	5.1	1082
	6	2.47	3126	2.04	433
	7+	1.95	2464	1.48	315
Number of Observations			126448		21227

Figure A.2: Retail sales, total excluding motor vehicles and parts (seasonally adjusted) constructed using Census Bureau estimates and Numerator panel data

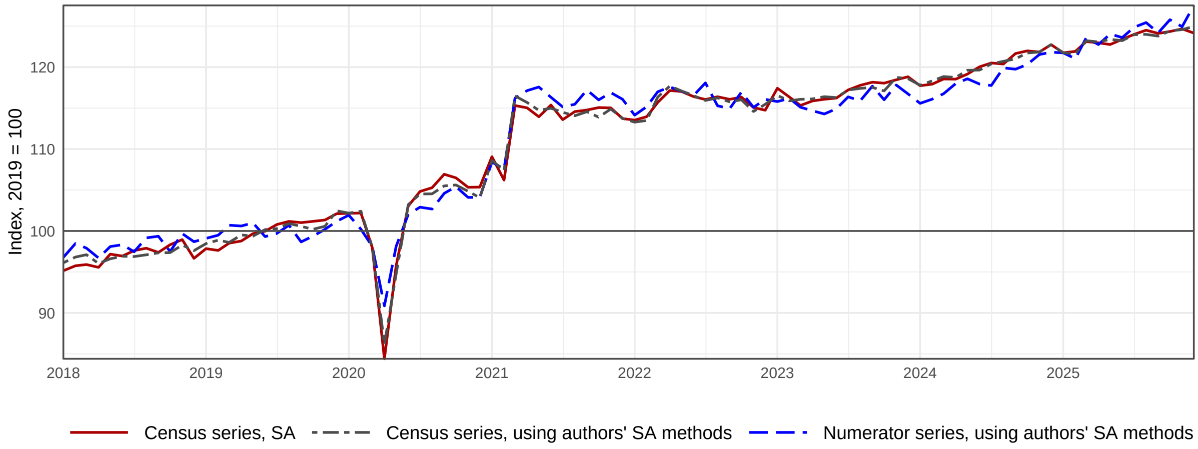

Note: The red solid line is the Census Bureau's seasonally-adjusted retail sales series excluding motor vehicles and parts. The blue dashed line is the retail spending measure we construct from the Numerator data, which we seasonally adjust. The gray dash-dotted line is the non-seasonally adjusted Census series that we then seasonally adjust. The data are monthly from January 2018 through December 2025. Authors' seasonal adjustment method uses X13-ARIMA-SEATS. All three lines are indexed to 100 in 2019.

Source: Census Bureau and Numerator, and authors' own calculations.

Figure A.3 compares the Numerator Fisher Price Index with the Bureau of Economic Analysis's (BEA) PCE deflator for food over 2018–2025, constructed as in Eq. (8). Panel A shows both measures moving in close alignment: stable through 2019, rising sharply from 2021 amid broad inflationary pressures (Shapiro, 2024; Braun et al., 2024), and reaching approximately 130 by mid-2025, a cumulative increase of roughly 30% since the base period.

FIGURE A.3: Numerator Fisher Price Index versus the PCE deflator for food

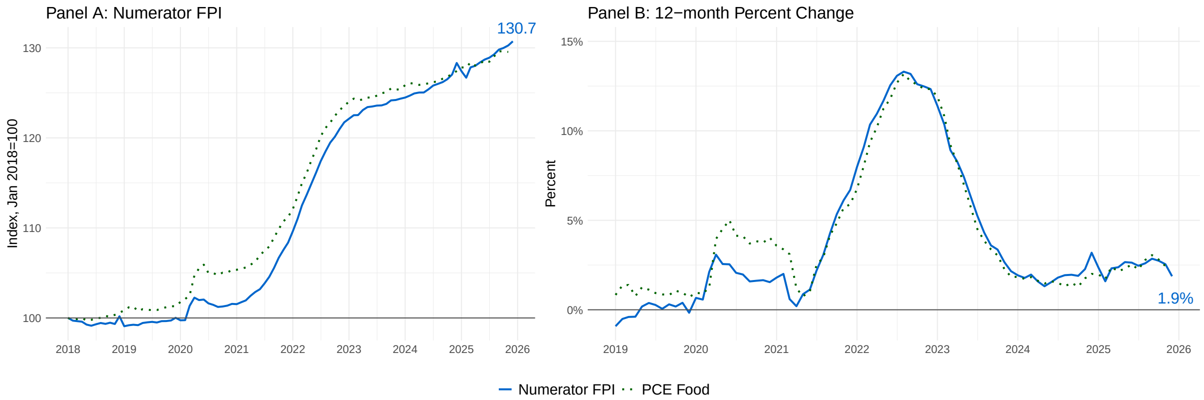

Note: Left panel: Numerator price index for January 2018 through December 2025; BEA PCE price index for food for January 2018 through November 2025. Fisher Price Index constructed using Numerator data in solid blue line, BEA's PCE for food in dotted green line; January 2018=100. The PCE data is seasonally adjusted by the BEA and the Fisher Price Index is seasonally adjusted using X13-ARIMA-SEATS. Right panel: Data for January 2019 through November 2025; the year-over-year percent change in the Fisher Price Index is constructed using Numerator data and shown as the solid blue line, the year-over-year percent change for the PCE Price Index for food is shown as the dotted green line.
Source: Numerator FPI calculated using Numerator Panel Data and based on authors' calculations. PCE Deflator for Food is from the Federal Reserve Bank of St. Louis FRED, series DFXARG3M086SBEA.

Panel B plots year-over-year changes, confirming the strong co-movement. Both series peak near 13% in 2021–2022 before decelerating, and both register a renewed pickup in food price inflation during 2025 relative to 2024. The correlation coefficients reported in Table A.2 indicate that although PCE food inflation is the most correlated with the Numerator FPI, other measures of inflation show high correlation as well.

TABLE A.2: Correlation between inflation measures and Numerator FPI

PCE Food & Bev	PCE Goods ex Auto	Core PCE	CPI	CPI Food
0.96	0.81	0.85	0.85	0.85

Note: Correlation of 12-month percent change of Numerator FPI and the 12-month percent changes of reported inflation measures. Underlying monthly data capture January 2019 to November 2025.
Source: Numerator Panel Data and Federal Reserve Bank of St. Louis FRED, authors' calculations. PCE Deflator for Food is from the Federal Reserve Bank of St. Louis FRED: DFXARG3M086SBEA. Core PCE is from FRED: PCEPILFE. CPI is from FRED: CUSR0000SAC. CPI Food is from FRED: CPIUFDSL. The PCE goods deflator that excludes motor vehicles is based on author's calculations.

Figure A.4 plots the income-specific Fisher Price Indices over the sample period.

Low-income households experience cumulatively higher price growth than middle- and high-income households, with the gap widening after the February 2025 tariff implementation. This pattern is consistent with the regression-based finding in Table 4 that price pass-through is monotonically decreasing in income, and with the broader literature on inflation inequality (Argente and Lee, 2021).

FIGURE A.4: Income-Specific Fisher Price Indices

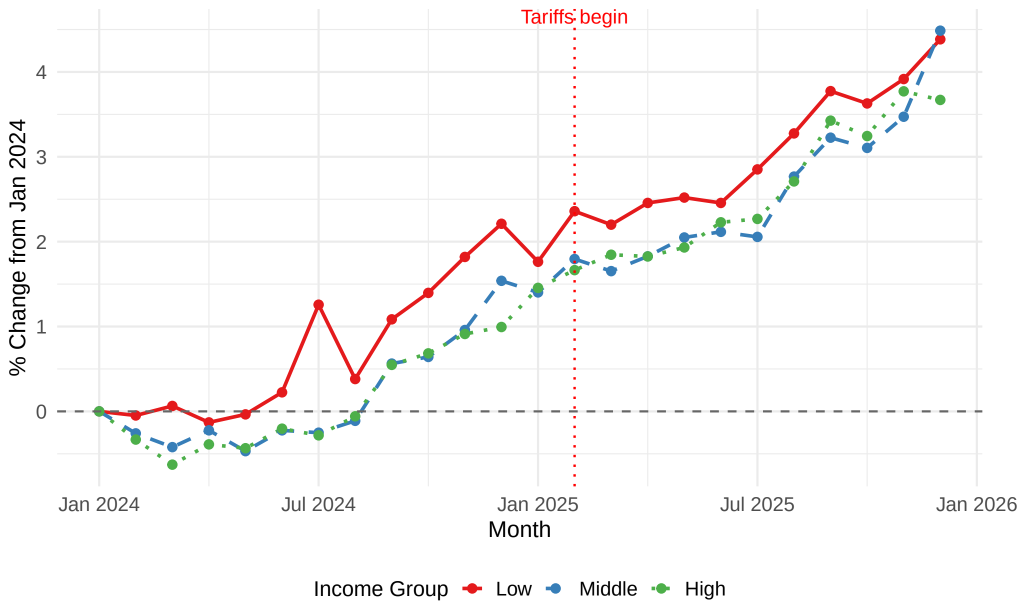

Notes: The figure plots the Fisher Price Index constructed separately for each income group (Low, Middle, High) from January 2024 through December 2025, expressed as cumulative percentage change from January 2024. The vertical dashed line marks February 2025, when tariffs were implemented. Income groups are defined as in Section 3.3.

B Validation of Tariff Exposure Measure

In the main analysis, we use two tariff measures. First, the tariff rate for HS code h is the import-share-weighted average of country-specific tariff rates: $\text{TariffRate}_h = \sum_C \text{ImportShare}_h^C \times \tau_h^C$. Second, the tariff exposure adjusts for import penetration as in Equation 3: $\text{TariffExposure}_h = \sum_C \text{ImportShare}_h^C \times \tau_h^C \times \text{ImportPenetration}_h$. To assess the validity of this approach, for a subset of items in our data, we can use Label Insight, a dataset of consumer packaged goods sold in the US that includes detailed product information such as country of origin linked to Universal Product Codes (UPCs).

The matching process between Numerator and Label Insight produces approximately 27,000 unique products, predominantly grocery items (over 75 percent), with Health & Beauty, Pet, and Household products comprising an additional 20 percent. Despite the concentration across a few sectors, we use country of origin information to construct direct tariff measures. For these products, the true tariff rate is simply the tariff rate applied to the product's actual origin country, $\tau_h^{C_{\text{actual}}}$. This exercise helps assess whether our main results helps to validate that the tariff measures we construct are reasonable proxies for the true origin composition of consumer purchases.

For each HS4 code h, we compute the actual tariff rate and actual tariff exposure as the spending-weighted average across all origin-known imported products within h:

$$\text{ActualRate}_h = \sum_C \omega_h^C \times \tau_h^C, \qquad \text{ActualExposure}_h = \sum_C \omega_h^C \times \tau_h^C \times \text{ImportPen}_h \tag{B.1}$$

where $\omega_h^C = \frac{\sum_{j\in(h,C)} r_j}{\sum_{j\in h} r_j}$ is the spending share of country C within HS code h among origin-known imported products, r_j is total 2024 revenue for product j, and τ_h^C is the average post-period tariff rate applied to country C for HS code h.

Figure B.1 plots actual against proxy measures for the 101 HS4 codes with origin-known imported products. Panel A shows tariff rates, while Panel B shows tariff exposure (rates adjusted for import penetration). For tariff exposure, the correlations are notably stronger than for tariff rates. The mean proxy exposure closely matches the mean actual exposure, and both proxy measures are approximately unbiased. To the extent that measurement error in the proxy attenuates coefficients, the baseline estimates reported in Sections 3 are conservative.

Figure B.1: Validation: Actual vs. Proxy Tariff Measures

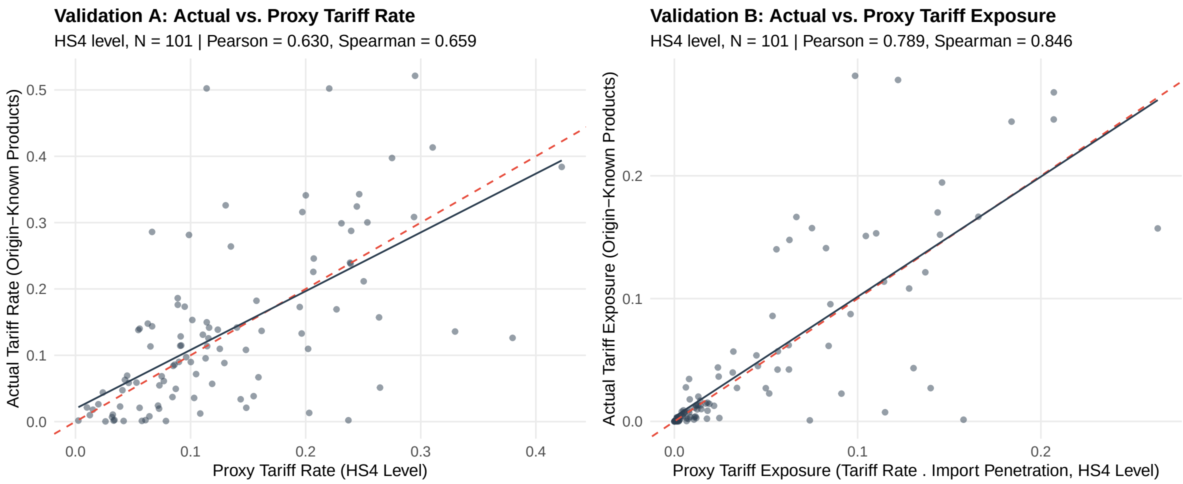

Notes: Each point represents one HS4 code ($N = 101$). Panel A shows tariff rates: the x-axis is the proxy tariff rate, computed as the duties paid over total imports for each HS4 code in 2024 using trade data; the y-axis is the actual tariff rate from products with verified country of origin using spending-weighted origin shares. Panel B shows tariff exposure: the x-axis multiplies the proxy tariff rate by import penetration (Equation 3); the y-axis multiplies the actual rate by import penetration (Equation B.1). The dashed line is the 45-degree line; the solid line is the OLS fit. Panel A: Pearson = 0.63, Spearman = 0.66. Panel B: Pearson = 0.79, Spearman = 0.85.

C Tariff Exposure Measure by Sector, Country, and Essentiality

Figure C.1 shows the different tariffs exposures for each country from January 2024 to December 2025. Figure C.2 shows the different tariff rates and exposures each sector. Figure C.3 presents the tariff rates and exposures by essentiality as defined in Section 3.

Figure C.1: Average Tariff Rate by Country

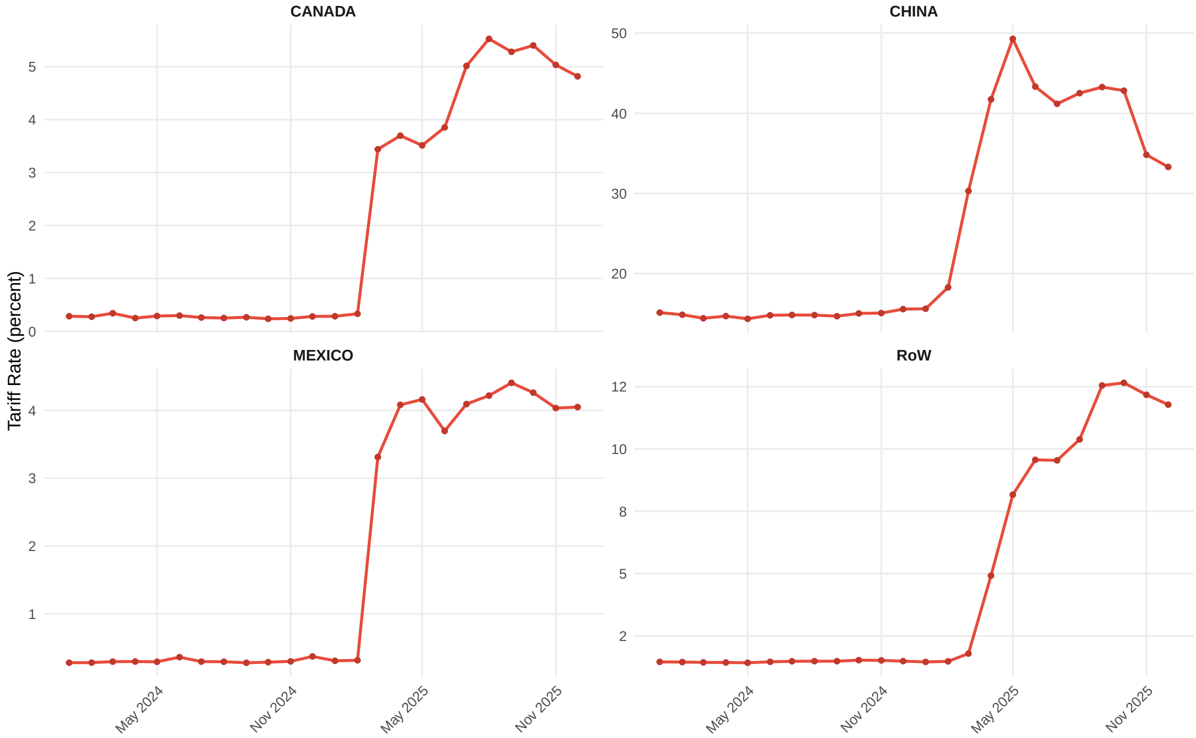

Note: This figure shows the monthly average tariff exposure over the sample period by country. For each year-month, we calculate the mean tariff exposure across all observations. Tariff exposure is measured in percent.

FIGURE C.2: Average Tariff Rate and Exposure by Sector

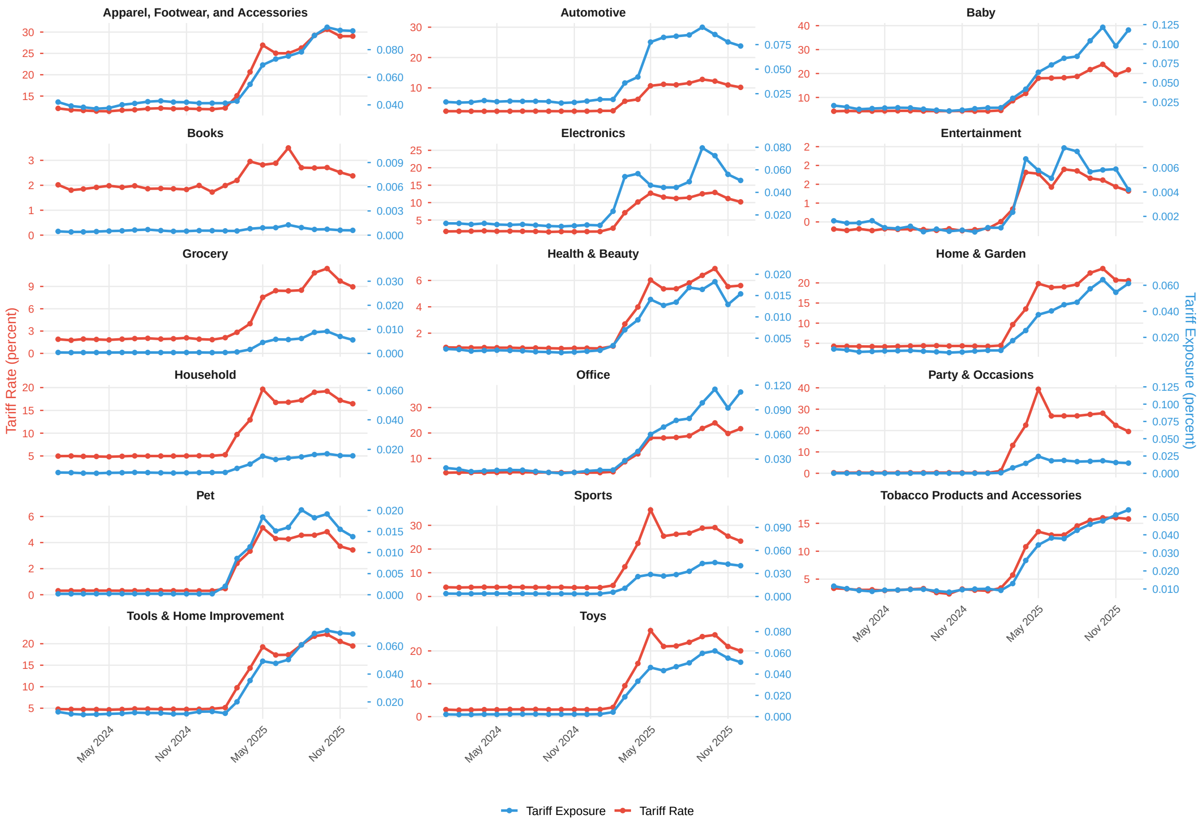

Note: This figure shows the monthly average tariff rate (left axes) and exposure (right axes) over the sample period by sector. For each year-month, we calculate the mean tariff rate and exposure across all observations. Both are measured in percent.

Figure C.3: Average Tariff Rate and Exposure by Essentiality

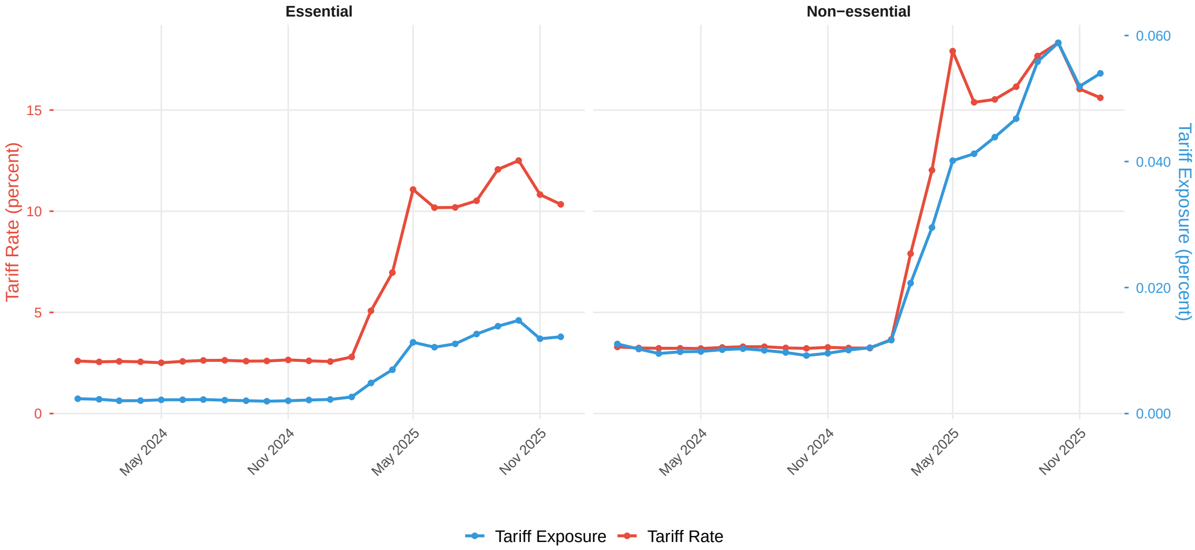

Note: This figure shows the monthly average tariff rate (left axis) and exposure (right axis) over the sample period by essentiality. Essential: Grocery, Household, Health & Beauty. Non-essential: everything else (Automotive, Electronics, Tools & Home Improvement, Home & Garden, Office, Baby, Pet, Tobacco, Apparel/Footwear/Accessories, Sports, Toys, Books, Entertainment, Party & Occasions). For each year-month, we calculate the mean tariff rate and exposure across all observations. Both are measured in percent.

D Stability of the Aggregate Results in Section 3

D.1 Excluding Grocery

Table D.1 replicates the analysis excluding grocery products, which constitute 46% of total spending. Grocery is dominated by domestic production so has low import penetration, as shown in Figure C.2 features intense retail competition on staple goods, and exhibits relatively inelastic demand. Including grocery may therefore attenuate estimated tariff effects by pooling high-exposure non-grocery categories with a large, low-exposure sector.

Results excluding grocery are qualitatively and quantitatively similar to the full sample. Price pass-through increases modestly in both specifications, consistent with grocery exhibiting lower pass-through due to competitive pressure and high domestic production shares. Spending and quantity responses remain large and significant, confirming that

Table D.1: Main Results: Price, Spending, and Quantity (Excluding Grocery)

	Price		Spending	Quantity
	Detrended (1)	Not Detrended (2)	(3)	(4)
Panel A: Tariff Rate				
Tariff Rate × Post	0.14**	0.13**	−0.41***	−0.39***
	(0.06)	(0.05)	(0.10)	(0.10)
Observations	81,352	81,352	81,352	81,352
Adjusted R^2	0.59	0.55	0.97	0.98
Panel B: Tariff Exposure				
Tariff Exposure × Post	0.16**	0.13*	−0.52***	−0.47***
	(0.08)	(0.07)	(0.15)	(0.14)
Observations	81,352	81,352	81,352	81,352
Adjusted R^2	0.58	0.54	0.97	0.98
Year-month FE	Yes	Yes	Yes	Yes
Panel ID × Month FE	Yes	Yes	Yes	Yes
Panel-specific pre-trends	Yes	No	No	No

Notes: This table replicates Table 3 on the sample excluding grocery products (which constitute 46% of total spending). Price results are reported for both detrended (column 1) and non-detrended (column 2) specifications. Spending and quantity use the detrended specification. Detrending includes panel-specific linear time trends in the price equation only. All specifications include year-month and panel-unit-by-calendar-month interaction fixed effects. Standard errors clustered at HS4 code level in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

the aggregate findings are not mechanically driven by grocery's dominant expenditure share.

D.2 Leave-One-Sector-Out Diagnostic

A natural question with the main specification in Equation (4) is that the estimated tariff coefficient may reflect not the causal effect of tariffs on prices and quantities, but rather the differential cyclical sensitivity of tariffed product categories to the 2025 macroeconomic environment. Because high-tariff categories are disproportionately durable and discretionary goods, any aggregate softening in durables demand that coincides with the post-tariff period risks loading onto β even after conditioning on year-month fixed effects, since γ_t absorbs only the common component of aggregate shocks and not their differential incidence across product categories. If the estimated response of quantities and spending to tariffs is in fact driven by a small number of cyclically sensitive sectors, the coefficient cannot be interpreted as a clean tariff pass-through or demand response.

To assess whether the aggregate β is driven by any single sector, we conduct a leave-one-sector-out (LOSO) diagnostic, re-estimating the benchmark specification seventeen times, each time excluding one of the seventeen sectors in the data. Figure D.1 reports the resulting point estimates and 95% confidence intervals for price, spending, and quantity under both the tariff rate specification (Panel A) and the tariff exposure specification (Panel B), with the full-sample benchmark indicated by the dashed reference line in each panel.

Figure D.1: Leave-one-sector-out estimates of β

(A) Tariff Rate

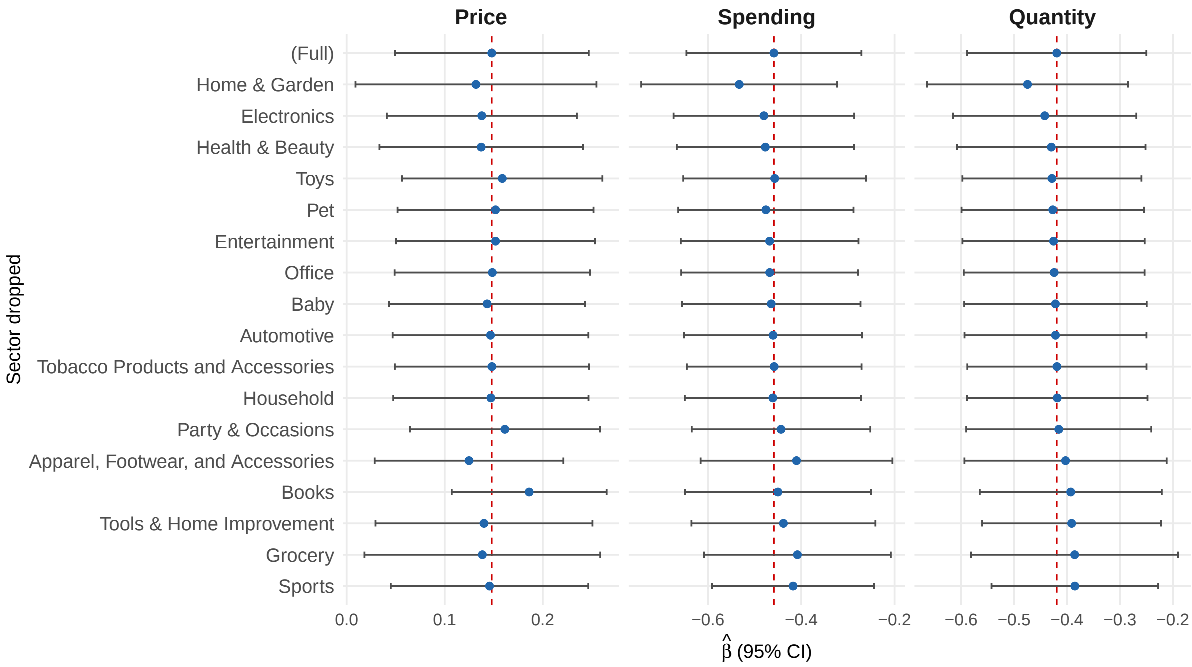

(B) Tariff Exposure

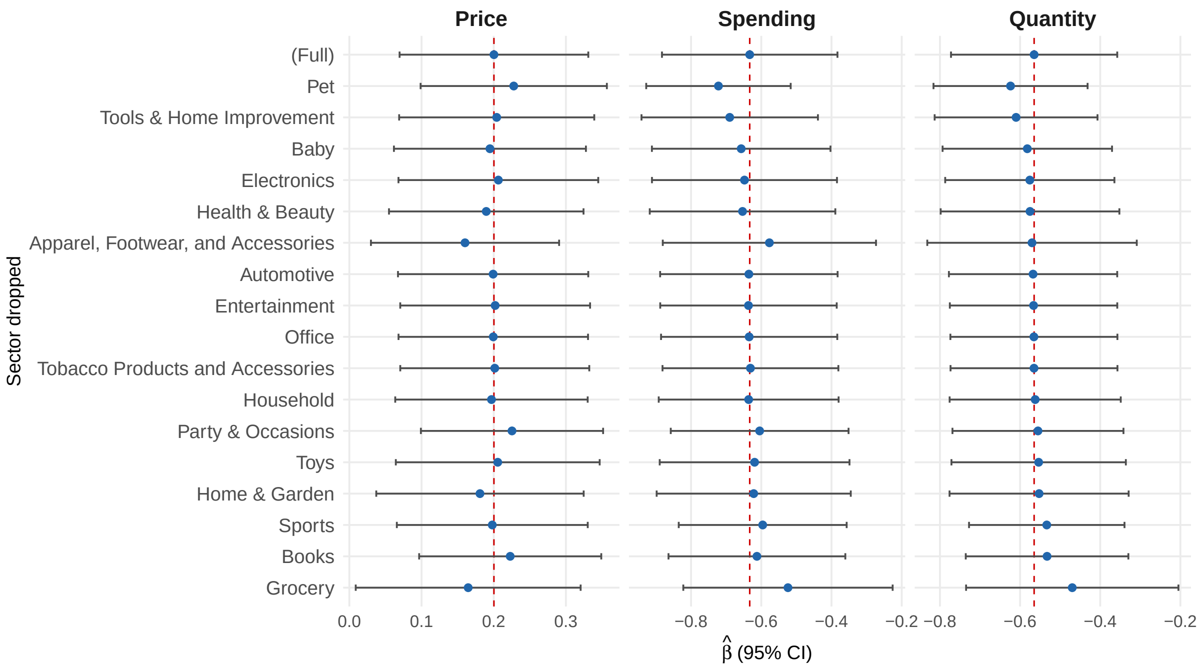

Note: Each row reports the benchmark tariff coefficient β from Equation (4) when the indicated sector is excluded from the estimation sample. The row labeled (Full) reports the full-sample estimate, also indicated by the red dashed vertical line in each panel. Whiskers denote 95% confidence intervals with standard errors clustered at the HS4 code level. Columns correspond to the three outcome variables. Panel A uses the realized tariff rate $\tau_{h(i),t}$ as the treatment variable; Panel B uses the tariff exposure measure defined in Equation (7).

The quantity coefficient moves in a range close to the full-sample estimates, and the spending coefficient exhibits comparable stability. The largest absolute deviation in both cases arises from the exclusion of Home & Garden, which shrinks $|\beta|$ by roughly thirteen percent on quantity and sixteen percent on spending, reflecting the sector's combination of high tariff exposure and substantial sample weight. No other sector alters the aggregate estimate by more than nine percent, and in no case does any single sector's removal overturn the sign, statistical significance, or order of magnitude of the aggregate response.

The pattern across sectors is informative about the cyclicality concern. If the 2025 weakness in durables demand were contaminating the aggregate coefficient, we would expect the removal of the most cyclically sensitive durable sectors to attenuate $|\beta|$ substantially. The opposite obtains. Dropping Electronics slightly increases $|\beta|$ on quantity, consistent with anticipatory stockpiling in consumer electronics dampening the post-period quantity response in that sector and thereby pulling the aggregate coefficient toward zero. Dropping Automotive leaves β essentially unchanged. Dropping Grocery, the largest nondurable staple category, attenuates $|\beta|$ by approximately eight percent, indicating that the untreated staple category plays a limited role as a de facto control group. Taken together, these patterns are inconsistent with a narrative in which macroeconomic weakness in durables is loading onto the tariff coefficient and are consistent with the interpretation of β as a tariff-driven response.

Price pass-through exhibits a similar degree of stability, with estimates ranging narrowly around the full-sample value and no sector's removal altering the qualitative conclusion of incomplete but economically meaningful pass-through.

E Survey Appendix

This appendix contains the survey questionnaire, the Q8-to-sector crosswalk, and supporting tables referenced in Section 5. The appendix is organized to follow the order of the main section: Appendix E.1 presents the full questionnaire and Appendix E.2 the Q8-to-sector crosswalk. Appendix E.3 reports the low-tariff precautionary-null tests referenced in the Section 5 preamble. Appendix E.4 provides the Q7 horserace robustness referenced in Section 5.3. Appendix E.7 reports the sentiment falsification test referenced in Section 5.4. Appendix E.8 presents the within-category trade-down item-level evidence referenced in Section 5.3. Appendix E.6 reports the stockpiling timing test referenced in Section 3.5. Appendix E.5 presents the income-group breakdown of the Q8 worry specification referenced in Section 5.2. Appendix E.9 reports the income-group breakdown of stated intent referenced in Section 5.4.

Notation throughout this appendix matches the paper: i indexes the household, h indexes the HS4 code, t indexes the year-month, $\text{TariffExposure}_{h,t}$ is the tariff exposure measure defined in Equation (3), Post_t is the post-February-2025 indicator, $\alpha_{i,h}$ denotes a household-by-HS4 fixed effect, and $\lambda_{h,t}$ denotes an HS4-by-year-month fixed effect. Standard errors are clustered at the household level throughout.

E.1 Survey Questionnaire

Q1. *Are you aware of any new or proposed tariffs on goods imported into the US?*

- Yes
- No
- Prefer not to answer

Q2. *How would you describe your opinion on tariffs and their potential impact on you and/or others?*

- Open ended

Q3. *How well do you understand what tariffs are and how they impact prices?*

- Very well – I understand how tariffs affect prices and the economy
- Somewhat – I have a general idea of how tariffs work
- Not much – I'm not really sure how tariffs work
- Not at all – I have no idea how tariffs work
- Prefer not to answer

Q4. *Based on what you know, do you generally support or oppose tariffs?*

- Strongly support
- Somewhat support
- Neutral/No opinion
- Somewhat oppose
- Strongly oppose
- Prefer not to answer

Q5. *What concerns, if any, do you have about the impact of tariffs on your personal finances or shopping habits?*

- Higher prices on everyday goods
- General inflation/Rising cost of goods
- Higher prices on non-essential items (e.g., electronics, entertainment, travel)
- Limited availability of certain products
- Impact on the stock market and/or my investments
- Slower economic growth
- Increased cost of healthcare or insurance
- Impact on my job or industry
- Other, please specify
- No concerns
- Prefer not to answer

Q6. *What impact do you think tariffs will have on the U.S. economy over the next year?*

- Very positive impact
- Somewhat positive impact
- No impact
- Somewhat negative impact
- Very negative impact

- Not sure

Q7. *What changes, if any, do you anticipate making to your finances or shopping habits in response to new or proposed tariffs?*

- Cut back spending on non-essentials
- Look for sales or coupons to offset price increases
- Delay non-essential or big-ticket purchases
- Switch to lower-priced retailers or discount stores
- Buy fewer imported goods
- Delay purchases until prices stabilize
- Switch to U.S.-made alternatives
- Buy or stock up on items earlier to avoid tariff-driven price increases
- Reallocate budget away from certain categories to prioritize essentials
- Increase my budget for certain items to account for price increases
- Not sure
- I don't expect to make any changes
- Prefer not to answer

Q8. *Are there specific items or categories you are particularly concerned about being impacted by tariffs?*

- Groceries
- Household goods
- Gasoline
- Medications or medical supplies
- Automobiles
- Electronics
- Personal care products

- Apparel
- Home appliances
- Building materials
- Electricity or natural gas
- Pet supplies
- Furniture
- Toys & games
- Other, please specify
- None of the above
- Prefer not to answer

Q9. *Looking ahead, how do you think the U.S. economy will perform over the next year?*

- It will significantly improve
- It will slightly improve
- It will stay about the same
- It will somewhat worsen
- It will significantly worsen
- Not sure

Q10. *How concerned are you about the possibility of a recession in the coming year?*

- Very concerned
- Moderately concerned
- Slightly concerned
- Not at all concerned
- Not sure

E.2 Q8-to-Sector Crosswalk

The crosswalk maps each Q8 category to the sectors in the household-data hierarchy described in Section 2. For each household-HS4-month observation, the indicator $\text{Worried}_{i,h}$ used in Equation (18) equals one if household i flagged any Q8 category whose mapped sector contains HS4 code h.

TABLE E.1: Q8 Categories to Household-Data Sectors

Q8 category	Mapped sector(s)
Groceries	Grocery
Household goods	Household
Gasoline	Automotive
Medications or medical supplies	Health & Beauty
Automobiles	Automotive
Electronics	Electronics
Personal care products	Health & Beauty
Apparel	Apparel
Home appliances	Appliances
Building materials	Home Improvement
Electricity or natural gas	(no mapped sector)
Pet supplies	Pet
Furniture	Home & Garden
Toys & games	Toys

E.3 Precautionary Channel: Null Tests on the Low-Tariff Basket

This subsection reports the null tests referenced in the Section 5 preamble that discipline the interpretation of the precautionary mechanism. A version of the precautionary story stronger than the one adopted in the paper would predict that pessimistic households cut spending differentially on goods that face no meaningful tariff shock, which would identify a contraction operating through mechanical pass-through on tariff-exposed goods specifically rather than through broad-basket belt-tightening. We test this stronger version directly and find no supporting evidence. We define the low-tariff basket as the set of HS4 codes in the bottom decile of post-period mean tariff rate $\tau_{h,t}$, where the cutoff is $\tau_{h,t}^{p10} \approx 0.045$. True zero-tariff HS4 codes do not exist in the data: the minimum post-period mean tariff rate is 0.002, reflecting the fact that even goods with low effective rates face some non-zero duties. The bottom-decile basket is the closest empirical analogue to a no-tariff basket.

Sentiment-amplified household-specific exposure on the low-tariff basket. This test combines the household-specific exposure measure HHExp_i with the low-tariff basket as the outcome. If pessimistic households with high baseline tariff exposure cut spending differentially on goods that face no meaningful tariff shock, that is a precautionary spillover that cannot be substitution along a stable demand curve and cannot be price incidence. We collapse the household-HS4-month panel restricted to low-tariff HS4 codes to the household-month level, with $Y_{i,t}^{\text{low}}$ denoting real household spending on the low-tariff basket in month t, and estimate

$$\ln Y_{i,t}^{\text{low}} = \alpha_i + \beta_1 \left(\text{HHExp}_i \cdot \text{Post}_t\right) + \beta_2 \left(\text{Sent}_i \cdot \text{Post}_t\right) + \beta_3 \left(\text{Sent}_i \cdot \text{HHExp}_i \cdot \text{Post}_t\right) + \theta\, t + \varepsilon_{i,t}. \tag{E.1}$$

The triple-interaction coefficient β_3 is the precautionary-spillover test of interest: it asks whether the response of low-tariff spending to baseline exposure is amplified for pessimistic households.

TABLE E.2: Precautionary Spillover Test: Sentiment × HHExp on Low-Tariff Basket

	Q6	Q9	Q10	PC1
$\text{HHExp}_i \cdot \text{Post}_t$	−0.786***	−0.760***	−0.742***	−0.777***
	(0.095)	(0.101)	(0.093)	(0.103)
$\text{Sent}_i \cdot \text{Post}_t$	0.013	0.012	0.007	0.015
	(0.012)	(0.013)	(0.012)	(0.013)
$\text{Sent}_i \cdot \text{HHExp}_i \cdot \text{Post}_t$	−0.314	−0.276	−0.156	−0.359
	(0.264)	(0.278)	(0.258)	(0.285)
t (linear trend)	0.000***	0.000***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)
Num. Obs.	454,130	399,330	467,662	379,783
R^2	0.524	0.523	0.523	0.524

Notes: The table reports estimates of Equation (E.1) on the household-month panel of survey respondents. The dependent variable is the natural log of real household spending on the low-tariff basket (HS4 codes in the bottom decile of post-period mean tariff rate). HHExp_i is the household-specific basket-weighted exposure measure. Sent_i is the household-level sentiment measure indicated in the column header; Q6, Q9, and Q10 are z-scored, and PC1 is the first principal component across Q6, Q9, Q10, and Q5. All specifications include household fixed effects α_i and a linear time trend; year-month fixed effects are omitted because they would absorb the time-invariant household-level treatment. Standard errors clustered at the household level in parentheses. *, **, *** denote significance at the 10, 5, and 1 percent levels.

Table E.2 reports the pooled estimates. The triple-interaction coefficient β_3 is negative-signed in every column but never statistically significant, with point estimates ranging from −0.156 (Q10) to −0.359 (PC1) and standard errors in the 0.26–0.29 range that comfortably cover zero. The headline precautionary-pass-through spillover is null.

The baseline-exposure coefficient β_1 is significantly negative and around -0.7 across all four sentiment specifications. High-baseline-exposure households cut spending on the low-tariff basket by about 0.8 percent per unit of HHExp_i. This means that the contraction in tariff-exposed categories spills over onto unexposed categories, identified independently of stated sentiment. The mechanism appears to be a basket-level income effect rather than a sentiment-driven uncertainty response.

TABLE E.3: Precautionary Spillover Test by Income Group: PC1 × HHExp on Low-Tariff Basket

	Low	Middle	High
$\text{HHExp}_i \cdot \text{Post}_t$	0.102	−0.890***	−1.264***
	(0.272)	(0.136)	(0.195)
$\text{Sent}_i \cdot \text{Post}_t$	0.045	−0.000	0.019
	(0.030)	(0.017)	(0.026)
$\text{Sent}_i \cdot \text{HHExp}_i \cdot \text{Post}_t$	−1.081	−0.017	−0.472
	(0.707)	(0.381)	(0.562)
t (linear trend)	−0.000	0.000***	0.000***
	(0.000)	(0.000)	(0.000)
Num. Obs.	80,546	207,619	87,605
R^2	0.518	0.504	0.511

Notes: The table reports estimates of Equation (E.1) estimated separately within each income group, using PC1 as the sentiment measure. Specifications, dependent variable, exposure measure, and fixed effects as in Table E.2. PC1 is the first principal component across Q6, Q9, Q10, and Q5. Standard errors clustered at the household level in parentheses. *, **, *** denote significance at the 10, 5, and 1 percent levels.

The income decomposition in Table E.3 shows that the baseline-exposure spillover (β_1) is concentrated in middle- and high-income households (Middle: −0.890, $p < 0.01$; High: −1.264, $p < 0.01$) and is null among low-income households. The sentiment-amplified spillover (β_3) remains null in every income group. Middle- and high-income households have the discretionary slack to contract spending on unexposed goods when their tariff-exposed basket is hit, while low-income households absorb the shock through forced level cuts on the tariff-exposed basket itself, with little room for additional contraction on the unexposed margin. What pessimism does not do, in any income group, is generate an additional contraction on unexposed goods beyond what baseline exposure already produces.

Sentiment-only test on the low-tariff basket. A weaker version of the same test asks whether pessimism by itself, without conditioning on baseline exposure, predicts differential cuts on the low-tariff basket relative to tariff-exposed goods. We estimate the within-HS4 specification restricted to low-tariff HS4 codes,

$$\ln Y_{i,h,t} = \beta\,(\text{Sent}_i \cdot \text{Post}_t) + \alpha_{i,h} + \lambda_{h,t} + \varepsilon_{i,h,t}, \tag{E.2}$$

and the pooled specification with a low-tariff contrast as a triple interaction.

TABLE E.4: Sentiment × Post on Low-Tariff Basket: Within-HS4 Specification

	Low-tariff only (PC1)	Pooled with contrast (PC1)
$\text{Sent}_i \cdot \text{Post}_t$ (PC1)	−0.002**	−0.002**
	(0.001)	(0.001)
$\text{Sent}_i \cdot \text{Post}_t \times \text{LowTariff}_h$ (PC1)		−0.001
		(0.001)
Num. Obs.	5,619,562	35,657,877
R^2	0.361	0.395

Notes: Column 1 reports estimates of Equation (E.2) on the household-HS4-month panel restricted to low-tariff HS4 codes (bottom decile of post-period mean tariff rate). Column 2 reports the pooled specification with $\text{LowTariff}_h \cdot \text{Post}_t \cdot \text{Sent}_i$ as a triple interaction estimated on the full panel; the main effect of $\text{LowTariff}_h \cdot \text{Post}_t$ is absorbed by $\lambda_{h,t}$ since LowTariff_h is constant within HS4. PC1 is the first principal component across Q6, Q9, Q10, and Q5. All specifications include household-by-HS4 ($\alpha_{i,h}$) and HS4-by-year-month ($\lambda_{h,t}$) fixed effects. Standard errors clustered at the household level in parentheses. *, **, *** denote significance at the 10, 5, and 1 percent levels.

Table E.4 reports the within-HS4 estimates. Pessimism has a small but significant negative effect on low-tariff spending in the restricted sample (−0.002, $p < 0.05$). The pooled specification with the low-tariff contrast in column 2 is the identifying test: the triple interaction is −0.001 and not statistically significant, indicating that pessimism's effect on the low-tariff basket is statistically indistinguishable from its effect on tariff-exposed goods. A precautionary contraction operating through mechanical pass-through on tariff-exposed goods would require pessimism to differentially reduce spending on those goods relative to unexposed goods. The null contrast rules out that mechanism and supports the interpretation adopted in Section 5.1: pessimistic households engage in broad-basket belt-tightening, with the compositional shift toward essentials identified at the household-month level.

In summary, Tables E.2–E.4 establish what the precautionary mechanism in this paper is not. Sentiment does not differentially predict cuts on tariff-exposed goods relative to unexposed goods. The version of the precautionary story that would require households to act as if they know the HS4-level tariff status of their basket fails. The version that survives, documented in Section 5.1, requires only that households recognize tariffs as a source of aggregate uncertainty and respond by reweighting their basket toward goods they regard as essentials. That mechanism operates on basket composition rather than on a differential cut on tariff-exposed goods, and it is supported by the share-essential shift, the asymmetric level response (essentials defended, non-essentials cut), and the income gradient (concentrated in middle-income households with discretionary slack). The basket-level income effect identified by β_1 in Table E.2 is a separate channel: high-baseline-exposure households contract spending on unexposed goods, but the contraction is driven by exposure itself, not by sentiment.

E.4 Q7 Reduction Intent: Horserace Robustness

This subsection reports the horserace specification referenced in Section 5.3, in which all five Q7 reduction sub-items are entered simultaneously as household-level indicators interacted with $\text{TariffExposure}_{h,t} \cdot \text{Post}_t$. Because the five sub-items are positively correlated across households (households who select one reduction-related item tend to select others), individual standard errors inflate relative to the single-item specifications in Table 11. We therefore report the individual sub-item regressions in the main text and the joint specification here as a robustness check.

TABLE E.5: Q7 Decomposition: Horserace Specification

	$\ln Y_{i,h,t}$
$\text{TariffExposure}_{h,t} \cdot \text{Post}_t \times \text{CutBack}_i$	−0.026
	(0.019)
$\text{TariffExposure}_{h,t} \cdot \text{Post}_t \times \text{Delay}_i$	−0.008
	(0.021)
$\text{TariffExposure}_{h,t} \cdot \text{Post}_t \times \text{DelayBigTkt}_i$	0.010
	(0.021)
$\text{TariffExposure}_{h,t} \cdot \text{Post}_t \times \text{FewerImports}_i$	0.019
	(0.019)
$\text{TariffExposure}_{h,t} \cdot \text{Post}_t \times \text{Reallocate}_i$	−0.040*
	(0.023)
Num. Obs.	47,381,907
R^2	0.395

Notes: The table reports estimates of Equation (19) with all five Q7 reduction sub-items (CutBack_i, Delay_i, DelayBigTkt_i, FewerImports_i, Reallocate_i) entered jointly as household-level indicators interacted with $\text{TariffExposure}_{h,t} \cdot \text{Post}_t$ in a single regression on the household-HS4-month panel. The dependent variable is $\ln Y_{i,h,t}$, the natural log of real spending on HS4 code h in month t. Sub-item definitions match Table 11. Specifications include household-by-HS4 ($\alpha_{i,h}$) and HS4-by-year-month ($\lambda_{h,t}$) fixed effects. Of the five sub-items, only Reallocate_i retains marginal significance when the five compete for variance, identifying stated reallocation toward essentials as the primary channel underlying the pooled reduction effect documented in Table 10. Standard errors clustered at the household level in parentheses. *, **, *** denote significance at the 10, 5, and 1 percent levels.

E.5 Q8 Worry by Income Group

This subsection reports the income-group breakdown of the pooled Q8 specification referenced in Section 5.2. The specification identifies the spending response of worried households relative to unworried households purchasing in the same HS4-month, pooling across essential and non-essential flagged categories.

Table E.6: Category-Specific Worry (Q8) by Income Group

	Low	Middle	High
$\text{Worried}_{i,h} \cdot \text{Post}_t$	-0.006^{**}	-0.006^{***}	-0.006^{***}
	(0.003)	(0.001)	(0.002)
Num. Obs.	9,969,954	25,513,184	11,384,672
R^2	0.386	0.392	0.395

Notes: The table reports estimates of the pooled specification of Equation (18) (imposing $\beta_2 = 0$) on the natural log of real (PCE-deflated) household spending on HS4 code h in month t, estimated separately within each income group. The household-HS4 indicator $\text{Worried}_{i,h}$ equals one if household i flagged any Q8 category whose mapped sector contains HS4 code h (mapping in Appendix E.2). All specifications include household-by-HS4 ($\alpha_{i,h}$) and HS4-by-year-month ($\lambda_{h,t}$) fixed effects. Standard errors clustered at the household level in parentheses. *, **, *** denote significance at the 10, 5, and 1 percent levels.

The pooled within-HS4 worry response is essentially uniform across the income distribution at −0.6 percent. The income uniformity of the pooled coefficient, combined with the income concentration of the level reallocation in Section 5.1, is consistent with the trade-down interpretation in Section 5.2: all income groups trade down within categories they worry about, which produces the within-HS4 spending cut without a quantity cut documented in the pooled column 2 of Table 9. Only middle-income households additionally execute the level cut on non-essentials that drives the basket-composition reallocation toward essentials.

E.6 Stockpiling Timing Test: Early versus Late Post-Period

This subsection reports the early-versus-late timing test referenced in Section 3.5. We partition the post-tariff period into an early window (February–April 2025), during which any pull-forward would plausibly materialize, and a late window (May–December 2025), and re-estimate Equation (19) with the tariff-exposure-by-intent interaction computed separately for each sub-period:

$$Y_{i,h,t} = \beta^{\text{early}} \left(\text{TariffExposure}_{h,t} \cdot \text{Post}_t^{\text{early}} \cdot \text{Intent}_i\right) + \beta^{\text{late}} \left(\text{TariffExposure}_{h,t} \cdot \text{Post}_t^{\text{late}} \cdot \text{Intent}_i\right) + \alpha_{i,h} + \lambda_{h,t} + \varepsilon_{i,h,t}, \tag{E.3}$$

where $Y_{i,h,t} \in \{\ln Q_{i,h,t}, \ln Y_{i,h,t}\}$ and $\text{Post}_t^{\text{early}}$ and $\text{Post}_t^{\text{late}}$ are indicators for the early and late sub-periods. Classic stockpiling predicts a front-loaded positive quantity response among stated stockpilers ($\beta^{\text{early}} > 0$, large) that fades or reverses in the late window.

TABLE E.7: Stated Intent × Tariff Exposure: Early versus Late Post-Period

	Stockup		Reduce		Precaution	
	Quantity	Spend	Quantity	Spend	Quantity	Spend
Intent × TariffExp × Early	0.034	−0.087*	−0.016	−0.059	−0.011	−0.061
	(0.025)	(0.052)	(0.021)	(0.045)	(0.020)	(0.043)
Intent × TariffExp × Late	0.018*	0.008	−0.002	−0.044**	0.004	−0.043**
	(0.010)	(0.021)	(0.008)	(0.018)	(0.008)	(0.017)
Household × HS4 FE	Yes	Yes	Yes	Yes	Yes	Yes
HS4 × year-month FE	Yes	Yes	Yes	Yes	Yes	Yes
Num. Obs.	47,381,907					
R^2 (qty / spend)	0.241 / 0.395					

Notes: Each column reports estimates of Equation (E.3) on the household-HS4-month panel of survey respondents, with the post-period split into an early window (February–April 2025) and a late window (May–December 2025). The dependent variable is $\ln Q_{i,h,t}$ (quantity columns) or $\ln Y_{i,h,t}$ (spend columns). Stockup_i is stated stockpile intent from Q7; Reduce_i is the union of the five Q7 reduction sub-items; Precaution_i is the four-item precautionary composite (cut back, delay, delay big-ticket, reallocate). All specifications include household-by-HS4 ($\alpha_{i,h}$) and HS4-by-year-month ($\lambda_{h,t}$) fixed effects. Standard errors clustered at the household level in parentheses. *, **, *** denote significance at the 10, 5, and 1 percent levels.

Table E.7 reports the results. Two patterns emerge. First, the stockpile-intent quantity response is not front-loaded: the early-period coefficient (+0.034) is positive but not significant, and the late-period coefficient (+0.018, $p < 0.10$) is of similar magnitude and marginally significant. Classic stockpiling would predict the reverse, with the quantity pull-forward concentrated in the early window and dissipating or reversing as inventories are drawn down. The absence of front-loading is inconsistent with anticipatory buying as the primary channel. Meanwhile, the stockpile-intent early-period spending coefficient is marginally negative (-0.087, $p < 0.10$), which is the opposite of what buying-ahead would produce and is more consistent with stated stockpilers being anxious, tariff-aware households whose early revealed behavior resembles the reduction group. Second, the reduce and precaution spending cuts are concentrated in the late post-period (-0.044 and -0.043, both $p < 0.05$), with the early-period coefficients negative but imprecisely estimated. The quantity responses for both reduce and precaution intents are null in both sub-periods, confirming that the trade-down interpretation holds in both the transition and steady-state windows. The persistence of the spending cut in the late period is consistent with sustained precautionary reallocation rather than a one-time response to the tariff announcement.

E.7 Sentiment Falsification

Table E.8 tests whether the Q6 sentiment amplification result is driven by sentiment proxying for realized price exposure. Column 1 reproduces the baseline specification using Q6. Column 2 adds a control for the household's basket-weighted realized log price change between the pre- and post-periods, constructed from within-HS4 unit price movements. The Q6 sentiment coefficient is essentially unchanged (-0.078 to -0.083), indicating that tariff-anchored pessimism operates independently of realized price shocks.

E.8 Within-Category Trade-Down: Item-Level Evidence

This subsection reports direct item-level evidence that households substitute toward cheaper varieties within tariff-exposed product categories, using the full household panel rather than the survey subsample. This evidence corroborates the trade-down interpretation of the spending-quantity wedge documented in Section 5: spending falls without a commensurate quantity decline, implying that unit prices paid decline through within-category item switching.

Table E.8: Sentiment Falsification: Controlling for Realized Price Changes

	Baseline (1)	With Realized Δp (2)
$\text{HHExp}_i \cdot \text{Post}_t$	-2.480^{***}	-1.977^{***}
	(0.069)	(0.111)
$\text{Sent}_i \cdot \text{HHExp}_i \cdot \text{Post}_t$ (Q6)	-0.078	-0.083^{*}
	(0.050)	(0.050)
Realized $\Delta \ln p_i \cdot \text{Post}_t$		-1.994^{***}
		(0.261)
t (linear trend)	0.0003^{***}	0.0003^{***}
	(0.000)	(0.000)
Household FE	Yes	Yes
Observations	457,863	457,863
R^2	0.639	0.639

Notes: The table reports estimates on the household-month panel of survey respondents. The dependent variable is the natural log of real household spending. Column 1 reproduces the baseline Q6 specification. Column 2 adds a control for the household's basket-weighted realized log price change, constructed as $\Delta \ln p_i = \sum_h s_{i,h}^{2024} \cdot (\ln \bar{p}_h^{\text{post}} - \ln \bar{p}_h^{\text{pre}})$ where $\bar{p}_h^{\text{post}}$ and $\bar{p}_h^{\text{pre}}$ are mean unit prices in HS4 code h in the post- and pre-periods. Standard errors clustered at the household level in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

For each item j within HS4 code h, we define its within-HS4 price percentile as $\pi_j = \text{rank}(p_j^{\text{pre}} \mid j \in h)/N_h$, where p_j^{pre} is item j's mean pre-tariff price and N_h is the number of distinct items in HS code h. We restrict the sample to HS4 codes with at least five items. We first assign items to tercile-based price tiers (cheap, middle, expensive), compute each tier's revenue share $s_{h,k,t}^{\text{tier}}$ within HS code h in month t, and estimate

$$\ln s_{h,k,t}^{\text{tier}} = \delta_{h \times k, m(t)} + \gamma_t + \sum_k \beta_k \left(\text{TariffExposure}_{h,t} \times \text{Post}_t \times \mathbf{1}[\text{Tier} = k] \right) + \varepsilon_{h,k,t}, \qquad \text{(E.4)}$$

with HS4 × tier × calendar-month fixed effects $\delta_{h \times k, m(t)}$ and year-month fixed effects γ_t.

Table E.9: Within-HS Substitution: Revenue Share Response by Price Tier

	$\ln s^{\text{tier}}_{h,k,t}$
Tariff Exp. × Post × Cheap	0.68***
	(0.09)
Tariff Exp. × Post × Middle	0.24***
	(0.06)
Tariff Exp. × Post × Expensive	−0.70***
	(0.07)
HS × Tier × month FE	Yes
Year-month FE	Yes
Observations	12,362
R^2	0.960

Notes: The dependent variable is the log revenue share of each price tier within HS code × month. The specification follows Equation (E.4). Items are classified into tercile-based price tiers using pre-tariff (pre-February 2025) average prices. Standard errors clustered at the HS4 code level in parentheses. *, **, *** denote significance at the 10, 5, and 1 percent levels.

Table E.9 shows that cheap items gained significant revenue share within tariff-exposed HS4 codes while expensive items lost share by a nearly symmetric magnitude, consistent with a zero-sum reallocation of revenue from expensive to cheap varieties.

To avoid arbitrary tercile cutoffs, we also estimate the substitution pattern continuously at the item level. Let $s_{j,t}$ denote item j's revenue share within its HS4 code in month t. We estimate

$$\ln s_{j,t} = \delta_{j,m(t)} + \gamma_t + \beta_1 \left(\text{TariffExposure}_{h(j),t} \times \text{Post}_t\right) + \beta_2 \left(\text{TariffExposure}_{h(j),t} \times \text{Post}_t \times \pi_j\right) + \varepsilon_{j,t}, \tag{E.5}$$

where $\delta_{j,m(t)}$ are item-by-calendar-month fixed effects, γ_t are year-month fixed effects, and $h(j)$ is the HS4 code containing item j. A negative $\hat{\beta}_2$ implies that cheaper items gained market share at the expense of more expensive items within tariff-exposed categories.

Table E.10: Within-HS Item Substitution: Continuous Specification

	$\ln s_{j,t}$
Tariff Exp. × Post	1.41***
	(0.39)
Tariff Exp. × Post × Price Percentile	−0.61***
	(0.13)
Item × month FE	Yes
Year-month FE	Yes
Observations	63,573,069
R^2	0.860

Notes: The dependent variable is the log revenue share of item j within its HS code in month t. The specification follows Equation (E.5). Price percentile π_j ranges from 0 (cheapest) to 1 (most expensive) within HS4 code. All specifications include item-by-calendar-month fixed effects $\delta_{j,m(t)}$ and year-month fixed effects γ_t. Standard errors clustered at the HS4 code level in parentheses. *, **, *** denote significance at the 10, 5, and 1 percent levels.

Table E.10 confirms the pattern across over over 60 million item-month observations. The interaction coefficient shows that for each step up in the within-category price distribution, the tariff-induced gain in market share falls, identifying a reallocation from expensive to cheap items within tariff-exposed categories.[19]

E.9 Stated Intent by Income Group

Table E.11 estimates Equation (19) separately for low-, middle-, and high-income households on three intent measures: stated reduction (any), the precautionary composite, and stated reallocation. Panel A reports spending responses; Panel B reports quantity responses.

[19] The positive main effect implies share gains across the entire price distribution, which cannot hold in a balanced panel and likely reflects differential item attrition: within tariff-exposed categories, disproportionately expensive items exit the sample in the post-period, mechanically raising the shares of surviving items. The interaction term is not affected by this compositional shift.

Table E.11: Stated Intent and Tariff Exposure, by Income Group

	Reduce (any)			Precautionary composite			Reallocate		
	Low	Middle	High	Low	Middle	High	Low	Middle	High
Panel A: Log spending									
$\text{TariffExposure}_{h,t} \times \text{Intent}_i$	−0.038	−0.065***	−0.007	−0.045	−0.053**	−0.025	−0.015	−0.061**	−0.043
	(0.043)	(0.025)	(0.034)	(0.042)	(0.023)	(0.032)	(0.055)	(0.028)	(0.041)
Panel B: Log quantity									
$\text{TariffExposure}_{h,t} \times \text{Intent}_i$	−0.003	−0.007	0.002	0.008	0.010	−0.015			
	(0.019)	(0.012)	(0.016)	(0.019)	(0.011)	(0.016)			
Household × HS4 FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
HS4 × Year-month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations (millions)	10.0	25.5	11.4	10.0	25.5	11.4	10.0	25.5	11.4

Notes: The table estimates Equation (19) separately by income group (low: \$0–40k; middle: \$40–125k; high: \$125k+). Each cell reports the coefficient on $\text{TariffExposure}_{h,t} \times \text{Intent}_i$, where Intent is stated reduction (any of the five reduce-or-delay items in Q7), the precautionary composite (cutback on non-essentials, reallocate budget, or buy fewer imports), or stated reallocation alone. Panel A uses log real spending as the dependent variable; Panel B uses log quantity. Reallocate quantity columns are omitted as they are subsumed by the precautionary composite quantity results. All specifications include household-by-HS4 ($\alpha_{i,h}$) and HS4-by-year-month ($\lambda_{h,t}$) fixed effects. Standard errors clustered at the household level in parentheses. *, **, *** denote significance at the 10, 5, and 1 percent levels.