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

First Draft: May 21, 2026
Revised: August 10, 2026

Abstract

We estimate the effects of the 2025 U.S. tariffs on households using transaction-level data linked to tariff exposure and a tariff sentiment survey. Using the realized tariff rate, we find retail price pass-through of about 0.15, while our benchmark tariff-exposure measure blends the tariff rate with import penetration and yields a larger price coefficient of 0.20. Low-income households bear a disproportionate welfare burden through regressive pass-through. Aggregate spending on affected goods falls about three times as much as prices rise. Quantity declines drive the contraction in spending and are concentrated entirely in non-essential categories, where households have the flexibility to cut back. The real income lost to tariffs explains at most a third of this contraction, leaving the greater part behavioral. We link stated intentions in the survey to households' spending, and identify the mechanism as a precautionary response driven by pessimism about the economic outlook. More pessimistic households reallocate their baskets toward essentials and the cut in non-essential spending behind this shift is concentrated among middle-income households. Within the essential categories they expect tariffs to affect, worried households at every income level trade down to cheaper varieties.

Keywords: Tariffs, Pass-through, Prices, Household Spending, Welfare
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 Rajashri Chakrabarti, Aaron Flaaen, Sharat Ganapati (discussant), Colin Hottman, Chris Kurz, Robbie Minton, Ekaterina Peneva, Erick Sager, Daniel Villar Vallenas, as well as the participants of the NBER SI 2026 Household Finance, SEM 2026 conference, SNDE 2026 conference, BSE Summer Forum Advances in Economic Research Using Bank Transactions Data 2026 workshop, Brookings Institution seminar, 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; Amiti 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 change their spending, and if so why, 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 broad product attributes for household purchases from brick-and-mortar and online retailers and also contain a rich set of household characteristics including income group. The spending 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. For a subset of over 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 behavioral responses to tariffs.

We link each purchase in the data to a measure of tariff exposure constructed at the 4-digit Harmonized System level, 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. A central advantage of the data we use is the joint observation of prices and quantities at the household-product level, which lets us estimate all three responses from a single source. We report all of our headline results for two tariff measures. The first uses the realized tariff rate directly, which is based on paid duties and weights all product categories equally regardless of their import dependence. The second, which is our benchmark specification, scales tariff rates by import penetration to better account for categories where tariffs are actually borne by domestic consumers.

We document three main findings. First, tariffs translate into substantially larger movements in spending than in prices. A one percentage point increase in the tariff rate raises retail prices by approximately 0.15 percentage points and this result has the traditional pass-through interpretation of around 15 percent. Our benchmark specification with the tariff exposure measure yields a larger price coefficient of 0.20. Our findings on the effects of tariffs on prices are consistent with the partial pass-through documented for the 2025 tariff episode. On the other hand, a one percentage point increase in the tariff measures reduces real household spending on affected categories by approximately 0.5 to 0.6 percent. Evaluated at each household's own basket exposure to the tariffs, these estimates imply that prices faced by the average household rise by around 0.9 to 1.7 percent while real spending falls by roughly 2.9 to 5.2 percent, a contraction about three times as large as the implied price increase.

The wedge between prices and spending is difficult to explain by a purely mechanical income effect, whereby tariffs erode real income and households with less real income spend less. A simple accounting identity built from our price and spending estimates places a ceiling on how much of the contraction such an income effect can explain, and that ceiling of a third leaves the greater part of the spending response unaccounted for. The larger part of the spending decline of around two-thirds is behavioral rather than a consequence of reduced purchasing power. 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). We do not find convincing evidence of stockpiling and its subsequent unwinding as an explanation even among households who state an intention to stockpile in the survey.

The second finding concerns the distribution of the burden across income groups. Low-income households bear a disproportionate price burden as the price response to tariff exposure monotonically decreases in income. When we combine higher pass through and lower income into welfare terms, low-income households bear losses of 0.72 percent of income against 0.07 percent for high-income households, roughly nine to ten times the burden as a share of income. We compute this cost using a standard consumer surplus formula and an exact cost-of-living index that captures substitution through observed spending reallocation rather than an assumed demand elasticity, and the two methods agree closely.

While the aggregate results document that spending falls by much more than prices rise, the third finding uses the survey to identify who drives it and how. The behavioral residual left by the accounting exercise turns out to reflect household pessimism about the economic outlook. Even when we account for a household's own tariff-induced income loss, the more pessimistic households still reallocate their baskets more sharply toward essentials, and this response is tied to concern about the future course of the economy rather than to the prices they have already paid. The strongest supporting evidence comes from the goods tariffs do not affect much. Pessimistic households cut spending on goods in the bottom decile of the tariff rate as sharply as they cut tariffed goods, even though prices of these items barely moved. In other words, it is a broad contraction that operates on the whole basket and reaches beyond the tariffed categories. The tariff-specific part of the behavioral response comes from our two other survey measures. A belief that a specific category will be affected and a stated plan of action both track exposure, since households buying heavily in exposed categories are the ones likely to notice.

All income groups contract spending on tariff-exposed goods by similar magnitudes, but the margin of adjustment differs. Within essential categories that households flag as likely to be affected by tariffs, spending falls while quantities hold so households trade down to cheaper varieties, and trade-down is common across the income distribution,

whereas in flagged non-essential categories they continue purchasing, spending rises as households have to absorb higher prices. The margins that do vary are across categories. The additional cut in the level of non-essential spending that drives the basket-composition shift toward essentials is concentrated among middle-income households. Low-income households absorb tariffs primarily through higher prices on continuing purchases. They do reallocate away from tariff-exposed categories, but that adjustment tracks the price increase they face rather than how pessimistic they are. High-income households contract by as much as the others, but their spending cut does not vary with pessimism. The price burden and the precautionary reallocation therefore land on different households. The price burden is regressive, while the reallocation is concentrated where households have the discretionary room to reweight their baskets. We use the term 'precautionary' throughout to refer to this forward-looking, defensive reallocation of the basket driven by pessimism, 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., 2021). 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; Amiti et al., 2026) and our tariff-rate pass-through finding sits within their range. The modest percentage pass-through we find partly reflects data composition toward everyday consumer goods, absorption of the tariff by foreign exporters rather than domestic buyers (Freund, 2026), and two conceptual points in the literature. First, Sangani (2026b) 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 and Bai et al. (2025) showing that price pass through of tariffs was substantial due to product replacement. We provide 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. Unlike the substitution-capacity finding in Ma et al. (2025) for the 2018-2019 tariffs, our evidence for 2025 tariffs 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 despite modest price increases, households reduce purchases even on goods facing negligible tariffs in 2025.

[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 (2019) 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.

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; Coibion and Gorodnichenko, 2026), but direct tariff-sentiment measures linked to transaction data are rare. We show that stated reduction intentions predict revealed spending cuts on tariff-exposed goods, that these cuts operate through trading down to cheaper varieties rather than through fewer units, and that stated substitution and stockpiling intentions do not translate into corresponding revealed behavior. The pessimism-driven cut in non-essential spending is concentrated among middle-income households.

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. Section 4 shows through a simple accounting exercise that the real-income effect of the tariffs can explain at most a third of the spending contraction, and quantifies the welfare cost of the price channel across income groups. Section 5 uses the linked household survey to identify forward-looking pessimism as the main channel behind the remainder of the contraction and to trace its distinct incidence across income groups. Section 6 concludes. Additional analyses are provided in the Appendix.

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 rate, 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 of various general merchandise and grocery stores, including Walmart, Amazon, and many 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 groups, education levels, age groups, 5-digit zip code, household composition, and race/ethnicity etc. The dataset provides 16 categorical income brackets that are self-reported by panelists.

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 rotating 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. Appendix A provides more information on data collection and representativeness.

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, other promotions 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 around 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. Groceries make up around half of the spending. Approximately 20 percent of spending is on health and beauty, and home and garden goods and around 12 percent is on apparel and household goods. Electronics spending is around 4 percent of the total

expenditure captured. In other words, the dataset primarily captures expenditures on physical goods purchased through both brick-and-mortar and online retailers.

Some spending categories have limited or no coverage. For instance, 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 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 of the data 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. Appendix A.5 demonstrates the close alignment of the aggregated Numerator spending series with U.S. Census Bureau's retail sales.

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. U.S. tariffs are generally imposed at the HS8 level and trade data can be aggregated to broader 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 Bureau. First, monthly

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

Notes: 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.

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 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 $\times$ month tariff rates spanning January 2024 through December 2025.

Second, import penetration ratios are constructed from 2022 trade and production data. 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

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

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.

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 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

[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.

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. Both $\tau_{h,t}$ and $\text{ImportPenetration}_h$ are ratios rather than percentages, so tariff exposure is also a ratio bounded between zero and one. For example, a 10 percent tariff on electronics (with import penetration near 0.9) yields a tariff exposure of 0.09, equivalent to 9 percent, while the same tariff on dairy products (with import penetration near 0.1) yields an exposure of 0.01, equivalent to 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 household 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.[6]

[6]Our approach is akin to Bai and Stumpner (2019); Bai et al. (2025) which use the NielsenIQ data. We map 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,

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. Appendix E replicates our qualitative results using a 6-digit HS code classification.

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	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.

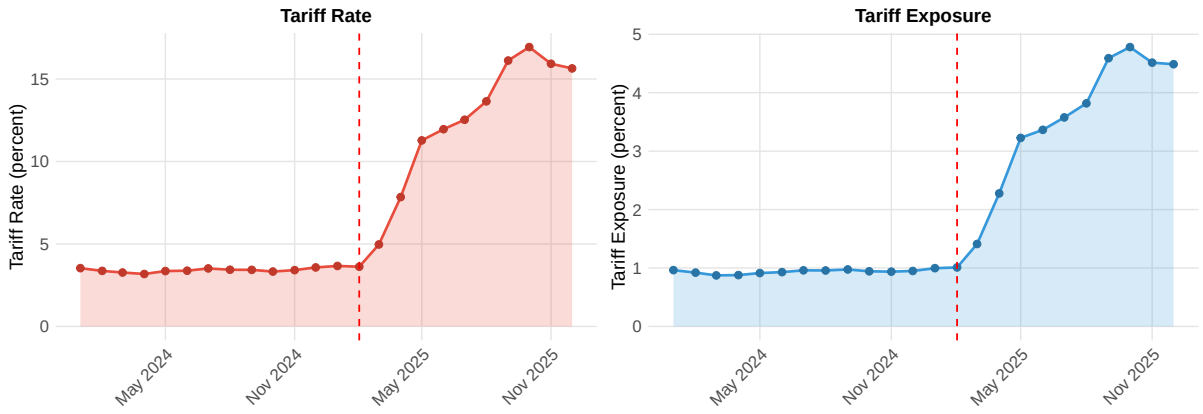
With a household-to-trade data crosswalk in place, we can visualize tariff rates and tariff exposure at the aggregate level for products in our household data. Figure 1 shows the evolution of both measures over the sample period. Unweighted average tariff rates are low and stable throughout 2024, then rise sharply after tariffs are imposed in February 2025. The mean post-period tariff rate across HS4 codes is around 11.9 percent, up from 3.4 in the pre-tariff period. The mean post-period tariff exposure is approximately 3.4 percent, up from 0.9 percent in the pre-period, reflecting the fact that many product categories have substantial domestic production that partially shields consumers from

where UPCs are only consistently available across a small share of products in the household data but the full retail hierarchy is universally available. Restricting the sample to UPC-identified products would sharply reduce coverage and compromise the representativeness of household consumption baskets. Finer classification of the Numerator data using UPCs or other product-level detail could in principle be achieved with internet-enabled AI, as in Cavallo et al. (2025). The proprietary nature of the dataset and institutional restrictions, however, prevent us from pursuing this route.

[7] A closed artificial intelligence system 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.

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). The red dashed vertical line marks February 2025, when tariffs were implemented.

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

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 (nominal 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. This approach follows Cavallo et al. (2025) and removes product-specific linear pre-trends estimated over the pre-tariff period before estimating 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

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

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 term $\delta_{i,m(t)}$ represents an interaction term that absorbs product-specific seasonality, as $m(t) \in \{1, \ldots, 12\}$ maps each year-month to its calendar month. Allowing each product category to have its own seasonal baseline accounts for unit-specific cyclical demand variation that might otherwise confound the tariff effect. The coefficient β shows 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.

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.

For period t relative to base period 0, which in our sample is January 2024, 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 Numerator Fisher Price Index for a longer time period of January 2018 to December 2025 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 and has low import penetration, see Appendix Figure C.2. We present both specifications to verify that its dominance does not drive the aggregate results. Additionally, we examine heterogeneity by income and essentiality by estimating the main specification using Fisher Price Indices constructed separately for each subgroup.

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 and Panel B uses the tariff exposure measure (tariff rate × import penetration). For each treatment definition, we report the results for prices (with and without pre-trends), spending and quantities.

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 shows larger effects across all outcomes, with a price coefficient of 0.20, indicating that the exposure measure better isolates categories where tariffs generate meaningful consumer impact.

With these coefficients in hand, we calculate each household's post-period tariff exposure and tariff rate, weighted by the household's 2024 expenditure shares across the categories they buy, and apply the benchmark coefficients to that level. The exposure weight here is the household basket exposure HHExp_i defined formally in Section 5.2. This yields an implied price increase and spending decline for each household, and we can focus on the average over the panel to gauge the price and spending effects of tariffs. For the average household, the tariff rate specification implies a price increase of about 1.7 percent and a real spending decline of about 5.2 percent, and the tariff exposure specification implies a price increase of about 0.9 percent and a spending decline of about 2.9 percent. In both cases, the spending contraction is about three times the implied price increase.[9] The not-detrended coefficients are slightly smaller, with slightly weaker

[9]The dispersion across households is wide, with the most exposed tenth facing spending cuts roughly twice as large as the least exposed tenth. The threefold wedge between spending and prices, however,

TABLE 3: Main Results: Price, Spending, and Quantity Responses to Tariffs

	Price		Spending	Quantity
	Detrended	Not Detrended		
	(1)	(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. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

model fit.[10]

The price coefficients in Table 3 might suggest that retail pass-through was modest. But a percentage change in the shelf price understates how much of the tariffs actually reached the shelf. A retail price is the sum of a foreign producer cost and a domestic distribution margin (Sangani, 2026b), and the tariffs fall only on the foreign-cost portion. When that portion is a small fraction s of the retail price, a tariff cost that passes through fully in dollars still moves the shelf price by only a small percentage. The percentage coefficient therefore understates the dollar incidence by the factor s, and the implied

holds across the distribution.

[10]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.

pass-through into the border cost is β_{retail}/s. We evaluate this at the tariff-rate coefficient of 0.15, the estimate that carries the percentage interpretation.

We can gauge the size of s in two ways, and both imply that dollar pass-through is substantial. Externally, Bai et al. (2025) estimate a foreign cost share near one third for tariff-affected consumer goods in 2018–2019, which places our implied dollar pass-through at roughly one half. From our own data, the foreign cost share is $s_h \approx (1 - \delta) \times \text{ImportPenetration}_h$, where δ is the domestic distribution margin, the share of the shelf price accounted for by wholesale and retail markups, transport, and in-store services, which retail studies place at 0.6 to 0.8. Our spending-weighted import penetration is 0.42, so imports account for about 42 percent of the value of the goods households actually buy, which implies a foreign cost share at or below one third and a dollar pass-through from roughly complete to well above it.[11] This uses category-level import intensity rather than the product-level border-to-retail unit-value ratio in Bai et al. (2025), so it is an approximation, but it points the same way. The two calculations bracket dollar pass-through between about one half and complete or more, all far above the face value of the retail coefficients. Our retail estimates are therefore best read as substantial dollar pass-through compressed into a small percentage by the foreign-cost share, not as evidence of low consumer incidence.

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.[12] The most noteworthy feature of Table 3, however, is that 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. In our context, this ratio is a comparison across equations estimated on the same treatment variable but with different dependent

[11] Applying $s_h \approx (1 - \delta) \times \text{ImportPenetration}_h$ to our spending-weighted import penetration of 0.42, a distribution margin of $\delta = 0.6$ gives $s \approx 0.17$, $\delta = 0.7$ gives $s \approx 0.13$, and $\delta = 0.8$ gives $s \approx 0.08$, all at or below one third. The implied dollar pass-through β_{retail}/s evaluated at the tariff-rate coefficient of 0.15 is then roughly 0.9, 1.2, and 1.9 across this margin range, that is broadly complete. Anchoring instead on the tariff-exposure coefficient of 0.20 would imply an even higher level pass-through, though that coefficient already embeds import penetration and so is not the appropriate numerator for a calculation that divides by a foreign cost share proportional to import penetration.

[12] 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.

variables and different levels of aggregation, and should not be read as a structural 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.

That the implied elasticity exceeds plausible values for e.g. groceries and household goods suggests that quantity adjustments operate through channels beyond price responses, such as 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, in Appendix D.2, leave one sector out at a time. Appendix E further shows that the aggregate results are robust to rebuilding the entire pipeline at the more granular six-digit (HS6) product level rather than our benchmark HS4, and that the same HS6 machinery applied to the earlier 2018–2019 tariff episode yields a comparable and partial price response but no detectable spending or quantity response.

Figure 2 presents event study estimates that trace 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 (percent) 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

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.

the end of the sample.[13] Under the Rambachan and Roth (2023) bounds on parallel-trends violations, allowing any pre-existing differential trend to continue linearly into the post-period leaves the price, spending, and quantity effects all significant.[14]

[13] 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.

[14] The Rambachan and Roth (2023) method constructs confidence sets that remain valid even if parallel trends does not hold exactly. We implement two classes of restrictions. The relative magnitudes approach

We show in Appendix D.3 that the spending and quantity decline reported in this section is not the unwinding of stockpiling behaviour in the wake of tariff announcements, including a timing test for front-loading in the early post-period reported in Appendix G.7. Also, note that weighting each HS4 by its fixed 2024 pre-tariff expenditure share in the regressions leaves our results qualitatively unchanged and if anything, produces a wider wedge between spending and price than the equal-weighted benchmark, which we therefore interpret as a conservative lower bound.

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. We continue to report tariff rate results wherever appropriate, 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.[15] We then estimate income-specific versions of Equations (4) and (5) where the

bounds the post-treatment trend violation as a fraction $\bar{M}$ of the largest pre-treatment violation. 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. At $\bar{M} = 0$ or $M = 0$ the spending and quantity confidence intervals exclude zero, and under the smoothness approach at $M = 0$ the average post-treatment price effect also excludes zero.

[15] 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.

price measure is the income-group specific FPIs.

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 Equations (4) and (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. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 4 presents price, spending, and quantity responses by income group.[16] Three patterns emerge from the benchmark tariff exposure specification in the second row of Panel A. First, price coefficient 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, lending support to the cheapflation phenomenon documented in Cavallo and Kryvtsov (2024) and Sangani (2026a). 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 com-

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

petitive pricing and consume products with fewer close substitutes would need to absorb higher costs. Consistent with this, our income-specific Fisher Price Indices show cumulatively higher price growth for the low-income basket than for the middle- and high-income baskets over the sample period (Appendix A.5, Figure A.4). 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.

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. The tariff rate specification tells the same story, with spending declines broadly similar across groups.

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.

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 the 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	0.09	0.09	−0.37***	0.08	−0.34**
	(0.04)	(0.06)	(0.11)	(0.12)	(0.08)	(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	0.10	0.10	−0.46**	0.03	−0.39**
	(0.05)	(0.08)	(0.14)	(0.19)	(0.13)	(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 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.

Our quantity measure counts the items purchased within an HS4-month, aggregated

across the distinct products mapped to each code, and is not a measure of physical volume. Package size is recorded only for products carrying a UPC, which do not span all spending, so we cannot build a consistent volume measure or isolate this margin directly. This bears on the concern of Bai et al. (2025), which show that tariffs can raise unit values through replacement toward smaller package sizes. Shrinkflation of this kind leaves a count of items unchanged and, if anything, causes our measure to understate the true fall in volume, so it works against the quantity and spending contraction we document rather than producing it. Moreover, we observe that the number of distinct products a household buys falls in our sample, whereas repackaging the same consumption into larger units would leave that set roughly similar. The reverse substitution, toward larger packs/bulk-buying, could in principle create as a spurious decline, but it does not drive our results as the contraction operates across products rather than within them. The quantity response is flat conditional on the individual product and unchanged when we restrict the sample to products purchased both before and after tariffs. In addition, both margins are most relevant for grocery and other frequently repackaged staples, and our headline results hold when grocery is excluded, with the contraction concentrated in non-essential goods where product turnover of this kind is a smaller concern.

Because our exposure measure has the algebraic form of a shift-share design, we follow Borusyak et al. (2025) and address the inference concerns that such designs are prone to, and find our headline results robust.[17]

[17]The estimated price, spending, and quantity responses are unchanged and remain statistically significant when we compute standard errors more conservatively, including clustering by broad sector and allowing for correlation across categories within each month, and the coarser groupings if anything tighten precision rather than loosen it. The results are also spread across many product categories rather than concentrated in a few, and they survive dropping any single sector as well as excluding the most China-exposed categories, so they are not driven by the small number of policy actions behind the 2025 tariffs. The instrumental-variables apparatus of that literature does not apply here, since tariffs are directly observed and our identification rests on standard difference-in-differences comparisons rather than an instrument. These results are available upon request.

4 Real-Income Effects and the Welfare Cost of Tariffs

This section takes the first step toward explaining the wedge between the spending and price responses and quantifying its cost. We begin with a simple accounting exercise that asks how much of the contraction can be a consequence of the real income the tariffs take away, and show that the answer is at most about one-third. The greater part of the spending response is therefore behavioral, and Section 5 uses the household survey to identify the channel behind it. We then quantify the welfare cost of the price increases households actually paid and document that this cost is sharply regressive, with low-income households bearing roughly nine to ten times the burden of high-income households as a share of income on the price-anchored measures, and about five times on the net spending effect. In this section, i indexes households and h indexes HS4 product codes.

4.1 How Much of the Spending Contraction Is a Real-Income Effect?

Why does spending on tariff-exposed goods contract by so much more than prices rise? One candidate explanation is that tariffs reduce households' real income, and a household with less real income spends less, so part of the contraction may simply be an income effect that requires no change in beliefs or in the way households shop. This subsection asks how large that mechanical channel can be.

The logic is an accounting identity rather than an estimate. At any level of tariff exposure, the real-income loss a household suffers on its exposed basket equals the price increase it pays, which our estimates put at $\hat{\beta}_{\text{price}}$ per unit of exposure. The observed contraction in real spending on the same basket is $\hat{\beta}_{\text{spend}}$ per unit of exposure. A household that responds to the income loss alone reduces spending by its marginal propensity to consume, κ, times that loss. The share of the observed contraction attributable to the

income effect is therefore

$$\text{Income-effect share} \;=\; \kappa\,\frac{\hat{\beta}_{\text{price}}}{\hat{\beta}_{\text{spend}}}. \tag{9}$$

Two features of this expression are worth noting. First, the level of tariff exposure cancels because it scales the income loss and the observed contraction by the same amount, so the share is the same whether we evaluate it at low or high exposure. The same cancellation applies to expenditure weights, since the income loss and the observed contraction are aggregated with the same basket weights and those weights drop out of their ratio. The magnitudes in dollars scale with exposure, but the share does not. Second, the ratio $\hat{\beta}_{\text{price}}/\hat{\beta}_{\text{spend}}$ places a ceiling on this channel, reached when the marginal propensity to consume equals one and the household reduces real spending one-for-one with its real income loss.[18] We evaluate that ratio at the benchmark Panel B estimates of Table 3, the detrended price coefficient in column 1, $\hat{\beta}_{\text{price}} = 0.20$, and the spending coefficient in column 3, $\hat{\beta}_{\text{spend}} = -0.63$, which gives $0.20/0.63 \approx 0.32$. Even in the extreme case where households cut spending one-for-one with their lost income, i.e. $\kappa = 1$, the income effect explains no more than a third of the contraction. Using the tariff rate specification results would lead to a similar value.

TABLE 6: Share of the Spending Contraction Attributable to the Income Effect

Marginal propensity to consume κ	0.1	0.3	0.5	0.7	0.9
Income-effect share (benchmark)	3%	10%	16%	22%	29%
Income-effect share (non-essentials)	2%	7%	11%	15%	20%

Notes: Each entry is the income-effect share of the spending contraction from Equation (9), $\kappa\,\hat{\beta}_{\text{price}}/\hat{\beta}_{\text{spend}}$, at the indicated marginal propensity to consume κ. The benchmark row uses the aggregate tariff-exposure coefficients from Table 3 ($\hat{\beta}_{\text{price}} = 0.20$, $\hat{\beta}_{\text{spend}} = 0.63$), giving a ceiling of 0.32 at $\kappa = 1$. The non-essentials row uses the non-essential coefficients from Table 5 ($\hat{\beta}_{\text{price}} = 0.10$, $\hat{\beta}_{\text{spend}} = 0.46$), giving a ceiling of 0.22. The share is invariant to the level of tariff exposure and to expenditure weighting.

Table 6 evaluates Equation (9) across a range of values for the marginal propensity

[18] Setting $\kappa = 1$ bounds the channel rather than the spending response itself. A household can in principle cut spending by more than its real income fell, saving the difference, which corresponds to an effective propensity to consume above one. Such a cut is precisely the precautionary behavior we seek to isolate and does not belong in the benchmark, so the ratio remains the relevant ceiling on the income-effect share.

to consume. The empirical consumption literature places the marginal propensity to consume out of transitory income shocks broadly between 0.2 and 0.7 (Johnson et al., 2006; Parker et al., 2013; Jappelli and Pistaferri, 2014). Across that range the income effect accounts for only about 6 to 22 percent of the spending contraction, and the ceiling reached at a marginal propensity to consume of one is still only about a third. If households absorb part of the income loss through saving or through categories of consumption not observed in our data, the share falling on the goods we measure is smaller still, so these figures are best read as upper bounds. Applying the non-essential coefficients, where pass-through is smaller but the spending response is larger, lowers the share further.

The accounting therefore leaves a large residual. The greater part of the spending contraction is not a response to lower real income but a behavioral one, whereby households change how much and what they buy by more than their reduced purchasing power alone requires. Characterizing that behavioral residual is the task of the survey analysis in Section 5, which identifies precautionary reallocation of the basket toward essentials and trade-down to cheaper varieties within categories as the margins on which it operates.

4.2 The Welfare Cost of the Price Channel

The accounting above bounds the income channel but does not price the burden households bear from the price increases they actually face. We now quantify that welfare cost. We present two price-anchored welfare measures alongside a net spending effect that serves as the dollar counterpart of the observed contraction. The gap between them is the welfare-accounting reflection of the wedge, and all measures deliver the same regressive pattern across income groups.

Before turning to the measures, we clarify what they calculate. 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. In other words, 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. 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. Our welfare totals therefore place a lower bound on the total household-level cost. Also 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 the three measures in turn.

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

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

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:

$$\text{PriceBurden}_{i,h} = \text{Spend}_{i,h}^{\text{pre}} \times |\hat{\beta}_{\text{price}}^{g(i)}| \times \overline{\text{TariffExposure}}_{h}^{\text{post}} \tag{11}$$

$$\text{DWL}_{i,h} = \frac{1}{2} \times \text{Spend}_{i,h}^{\text{pre}} \times |\hat{\beta}_{\text{price}}^{g(i)} \times \overline{\text{TariffExposure}}_{h}^{\text{post}}| \times |\hat{\beta}_{\text{qty}}^{g(i)} \times \overline{\text{TariffExposure}}_{h}^{\text{post}}| \tag{12}$$

$$\text{WelfareLoss}_{i,h}^{\text{DWL}} = \text{PriceBurden}_{i,h} + \text{DWL}_{i,h} \tag{13}$$

The Harberger formulation requires only the local price and quantity responses and a linear demand approximation around the pre-tariff equilibrium (Hausman, 1981). It is valid when foreign supply is perfectly elastic, which recent evidence supports, since the bulk of tariff-induced price changes is borne by domestic buyers rather than absorbed by foreign exporters (Amiti et al., 2026; Hinz et al., 2026).[19]

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

$$\Delta \log P_g^{SV} = \sum_h w_{g,h}^{SV} \cdot \hat{\beta}_{\text{price}}^{g} \cdot \overline{\text{TariffExposure}}_h^{\text{post}} \tag{14}$$

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

$$w_{g,h}^{SV} = \frac{L(\omega_{g,h}^{post}, \omega_{g,h}^{pre})}{\sum_{h'} L(\omega_{g,h'}^{post}, \omega_{g,h'}^{pre})}, \qquad L(a,b) = \frac{a-b}{\log a - \log b} \tag{15}$$

Pre-tariff shares $\omega_{g,h}^{pre}$ are computed from 2024 spending and post-tariff shares $\omega_{g,h}^{post}$ are computed symmetrically from observed post-period spending and renormalized within income group. The index measures the minimum expenditure needed to maintain pre-tariff utility at post-tariff prices, allowing for substitution. The logarithmic-mean weights deliver this property such that each category enters with the log mean of its pre- and post-tariff expenditure shares, so its contribution reflects its weight in the basket both before and after households substitute, which is what makes the index an exact cost-of-living measure rather than a fixed-basket approximation. Only $\hat{\beta}_{\text{price}}^{g}$ enters the Sato-Vartia construction, and substitution is captured non-parametrically through observed expenditure reallocation rather than through an estimated elasticity. The household-level welfare loss is:

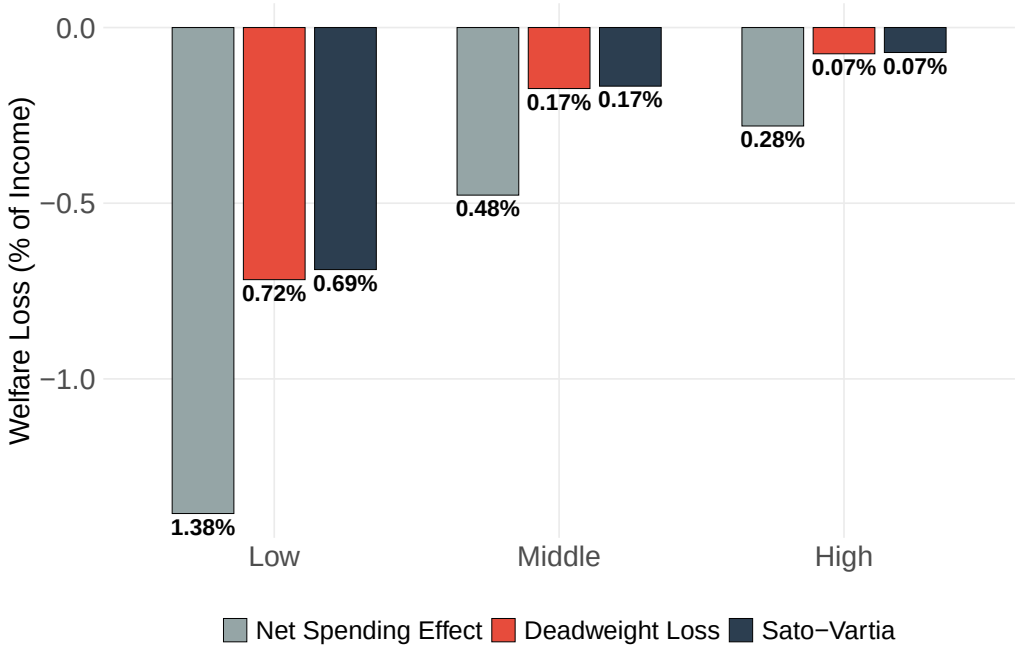
$$\text{Welfare}_i^{SV} = \text{TotalSpend}_i^{\text{pre}} \cdot \Delta \log P_{g(i)}^{SV} \tag{16}$$

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

The deadweight loss and the Sato-Vartia measure are complementary. The DWL uses the local quantity elasticity together with a linear demand approximation and 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: $\text{WelfareLoss}_i = \sum_h \text{WelfareLoss}_{i,h}$.

Figure 3 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.[20] The welfare burden of the 2025 tariffs is sharply regressive when measured as a share of household income.

FIGURE 3: 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 (10)–(16) 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.

Under the deadweight loss measure, low-income households bear a welfare cost of 0.72

[20]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.

percent of income, compared with 0.17 percent for middle-income and 0.07 percent for high-income households. The Sato-Vartia measure yields 0.69, 0.17, and 0.07 percent respectively. Under both price-anchored welfare measures, low-income households bear roughly nine to ten 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 1.38 percent of income versus 0.48 and 0.28 percent for middle and high-income groups, respectively, a ratio closer to five. High-income households reduce spending by more in dollar terms but can readily afford to do so and 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 factors. 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, as documented in Table 4. Second, low-income households have the smaller 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 7 presents all three measures side by side, expressed both in dollars per month and as a share of spending on tariff-affected goods. The net spending effect averages about $33 per month but varies by income. High-income households cut spending by about $44 per month while low-income households cut by about $23 per month. In contrast, the deadweight loss and Sato-Vartia measures show similar costs across the income distribution, averaging approximately $12 per month for all groups.

Two further points stand out from Table 7. First, the deadweight loss and Sato-Vartia measures agree closely in every income group despite sharing only the price coefficient in their construction, indicating that the price-anchored welfare cost is robust to the measurement assumption.[21] Second, the net spending effect substantially exceeds both price-anchored measures across all income groups. This gap is the welfare-accounting counterpart of the wedge documented in Section 3. The price-anchored measures capture the cost of the price increases households paid, while the net spending effect captures the full behavioral contraction. The DWL and Sato-Vartia measures should therefore be read as a lower bound on the consumer welfare cost of the 2025 tariffs.

[21] The welfare analysis applies income-specific but not essentiality-specific coefficients.

TABLE 7: 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	23.05	11.96	11.48	2.31	1.20	1.15
Middle	64,943	32.79	11.94	11.47	2.62	0.95	0.91
High	31,739	43.78	11.68	11.11	2.81	0.75	0.71
All	124,833	33.39	11.88	11.38	2.62	0.93	0.89

Notes: Net spending effect computed using Equation (10); DWL using Equations (11)–(13); Sato-Vartia using Equations (14)–(16). 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. 1,615 do not have an income group. Their mean annual spending is close to that of the retained high-income households, so the regressivity comparison stands on a representative sample.

The deadweight loss calculation assumes perfectly elastic foreign supply, so that domestic consumers bear the entire price change. What matters for our welfare estimates is how far their magnitude would move if this assumption were relaxed. Suppose consumers bear only a share ϕ of the price change, with foreign producers absorbing the remaining $1 - \phi$, so that the benchmark is the case $\phi = 1$. Because ϕ multiplies every household's loss, the welfare level scales one-for-one with it, and the size of any revision is governed entirely by how far ϕ falls below one. The evidence places ϕ close to one, with the bulk of tariff-induced price changes borne by domestic buyers rather than absorbed by foreign exporters (Amiti et al., 2026; Hinz et al., 2026), so the welfare estimates would fall by at most about a tenth relative to the benchmark. A common ϕ leaves each income group's relative burden intact, so the regressive incidence is unchanged.

Overall, the welfare burden of the 2025 tariffs is regressive, with low-income households bearing a disproportionate cost. In the next section, we turn to the survey evidence that identifies the channel behind the large wedge between the price and spending responses to tariffs.

5 Evidence from Household Survey Responses

The accounting in Section 4.1 establishes that an income effect can explain at most a third of the spending contraction on tariff-exposed goods. This section uses a survey run on a subset of households in the panel to identify what drives the remainder. The headline

finding is that the behavioral residual is a forward-looking precautionary response driven by pessimism. More pessimistic households reallocate their baskets toward essentials and trade down to cheaper varieties within the essential categories they worry about. The wedge between modest price pass-through and the large spending contraction thus reflects household beliefs about the economic outlook rather than the purchasing power households have already lost.

The survey-based estimates answer a narrower question than the aggregate regressions. Where the aggregate regressions estimate the total effect of tariff exposure averaging across all households in an income group, the survey-based estimates ask whether those who are more worried or who report specific intent behave differently. This isolates the portion of the 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.

It is useful to organize the survey findings around three margins of adjustment. Households can buy less, purchasing fewer units of the goods they were already buying. They can buy cheaper, holding units fixed while switching to lower-priced varieties within a category. Or they can buy different, reweighting the basket across categories toward essentials. The aggregate difference-in-differences results in Section 3 identify the buy-less margin, which is common to all households. In this section that common response is absorbed by the time and category-by-time fixed effects, so the survey estimates identify the other two margins.

After we introduce the survey design in Section 5.1, Section 5.2 takes up the buy-different margin and documents the headline basket-reallocation result. More pessimistic households cut non-essential spending while holding essential spending flat, even after netting out the income loss they experience through their tariff exposure, with the significant sentiment differential appearing among middle-income households.

Section 5.3 asks where these cuts land. Pessimistic households cut spending on goods facing almost no tariffs as much as they cut tariffed goods, which points to beliefs rather than to price effects. Sections 5.4 and 5.5 then take up the buy-cheaper margin. Households who flag an essential category as likely to be affected by tariffs trade down to

cheaper varieties, and this response runs across the income distribution. Households who state an intent to reduce spending trade down in the same way on tariff-exposed goods, though we observe this only among middle-income households. Section 5.6 summarizes our findings.

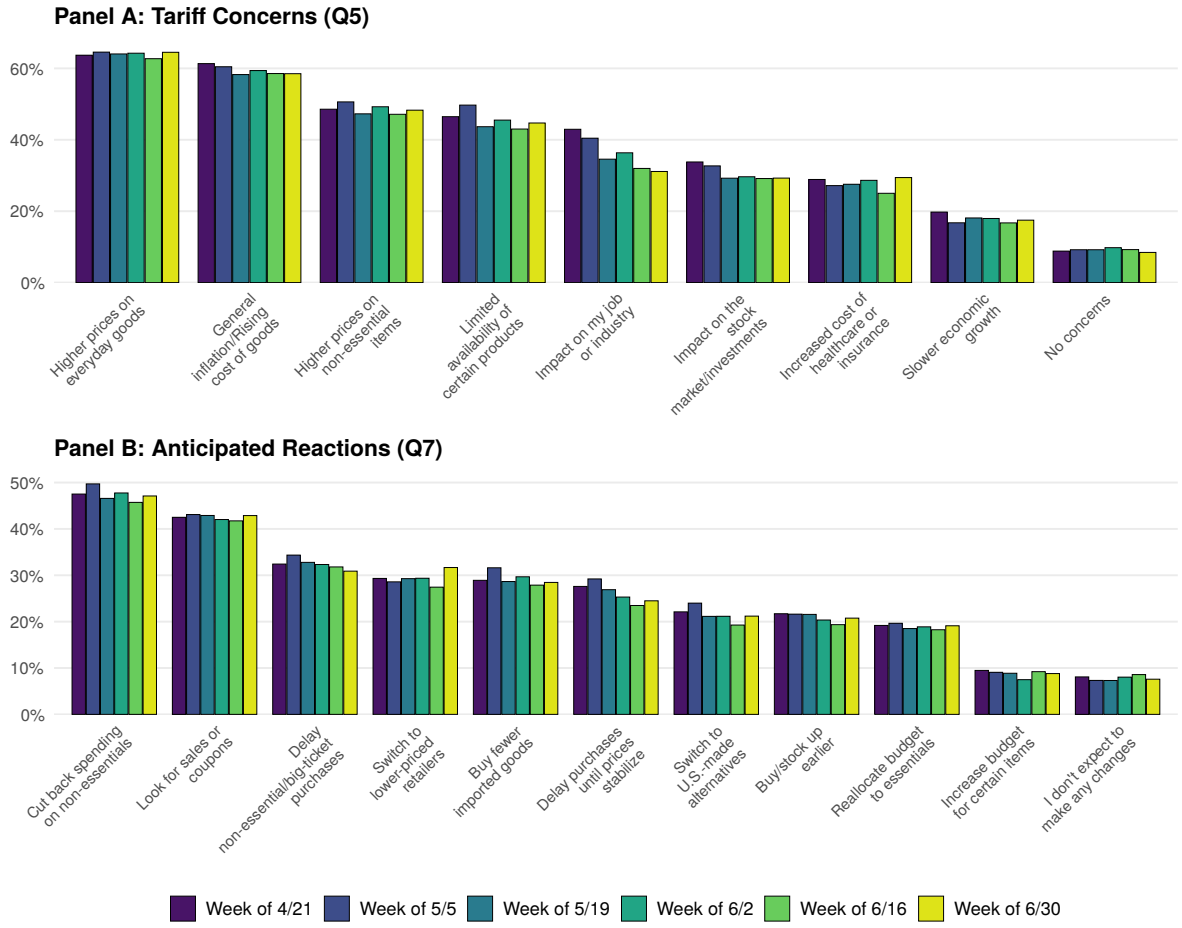
5.1 Survey Design and Measurement

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 G.1 lists the full questionnaire.

Figure 4 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. We provide a short discussion on this in Section 5.6.

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), and 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

FIGURE 4: 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.

effects α_i and year-month fixed effects γ_t. Standard errors are clustered at the household level throughout.

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. 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.[22] We use Sent_i as

[22]For each of the four sentiment items, we subtract the cross-household mean and divide by the

a generic notation for the household-level sentiment measure, and 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.

5.2 Reallocation Toward Essentials Is Not Just Lower Income

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. 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. This prediction is also what recent work on the 2025 episode would lead one to expect. Candia et al. (2026) show that expected tariff levels and tariff uncertainty push durable spending in opposite directions, the former encouraging intertemporal substitution toward current purchases and the latter discouraging spending, so that in an episode of exceptionally volatile policy the precautionary force is the one likely to dominate.

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

$$\text{ShareEssential}_{i,t} = \alpha_i + \gamma_t + \beta\,(\text{Sent}_i \cdot \text{Post}_t) + \varepsilon_{i,t}, \tag{17}$$

$$\ln Y_{i,t}^{k} = \alpha_i + \gamma_t + \beta\,(\text{Sent}_i \cdot \text{Post}_t) + \varepsilon_{i,t}, \quad k \in \{\text{ess}, \text{noness}\}, \tag{18}$$

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 economic-outlook 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. Households answering "Not sure" on the outlook items lack a defined component score, so the PC1 specifications retain about 84 percent of the household-months available for the Q6 binary measure. The Q6 measure spans the wider sample and returns the same reallocation pattern, so the principal-component sample is representative of the population the mechanism describes.

where α_i is a household fixed effect and γ_t is a year-month fixed effect. Coefficient β identifies whether pessimistic households shift their behavior in the post-period differentially relative to less-pessimistic households in the same month.

TABLE 8: 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 (17) (columns 1–2) and (18) (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 household-months with positive essential and non-essential spending. 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 8 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 of non-essential spending falls by 0.9 percent per standard deviation of PC1, while essential spending does not move. 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, because households protect essential consumption when they contract.

Because tariffs reduce real income, and a poorer household devotes a larger share of its budget to necessities through a standard income effect, the reallocation documented above could in principle reflect lower income rather than beliefs. The welfare accounting of Section 4.1 bounds the income channel well below the size of the ag-

gregate contraction, but that is an aggregate statement and does not by itself establish that the household-level reallocation survives the income effect. We therefore measure each household's own tariff-induced income shock and include it in the reallocation specification alongside sentiment. Our baseline measure is the predetermined tariff burden $\text{HHExp}_i = \sum_h \omega_{i,h} \overline{\text{TariffExposure}}_h$, the household's 2024 expenditure-share-weighted post-period tariff exposure, where $\omega_{i,h}$ is household i's share of 2024 spending on HS4 code h and $\overline{\text{TariffExposure}}_h$ is the mean post-period exposure of code h. We augment Equations (17) and (18) with this measure interacted with the post-period indicator,

$$Y_{i,t}^{k} = \alpha_i + \gamma_t + \beta \left(\text{Sent}_i \cdot \text{Post}_t\right) + \phi \left(\text{HHExp}_i \cdot \text{Post}_t\right) + \varepsilon_{i,t}, \qquad k \in \{\text{ess}, \text{noness}\}, \tag{19}$$

where ϕ captures the income channel and β captures the belief channel that remains after the income effect is accounted for. If the reallocation were purely an income effect, the sentiment coefficient β would fall toward zero once the income effect is included.

TABLE 9: Reallocation Toward Essentials: Sentiment and the Income Effect

	Share essential (pp)	$\ln Y^{\text{noness}}$ (%)
$\text{Sent}_i \cdot \text{Post}_t$ (PC1)	0.145**	−0.008**
	(0.063)	(0.004)
$\text{HHExp}_i \cdot \text{Post}_t$	1.403***	−0.048***
	(0.055)	(0.003)
Num. Obs.	382,866	377,845

Notes: Estimates of Equation (19) on the household-month panel of survey respondents. Column 1 uses $\text{ShareEssential}_{i,t}$ in percentage points and column 2 uses log real non-essential spending. The specifications are those of columns 1 and 4 of Table 8 augmented with $\text{HHExp}_i \cdot \text{Post}_t$, and the sentiment-only estimates reported there are the relevant point of comparison. HHExp_i is expressed in percentage points, so its coefficients give the response per percentage point of household tariff exposure, and in column 2 the entry multiplied by 100 is the approximate percent change. Sent_i is the first principal component of the four sentiment items (PC1). 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 9 shows that households whose baskets were hit harder raise the share of the budget going to essentials and cut non-essential spending, and this response is large and precisely estimated.[23] At a high-exposure basket with HHExp_i of 7 percent, around

[23] HHExp_i has a mean of 4.5 percent and a standard deviation of 1.2 percent, with most households between about 2 and 7 percent and a maximum of 13 percent.

the 97th percentile of the distribution, households shift roughly 10 percentage points of their budget toward essentials and cut non-essential spending by about a third.[24] When we account for the size of the income hit, the essential-share response to pessimism moves only from 0.18 to 0.15 percentage points per standard deviation of PC1 relative to the sentiment-only estimates of Table 8, and the non-essential spending cut is very little changed at 0.8 percent. In other words, among households with comparable tariff-induced income losses, the more pessimistic one still shifts its basket more sharply toward essentials. The stability of the sentiment coefficient reflects that pessimism and the size of the income shock are close to independent across households, i.e. households do not become more pessimistic simply because their own basket became more expensive.

Table 10 estimates the reallocation specifications separately by income group and shows that the non-essential level cut is prevalent in middle-income households, who reduce non-essential spending by 1.0 percent per standard deviation of PC1, while the low-income and high-income estimates are insignificant. This finding fits the "slack, not budget" interpretation. Middle-income households have the slack to execute the precautionary contraction on the discretionary margin. Low-income households face binding budget constraints that leave little room for additional discretionary contraction in response to pessimism, whereas high-income households have enough room that pessimism does not force a discretionary cut.

The conclusion is unchanged when we measure the income shock by the price change households actually paid rather than by predetermined exposure. Appendix G.3 re-estimates Equation (19) using the household's 2024-basket-weighted realized change in unit prices, alone and jointly with the exposure measure, and the sentiment coefficient retains its magnitude throughout.

Taken together with the welfare accounting, a tariff-induced income effect is real but can account for only a minority of the spending contraction. The larger, behavioral part is a precautionary response, in which households defend essential consumption and reallocate away from discretionary categories by more than their lost income requires,

[24] This exposure coefficient is not a pure measure of the income channel. Because high-exposure baskets are by construction tilted toward non-essential goods, the exposure interaction absorbs the direct effect of the tariffs on those baskets alongside any income effect, and its magnitude should not be read against the accounting ceiling of Section 4.1, which bounds the share of the contraction that a real-income effect can explain rather than the size of the cut itself.

TABLE 10: 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)	0.215	0.157*	0.185	0.006	−0.001	−0.001	−0.010	−0.010**	−0.006
	(0.150)	(0.086)	(0.127)	(0.006)	(0.003)	(0.005)	(0.009)	(0.005)	(0.008)
Num. Obs.		382,866			377,845			377,845	
R^2		0.460			0.635			0.551	

Notes: The table reports estimates of Equations (17) and (18) on the Numerator panel from a single pooled specification in which sentiment is fully interacted with 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 8. 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.

concentrated among the middle-income households with the discretionary slack to do so.

5.3 Pessimistic Households Cut Low-Tariff Goods Too

We now ask whether the contraction documented in Section 5.2 reflects beliefs about the economy or a stronger response to the prices tariffs raised. Households responding to prices would cut most where prices rose most. Households reading the tariffs as a signal about the economy would cut across the basket, whether or not particular goods were hit.

We test this directly on the goods the tariffs barely affected. As the U.S. had existing tariffs in place before February 2025, the bottom decile is the closest empirical analogue to an untariffed basket. Define the low-tariff basket as the HS4 codes in the bottom decile of post-period mean tariff rate $\tau_{h,t}$, with a cutoff of $\tau_{h,t}^{p10} \approx 0.045$. We first estimate the within-HS4 specification on these codes alone,

$$\ln Y_{i,h,t} = \beta \left(\text{Sent}_i \cdot \text{Post}_t\right) + \alpha_{i,h} + \lambda_{h,t} + \varepsilon_{i,h,t}, \tag{20}$$

and then estimate the same specification on the full panel, adding the low-tariff contrast,

$$\ln Y_{i,h,t} = \beta_1 \left(\text{Sent}_i \cdot \text{Post}_t\right) + \beta_2 \left(\text{Sent}_i \cdot \text{Post}_t \cdot \text{LowTariff}_h\right) + \alpha_{i,h} + \lambda_{h,t} + \varepsilon_{i,h,t}. \tag{21}$$

where $\text{LowTariff}_h \cdot \text{Post}_t$ is absorbed by $\lambda_{h,t}$ and $\text{Sent}_i \cdot \text{LowTariff}_h$ by $\alpha_{i,h}$, since both are constant within HS4. Coefficient β_2 tests whether pessimistic households cut differently on goods the tariffs did not reach.

TABLE 11: Sentiment and Spending on Low-tariff Goods

	Low-tariff goods (PC1)	All goods (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 (20) 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 $\text{LowTariff}_h \cdot \text{Post}_t \cdot \text{Sent}_i$ estimated on the full panel. 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 11 reports the estimates. A one standard deviation increase in pessimism reduces spending on the low-tariff basket by 0.2 percent, the same magnitude as the effect on tariff-exposed goods in the pooled specification. The contrast between the two is small and statistically indistinguishable from zero, so pessimistic households cut low-tariff and tariff-exposed goods by the same amount.[25] Pessimism reduces spending on goods the tariffs barely affected, whereas a demand response to tariff-driven price increases would appear only where the tariffs raised prices. The term Sent_i records a household's assessment of where the economy is heading, and a household holding such a view has no reason to know which HS4 codes were tariffed, so their cuts fall on the basket as a whole rather than on the goods tariffs reached.

Appendix G.4 reports a companion test built on household-specific exposure rather than on the HS4 tariff rate, and finds the same result. The two measures we turn to next are tariff-specific, and both track exposure because households buying heavily in exposed categories are the ones likely to notice.

[25] The 95 percent confidence interval on the contrast runs from roughly −0.0023 to +0.0013. The upper bound is smaller in magnitude than the direct pessimism effect on the low-tariff basket, so tariff-exposed goods are at most marginally more sensitive to pessimism than low-tariff goods, well short of the gap a demand response to prices would produce.

5.4 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 G.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

$$Y_{i,h,t} = \beta_1 \left(\text{Worried}_{i,h} \cdot \text{Post}_t\right) + \beta_2 \left(\text{Worried}_{i,h} \cdot \text{Post}_t \cdot \text{Essential}_h\right) + \alpha_{i,h} + \lambda_{h,t} + \varepsilon_{i,h,t}, \quad (22)$$

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 12 reports the results. In the pooled specification, worried households cut spending on flagged categories by 1.0 percent relative to unworried households in the same HS4-month, with no quantity response. In essential flagged categories, the implied spending response is a 1.4 percent cut, and quantities again do not move. Spending falls, units hold, so unit prices paid fall, which imply trading-down to cheaper varieties within the flagged category. In non-essential flagged categories the pattern inverts. Spending rises by 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 12 operates on the intensive margin conditional on continued category participation.

The findings in this section shed light on the aggregate essentiality result in Section 3, where essentials show no spending or quantity response to tariff exposure. The estimate

TABLE 12: 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 (22). 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 G.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.

here compares worried and unworried households buying in the same category and month, and it shows that worried households shift toward cheaper varieties while their unit counts hold steady. What looks like no response at the category level is a compositional shift within it. The non-essential result works the other way, with worried households paying higher unit prices rather than trading down or buying less, which is the household-level version of the pass-through we document in the aggregate. Appendix G.6 shows the trade-down is uniform across the income distribution, and Appendix F documents the same substitution behavior in the full panel.

5.5 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 (a multi-select including reduce, substitute to U.S.-made, switch retailers, reallocate across categories,

stockpile, and buy fewer imports). For each intent, we estimate

$$Y_{i,h,t} = \beta \left(\text{TariffExposure}_{h,t} \cdot \text{Post}_t \cdot \text{Intent}_i \right) + \alpha_{i,h} + \lambda_{h,t} + \varepsilon_{i,h,t} \tag{23}$$

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 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.

We construct a pooled Reduce_i indicator for intent which 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.

TABLE 13: 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 (23) 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 13 reports the results. Only stated reduction intent translates into a differential cut on tariff-exposed goods, a 4.4 percent decline per unit exposure, while stated substitution and stockpile intent do not affect spending. The stated reducer's spending cut is accompanied by no quantity response. So stated reducers trade-down to cheaper varieties

within exposed HS4 codes, parallel to the within-worried-category pattern documented in Section 5.4. 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, as documented in Appendix F.2, Tables F.2–F.3. Stated substitution intent moves neither margin. Substitution-related responses dominate stated intentions in the survey, as Panel B of Figure 4 shows, yet produce no differential response on the exposed-good revealed margin. Stated stockpile intent leaves spending unchanged but delivers a small positive quantity response of 1.8 percent per unit exposure. The coefficient is an order of magnitude below the aggregate quantity contraction documented in Section 3 and cannot account for the wedge. Nor is it front-loaded. Appendix G.7 splits the post-period and finds no precisely estimated early-window response, which is where pull-forward would have to appear.

We construct an additional index Precaution_i that aggregates the four sub-items of Reduce_i most directly consistent with a precautionary response: cut back, delay, delay big-ticket, and reallocate, excluding buy-fewer-imports. We then re-estimate Equation (23) with each sub-item and the composite entered separately.

TABLE 14: 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 (23) with Intent_i set to the indicated Q7 sub-item. "Cut back" is Q7 item 1 (cut back spending on non-essentials); "Delay" is item 6 (delay until prices stabilize); "Delay big-tkt" is item 3 (delay non-essential or big-ticket purchases); "Fewer imports" is item 5 (buy fewer imported goods); "Reallocate" is item 9 (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 G.5, Table G.4. Standard errors clustered at the household level in parentheses. *, **, *** denote significance at the 10, 5, and 1 percent levels.

Table 14 reports the results. First, stated reallocation is the single largest and most precisely estimated driver of the exposed-good spending cut at 4.5 percent, with stated cutback on non-essentials providing a smaller secondary contribution of 3.0 percent. The precaution composite has a coefficient of 4.3 percent that is close to the reallocate estimate and within sampling variation of the pooled Reduce$_i$ coefficient in Table 13, confirming that the pooled reduction effect is driven almost entirely by the precautionary-reallocation subset of the five sub-items. A horserace specification that enters all five sub-items simultaneously, reported in Appendix G.5, confirms that reallocate is the only sub-item to survive. Second, the two timing-related sub-items, delay and delay of big-ticket purchases, move neither margin. 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 Appendix D.3. Third, the buy-fewer-imports sub-item exhibits a distinct pattern, with no spending response alongside a small negative quantity response of 1.6 percent. Stated import-substituters reduce units of tariff-exposed goods while paying roughly the same aggregate amount, which might be consistent with substitution toward U.S.-produced alternatives at higher unit prices.

A natural question is why high-income households, who exhibit sizable aggregate spending and quantity contractions in Section 3, show no differential response on any survey-based margin. Appendix Table G.7 estimates Equation (23) separately by income group for stated reduction intent, the precautionary composite, and stated reallocation. Middle-income stated reducers cut spending on tariff-exposed goods significantly by trading down. High-income estimates are small and imprecise 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.

5.6 Summary and Reverse Causality

The survey evidence supports the precautionary interpretation on five margins. First, pessimistic households shift their basket toward essentials, cutting non-essential spending while holding essential spending flat, with the significant sentiment differential among middle-income households. Second, this contraction is not limited to tariffed goods. Pessimistic households cut goods in the bottom decile of the tariff rate as sharply as they cut tariffed goods, so the cuts reflect beliefs rather than the prices households paid. Third, households who flag a category as likely to be affected by tariffs cut spending on it within HS4 while quantities hold, so unit prices paid fall which indicates trading-down to cheaper varieties within the category of concern. This holds in flagged essential categories. In flagged non-essential categories spending instead rises with quantities flat, which we read as absorption of higher unit prices where the within-category substitution margin is narrower. Fourth, stated reduction intent predicts a differential spending cut on tariff-exposed goods with no matching quantity cut, again trade-down rather than contraction, while stated substitution intent predicts nothing on the matched margin and stated stockpile intent predicts only a small positive quantity response an order of magnitude below the aggregate quantity decline. The one exception is buy-fewer-imports, which moves quantities down without moving spending, a pattern we interpret as substitution toward domestic alternatives rather than as reduction. Fifth, decomposing stated reduction intent identifies reallocation toward essentials as the significant driver, with non-essential cutbacks a secondary contributor and the two delay items flat.

One potential concern with this evidence is that sentiment is measured contemporaneously with the treatment, since 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.

Four features of the results presented in this section and of the policy environment mitigate this concern. First, the item that loads most heavily on our pessimism index is Q9, a forward-looking assessment of how the economy will perform over the coming year, rather than a measure of already-realized conditions. If sentiment were simply

proxying for realized exposure, the response would be driven by households' backward-looking experience of the price shocks they had already absorbed. That the dominant component is instead a forward-looking expectation about the future, rather than realized price changes, is consistent with the belief-driven contraction we document. 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 G.3 we re-estimate the reallocation specification controlling for the household's 2024-basket-weighted realized change in unit prices between the pre- and post-periods. The sentiment coefficient holds its magnitude and its significance, so tariff-anchored pessimism is not standing in for the price change households actually paid. Fourth, although our post-period begins in February 2025, tariff implementation remained volatile throughout the 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 trade data and a subset with survey-reported tariff sentiment, to study how the 2025 U.S. tariffs reshape household spending. The joint observation of prices, quantities, expected economic outlook, and stated intentions and revealed behavior in response to tariffs from a single source lets us document both the aggregate response and the behavioral channel behind it. Tariff exposure raises retail prices modestly, with pass-through decreasing in income, but reduces real spending by a far larger margin. Quantities and spending both decline sharply, and the contraction is sustained rather than front-loaded, weighing against intertemporal substitution as the primary channel. This wedge between modest pass-through and large contractions cannot be rationalized by pass-through alone, and a simple accounting exercise shows that the real income the tariffs take away ex-

plains at most a third of it. The survey identifies forward-looking pessimism as the main channel behind the remainder. Even after accounting for the tariff-induced income loss, pessimistic households reallocate toward essentials and trade down to cheaper varieties within the essential tariff-exposed categories they flag. Pessimistic households cut spending on goods facing low tariffs, with no detectable difference from their cuts on tariffed goods. An elastic demand response to the prices tariffs raised cannot generate that, which leaves beliefs as the operative channel. Stated stockpile and delay intentions, meanwhile, produce no evidence of front-loading.

Tariff effects operate through two channels with different income profiles. The price channel dominates at the bottom, where households face the highest pass-through, producing a welfare cost far greater as a share of income than at the top. The precaution channel driven by forward-looking pessimism operates through the composition of spending. Within-category trade-down to cheaper varieties in flagged essentials runs across the income distribution, while the additional cut in non-essential spending that generates the shift toward essentials is concentrated among middle-income households. Low- and high-income households show no sentiment differential.

In closing, our sample captures the tariff response through 2025. Although various reports and commentary suggest retail pass-through peaked around the fall of 2025 (Hacıoğlu Hoke et al., 2026; Minton et al., 2026), pass-through may not be complete, e.g. see Flaaen et al. (2020, 2025) on the pass-through timeline in some specific cases of tariffs. More retail pass-through from the 2025 tariffs may therefore still be in train beyond our window. This does not weaken our central point. The spending and quantity responses we document arose in an environment of exceptionally high policy uncertainty and were driven in large part by forward-looking pessimism rather than by realized prices alone. Isolating this belief-driven channel, and separating it from the income effect of higher prices, is where the novelty and the policy relevance of our results lie, and it is a channel that operates whether or not the price pass-through is yet complete.

References

Acosta, M. and L. Cox (2024). The Regressive Nature of the U.S. Tariff Code: Origins and Implications. Manuscript.

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. February 12.

Amiti, M., S. Heise, and D. E. Weinstein (2026, July). The Anatomy of Tariff Pass-through into Consumer Prices. Working Paper 35561, National Bureau of Economic Research.

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). Cost of Living Inequality During the Great Recession. *Journal of the European Economic Association 19*(2), 913–952.

Arora, A. and M. Dell (2024, August). LinkTransformer: A Unified Package for Record Linkage with Transformer Language Models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, Bangkok, Thailand, pp. 221–231. Association for Computational Linguistics.

Autor, D., D. Dorn, and G. H. Hanson (2021). On the Persistence of the China Shock. *Brookings Papers on Economic Activity 2021*(2), 381–447.

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, Federal Reserve Bank of Richmond.

Bai, L., T. S. Jaccard, and S. Stumpner (2025). Tariff Pass-through into Retail Prices: Measurement, Replacement, and Shrinkflation. Working paper. SSRN Working Paper No. 5985094.

Bai, L. and S. Stumpner (2019). Estimating US Consumer Gains from Chinese Imports. *American Economic Review: Insights 1*(2), 209–224.

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., P. Hull, and X. Jaravel (2025). A Practical Guide to Shift-Share Instruments. *Journal of Economic Perspectives 39*(1), 181–204.

Borusyak, K. and X. Jaravel (2021). The Distributional Effects of Trade: Theory and Evidence from the United States. Working Paper 28957, National Bureau of Economic Research.

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, February 16.

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 (2021). 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. M. Vazquez (2025, October). Tracking the Short-Run Price Impact of U.S. 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.

Coibion, O. and Y. Gorodnichenko (2026). *Expectations Matter: The New Causal Macroeconomics of Surveys and Experiments*. Princeton, NJ: Princeton University Press.

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. Working paper, University of Akron.

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. On the Economy Blog, Federal Reserve Bank of St. Louis. October 16.

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. *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, June). Who Pays for the Tariffs and Why? A Tale of Two Countries. CESifo Working Paper 10497, CESifo.

Flaaen, A., A. Hortaçsu, and F. Tintelnot (2020). The Production Relocation and Price Effects of US Trade Policy: The Case of Washing Machines. *American Economic Review 110*(7), 2103–2127.

Flaaen, A. B., A. Hortaçsu, F. Tintelnot, N. Urdaneta, and D. Xu (2025). Who Pays for Tariffs Along the Supply Chain? Evidence from European Wine Tariffs. Working Paper 34392, National Bureau of Economic Research.

Freund, C. (2026). Tariff incidence and market power: Evidence from the 2025 us tariff shock. CEPR Discussion Paper 21798, Centre for Economic Policy Research (CEPR), Paris and London.

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., 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, October 11.

Hacıoğlu Hoke, S., L. Feler, and J. Chylak (2025, July). Lost in Aggregation: Geographic Mismeasurement of Income and Spending. Finance and Economics Discussion Series 2025-050, Board of Governors of the Federal Reserve System.

Hacıoğlu Hoke, S., S. Malladi, and L. Feler (2026, March). The Slow Climb: How Tariffs Gradually Raised Retail Prices in 2025. FEDS Notes. Washington: Board of Governors of the Federal Reserve System, March 5.

Hausman, J. A. (1981). Exact Consumer's Surplus and Deadweight Loss. *American Economic Review 71*(4), 662–676.

Hinz, J., A. Lohmann, H. Mahlkow, and A. Vorwig (2026, January). America's Own Goal: Who Pays the Tariffs? Kiel Policy Brief 201, Kiel Institute for the World Economy.

Jappelli, T. and L. Pistaferri (2014). Fiscal Policy and MPC Heterogeneity. *American Economic Journal: Macroeconomics 6*(4), 107–136.

Jaravel, X. and E. Sager (2019). What Are the Price Effects of Trade? Evidence from the U.S. and Implications for Quantitative Trade Models. Finance and Economics Discussion Series 2019-068, Board of Governors of the Federal Reserve System.

Johnson, D. S., J. A. Parker, and N. S. Souleles (2006). Household Expenditure and the Income Tax Rebates of 2001. *American Economic Review 96*(5), 1589–1610.

Ma, H., L. Macedoni, J. Ning, and M. Xu (2025). Tariffs Tax the Poor More: Evidence from Household Consumption During the US-China Trade War. CESifo Working Paper 11610, CESifo.

Minton, R., M. Ray, and M. Somale (2026, April). Detecting Tariff Effects on Consumer Prices in Real Time – Part II. FEDS Notes. Washington: Board of Governors of the Federal Reserve System, April 8.

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, May 9.

Parker, J. A., N. S. Souleles, D. S. Johnson, and R. McClelland (2013). Consumer Spending and the Economic Stimulus Payments of 2008. *American Economic Review 103*(6), 2530–2553.

Pierce, J. R. and P. K. Schott (2016). The Surprisingly Swift Decline of U.S. 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. (2026a, May). Cheapflation Cycles. Working Paper 35235, National Bureau of Economic Research.

Sangani, K. (2026b, May). Complete Pass-Through in Levels. *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. (2026). Decomposing Supply- and Demand-Driven Inflation. *Journal of Money, Credit and Banking 58*(2), 365–388.

Topalova, P. (2010). Factor Immobility and Regional Impacts of Trade Liberalization: Evidence on Poverty from India. *American Economic Journal: Applied Economics 2*(4), 1–41.

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. Working Paper 26353, National Bureau of Economic Research.

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.

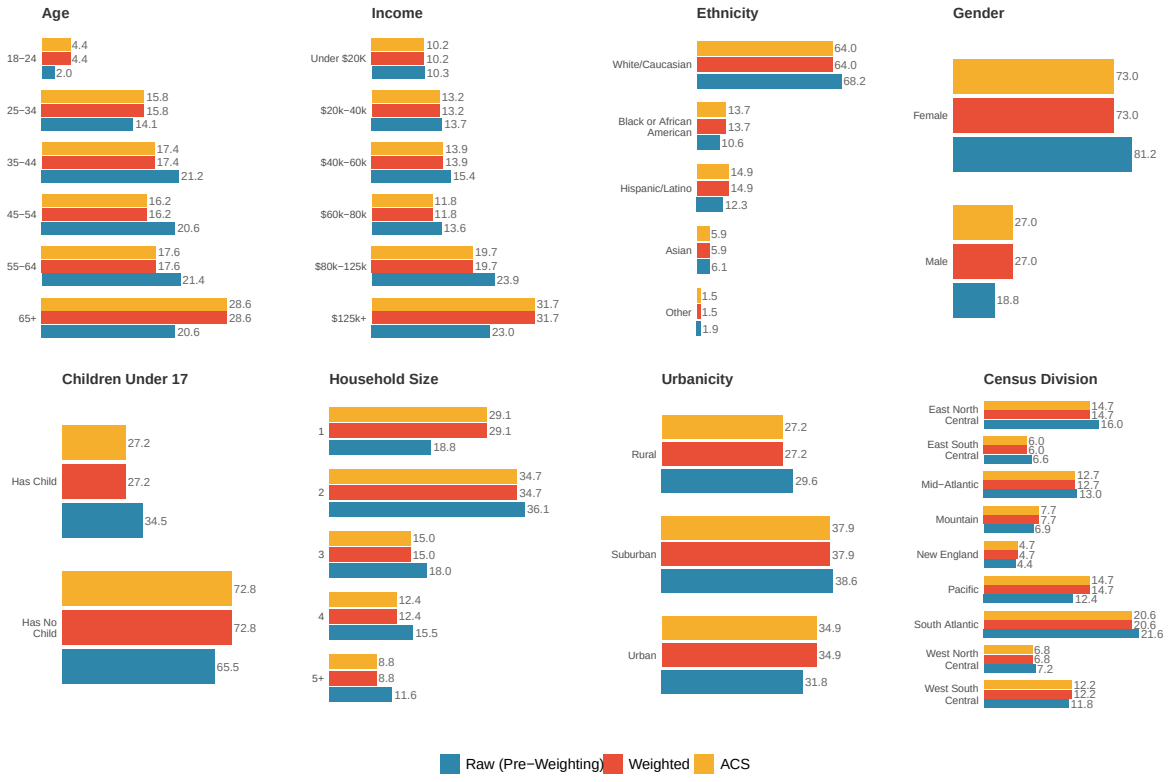
Our analysis uses Numerator's Static Panel households, a subset of the broader panel selected for consistent reporting behavior. Static Panelists must meet the following criteria: (1) maintain active status for a minimum of 12 consecutive months, (2) submit purchases consistently across the period with 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 simul-

taneously: 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



Notes: 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.

Note that the 20% male share in Numerator's panel refers to the gender of the app owner, who is treated as the household's Head of Household rather than as the sole shopper. Numerator deliberately sets female/male recruitment targets of 56/44 for single-member households and 80/20 for households of two or more, reflecting that grocery/shopping-trip logging skews female. The apparent underrepresentation of males is therefore a feature of the panel-registration unit (households represented by their app owner), not of who actually shops: because panelists report who was present on each trip, as of March 2025 roughly 39% of all trips include a male (25.2% male-only plus 13.8% male-and-female), so male shopping activity is far better captured than the 20% app-owner figure suggests.

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 struc-

ture 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 shows 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

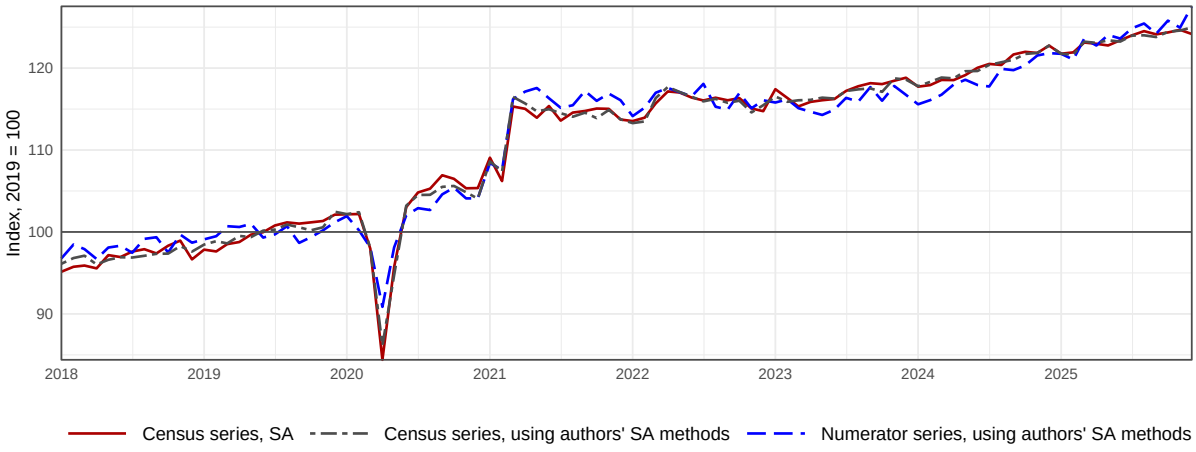
TABLE A.1: Descriptive Statistics: Full Sample vs. Survey Sample

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

Notes: The income rows assign all 126,448 households to their reported bracket. The income-group regressions and the welfare estimates instead use a three-way grouping that leaves 1,615 households unclassified, the large majority reporting no income and a small remainder in the top bracket, so those analyses rest on 124,833 households.

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.

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

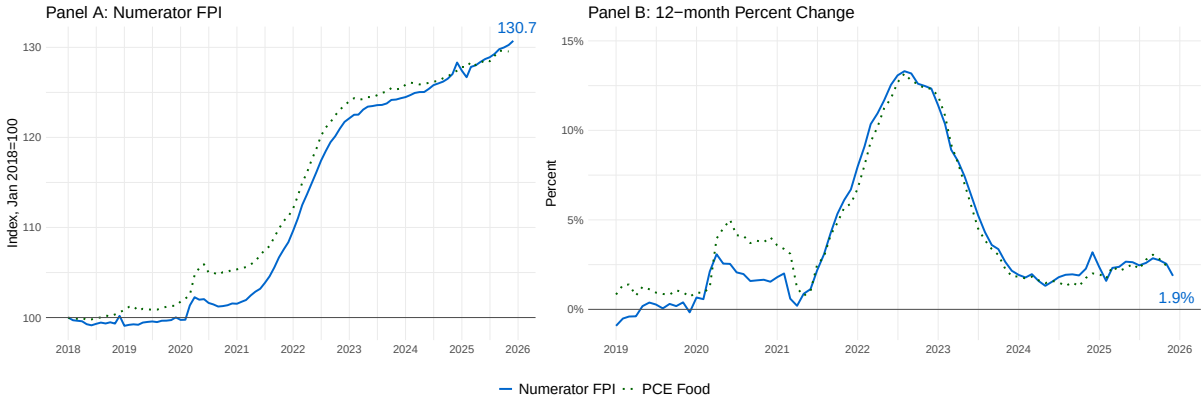


Notes: The red solid line is the Census Bureau's seasonally-adjusted retail sales series excluding motor vehicles and parts. The blue dashed line is the retail spending measure we construct from the Numerator data, which we seasonally adjust. The gray dash-dotted line is the non-seasonally adjusted Census 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, 2026; 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



Notes: 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.

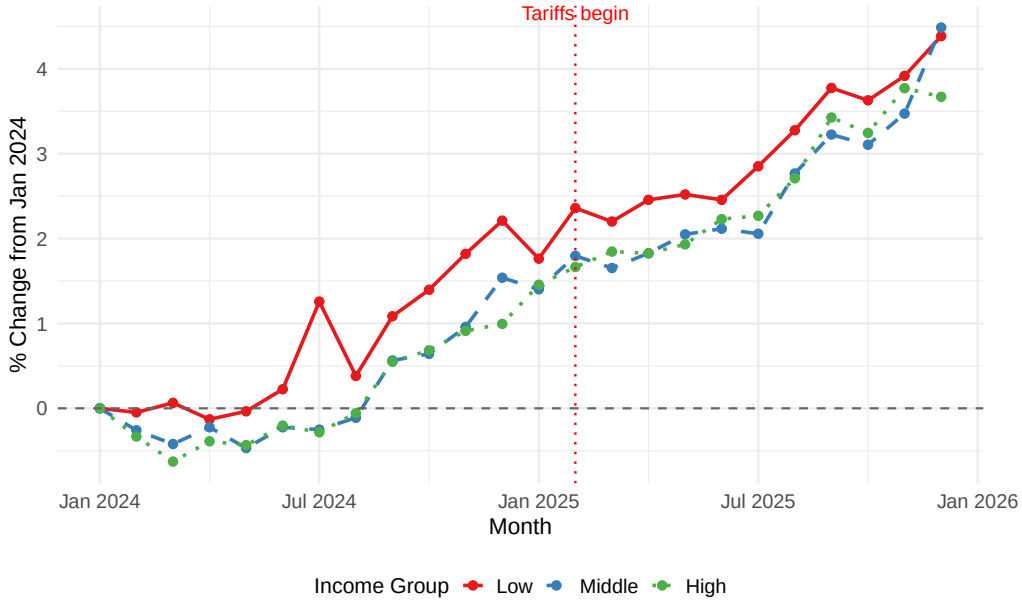
TABLE A.2: Correlation between inflation measures and Numerator FPI

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

Notes: 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 Country-of-Origin Validation of the Tariff Exposure Measure

In the main analysis, we use two tariff measures. For HS4 code h, the realized effective tariff rate τ_h is total duties divided by total import value, pooled across all 10-digit codes within h and across all source countries, as in Equation 1, and the tariff exposure

multiplies that rate by import penetration, $\text{TariffExposure}_h = \tau_h \times \text{ImportPenetration}_h$, as in Equation 3. Because the rate pools duties and imports over all source countries of a code, it replaces the unobserved country of origin of each purchased good with an implicit import-weighted average across those countries. This appendix assesses that approximation directly, by comparing it against tariff measures constructed for the subset of products whose true country of origin we observe.

To recover the country of origin of individual products, we combine several external reference sources that report it at the Universal Product Code (UPC) level. The first is Label Insight, a dataset of consumer packaged goods sold in the U.S. that links detailed product information, including country-of-origin, to UPCs. The snapshot of Label Insight we have access to is from 2017 so it is only helpful for UPC codes that are still active. We supplement it with a few crowd-sourced product databases, namely Open Food Facts, Open Beauty Facts, Open Pet Food Facts, and Open Product Facts (in short Open Facts), all of which also report country of origin for some barcoded products.

Our goal is to attach a country of origin to as many Numerator products as we can. We do this in two steps. First, we find direct matches between UPCs reported in Numerator data and a reference source, either Label Insight or Open Facts. These matches are unambiguous, and this step alone connects 59,685 Numerator products to country of origin.

The second step handles the products whose barcode we cannot find in the reference data. For these we rely instead on the text description of the product, for example "Heinz Tomato Ketchup 32 oz," and search for the description in the reference data that refers to the same product. The difficulty is that the same item is often described in slightly different words across datasets, so a literal, character-by-character comparison of the two strings misses obvious matches. We therefore use a method that compares descriptions by meaning rather than by exact spelling. A sentence transformer model embeds each description as a point in a high-dimensional space, positioned so that descriptions with similar meanings sit close together and unrelated ones sit far apart, whether or not they share the same words. This is an embedding model that maps each description to a fixed vector, not a generative large language model. The distance between two descriptions in this space is measured by their cosine similarity, a standard closeness score that is largest

when two descriptions are near-identical in meaning. This kind of meaning-based, or "fuzzy," matching substantially outperforms literal string comparison (Arora and Dell, 2024).

For each Numerator description we take its ten closest reference descriptions by this score as candidate matches. To guard against spurious pairings of unrelated products that merely sound alike, we then keep only candidates made by the same company as the Numerator product, which we identify from the manufacturer code embedded in the barcode (the GS1 company prefix). Requiring a candidate to be both close in meaning and from the same manufacturer sharply reduces false matches. Among the surviving candidates we keep the single closest one as the match. We implement this procedure with the `LinkTransformer` package of Arora and Dell (2024).

Matching on meaning rather than on the barcode does more than absorb differences in how the same product is described. It also lets us link products that never share a UPC. One might reasonably ask why two records for the same product would not carry the same barcode. UPCs are highly specific and change over time, so even minor variants of a product are assigned distinct codes. "Heinz Tomato Ketchup 32 oz." and "Heinz Tomato Ketchup 16 oz." differ only in size yet have different UPCs. Because such close variants are likely manufactured in the same place, meaning-based matching lets us assign a country of origin to a Numerator product from a near-identical reference variant even without an exact barcode match, which expands the set of origin-known observations well beyond exact-barcode overlap.

A UPC sometimes appears in both the Open Facts data and Label Insight, and in the rare cases where two sources report a different origin for the same code, which happens for well under one percent of overlapping codes, we keep the Label Insight claim. Matching the final UPC-to-country-of-origin classifications with Numerator products that carry UPCs gives us country of origin for roughly 230,000 unique products spanning 81 origin countries. These products are concentrated in grocery, health and beauty, pet, and household goods.

For each origin-known product we observe both its country of origin and, through the household-to-trade crosswalk described in Section 2, its HS4 code. Knowing the true origin removes the country-composition approximation entirely: rather than weighting

across all potential source countries, we assign each product the tariff rate that actually applies to its origin country, $\tau_h^{C_{\text{actual}}}$, using the same HS4 × country × month realized effective tariff rates constructed from Census import microdata for the main analysis. We then aggregate to the HS4 level as the spending-weighted average across origin-known products within each code,

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

where $\omega_h^C = \left(\sum_{j\in(h,C)} r_j\right)\Big/\left(\sum_{j\in h} r_j\right)$ is the spending share of origin country C within HS4 code h among origin-known products, r_j is total post-period spending on product j, and ImportPen_h is the same fixed 2022 import penetration ratio used in the main analysis. The actual exposure applies import penetration to the origin-specific rate, exactly paralleling the proxy in Equation 3, so that the two series differ only in how the country dimension is resolved.

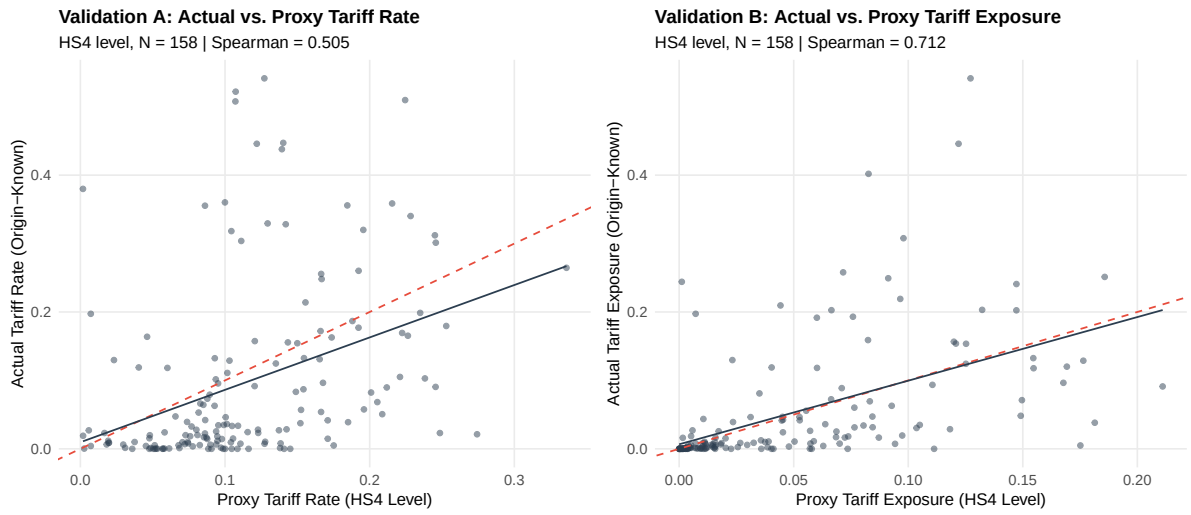
This exercise removes one of the two unknowns embedded in the proxy. Once country of origin is observed, the only remaining source of approximation is the mapping from a product to its HS code, that is, whether the goods households actually buy within a category face the tariffs we assign to that category. The close correspondence documented below therefore validates both the country-composition approximation in the proxy and the accuracy of the HS4 crosswalk for these products.

Figure B.1 plots the actual against the proxy measures across the HS4 codes with origin-known products, using the full merged sample of Label Insight and the Open Facts sources. Panel A shows tariff rates and Panel B shows tariff exposure. We summarize each panel with the Spearman rank correlation between the actual and proxy measures, which captures how well the proxy preserves the ordering of HS4 codes from least to most tariffed and is less sensitive to a handful of high-tariff outliers than the Pearson correlation. The proxy tracks the origin-based measure well: the rank correlation is 0.71 for tariff exposure and 0.51 for the tariff rate. That exposure is the more closely tracked of the two is reassuring, because tariff exposure is our benchmark treatment. The points cluster around the 45-degree line with no systematic bias, so the proxy is close to unbiased for the origin-based measure. These 158 codes account for about 99 percent of panel

spending and span the full range of post-period tariff exposure, so the validation covers the spending and the treatment variation that identify our estimates. Grocery makes up about a third of the codes, and among the non-grocery codes alone the exposure rank correlation holds at 0.53, so the correspondence is a broad-based property of the origin-known sample. To the extent that any residual measurement error in the proxy attenuates coefficients, the baseline estimates reported in Section 3 are conservative.

The figure pools Label Insight with the Open Facts sources to maximize coverage, but the two differ in how origin is recorded: Label Insight reports a verified country of origin, whereas the Open Facts entries are crowd-sourced and in principle noisier. As a check, we repeat the validation on the Label-Insight-only subsample, for which origin is known with the most confidence, and the results are very similar. We also repeat the entire exercise at the 6-digit (HS6) level and obtain correlations at least as strong, indicating that the validation is not an artifact of the HS4 aggregation.

FIGURE B.1: Country-of-Origin Validation: Actual vs. Proxy Tariff Measures

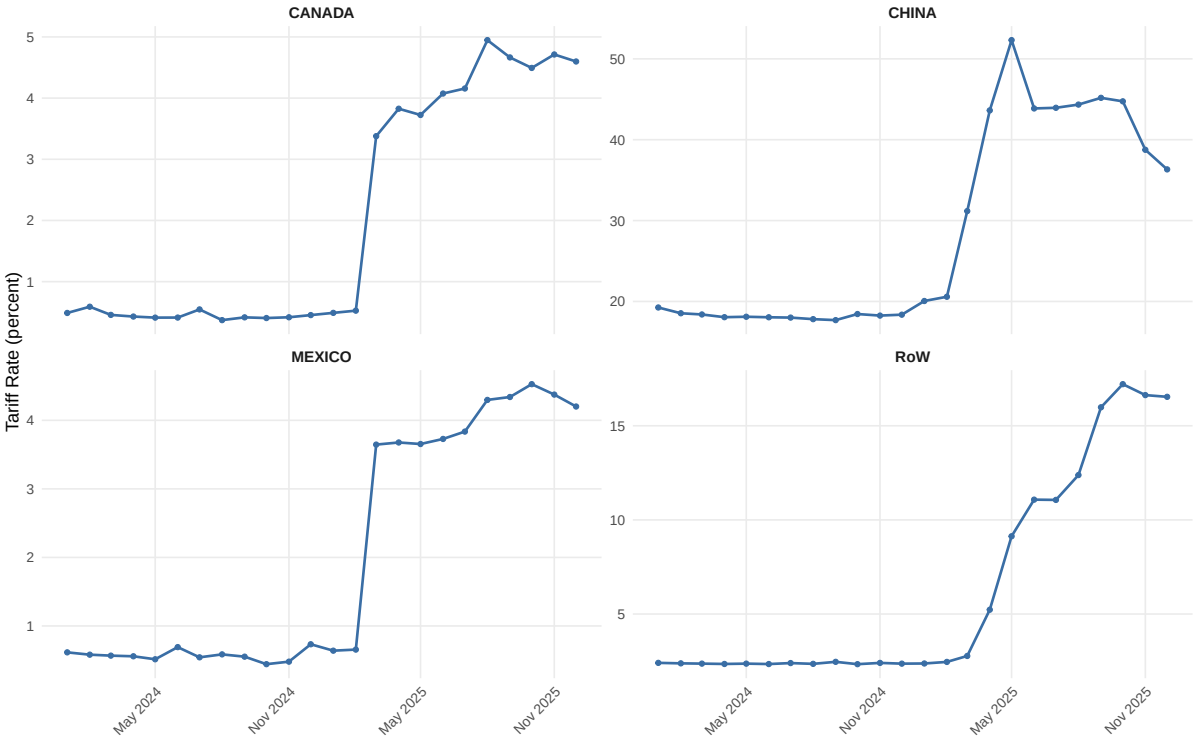


Notes: Each point is one HS4 code ($N = 158$), using all products with an observed country of origin from the merged Label Insight and Open Facts sources. Panel A shows tariff rates: the x-axis is the proxy tariff rate (duties over total imports for each HS4 code, from trade data); the y-axis is the actual tariff rate built from products with verified origin using spending-weighted origin shares (Equation B.1). Panel B shows tariff exposure: the x-axis multiplies the proxy rate by import penetration (Equation 3); the y-axis multiplies the actual rate by the same import penetration (Equation B.1). The dashed line is the 45-degree line and the solid line is the OLS fit. The subtitle of each panel reports the Spearman rank correlation between the actual and proxy measures. Panel A: Spearman = 0.51. Panel B: Spearman = 0.71.

C Tariff Exposure Measure by Sector, Country, and Essentiality

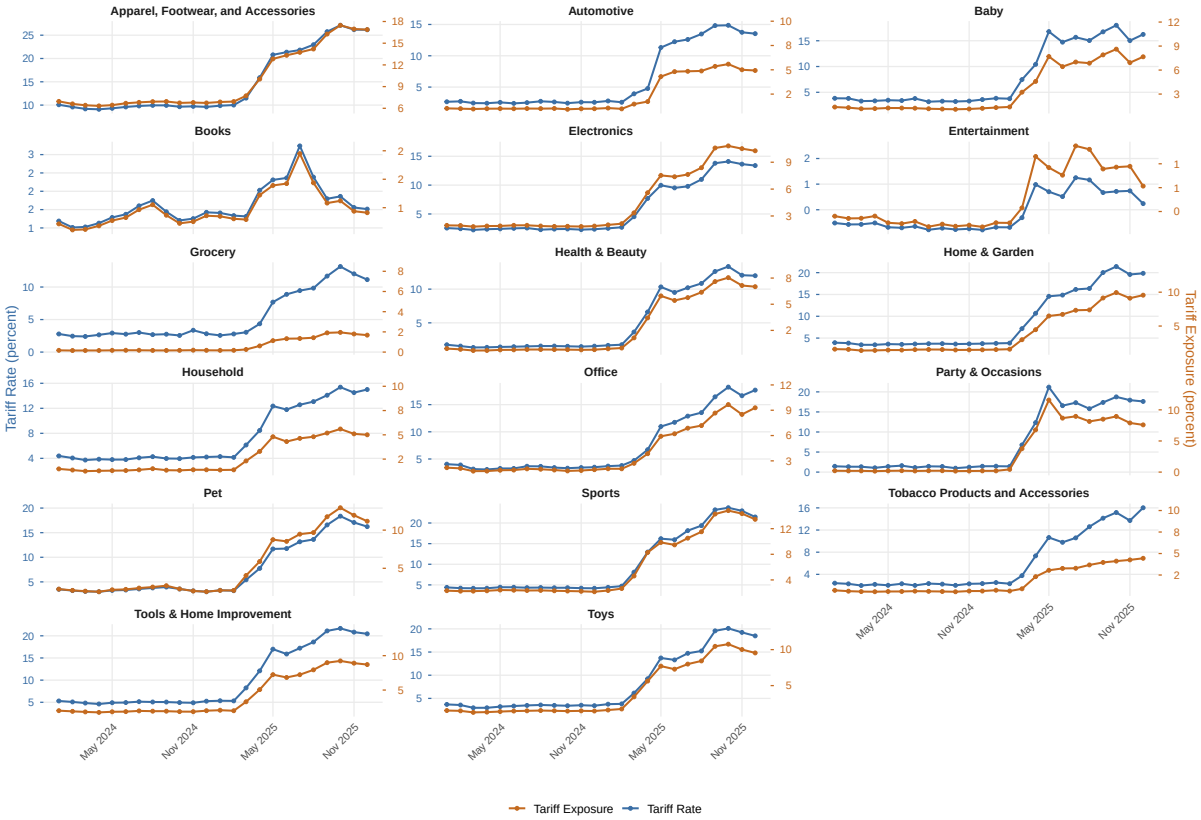
Figure C.1 shows the different tariff rates for each country from January 2024 to December 2025. Figure C.2 shows the different tariff rates and exposures each sector. Figure C.3 presents the tariff rates and exposures by essentiality as defined in Section 3.

FIGURE C.1: Average Tariff Rate by Country



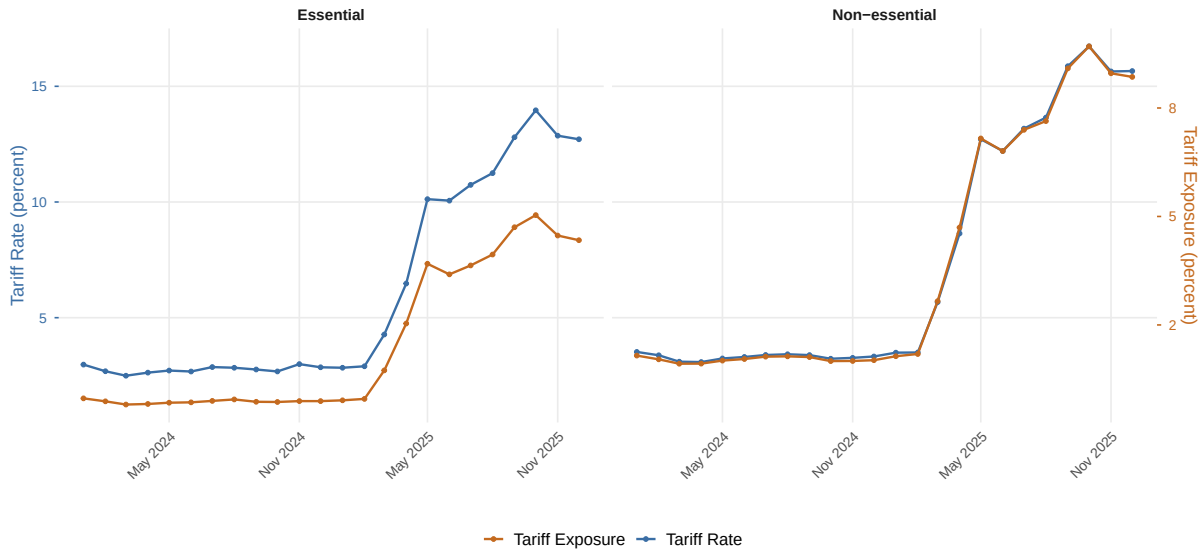
Notes: This figure shows the monthly average tariff rate over the sample period by country. For each year-month, we calculate the mean tariff rate across all observations. The tariff rate is measured in percent.

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



Notes: 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



Notes: 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 and so has low import penetration, as shown in Figure C.2. It also features intense retail competition on staple goods and exhibits relatively inelastic demand. We exclude it to verify that the aggregate results reflect the high-exposure categories where the tariffs hit rather than being mechanically shaped by this large, low-exposure sector.

Results excluding grocery are qualitatively and quantitatively similar to the full sample. Price pass-through edges down only slightly under both treatments. Spending and quantity responses remain large and significant, confirming that the aggregate findings are not mechanically driven by grocery's dominant expenditure share.

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

	Price		Spending	Quantity
	Detrended	Not Detrended		
	(1)	(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. Detrending includes panel-specific linear time trends in the price equation only, so the spending and quantity equations are estimated without 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. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

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.

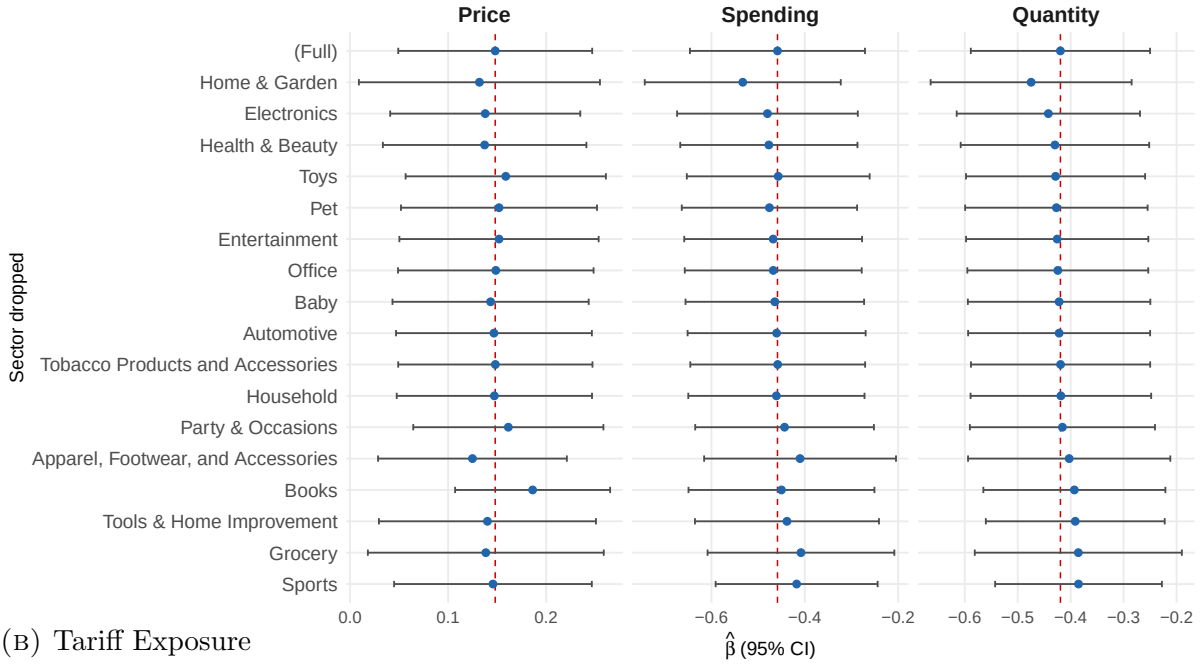
The quantity coefficient moves in a range close to the full-sample estimates, and the spending coefficient exhibits comparable stability. The largest movers differ across the two treatments. Under the tariff rate the exclusion of Home & Garden raises $|\beta|$ by roughly thirteen percent on quantity and sixteen percent on spending, reflecting the sector's combination of high tariff exposure and substantial sample weight, and no other sector shifts the estimate by more than about eleven percent. Under the tariff exposure measure the largest mover is Grocery, whose removal attenuates $|\beta|$ by about seventeen percent. 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 $|\beta|$ on quantity by about eight percent under the tariff rate and by about seventeen percent under the tariff exposure measure, 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.

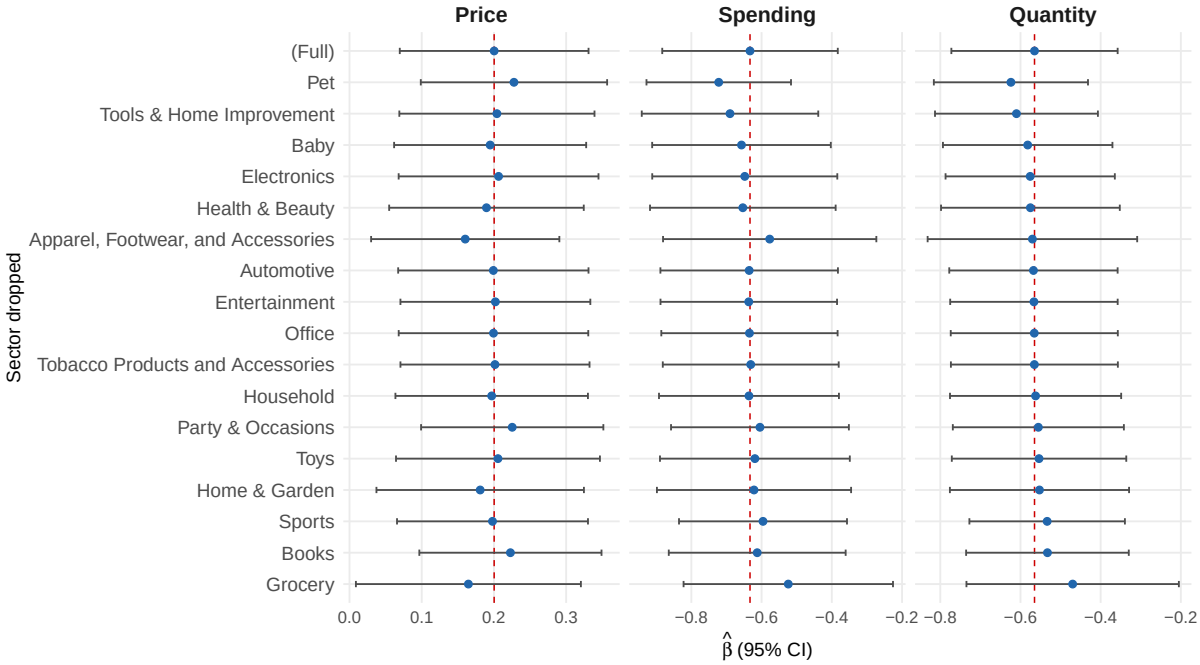
Price pass-through exhibits a similar degree of stability, with estimates ranging narrowly around the full-sample value and no sector's removal altering the qualitative conclusion of incomplete but economically meaningful pass-through.

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

(A) Tariff Rate



(B) Tariff Exposure



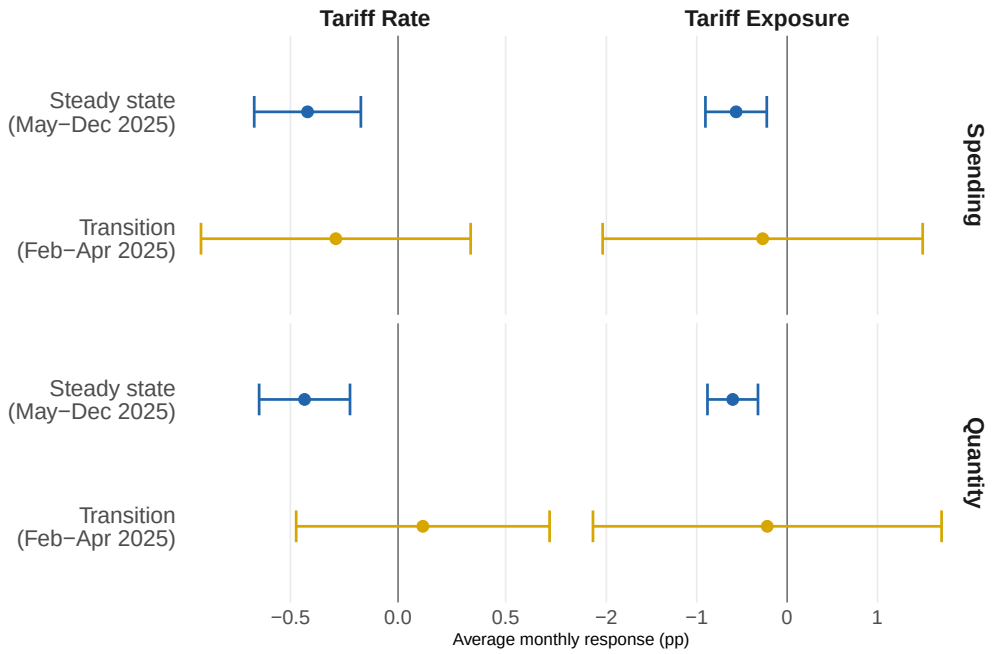
Notes: Each row reports the benchmark tariff coefficient β from Equation (4) when the indicated sector is excluded from the estimation sample. The row labeled (Full) reports the full-sample estimate, also indicated by the red dashed vertical line in each panel. Whiskers denote 95% confidence intervals 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).

D.3 Evidence Against Stockpiling

We address a natural alternative explanation for the observed contractions in our benchmark results presented in Section 3.2: 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. This pull-forward is the intertemporal-substitution response to expected tariff increases emphasized by Candia et al. (2026), and the tests below assess whether it can account for the contraction we document. 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 D.2 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 for spending and quantity alike, 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 D.2: 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 13). A further timing test confirms no front-loading signature in the early post-period (Appendix G.7, Table G.6).

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 an-

ticipatory buying. The wedge between the modest price pass-through and the much larger spending contraction that motivates the welfare and survey analyses in Sections 4 and 5 therefore requires an explanation beyond intertemporal substitution.

E Robustness to HS6 Product Granularity and the 2018–2019 Episode

E.1 HS6 Classification and Results

Our baseline maps household expenditure categories to four-digit HS4 codes, but the AI-assisted crosswalk of Section 2 often resolves product categories down to the six-digit HS level. A natural question is whether the results survive at this finer granularity. We rebuild the analysis at HS6 in this section, though we flag at the outset that HS6 is not our preferred level. Category-to-code classification error rises as the mapping grows more granular, since many household categories do not map cleanly to a single HS6 subheading, so aggregating to HS4 is the more conservative and defensible choice for the headline results. A closed version of Anthropic Opus 4.8 was used to create the HS6 crosswalk.

We rebuild the entire pipeline at the six-digit level. The shipments aggregation, the tariff-exposure measure of Equation (3), the Fisher price index, and the panel underlying Equations (4) and (5) are all reconstructed keyed on HS6 rather than HS4, with standard errors clustered on the HS6 code level. Table E.1 reports the HS6 analog of the benchmark aggregate results in Table 3 and the HS6 estimates line up closely with the HS4 benchmark. Price pass-through is a touch lower at HS6 under both the tariff-rate and tariff-exposure treatments, while the spending and quantity contractions are close to their HS4 counterparts. Overall the finer classification leaves the interpretation of our results unchanged.

TABLE E.1: HS6 Robustness of the Aggregate Results (All Sectors, 2024–2025)

	Price		Spending	Quantity
	Detrended	Not Detrended		
	(1)	(2)	(3)	(4)
Panel A: Tariff Rate				
Tariff Rate × Post	0.11***	0.08***	−0.37***	−0.33***
	(0.04)	(0.03)	(0.05)	(0.05)
Observations	108,334	108,334	108,334	108,334
Panel B: Tariff Exposure				
Tariff Exposure × Post	0.13***	0.08**	−0.65***	−0.55***
	(0.04)	(0.03)	(0.07)	(0.07)
Observations	108,334	108,334	108,334	108,334
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 with the entire pipeline rebuilt at the six-digit (HS6) product level rather than HS4. Dependent variables are the 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; Panel B uses the tariff exposure measure (tariff rate × import penetration). Price is reported for both detrended (column 1) and non-detrended (column 2) specifications; detrending includes panel-specific linear time trends in the price equation only, so the spending and quantity columns are common to both price specifications. All specifications include year-month and panel-unit-by-calendar-month interaction fixed effects. Standard errors clustered at the HS6 code level in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

E.2 The 2018–2019 Episode at HS6

A natural follow up question our analyses leads to is whether and how households responded to the earlier 2018–2019 tariffs. This section answers that question while also putting the HS6 classification to work on a second episode. We run the analysis at HS6 for consistency with the rest of this appendix, and the same analysis at HS4 delivers similar results which are unreported and available upon request.

The 2018–2019 episode comprises the Section 232 tariffs on steel and aluminum and the Section 301 tranches on Chinese imports. We construct a monthly household panel spanning January 2018 through December 2019 and take September 2018 as the treatment onset following Bai et al. (2025). The earlier 2018 actions, namely the March Section 232 metals tariffs and the July–August Section 301 tranches, fell largely on intermediate inputs and capital goods, whereas the $200 billion September 2018 tranche was the first

to reach a broad range of consumer goods at first order and is therefore the relevant date for household spending.

The sample is filtered from Numerator's static panelists at the time, which was a smaller sample. The final sample includes 57,040 panelists present from January 2018 to December 2019. Numerator data that is made available for research starts in 2018 so we cannot extend the pre-tariff period further back.

The category-to-HS crosswalk, the tariff-exposure measure, and the difference-in-differences and event-study specifications are identical to those in the main text, only re-anchored to the earlier tariff period at the HS6 level. We refer the reader there for the crosswalk construction (Section 2), the tariff and exposure formulas (Equations (1), (2), and (3)), and the estimating equations (Equations (4) and (5)). As in the main analysis, the tariff measures combine a fixed pre-tariff baseline with a time-varying tariff rate. Import share and import penetration are built from 2017 trade microdata (imports and exports) and the 2017 Economic Census of domestic shipments, the pre-tariff baseline that plays the role of the 2022 base year in the main text, while the tariff rate is constructed from 2018–2019 monthly import microdata as customs duties over customs value, analogous to the 2024–2025 monthly series.

Similar to the 2024–2025 episode, the 2018–2019 data show modest price pass-through but in contrast, shows no detectable response of spending or quantity to tariff exposure. The point estimate for the tariff-rate specification is 0.18 and statistically significant, but the point estimates of spending and quantity in all specifications are imprecisely estimated. The event study in Figure E.1 shows marginally significant results in some months, mainly for prices, in mid-2019, but no clean or consistent break in spending or quantity at the September 2018 onset. We read this as consistent with the more diffuse incidence of the 2018–2019 measures relative to the broad 2024–2025 tariffs and the absence of the precautionary behavior driven by tariff-policy uncertainty. Within this episode we nonetheless find that low-income households face higher pass-through than higher-income households, consistent with our main-sample results (Section 3.3) and with the findings of Ma et al. (2025).

FIGURE E.1: HS6 Event Study for the 2018–2019 Tariff Episode



Notes: The figure plots coefficients from event-study versions of Equations (4) and (5) estimated on the 2018–2019 Numerator panel rebuilt at the HS6 level, interacting the treatment variable with year-month indicators (the month before onset omitted as reference). The top row uses the tariff rate; the bottom row uses the tariff exposure measure. The left column shows the price response (log Fisher Price Index); the right column shows spending and quantity (both in logs, with spending in real terms). 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 are 95% confidence intervals from standard errors clustered at the HS6 code level. The vertical dashed line marks September 2018, the tariff onset.

F Substitution

The paper establishes the large quantity adjustments by households in response to tariffs, despite small price pass-through. This section examines how households adjust their spending, and whether these adjustment margins differ across the income distribution. We consider two margins of adjustment: an extensive margin, where households shift spending away from tariff-exposed product categories, and an intensive margin, where

households substitute toward cheaper items within the same category. As in the welfare analysis, all coefficients reported in this section are identified from cross-sectional variation across HS4 codes within the same month through year-month fixed effects: they capture how households *differentially* reallocated spending toward or away from high-exposure categories relative to low-exposure ones, not the total change in any category's share.

F.1 Cross-Category Substitution (Extensive Margin)

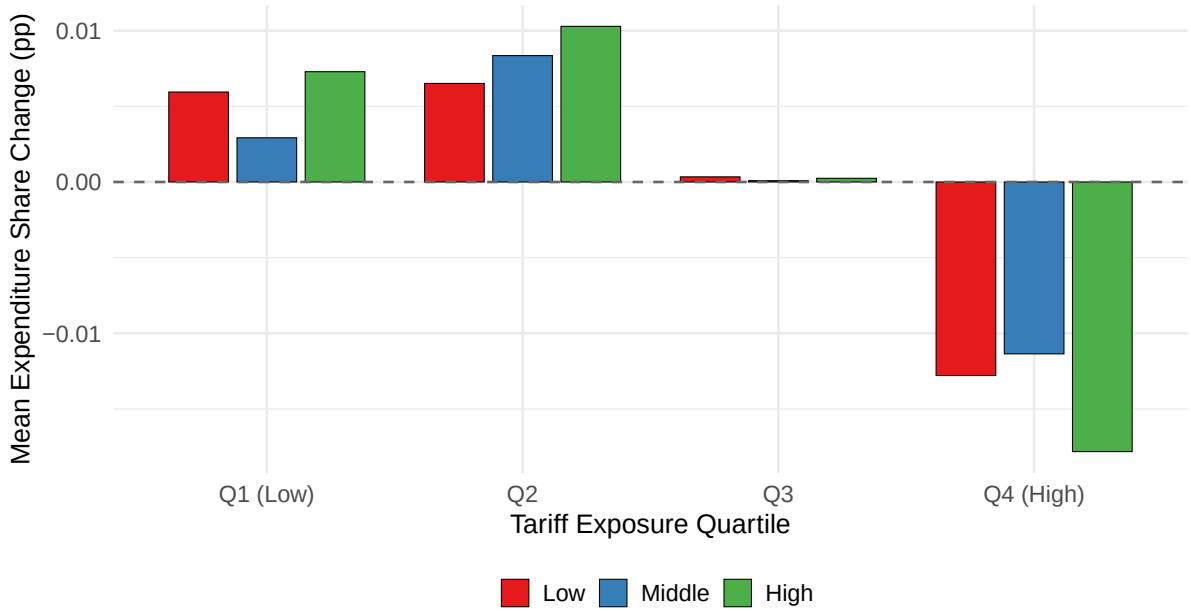
We first test whether households shifted spending away from tariff-exposed product categories, and whether the magnitude of reallocation differs by income. Figure F.1 provides descriptive evidence of cross-category substitution by comparing expenditure shares before and after the tariff shock. For each HS4 product code, we compute the change in its expenditure share between the pre- and post-tariff periods, then group products into quartiles of tariff exposure and average these share changes within each income group. Across all income groups, expenditure shares increase for low-tariff goods in the first quartile and decline for high-tariff goods in the third and fourth quartiles, consistent with households reallocating spending away from tariff-exposed categories. The most notable finding is that high-income households exhibit larger share gains in the first quartile compared to their lower-income counterparts.

To approach the question more rigorously, we estimate the following regression by defining $s_{h,g,t}$ as the expenditure share of HS4 code h in income group g's total spending in month t:

$$\ln s_{h,g,t} = \delta_{h\times g,m(t)} + \gamma_t + \beta_g \left(\text{TariffExposure}_{h,t} \times \text{Post}_t\right) + \varepsilon_{h,g,t}, \tag{F.1}$$

separately for each income group g. The HS4 × income × calendar-month fixed effects $\delta_{h\times g,m(t)}$ absorb both time-invariant differences in baseline expenditure shares across income groups for the same product and product-by-income-specific seasonal patterns; γ_t are year-month fixed effects that absorb common shocks. The measure $\text{TariffExposure}_{h,t}$ enters interacted with the post-tariff indicator, so identification comes from differential changes in spending shares across high- and low-exposure HS4 codes within income group

FIGURE F.1: Expenditure Share Changes by Tariff Exposure Quartile



Notes: The figure plots mean expenditure share changes (in percentage points) from the pre- to post-tariff period, by tariff exposure quartile and income group. Q1 corresponds to HS codes with the lowest tariff exposure; Q4 corresponds to the highest. Negative values indicate households shifted spending away from those categories.

g after February 2025. A negative coefficient $\hat{\beta}_g$ implies households shifted spending away from tariff-exposed categories.

TABLE F.1: Cross-Category Substitution by Income Group

	(1) Low	(2) Middle	(3) High
Tariff Exp. × Post	−0.71***	−0.61**	−0.82***
	(0.26)	(0.24)	(0.25)
Year-month FE	Yes	Yes	Yes
HS × Income × month	Yes	Yes	Yes
Observations	4,123	4,127	4,128
Adjusted R^2	0.99	0.99	0.99

Notes: The dependent variable is the log expenditure share of HS code h in income group g's total spending in month t. Coefficients are multiplied by 100 for readability. The specification follows Equation (F.1), estimated separately by income group. Standard errors clustered at the HS4 code level are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table F.1 presents the regression results. All three income groups exhibit statistically significant reallocation away from tariff-exposed categories. High-income households show the largest magnitude response, followed closely by low-income households, while middle-income households exhibit a somewhat smaller but still significant response. This result indicates that all income groups respond to tariffs in part by reallocating spending across product categories.

F.2 Within-Category Item Substitution (Intensive Margin)

We now ask a different question: conditional on continuing to purchase within a given category, do households substitute toward cheaper items? Because the item-level data are aggregated across households, the within-category tests in this subsection capture market-level substitution patterns rather than income-group-specific behavior.

For each item j within HS4 code h, we classify items into price tiers based on their pre-tariff (pre-January 2025) average price relative to other items in the same HS4 code. Let p_j^{pre} denote item j's mean pre-tariff price and define its within-HS4 price percentile as

$$\pi_j = \frac{\text{rank}(p_j^{\text{pre}} \mid j \in h)}{N_h}, \tag{F.2}$$

where N_h is the number of distinct items in HS code h. We restrict the sample to HS4 codes with at least three distinct items, ensuring meaningful within-category variation.

First, we assign items to tercile-based price tiers with cheap ($\pi_j \leq 1/3$), middle, and expensive ($\pi_j > 2/3$), and compute the revenue share of each tier k within HS code h in month t:

$$s_{h,k,t}^{\text{tier}} = \frac{\sum_{j \in (h,k)} r_{j,t}}{\sum_{j \in h} r_{j,t}}, \tag{F.3}$$

where $r_{j,t}$ is item j's revenue (price × quantity). We then estimate

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

with HS4 × tier × calendar-month fixed effects and year-month fixed effects. A positive $\hat{\beta}_{\text{Cheap}}$ paired with a negative $\hat{\beta}_{\text{Expensive}}$ indicates that cheap items gained revenue share while expensive items lost share within high-tariff HS4 codes after the tariff shock,

providing direct evidence of within-category substitution toward cheaper varieties.

TABLE F.2: 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 FE	Yes
Year-Month FE	Yes
HS × month	Yes
Observations	12,362
Adjusted R^2	0.96

Notes: The dependent variable is the log revenue share of each price tier within HS code × month. The specification follows Equation (F.4). Items are classified into tercile-based price tiers using pre-tariff (pre January 2025) average prices. Standard errors clustered at the HS4 code level are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table F.2 reports the results. Cheap items gained significant revenue share within tariff-exposed HS4 codes, while expensive items lost share by a nearly symmetric magnitude. Middle-tier items also gained share, though by less than cheap items. The near-symmetry between cheap-tier gains and expensive-tier losses is consistent with a zero-sum reallocation of revenue shares within HS codes: tariffs induced consumers to substitute away from expensive varieties toward cheaper alternatives within the same product category.

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

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

where $\delta_{j,m(t)}$ are item-by-calendar-month fixed effects (which absorb time-invariant item heterogeneity) and γ_t are year-month fixed effects, with $h(j)$ denoting the HS4 code

containing item j. The coefficient $\hat{\beta}_2$ traces how the tariff-induced share reallocation varies across the within-HS price distribution. A negative $\hat{\beta}_2$ implies that, within high-tariff categories, cheaper items (low π_j) gained market share at the expense of more expensive items.

TABLE F.3: 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 FE	Yes
Year-month FE	Yes
Item × month FE	Yes
Observations	63,573,069
Adjusted R^2	0.86

Notes: The dependent variable is the log revenue share of item j within its HS code in month t. The specification follows Equation (F.5). Price percentile π_j is defined in Equation (F.2) and ranges from 0 (cheapest) to 1 (most expensive). All specifications include item-by-calendar-month fixed effects. Standard errors clustered at the HS4 code level are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table F.3 reports the results from over 63.5 million item-month observations. The coefficient on Tariff Exp. × Post captures the baseline tariff-induced share gain for the cheapest items ($\pi_j = 0$). The interaction coefficient shows that this advantage diminishes linearly with price. For each step up in the within-category price distribution, the tariff-induced gain in market share falls. Since shares within an HS4 code must sum to 100 percent, this wedge represents a reallocation from expensive to cheap items within tariff-exposed categories, relative to the control group, consistent with consumers trading down to save money.[26]

The within-category (intensive margin) substitution documented above is estimated

[26] A notable feature of the continuous specification is that the main effect implies positive share gains across the entire price distribution, including the most expensive items. Since revenue shares within an HS4 code must sum to one, this pattern cannot hold in a balanced panel and likely reflects differential item attrition: within tariff-exposed categories, some items, mostly disproportionately expensive ones, exit the sample in the post-period, mechanically raising the shares of surviving items. The interaction term is not affected by this compositional shift and correctly identifies the relative reallocation from expensive to cheap items, consistent with the tercile-based results in Table F.2, where expensive items explicitly lose share.

at the market level, pooling across all households. To examine whether the intensity of within-category trading down varies by income, we exploit the wedge between unit values and fixed-basket price indices. The Fisher Price Indices used in Section 3 hold the item basket fixed, while unit values reflect actual purchasing behavior including item switching within HS codes. The wedge between the two measures therefore captures within-category quality adjustment. For each HS4 code h, income group g, and month t, we compute

$$Q_{h,g,t} = \ln\left(\frac{\text{UV}_{h,g,t}}{\text{UV}_{h,g,0}}\right) - \ln\left(\frac{\text{FPI}_{h,g,t}}{\text{FPI}_{h,g,0}}\right), \tag{F.6}$$

where $\text{UV}_{h,g,t}$ is the unit value (total expenditure divided by total quantity) and $\text{FPI}_{h,g,t}$ is the income-specific Fisher Price Index, both indexed to the base period $t = 0$ (January 2024). A negative quality wedge indicates that unit values grew less than the fixed-basket price index, meaning that households switched to cheaper items within the category. We estimate

$$Q_{h,g,t} = \delta_{h\times g,m(t)} + \gamma_t + \sum_{g'} \beta_{g'} \left(\text{TariffExposure}_{h,t} \times \text{Post}_t \times \mathbf{1}[G = g']\right) + \varepsilon_{h,g,t}, \tag{F.7}$$

where $Q_{h,g,t}$ is the quality wedge measuring within-category item substitution defined in Equation (F.6); $\delta_{h\times g,m(t)}$ denotes HS4 × income × calendar-month fixed effects that absorb seasonal purchasing patterns specific to each product-income cell as well as time-invariant differences across products and income groups; γ_t are year-month fixed effects; $\text{TariffExposure}_{h,t}$ is the time-varying tariff exposure for HS4 code h; Post_t is an indicator for months from February 2025 onward; and $\mathbf{1}[G = g']$ indicates income group $g' \in \{\text{Low, Middle, High}\}$. The coefficients $\beta_{g'}$ capture the differential quality downgrading response to tariffs across income groups. Negative values indicate that higher tariff exposure led to greater within-category trading down to cheaper items.

Table F.4 reports the regression results. All three income groups exhibit negative coefficients, with low-income households showing the largest and only statistically significant response. The low-income coefficient is approximately twice the magnitude of middle- and high-income households, neither of which is statistically distinguishable from zero. This pattern suggests that low-income households rely more heavily on within-category trading down as their primary adjustment mechanism, while higher-income households

TABLE F.4: Quality Wedge Regression Results by Income Group

	Quality Wedge ($Q_{h,g,t}$)
Tariff Exposure × Post × Low Income	−0.53**
	(0.23)
Tariff Exposure × Post × Middle Income	−0.23
	(0.25)
Tariff Exposure × Post × High Income	−0.25
	(0.16)
HS × Income × month FE	Yes
Year-Month FE	Yes
Observations	12,313
Adjusted R^2	0.57

Notes: The dependent variable is the quality wedge defined in Equation (F.6). The specification follows Equation (F.7). Standard errors clustered at the HS4 code level are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

have greater flexibility to reallocate across categories as documented in Table F.1. The significant market-level results in Tables F.2 and F.3, combined with the income-specific quality wedge patterns, provide robust evidence that quality downgrading is an active response to tariffs, with particularly strong effects among low-income households.

Comparing the cross-category results in Table F.1 with the within-category quality wedge in Table F.4 reveals distinct adjustment margins across the income distribution. All income groups show significant cross-category reallocation (extensive margin), but the relative importance of extensive versus intensive margin adjustments differs systematically by income. Low-income households exhibit both significant cross-category reallocation and the strongest quality downgrading response, suggesting they rely heavily on intensive margin adjustments even while also reallocating across categories. High-income households show the largest cross-category reallocation but more modest quality downgrading, indicating a more balanced use of both margins with relatively greater reliance on the extensive margin. Middle-income households exhibit intermediate patterns. This is consistent with lower-income households facing more binding budget constraints that necessitate intensive trading down within essential categories, while higher-income households have greater flexibility to exit discretionary categories.

These estimates answer a different question from the income patterns in Section 5. The

coefficients here are total responses to tariff exposure, averaged across all households in an income group, so they describe how each group absorbs the shock it faces. The survey coefficients are differentials between more and less pessimistic households within a group, so they describe how much of the response scales with beliefs. The two can diverge. Low-income households absorb the shock mainly by trading down, as the quality wedge shows, while their response holds steady across levels of pessimism. Middle-income households make a modest discretionary cut in total, and that cut is the margin that moves with pessimism.

G Survey Appendix

This appendix collects the survey materials and supporting tables for Section 5, in order of appearance. Appendix G.1 reproduces the full survey questionnaire. Appendix G.2 gives the Q8-to-sector crosswalk. Appendix G.3 adds each household's own tariff-induced income shock to the reallocation specification and shows that the sentiment response survives it. Appendix G.4 reports the household-specific-exposure version of the low-tariff test in Section 5.3. Appendix G.5 provides the Q7 reduction-intent horserace. Appendix G.6 breaks the Q8 worry specification down by income group. Appendix G.7 reports the stockpiling timing test for the early versus late post-period. Appendix G.8 breaks stated reduction intent down by income group.

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.

G.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

G.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 (22) equals one if household i flagged any Q8 category whose mapped sector contains HS4 code h.

TABLE G.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	Home & Garden
Building materials	Tools & Home Improvement
Electricity or natural gas	(no mapped sector)
Pet supplies	Pet
Furniture	Home & Garden
Toys & games	Toys

G.3 Reallocation and the Real-Income Effect

The specification that separates the belief channel from the household's tariff-induced income shock is presented in Section 5.2, together with the baseline estimates in Equation (19) and Table 9. This appendix reports the supporting evidence referenced there. Throughout, the exposure measure HHExp_i is the same household-level measure used in Appendix G.4.

Baseline exposure captures how heavily a household's basket was tilted toward tariffed goods, but it does not capture the price change that the household actually paid. We therefore construct a tighter measure of the realized income shock, HHRealizedDP_i, defined as the household's 2024-basket-weighted realized change in HS4-level log unit prices between the pre- and post-periods. Table G.2 re-estimates Equation (19) with the realized-price measure in place of the exposure measure, and then with both measures

entered jointly. The conclusion is unchanged and if anything sharper. Controlling for the price change that households actually experienced leaves the sentiment effect similar, and the effect survives when the realized price change and the predetermined exposure are included together.

Both measures capture the realized income shock, that is, the loss of purchasing power that had already occurred over the sample, and neither captures a household's expectation of future income losses. The evidence in Table G.2 therefore establishes that the belief-driven reallocation is not a relabeling of the realized income effect, but it does not on its own separate a precautionary response to future uncertainty from a response to an anticipated future decline in income.

TABLE G.2: Reallocation and the Realized Price Change

$\ln Y^{\text{noness}}$ (%)	Baseline	+ Realized Δp	+ Exposure and Δp
Panel A: Sent_i = PC1			
$\text{Sent}_i \cdot \text{Post}_t$	−0.0093**	−0.0100***	−0.0090**
	(0.0038)	(0.0038)	(0.0038)
$\text{HHRealizedDP}_i \cdot \text{Post}_t$		10.07***	6.27***
		(0.69)	(0.83)
$\text{HHExp}_i \cdot \text{Post}_t$			−0.0312***
			(0.0040)
Panel B: Sent_i = Q6 negative			
$\text{Sent}_i \cdot \text{Post}_t$	−0.0199***	−0.0202***	−0.0176**
	(0.0073)	(0.0072)	(0.0072)
$\text{HHRealizedDP}_i \cdot \text{Post}_t$		9.57***	5.56***
		(0.62)	(0.76)
$\text{HHExp}_i \cdot \text{Post}_t$			−0.0330***
			(0.0038)
Num. Obs. (Panel A / B)		377,845 / 451,743	

Notes: Estimates of Equation (19) on the household-month panel, with log real non-essential spending as the dependent variable. Column 1 is the sentiment-only baseline; column 2 adds the household's realized basket price change $\text{HHRealizedDP}_i \cdot \text{Post}_t$; column 3 adds both the realized price change and the predetermined exposure $\text{HHExp}_i \cdot \text{Post}_t$. Panel A uses PC1 as Sent_i; Panel B uses the Q6-negative indicator. HHRealizedDP_i is the 2024-basket-weighted realized change in HS4-level log unit prices between the pre- and post-periods. HHExp_i is expressed in percentage points as in Table 9. All specifications include household and year-month fixed effects. Standard errors clustered at the household level in parentheses. *, **, *** denote significance at the 10, 5, and 1 percent levels.

G.4 Low-Tariff Basket with Household-Specific Exposure

Section 5.3 shows that pessimism also affects spending on low-tariff goods using the within-HS4 specification. This subsection reports the companion test built on household-specific exposure rather than on the HS4 tariff rate, and reaches the same conclusion. We follow the definition in Section 5.3 and 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.

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

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

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

TABLE G.3: Precautionary Spillover Test: Sentiment × HHExp on Low-Tariff Basket

	Q6	Q9	Q10	PC1
$\text{HHExp}_i \cdot \text{Post}_t$	−0.00786***	−0.00760***	−0.00742***	−0.00777***
	(0.00095)	(0.00101)	(0.00093)	(0.00103)
$\text{Sent}_i \cdot \text{Post}_t$	0.013	0.012	0.007	0.015
	(0.012)	(0.013)	(0.012)	(0.013)
$\text{Sent}_i \cdot \text{HHExp}_i \cdot \text{Post}_t$	−0.00314	−0.00276	−0.00156	−0.00359
	(0.00264)	(0.00278)	(0.00258)	(0.00285)
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 (G.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, expressed in percentage points as in Table 9. 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. The linear-trend coefficient is positive but rounds to zero at three decimals. Standard errors clustered at the household level in parentheses. *, **, *** denote significance at the 10, 5, and 1 percent levels.

Table G.3 reports the estimates. The triple-interaction coefficient β_3 is negative-signed in every column but never statistically significant. Pessimistic households with heavily exposed baskets do not cut the low-tariff basket by any more than other exposed households do, which is the household-specific analogue of the null contrast in Table 11.

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

finding from the β_3 null above.

The sentiment main effect β_2 here is positive and imprecise, while the within-HS4 estimate in Table 11 is negative and precisely estimated. The two terms draw on different variation. This specification runs at the household-month level with a linear trend, so the sentiment term is identified from time variation within a household. The within-HS4 specification identifies it against other households buying the same HS4 in the same month, and that is the estimate Section 5.3 interprets.

G.5 Q7 Reduction Intent: Horserace Robustness

This subsection reports the horserace specification referenced in Section 5.5, 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 14. We therefore report the individual sub-item regressions in the main text and the joint specification here as a robustness check.

TABLE G.4: 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 (23) 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 14. 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 13. Standard errors clustered at the household level in parentheses. *, **, *** denote significance at the 10, 5, and 1 percent levels.

G.6 Q8 Worry by Income Group

This subsection reports the income-group breakdown of the pooled Q8 specification referenced in Section 5.4. 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 G.5: Category-Specific Worry (Q8) by Income Group

	Low	Middle	High
Panel A: Spending			
$\text{Worried}_{i,h} \cdot \text{Post}_t$	−0.010***	−0.009***	−0.011***
	(0.003)	(0.001)	(0.002)
Panel B: Quantity			
$\text{Worried}_{i,h} \cdot \text{Post}_t$	−0.001	0.002**	0.001
	(0.002)	(0.001)	(0.001)
Num. Obs.	9,969,954	25,513,184	11,898,769

Notes: The table reports estimates of the pooled specification of Equation (22) (imposing $\beta_2 = 0$), estimated separately within each income group. Panel A uses the natural log of real (PCE-deflated) household spending on HS4 code h in month t, and Panel B uses the natural log of household quantity. 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 G.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 within-HS4 worry response is uniform across the income distribution. Worried households cut spending by about 1 percent in every income group, matching the pooled estimate in Table 12, while quantities do not fall in any group. This specification pools flagged essential and non-essential categories. Since the full-sample split in Table 12 shows spending rising in flagged non-essentials, a uniform pooled cut with flat quantities points to the essential-side trade-down operating in every income group rather than being driven by any one of them. The level cut on non-essentials that drives the basket-composition reallocation toward essentials is, by contrast, concentrated among middle-income households.

G.7 Stockpiling Timing Test: Early versus Late Post-Period

This subsection reports the early-versus-late timing test referenced in Appendix D.3. 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 (23) with the tariff-exposure-by-intent interaction computed

separately for each sub-period:

$$Y_{i,h,t} = \beta^{\text{early}} \left(\text{TariffExposure}_{h,t} \cdot \text{Post}_t^{\text{early}} \cdot \text{Intent}_i \right) + \beta^{\text{late}} \left(\text{TariffExposure}_{h,t} \cdot \text{Post}_t^{\text{late}} \cdot \text{Intent}_i \right) + \alpha_{i,h} + \lambda_{h,t} + \varepsilon_{i,h,t}, \tag{G.2}$$

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 G.6: 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 (G.2) 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 G.6 reports the results. Two patterns emerge. First, the stockpile-intent quantity response is not front-loaded: the early-period coefficient is positive but not significant, and the late-period coefficient is marginally significant, with neither distinguishable from the other. 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, 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 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.

G.8 Stated Intent by Income Group

Table G.7 estimates Equation (23) 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.

TABLE G.7: 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***	−0.003	−0.045	−0.053**	−0.019	−0.015	−0.061**	−0.027
	(0.043)	(0.025)	(0.033)	(0.042)	(0.023)	(0.031)	(0.055)	(0.028)	(0.040)
Panel B: Log quantity									
$\text{TariffExposure}_{h,t} \times \text{Intent}_i$	−0.003	−0.007	0.003	0.008	0.010	−0.015			
	(0.019)	(0.012)	(0.016)	(0.019)	(0.011)	(0.015)			
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	11.9	10.0	25.5	11.9	10.0	25.5	11.9

Notes: The table estimates Equation (23) 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 (cut back on non-essentials, delay purchases, delay big-ticket purchases, or reallocate budget toward essentials, excluding 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.