Paying More, Expecting the Same: Supply Shocks, Household Consumption, and Inflation Beliefs

Leandro Sanz[†], Paul Decaire[‡], Felix von Meyerinck[§], Michael Weber[¶] and Michael Wittry[‖]

March 16, 2026

Abstract

The prices households pay shape what they buy and what they expect about inflation. But does it matter *why* prices changed? We study this using supply chain disruptions—a cleanly isolated supply-side shock to retail prices—and a novel dataset spanning 2010–2023 that links satellite-based vessel-level shipping delays to barcode-level retail prices of delayed goods, household purchase records, and inflation expectations surveys. We first establish that shipping delays lead to significant increases in retail prices: a ten percentage point jump in delayed import volume is associated with 1.4% higher prices, a result that holds both during and outside the COVID-19 pandemic. Households exposed to these delay-driven price increases reduce product purchase probability by 0.41 percentage points and total expenditure by 0.53%, indicating that households notice and respond to the higher prices. Yet despite these significant consumption responses, household inflation expectations do not adjust: exposed households are no more likely to revise their inflation beliefs upward than unexposed households. Supply-side price shocks impose real welfare costs on households but, unlike demand-driven price experiences documented in prior work, do not propagate to inflation beliefs.

JEL classification: E31, F14, D12, E21, D84

Keywords: Supply Chain Disruptions; Shipping Delays; Retail Price Pass-Through; Household Consumption; Inflation Expectations

Felix gratefully acknowledges financial support from the Swiss National Science Foundation (grant number 207668).

[†]Mendoza College of Business, University of Notre Dame, Email: lsanz@nd.edu

[‡]W. P. Carey School of Business, Arizona State University, Email: paul.decaire@asu.edu

[§]Lake Lucerne Institute and University of Zurich, Email: felix.vonmeyerinck@llui.org

[¶]Daniels School of Business, Purdue University, Email: weber366@purdue.edu

[‖] Fisher College of Business, Ohio State University, and European Corporate Governance Institute, Email: wittry.2@osu.edu

1 Introduction

How households respond to price changes—in both their spending and their beliefs—is central to understanding how shocks propagate through the economy. Prior work has established that the prices households encounter lead to updates in their inflation beliefs (D'Acunto, Malmendier, Ospina, and Weber, 2021; Cavallo, Cruces, and Perez-Truglia, 2017; Malmendier and Nagel, 2016). Whether it matters *why* those prices changed is less understood. We study this question by exploiting supply chain disruptions as a cleanly identified supply-side shock to retail prices.

Our analysis relies on a novel dataset that integrates four sources, spanning 2010–2023. The first is satellite-based Automatic Identification System (AIS) data from Spire Maritime, which continuously tracks the position of every ocean-going container vessel in the global fleet. Unlike aggregate shipping indices such as the Baltic Dry Index or the Global Supply Chain Pressure Index, the AIS data allow us to reconstruct the complete itinerary of each individual voyage—port of origin, actual arrival time, and scheduled arrival time—yielding a precise, continuous, vessel-level measure of delay for every container shipment arriving at a U.S. port. We merge these voyage records to the universe of U.S. import Bill of Lading (BoL) filings, which are administrative records collected by U.S. Customs for every incoming containerized shipment and contain the receiving firm's identity, the product's Harmonized System (HS) code, and the shipment's TEU volume. Matching vessel voyages to BoL entries allows us to assign a delay duration to each individual shipment and aggregate these to a firm$\times$product$\times$week delay exposure measure.

We then link firm identities to the Nielsen Retail Scanner (RMS) database—one of the most comprehensive retail transaction datasets in existence, covering more than 90,000 stores across all major U.S. retail chains and recording the actual transaction price and quantity sold at the barcode (UPC) level on a weekly basis—and to the Nielsen Consumer Panel (HMS), a nationally representative longitudinal panel tracking the complete purchase histories of tens of thousands of U.S. households across all product categories and retail outlets from 2010 to 2023. Crucially, HMS households also participate in a quarterly survey of one-year-ahead inflation expectations following the methodology of D'Acunto et al. (2021), enabling us to link each household's realized supply chain shock exposure directly to their forward-looking inflation beliefs. The resulting four-way merged dataset is, to our knowledge, the first to connect vessel-level shipping delays through BoL-identified importing firms to both barcode-level retail prices and household-level purchase decisions and expectations within a single unified framework.

Through the combination of these data, we provide evidence that the source of price changes does matter. We establish this through three linked empirical analyses that trace supply-side price shocks from the shipping network through retail prices to household consumption and inflation beliefs.

First, we establish that shipping delays cause significant retail price increases. We document this association in a firm×product×week panel spanning 2.2 million observations. Our key variable, *Fraction in delay*, measures the share of a firm's total in-transit import volume (in TEUs) that is delayed in a given week. We restrict the sample to regularly shipped firm×product pairs and estimate specifications with progressively richer fixed effects, from week fixed effects alone up to HS code×week and firm×week fixed effects that absorb industry-wide time trends and firm-specific shocks, respectively. Across all specifications, we find a positive and statistically significant relationship between the fraction of delayed volume and log retail price (Table 2). In our most demanding specification—which saturates the model with HS code×week, firm×week, and UPC prefix fixed effects—a ten percentage point increase in the fraction of import volume delayed is associated with a 1.4% increase in retail price ($\hat{\beta} = 0.139$, $p < 0.01$). The price effect is further amplified by the scale of the firm's import pipeline: the interaction of *Fraction in delay* with log total volume in transit is positive and significant across all specifications (Table 3), consistent with delays compressing inventory more severely when a larger share of a firm's supply is in transit. We further examine whether delays predict abnormal new shipment arrivals in the following week (Table 4) and find a small positive and significant association ($\hat{\beta} = 0.043$, $p < 0.01$), consistent with firms actively replenishing inventory in response to shortfalls. This restocking response, if anything, would attenuate estimated price effects, lending our pass-through estimates a conservative interpretation.

Delay-based shocks are conceptually distinct from the tariff and exchange rate shocks that dominate the existing pass-through literature (Fajgelbaum, Goldberg, Kennedy, and Khandelwal, 2020; Cavallo, Gopinath, Neiman, and Tang, 2021). Tariff increases raise the cost of a shipment directly; delays do not. Instead, they compress inventory buffers, create restocking uncertainty, and force firms to choose between absorbing a temporary disruption and preemptively adjusting prices. Because the fraction of import volume delayed may nonetheless be correlated with demand conditions or other product-level confounders, we do not interpret the panel estimates as causal. Nevertheless, the consistency of the positive coefficient across seven specifications with progressively stricter controls—absorbing time trends at the week, industry, and firm level—suggests that the association is not an artifact of common demand shocks, and the volume amplification pattern points to an inventory-compression mechanism: delays impose the largest price pressure when a firm's entire supply pipeline is at risk, leaving little buffer against stockouts.

To sharpen identification, we implement a stacked difference-in-differences design that exploits the staggered timing of first delay events across firm×product units. For each cohort—defined as the set of firm×products experiencing their first delay in a given week—we construct an event window of $[-6, +12]$ weeks around the delay onset and match treated units to not-yet-treated controls on pre-treatment log average price using propensity score matching with five nearest neighbors. All specifications include cohort×product and cohort×week fixed effects, absorbing unit-specific price levels and common time trends within each event window. The event study in Figure 1 documents flat pre-trends in the weeks preceding delay onset and a persistent upward shift in prices thereafter, supporting the parallel trends assumption. In the static specification, experiencing a delay raises retail prices by approximately 1.1% relative to matched controls ($\hat{\beta} = 0.011$, $p < 0.01$), a result that is stable across specifications and robust to controlling for total volume in transit and total units sold (Table 5). This effect is economically meaningful: during the height of the COVID-19 disruptions, the fraction of import volume delayed reached levels far exceeding the ten percentage point benchmark, implying that aggregate retail price inflation from the supply chain channel alone was substantial. More broadly, the finding that delay-driven price increases are persistent—not merely a short-lived spike—suggests that supply chain disruptions can sustain price pressure well beyond the initial disruption event. Importantly, the delay–price relationship holds outside the COVID-19 pandemic period, establishing it as a general feature of supply chain dynamics rather than an artifact of the extraordinary 2020–2021 disruptions.

Second, we study how household consumption responds to these product-level supply shocks. Using the Nielsen Consumer Panel, we construct a household-level stacked DiD that mirrors the price analysis, tracking the same delay events over an event window spanning 6 weeks before and 12 weeks after delay onset, matching treated households—those who purchased a delayed product in the 12 weeks preceding the event—to control households via propensity score matching on income, household size, age, and education of the household head. We find that delay events cause a statistically significant decline across all four consumption margins: the probability of purchasing any unit falls by 0.41 percentage points ($\hat{\beta} = -0.0041$, $p < 0.01$), the share of matched UPCs purchased falls by 0.27 percentage points ($\hat{\beta} = -0.0027$, $p < 0.01$), log quantity falls by 0.38% ($\hat{\beta} = -0.0038$, $p < 0.01$), and log expenditure falls by 0.53% ($\hat{\beta} = -0.0053$, $p < 0.01$; Table 6 and Figure 2).

These results reveal that the welfare cost of supply disruptions extends well beyond posted price increases: households bear real consumption losses through reduced quantities and reduced access, and conventional inflation measures—which track prices of available goods but not the cost of quantity rationing—understate the full burden. The consumption responses also vary across the income distribution in ways consistent with the predictions of heterogeneous agent models (Kaplan, Moll, and Violante, 2018; Auclert, 2019). Preliminary evidence suggests that these consumption responses are larger for lower-income households—who hold fewer liquid assets and have less ability to substitute across goods or retailers—implying that the aggregate welfare cost of supply disruptions may understate the burden faced by the most exposed households. The consumption results also serve a key interpretive role for the analysis that follows: they confirm that households directly notice and respond to delay-driven price increases, ruling out the possibility that the expectations results simply reflect households failing to encounter the shock.

Third, and most central to our motivating question, we analyze whether these supply-driven price increases propagate into household inflation expectations. We link each household's delay exposure to their reported one-year-ahead inflation expectations from the quarterly HMS survey, controlling for household and survey-wave fixed effects. Despite facing higher prices and significantly curtailing their purchases of affected products, households do not revise their inflation expectations upward in response to delay exposure. This null result stands in direct contrast to the experiential learning channel documented in prior work, where personally encountered price changes—particularly on frequently purchased staples such as milk, bread, and gasoline—drive expectation updating (D'Acunto et al., 2021; Cavallo et al., 2017; Malmendier and Nagel, 2016).

A natural interpretation combines two mechanisms. First, the same literature that motivates this channel shows that households form inflation beliefs disproportionately from goods they purchase frequently and repeatedly. The products most affected by shipping delays are durable and semi-durable imports that constitute a small and irregular share of most households' weekly expenditure. The price signal from delay-affected products is therefore too infrequent and too narrow to shift aggregate inflation beliefs. Second, even households that do observe higher prices on imported goods may attribute the increase to a transitory and product-specific disruption rather than to a broad inflation regime change, dampening the updating response. Together, these mechanisms suggest that the conditions under which price experiences update inflation beliefs depend critically on the nature and perceived source of the price change—not merely on its magnitude.

This finding has direct policy relevance. A key concern during the pandemic-era inflation was that supply-driven price increases might become self-fulfilling through an expectations channel: if households exposed to higher prices revise their inflation beliefs upward, these beliefs can feed into wage demands and pricing decisions, generating a secondary inflationary spiral. Our evidence suggests this channel did not activate for supply-chain-driven price shocks, even among households bearing real consumption costs. Supply-side price pressures of the kind observed during the pandemic need not become self-fulfilling—a finding that speaks to the debate on expectation anchoring and the conditions under which supply shocks pass into longer-run inflation beliefs (Coibion and Gorodnichenko, 2015).

Overall, our results trace the full downstream propagation of a real supply shock—from vessel-level shipping delays to firm-level retail price increases to household-level consumption losses—and show that this transmission is statistically significant and economically meaningful at each step. Importantly, our 13-year sample (2010–2023) allows us to establish this transmission chain both before and during the COVID-19 pandemic. At the same time, the chain stops short of household inflation expectations: despite facing higher prices and reduced product availability, households do not revise their inflation beliefs upward, suggesting that supply-driven price episodes are perceived as temporary and need not become self-fulfilling. These findings provide micro-level evidence that the source of price changes matters—supply chain disruptions reshape household consumption in ways that are consequential for welfare but distinct from the demand-driven price experiences that update inflation beliefs.

We contribute to at least three streams of the literature. First, we provide the first evidence on whether supply-side price shocks propagate to household inflation expectations. A well-established literature documents that personally experienced price changes shape inflation beliefs (D'Acunto et al., 2021; Cavallo et al., 2017; Malmendier and Nagel, 2016), but the identifying variation in these studies does not distinguish the source of the price change. We exploit a cleanly identified supply shock and find that exposed households do not revise their inflation expectations upward, despite significant consumption responses to the same price increases. This suggests that the experiential learning channel has boundary conditions—not all price increases are equal for belief formation—with direct relevance for expectation anchoring and central bank communication during supply disruptions (Coibion and Gorodnichenko, 2015).

Second, we document that supply chain disruptions impose real and unequal welfare costs on households. Exposed households reduce purchase probability, quantity, and total expenditure on affected products, with preliminary evidence suggesting that lower-income households bear a disproportionate share of the consumption loss. These results confirm that households directly experience delay-driven price increases—ruling out the trivial explanation for the expectations null—and connect to the redistribution mechanisms in heterogeneous agent models (Kaplan et al., 2018; Auclert, 2019).

Third, we provide direct estimates of how shipping delays pass through to retail consumer prices—a necessary first step for the household analysis above. The most closely related paper is Baslandze and Fuchs (2025), who study delay pass-through using Bill of Lading data and consumer panel prices. We differ in using vessel-level AIS data to measure actual versus scheduled arrival times, a sample spanning 2010–2023 that extends beyond the pandemic, and the addition of household consumption and expectations dimensions. Our results speak to a long-standing question about timing frictions in international trade (Hummels and Schaur, 2013) and suggest that the price consequences of supply chain disruptions extend beyond the direct cost channel emphasized in aggregate analyses (Carrière-Swallow, Deb, Furceri, Jiménez, and Ostry, 2023; Bai, Fernández-Villaverde, Li, and Zanetti, 2024).

2 Data

This section describes the shipping delay measures, the retail pricing data, and the household consumption panel. We create a unique dataset that links vessel-level shipping delays to both barcode-level retail prices and household-level purchase decisions and expectations from 2010 to 2023. We do so by combining four data sources: satellite-based vessel tracking data from Spire

Maritime, bill of lading (BoL) import filings from Panjiva, and retail prices and household purchase records from the Nielsen Retail Scanner (RMS) data and the Nielsen Consumer Panel (HMS) from the Kilts Center at the University of Chicago.

2.1 Shipping Delays and Supply Chain Disruptions

We measure supply chain disruptions using satellite-based Automatic Identification System (AIS) data from Spire Maritime, which continuously tracks the position of every ocean-going container vessel in the global fleet. Unlike aggregate shipping indices such as the Baltic Dry Index or the Global Supply Chain Pressure Index, the AIS data allow us to reconstruct each voyage's complete itinerary, including the port of origin, actual arrival time, and scheduled arrival time. The AIS data allows us to construct a vessel-level measure of delay for every container shipment arriving at a U.S. port. For each voyage, we measure delays as the difference between actual and scheduled arrival at the U.S. destination port.

We link voyage records to import filings using a crosswalk between the Maritime Mobile Service Identity (MMSI) number broadcast in real-time AIS signals and the International Maritime Organization (IMO) number recorded in BoL filings. The IMO is a permanent vessel identifier, whereas the MMSI can change when a vessel is re-flagged or re-registered. Combining trip-level delay with each shipment's TEU volume and the receiving firm's *gvkey* from the BoL allows us to construct a firm×product×week delay exposure measure. Our main explanatory variable, *Fraction in delay*, is the share of a firm's total in-transit import volume that is delayed in a given week.

2.2 Bills of Lading and Product Mapping

U.S. Customs and Border Protection collects a Bill of Lading filing for every containerized shipment arriving at a U.S. port. Each record identifies the receiving firm, the vessel and voyage, the shipment volume in twenty-foot equivalent units (TEUs), the arrival date, and the six-digit Harmonized System (HS) code of the imported goods. We obtain these records from Panjiva. Importantly, Panjiva provides Capital IQ identifiers for recipient firms, which allow us to identify them in Compustat and obtain their *gvkey* identifiers.

To connect import shipments to specific retail products, we map Nielsen product module codes to four-digit HS codes. Nielsen classifies all UPCs into roughly 1,000 product modules based on product descriptions. We match these to Harmonized System (HS) product categories using an LLM-based fuzzy-matching procedure that leverages the natural-language descriptions of both taxonomies. We use the four-digit HS codes because they are the finest aggregation at which BoL records reliably identify product types across importers. The resulting crosswalk links each firm's delay measures from its BoL shipments for a given HS code to the retail prices of its Nielsen products in the corresponding product module.

2.3 Retail Prices

Retail price data come from the Nielsen RMS database, which records weekly transaction prices at the barcode (UPC) level from more than 90,000 stores across all major U.S. retail chains from 2010 through 2023. We aggregate to the firm×UPC×week level by averaging prices and summing units sold across stores within each firm-product-week, yielding a weekly price series for each firm–product pair. The main outcome variable we obtain from these data is the log average retail price, $\ln(\bar{p}_{ipt})$.

For the panel regressions, we restrict the sample to *regularly shipped* firm×product pairs, which we define as products with at least 26 weeks of price observations, a shipment arrival rate of at least 10 percent of observed weeks, a 90th-percentile inter-shipment gap of no more than 26 weeks, and a coefficient of variation of inter-shipment gaps of at most one. These restrictions exclude firm–product pairs imported only sporadically or on highly irregular schedules, for which the delay exposure measure is unlikely to capture meaningful supply chain conditions. The estimation sample is held fixed to the support of the most demanding specification to ensure comparability across columns.

2.4 Household Consumption and Inflation Expectations

Data on household purchasing behavior come from the Nielsen Consumer Panel (HMS), a nationally representative longitudinal panel that tracks the complete purchase histories of approximately 60,000 U.S. households across all product categories and retail outlets from 2010 to 2023. Each purchase record contains the household identifier, UPC, purchase date, quantity, and total price paid. We aggregate to the household×UPC×week level.

HMS households also participate in a quarterly survey of one-year-ahead inflation expectations following the methodology of D'Acunto et al. (2021). We link survey responses to purchase histories via a household identifier bridge and construct each household's delay exposure from their pre-survey purchase basket. This linkage allows us to trace the full chain from a firm-level shipping delay to a household's realized price experience to their forward-looking inflation beliefs.

2.5 Descriptive Statistics

Table 1 reports summary statistics for the main analysis sample at the firm×product×week level. The sample contains 12.3 million observations.

Panel A describes the delay measures. In-transit import volume averages 11.7 TEUs per firm–product–week, with a right-skewed distribution (median of zero) reflecting that most product–weeks have no active shipment in transit. In 20 percent of observations, at least some fraction of in-transit volume is delayed; unconditionally, the average share of in-transit volume delayed is 13 percent. Among observations with positive delay, the volume-weighted average delay is 17.8 days, with considerable dispersion (standard deviation of 62 days) reflecting the coexistence of minor scheduling adjustments and severe disruptions during the COVID-19 pandemic.

Panel B describes the retail pricing outcomes. The average retail price in the sample is \$11.32, with a median of \$5.99 and substantial right skew, consistent with the heterogeneous product mix spanning food, personal care, and household goods.

3 Empirical Strategy

This section describes the two empirical approaches used to estimate the effect of shipping delays on retail prices and household consumption outcomes. We first document the association between shipping delays and prices in a firm×product×week panel, and then identify causal effects using a stacked difference-in-differences design that exploits the staggered timing of delay events.

3.1 Panel regressions

We estimate the relationship between shipping delays and retail prices using the following panel regression:

$$\ln(\bar{p}_{ipt}) = \alpha_t + \beta \, \textit{Fraction in delay}_{ipt} + \gamma' \mathbf{Z}_{ipt} + \varepsilon_{ipt}, \tag{1}$$

where i indexes firms, p products (UPCs), and t weeks. The dependent variable is the log average retail price across stores for a given firm–product–week. $\textit{Fraction in delay}_{ipt}$ is the share of firm i's total in-transit import volume (in TEUs) for product p that is delayed in week t. The vector $\mathbf{Z}_{ipt}$ includes log total volume in transit and log total units sold. α_t denotes fixed effects, which we progressively enrich across specifications. Specifically, we include week fixed effects, week and HS code fixed effects, HS code×week, firm, firm×week, and UPC prefix fixed effects. The coefficient of interest, β, captures the correlation between the fraction of delayed import volume and retail prices.

Because the fraction of import volume delayed may be correlated with demand conditions or other product-level confounders, we do not interpret the panel estimates as causal. To isolate causal effects, we turn to a stacked difference-in-differences design.

3.2 Stacked difference-in-differences

Delay events create a staggered treatment setting, with treatment timing varying across firm×product units. Given that the traditional difference-in-differences estimator can be biased under staggered adoption (e.g., Goodman-Bacon, 2021), we adopt a stacked regression approach (Gormley and Matsa, 2011; Cengiz, Dube, Lindner, and Zipperer, 2019; Baker, Larcker, and Wang, 2022). This method constructs separate cohorts for each delay event, ensuring that control units are not yet treated within each cohort, which avoids the contamination bias that arises when already-treated units serve as controls. The baseline stacked difference-in-differences specification is:

$$Y_{iptj} = \alpha_{ipj} + \alpha_{tj} + \beta \, \textit{Treated}_{ipj} \times \textit{Post}_{tj} + \gamma' \mathbf{Z}_{ipt} + \varepsilon_{iptj}, \tag{2}$$

where i indexes firms, p products (UPCs), t weeks, and j cohorts. A cohort is defined as the set of firm×product units experiencing their first delay in a given week. $\textit{Treated}_{ipj}$ equals one for firm×product units that experience a delay in cohort j. $\textit{Post}_{tj}$ equals one during the post-delay period within each cohort's event window. $\mathbf{Z}_{ipt}$ includes log total volume in transit and log total units sold.

The specification includes product×cohort fixed effects α_{ipj} and calendar week×cohort fixed effects α_{tj}, which absorb cohort-specific unit heterogeneity and common time trends within each event window. By interacting fixed effects with the cohort indicator j, the model estimates treatment effects separately within each cohort and pools these estimates via OLS. The coefficient of interest, β, captures the average treatment effect of a delay event on retail prices. Standard errors are clustered at the cohort×product level.

We construct an event window of $[-6, +12]$ weeks around delays and match treated units to not-yet-treated controls on pre-treatment log average price using propensity score matching. To test for differential pre-trends and characterize the dynamics of the treatment effect, we also estimate an event-study version of Equation 2:

$$Y_{iptj} = \alpha_{ipj} + \alpha_{tj} + \sum_{k \in \{-6,\ldots,+12\},\, k \neq -1} \beta_k \left(\mathbb{1}[Rel_t = k] \times \mathit{Treated}_{ipj} \right) + \gamma' \mathbf{Z}_{ipt} + \varepsilon_{iptj}, \qquad (3)$$

where Rel_t is event time in weeks relative to shipping delays. We omit $k = -1$ as the reference period. The pre-event coefficients $\{\beta_k\}_{k<0}$ provide a direct test of the parallel-trends assumption, while the post-event coefficients characterize the timing and persistence of treatment effects.

We apply the same stacked DiD framework to household consumption outcomes, replacing the unit of observation with household×product pairs. In those regressions, the dependent variable Y_{hptj} is one of the following: an indicator for whether household h purchases product p in week t, the UPC's share of the household's within-category expenditure, log quantity purchased, and log expenditure. Fixed effects become household×cohort (α_{hj}) and week×cohort (α_{tj}), absorbing household-specific purchasing patterns within each event window. We match treated households to not-yet-treated controls using propensity scores estimated on pre-treatment demographics, including household income, household size, age of household head, and education, rather than pre-treatment prices. The coefficient β then captures the average effect of a delay event on each consumption margin.

4 Shipping Delays and Retail Prices

4.1 Panel regressions

Table 2 presents the results from estimating Equation 1. Across all seven specifications, the coefficient on *Fraction in delay* is positive. In column (1), which includes only week fixed effects, the estimated coefficient is 0.170, statistically significant at the 5% level. Adding controls for log total volume in transit and log total units sold in column (2) and HS code fixed effects in column (3) leaves the estimate largely unchanged. In our most demanding specification, column (7), which includes HS code×week, firm×week, and UPC prefix fixed effects, the estimated coefficient is 0.139, statistically significant at the 1% level. Based on this specification, a ten percentage point increase in the fraction of import volume delayed is associated with a 1.4% increase in retail prices. The stability of the coefficient across specifications that progressively absorb time trends at the week, industry, and firm level suggests that the association is not an artifact of common demand shocks.

We further examine whether the price effect depends on the scale of a firm's import pipeline. Table 3 reports results from specifications that interact *Fraction in delay* with log total volume in transit. The interaction term is positive and statistically significant across all six specifications, indicating that the price effect of delays is amplified when a firm has more volume in transit. This pattern is consistent with a supply-exposure channel whereby firms with larger shipment volumes at sea face greater disruption to their incoming supply when delays occur, translating into larger price increases.

We next examine whether firms respond to delays by accelerating restocking. Table 4 regresses an indicator for a new shipment arrival in week $t+1$ on the fraction of volume currently delayed. The coefficient is positive and statistically significant across most specifications, indicating that firms are more likely to receive a new shipment in the week following a delay. This restocking response, if anything, would attenuate the estimated price effects, lending our pass-through estimates a conservative interpretation.

4.2 Stacked difference-in-differences

Table 5 presents the results from estimating Equation 2. In column (1), experiencing a delay increases retail prices by 0.9% relative to matched controls. Column (2) adds log total volume in transit as a control, and the estimated treatment effect increases slightly to 1.2%. In the most saturated specification, column (3), which additionally controls for log total units sold, the estimate is 0.011, corresponding to a 1.1% price increase. All three estimates are statistically significant at the 1% level. The stability across specifications suggests that the treatment effect is not sensitive to controlling for potentially endogenous volume and quantity measures.

Figure 1 presents the corresponding event-study estimates from Equation 3. The pre-treatment coefficients are close to zero and statistically insignificant in the weeks before delays, supporting the parallel-trends assumption. The treatment effect emerges sharply at the time of delays and persists through the end of the event window, suggesting that delay-driven price increases are not merely short-lived spikes but rather sustained adjustments. This persistence is consistent with firms adjusting prices in response to inventory depletion that takes time to resolve through restocking.

5 From Prices to Household Consumption Behavior

We now examine whether households adjust their purchasing behavior in response to delay-driven price increases. We construct a household-level stacked DiD that mirrors the price analysis, exploiting the same delay events.

5.1 Sample construction

Using the Nielsen Consumer Panel, we identify households that purchased a product from a firm experiencing a shipping delay during a qualification window spanning weeks -24 to -13 relative to the delay event. These households constitute the treated group. Control households are those that purchased products from non-delayed firms in the same product categories during the same qualification window. Treated and control households are matched via propensity score matching on income, household size, age of household head, and education of household head. The event window spans 6 weeks before and 12 weeks after the start of delays. All specifications include household$\times$cohort and calendar week$\times$cohort fixed effects.

5.2 Results

Table 6 presents the results across four consumption margins. In column (1), the probability of purchasing any matched UPC in a given week falls by 0.41 percentage points for treated households relative to controls. Column (2) shows that the expenditure share allocated to matched UPCs declines by 0.27 percentage points. Columns (3) and (4) report effects on log quantity and log expenditure, respectively: quantity falls by 0.38% and expenditure falls by 0.53%. All four estimates are statistically significant at the 1% level. The fact that the expenditure decline exceeds the quantity decline is consistent with households reducing both the number of units purchased and their total spending on affected products.

Figure 2 presents the corresponding event-study estimates. Panel A shows that the probability of purchase is stable in the pre-treatment period and declines sharply after delays. Panels B through D document similar patterns for expenditure, quantity, and expenditure share, respectively.

The welfare cost of supply disruptions therefore extends beyond posted price increases. Households reduce purchase frequency, quantities, and total expenditure on affected products. The pattern is consistent with substitution away from delay-affected products rather than absorption of higher prices on the same goods. Because conventional inflation measures track prices of available goods but not the cost of quantity rationing or product substitution, they may understate the full burden of supply chain disruptions on consumer welfare.

6 Household Inflation Expectations

We next ask whether the price increases and consumption responses documented above propagate into household inflation expectations. Supply chain disruptions may be perceived as transitory and product-specific rather than as signals of broad inflationary pressure, in which case households need not revise their inflation beliefs upward even when facing higher prices on affected goods.

6.1 Realized inflation and expectations

We link each household's delay exposure to their reported one-year-ahead inflation expectations from the quarterly HMS survey, controlling for household and survey-wave fixed effects. We construct a household-specific realized inflation measure by weighting product-level price changes by the household's lagged consumption shares. This measure captures each household's personal inflation experience, which the literature has shown to be the dominant driver of expectation formation (D'Acunto et al., 2021; Malmendier and Nagel, 2016).

Table 7 presents the results. Column (1) regresses household-level realized inflation on delay exposure. The coefficient is 0.057, statistically significant at the 5% level. It indicates that households more exposed to shipping delays experience higher realized inflation. Column (2) tests whether this realized inflation predicts expectations. The coefficient on household CPI is 0.14 but statistically insignificant, indicating that, conditional on household and wave fixed effects, cross-sectional variation in realized inflation does not translate into higher expected inflation. Column (3) tests whether delay exposure directly affects expectations. The coefficient is small (0.10) and statistically insignificant.

The null results in columns (2) and (3) are consistent with two mechanisms. First, the products most affected by shipping delays are durable and semi-durable imports that constitute a small and irregular share of most households' weekly expenditure; the price signal is too weak and too infrequent to shift aggregate inflation beliefs. Second, households that observe higher prices on imported goods may attribute the increase to a transitory, product-specific shock rather than to a broad inflation regime change. Supply chain-driven price increases therefore need not become self-fulfilling through an expectations channel, a finding with direct relevance for central bank communication during episodes of supply disruption.

7 Conclusion

This paper asks whether the source of a price change matters for how households respond—in both their consumption and their beliefs. Using a novel dataset that traces supply-side shocks from vessel-level shipping delays through retail prices to household purchasing behavior and inflation expectations, we find that it does. Shipping delays raise retail prices significantly, and households respond by curtailing purchases of affected products. Yet despite these real consumption costs, households do not revise their inflation expectations upward. The disconnect between behavioral response and belief formation suggests that households implicitly distinguish supply-driven price increases from signals of broad inflationary pressure.

These findings carry implications for both policy and research. For central banks, the results suggest that supply-side price pressures need not become self-fulfilling through an expectations channel—even when they impose meaningful welfare costs on consumers. For the household finance literature, our evidence points to an important boundary condition on experiential learning: not all price increases are equal for belief updating, and understanding which price signals households generalize from remains an open question. Future work examining whether this pattern extends to other supply-side shocks, to different product categories, and to settings where supply disruptions are more salient in public discourse would help delineate the conditions under which supply-driven inflation does—and does not—feed back into expectations.

References

Auclert, A. 2019. Monetary Policy and the Redistribution Channel. *American Economic Review* 109:2333–2367.

Bai, X., J. Fernández-Villaverde, Y. Li, and F. Zanetti. 2024. The Causal Effects of Global Supply Chain Disruptions on Macroeconomic Outcomes: Evidence and Theory. Working paper.

Baker, A. C., D. F. Larcker, and C. C. Y. Wang. 2022. How much should we trust staggered difference-in-differences estimates? *Journal of Financial Economics* 144:370–395.

Baslandze, S., and W. Fuchs. 2025. The Price of Delay: Supply Chain Disruptions and Pricing Dynamics. Federal Reserve Bank of Atlanta Working Paper 2025-8.

Carrière-Swallow, Y., P. Deb, D. Furceri, D. Jiménez, and J. D. Ostry. 2023. Shipping Costs and Inflation. *Journal of International Money and Finance* 130:102771.

Cavallo, A., G. Cruces, and R. Perez-Truglia. 2017. Inflation Expectations, Learning, and Supermarket Prices: Evidence from Field Experiments. *American Economic Journal: Macroeconomics* 9:1–35.

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* 3:19–34.

Cengiz, D., A. Dube, A. Lindner, and B. Zipperer. 2019. The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics* 134:1405–1454.

Coibion, O., and Y. Gorodnichenko. 2015. Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts. *American Economic Review* 105:2644–2678.

D'Acunto, F., U. Malmendier, J. Ospina, and M. Weber. 2021. Exposure to Grocery Prices and Inflation Expectations. *Journal of Political Economy* 129:1615–1639.

Fajgelbaum, P. D., P. K. Goldberg, P. J. Kennedy, and A. K. Khandelwal. 2020. The Return to Protectionism. *Quarterly Journal of Economics* 135:1–55.

Goodman-Bacon, A. 2021. Differences-in-differences with variation in treatment timing. *Journal of Econometrics* 225:254–277.

Gormley, T. A., and D. A. Matsa. 2011. Growing out of trouble? Corporate responses to liability risk. *The Review of Financial Studies* 24:2781–2821.

Hummels, D. L., and G. Schaur. 2013. Time as a trade barrier. *American Economic Review* 103:2935–2959.

Kaplan, G., B. Moll, and G. L. Violante. 2018. Monetary Policy According to HANK. *American Economic Review* 108:697–743.

Malmendier, U., and S. Nagel. 2016. Learning from Inflation Experiences. *Quarterly Journal of Economics* 131:53–87.

Figure 1: Price level event study

This figure plots the dynamic treatment effects and 95% confidence intervals from stacked difference-in-differences regressions of the form: $\ln(\text{Average sales price})_{ipt} = \alpha_{ips} + \alpha_{st} + \sum_{k \neq -1} \beta_k \mathbb{1}[\text{Event time} = k] \times \text{Treated}_{ip} + \epsilon_{ipt}$, where i denotes firms, p products (UPCs), and t weeks. The dependent variable is the log of the average retail price across stores for a given firm, product, and week. *Treated* equals one for firm×product units that experience a shipment delay in the cohort week. The sample is constructed by stacking cohort-specific event windows of $[-6, +12]$ weeks around the first delay event for each firm×product unit, together with not-yet-treated controls. Treated and control units are matched within each cohort on pre-treatment log average price using propensity score matching. Event time $k = 0$ denotes the week of the delay, with $k = -1$ as the omitted base period. α_{ips} and α_{st} denote product×cohort and calendar week×cohort fixed effects, respectively. Standard errors are clustered at the cohort×product level.

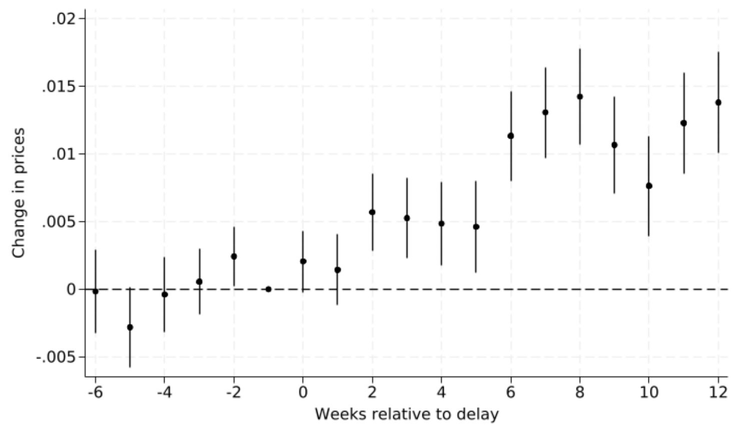

Figure 2: Household consumption event study

This figure plots the dynamic treatment effects and 95% confidence intervals from stacked difference-in-differences regressions of the form: $Y_{ipt} = \alpha_{is} + \alpha_{st} + \sum_{k \neq -1} \beta_k \mathbb{1}[\text{Event time} = k] \times \text{Treated}_{ip} + \epsilon_{ipt}$, where i denotes households, p products (UPCs), and t weeks. The dependent variable Y_{ipt} is a household-product consumption outcome, including the probability of purchase (Panel A), log expenditure (Panel B), log quantity purchased (Panel C), and the share of products in the pre-delay household consumption basket that are purchased during delays (Panel D). *Treated* equals one for household-product pairs where the product is supplied by an importing firm that experiences a shipping delay in the cohort week. Control household-product pairs purchased the same product from firms that are either never treated or not yet treated in the event window. The sample is constructed by stacking cohort-specific event windows of $[-6, +12]$ weeks around the first delay event for each firm×product unit, together with not-yet-treated controls. Households are included if they purchased any UPC within the same product module during weeks $[-24, -13]$ relative to the delay event. Treated and control households are matched within cohorts via propensity score matching on demographic characteristics including household income, household size, household head age, and household head education. All specifications include household×cohort and calendar week×cohort fixed effects. Standard errors are clustered at the household level.

Panel A: Probability of purchase

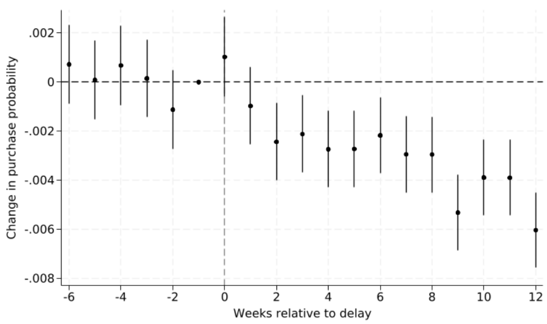

Panel B: ln(Expenditure)

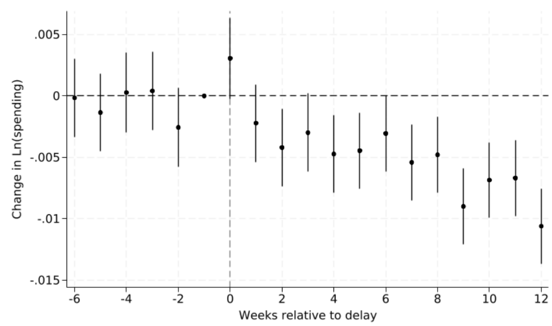

Panel C: ln(Quantity)

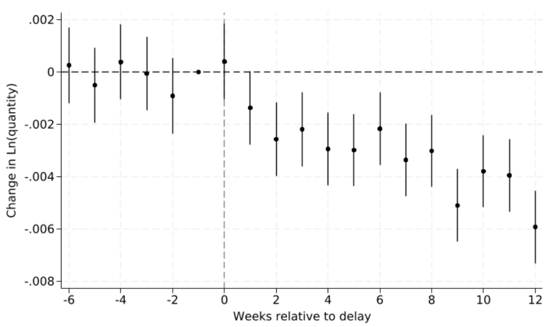

Panel D: Expenditure share

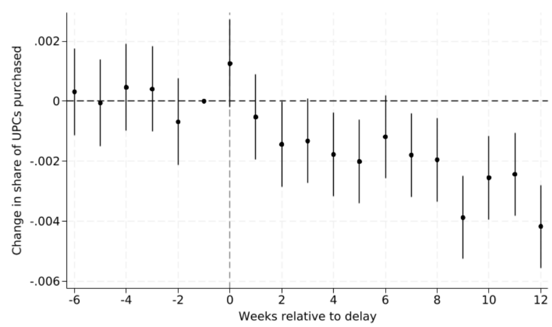

Table 1: Summary Statistics

This table reports summary statistics at the firm×product×week level. Panel A reports delay measures constructed from vessel tracking and bill of lading data. Panel B reports product prices from the Nielsen Retail Scanner data.

	Obs (1)	Mean (2)	Std. Dev. (3)	25th Pct. (4)	Median (5)	75th Pct. (6)
Panel A: Delay measures						
Vol. (TEU)	12,312,610	11.70	55.08	0.00	0.00	2.00
Delayed vol. (TEU)	12,312,610	5.98	34.96	0.00	0.00	0.00
Vol. delayed (%)	12,312,610	0.13	0.29	0.00	0.00	0.00
Vol. delayed > 0 (d)	12,312,610	0.20				
Days delayed (VW)	12,312,610	3.66	28.99	0.00	0.00	0.00
Days delayed (VW) $\mid$ delayed > 0	2,522,408	17.84	62.04	0.95	3.23	8.76
Panel B: Product prices						
Average product price	12,312,610	11.32	26.93	3.06	5.99	11.77

Table 2: Shipping delays and retail prices

This table presents results from panel regressions of log retail price on the fraction of import volume in transit that is delayed. The sample is a firm×product×week panel linking U.S. publicly listed firms to Nielsen retail scanner data. The dependent variable is the log of the average retail price across stores for a given firm, product, and week. *Fraction in delay* is the share of a firm's in-transit import volume that is experiencing a delay relative to its expected arrival date. The sample is restricted to regularly shipped firm×product pairs, defined as those with at least 26 weeks of observations, a shipment rate of at least 10%, a 90th-percentile inter-shipment gap of at most 26 weeks, and a coefficient of variation of inter-shipment gaps of at most one. The estimation sample is held fixed to the support of the most demanding specification. Columns (1) and (2) include week fixed effects; column (3) adds HS code fixed effects; columns (4)–(7) replace these with HS code × week fixed effects. Column (5) adds firm fixed effects; columns (6) and (7) replace these with firm × week fixed effects. Column (7) further includes UPC prefix fixed effects. Columns (2)–(7) control for log total volume in transit and log total units sold. Standard errors clustered at the four-digit HS code level are reported in parentheses. Statistical significance at the 1, 5, and 10 percent significance levels is denoted by ***, **, and *, respectively.

	ln(Average sales price)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Fraction in delay	0.170**	0.143*	0.044	0.149*	0.060*	0.369**	0.139***
	(0.082)	(0.075)	(0.041)	(0.075)	(0.032)	(0.152)	(0.041)
Ln(Total volume in transit)		0.093	0.012	0.015	−0.013**	−0.017	0.015
		(0.061)	(0.011)	(0.018)	(0.006)	(0.014)	(0.017)
Ln(Total units sold)		−0.019	0.011**	0.005	0.009	0.011	0.017*
		(0.017)	(0.005)	(0.004)	(0.008)	(0.008)	(0.009)
Fixed effects							
Week	✓	✓	✓				
HS Code			✓				
Firm					✓		
HS Code × Week				✓	✓	✓	✓
Firm × Week						✓	✓
UPC Prefix							✓
Observations	2,209,667	2,209,667	2,209,667	2,209,667	2,209,667	2,209,667	2,209,667
Adj. R^2	0.01	0.04	0.43	0.46	0.53	0.54	0.61

Table 3: Shipping delays, retail prices, and volume in transit

This table presents results from panel regressions of log retail price on the fraction of import volume in transit that is delayed, interacted with log total volume in transit. The dependent variable is the log of the average retail price across stores for a given firm, product, and week. The interaction term captures whether the price effect of delays is amplified or attenuated by the scale of a firm's import activity. *Fraction in delay* and *Ln(Total volume in transit)* are defined as in Table 2. The sample, controls, and fixed effect specifications mirror those in Table 2, with the interaction included in all specifications. Standard errors clustered at the four-digit HS code level are reported in parentheses. Statistical significance at the 1, 5, and 10 percent significance levels is denoted by ***, **, and *, respectively.

	ln(Average sales price)					
	(1)	(2)	(3)	(4)	(5)	(6)
Fraction in delay	−0.070	−0.033	0.008	0.001	0.136	0.052
	(0.072)	(0.029)	(0.065)	(0.014)	(0.091)	(0.033)
Ln(Total volume in transit)	0.026	−0.012	-0.034^{*}	-0.034^{***}	-0.097^{***}	−0.015
	(0.048)	(0.010)	(0.019)	(0.009)	(0.024)	(0.022)
Fraction in delay × Ln(Volume)	0.140^{***}	0.055^{**}	0.101^{***}	0.045^{***}	0.189^{***}	0.080^{***}
	(0.038)	(0.022)	(0.032)	(0.014)	(0.039)	(0.026)
Ln(Total units sold)	−0.019	0.011^{**}	0.005	0.009	0.011	0.017^{*}
	(0.017)	(0.004)	(0.004)	(0.008)	(0.008)	(0.009)
Fixed effects						
Week	✓	✓				
HS Code		✓				
Firm				✓		
HS Code × Week			✓	✓	✓	✓
Firm × Week					✓	✓
UPC Prefix						✓
Observations	2,209,667	2,209,667	2,209,667	2,209,667	2,209,667	2,209,667
Adj. R^2	0.05	0.43	0.47	0.53	0.54	0.61

Table 4: Shipping delays and inventory restocking

This table validates the delay measure by testing whether firms actively replenish inventory in response to shipping disruptions. The dependent variable is the deviation of new shipment arrivals at $t+1$ from the trailing 12-week mean (abnormal new shipments). A positive coefficient confirms that delays reflect genuine supply disruptions to which firms respond by placing additional orders. *Fraction in delay* is defined as in Table 2. The sample, controls, and fixed effect specifications mirror those in Table 2. Because this restocking response increases incoming supply, it works against price increases; the price effects reported in Table 5 are therefore conservative. Standard errors clustered at the four-digit HS code level are reported in parentheses. Statistical significance at the 1, 5, and 10 percent significance levels is denoted by ***, **, and *, respectively.

	I(New Shipment$_{t+1}$)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Fraction in delay	0.021	0.031***	0.032***	0.025***	0.021**	−0.004	0.043***
	(0.014)	(0.009)	(0.008)	(0.006)	(0.009)	(0.025)	(0.013)
Ln(Total volume in transit)		−0.026***	−0.026***	−0.030***	−0.038***	−0.044***	−0.047***
		(0.004)	(0.004)	(0.004)	(0.007)	(0.005)	(0.004)
Ln(Total units sold)		0.002*	0.001**	0.001**	0.000*	0.000	0.000
		(0.001)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Fixed effects							
Week	✓	✓	✓				
HS Code			✓				
Firm					✓		
HS Code × Week				✓	✓	✓	✓
Firm × Week						✓	✓
UPC Prefix							✓
Observations	2,193,343	2,193,343	2,193,343	2,193,322	2,193,322	2,193,314	2,193,314
Adj. R^2	0.09	0.13	0.15	0.57	0.58	0.91	0.91

Table 5: Retail prices around delay shocks

This table presents results from stacked difference-in-differences regressions of log retail price around shipment delay events. The sample is constructed by stacking cohort-specific event windows, where each cohort consists of firm×product units that experience their first delay in a given week, together with not-yet-treated controls. Units are matched within each cohort on pre-treatment log average price using propensity score matching with five nearest neighbors. The event window spans $[-6, +12]$ weeks around the delay onset. *Treated* × *Post* equals one for treated units in the post-delay period. All specifications include cohort×product and cohort×week fixed effects, which absorb cohort-specific unit heterogeneity and common time trends within each event window. Column (1) includes only cohort×product and cohort×week fixed effects. Column (2) adds log total volume in transit as a control. Column (3) additionally controls for log total units sold. Standard errors clustered at the cohort×product level are reported in parentheses. Statistical significance at the 1, 5, and 10 percent significance levels is denoted by ***, **, and *, respectively.

	ln(Average sales price)		
	(1)	(2)	(3)
Treated × Post	0.009***	0.012***	0.011***
	(0.001)	(0.002)	(0.002)
Ln(1+Total volume in transit)		−0.006***	−0.006***
		(0.001)	(0.001)
Ln(Total units sold)			−0.011***
			(0.001)
Fixed effects			
UPC × Cohort	✓	✓	✓
Calendar week × Cohort	✓	✓	✓
Observations	1,495,712	1,495,712	1,495,712
Adj. R^2	0.98	0.98	0.98

Table 6: Household consumption around delay shocks

This table presents results from stacked difference-in-differences regressions of household consumption outcomes around shipment delay events. The sample is constructed by matching households to treated and control firm×product units from the stacked DiD panel, using the Nielsen Consumer Panel. Households are classified as treated if they purchased a delayed product, or as control if they purchased a non-delayed product, during a qualification window spanning weeks −24 to −13 relative to the delay event. Treated and control households are then matched via propensity score matching on income, household size, age of household head, and education of household head. The event window spans 6 weeks before and 12 weeks after the delay onset. *Treated* × *Post* equals one for treated households in the weeks following delay onset. Column (1) reports the effect on the probability of purchasing any UPC in a given week. Column (2) reports the fraction of UPCs purchased. Column (3) reports the log of one plus total quantity purchased across UPCs. Column (4) reports the log of one plus total expenditure across UPCs. All specifications include cohort×household and cohort×week fixed effects. Standard errors clustered at the household level are reported in parentheses. Statistical significance at the 1, 5, and 10 percent significance levels is denoted by ***, **, and *, respectively.

	Any purchased	Share purchased	Ln(Quantity)	Ln(Spending)
	(1)	(2)	(3)	(4)
Treated × Post	-0.0041^{***}	-0.0027^{***}	-0.0038^{***}	-0.0053^{***}
	(0.0007)	(0.0006)	(0.0006)	(0.0013)
Fixed effects				
HH × Cohort	✓	✓	✓	✓
Calendar week × Cohort	✓	✓	✓	✓
Observations	4,640,730	4,640,730	4,640,730	4,640,730
Adj. R^2	0.14	0.13	0.15	0.14

Table 7: Shipping delays, realized inflation, and inflation expectations

This table presents results from household-level regressions linking shipping delays to realized inflation and inflation expectations, using the Nielsen Consumer Panel matched to the quarterly HMS inflation expectations survey. The dependent variable in column (1) is household-level CPI, computed by weighting product-level price changes by the household's lagged consumption shares. The dependent variable in columns (2) and (3) is the household's one-year-ahead inflation expectation reported in the HMS survey. *Household Delays* is the consumption-weighted share of a household's expenditure allocated to products experiencing shipping delays. *Household CPI* is the household-level realized inflation measure. Consumption controls include the fraction of expenditures on shipped products and log total expenditures. All specifications include household and survey-wave fixed effects. Standard errors two-way clustered by household and survey wave are reported in parentheses. Statistical significance at the 1, 5, and 10 percent significance levels is denoted by ***, **, and *, respectively.

	Household CPI	Expected inflation (%)	
	(1)	(2)	(3)
Household Delays	0.0570**		0.0954
	(0.0234)		(2.6877)
Household CPI		0.1401	
		(0.3822)	
Fixed effects			
Consumption controls	✓	✓	✓
Household	✓	✓	✓
Wave	✓	✓	✓
Observations	215,399	215,399	215,399
Adj. R^2	0.471	0.373	0.373

Appendix A: Variable Definitions

Variable	Definition	Source
Panel A: Dependent variables		
ln(Average price)	The natural log of the average unit sales price of a firm-product (UPC) in a given week, computed as total revenue divided by total units sold across stores.	Nielsen Retail Scanner (RMS)
Abnormal shipment	The difference between a firm-product's new shipment indicator in week t and its trailing 12-week moving average. Used to measure inventory restocking responses to delay shocks.	Bill of Lading; Spire AIS
Panel B: Treatment and exposure variables		
Fraction of volume delayed	The fraction of a firm-product's total import volume (in TEUs) that is delayed in a given week, where delay is defined as actual arrival date exceeding the expected arrival date.	Spire AIS; Bill of Lading
Volume delayed (> 0)	An indicator variable equal to one if any import volume for a firm-product is delayed in a given week.	Spire AIS; Bill of Lading
VW delay days	The volume-weighted average number of days of delay across all shipments for a firm-product in a given week. Weights are shipment-level TEU volumes.	Spire AIS; Bill of Lading
VW delay days | delayed	Volume-weighted average delay days, conditional on having positive delay.	Spire AIS; Bill of Lading
Panel C: Shipment and volume variables		
Total volume in transit	The total volume (in TEUs) of imports in transit for a firm-product in a given week, including both delayed and on-time shipments.	Spire AIS; Bill of Lading
Total volume delayed	The total volume (in TEUs) of delayed imports for a firm-product in a given week.	Spire AIS; Bill of Lading
Total units sold	The total number of units of a UPC sold across all stores in a given week.	Nielsen Retail Scanner (RMS)

Variable	Definition	Source
Panel D: Household consumption variables		
$\mathbb{P}$(Purchase)	An indicator equal to one if the household purchased any of its matched UPCs in a given week.	Nielsen Consumer Panel (HMS)
Fraction of UPCs purchased	The fraction of matched UPCs that the household purchased in a given week.	Nielsen Consumer Panel (HMS)
ln(1 + Quantity)	The natural log of one plus total quantity purchased across all matched UPCs by the household in a given week.	Nielsen Consumer Panel (HMS)
ln(1 + Expenditure)	The natural log of one plus total expenditure (in dollars) across all matched UPCs by the household in a given week.	Nielsen Consumer Panel (HMS)
Panel E: Household demographics (matching variables)		
Income	Household income bracket. Used as a matching variable in propensity score matching of treated and control households.	Nielsen Consumer Panel (HMS)
Household size	Number of members in the household. Used as a matching variable in propensity score matching.	Nielsen Consumer Panel (HMS)
Age of household head	Age of the household head. Used as a matching variable in propensity score matching.	Nielsen Consumer Panel (HMS)
Education of household head	Education level of the household head. Used as a matching variable in propensity score matching.	Nielsen Consumer Panel (HMS)