More than a Late Shipment: How Supply Chain Disruptions Affect Households

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

October 14, 2026

Abstract

We study how supply chain disruptions affect households, in the prices they pay, the goods they can buy, and the inflation they expect, using shipping delays as a well-measured, supply-driven shock and a novel dataset for 2010–2024 that links satellite-based vessel-level delays to importing firms' barcode-level retail prices, household purchase records, and inflation expectations surveys. Shipping delays raise retail prices: an increase of one standard deviation in the share of a firm's imports arriving on delayed vessels raises prices for identical products by 0.6%, and prices remain elevated for months after individual delay events. Delays also remove products from shelves and shift the mix of products that remain available toward premium varieties, so the price of the goods that remain available increases by more than four times the same-product response. Exposed households do not exit affected product markets but shift spending toward cheaper and more familiar products, while stockouts push part of their bundles into replacement purchases. Households with more delay-exposed baskets also report higher expected inflation over the following 12 months, and although bundle composition accounts for much of this cross-sectional relationship, households whose exposure rises more also tend to revise their expectations upward. Supply-side shocks thus raise the prices households pay and change the products available to them, margins that posted-price measures alone do not fully capture.

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

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

We thank Xi Wang for excellent research assistance. Researcher(s)' own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researcher(s) and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

[†]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

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

[¶]Daniels School of Business, Purdue University, ESMT Berlin, NBER, CEPR, CESifo, and ifo Institute, Email: weber366@purdue.edu

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

1 Introduction

Supply shocks routinely disrupt global product markets. Congested ports and delayed container ships left retailers waiting on goods during the COVID-19 pandemic. At the early-2022 peak, more than 70 percent of U.S.-bound container voyages ran late, alongside the sharpest U.S. inflation in four decades. Yet such disruptions are no artifact of the pandemic: container shipping ran late through the 2015 West Coast port slowdown before COVID-19, and through the Panama Canal drought and Red Sea attacks after it, making them likely to recur. Nor are such disruptions confined to ocean shipping or to the United States: extreme heat has repeatedly pushed water levels on the Rhine, Europe's most important inland waterway, to near-record lows, cutting the cargo vessels can carry and slowing the flow of goods. Moreover, a late shipment is not simply a matter of timing. Additional time in transit is itself costly for the importing firm (Hummels and Schaur, 2013), and a delayed arrival leaves it short of inventory until stock is replenished. Firm-level supply shocks of this kind are well-known to propagate to prices and quantities through production networks (e.g., Acemoglu, Carvalho, Ozdaglar, and Tahbaz-Salehi, 2012; Barrot and Sauvagnat, 2016; Baqaee and Farhi, 2019; Carvalho, Nirei, Saito, and Tahbaz-Salehi, 2021). How they affect households, however, is less understood. Standard inflation measures track posted prices but miss the stockouts that change what households find on the shelf, so they capture the household burden only in part. We use shipping delays as a well-measured, supply-driven shock, and ask whether the costs of a delayed shipment affect households, in the prices they pay, the goods they can buy, and the inflation they expect.

Tracing a supply shock from vessels through to households requires linking shipping delays to retail prices and product availability, household consumption, and household beliefs within a single dataset, which we build from five sources. First, we use satellite-based Automatic Identification System (AIS) data to track the positions of every container ship. Unlike aggregate indices of shipping conditions, such as the Baltic Dry Index or the Global Supply Chain Pressure Index, the AIS data allow us to reconstruct the entire journey of a vessel during a particular voyage, including the origin and destination ports, departure times, and the actual and expected arrival times. With this information, we construct a precise vessel-level measure of delays for every container shipment

arriving at a U.S. port. Second, we match each of these voyages to the universe of U.S. import bill of lading (BoL) filings. These documents, collected by U.S. Customs for each incoming shipment, contain key details such as the name of the receiving company, the Harmonized System (HS) code of imported products, and the shipment's volume in twenty-foot equivalent units (TEUs), allowing us to assign a delay duration to each shipment and aggregate these into a measure of delay exposure for each firm and product category, week by week. Third, we identify the barcodes (UPCs) of the products sold by importing firms and obtain pricing information for these products from the NielsenIQ Retail Scanner (RMS).[1] The RMS dataset is one of the most comprehensive retail transaction datasets available, covering more than 90,000 stores across all major U.S. retail chains and recording the actual transaction price at the UPC level on a weekly basis. Fourth, we obtain product-level consumption information from the NielsenIQ Consumer Panel (HMS). The HMS dataset is a representative longitudinal panel that tracks the complete purchase histories of approximately 60,000 U.S. households across all product categories and retail outlets. Fifth, a subset of HMS households participates in a quarterly survey of expected inflation over the following 12 months, following the methodology of D'Acunto, Malmendier, Ospina, and Weber (2021). This allows us to link each household's realized supply chain shock exposure directly to its forward-looking inflation beliefs. The resulting 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 in a single integrated framework.

We present three main findings. First, we document that shipping delays raise retail prices and reduce product availability. We begin with the price of a good that stays on the shelf, which isolates repricing from any change in what is available. We measure a firm's delay exposure as the share of its import volume over the trailing 13 weeks that arrives on shipments running more than a week behind the time the same route normally takes, and we follow the price of the same product within the same firm from week to week, asking whether it increases when the firm's import pipeline

[1] Throughout the paper, a UPC (Universal Product Code) is an individual barcoded product, and we use the terms UPC and product interchangeably. A product market is what NielsenIQ calls a product module, the group of interchangeable products that compete with one another, such as shampoo. Exhibits and variable definitions use the NielsenIQ term product module. Our product markets are national and abstract from geography. Uniform pricing within retail chains (DellaVigna and Gentzkow, 2019) limits the price variation this abstraction discards, although local product availability may differ across markets. Section 2 describes both classifications.

becomes delayed relative to less-exposed firms selling in the same product market and week. An increase of one standard deviation in delay exposure raises the price of a continuing product by roughly 0.6 percent. Because the comparison holds the product, firm, product market, and week fixed, the estimate is difficult to attribute to demand shocks common to the product market.

The price of a continuing product is only one margin. What a shopper pays also depends on which goods are on the shelf, so we ask whether delays reduce product availability and change the mix of products that remain. Delays raise the chance that a product is absent from stores, reduce the number of distinct products a firm offers in a category, and shift the mix of available products toward premium items. As a result, a one standard deviation increase in delay exposure raises the average price of the products a firm keeps available by 2.6%, against 0.6% for the identical continuing product. That gap is the central quantitative fact of the paper. Because the two estimates share the same treatment and sample family and differ only in whether the set of observed products is held fixed, the difference between them implies that most of the consumer-facing price response operates through changes in which products are available rather than through the repricing of products that remain. A price index built from continuing goods, which is how price statistics are normally constructed, therefore captures roughly a fifth of what a shopper faces.

What makes a delayed shipment raise prices is not an obvious cost increase, since a late shipment need not arrive with a higher invoice cost. Delays are instead likely to work through inventory and supply uncertainty, by drawing down inventory buffers, raising restocking risk, and increasing the option value of holding inventory. We test this scarcity mechanism using a predetermined measure of total inventory relative to cost of goods sold. Firms that enter a disruption with thin inventory buffers raise the prices of continuing products more when delay exposure increases. Delays across other firms in the same product market are associated with reduced availability and a shift in the remaining mix toward premium varieties. Products without delayed shipments of their own do not increase in price when rivals are delayed and, if anything, decrease. These patterns fit a scarcity channel rather than the direct cost pass-through emphasized for the tariff and exchange rate shocks that the literature typically studies (e.g., Gopinath and Itskhoki, 2010; Amiti, Itskhoki, and Konings, 2019; Amiti, Redding, and Weinstein, 2019; Fajgelbaum, Goldberg, Kennedy, and

Khandelwal, 2020; Cavallo, Gopinath, Neiman, and Tang, 2021). They also speak to how retail prices adjust in the first place. The micro pricing literature documents that individual prices change infrequently and often revert to reference levels (e.g., Bils and Klenow, 2004; Nakamura and Steinsson, 2008; Eichenbaum, Jaimovich, and Rebelo, 2011), and that chains set prices with limited variation across their stores (DellaVigna and Gentzkow, 2019). Both features imply that a supply disruption need not show up promptly in the posted price of a continuing good, which is consistent with our finding that the larger consumer-facing response runs through which products are available. For the 2018–2019 tariffs, Bai, Jaccard, and Stumpner (2026) document the same pattern: pass-through into retail prices operates largely through product replacement rather than through the repricing of continuing barcodes. The identifying variation comes from the difference between each voyage's realized duration and its route benchmark. This difference is determined after the shipment departs and is unlikely to reflect the firm's retail pricing decisions. The panel regressions also compare firms selling in the same product market and week. To address the main concern that COVID-19 simultaneously disrupted shipping and retail demand, we reestimate the baseline outside the pandemic period. The product module×week estimate is 0.024 excluding 2020–2021 and 0.030 in the 2015–2019 sample alone, both statistically significant at the 1% level. The estimates are also similar when we measure delay exposure over shorter windows.

To complement the panel estimates, we implement a stacked difference-in-differences design that exploits the staggered timing of the first delay event across the individual products firms sell. The event-time design gives each delay a clear onset date, allowing us to examine whether price changes precede the disruption and when the response emerges. We match treated products to not-yet-treated controls on their pre-delay prices and units sold, and estimate difference-in-differences regressions over a window of twelve weeks before and after each delay. Prices exhibit no trend before the delay and remain unchanged while the shipment is still at sea. They begin to increase only when the delayed shipment should have arrived and the missed replenishment begins to bind. Delay events raise average post-delay retail prices by 0.4 to 0.5 percent relative to matched controls, with dynamic effects increasing to roughly 1.1 percent later in the event window. In longer-window versions of the design, average post-delay prices remain about 1.0 percent higher a full year after

the delay. These estimates indicate that delay-induced price pressure persists beyond the initial disruption rather than reversing quickly.

Second, we study how these product-level supply shocks affect household consumption, the step from firm-level pass-through to the household basket. Because households hold different baskets, the same disruptions imply different household exposures. Our analysis rests on three objects: the baseline bundle a household usually buys, the household's delay exposure through the product categories in that bundle, and the outcomes of what the bundle costs and what the household actually buys. We begin by fixing each household's bundle using purchases in one calendar year and tracking its price in the next, so any price change reflects repricing of the goods the household already buys. An increase of one standard deviation in household delay exposure raises the price of the fixed bundle by about 0.4 percent. An arrival-timed stacked event study yields average post-delay effects of 0.36 to 0.42 percent, with the price response building over the following twelve weeks. Together, these estimates show that the retail price response affects households through higher prices for the goods they were already buying.

The fixed bundle rules out any adjustment. We therefore let the bundle roll forward with the household's own purchases and capture how the basket responds. An increase of one standard deviation in bundle exposure raises the household's bundle price index by about 1.9 percent. Households obtain some relief by shifting spending toward the goods whose prices increase least, so the price they actually pay increases by less than the price of a fixed basket would. These household estimates are similar when we control for how import-intensive the household's bundle is, so they reflect exposure to delays rather than a generally import-heavy basket.

We next ask what households change inside their bundles. They do not exit the product markets exposed to delays. Instead, they buy slightly more physical quantity while spending slightly less, and they tilt spending toward cheaper products and toward the specific items and brands they usually buy. The adjustment is concentrated within product markets rather than across them, so households absorb the shock by reallocating inside their familiar categories.

Finally, we separate the sources of the household price increase. Almost the entire response reflects higher prices for goods that remain available, while voluntary substitution and the forced

replacement of stocked-out goods roughly offset each other. An increase of one standard deviation in household delay exposure raises the share of the basket exposed to stockouts by 5.0 percentage points and the share of spending on products the household has not bought before by 2.1 percentage points. The availability margin is thus large relative to the repricing of goods that stay on the shelf.

These results indicate that the burden of supply chain disruptions extends beyond posted prices. Part of it takes the form of buying different goods rather than paying more for the same ones, a margin that standard inflation measures, which track the prices of available goods, capture only partly. The consumption results also provide the household-specific bundle weights used to measure exposure in the expectations survey, so the expectations analysis follows disruptions through product modules in each household's observed bundle.

Third, and more tentatively, we ask whether households with more delay-exposed bundles report different inflation expectations. Comparing households within the same survey week, a one standard deviation increase in exposure is associated with 0.40 percentage points higher expected inflation in our fully controlled specification.

This cross-sectional comparison cannot separate what a household experiences from the kind of bundle it holds, because households choose their bundles and placebo reassignments show that much of the association is shared by compositionally similar baskets. Section 5 reports that evidence in full.

We therefore compare each household with itself over time, using the pandemic disruption, which raised delay exposure far more for some bundles than for others. Households whose exposure rose most were also the ones that revised their expectations upward: a one standard deviation increase in the exposure change is associated with a 2.4 percentage point higher probability of an upward revision, and the estimate is essentially unchanged when each household's bundle is held fixed at its pre-pandemic purchases. Exposure during the pandemic drives the association, while pre-pandemic exposure is unrelated to the revision. We read this as suggestive of experienced delays shifting beliefs rather than as a causal estimate of that effect.

Overall, we assemble a dataset that follows a supply shock along its entire path, from individual vessels through importing firms and retail prices to household consumption and inflation beliefs,

across a 2010–2024 sample spanning the period before, during, and after the COVID-19 pandemic. Shipping delays are one instance of such a shock, and the same linkage, from a measurable trade disruption to the firms, products, and households it affects, can trace the household incidence of other supply shocks transmitted through imports. These findings provide micro-level evidence that the source of a price change, and not only its magnitude, matters for how households respond in what they buy. A supply-side shock does more than raise posted prices: it changes which goods are on the shelf. The price of the products a store keeps available increases by more than four times the same-product response, and rather than pay that increase in full, households substitute toward the goods that stay cheap and familiar. The resulting change in what households buy, not only what they pay, is a margin that standard inflation measures largely miss. The expectations evidence points the same way, with more caution. Much of the cross-sectional association between exposure and expected inflation is shared by compositionally similar placebo bundles, while households whose exposure rises more also tend to revise their expectations upward, including when their bundle is held fixed, suggesting that beliefs may respond as well.

The availability margin carries a policy implication. Because supply shocks change which products households can buy as well as the prices of the products that remain available, posted-price measures capture only part of their household burden. A price statistic constructed from continuing goods is, by construction, blind to the margin that accounts for most of the response we measure.

We contribute to at least three strands of the literature. First, we document that supply chain disruptions impose real costs on households beyond posted prices. Exposed households shift spending toward cheaper and more familiar products within affected product markets, and stockouts push part of their bundles into replacement purchases, a within-basket adjustment margin in the spirit of Hastings and Shapiro (2013) and of the trading-down response documented over the business cycle by Jaimovich, Rebelo, and Wong (2019). These results add a supply-chain shock to the literature that measures the cost of living and household adjustment in micro data (e.g., Kaplan and Schulhofer-Wohl, 2017; Jaravel, 2019; Redding and Weinstein, 2020; Argente and Lee, 2021; Michelacci, Paciello, and Pozzi, 2022; Atkin, Faber, Fally, and Gonzalez-Navarro, 2024; Baqaee, Burstein, and Koike-Mori, 2024; Jaravel and Lashkari, 2024) and to the evidence on the unequal

household incidence of price and trade shocks (e.g., Cravino and Levchenko, 2017; Adão, Carrillo, Costinot, Donaldson, and Pomeranz, 2022; Auer, Burstein, Lein, and Vogel, 2024; Colicev, Hoste, and Konings, 2024; Canto, Grigsby, Qian, and Walsh, 2025; Chen, Levell, and O'Connell, 2026). The closest paper is Hoke and Feler (2026), who study household spending responses to the 2025 tariffs. We differ in studying a physical supply disruption rather than a tariff and in tracing the shock from individual vessels to firms, products, and households.

Second, we examine whether exposure to supply-side disruptions is associated with household inflation expectations. A well-established literature documents that the prices households personally encounter shape their inflation beliefs (e.g., Cavallo, Cruces, and Perez-Truglia, 2017; D'Acunto et al., 2021), but the identifying variation in these studies does not distinguish the source of the price change. We exploit a well-measured, supply-driven shock and find that households with more exposed baskets report higher inflation expectations, and that households whose exposure rises more tend to revise their expectations upward. Placebo comparisons show that much of the cross-sectional association is shared by compositionally similar baskets, so we rely on the within-household comparison, which we read as suggestive rather than causal. The closest related evidence is Albagli, Grigoli, and Luttini (2025), in which firms extrapolate from the prices of their suppliers. We study the household counterpart, in which measured exposure runs through the consumption bundle.

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 (2026), 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 delays, a sample spanning 2010 through 2024 that extends beyond the pandemic, and the addition of household consumption and expectations dimensions. Related work links port congestion to consumer price inflation (Jiao, Lan, Liu, and Zhou, 2026), uses the bankruptcy of Hanjin Shipping as a shipping-disruption shock (Kwon, 2021), and studies how firms build resilience to such disruptions (Sanz, 2026). Fleet concentration itself can be a source of vulnerability: Fuchs, Leibovici, and Wong (2026) study countries' dependence on foreign-built and foreign-operated containerships and the exposure to shipping

disruptions this dependence creates. Our results speak to a long-standing question about timing frictions in international trade (e.g., Djankov, Freund, and Pham, 2010; Hummels and Schaur, 2013) and suggest that the price consequences of supply chain disruptions extend beyond the direct cost channel emphasized in aggregate analyses (e.g., Carrière-Swallow, Deb, Furceri, Jiménez, and Ostry, 2023; Bai, Fernández-Villaverde, Li, and Zanetti, 2024; Känzig and Raghavan, 2026).

2 Data

We create a unique dataset that links vessel-level shipping delays to barcode-level retail prices through 2024 and to household purchase decisions and inflation expectations through 2023. We combine five data sources: satellite-based vessel tracking data from Spire Maritime, bill of lading (BoL) import filings from Panjiva, retail prices from the NielsenIQ Retail Scanner (RMS) data, household purchase records from the NielsenIQ Consumer Panel (HMS), and a quarterly survey of expected inflation over the following 12 months answered by a subset of HMS households. The RMS and HMS data come from the Kilts Center at the University of Chicago, and the survey provides the belief data for the expectations analysis.

2.1 Shipping Delays and Supply Chain Disruptions

To measure supply chain disruptions, we use satellite-based Automatic Identification System (AIS) data from Spire Maritime, which track the position of container vessels worldwide, at sea and in port (e.g., Brancaccio, Kalouptsidi, and Papageorgiou, 2020; Heiland, Moxnes, Ulltveit-Moe, and Zi, 2025). The data record each voyage's foreign departure port, departure time, U.S. arrival port, arrival time, and route. Unlike aggregate shipping indices such as the Baltic Dry Index or the Global Supply Chain Pressure Index, the AIS data measure realized transit time for every trip at the origin–destination port-pair level.

Motivated by evidence that each additional day in transit imposes economically meaningful trade costs (e.g., Evans and Harrigan, 2005; Hummels and Schaur, 2013), we measure disruptions as excess transit time relative to normal conditions on the same route. We define a voyage's delay as realized transit time minus the 2015–2017 median transit time for the same origin–destination

port pair (the route pair) and classify a voyage as delayed when this difference exceeds seven days. This route-specific benchmark nets out persistent differences in distance, geography, and routing across port pairs, so that our measure captures unusually slow voyages.[2]

Figure 1 shows the foreign origins, shipping corridors, and U.S. arrival ports represented in the matched shipping data. Panel A shows dense transpacific connections between East Asia and West Coast ports, together with transatlantic connections between Europe and the Mediterranean and the East Coast. Additional routes originate in South and Southeast Asia and the Americas. This pattern is consistent with Pacific ports serving as the principal gateways for Asian containerized imports, while Atlantic ports connect the United States to a more diversified set of trading regions. Panel B shows that arrivals are concentrated at major ports. For example, Los Angeles–Long Beach and New York–Newark receive 25.8% and 18.2% of matched voyages arriving at contiguous U.S. ports, respectively, and together account for 44.1%. Overall, the figure shows that our data captures the central routes and ports of U.S. container trade. Figure 2 plots the evolution of three measures of shipping conditions constructed from the AIS data. Panel A plots the mean deviation of voyage duration from the 2015–2017 port-pair benchmarks for all U.S.-bound container voyages. Deviations hover near zero through 2019 and increase sharply from mid-2020 as pandemic-related port congestion and container shortages lengthen voyages. The March 2021 blockage of the Suez Canal by the *Ever Given* occurs during this broader deterioration in shipping conditions. Deviations peak at approximately 28 days above the port-pair median in early 2022 and return toward baseline through 2022–2023. They begin increasing again in late 2023 as drought-related restrictions reduce traffic through the Panama Canal. Houthi attacks on commercial shipping in the Red Sea, beginning in November 2023, create an additional disruption as carriers reroute around the Cape of Good Hope.[3] In Panel B, the fraction of U.S. bound voyages classified as delayed mirrors this pattern on the extensive margin, increasing from approximately 5–10% in 2018–2019 to over 70% at the peak of the supply chain crisis in early 2022. Panel C plots mean anchorage time near the ports

[2] We require at least 10 voyages per port pair over 2015–2017 to form the benchmark and drop voyages whose transit time exceeds three times the port-pair median as mismatches.

[3] The drought reduced daily Panama Canal vessel transits from 36 to as few as 18 between July 2023 and August 2024 (Pankratz, 2026). The attacks, beginning in November 2023, forced major carriers to suspend Suez Canal transits and reroute around the Cape of Good Hope, increasing average delivery times by 10 days or more (Kamali, Koepke, Sozzi, and Verschuur, 2024).

of Los Angeles and Long Beach. Anchorage time was close to zero through 2019, increased to approximately 235 hours, or nearly 10 days, in early 2022, and returned to near zero by mid-2023.

Figure 3 examines the voyage-level variation underlying these aggregate disruption episodes. Because normal transit times differ across routes, we define a voyage's delay as its realized duration minus the median duration of voyages between the same origin and destination ports during 2015–2017. In Panel A, zero means that a voyage took the usual amount of time for its route, while positive values report additional days beyond that route's typical duration. The median voyage lasts only 0.3 days longer than its route benchmark, while 12.7% of voyages exceed the benchmark by more than seven days and are classified as delayed. This comparison captures unusually slow realizations within a route rather than mechanically selecting routes with longer normal transit times. Panel B shows that the resulting disruptions are economically meaningful. Conditional on being classified as delayed, a voyage lasts 12.6 days longer than its route benchmark at the median and 16.2 days longer on average, while 10% of delayed voyages exceed the benchmark by more than 29 days.

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 (Flaaen, Haberkorn, Lewis, Monken, Pierce, Rhodes, and Yi, 2023). 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 data from Panjiva.[4]

The receiving firm on a bill of lading is the importer of record, and in our matched sample it is almost always the owner of the brand rather than the retailer that sells it. Classifying the

[4] Firms may request that U.S. Customs and Border Protection withhold their identity from the publicly disseminated bill of lading records, so we observe the subset of importers that do not request confidential treatment (e.g., Flaaen et al., 2023; Pankratz, 2026; Sanz, 2026; Schafhäutle and Sran, 2026). Requests are granted for two years before requiring renewal, and the share of U.S. import records with a missing consignee identifier ranges from 9.8% to 36.6% between 2007 and 2021 (Flaaen et al., 2023). Schafhäutle and Sran (2026) document systematic variation in redaction across time, origin countries, and goods, and show that redaction becomes more common as forced labor scrutiny intensifies, so the importers we observe are a selected sample. A request applies to a firm name rather than to individual shipments, however, so within an observed firm the records we see are not selected on whether a particular voyage ran behind its route benchmark, and the resulting measurement error in delay exposure attenuates our estimates toward zero.

314 matched firms by their primary SIC code, 79 percent are manufacturers, and they account for 98 percent of matched import volume; retailers and wholesalers together represent 11 percent of firms but less than 1 percent of volume.[5] The distinction matters because we observe prices at the retail shelf, one vertical layer below the firm whose imports are delayed. Brand owners, however, hold direct authority over shelf prices: they place floors under retail prices through minimum-advertised-price and unilateral resale-price policies, condition promotional funding and rebates on price compliance, and discipline deviations through chargebacks and termination (Chen, Sanz, and Wittry, 2026). In that evidence, more than three-quarters of brand owners disclose arrangements in their 10-K filings that tie their payoffs to their products' retail sales, fewer than 2 percent describe a wholesale-only relationship, and manufacturer identity accounts for 91 percent of the systematic variation in retail price levels and 73 percent of the systematic timing of price changes, against 22 percent for retailers. In the minority of cases in which the importer is itself a retailer, the importer and the price-setter simply coincide. In either case, the firm bearing the delay holds authority over the price we observe.

We link each shipment to its AIS voyage with a crosswalk from the International Maritime Organization (IMO) number in the BoL filings to the Maritime Mobile Service Identity (MMSI) number in the AIS data. The crosswalk is needed because the IMO number is a permanent vessel identifier, whereas the MMSI can change when a vessel is re-flagged or re-registered.

To link import shipments to retail products, we map NielsenIQ product modules to HS codes. NielsenIQ classifies all UPCs into roughly 1,000 product modules based on product descriptions. A UPC identifies a specific barcoded product, such as a 12-ounce bottle of a particular brand shampoo, while a product module collects the interchangeable products that compete with it, such as shampoo from other brands. We treat the product module as the relevant product market and use the two terms interchangeably. We build the mapping with a fuzzy matching procedure based on a large language model that draws on the textual descriptions of both taxonomies. We use 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 BoL shipments for a given

[5] The remaining firms are non-manufacturing brand owners, such as consumer-technology and media firms importing hardware and branded merchandise. Industry classifications are current Compustat header codes.

HS code to its NielsenIQ products in the corresponding product module.

After linking each shipment to a firm and product module, we aggregate voyage-level delays into a firm×product module×week measure of delay exposure. *Fraction in delay (13w)* is the TEU-weighted share of the firm's import volume in the product module carried on delayed voyages over the trailing 13 weeks.[6] For example, if half of the firm's import volume in a product module was carried on delayed voyages, *Fraction in delay (13w)* equals 0.5. For the retail analyses, we date delayed volume to the vessel's departure week, reflecting firms' earlier information about delay conditions, and assign the resulting measure to the firm's UPCs in the corresponding module. For the household consumption and inflation-expectations analyses, we instead date delayed volume to the vessel's arrival week, when households encounter and learn about the disruption.[7]

2.3 Retail Prices

Retail price data come from the NielsenIQ RMS database, which records weekly transaction prices at the UPC level from more than 90,000 stores across all major U.S. retail chains (e.g., Jaravel, 2019). The recorded price is the volume-weighted average price at which the UPC actually sold during the week, so temporary price reductions and retailer discounts are reflected in the measure, while manufacturer coupons are not. We deliberately use this transaction price rather than a regular-price series that strips out promotions, both because it is the price consumers actually pay and because firms can respond to a delay by withdrawing promotional support rather than raising list prices, a margin a regular-price series would miss. The linked retail sample covers 2015 through 2024, the period over which scanner prices can be matched to delay exposure. We aggregate to the firm×UPC×week level by summing units sold and computing the average price per item across stores, weighted by unit sales. Because retail chains set prices largely uniformly

[6] The 13-week window corresponds to one quarter of import activity, long enough to smooth week-to-week lumpiness in shipment arrivals while remaining short enough to capture the disruption a firm currently faces. Estimates using a weekly flow measure and a 4-week rolling window are similar, as Table IA.4 in the Internet Appendix shows.

[7] We do not observe the time from port arrival to store shelf. Data from the Pacific Merchant Shipping Association indicate that the port leg of this lag is measured in days: import containers dwelled at the terminals of Los Angeles and Long Beach for two to three days on average before the pandemic, and truck-bound containers moved inland within roughly eight days even at the height of the congestion (Schuler, 2023). The event-study timing in Section 3.7 aligns with a short lag, as the price response begins in the week the average delayed shipment arrives and peaks three weeks later.

across their stores (DellaVigna and Gentzkow, 2019), averaging across stores discards little of the price variation relevant for our analysis. Our main outcome, *ln(Price)*, is the natural log of this average unit price.

2.4 Household Consumption

We obtain data on household purchases from HMS, a nationally representative panel that tracks the complete purchase histories of approximately 60,000 U.S. households across product categories and retail outlets from 2010 through 2023 (e.g., Broda and Weinstein, 2010; Kaplan and Schulhofer-Wohl, 2017). Each purchase record identifies the household, UPC, purchase date, quantity, and total price paid. We aggregate these purchases to the household×UPC×week level.

We use purchases observed before the outcome period to define the products each household normally buys and their baseline expenditure weights. These baseline bundles allow us to measure changes in bundle cost, spending patterns, and product availability using current prices and purchases. We measure each household's exposure to shipping delays by averaging delay rates across the product modules in its bundle, weighting each module by its share of baseline spending. Depending on the analysis, the baseline is either fixed over the following calendar year or updated using a rolling window. We describe the corresponding construction with each set of results. For the household consumption analyses, delayed import volume is assigned to the vessel arrival week before constructing exposure.

2.5 Inflation Expectations

Our expectations data come from a customized survey fielded to a subset of HMS panel members, described in D'Acunto et al. (2021), Coibion, Gorodnichenko, and Weber (2022), Coibion, Georgarakos, Gorodnichenko, and Weber (2023), and Weber et al. (2025). The survey is not a NielsenIQ product. It is commissioned and paid for by the researchers on a per-respondent basis and administered quarterly through the HMS platform, which allows each respondent's answers to be linked to their recorded purchase history. Respondents report the rate at which they expect prices to change over the following 12 months. Waves in 2015 and 2016 randomized the elicitation format

between the prices of the items on which the household spends money and the overall inflation rate; since 2018 the survey has consistently elicited expected inflation. The cross-sectional analysis uses 17 survey waves with direct numerical forecasts from April 2018 through December 2023. The long-difference analysis also uses three pre-pandemic surveys beginning in September 2018, June 2019, and September 2019. These surveys elicit subjective probability distributions across future inflation ranges, from which we calculate the implied mean forecast. The broader questionnaire also asks about perceived inflation over the previous 12 months.

We link each household survey observation to the household's rolling consumption bundle. We measure delay exposure over the 13 weeks ending on the Saturday before the survey. The measure averages delay rates based on vessel arrivals across the product modules in the bundle, weighting each module by its share of household spending. Modules with no observed import arrivals during the window contribute zero to the household's measured delay exposure. We restrict the sample to observations for which modules with observed import arrivals account for at least 50 percent of bundle spending.[8]

2.6 Descriptive Statistics

Table 1 summarizes the samples used in the firm, household consumption, and household expectations analyses. Panel A describes the 18.7 million firm×UPC×week observations used in the retail analysis. On average, 10 percent of a firm's import volume within a product module was carried on delayed voyages during the preceding 13 weeks. Products are absent in 35 percent of observations. The inventory-buffer sample contains 11.0 million observations, 41 percent of which belong to firms with low inventory buffers. The average product price is $16.77 among the 12.1 million observations with observed sales.

Panel B summarizes the 35.3 million household×week observations used to measure bundle

[8] Because the expectations analysis compares households, bundle observability is a potential concern. Households with more observable bundles can mechanically have higher measured delay exposure. Products without measurable exposure are concentrated in produce, dairy products, eggs, and fuel, categories often sourced domestically or imported outside the containerized shipping channel we measure. We nevertheless require at least 50 percent coverage so that differences across households reflect shipping delays rather than differences in our ability to measure exposure. The household panel regressions instead include household fixed effects, which absorb persistent differences in bundle observability. Table IA.10 reports the composition of the products without measurable exposure.

inflation and reallocation. The baseline bundle contains 256 UPCs and represents $3,759 of spending over the 52-week baseline window, on average. The spending-weighted delay exposure across the product modules in a household's baseline bundle averages 7 percent. Delay exposure can be measured for modules representing 63 percent of baseline bundle spending. UPCs that are not observed in the outcome week account for 42 percent of baseline bundle spending on average.

Panel C describes the sample used across the expectations regressions in Table 14. The sample contains 116,009 household survey observations from 36,871 households across 17 survey waves through December 2023. Expected inflation over the following 12 months is dispersed and right skewed. The mean is 8.30 percent, the median is 5 percent, and the interquartile range is 2 to 10 percent. Delay exposure coverage is the share of baseline bundle spending represented by product modules with measurable delay exposure. By construction, coverage is at least 50 percent in every observation and averages 70 percent in the sample. The spending-weighted delay exposure across modules averages 12 percent. The average respondent is 59 years old, 73 percent of respondents are women, and children are present in 14 percent of households. The corresponding bundles contain 280 UPCs and represent $4,127 of spending over the 52-week baseline window, on average. Throughout the paper, we scale delay-exposure coefficients using the rounded standard deviation of *Fraction in delay (13w)* reported in the corresponding panel: 0.20 for the retail analysis, 0.08 for the household analysis, and 0.10 for the expectations analysis. The long-difference specification instead uses the standard deviation of the exposure change reported in its caption.

3 Shipping Delays, Retail Prices, and Product Availability

Shipping delays can raise the cost of a household's shopping basket through two margins. The first is the average transaction price of a continuing UPC. When a shipment arrives late, the firm has less inventory available today, and inventory constraints may bind before the next replenishment arrives. The second is product availability. If some UPCs disappear from shelves or the mix of products that remain changes toward higher-priced goods, the average price of the products available to consumers increases even when the prices of continuing UPCs change little. We therefore examine separately the effects of shipping delays on within-UPC prices, product availability, and the mix of

available products. We then test whether thin inventory buffers amplify direct repricing, whether delays shared across firms generate a product market scarcity margin, and whether the price effect extends beyond directly affected products and persists after individual delay events.

3.1 Empirical Specification

We first estimate the following specification:

$$Y_{ipt} = \alpha_{ip} + \delta_{mt} + \beta \, \textit{Fraction in delay (13w)}_{imt} + \varepsilon_{ipt}, \tag{1}$$

where i indexes firms, p UPCs, m product modules, and t weeks. The dependent variable Y_{ipt} is one of five outcomes. The two price outcomes are the log retail price of continuing products and the log retail price of available products. For the second outcome, the set of observed UPCs may change. The remaining outcomes are an indicator for a UPC being absent, the log number of active UPCs, and the share of premium UPCs. We define each outcome in the subsection that studies it. For outcomes measured at the firm×product module×week level, the unit subscript p is replaced by m. $\textit{Fraction in delay (13w)}_{imt}$ is the share of firm i's import volume in product module m carried by delayed shipments over the preceding 13 weeks.

Identification comes from realized transit-time overruns rather than persistent differences in the routes or products firms import. We classify a voyage as delayed when its realized duration exceeds the 2015–2017 median duration for the same origin–destination port pair by more than seven days. This route benchmark removes persistent differences in normal transit times. The remaining variation reflects whether a shipment takes unusually long after departure and is unlikely to be determined by the importing firm's retail pricing decisions. The fixed effects described below further absorb persistent differences in the products firms sell and contemporaneous shocks common to firms in the same product market and week.

Specifications that follow the same UPC over time include firm×UPC fixed effects, which compare a product to itself within the same firm. Specifications that allow the set of observed UPCs to change include firm×product module fixed effects. These fixed effects absorb persistent differences in the number and types of products a firm offers in a product market, so identification comes from

changes within the same firm and product market over time.

The time fixed effects depend on the source of variation used in each test. Product module×week fixed effects compare firms within the same product market and calendar week and absorb contemporaneous shocks common to that market. We use them when exposure varies across firms within the same product module and week. Product module×delay episode fixed effects instead absorb product market changes during major disruption episodes while preserving weekly variation within episodes. We use them when the test requires weekly variation within an episode or when product module×week fixed effects would absorb the variation of interest. Because these specifications compare firms within the same product market, the estimates capture relative price responses across firms. Any component of a disruption common to all firms in a market, including market-wide price increases and spillovers from exposed to unexposed firms, is absorbed by the time fixed effects, so our estimates understate the aggregate price response to the extent disruptions raise all prices in a market. Standard errors are clustered at the firm×product module level. The coefficient of interest, β, captures the relationship between delayed import exposure and each outcome. If delayed replenishment makes inventory scarce before restocking, we expect β to be positive for prices, absences, and the premium share, and negative for the number of active UPCs.

3.2 Direct Price Effects on Continuing Products

We begin with the first margin, which captures changes in the average transaction prices of continuing products. The outcome is *ln(Price)*, the natural log of the average price per item across stores, weighted by unit sales. Because the comparison is within the same firm×UPC, the estimate removes permanent differences in product quality, brand position, and firm pricing policy, and captures how the average transaction price of an available product changes when the firm's import pipeline for its product module becomes delayed.

Table 2 reports estimates of Eq. 1. All three columns include firm×UPC fixed effects. Column (1) also includes indicators for the six major supply chain disruption episodes listed in the table caption. These indicators absorb price changes that are common across product markets during each disruption period. The coefficient on *Fraction in delay (13w)* is 0.028 and statistically significant

at the 1% level. Column (2) instead includes product module×delay episode fixed effects, allowing price changes during the disruption periods to differ across product markets. The estimate is 0.029. Column (3) includes product module×week fixed effects, which absorb shocks shared by firms selling in the same product market and week. The estimate is 0.030. Hence, an increase of one standard deviation in delayed import exposure raises average transaction prices for the same UPC within the same firm by roughly 0.6%.

Table IA.2 in the Internet Appendix reports the direct price estimates using different time fixed effects. The coefficient is 0.031 with only firm×UPC fixed effects and 0.028 after adding delay episode indicators. Allowing time effects to vary across product markets yields a coefficient of 0.029 with product module×delay episode fixed effects and 0.030 with product module×week fixed effects. The last specification compares firms with different delayed import exposure selling in the same product market and week and absorbs price changes common to that market. The stable estimates across specifications are difficult to reconcile with a demand shock common to all firms in a product module. Table IA.3 in the Internet Appendix shows that this relationship is not specific to the pandemic years. The product module×week estimate is 0.024 excluding 2020–2021 and 0.030 in the 2015–2019 sample alone, both statistically significant at the 1% level. The stability of the estimate indicates that the baseline result is not driven by COVID-era changes in retail demand.

3.3 Product Availability

We next turn to product availability. The price consumers face also depends on which UPCs appear in stores (Handbury, 2021). When inventories tighten, lower-price varieties may stock out first and firms may allocate scarce units toward higher-margin products. The average price of the products that remain available can therefore increase even if continuing UPCs are repriced only modestly. Product availability moved sharply during recent disruptions and carries information about inflation pressure (Cavallo and Kryvtsov, 2023).

We estimate Eq. 1 for the prices of products that remain available and three availability outcomes. The price specification keeps *ln(Price)* as the outcome but replaces firm×UPC with firm×product module fixed effects. It therefore compares the prices of the firm's observed UPCs

within a product market without holding the UPC fixed, allowing changes in the set of available products to contribute to the estimate. *Absent* equals one when a UPC does not appear in NielsenIQ in a given week, measured in a balanced firm×UPC×week panel over each UPC's observed lifetime. The log number of active UPCs and the share of premium UPCs are measured at the firm×product module×week level. A UPC is premium when its average log price over the sample lies above the median across the firm's own products in the product module, so the premium share is a fixed classification that moves only with the set of available products. These availability outcomes can respond to supply chain disruptions shared across firms and product markets, not only to relative differences in firm exposure. Week fixed effects would absorb the aggregate component of these shortages, while product module×week fixed effects would additionally absorb scarcity shared within a product market. We therefore include product module×delay-episode fixed effects, which absorb broad differences across product markets during each disruption while preserving the weekly variation through which shortages develop. Table 5 provides direct evidence that this common component is economically important for availability specifically: delay exposure among the *other* firms in a product market, constructed by excluding the focal firm and weighting the remaining firms by fixed 2015–2017 sales shares, is associated with a 3.28 percentage point increase in UPC absence and 2.66% fewer active UPCs, magnitudes comparable to or larger than the own-exposure estimates reported here. Because that variation is shared within a product market and week, week and product module×week fixed effects absorb it.

Table 3 reports the results. In Column (1), the coefficient on *Fraction in delay (13w)* is 0.131, statistically significant at the 1% level, so an increase of one standard deviation in delayed import exposure is associated with 2.6% higher prices among the UPCs the firm has available in that product market. The estimate can be compared directly with the same-UPC estimate of 0.030 in Table 2: the two regressions share the same treatment and sample family, and the regression in Column (1) additionally lets the set of observed UPCs change. The gap between the two estimates implies that most of the consumer-facing price response operates through changes in which products are available rather than through repricing of continuing products.

Columns (2) to (4) report the availability outcomes. In Column (2), an increase of one standard

deviation in *Fraction in delay (13w)* raises the probability that a UPC is absent by 2.0 percentage points. In Column (3), the same increase reduces the number of UPCs the firm actively offers in a product market by 1.2%. In Column (4), the premium share of available UPCs increases by 1.1 percentage points. All three are statistically significant at the 1% level. Delays thus remove products from the choice set, reduce the number of available products, and shift the products that remain toward more expensive varieties. Together, the estimates indicate that shipping delays affect consumers through both the prices of continuing UPCs and the availability and mix of products on the shelf.[9]

3.4 Inventory Buffers and Firm Pricing

We now return to the direct price response and examine whether it reflects inventory scarcity. A delay can increase prices even if the goods in transit cost no more than they otherwise would. By postponing replenishment, the delay temporarily reduces the inventory available for sale. This creates intertemporal scarcity and increases the value of conserving the remaining stock (e.g., Alessandria, Kaboski, and Midrigan, 2010; Kryvtsov and Midrigan, 2013; Alessandria, Khan, Khederlarian, Mix, and Ruhl, 2023).

We assign delay exposure to the importing firm, typically the brand owner, but observe the price posted at the retail shelf. The pricing decision that connects the two need not wait on retailer-by-retailer wholesale pass-through. As Section 2.2 describes, brand owners exercise direct authority over shelf prices through resale-price floors and compliance-linked promotional funding, and the timing of retail price changes is largely a manufacturer-level event: manufacturer effects account for roughly three-quarters of the systematic variation in when prices change, against about one-fifth for retailers, so a brand's repricing propagates across retailers at once (Chen, Sanz, and Wittry, 2026). A firm short of inventory after a delayed shipment can therefore raise the shelf prices of its products directly, by lifting its floor and withdrawing promotional support, rather than only

[9] The shift toward premium varieties need not reflect a deliberate allocation choice by firms. If lower-priced varieties sell faster, they stock out first when replenishment pauses, and the mix remaining on the shelf tilts toward premium items even without any reallocation of scarce units. Both readings carry the same implication for consumers, namely that the average price of the available products increases, so we read the premium share as describing the choice set households face rather than as evidence on how firms allocate inventory.

through a wholesale price that retailers may or may not pass through. Firms with less inventory relative to cost of goods sold are less able to absorb a shipping delay, so a given increase in delay exposure should lead to a larger price increase.

We measure the inventory buffer using days of total inventory coverage. We construct this measure from Compustat by dividing quarterly total inventory by quarterly cost of goods sold and multiplying the ratio by 91.25, the average number of days in a quarter. We assign zero total inventory when quarterly total inventory is missing and cost of goods sold is positive.[10] The measure uses the latest fiscal quarter ending before a reference date. When no disruption indicator is active, the reference date is the beginning of the 13-week delay window. When one or more disruption indicators are active, the reference date is the first week in which that combination of indicators appears, and the classification remains fixed for all weeks assigned to that combination. This construction prevents subsequent delays under the same set of disruptions from changing the inventory classification. *Low inventory buffer* equals one when coverage is below the median across firms at the same reference date. We then interact *Low inventory buffer* with *Fraction in delay (13w)*. All three specifications include firm×UPC fixed effects. Column (1) includes the six delay-episode indicators, Column (2) includes product module×delay-episode fixed effects, and Column (3) includes product module×week fixed effects.

Table 4 reports the results. In Column (1), the interaction between delay exposure and a low inventory buffer is positive but imprecisely estimated at 0.013. Column (2) allows each product market to experience a different change during each disruption episode. The interaction increases to 0.040 and is statistically significant at the 5% level. An increase of one standard deviation in delay exposure is associated with a 0.36% price increase among firms with high inventory buffers and a 1.14% increase among firms with low buffers. Column (3) absorbs product market conditions in each calendar week. The corresponding price increases are 0.30% and 0.76%, and their difference is statistically significant at the 10% level.

All three specifications compare prices within the same firm×UPC pair over time. Columns

[10] The Compustat measure captures a firm's total inventories rather than inventories of the specific delayed products, and for diversified firms it includes goods unrelated to the delayed shipments. This measurement error blurs the distinction between low- and high-buffer firms and works against finding differences between the two groups.

(2) and (3) additionally account for conditions specific to each product market during disruption episodes or in each calendar week. In both specifications, the same increase in shipping delays is associated with a larger price increase among firms with lower predetermined inventory buffers. This pattern is consistent with inventory scarcity amplifying the price response to shipping delays.

3.5 Product Market Shipping Delays

Shipping delays can constrain the focal firm and its competitors at the same time. When delays are shared across firms, the product market may lose varieties even if no firm changes the price of its continuing UPCs. We measure this common pressure using *Product market delay exposure (13w)*. We calculate the measure separately for each firm by excluding its own delay exposure and averaging delay exposure across the other firms in the same product module and week. We weight the other firms by their 2015–2017 unit sales shares, rescaled to sum to one. The fixed preperiod weights prevent current sales or product availability from mechanically determining the measure. The estimation sample begins in 2018 and requires at least two other firms with positive preperiod sales weights.

We first examine whether delays affecting competitors are associated with changes in a firm's own prices. Column (1) of Table 5 compares prices within the same firm×UPC pair over time. Scaling the coefficient on *Product market delay exposure (13w)* by the 0.20 retail benchmark implies a 0.04% price increase. The estimate is economically small and not statistically significant.

Column (2) also uses *ln(Price)* as the outcome but compares the average price among UPCs sold by the same firm in the same product module over time. The estimate can therefore reflect both price changes within UPCs and changes in the set of UPCs available. Under the same 0.20 benchmark, product market delay exposure is associated with 2.84% higher prices among available products. The availability margins respond similarly. The same increase is associated with a 3.28 percentage point increase in UPC absence, 2.66% fewer active UPCs, and a 1.40 percentage point increase in the premium share. All four estimates are statistically significant at the 1% level.

These results complement the own-delay estimates. A firm's own delay exposure predicts price increases for the same UPC, whereas delay exposure among other firms in the product market does

not. Both own and product market delay exposure, however, predict reduced availability and a higher premium share. The pattern suggests that within-UPC repricing is tied to the firm's own delayed shipments, while an important component of the availability response arises when delay pressure extends across firms in the same product market.

3.6 Spillovers to Unaffected Products

Whether shipping delays affect only directly exposed products or spill over to unaffected products determines how broadly these disruptions affect consumer prices. Price increases in the same firm's unaffected modules would indicate portfolio-wide repricing, while increases by firms whose rivals are delayed would suggest that reduced competing supply increases residual demand and pricing power. Either spillover would extend the price effects of shipping disruptions beyond directly constrained products. We test these possibilities using a panel of firm×product module×week observations with zero own delay exposure over the preceding 13 weeks.

Table 6 reports three spillover tests. Each pair of columns uses *ln(Price)* as the outcome. The first column in each pair includes firm×UPC fixed effects and captures price changes for the same UPC, while the second includes firm×product module fixed effects and allows changes in the set of available UPCs to contribute to the estimate. Columns (1) and (2) measure exposure as the average delay exposure across the same firm's other product modules in the same week. Product module×week fixed effects compare firms with zero own delay exposure in the same product market and week. Under the 0.20 retail benchmark, delay exposure across the firm's other modules is associated with 0.20% lower prices for the same UPC and 0.98% lower prices among available products. The first estimate is statistically insignificant, while the second is significant at the 5% level. Neither estimate indicates an upward price spillover.

Columns (3) and (4) replace the average exposure with an indicator equal to one when any other product module in the same firm and week has positive delay exposure. Both estimates are −0.002 and statistically insignificant. A delay elsewhere in the firm's portfolio therefore does not predict a broader increase in its prices.

Columns (5) and (6) test whether the loss of competing supply increases the pricing power

of firms with zero own delay exposure. *Rival delay* is the equally weighted average delay exposure among the other firms in the product module and week. Within the zero-own-delay sample, this measure is common across focal firms in the same product module and week, so product module×week fixed effects would absorb it. We instead include firm×week and product module×delay-episode fixed effects, comparing the same firm's product modules with zero own delay exposure in the same week. Under the 0.20 retail benchmark, rival delay is associated with 0.48% lower prices for the same UPC and 0.46% lower prices among available products. The estimates are significant at the 1% and 10% levels, respectively, and provide no evidence that firms increase prices when their competitors are constrained.

Taken together, the estimates provide no evidence that firms increase prices in product modules with zero own delay exposure when delays occur elsewhere in their portfolio or among competitors. Combined with the positive own-delay estimates in Table 2, the results indicate that within-UPC price increases are concentrated in product modules whose own import pipeline is delayed. The negative rival delay estimate also reinforces the finding in Table 5 that delay exposure among other firms does not predict higher within-UPC prices at the focal firm. The evidence is therefore more consistent with localized inventory scarcity than with portfolio-wide repricing or higher markups by firms whose competitors are delayed.

3.7 Retail Prices around Delay Shocks

We complement the continuous panel estimates with a difference-in-differences analysis centered on clean onsets of delay exposure. A clean onset occurs when a firm×UPC unit has positive weekly delay exposure after at least 26 weeks without exposure. Because clean onsets occur at different times across firm×UPC units, the traditional difference-in-differences estimator can be biased under such staggered adoption (e.g., de Chaisemartin and D'Haultfœuille, 2020; Callaway and Sant'Anna, 2021; Goodman-Bacon, 2021; Sun and Abraham, 2021). We therefore adopt a stacked regression approach (e.g., Gormley and Matsa, 2011; Cengiz, Dube, Lindner, and Zipperer, 2019; Baker, Larcker, and Wang, 2022), which compares newly exposed units with matched not-yet-treated controls within the same cohort and calendar week. The resulting event-time coefficients

summarize the average price path across cohorts (e.g., Roth, Sant'Anna, Bilinski, and Poe, 2023; Borusyak, Jaravel, and Spiess, 2024). A cohort comprises the firm×UPC units whose clean onset occurs in a given week. These units constitute the treated observations in the cohort. For each cohort, we construct a $[-12, +12]$ week event window around the onset in which only not-yet-treated units serve as controls, avoiding contamination from already-treated units, and we match treated units to controls on pre-treatment *ln(Price)* and the log of one plus total units sold using propensity score matching. We repeat this construction for every cohort and stack the samples to estimate the average price response around clean delay onsets.

We estimate the following dynamic stacked difference-in-differences regression:

$$Y_{iptj} = \alpha_{ipj} + \alpha_{tj} + \sum_{k \in \{-12,\ldots,+12\},\, k \neq -1} \beta_k \left(\mathbb{1}[Rel_t = k] \times Treated_{ipj} \right) + \gamma' \mathbf{Z}_{ipt} + \varepsilon_{iptj}, \tag{2}$$

where i indexes firms, p products (UPCs), t weeks, and j cohorts. The dependent variable is the log retail price. $Treated_{ipj}$ equals one for firm×UPC units whose clean onset of delay exposure occurs in cohort j, Rel_t is event time in weeks relative to the onset, and $k = -1$ is the omitted reference period. The vector $\mathbf{Z}_{ipt}$ includes the log of one plus total volume in transit and the log of one plus total units sold. The firm×UPC×cohort fixed effects α_{ipj} and cohort×week fixed effects α_{tj} absorb unit heterogeneity and common time trends within each event window, so we compare treated units with their matched controls in the same calendar week of the same event window. We cluster standard errors at the firm×UPC×cohort level. The coefficients of interest are the β_k, which trace the retail price of treated units relative to their matched controls in each event week. If delay exposure constrains replenishment and increases prices, we expect the coefficients to be close to zero before the clean onset and to become positive as delayed shipments fail to arrive on schedule. The pre-onset coefficients provide evidence on the parallel-trends assumption (Rambachan and Roth, 2023), while the post-onset coefficients summarize the average price path after exposure begins.

Figure 4, Panel A plots the estimates from Eq. 2 for weeks −12 through +12 relative to the clean onset of delay exposure. We classify delay using the completed voyage and date the resulting delayed volume to the vessel's departure week. Event time zero therefore marks departure, before

the realized delay can be observed. Consistent with this timing, the coefficients are small and exhibit no trend before the onset of delays. Voyage durations provide a benchmark for when the delay should become observable. The median voyage in our matched sample lasts roughly four weeks, and a shipment is classified as delayed when its duration exceeds its route benchmark by more than seven days. A delay therefore cannot typically become observable before week 5. Delayed volume arrives 17.8 days late on average, placing the average delayed arrival around week 7. The price response remains close to zero through week 6, begins to increase in week 7, peaks at approximately 1.1% in week 10, and remains positive through week 12. Prices therefore increase when the delayed shipment should have arrived and the missed replenishment begins to bind, rather than when the shipment departs. The gradual and persistent increase is consistent with inventory depletion and restocking frictions.

Figure 4, Panel B plots binned estimates through week 52 to examine whether the price response persists over a longer horizon. The average post-delay price effect is 0.7% through week 12, 1.0% through week 26, and 1.0% through week 52. In the binned estimates, treated products remain 1.3% above matched controls in weeks 27–39 and 0.8% in weeks 40–52. Hence, delays are associated with a sustained elevation of prices, consistent with a scarcity channel in which delayed replenishment affects pricing while firms rebuild inventories and restore availability, and with the persistence of temporary upstream shocks in production networks (Liu and Tsyvinski, 2024).

To complement the dynamic estimates with an average post-onset effect, we define $Post_{tj}$ to equal one in weeks 0 through 12 of the $[-12, +12]$ window and estimate the following static version of Eq. 2.

$$Y_{iptj} = \alpha_{ipj} + \alpha_{tj} + \beta \, Treated_{ipj} \times Post_{tj} + \gamma' \mathbf{Z}_{ipt} + \varepsilon_{iptj}. \tag{3}$$

The fixed effects, controls, and clustering are the same as in Eq. 2. The coefficient of interest, β, captures the average post-onset price effect.

Table 7 reports the estimates. Column (1) estimates a 0.5% increase in retail prices after the clean onset of delay exposure relative to matched controls. Column (2) adds the log of one plus total volume in transit, and the estimate remains 0.5%. Column (3) additionally controls for the log of one plus total units sold, and the estimate declines slightly to 0.4%. All three estimates

are statistically significant at the 1% level. The estimated price response is stable across the three specifications.

To compare magnitudes across the two designs, we scale the panel estimate by the dose of delay exposure an event implies. Within the stacked sample, treated units' trailing 13-week delay exposure increases from zero before the event to 0.34 afterwards, compared with 0.08 among matched controls, a dose gap of roughly 0.26. Scaled by the panel estimate of 0.030 from Table 2, this gap implies a price response of about 0.8% once prices fully adjust, close to the event-study estimates of 0.9% to 1.1% in weeks 7 through 12 after the onset of delays. The panel and event-time estimates are thus consistent in magnitude. The smaller average post-onset effect in Table 7 reflects averaging over the adjustment period. We assign a shipment's eventual delay status to its departure week, so measured exposure increases at event time zero, before the delay can be observed. Prices respond only once the shipment fails to arrive on time, which explains why the static estimate is smaller than the later weekly estimates.

The results in this section show that shipping delays increase the cost of affected products through two margins. The first is the direct price margin, which captures increases in the prices of continuing UPCs. The second is the availability margin, which reflects both the loss of products and a shift in the available mix toward higher-priced varieties. The direct price response is stronger when inventory buffers are thin, while delays shared across firms in the same product market predict broader changes in availability and product mix. Direct price increases remain concentrated in product modules with own delay exposure and persist for months after exposure begins. For households, shipping delays therefore change both the prices of products they continue to buy and the set of products available to them.

4 Shipping Delays and Household Consumption

The retail effects need not translate uniformly into household consumption because households buy different products and can adjust their purchases when prices increase or products become unavailable. The same shipping disruption therefore produces different household-level exposures, consistent with evidence on heterogeneous household inflation rates (e.g., Kaplan and Schulhofer-

Wohl, 2017; Jaravel, 2019; Argente and Lee, 2021). We use each household's past purchases to define a baseline household bundle and link the product modules in that bundle to shipping delays. We then examine how household-level delay exposure relates to the cost of the household's bundle and the composition of its purchases. We begin by holding the baseline bundle fixed and measuring how its cost varies with delay exposure. We then allow the bundle to update with recent purchases, examine reallocation within product markets, and use an additive decomposition to separate price changes from voluntary and forced substitution.

4.1 Empirical Specification

We first estimate the following panel regression at the household×product module×week level:

$$Y_{hmt} = \alpha_{hm} + \alpha_{ht} + \delta_{me} + \beta \, \textit{Fraction in delay (13w)}_{hmt} + \varepsilon_{hmt}, \tag{4}$$

where h indexes households, m product modules, t weeks, and e delay episodes. The dependent variable Y_{hmt} is a bundle price or a purchase outcome, defined in the subsection that studies it. $\textit{Fraction in delay (13w)}_{hmt}$ is the share of import volume delayed over the 13 weeks ending in week t among the products household h normally buys in product module m. For household outcomes, we assign delayed volume to the vessel's arrival week because this timing more closely aligns exposure with when a shipping disruption affects domestic product markets and can affect the prices and availability households encounter. At the household×product module level, the measure averages firm-level delay exposure within the product market using the household's baseline spending on products supplied by each firm.

For outcomes measured at the household×product module×week level, the specification includes household×product module, household×week, and product module×delay-episode fixed effects. Household×product module fixed effects absorb stable tastes for a product category. Household×week fixed effects absorb everything that happens to a household in a given week, so identification compares differently exposed product markets within the same household and week. Product module×delay episode fixed effects absorb broad product market changes during major disruptions while preserving weekly variation within them, so we compare weeks within the

same product market and disruption episode. For outcomes measured at the household×week level, we aggregate exposure across the bundle using baseline module shares and include household and week fixed effects. We cluster standard errors by household×product module, and by household in regressions at the household×week level. The coefficient of interest, β, captures the relationship between delay exposure and each outcome. If the retail price responses affect households, we expect β to be positive for bundle prices.

4.2 Household Bundle Prices

To isolate price changes from changes in household purchasing behavior, we define each household's bundle using its purchases in one baseline calendar year and track the weekly price of that same bundle in the next calendar year. As before, we retain product modules purchased in at least four weeks of the baseline year and households active in all 12 months of both years, so that zeros reflect nonpurchase rather than attrition. Purchases in the baseline year determine the UPCs in the bundle and their expenditure weights, which remain fixed throughout the following year. In each week, we average observed UPC log prices using these fixed weights, renormalized over UPCs with observed prices. Within each module, household delay exposure averages firm×product module *Fraction in delay (13w)* using the household's spending shares across firms in the first year. We estimate Eq. 4 at the household×product module×week level, with the log price of the fixed bundle in the product market as the outcome. Because purchases in the following year do not update the bundle, the estimate captures repricing of products already in the household's bundle rather than household substitution.

Table 8 reports the results. Because all columns include household×product module fixed effects, they compare the price of the same household's baseline bundle within a product market as its delay exposure changes. Column (1) includes delay episode fixed effects and yields a coefficient of 0.180. Column (2) includes product module×delay episode fixed effects, and the estimate is 0.201. Column (3) includes product module×week fixed effects, which absorb price movements common to a product market in a given week, and the estimate is 0.047. All three estimates are statistically significant at the 1% level. In Column (3), an increase of one standard deviation in

household delay exposure is associated with a 0.38% higher bundle price. This estimate indicates that the price response to shipping delays extends to products already in household bundles.

4.3 Household Bundle Prices around Delay Shocks

To examine how bundle prices evolve around discrete increases in delay exposure, we adapt the stacked difference-in-differences design used for retail prices. A household×product module becomes treated when its fixed bundle first has positive exposure to a delay dated to the vessel's arrival week, following 26 consecutive weeks with no exposure. Within each cohort, we match each treated unit to five not-yet-treated units in the same product module using average log bundle prices and average log units sold during event weeks -8 through -2. The event study spans weeks -12 through $+12$, omits week -1, and includes household×product module×cohort and cohort×week fixed effects.

Figure 5 plots the dynamic estimates. The coefficients before treatment show no clear trend. The estimate is 0.06% in event week zero and increases to 0.71% by week twelve. Table 9 reports the corresponding average post-event effects. Bundle prices are 0.42% higher after a delay event in Column (1) without additional controls and in Column (2) after controlling for volume in transit. The estimate is 0.36% in Column (3), which also controls for units sold. All three estimates are statistically significant at the 1% level. The dynamic and average difference-in-differences estimates therefore reinforce the panel evidence that higher delay exposure is followed by higher household bundle prices.

4.4 Household Bundle Inflation

The analysis of fixed household bundles isolates repricing but does not capture changes in what households buy or the role of product availability. We therefore construct a rolling bundle for each household and calendar week t. We measure current prices and purchases in week t and delay exposure over the 13 weeks ending in that week. We define the bundle using purchases during weeks $t-64$ through $t-13$, the 52 weeks immediately preceding the delay exposure window. This timing fixes the products and spending weights before the delay exposure window begins. We first rank product modules by baseline spending and select modules through the one that brings

cumulative spending to at least 80 percent. We then retain selected modules that appear in at least two baseline weeks. These restrictions focus the bundle on economically important, recurring purchases and limit noise from small or one-time purchases.

Within the retained modules, baseline expenditures determine the UPC weights. When a price is missing in week t, we impute it through a sequence of HMS and RMS prices without assigning a zero price to a missing or stocked-out product. Household delay exposure averages delay rates across the retained modules using baseline module spending shares. Within each module and week, the delay rate averages firm-level exposure using delayed import volume. This module-level construction reduces sensitivity to noise when linking individual UPCs to importing firms.

We use these rolling bundles to construct three household price indices that differ in the expenditure shares used to weight UPC-level price changes (Kaplan and Schulhofer-Wohl, 2017; Argente and Lee, 2021). Let p_{uhb} and s_{uhb} denote the price and expenditure share of UPC u in household h's baseline bundle. Let p_{uht} and s_{uht} denote the corresponding price and expenditure share in week t, and define $\Delta \ln(p_{uht}) = \ln(p_{uht}/p_{uhb})$. The three indices are

$$\begin{aligned} \pi_{ht}^{L} &= \sum_{u} s_{uhb} \Delta \ln(p_{uht}), \\ \pi_{ht}^{P} &= \sum_{u} s_{uht} \Delta \ln(p_{uht}), \\ \pi_{ht}^{T} &= \sum_{u} \frac{s_{uhb} + s_{uht}}{2} \Delta \ln(p_{uht}). \end{aligned} \tag{5}$$

π_{ht}^{L} is the Laspeyres index, which weights price changes by baseline expenditure shares and therefore holds the household's baseline allocation across UPCs fixed. π_{ht}^{P} is the Paasche index, which uses current expenditure shares and therefore incorporates the household's reallocation across UPCs by week t. π_{ht}^{T} is the Tornqvist index, which weights each UPC by the average of its baseline and current expenditure shares and therefore treats the two allocations symmetrically.

Table 10 reports household-week regressions of these indices on delay exposure.[11] The estimate

[11] Household delay exposure may partly capture how heavily the baseline bundle relies on imports. Product modules with greater import penetration may experience different price changes for reasons unrelated to realized shipping delays. To address this concern, Table IA.6 in the Internet Appendix controls for mean import propensity across the classified product modules in the rolling baseline bundle, weighted by the household's spending in each module. The estimates remain similar.

is 0.240 for the Laspeyres index in Column (1), 0.191 for the Tornqvist index in Column (2), 0.120 for the Paasche index in Column (3), and 0.238 for the rolling Laspeyres index in Column (4). All four estimates are statistically significant at the 1% level. The Tornqvist estimate lies between the Laspeyres and Paasche estimates, as expected from its weighting scheme. An increase of one standard deviation in household delay exposure implies about a 1.9% increase in the household bundle price index at baseline weights. The Paasche estimate is approximately half the Laspeyres estimate, consistent with households shifting spending away from the products whose prices increase most. Taken together, the four estimates show that shipping delays raise the cost of household bundles even when the index accounts for household reallocation across products. The household regressions therefore complement the retail analysis by aggregating delay conditions and price changes across the product markets in each household's bundle.

4.5 Household Consumption Reallocation

Households may respond to higher prices and reduced product availability by changing both how much they buy in a product market and which products they choose within that market. We examine these responses to shipping delays using bundles defined by purchases in the prior calendar year. We rank product modules by baseline spending and retain modules through the one that brings cumulative spending to at least 80 percent. This restriction focuses the analysis on product markets that account for most of the household's spending and limits noise from categories that make up only a small share of the bundle. We require each household to record at least one purchase in every month of both the calendar year used to define the bundle and the following outcome year so that measured changes do not reflect gaps in panel participation. Within each module, household delay exposure averages firm×product module *Fraction in delay (13w)* using baseline firm shares, while the outcomes measure actual purchases in the outcome week. The regressions use Eq. 4 and include all three sets of fixed effects. We expect households to adjust primarily by reallocating spending toward cheaper and more familiar products within exposed markets rather than by reducing purchases in those markets.

Table 11, Panel A reports purchase, spending, and quantity outcomes. An increase of one

standard deviation in household delay exposure increases the probability of purchasing the module by 0.03 percentage points in Column (1), decreases log spending by 0.02% in Column (2), and increases log quantity by 0.05% in Column (3). Columns (4) and (5) show that the module's share of total spending and of spending within the selected bundle decreases by 0.002 and 0.004 percentage points, respectively. In Column (6), quantity relative to the household's baseline average increases by 0.002 units. All six estimates are statistically significant at the 1% level, but economically modest. Taken together, they indicate little change in whether or how much households buy in exposed product markets.

Panel B examines how households allocate spending within the module. In Column (1), an increase of one standard deviation in household delay exposure increases the expenditure share of the baseline UPC, the product purchased during the bundle's baseline window, by 0.24 percentage points. The share spent on other UPCs decreases by the same amount in Column (2). Column (3) shows that the share spent on the baseline brand increases by 0.04 percentage points. The shares spent on UPCs priced below the baseline module's 25th percentile and median increase by 0.05 and 0.12 percentage points in Columns (4) and (5), respectively. The share allocated to large-quantity purchases increases by 0.01 percentage points in Column (6). All estimates are statistically significant at the 1% level.

These within-module reallocations are larger than the changes in Panel A but remain economically modest. The largest estimate is the 0.24 percentage point increase in the baseline UPC share, while the shares allocated to cheaper products increase by no more than 0.12 percentage points. Household adjustment is therefore concentrated within product markets, but average changes are small. Households continue purchasing within exposed product markets while reallocating spending toward cheaper and more familiar products. The fixed effects already control for baseline import propensity. Household×week fixed effects absorb the import propensity of the household's overall baseline bundle, while household×product module fixed effects absorb the fixed Census import propensity of each product module.

4.6 Additive Decomposition and Stockouts

The rolling bundle price index combines price changes for individual products with changes in household expenditure weights. Reallocation may be voluntary, when households change spending among available products, or forced, when they replace an unavailable baseline product. We separate total bundle price change into a price component, which averages each UPC's log price change using the mean of its baseline and current expenditure levels, a voluntary reallocation component for changes among available products, and a forced stockout component for replacements of unavailable products.[12] We estimate Eq. 4 at the household×week level with each component as the outcome.

Table 12 reports the decomposition. The total response to delay exposure is 0.240 in Column (1). The price component is 0.239 in Column (2), the voluntary reallocation component is 0.029 in Column (3), and the forced stockout component is −0.029 in Column (4). All four estimates are statistically significant at the 1% level and add exactly by construction. Almost the entire net response therefore reflects price changes for individual products evaluated using the mean of baseline and current expenditure weights, while the positive voluntary reallocation component and negative forced stockout component offset each other.

Constructing the decomposition requires a current price for baseline products that the household does not purchase in the outcome week, including products that are unavailable. Table IA.5 in the Internet Appendix examines whether the results depend on how we assign these missing prices. The table replaces the preferred sequence of price sources with each source separately, restricts the index to products with observed prices in both periods, and varies the price change assigned to unavailable products. Across these alternatives, the price component continues to account for nearly the entire response. The decomposition also remains similar when we restrict the sample to households whose bundle import propensity is observed and include the control, as Table IA.7 shows.

Table 13 examines the changes in availability that underlie forced replacement. An increase

[12] Voluntary and forced reallocation can occur simultaneously, so the contribution assigned to each depends on which change in expenditure weights is introduced first. We calculate the contribution of each form of reallocation under both possible orderings and average the two values. This Shapley split removes the arbitrary ordering and makes the decomposition exactly additive, so that $\pi_{\text{total}} = \pi_{\text{prices}} + \pi_{\text{voluntary}} + \pi_{\text{forced}}$.

of one standard deviation in household delay exposure increases the share of the baseline bundle exposed to stockouts by 5.0 percentage points in Column (1). The weight on sparsely priced UPCs, which have no usable price in the outcome week but are not classified as stockouts, increases by 0.1 percentage points in Column (2). The shares of spending on UPCs and product markets that are new to the household increase by 2.1 and 0.6 percentage points in Columns (3) and (4), respectively. The share of the bundle for which prices can be observed or imputed decreases by 0.3 percentage points in Column (5). All estimates are statistically significant at the 1% level. These results show that delay shocks change the set of products households can purchase, not only the transaction prices of products that remain available (e.g., Broda and Weinstein, 2010; Michelacci, Paciello, and Pozzi, 2022; Atkin et al., 2024). The estimates remain similar after restricting the sample to households whose bundle import propensity is observed and including the control, as Table IA.8 shows.

Together, these results show that shipping delays increase the prices of products households already buy and alter the composition of their purchases. Households reallocate spending within product markets toward cheaper and more familiar products, while stockouts force them to replace unavailable goods. Shipping delays therefore affect households through both the prices they pay and the products they can buy.

5 Shipping Delays and Household Inflation Expectations

The changes in household prices and product availability documented above may also shape beliefs about aggregate inflation. We therefore examine whether households with greater shipping delay exposure through their consumption bundles report higher expected inflation over the following 12 months. We present this section as a descriptive extension rather than as a causal analysis. Households choose the bundles that determine their measured exposure, so the cross-sectional comparison identifies an association between basket composition and beliefs; the within-household comparison narrows this considerably but still relies on the pandemic disruption as the source of variation. We report both, together with the placebo evidence that delimits them, and draw no causal conclusion about the strength of belief responses.

5.1 Empirical Specification

Household inflation expectations come from customized surveys administered to members of the NielsenIQ Consumer Panel (D'Acunto et al., 2021; Weber, Gorodnichenko, and Coibion, 2023; Weber et al., 2025). The surveys follow the question formats used by the Michigan Survey of Consumers and the New York Fed Survey of Consumer Expectations. The direct question asks respondents for a numerical forecast of the expected percentage change in prices or inflation over the following 12 months. The distributional question asks them to assign probabilities across possible inflation and deflation ranges over the same horizon.

NielsenIQ fields these surveys quarterly, and the waves used in our analysis span April 2018 through December 2023. Most waves were open for about three weeks, so respondents in the same wave could complete the questionnaire in different calendar weeks. Households participated in some waves but not others, producing an unbalanced panel.

Our unit of observation is a household in one of these survey waves. When more than one member of a household responds in the same wave, we average their forecasts. We define the survey week as the NielsenIQ calendar week ending on the Saturday before the response. The rolling baseline bundle covers the 52 weeks ending 13 weeks before the survey week. We retain product modules that account for the top 80 percent of baseline spending and appear in at least two baseline weeks, then weight them by baseline expenditure. Exposure is the resulting bundle-weighted average of product module delay rates, dated to the vessel's arrival week, over the 13 completed weeks ending in the survey week. To ensure that exposure can be measured for a substantial portion of the bundle, we require product modules with positive observed import arrival volume during the same window to represent at least 50 percent of bundle spending.

Only a minority of households appear in at least three consecutive waves, and many repeat respondents have long gaps between interviews. In addition, because shipping disruptions are transient, the effect of a disruption on inflation expectations may dissipate before the household's next response. We therefore first focus on cross-sectional variation in delay exposure among households interviewed during the same survey week.

We estimate the following household-level regression:

$$Expectation_{h,w} = \beta\, Fraction\ in\ delay\ (13w)_{h,w} + X'_{h,w}\gamma + \delta_w + \lambda_s + \varepsilon_{h,w}, \tag{6}$$

where h indexes households, w survey weeks, and s states. The dependent variable, *Expectation*, is the household's direct numerical forecast of inflation over the following 12 months. We observe this outcome in 17 survey waves fielded between April 2018 and December 2023. These include four waves between April 2018 and March 2019, three in the second half of 2020, four in 2021, and three in each of 2022 and 2023. The survey fixed effects δ_w absorb conditions common to households responding during the same week, while the state fixed effects λ_s absorb persistent differences across states. The vector $X_{h,w}$ includes respondent age and gender composition, the presence of children, household size, race, Hispanic origin, income, log bundle spending, log bundle breadth, and indicators for the household's survey order, the ordinal position of the current wave among the waves the household has answered, which capture the household's experience with the survey. Standard errors are clustered by household and survey week.

Persistent differences across households may still contribute to the cross-sectional association. To address this concern, we compare changes in each household's average forecast and bundle-specific delay exposure across two widely separated survey windows. The COVID-19 period provides substantial variation because the disruption to container shipping was both large and sustained. Figure 2 shows that shipping delays increased sharply in 2020 and remained elevated throughout 2021. We compare the six pre-pandemic waves from April 2018 through September 2019 with the four waves from February through December 2021. The 12-month expectation horizons of the two windows do not overlap. The design therefore compares households that experienced larger and smaller changes in exposure during the same aggregate disruption, rather than attributing the common change in expectations during COVID-19 to shipping delays.

Restricting the revision analysis to the direct numerical forecasts used in the cross-sectional analysis would omit three of the six pre-pandemic survey waves and leave relatively few households observed both before the pandemic and in 2021. These are the surveys fielded in 2018Q3, 2019Q2, and 2019Q3. The latter two are the latest surveys in our pre-pandemic comparison window. The

survey designers omitted the direct point-forecast question from these waves because of questionnaire space constraints, while retaining the distributional question (Weber, Gorodnichenko, and Coibion, 2023). Excluding them would therefore end the pre-pandemic comparison in 2019Q1 and sharply reduce the number of households observed in both periods. We include these surveys and calculate the mean forecast implied by each respondent's reported probabilities. We then average each household's eligible forecasts and corresponding delay exposure measures within each window.

Most households in the long-difference sample have at least one probability-derived forecast in their pre-pandemic average, while all four 2021 surveys elicit direct numerical forecasts. Direct numerical forecasts are typically higher than the mean forecasts implied by reported probability distributions, so variation across households in the mix of pre-pandemic responses could bias the exposure coefficient if that mix is correlated with changes in exposure.[13] To reduce sensitivity to differences in scale and extreme responses, we use the direction of the revision as the main outcome. *Upward revision* equals one when the household's average forecast is higher in 2021 than before the pandemic. This indicator places revisions on a common directional scale, although it does not make the two elicitation formats identical. We estimate the following linear probability model:

$$\textit{Upward revision}_h = \alpha + \beta\, \Delta \textit{Fraction in delay (13w)}_h + X_h'\gamma + \varepsilon_h. \tag{7}$$

The change in delay exposure, Δ*Fraction in delay (13w)*, is the difference between the household's average exposure in 2021 and its pre-pandemic average. Taking the difference removes any household-specific tendency to report consistently high or low inflation expectations, provided that tendency remains stable across the two periods. The vector X_h includes the household's pre-pandemic average forecast and the change in the share of bundle spending covered by observed import arrivals. We report heteroskedasticity-robust standard errors.

[13] In the long-difference sample, pre-pandemic forecasts average 7.7 percent among households answering only direct questions and 2.1 percent among households whose forecasts are derived only from probability distributions, a comparison that mixes the format difference with differences in which households answer each format. The exposure association shares this pattern. A 10 percentage point increase in the exposure change is associated with a 7.9 percentage point higher probability of an upward revision among households with only probability-derived pre-pandemic forecasts and 5.4 percentage points among households with both formats, while the association is statistically indistinguishable from zero among the 8 percent of households with only direct pre-pandemic forecasts.

5.2 Delay Exposure and Expected Inflation

Table 14 reports estimates from Eq. 6. The estimates are positive across all four specifications. With only survey week fixed effects, an increase of one standard deviation in delay exposure is associated with 0.621 percentage points higher expected inflation over the following 12 months in Column (1). Adding household characteristics reduces the estimate to 0.471 in Column (2). Both estimates are statistically significant at the 1% level.

The relationship remains positive after adding bundle and survey controls. Adding bundle spending and bundle breadth reduces the effect of a one-standard-deviation increase in exposure to 0.417 percentage points in Column (3), while adding survey experience reduces it to 0.398 percentage points in Column (4). The controls also reproduce established cross-sectional survey patterns. Female respondents report higher expectations, middle- and high-income households report lower expectations, and reported expectations decrease with survey experience. Taken together, the estimates indicate that households with greater delay exposure through their consumption baskets report higher expected inflation than other households interviewed during the same survey week.

If the positive cross-sectional association reflects delays affecting the products a household buys, the coefficient based on the household's own basket should be larger than coefficients based on other households' baskets. If instead the association reflects persistent differences in household types or broad basket composition, exposure measured through other households with similar baskets should produce comparable estimates. We test these alternatives by assigning each household the delay exposure calculated from another household's basket in the same survey week, while keeping the household's inflation forecast, controls, fixed effects, and estimation sample unchanged. Table IA.11 in the Internet Appendix reports two versions of this placebo test. Panel A reassigns baskets among all households interviewed during the same survey week. Panel B reassigns baskets only among households that are similar in their pre-2018 consumption of imported goods, share of spending in product modules with measurable delay exposure, and bundle breadth. Each panel repeats the reassignment 999 times. The upper tail reference probability equals one plus the number of placebo coefficients at least as large as the coefficient based on the household's own basket, divided by 1,000. None of the unrestricted reassignments in Panel A produces a coefficient as large as the own-basket

estimate, yielding a reference probability of 0.001 in every specification. The reference probabilities range from 0.035 to 0.077 when Panel B restricts the reassignments to similar baskets.

Exposure measured through the household's own basket is therefore more informative than exposure measured through another household's basket. The higher reference probabilities among similar baskets indicate that persistent differences in basket composition contribute to the cross-sectional relationship. This concern motivates the long-difference analysis in the next subsection. By comparing changes in exposure and expectations within the same household, the design removes stable tendencies to report high or low expectations and asks whether households whose delay exposure increased more were also more likely to revise their expectations upward.

5.3 Delay Exposure and Expectation Revisions

Table 15 reports estimates from Eq. 7. The sample extends the pre-pandemic window beyond the waves with direct numerical forecasts used in Table 14. Changes in delay exposure are positively related to the direction of household expectation revisions. Adding the pre-pandemic forecast and the change in exposure coverage increases the coefficient from 0.470 in Column (1) to 0.566 in Column (3), and all three estimates are statistically significant at the 1% level. An increase of one standard deviation in the exposure change is associated with a 2.4 percentage point increase in the probability that a household revises its forecast upward, against an unconditional share of 74 percent of households revising upward. Column (4) enters exposure in each period separately. The coefficient on 2021 exposure is 0.527 and statistically significant at the 1% level, while pre-pandemic exposure is not statistically distinguishable from zero. Households with greater delay exposure during the 2021 supply chain disruption therefore tend to revise their expectations upward, while higher pre-pandemic delay exposure is unrelated to an upward revision.

The result does not depend on changes in the household's product module spending shares. Column (5) fixes these shares using 2019 purchases. The coefficient is 0.655, which implies a 2.3 percentage point increase in the probability of an upward revision for an increase of one standard deviation in fixed-basket exposure. The fixed-basket estimate indicates that the association does not reflect households changing the product composition of their baskets between periods. Taken

together, the estimates suggest that households experiencing larger increases in delay exposure tend to revise their inflation expectations upward, including when exposure is constructed from a fixed pre-pandemic basket. The estimates also remain positive and statistically significant when we control for household characteristics and local pandemic exposure, as Table IA.9 shows.

Overall, the cross-sectional and revision estimates are consistent with delay exposure shifting household inflation expectations. Because the revision comparison spans the pandemic, we read it as descriptive of that relationship rather than as a causal estimate of its magnitude.

6 Conclusion

Shipping delays affect households through higher prices and changes in product availability. Using vessel tracking records linked to retail and household purchase data, we show that shipping delays raise the prices of continuing products and reduce the number of products available to consumers. Households respond by reallocating spending within affected product markets toward cheaper and more familiar products, while stockouts force them to replace unavailable goods. The evidence on inflation expectations points in the same direction, but is more tentative. Households with more exposed bundles report higher expected inflation across the cross-sectional specifications, but compositionally similar placebo bundles reproduce much of this association. Comparing each household with itself before and during the pandemic disruption, households experiencing larger increases in delay exposure tend to revise their expectations upward, including when their bundle is held fixed at pre-pandemic purchases. We read this pattern as suggesting that beliefs respond, rather than as a causal estimate of that response.

Measures based on posted prices capture the repricing of products that remain available but not the changes in consumption caused by product shortages. When supply disruptions remove products from shelves, households must change what they buy even if the prices of the remaining products are fully observed. The inflation expectations results also show that bundle composition can generate substantial cross-sectional relationships between measured exposure and reported beliefs. The directional long difference reduces this concern but cannot eliminate differences created by survey elicitation. Researchers must therefore separate persistent differences in bundle composition

from the shocks households experience before drawing conclusions about belief updating.

References

Acemoglu, D., V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi. 2012. The Network Origins of Aggregate Fluctuations. *Econometrica* 80:1977–2016.

Adão, R., P. Carrillo, A. Costinot, D. Donaldson, and D. Pomeranz. 2022. Imports, Exports, and Earnings Inequality: Measures of Exposure and Estimates of Incidence. *Quarterly Journal of Economics* 137:1553–1614.

Albagli, E., F. Grigoli, and E. Luttini. 2025. Inflation Expectations and the Supply Chain. *IMF Economic Review* 73:819–850.

Alessandria, G., J. P. Kaboski, and V. Midrigan. 2010. Inventories, Lumpy Trade, and Large Devaluations. *American Economic Review* 100:2304–2339.

Alessandria, G., S. Y. Khan, A. Khederlarian, C. Mix, and K. J. Ruhl. 2023. The Aggregate Effects of Global and Local Supply Chain Disruptions. *Journal of International Economics* 146:103740.

Amiti, M., O. Itskhoki, and J. Konings. 2019. International Shocks, Variable Markups, and Domestic Prices. *Review of Economic Studies* 86:2356–2402.

Amiti, M., S. J. Redding, and D. E. Weinstein. 2019. The Impact of the 2018 Tariffs on Prices and Welfare. *Journal of Economic Perspectives* 33:187–210.

Argente, D., and M. Lee. 2021. Cost of Living Inequality During the Great Recession. *Journal of the European Economic Association* 19:913–952.

Atkin, D., B. Faber, T. Fally, and M. Gonzalez-Navarro. 2024. Measuring Welfare and Inequality with Incomplete Price Information. *Quarterly Journal of Economics* 139:419–475.

Auer, R., A. Burstein, S. M. Lein, and J. Vogel. 2024. Unequal Expenditure Switching: Evidence from Switzerland. *Review of Economic Studies* 91:2572–2603.

Bai, L., T. S. Jaccard, and S. Stumpner. 2026. Tariff Pass-through into Retail Prices: Measurement, Replacement, and Shrink-flation. Working Paper.

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. NBER Working Paper 32098.

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.

Baqaee, D., A. Burstein, and Y. Koike-Mori. 2024. Measuring Welfare by Matching Households across Time. *Quarterly Journal of Economics* 139:533–573.

Baqaee, D., and E. Farhi. 2019. The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem. *Econometrica* 87:1155–1203.

Barrot, J.-N., and J. Sauvagnat. 2016. Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks. *Quarterly Journal of Economics* 131:1543–1592.

Baslandze, S., and S. Fuchs. 2026. The Price of Delay: Supply Chain Disruptions and Pricing Dynamics. *Journal of International Economics* 162:104265.

Bils, M., and P. J. Klenow. 2004. Some Evidence on the Importance of Sticky Prices. *Journal of Political Economy* 112:947–985.

Borusyak, K., X. Jaravel, and J. Spiess. 2024. Revisiting Event-Study Designs: Robust and Efficient Estimation. *Review of Economic Studies* 91:3253–3285.

Brancaccio, G., M. Kalouptsidi, and T. Papageorgiou. 2020. Geography, transportation, and endogenous trade costs. *Econometrica* 88:657–691.

Broda, C., and D. E. Weinstein. 2010. Product Creation and Destruction: Evidence and Price Implications. *American Economic Review* 100:691–723.

Callaway, B., and P. H. C. Sant'Anna. 2021. Difference-in-Differences with Multiple Time Periods. *Journal of Econometrics* 225:200–230.

Canto, F. D., J. Grigsby, E. Qian, and C. J. Walsh. 2025. Are Inflationary Shocks Regressive? A Feasible Set Approach. *Quarterly Journal of Economics* 140:2685–2747.

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.

Carvalho, V. M., M. Nirei, Y. Saito, and A. Tahbaz-Salehi. 2021. Supply Chain Disruptions: Evidence from the Great East Japan Earthquake. *Quarterly Journal of Economics* 136:1255–1321.

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.

Cavallo, A., and O. Kryvtsov. 2023. What Can Stockouts Tell Us About Inflation? Evidence from Online Micro Data. *Journal of International Economics* 146:103769.

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

Chen, A., L. Sanz, and M. D. Wittry. 2026. Who Sets the Price? Information and the Delegation of Pricing Authority. Working Paper.

Chen, T., P. Levell, and M. O'Connell. 2026. Measuring cost of living inequality during an inflation surge. Working Paper.

Coibion, O., D. Georgarakos, Y. Gorodnichenko, and M. Weber. 2023. Forward Guidance and Household Expectations. *Journal of the European Economic Association* 21:2131–2171.

Coibion, O., Y. Gorodnichenko, and M. Weber. 2022. Monetary Policy Communications and Their Effects on Household Inflation Expectations. *Journal of Political Economy* 130:1537–1584.

Colicev, A., J. Hoste, and J. Konings. 2024. The Impact of a Large Depreciation on the Cost of Living of Rich and Poor Consumers. *International Economic Review* 65:1625–1656.

Cravino, J., and A. A. Levchenko. 2017. The Distributional Consequences of Large Devaluations. *American Economic Review* 107:3477–3509.

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.

de Chaisemartin, C., and X. D'Haultfœuille. 2020. Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects. *American Economic Review* 110:2964–2996.

DellaVigna, S., and M. Gentzkow. 2019. Uniform Pricing in U.S. Retail Chains. *Quarterly Journal of Economics* 134:2011–2084.

Djankov, S., C. Freund, and C. S. Pham. 2010. Trading on Time. *Review of Economics and Statistics* 92:166–173.

Eichenbaum, M., N. Jaimovich, and S. Rebelo. 2011. Reference Prices, Costs, and Nominal Rigidities. *American Economic Review* 101:234–262.

Evans, C. L., and J. Harrigan. 2005. Distance, Time, and Specialization: Lean Retailing in General Equilibrium. *American Economic Review* 95:292–313.

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.

Flaaen, A., F. Haberkorn, L. Lewis, A. Monken, J. Pierce, R. Rhodes, and M. Yi. 2023. Bill of lading data in international trade research with an application to the COVID-19 pandemic. *Review of International Economics* 31:1146–1172.

Fuchs, S., F. Leibovici, and W. F. Wong. 2026. Geopolitical Risk and Foreign Dependence in Global Shipping. Working Paper.

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

Gopinath, G., and O. Itskhoki. 2010. Frequency of Price Adjustment and Pass-Through. *Quarterly Journal of Economics* 125:675–727.

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

Handbury, J. 2021. Are Poor Cities Cheap for Everyone? Non-Homotheticity and the Cost of Living Across U.S. Cities. *Econometrica* 89:2679–2715.

Hastings, J. S., and J. M. Shapiro. 2013. Fungibility and Consumer Choice: Evidence from Commodity Price Shocks. *Quarterly Journal of Economics* 128:1449–1498.

Heiland, I., A. Moxnes, K. H. Ulltveit-Moe, and Y. Zi. 2025. Trade From Space: Shipping Networks and The Global Implications of Local Shocks. *Review of Economics and Statistics* 1–45.

Hoke, S. H., and L. Feler. 2026. Paying More and Buying Less: 2025 Tariffs and U.S. Household Spending. Working Paper.

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

Jaimovich, N., S. Rebelo, and A. Wong. 2019. Trading down and the business cycle. *Journal of Monetary Economics* 102:96–121.

Jaravel, X. 2019. The Unequal Gains from Product Innovations: Evidence from the U.S. Retail Sector. *Quarterly Journal of Economics* 134:715–783.

Jaravel, X., and D. Lashkari. 2024. Measuring Growth in Consumer Welfare with Income-Dependent Preferences: Nonparametric Methods and Estimates for the United States. *Quarterly Journal of Economics* 139:477–532.

Jiao, Y., T. Lan, Y. Liu, and X. Zhou. 2026. From Ports to Prices: The Inflationary Effects of Global Supply Chain Disruptions. Working Paper.

Kamali, P., R. Koepke, A. Sozzi, and J. Verschuur. 2024. Red Sea Attacks Disrupt Global Trade. IMF Blog, March 7, 2024.

Känzig, D. R., and R. Raghavan. 2026. Supply Chain Shocks and the Macroeconomy: Evidence from Global Shipping Disruptions. Working Paper.

Kaplan, G., and S. Schulhofer-Wohl. 2017. Inflation at the household level. *Journal of Monetary Economics* 91:19–38.

Kryvtsov, O., and V. Midrigan. 2013. Inventories, Markups, and Real Rigidities in Menu Cost Models. *Review of Economic Studies* 80:249–276.

Kwon, C. 2021. Supply Chain Disruptions: Evidence from the Bankruptcy of Hanjin Shipping. Working Paper.

Liu, E., and A. Tsyvinski. 2024. A Dynamic Model of Input–Output Networks. *Review of Economic Studies* 91:3608–3644.

Michelacci, C., L. Paciello, and A. Pozzi. 2022. The Extensive Margin of Aggregate Consumption Demand. *Review of Economic Studies* 89:909–947.

Nakamura, E., and J. Steinsson. 2008. Five Facts about Prices: A Reevaluation of Menu Cost Models. *Quarterly Journal of Economics* 123:1415–1464.

Pankratz, N. M. 2026. When Trade Dries Up: Infrastructure Constraints and Firm Performance. Working Paper.

Rambachan, A., and J. Roth. 2023. A More Credible Approach to Parallel Trends. *Review of Economic Studies* 90:2555–2591.

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* 135:503–560.

Roth, J., P. H. C. Sant'Anna, A. Bilinski, and J. Poe. 2023. What's Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature. *Journal of Econometrics* 235:2218–2244.

Sanz, L. 2026. Building Corporate Resilience to Supply Chain Disruptions. Working Paper (SSRN 4635488).

Schafhäutle, S. G., and G. S. Sran. 2026. Redacted Identities in Shipment Records: Evidence from Forced Labor Scrutiny in Supply Chains. Wharton School Research Paper.

Schuler, M. 2023. Container Dwell Times at Los Angeles and Long Beach Ports Return to Normal. gCaptain, February 21, 2023, https://gcaptain.com/container-dwell-times-at-los-angeles-and-long-beach-ports-return-to-normal/ (Accessed: August 31, 2026).

Sun, L., and S. Abraham. 2021. Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects. *Journal of Econometrics* 225:175–199.

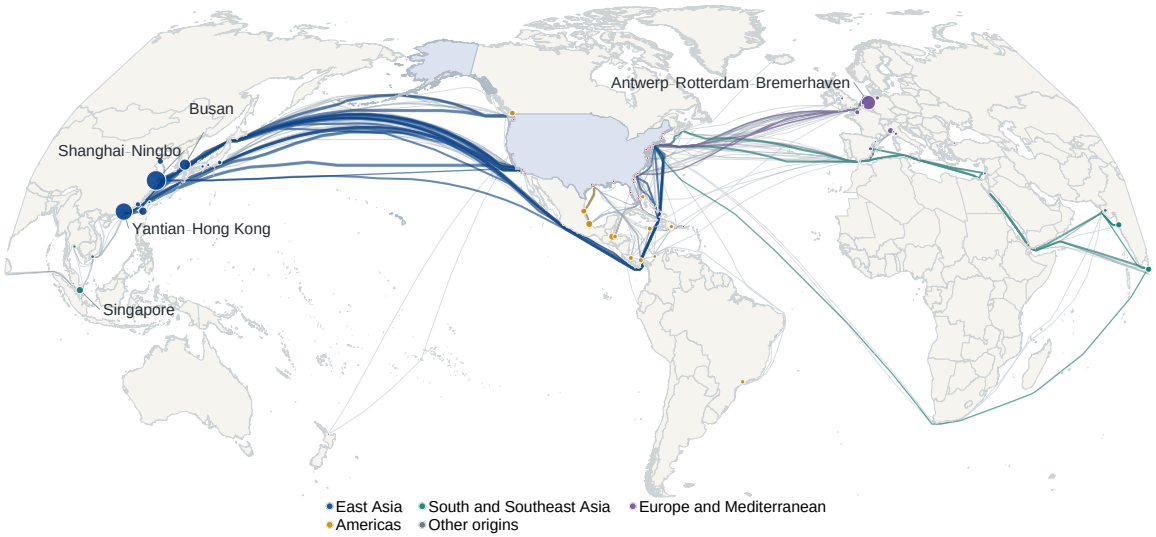
Weber, M., B. Candia, H. Afrouzi, T. Ropele, R. Lluberas, S. Frache, B. Meyer, S. Kumar, Y. Gorodnichenko, D. Georgarakos, O. Coibion, G. Kenny, and J. Ponce. 2025. Tell Me Something I Don't Already Know: Learning in Low- and High-Inflation Settings. *Econometrica* 93:229–264.

Weber, M., Y. Gorodnichenko, and O. Coibion. 2023. The Expected, Perceived, and Realized Inflation of U.S. Households before and during the COVID-19 Pandemic. *IMF Economic Review* 71:326–368.

Figure 1: Shipping routes and U.S. arrival ports

This figure shows the shipping routes and U.S. arrival ports in the matched AIS and bill of lading sample from 2015 through 2024. Panel A ranks origin–destination port pairs by the number of matched voyages and plots representative observed AIS trajectories for the pairs that cumulatively account for 70.0% of those voyages. Line width increases with the number of matched voyages assigned to each displayed route. The 25 port pairs with the largest numbers of matched voyages are colored by origin region, and origin-port marker radius increases with voyage volume. Panel B plots all 40 contiguous U.S. destination port codes in the full matched sample. Marker area is proportional to the number of matched voyages. The eight labeled arrival areas account for 85.9% of contiguous U.S. matched voyages and 83.3% of all matched voyages.

Panel A: Representative AIS shipping routes



Panel B: Contiguous U.S. arrival ports

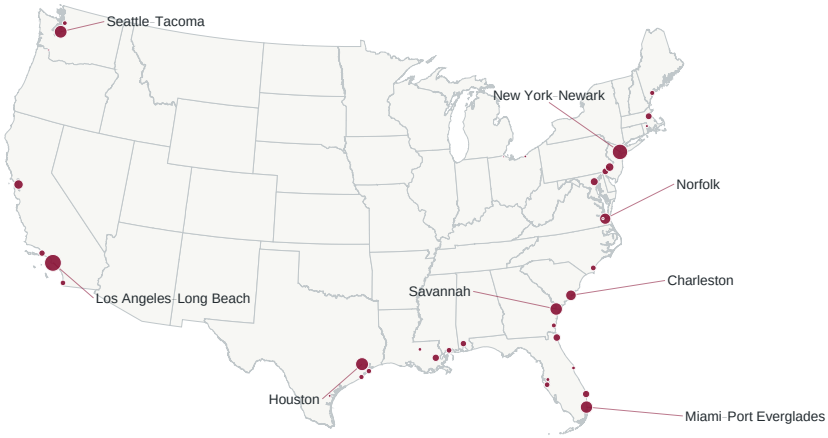


Figure 2: Shipping delays

This figure plots several measures of shipping delays using AIS vessel tracking data from 2018 through 2024. Panel A plots the 4-week moving average of the mean deviation of voyage duration from port-pair historical benchmarks. The benchmark is the median voyage duration for each origin–destination port pair computed over 2015–2017. Panel B plots the 4-week moving average of the fraction of voyages classified as delayed (voyage duration exceeding the port-pair benchmark by more than seven days). Panel C plots the 4-week moving average of mean anchorage hours near the ports of Los Angeles and Long Beach, computed from AIS navigation status broadcasts. A vessel is classified as at anchor when its navigation status code equals 1 and it is within 50 kilometers of the destination port. Table A.1 provides detailed variable definitions.

Panel A: Trip duration



Panel B: Fraction of voyages delayed



Panel C: Port delays

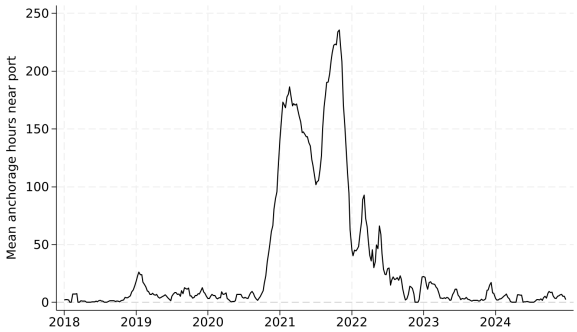
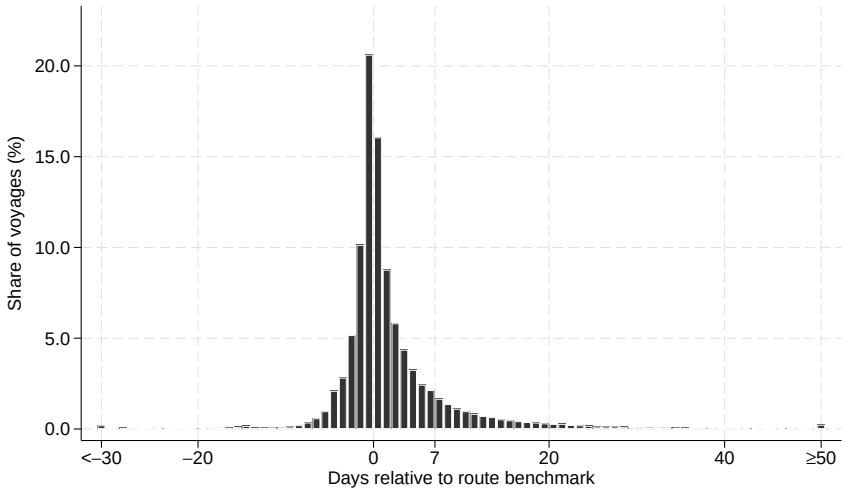


Figure 3: Voyage and delay durations

This figure plots distributions of voyage duration in the matched AIS and bill of lading sample used to construct the delay measure from 2015 through 2024. Each unique voyage receives equal weight. Panel A plots realized voyage duration minus the 2015–2017 median duration for the same origin–destination port pair. Port-pair benchmarks require at least 10 voyages during 2015–2017, and voyages whose duration exceeds three times the route median are excluded as likely mismatches. Panel B restricts the sample to voyages classified as delayed, defined as voyages whose duration exceeds the port-pair benchmark by more than seven days, and plots their excess duration relative to the benchmark. The sample contains 173,610 voyages, of which 22,081, or 12.7%, are classified as delayed. Bars report the share of voyages in one-day bins. In Panel A, the first bin includes voyages completed more than 30 days faster than the route benchmark, and the last bin includes voyages taking at least 50 days longer than the benchmark. In Panel B, the last bin includes delayed voyages taking at least 60 days longer than the route benchmark. Table A.1 provides detailed variable definitions.

Panel A: Duration relative to route benchmark



Panel B: Excess duration among delayed voyages

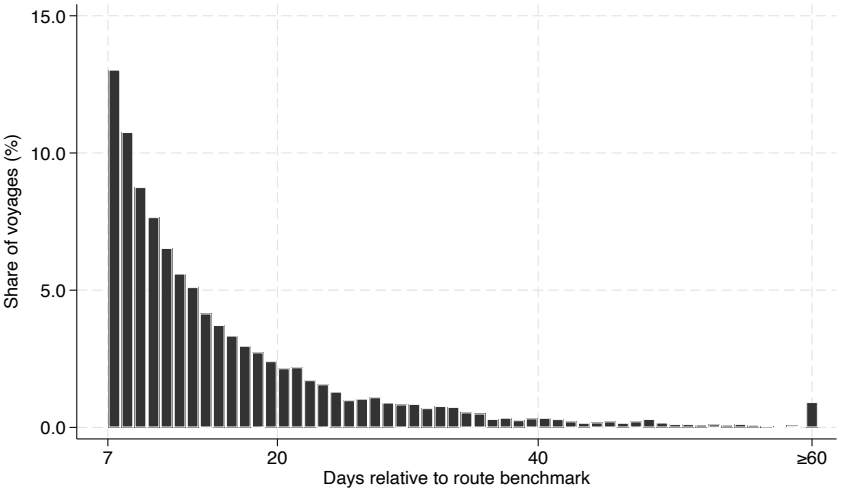
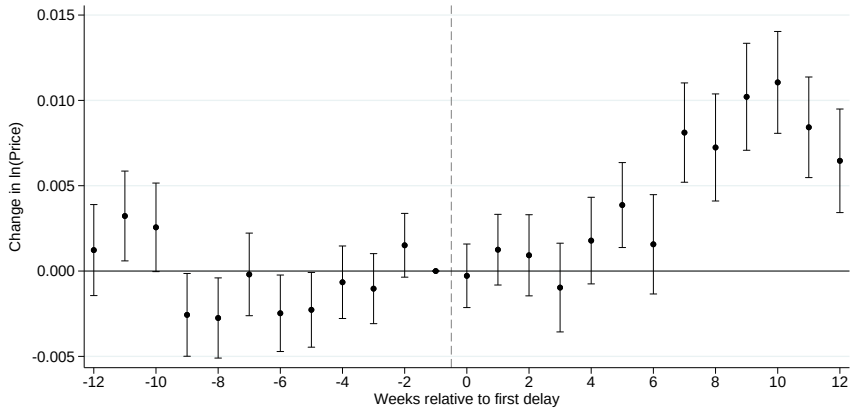


Figure 4: Retail prices around delay shocks

This figure plots dynamic treatment effects and 90% confidence intervals from stacked difference-in-differences regressions of *ln(Price)* around shipment delay events. Panel A plots weekly event-time estimates from the baseline event window. Panel B plots binned event-time estimates from a longer-window stack that requires each cohort to have support through 52 weeks after the delay event. *Treated* equals one for firm×UPC units that experience their first clean delay event in the cohort week. Delayed exposure is the weekly share of a firm's import volume carried by shipments whose matched vessel trip duration exceeds the 2015–2017 route-pair median by more than seven days. Delays are assigned to the vessel departure week. A clean event is a positive delayed exposure after 26 weeks without delayed exposure. The sample stacks cohort-specific event windows around delay events and includes not-yet-treated controls, drawn from a firm×UPC×week panel linking U.S. publicly listed firms to NielsenIQ retail scanner data from January 2015 through December 2024. Treated and control units are matched within cohort using propensity score matching with five nearest neighbors on pre-treatment *ln(Price)* and the log of one plus total units sold. Event time $k = 0$ denotes the beginning of the first delay, the week in which the first delayed shipment departs, and $k = -1$ is the omitted reference period. The specification includes firm×UPC×cohort and cohort×week fixed effects. Standard errors are clustered at the firm×UPC×cohort level. Table A.1 provides detailed variable definitions.

Panel A: Baseline event window



Panel B: Longer-window persistence

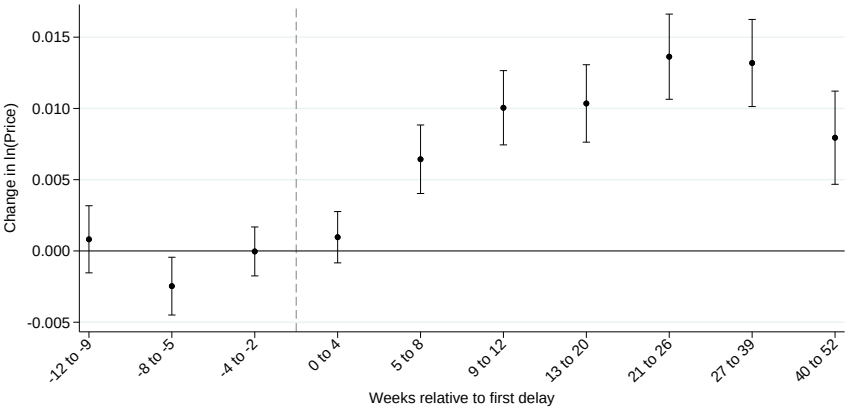


Figure 5: Household fixed-bundle prices around delay shocks

This figure plots dynamic treatment effects and 90% confidence intervals from stacked difference-in-differences regressions of *ln(Bundle price)* around shipment delay events. A treatment event occurs when a household's fixed bundle in a product module first has positive arrival-timed delay exposure after 26 weeks without exposure. Controls are not-yet-treated household×product module units from the same product module that remain untreated through the event window. Treated and control units are matched within cohort using propensity score matching with five nearest neighbors on average *ln(Bundle price)* and the average log of one plus fixed-bundle units sold over event weeks -8 through -2. The regression window spans weeks $[-12, +12]$, and event week -1 is the omitted reference period. Weeks -12 through -9 are held out from matching. The specification includes household×product module×cohort and cohort×week fixed effects. Standard errors are clustered at the household×product module×cohort level. Table A.1 provides detailed variable definitions.

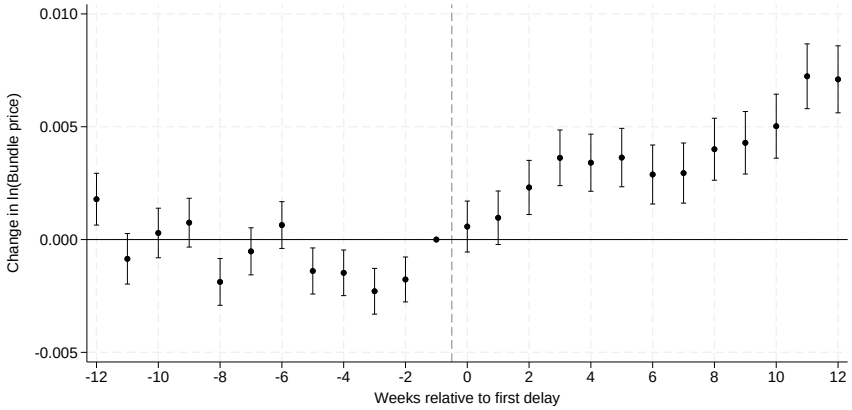


Table 1: Summary statistics

This table reports summary statistics for the paper's main analysis samples. Panel A reports summary statistics for a firm×UPC×week sample. Panel B reports summary statistics for a household×week sample. For each observation, week t is the week in which household inflation and reallocation outcomes are measured. The baseline bundle is constructed from purchases during the 52 weeks from $t-64$ through $t-13$. The 12 weeks immediately preceding t are excluded from bundle construction because they enter the 13-week delay exposure measure for t, which covers weeks $t-12$ through t. Product modules are ranked by baseline expenditure, and modules up to and including the one that brings cumulative spending to at least 80 percent of total spending are included in the bundle. In addition, the bundle retains modules purchased in at least two baseline weeks. Panel C reports summary statistics for the common household×survey wave sample used in the inflation expectations regressions. Direct point forecasts are averaged within household and survey wave. The last completed week is the NielsenIQ calendar week ending on the Saturday before the response, and delay exposure is measured over the 13 completed weeks ending in that week. The sample requires product modules with positive observed import-arrival volume during the same window to represent at least 50 percent of bundle spending. Table A.1 provides detailed definitions of variables used in the regressions.

	Obs. (1)	Mean (2)	Std. dev. (3)	25th pct. (4)	Median (5)	75th pct. (6)
Panel A: Firm × UPC × week sample						
Fraction in delay (13w)	18,709,303	0.10	0.20	0.00	0.00	0.09
Average product price	12,117,663	16.77	38.67	4.49	8.66	17.32
Absent	18,709,303	0.35	0.48	0.00	0.00	1.00
Low inventory buffer	10,978,543	0.41	0.49	0.00	0.00	1.00
Rival delay	18,709,303	0.08	0.14	0.00	0.02	0.09
Panel B: Household × week sample						
Fraction in delay (13w)	35,347,026	0.07	0.08	0.02	0.04	0.11
Delay exposure coverage	35,347,026	0.63	0.21	0.50	0.66	0.79
Baseline bundle expenditure ($)	35,347,026	3758.81	2885.47	1726.85	3179.06	5090.61
Baseline bundle UPCs	35,347,026	255.61	175.85	127.00	227.00	349.00
Laspeyres bundle inflation	35,347,026	−0.03	0.16	−0.08	−0.01	0.05
Paasche bundle inflation	22,627,559	0.03	0.27	−0.05	0.02	0.11
Stockout weight	35,347,026	0.42	0.16	0.31	0.42	0.53
New UPC expenditure share	24,259,060	0.55	0.26	0.37	0.55	0.74
Panel C: Household × survey wave sample						
Expected inflation, next 12 months (%)	116,009	8.30	15.54	2.00	5.00	10.00
Fraction in delay (13w)	116,009	0.12	0.10	0.03	0.09	0.17
Delay exposure coverage	116,009	0.70	0.13	0.60	0.69	0.79
Bundle spending ($)	116,009	4127.47	2925.63	2124.30	3586.70	5472.27
Bundle UPCs	116,009	279.83	180.23	153.00	252.00	373.00
Mean respondent age	115,790	58.95	13.49	50.00	61.00	69.00
Female respondent share	116,009	0.73	0.43	0.00	1.00	1.00
Children present	116,009	0.14	0.35	0.00	0.00	0.00
Household size	89,602	2.03	1.15	1.00	2.00	2.00
Hispanic origin	116,009	0.06	0.23	0.00	0.00	0.00
Middle income	116,009	0.38	0.48	0.00	0.00	1.00
High income	116,009	0.17	0.37	0.00	0.00	0.00
Survey order 2–4	116,009	0.33	0.47	0.00	0.00	1.00
Survey order 5–9	116,009	0.24	0.43	0.00	0.00	0.00
Survey order 10 or more	116,009	0.26	0.44	0.00	0.00	1.00

Table 2: Shipping delays and the retail prices of continuing products

This table reports estimates from panel regressions of *ln(Price)* on *Fraction in delay (13w)*. *ln(Price)* is the natural log of the average price per item across stores, weighted by unit sales, for a firm×UPC×week. *Fraction in delay (13w)* is the share of a firm's import volume over the preceding 13 weeks carried by shipments whose matched vessel trip duration exceeds the 2015–2017 route-pair median by more than seven days. Delayed volume is assigned to the vessel departure week before constructing the rolling share. The sample is a firm×UPC×week panel linking U.S. publicly listed firms to NielsenIQ retail scanner data from January 2015 through December 2024. All columns include firm×UPC fixed effects. Column (1) includes delay-episode fixed effects for the 2015 West Coast port labor slowdown, COVID port congestion (Feb 2020–Dec 2021), Suez Canal blockage (Mar 2021), post-COVID backlog (Jan 2021–Jun 2022), Panama Canal drought (Aug 2023–Mar 2024), and Red Sea disruption (Dec 2023–Dec 2024). Column (2) includes product module×delay-episode fixed effects. Column (3) includes product module×week fixed effects. Standard errors are clustered at the firm×product module level and reported in parentheses. ***, **, * denote statistical significance at the 1%, 5%, and 10% level, respectively. Table A.1 provides detailed variable definitions.

	ln(Price)		
	(1)	(2)	(3)
Fraction in delay (13w)	0.028***	0.029***	0.030***
	(0.009)	(0.009)	(0.009)
Fixed effects			
Firm × UPC	✓	✓	✓
Delay episodes	✓		
Product module × Delay episodes		✓	
Product module × Week			✓
Observations	17,324,084	17,324,084	17,324,084
Adj. R^2	0.898	0.900	0.905

Table 3: Shipping delays, stockouts, and the composition of available products

This table reports estimates from panel regressions of the log prices of available UPCs and three outcomes capturing product availability and composition on *Fraction in delay (13w)*. *Fraction in delay (13w)* is the share of a firm's import volume over the preceding 13 weeks carried by shipments whose matched vessel trip duration exceeds the 2015–2017 route-pair median by more than seven days. Delayed volume is assigned to the vessel departure week before constructing the rolling share. The sample links U.S. publicly listed firms to NielsenIQ retail scanner data from January 2015 through December 2024. Column (1) uses *ln(Price)* as the dependent variable in the firm×UPC×week price panel and includes firm×product module fixed effects. Column (2) uses a balanced firm×UPC×week panel over each UPC's observed lifetime and reports *Absent*, an indicator equal to one if the UPC does not appear in NielsenIQ that week. Columns (3) and (4) use a firm×product module×week panel and report *ln(# UPCs)*, the log number of active UPCs, and *Premium*, the share of premium UPCs. A premium UPC is one whose mean log price over the sample exceeds the median of UPC mean log prices within the same firm and product module, so the premium share moves only with the set of available UPCs. All columns include product module×delay-episode fixed effects. Column (2) includes firm×UPC fixed effects, and Columns (1), (3), and (4) include firm×product module fixed effects. Standard errors are clustered at the firm×product module level and reported in parentheses. ***, **, * denote statistical significance at the 1%, 5%, and 10% level, respectively. Table A.1 provides detailed variable definitions.

	ln(Price)	Absent	ln(# UPCs)	Premium
	(1)	(2)	(3)	(4)
Fraction in delay (13w)	0.131***	0.101***	−0.060***	0.055***
	(0.016)	(0.013)	(0.016)	(0.006)
Fixed effects				
Firm × UPC		✓		
Firm × Product module	✓		✓	✓
Product module × Delay episodes	✓	✓	✓	✓
Unit of observation	$i \times p \times t$	$i \times p \times t$	$i \times m \times t$	$i \times m \times t$
Observations	17,381,022	26,186,282	739,176	739,176
Adj. R^2	0.522	0.498	0.929	0.664

Table 4: Shipping delays and inventory buffers

This table reports estimates from panel regressions of *ln(Price)* that interact *Fraction in delay (13w)* with *Low inventory buffer*. *ln(Price)* is the natural log of the average price per item across stores, weighted by unit sales, for a firm×UPC×week. *Fraction in delay (13w)* is the share of a firm's import volume over the preceding 13 weeks carried by shipments whose matched vessel trip duration exceeds the 2015–2017 route-pair median by more than seven days. Delayed volume is assigned to the vessel departure week before constructing the measure. The sample is a firm×UPC×week panel linking U.S. publicly listed firms to NielsenIQ retail scanner data from January 2015 through December 2024. Days of inventory coverage equals quarterly total inventory divided by quarterly cost of goods sold, multiplied by 91.25, the average number of days in a quarter. Missing total inventory is assigned zero when quarterly cost of goods sold is positive. *Low inventory buffer* equals one when coverage is below the median across firms at the same reference date. The measure uses the latest fiscal quarter ending before the reference date. When no disruption indicator is active, the reference date is the beginning of the 13-week delay window. When one or more disruption indicators are active, the reference date is the first week in which that combination of indicators appears, and the classification remains fixed for all weeks assigned to that combination. All columns include firm×UPC fixed effects. Column (1) includes the six delay-episode indicators, Column (2) includes product module×delay-episode fixed effects, and Column (3) includes product module×week fixed effects. The row *Effect of Fraction in delay for low-inventory firms* reports the sum of the coefficients on *Fraction in delay (13w)* and its interaction with *Low inventory buffer*. Its standard error uses the covariance between the two estimates. Standard errors are clustered at the firm×product module level and reported in parentheses. ***, **, * denote statistical significance at the 1%, 5%, and 10% level, respectively. Table A.1 provides detailed variable definitions.

	ln(Price)		
	(1)	(2)	(3)
Fraction in delay (13w)	0.018**	0.018*	0.015*
	(0.009)	(0.010)	(0.008)
Low inventory buffer	−0.013	−0.020*	−0.001
	(0.012)	(0.012)	(0.006)
Fraction in delay (13w) × Low inventory buffer	0.013	0.040**	0.023*
	(0.020)	(0.016)	(0.013)
Effect of Fraction in delay for low-inventory firms	0.031*	0.057***	0.038***
	(0.018)	(0.014)	(0.012)
Fixed effects			
Firm × UPC	✓	✓	✓
Delay episodes	✓		
Product module × Delay episodes		✓	
Product module × Week			✓
Observations	10,978,543	10,978,543	10,978,543
Adj. R^2	0.892	0.894	0.899

Table 5: Product market shipping delays

This table reports estimates from panel regressions of prices and product availability on *Product market delay exposure (13w)*. *Product market delay exposure (13w)* is the mean delay exposure among all other firms in the same product module and week, weighted by their unit sales shares during 2015–2017. Firm delay exposure is the share of import volume over the preceding 13 weeks carried by shipments whose matched vessel trip duration exceeds the 2015–2017 route-pair median by more than seven days. Delayed volume is assigned to the vessel departure week before constructing the measure. To ensure that the exposure measure is based on multiple peers, observations are retained only when at least two other firms in the same product module have positive sales during 2015–2017. The sample links U.S. publicly listed firms to NielsenIQ retail scanner data from January 2018 through December 2024. Columns (1) and (2) use *ln(Price)* as the dependent variable in the firm×UPC×week price panel. Column (1) includes firm×UPC fixed effects and captures direct price changes for continuing UPCs. Column (2) includes firm×product module fixed effects and therefore also captures changes in the composition of available UPCs. Column (3) uses a balanced firm×UPC×week panel over each UPC's observed lifetime and reports *Absent*, an indicator equal to one if the UPC does not appear in NielsenIQ that week. Columns (4) and (5) use a firm×product module×week panel and report *ln(# UPCs)*, the log number of active UPCs, and *Premium*, the share of premium UPCs. A premium UPC is one whose mean log price over the sample exceeds the median of UPC mean log prices within the same firm and product module, so the premium share moves only with the set of available UPCs. All columns include product module×delay episode fixed effects. Standard errors are clustered at the firm×product module level and reported in parentheses. ***, **, * denote statistical significance at the 1%, 5%, and 10% level, respectively. Table A.1 provides detailed variable definitions.

	ln(Price)		Absent	ln(# UPCs)	Premium
	Direct effect	Compositional effect			
	(1)	(2)	(3)	(4)	(5)
Product market delay exposure (13w)	0.002	0.142***	0.164***	−0.133***	0.070***
	(0.018)	(0.013)	(0.020)	(0.025)	(0.009)
Fixed effects					
Firm × UPC	✓		✓		
Firm × Product module		✓		✓	✓
Product module × Delay episodes	✓	✓	✓	✓	✓
Unit of observation	$i \times p \times t$	$i \times p \times t$	$i \times p \times t$	$i \times m \times t$	$i \times m \times t$
Observations	9,329,647	9,337,277	14,392,196	359,046	359,046
Adj. R^2	0.911	0.515	0.544	0.939	0.692

Table 6: Spillovers to unaffected products

This table reports panel regressions testing whether shipping delays spill over to firm×product module×week observations with zero own delay exposure over the preceding 13 weeks. *Fraction in delay (13w)* is the share of a firm's import volume over the preceding 13 weeks carried by shipments whose matched vessel trip duration exceeds the 2015–2017 route-pair median by more than seven days. Delayed volume is assigned to the vessel departure week before constructing the measure. The sample links U.S. publicly listed firms to NielsenIQ retail scanner data from January 2015 through December 2024. In Columns (1) and (2), *Other product modules' delay* is the average of *Fraction in delay (13w)* across the same firm's other product modules in the same week, excluding the focal module. In Columns (3) and (4), *Any other module delayed* equals one if at least one other product module in the same firm and week has positive *Fraction in delay (13w)*. In Columns (5) and (6), *Rival delay* is the equally weighted average of *Fraction in delay (13w)* among other firms in the same product module and week. These columns require each firm×week to contain at least two product modules with zero own delay exposure and different values of *Rival delay*. All columns use *ln(Price)* as the dependent variable in the firm×UPC×week price panel. Columns (1), (3), and (5) include firm×UPC fixed effects and capture price changes for the same UPC. Columns (2), (4), and (6) include firm×product module fixed effects and allow changes in the set of available UPCs to contribute to the estimate. Columns (1) to (4) include product module×week fixed effects. Columns (5) and (6) include firm×week and product module×delay-episode fixed effects. Standard errors are clustered at the firm×product module and product module×week levels and reported in parentheses. ***, **, * denote statistical significance at the 1%, 5%, and 10% level, respectively. Table A.1 provides detailed variable definitions.

	ln(Price)					
	Within firm		Firm-wide		Across firms	
	Direct effect	Compositional effect	Direct effect	Compositional effect	Direct effect	Compositional effect
	(1)	(2)	(3)	(4)	(5)	(6)
Other product modules' delay	−0.010	−0.049**				
	(0.012)	(0.019)				
Any other module delayed			−0.002	−0.002		
			(0.002)	(0.005)		
Rival delay					−0.024***	−0.023*
					(0.009)	(0.014)
Fixed effects						
Firm × UPC	✓		✓		✓	
Firm × Product module		✓		✓		✓
Product module × Week	✓	✓	✓	✓		
Firm × Week					✓	✓
Product module × Delay episode					✓	✓
Unit of observation	$i \times p \times t$	$i \times p \times t$	$i \times p \times t$	$i \times p \times t$	$i \times p \times t$	$i \times p \times t$
Observations	9,862,985	9,872,328	9,862,985	9,872,328	7,315,332	7,323,890
Adj. R^2	0.918	0.546	0.918	0.546	0.909	0.525

Table 7: Retail prices around delay shocks

This table reports results from stacked difference-in-differences regressions of log retail prices around shipment delay events. The sample stacks cohort-specific event windows around delay events and includes not-yet-treated controls, drawn from a firm×UPC×week panel linking U.S. publicly listed firms to NielsenIQ retail scanner data from January 2015 through December 2024. *Treated* equals one for firm×UPC units that experience their first clean delay event in the cohort week. Delayed exposure is the weekly share of a firm's import volume carried by shipments whose matched vessel trip duration exceeds the 2015–2017 route-pair median by more than seven days, assigned to the vessel departure week. A clean event is a positive delayed exposure after 26 weeks without delayed exposure. Treated and control units are matched within cohort using propensity score matching with five nearest neighbors on pre-treatment *ln(Price)* and the log of one plus total units sold. The regression window spans weeks $[-12, +12]$ around delay onset. *Treated* × *Post* equals one for treated units in the post-delay period. All specifications include firm×UPC×cohort and cohort×week fixed effects. Column (1) includes only these fixed effects. Column (2) adds the log of one plus total volume in transit. Column (3) additionally controls for the log of one plus total units sold. Standard errors are clustered at the firm×UPC×cohort level and reported in parentheses. ***, **, * denote statistical significance at the 1%, 5%, and 10% level, respectively. Table A.1 provides detailed variable definitions.

	ln(Price)		
	(1)	(2)	(3)
Treated × Post	0.005^{***}	0.005^{***}	0.004^{***}
	(0.001)	(0.001)	(0.001)
ln(1+Total volume in transit)		-0.001^{***}	-0.001^{***}
		(0.000)	(0.000)
ln(1+Total units sold)			-0.015^{***}
			(0.001)
Fixed effects			
Firm × UPC × Cohort	✓	✓	✓
Cohort × Week	✓	✓	✓
Observations	3,207,265	3,207,265	3,207,265
Adj. R^2	0.969	0.969	0.970

Table 8: Shipping delays and household fixed-bundle prices

This table reports estimates from panel regressions of log household fixed-bundle prices on *Fraction in delay (13w)*. The dependent variable is the log price of a household×product module fixed bundle, computed from HMS UPC-week prices using within-module UPC expenditure shares from the preceding calendar year. The sample is a household×product module×week panel. A household-module enters the sample if the household purchased the module in at least four weeks of the preceding calendar year, and the household is active in all 12 months of both the preceding and outcome years. *Fraction in delay (13w)* averages firm×product module delay exposure using the household's firm spending shares from the preceding calendar year. Delayed import volume is assigned to the vessel arrival week before constructing exposure. All columns include household×product module fixed effects. Column (1) adds delay-episode fixed effects. Column (2) replaces these fixed effects with product module×delay-episode fixed effects. Column (3) includes product module×week fixed effects. Standard errors are clustered at the household×product module level and reported in parentheses. ***, **, * denote statistical significance at the 1%, 5%, and 10% level, respectively. Table A.1 provides detailed variable definitions.

	ln(Bundle price)		
	(1)	(2)	(3)
Fraction in delay (13w)	0.180***	0.201***	0.047***
	(0.002)	(0.002)	(0.002)
Fixed effects			
Household × product module	✓	✓	✓
Delay episodes	✓		
Product module × Delay episodes		✓	
Product module × Week			✓
Observations	1,356,251,376	1,356,251,348	1,356,240,541
Adj. R^2	0.884	0.885	0.898

Table 9: Household fixed-bundle prices around delay shocks

This table reports results from stacked difference-in-differences regressions of *ln(Bundle price)* around shipment delay events. A treatment event occurs when a household's fixed bundle in a product module first has positive arrival-timed delay exposure after 26 weeks without exposure. Controls are not-yet-treated household×product module units from the same product module that remain untreated through the event window. Treated and control units are matched within cohort using propensity score matching with five nearest neighbors on average *ln(Bundle price)* and the average log of one plus fixed-bundle units sold over event weeks -8 through -2. The regression window spans weeks $[-12, +12]$. *Treated* × *Post* equals one for treated units in event weeks zero through twelve. All specifications include household×product module×cohort and cohort×week fixed effects. Column (1) includes only these fixed effects. Column (2) adds the log of one plus the household's fixed-bundle firm-share-weighted volume in transit. Column (3) additionally controls for the log of one plus fixed-bundle units sold. Standard errors are clustered at the household×product module×cohort level and reported in parentheses. ***, **, * denote statistical significance at the 1%, 5%, and 10% level, respectively. Table A.1 provides detailed variable definitions.

	ln(Bundle price)		
	(1)	(2)	(3)
Treated × Post	0.0042***	0.0042***	0.0036***
	(0.0005)	(0.0005)	(0.0005)
ln(1 + Total volume in transit)		0.0000	−0.0002**
		(0.0001)	(0.0001)
ln(1 + Total units sold)			0.0018***
			(0.0003)
Fixed effects			
Household × Product module × Cohort	✓	✓	✓
Cohort × Week	✓	✓	✓
Observations	14,419,215	14,419,215	14,419,215
Adj. R^2	0.948	0.948	0.948

Table 10: Shipping delays and rolling household bundle inflation

This table reports estimates from household-week panel regressions of log household bundle price indices on *Fraction in delay (13w)*. The sample is a household×week panel constructed from rolling household UPC bundles. The rolling baseline bundle is measured over 52 weeks ending 13 weeks before the outcome week and keeps product modules that account for the top 80 percent of household baseline spending and are purchased in at least two baseline weeks. *Fraction in delay (13w)* is the household's rolling baseline product-module-bundle-weighted average of product-module delay exposure, measured using the 13-week rolling delay calendar. Product-module delay exposure is volume-weighted across firms within each product module and week, with delayed import volume assigned to the vessel arrival week. Columns differ in the price-index formula used to aggregate UPC-level price changes within the household bundle. Column (1) reports a Laspeyres index that weights price changes by baseline UPC expenditure shares. Column (2) reports a Tornqvist index that weights price changes by the arithmetic mean of each UPC's baseline and current expenditure shares. Column (3) reports a Paasche index that weights price changes by current UPC expenditure shares. Column (4) reports the rolling Laspeyres index using the price-window imputation cascade for UPCs without exact-week prices. Columns (2) and (3) require current expenditure shares and therefore use fewer observations. All specifications include household and week fixed effects. Standard errors are clustered at the household level and reported in parentheses. ***, **, * denote statistical significance at the 1%, 5%, and 10% level, respectively. Table A.1 provides detailed variable definitions.

	ln(Bundle price)			
	Laspeyres	Tornqvist	Paasche	Rolling Laspeyres
	(1)	(2)	(3)	(4)
Fraction in delay (13w)	0.240***	0.191***	0.120***	0.238***
	(0.004)	(0.003)	(0.003)	(0.004)
Fixed effects				
Household	✓	✓	✓	✓
Week	✓	✓	✓	✓
Observations	35,347,026	22,627,559	22,627,559	35,347,026
Adj. R^2	0.418	0.099	0.025	0.418

Table 11: Household reallocation within product modules

This table reports estimates from household×product module×week panel regressions of household purchase and reallocation outcomes on *Fraction in delay (13w)*. The sample uses product modules that account for the top 80 percent of a household's baseline spending bundle. *Fraction in delay (13w)* is the baseline firm-share-weighted average of firm×product module delay exposure in the household's module bundle. Delayed import volume is assigned to the vessel arrival week before constructing exposure. Panel A reports module-level purchase, spending, quantity, and spending-share outcomes. Panel B reports within-module reallocation outcomes, including spending shares allocated to baseline UPCs and brands, cheaper UPCs, and large-quantity purchases. All columns include household×product module, household×week, and product module×delay-episode fixed effects. Standard errors are clustered at the household×product module level and reported in parentheses. ***, **, * denote statistical significance at the 1%, 5%, and 10% level, respectively. Table A.1 provides detailed variable definitions.

Panel A: Module purchase, spending, and quantity

	Any purchase	ln(Spending)	ln(Quantity)	Module share	Bundle share	Excess quantity
	(1)	(2)	(3)	(4)	(5)	(6)
Fraction in delay (13w)	0.004***	−0.003***	0.006***	−0.000***	−0.000***	0.028***
	(0.000)	(0.001)	(0.000)	(0.000)	(0.000)	(0.004)
Fixed effects						
Household × module	✓	✓	✓	✓	✓	✓
Household × week	✓	✓	✓	✓	✓	✓
Module × delay episode	✓	✓	✓	✓	✓	✓
Observations	1,133,022,793	1,133,022,793	1,133,022,793	1,133,022,793	1,113,500,830	1,133,022,793
Adj. R^2	0.232	0.243	0.268	0.206	0.189	0.528

Panel B: Within-module reallocation

	Baseline UPC	Other UPCs	Baseline brand	Cheap p25	Cheap p50	Large size
	(1)	(2)	(3)	(4)	(5)	(6)
Fraction in delay (13w)	0.031***	−0.031***	0.005***	0.006***	0.014***	0.001***
	(0.002)	(0.002)	(0.001)	(0.001)	(0.002)	(0.000)
Fixed effects						
Household × module	✓	✓	✓	✓	✓	✓
Household × week	✓	✓	✓	✓	✓	✓
Module × delay episode	✓	✓	✓	✓	✓	✓
Observations	183,262,759	183,262,759	183,262,759	183,262,759	183,262,759	1,133,022,793
Adj. R^2	0.370	0.370	0.510	0.475	0.446	0.036

Table 12: Shipping delays and the additive household inflation decomposition

This table reports estimates from household-week panel regressions of the additive household inflation decomposition on *Fraction in delay (13w)*. The decomposition separates total household price change into a price component that averages each UPC's log price change using the mean of its baseline and current expenditure levels, a voluntary reallocation component from changes in expenditure weights among available products, and a forced stockout component from replacing unavailable baseline products. Because voluntary and forced reallocation can occur simultaneously, the contribution assigned to each depends on which change in expenditure weights is introduced first. The Shapley split computes each contribution under both possible orderings and averages the two values, so $\pi_{\text{total}} = \pi_{\text{prices}} + \pi_{\text{voluntary}} + \pi_{\text{forced}}$. The rolling baseline bundle is measured over 52 weeks ending 13 weeks before the outcome week and keeps product modules that account for the top 80 percent of household baseline spending and are purchased in at least two baseline weeks. *Fraction in delay (13w)* is the household's rolling baseline product-module-bundle-weighted average of product-module delay exposure, measured using the 13-week rolling delay calendar. Product-module delay exposure is volume-weighted across firms within each product module and week, with delayed import volume assigned to the vessel arrival week. All specifications include household and week fixed effects. Standard errors are clustered at the household level and reported in parentheses. ***, **, * denote statistical significance at the 1%, 5%, and 10% level, respectively. Table A.1 provides detailed variable definitions.

	ln(Bundle price)			
	Total price	Price component	Voluntary	Forced stockout
	(1)	(2)	(3)	(4)
Fraction in delay (13w)	0.240***	0.239***	0.029***	−0.029***
	(0.004)	(0.004)	(0.001)	(0.000)
Fixed effects				
Household	✓	✓	✓	✓
Week	✓	✓	✓	✓
Observations	35,347,026	35,347,026	35,347,026	35,347,026
Adj. R^2	0.418	0.408	0.312	0.315

Table 13: Shipping delays, rolling-bundle reallocation, and stockouts

This table reports estimates from household-week panel regressions of rolling-bundle reallocation and stockout outcomes on *Fraction in delay (13w)*. Outcomes measure the baseline bundle weight exposed to stockouts, sparse-price support, spending shares on new UPCs and new modules, and imputed price coverage. The rolling baseline bundle is measured over 52 weeks ending 13 weeks before the outcome week and keeps product modules that account for the top 80 percent of household baseline spending and are purchased in at least two baseline weeks. *Fraction in delay (13w)* is the household's rolling baseline product-module-bundle-weighted average of product-module delay exposure, measured using the 13-week rolling delay calendar. Product-module delay exposure is volume-weighted across firms within each product module and week, with delayed import volume assigned to the vessel arrival week. All specifications include household and week fixed effects. Standard errors are clustered at the household level and reported in parentheses. ***, **, * denote statistical significance at the 1%, 5%, and 10% level, respectively. Table A.1 provides detailed variable definitions.

	Stockout weight	Sparse weight	New UPC share	New module share	Price coverage
	(1)	(2)	(3)	(4)	(5)
Fraction in delay (13w)	0.622^{***}	0.014^{***}	0.263^{***}	0.080^{***}	-0.039^{***}
	(0.005)	(0.001)	(0.004)	(0.003)	(0.001)
Fixed effects					
Household	✓	✓	✓	✓	✓
Week	✓	✓	✓	✓	✓
Observations	35,347,026	35,347,026	24,259,060	24,259,060	35,347,026
Adj. R^2	0.670	0.202	0.229	0.128	0.254

Table 14: Shipping delays and household inflation expectations

This table reports cross-sectional regressions of expected inflation over the following 12 months on *Fraction in delay (13w)*. The sample includes the 17 survey waves fielded between April 2018 and December 2023 that elicit direct numerical forecasts. These forecasts are averaged within household and survey wave. *Fraction in delay (13w)* is the household's bundle-weighted average of arrival-timed product-module delay exposure during the 13 completed weeks ending in the last completed week. The last completed week is the NielsenIQ calendar week ending on the Saturday before the response. All columns use the same sample and require product modules with positive observed import-arrival volume during the same window to represent at least 50 percent of bundle spending. All columns include survey fixed effects that group responses by survey week. Column (2) adds respondent age and gender composition, children, household size, race, Hispanic origin, income, and state. Column (3) adds log bundle spending and log bundle breadth. Column (4) adds survey-order indicators. Standard errors are clustered at the household and survey week levels and reported in parentheses. ***, **, * denote statistical significance at the 1%, 5%, and 10% level, respectively. Table A.1 provides detailed variable definitions.

	Expected inflation, next 12 months			
	(1)	(2)	(3)	(4)
Fraction in delay (13w)	6.211***	4.714***	4.171***	3.985***
	(1.081)	(0.933)	(0.896)	(0.878)
Female respondent share		1.529***	1.621***	1.558***
		(0.286)	(0.294)	(0.291)
Middle-income household		−1.981***	−1.964***	−1.993***
		(0.252)	(0.252)	(0.254)
High-income household		−3.632***	−3.616***	−3.646***
		(0.310)	(0.306)	(0.307)
Survey order 2–4				−0.507**
				(0.251)
Survey order 5–9				−1.691***
				(0.311)
Survey order 10 or more				−2.526***
				(0.372)
Additional controls				
Household characteristics		✓	✓	✓
Bundle spending and breadth			✓	✓
Survey experience				✓
Fixed effects				
State		✓	✓	✓
Survey	✓	✓	✓	✓
Observations	116,009	116,009	116,009	116,009
Households	36,871	36,871	36,871	36,871
Adj. R^2	0.024	0.042	0.042	0.044

Table 15: Shipping delays and upward revisions in household inflation expectations

This table reports household-level linear probability regressions in which the dependent variable equals one when the household's average forecast of inflation over the following 12 months is higher in 2021 than before the pandemic. The pre-pandemic window uses six surveys from April 2018 through September 2019, and the later window uses four surveys from February through December 2021. The inflation forecast is the respondent's expected inflation rate over the following 12 months, measured using the direct numerical response when available and the mean implied by the reported probability distribution in three pre-pandemic surveys. Forecasts and *Fraction in delay (13w)* are first averaged within household and survey and then across the surveys in each window. Columns (1) to (4) measure delay exposure using rolling household baskets. Column (5) holds calendar-2019 household product-module spending shares fixed at every survey anchor. The rolling-basket columns contain 8,200 households, and the fixed-basket column contains 9,458 households observed in both windows. For each basket definition, product modules with measurable delay exposure must represent at least 50 percent of bundle spending in each included household survey. Unsupported modules are assigned zero delay and remain in the basket denominator. Column (2) adds the household's pre-pandemic average expected inflation. Column (3) adds the change in the share of bundle spending represented by product modules with measurable delay exposure. Column (4) enters exposure in the two windows separately. Column (5) repeats Column (3) using the fixed 2019 basket. Heteroskedasticity-robust standard errors are reported in parentheses. ***, **, * denote statistical significance at the 1%, 5%, and 10% level, respectively. Table A.1 provides detailed variable definitions.

	Upward revision in expected inflation				
	Rolling basket				Fixed 2019 basket
	(1)	(2)	(3)	(4)	(5)
Δ Fraction in delay (13w)	0.470***	0.512***	0.566***		0.655***
	(0.115)	(0.108)	(0.111)		(0.116)
Fraction in delay (13w), 2021				0.527***	
				(0.108)	
Fraction in delay (13w), pre-pandemic				0.456	
				(0.423)	
Pre-pandemic expected inflation		−0.023***	−0.023***	−0.023***	−0.023***
		(0.001)	(0.001)	(0.001)	(0.001)
Δ Delay exposure coverage			−0.110**		−0.134
			(0.046)		(0.221)
Observations	8,200	8,200	8,200	8,200	9,458
Adj. R^2	0.002	0.180	0.181	0.181	0.180

Appendix A: Variable Definitions

Table A.1: Variable definitions

This table defines the variables used in the reported regressions. Panel A describes variables used in the retail pricing analyses. Panel B describes variables used in the household analyses. Panel C describes variables used in the inflation expectations analyses. For variables based on bundles, the fixed bundle construction period for each outcome year is the preceding calendar year, and the rolling bundle construction period for week t covers weeks $t-64$ through $t-13$.

Variable	Definition	Source
Panel A: Retail pricing analyses		
Fraction in delay (13w)	The share of matched import volume on voyages delayed by more than seven days relative to the 2015–2017 median for the route pair, measured over 13 weeks. In the retail analyses, the share is computed for the firm and four digit HS code mapped to the product module, with delayed volume assigned to the departure week. In the household analyses, delayed volume is assigned to the arrival week and firm or product-module delay rates are weighted by the household's bundle. In the expectations analysis, exposure is measured over the 13 completed weeks ending in the last fully completed week before the survey response, and product modules without observed import-arrival volume during the window receive zero delay exposure.	Spire AIS; Panjiva BoL; NielsenIQ HMS
ln(Price)	The natural log of the average price per item across stores, weighted by unit sales, for a firm×UPC×week.	NielsenIQ RMS
Absent	An indicator equal to one when a UPC is not observed in a week, conditional on the firm and product module having another observed UPC that week.	NielsenIQ RMS
ln(# UPCs)	The natural log of the number of active UPCs in a firm×product module×week.	NielsenIQ RMS
Premium	The share of active UPCs in a firm×product module×week whose mean log price over the full sample exceeds the firm×product module median.	NielsenIQ RMS
Low inventory buffer	An indicator equal to one when a firm's days of total inventory coverage are below the median across firms at the same reference date. Coverage is quarterly total inventory divided by quarterly cost of goods sold and multiplied by 91.25, the average number of days in a quarter. Missing total inventory is assigned zero when quarterly cost of goods sold is positive. The measure uses the latest fiscal quarter ending before the reference date. When no disruption indicator is active, the reference date is the beginning of the 13-week delay window. When one or more disruption indicators are active, the reference date is the first week in which that combination of indicators appears, and the classification remains fixed for all weeks assigned to that combination.	Compustat Quarterly
Product market delay exposure (13w)	The mean *Fraction in delay (13w)* among all other firms in the same product module and week, weighted by their unit sales shares during 2015–2017. The focal firm is excluded and the remaining weights are renormalized. The estimation sample begins in 2018 and requires at least two other firms with positive preperiod sales weights.	Spire AIS; Panjiva BoL; NielsenIQ RMS

Variable	Definition	Source
Rival delay	The equally weighted mean of *Fraction in delay (13w)* among other firms observed in the same product module and week.	Spire AIS; Panjiva BoL; NielsenIQ RMS
Treated × Post	In the retail event study, *Treated* equals one for firm×UPC units first exposed in the cohort week after at least 26 weeks without exposure. In the household event study, it equals one for household×product module units whose fixed bundle first has positive arrival-timed delay exposure after 26 weeks without exposure. *Post* equals one in event weeks zero through twelve.	Authors' calculations
ln(1 + Total volume in transit)	The natural log of one plus import volume in TEUs, dated to the vessel departure week. The retail event study uses the firm's volume for the corresponding four digit HS code. The household event study weights firm×product module volume by the fixed bundle's firm shares.	Spire AIS; Panjiva BoL
ln(1 + Total units sold)	The natural log of one plus total units sold across RMS stores. The retail event study measures units for a UPC in a week, and the household event study weights UPC units by the fixed bundle's UPC shares.	NielsenIQ RMS
Panel B: Household analyses		
ln(Bundle price)	The weighted average of weekly log UPC prices in a household's product module bundle, using UPCs and expenditure weights from the fixed bundle construction period.	NielsenIQ HMS
Laspeyres	The average log price change of rolling bundle UPCs from the bundle construction period to week t, weighted by expenditure during the construction period.	NielsenIQ HMS; NielsenIQ RMS
Tornqvist	The average log price change for UPCs in the rolling bundle from the bundle construction period to week t, weighted by the arithmetic mean of each UPC's normalized expenditure share in the two periods.	NielsenIQ HMS; NielsenIQ RMS
Paasche	The average log price change for UPCs in the rolling bundle from the bundle construction period to week t, weighted by expenditure in week t.	NielsenIQ HMS; NielsenIQ RMS
Rolling Laspeyres	The sum of each UPC's bundle construction period expenditure share multiplied by its log price change to week t.	NielsenIQ HMS; NielsenIQ RMS
Total price	The Laspeyres price change for the rolling bundle used in the additive inflation decomposition.	NielsenIQ HMS; NielsenIQ RMS
Price component	The average log price change for UPCs in the rolling baseline bundle, weighted in proportion to the mean of baseline and current expenditure levels. Current expenditure is zero when the household does not purchase the UPC in week t, and its current price is assigned using the price hierarchy described in the text.	NielsenIQ HMS; NielsenIQ RMS
Voluntary reallocation	The Shapley component of the difference between total price and the price component attributable to changes in expenditure weights from the bundle construction period to week t among UPCs not classified as stockouts.	NielsenIQ HMS; NielsenIQ RMS

Variable	Definition	Source
Forced stockout	The Shapley component of the difference between total price and the price component attributable to changes in expenditure weights for rolling bundle UPCs classified as stockouts in week t.	NielsenIQ HMS; NielsenIQ RMS
Any purchase	An indicator equal to one when the household purchases the product module in week t.	NielsenIQ HMS
ln(Spending)	The natural log of one plus household expenditure in the product module in week t.	NielsenIQ HMS
ln(Quantity)	The natural log of one plus household quantity purchased in the product module in week t.	NielsenIQ HMS
Module share	Product module expenditure divided by the household's total expenditure in week t.	NielsenIQ HMS
Bundle share	The product module's share of expenditure in week t on modules selected in the fixed bundle construction period.	NielsenIQ HMS
Excess quantity	Module quantity in week t minus the household's mean quantity in purchase weeks for that module during the fixed bundle construction period.	NielsenIQ HMS
Baseline UPC	The share of module expenditure in week t on UPCs purchased during the fixed bundle construction period.	NielsenIQ HMS
Other UPCs	The share of module expenditure in week t on UPCs not purchased during the fixed bundle construction period.	NielsenIQ HMS
Baseline brand	The share of module expenditure in week t on brands represented among UPCs purchased during the fixed bundle construction period.	NielsenIQ HMS
Cheap p25	The share of product module expenditure in week t allocated to UPCs whose average price in the fixed bundle construction period is at or below the 25th percentile for the product module.	NielsenIQ HMS
Cheap p50	The share of product module expenditure in week t allocated to UPCs whose average price in the fixed bundle construction period is at or below the median for the product module.	NielsenIQ HMS
Large size	An indicator equal to one when quantity purchased in the product module in week t exceeds the household's 90th percentile of purchase week quantities for that module in the fixed bundle construction period.	NielsenIQ HMS
Stockout weight	The share of bundle construction period expenditure allocated to UPCs without a price in week t and with a sharp decline in aggregate availability.	NielsenIQ HMS; NielsenIQ RMS
Sparse weight	The share of bundle construction period expenditure allocated to UPCs without a usable price in week t that are not classified as stockouts.	NielsenIQ HMS; NielsenIQ RMS
New UPC share	The share of total household expenditure in week t allocated to UPCs outside the rolling bundle.	NielsenIQ HMS
New module share	The share of total household expenditure in week t allocated to product modules outside the rolling bundle.	NielsenIQ HMS
Price coverage	The share of bundle construction period expenditure allocated to UPCs with a usable price in week t.	NielsenIQ HMS; NielsenIQ RMS

Variable	Definition	Source
Panel C: Inflation expectations analyses		
Expected inflation, next 12 months	The mean direct forecast of inflation over the next year within a household and survey wave.	NielsenIQ inflation expectations survey
ln(Bundle spending)	The natural log of household spending on the UPCs retained in the bundle during the rolling bundle construction period.	NielsenIQ HMS
ln(1 + Bundle UPCs)	The natural log of one plus the number of distinct UPCs retained in the household's bundle.	NielsenIQ HMS
Mean respondent age	The mean age of respondents from the same household in a survey wave.	NielsenIQ inflation expectations survey
Female respondent share	The share of respondents coded as female within the same household and survey wave.	NielsenIQ inflation expectations survey
Children present	An indicator equal to one when children are present in the household.	NielsenIQ HMS
Household size	The number of household members.	NielsenIQ HMS
Hispanic origin	An indicator equal to one when the household reports Hispanic origin.	NielsenIQ HMS
Race	The household's reported race category.	NielsenIQ HMS
Middle income	An indicator equal to one when the midpoint of the household's NielsenIQ income bracket is in the middle third of the common regression sample's household income distribution. Low income is the omitted group.	NielsenIQ HMS
High income	An indicator equal to one when the midpoint of the household's NielsenIQ income bracket is in the upper third of the common regression sample's household income distribution. Low income is the omitted group.	NielsenIQ HMS
State	The household's state of residence.	NielsenIQ HMS
Survey order 2–4	An indicator equal to one when the current survey wave is the household's second through fourth observed wave. The first observed wave is the omitted group.	NielsenIQ inflation expectations survey
Survey order 5–9	An indicator equal to one when the current survey wave is the household's fifth through ninth observed wave. The first observed wave is the omitted group.	NielsenIQ inflation expectations survey
Survey order 10 or more	An indicator equal to one when the current survey wave is at least the household's tenth observed wave. The first observed wave is the omitted group.	NielsenIQ inflation expectations survey