Tariff Pass-through into Retail Prices: Measurement, Replacement, and Shrink-flation*

Liang Bai
King's College London

Torsten Søchting Jaccard
Vancouver School of Economics

Sebastian Stumpner
Banque de France

This Draft: July 13, 2026
First Draft: Dec 22, 2025

Abstract

To what extent do import tariffs pass-through into retail prices? Existing estimates for the U.S.-China trade war present a puzzle: pass-through into U.S. border prices was complete, with seemingly no effect on retail prices. We combine barcode country-of-origin data with retail scanner data to address this apparent discrepancy in the context of consumer goods. Within barcodes, the tariff elasticity of retail prices is 0.14, yet the tariff elasticity of aggregated prices per unit weight, a measure similar to the unit values recovered from customs data, is 0.29. This discrepancy is attributable to product replacement bias: pass-through occurs predominantly via product turn-over towards barcodes with higher price per unit weight, half of which is achieved via smaller package size. We estimate a foreign cost share of U.S. retail imports consistent with high, if not complete, pass-through into retail prices. U.S. consumers therefore bore almost the entire incidence of tariffs levied on consumer goods.

JEL Categories: E31, F13, F14, F31, L11.

*Bai: liang.bai@kcl.ac.uk, Jaccard: tjaccard@mail.ubc.ca, Stumpner: sebastian.stumpner@banque-france.fr.

We thank Pierre Lambelet for research assistance. For many comments and suggestions we are grateful to conference and seminar participants at Bristol, CBS, Harvard/MIT, UCLA, China Economic Study Group, EcHO, ETSG, Kent, KCL, Chinese University of Hong Kong, Kiel/Goettingen CEPR workshop, Oslo, Patras Applied Micro workshop, Swansea, RES, Banque de France, the ECB-Bank of Canada conference on the future of global trade, the BdI-BOE-BdF international macro workshop, ECB Trade Expert Network, ESCB research cluster 2 workshop, RMET 2026, and the Venice Summer Institute. The analysis in this paper is our own, calculated 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 our own 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. All errors are our own.

1 Introduction

Who pays for import tariffs? Basic economic theory suggests that the rate of tariff pass-through - the extent to which the buyer price reflects a changing tariff - is the key metric for quantifying tariff incidence across foreign producers and domestic consumers. Empirical estimates of tariff pass-through, however, are scarce, given the remarkable liberalization of global trade policy in the post-war period.

The U.S.-China trade war of 2018-2019 therefore offers a unique setting in which to study both the outcomes of this particular policy episode, but also the broader incidence of tariffs across importers and exporters, foreign producers and domestic consumers. Almost all research in this vein convincingly points to the same finding: U.S. tariffs applied to Chinese exports in 2018-2019 were passed through completely into prices paid by U.S. importers (Amiti et al. 2019; Fajgelbaum et al. 2020; Cavallo et al. 2021; Gopinath and Neiman 2026; Hinz et al. 2026). And yet there remains remarkably little evidence for any tariff pass-through into prices faced by U.S. households at U.S. stores.[1]

In this paper we use barcode-level scanner data to estimate the effect of these tariffs on final consumer prices. We adopt a broad definition of pass-through that encompasses not only tariff-induced price changes for existing barcodes, but also the extent to which tariffs might generate higher prices via product replacement. For every 10% increase in tariffs, retail prices (per unit weight) increase by 2.9%. To the best of our knowledge, this is the first paper to document price effects of this magnitude for consumer goods during the U.S.-China trade war. We show that our estimate of the tariff elasticity of retail prices is compatible with complete pass-through in price levels, despite being incomplete in terms of percentage changes. Since tariffs are only applied to import values recorded at the border, a low foreign cost share of retail imports generates muted estimates of tariff pass-through into retail prices. By estimating this foreign cost share directly, we find that tariff pass-through into retail prices was high, if not complete, such that the incidence of import tariffs fell largely on U.S. consumers.[2]

[1] Cavallo et al. (2021) estimate pass-through into U.S. retailer prices of only 3.5%, and this estimate is not distinguishable from zero at standard confidence levels. See Fajgelbaum and Khandelwal (2022) for more discussion.

[2] Throughout this analysis, we focus exclusively on tariff waves 3 and 4a associated with the U.S.-China

We combine several datasets to estimate the response of consumer prices to tariffs. Our primary source is NielsenIQ data on barcode-level prices and sales across approximately 70 retail chains. We merge these data with firm identifiers and barcode-level country-of-origin information and use a crosswalk between NielsenIQ product modules and HS8 tariff codes to link products to the 2018/19 waves of tariffs on Chinese imports. A distinguishing feature of the dataset used in this paper is that we observe both price and quantity data for detailed product definitions – barcodes – over time, such that we can compare tariff effects on prices at varying levels of aggregation. Specifically, we estimate both within-barcode tariff elasticities and within-product elasticities, in which we define a product as an aggregation at the firm-category-country level. For example, the average price per unit weight of a specific firm's barcodes of deodorant produced in China and sold at a specific retail chain in the U.S.

We begin by estimating the retail price elasticity of tariffs measured using within-product unit values via a triple-difference design comparing changes in unit values within production origins and within product categories. While our headline estimate of 0.29 is considerably larger in magnitude than existing tariff elasticity estimates associated with U.S. retail prices, it still appears to fall short of the full pass-through consistently recovered using customs data. We spend the following sections of this paper providing evidence that in fact our estimate of the retail unit value tariff elasticity is entirely compatible with both sets of existing estimates, and that all results are compatible with high, if not full, pass-through into U.S. retail prices.

First, we show that the tariff elasticity of retail unit values may be smaller in magnitude than similar percentage pass-through estimates at the border despite reflecting the exact same magnitude of pass-through in levels. In the presence of some cost wedge associated with post-tariff distribution, such as retailing mark-ups, retailing costs and/or intra-national distribution costs, estimates of pass-through into retail prices must be scaled by the value share of retail imports that is attributable to pre-tariff activity—the foreign cost share—in order to be comparable to estimates of pass-through at the border.[3]

trade war. Waves 1 and 2 primarily targeted goods with poor representation in scanner data, such as intermediates and capital goods. Section 2 provides more detail.

[3]For example, if the cost share of domestic economic activity associated with retail goods, such as intra-national distribution costs and retailing costs (land, labor, etc.), is 70%, a tariff is only applied to 30% of the value of the good. Full pass-through at the border would imply 30% pass-through into retail prices. See Sangani (2025) for a similar argument regarding commodity shock pass-through.

Despite the crucial role played by the foreign cost share in shaping pass-through of international shocks to retail prices, estimates are scarce. Existing estimates tend to place the foreign cost share of retail imports between 30%-50% and between 30%-40% for U.S. retail imports of consumer goods (Berger et al. 2012; Flaaen et al. 2025; Errico 2026). However, the product categories reported in (Berger et al. 2012) map only imperfectly into a subset of our product categories and are based on older data (1994-2007). We therefore carry out our own estimation of the foreign cost share, by mapping U.S. customs unit values to scanner data unit values across product categories and origin countries, with an average recovered estimate of 32% for consumer retail goods imported from China.[4] We provide a qualitative validation of this estimate by documenting a positive relationship between our foreign cost share estimates at the product category level and category-specific levels of tariff pass-through into retail unit values. We therefore provide evidence that our unit value elasticity of 0.29 overlaps considerably with our estimated range of the foreign cost share of U.S. retail imports, such that this estimate is entirely compatible with the estimates of full pass-through recorded using customs data.

Second, we show that while our headline estimate aligns with existing estimates of full pass-through into border prices, it is also entirely compatible with the null estimate of product-specific tariff pass-through reported by Cavallo et al. (2021). Specifically, we find that the retail unit value tariff elasticity is only 0.11 when we restrict the sample to continuing barcodes.[5] Similarly, when we estimate the tariff elasticity using within-barcode variation for this continuing panel of barcodes, we recover an estimate of 0.14.[6] However this within-barcode estimate is reduced to zero once we expand our sample to an unbalanced panel of barcodes, similar to Cavallo et al. (2021).

Taken together, these estimates suggest that a significant share of the tariff elasticity of retail prices occurs via product replacement, such that within-barcode estimates differ considerably from unit value estimates. We decompose our unit value tariff elasticity esti-

[4]In contemporaneous work, Flaaen et al. (2025) also consider the importance of the foreign cost share for tariff pass-through in their study of EU wine exports to the U.S. They estimate a foreign cost share for EU wine that is similar to our estimate, at 38%.

[5]We define a continuing barcode as one with non-zero sales throughout the entirety of our sample period.

[6]We confirm that our relatively low estimates of within-barcode pass-through are not driven by endogenous price spillover effects to non-tariff goods, tariff "smoothing" by retailers, or the effect of tariff shocks raising all production costs and thereby reducing the effect of product-specific tariffs.

mate into changes in the numerator—what we call "sticker prices"—and the denominator, or average package size.[7] We recover a tariff elasticity of average package size of -0.14, accounting for just under half of the unit value tariff elasticity. We refer to this effect as "shrink-flation", and confirm a lack of pre-trends, suggesting this decrease in size is causally attributable to the tariff change.[8] Our findings therefore suggest that within-barcode measures of pass-through are potentially subject to product replacement bias (Nakamura and Steinsson 2012), and such estimates may under-state the true pass-through of import tariffs into consumer prices.

Still, within-barcode estimates have the benefit of a clear counterfactual price, whereas entry and exit could lead to compositional shifts in the set of products offered, thereby generating price elasticity estimates which are not necessarily relevant for consumer welfare.[9] There are several reasons to think that the unit value tariff elasticity estimates we recover are likely welfare relevant, however. First, we estimate these effects within narrow firm-country-category-retailer cells, such that changes in average package size occur within fairly narrow product definitions. Second, it would have to be that quality is almost perfectly correlated with the inverse of package size in order for our "shrink-flation" estimates to be driven by an endogenous quality response, rather than product replacement. Finally, we document a positive tariff elasticity of both barcode entry and exit, suggesting that tariffs led to greater barcode churn, consistent with a story of product replacement.

Finally, we consider the role of buyer and seller power in shaping heterogeneity in observed tariff elasticities across products and retailers. Specifically, we interact tariff changes with both the share of expenditure within a given retailer attributable to a specific firm (the seller share) and the share of expenditure within a given firm attributable to a specific retailer (the buyer share). In line with notions of monopoly and monopsony power, we find elevated tariff

[7] By "sticker price", we refer to the average price per barcode sold, rather than the average price per unit of weight sold, which we refer to as unit values. Section 4 provides an exact decomposition of unit value pass-through into a sticker price component and an average package size component.

[8] We provide heterogeneity analysis to compare the use of "shrink-flation" across product modules with high and low initial variation in size offered: while both module types exhibit similar unit value tariff elasticity estimates, package size reduction accounts for less than a fifth of this elasticity in modules with below-median pre-tariff variance in sizes, and just over two thirds of the unit value tariff elasticity in above-median modules.

[9] As discussed in Metzler (1949) and documented by Ludema and Yu (2016); Chen and Juvenal (2016); Freitag and Lein (2023), tariffs might generate an endogenous exit of low-quality varieties, thereby mechanically increasing unit values even in the absence of any within-variety pass-through.

price effects for above-median seller shares and reduced tariff price effects for above-median buyer shares. Importantly, the unit value tariff elasticity associated with small producers selling at large retailers is effectively zero, while the same elasticity associated with large producers selling at small retailers is roughly 0.40, suggesting a broad range of tariff pass-through levels across producers and retailers in the U.S. market according to producer and retailer size.

Related literature: This paper fits into a growing empirical literature on the economic consequences of the U.S.-China trade war, with a particular focus on the price effects of this event (Amiti et al. 2019; Fajgelbaum et al. 2020; Flaaen et al. 2020; Cavallo et al. 2021).[10] The closest papers to ours are Cavallo et al. (2021) and Flaaen et al. (2025) who study pass-through in the context of, respectively, consumer retail goods and E.U. wine exports. Our null estimate of within-barcode pass-through for an unbalanced panel matches findings in Cavallo et al. (2021), who consider a similar specification. By estimating pass-through into retail unit values, we show that product replacement can align this result with full pass-through at the border. Similar to Flaaen et al. (2025) and Errico (2026), we show that domestic wedges associated with inter-mediation of imported goods generate muted price effects in retail data that are not necessarily inconsistent with full pass-through at the border.[11]

Our results build on the product replacement bias documented in Nakamura and Steinsson (2012).[12] Producers and/or retailers may perceive product replacement as a mechanism for increasing prices which generates less negative sentiment among consumers (Rotemberg 2005; Nakamura and Steinsson 2011), and a growing literature suggests that "shrink-flation"

[10] See Fajgelbaum and Khandelwal (2022) for a full literature review of the first U.S.-China trade war. Key papers include Charbonneau and Landry (2018); Fetzer and Schwarz (2021); Benguria et al. (2022); Blanchard et al. (2024); Handley et al. (2025); Hankins et al. (2025). See Bown (2021) for a full timeline of the events which constitute the U.S.-China trade war. See Cavallo et al. (2025) and Amiti et al. (2026) for estimates of tariff pass-through following the "Liberation Day" tariffs announced on April 1, 2025.

[11] Ganapati and Hottman (2026) argue that pass-through to U.S. importers was incomplete despite import prices increasing by the tariff change: scale economies generate price effects which exceed the incidence of import tariffs. From the perspective of retail prices, the cost shock remains equivalent to full tariff pass-through at the border, and we therefore maintain full pass-through at the border as our benchmark.

[12] While product entry is beneficial to consumers (Feenstra 1994; Broda and Weinstein 2006), product replacement bias suggests not all entry is equal. If product entry serves only as a means to increasing the price of an otherwise identical product, such "entry" should not be considered relevant for consumer welfare.

may be a particularly common method by which firms achieve higher prices when replacing products (Çakır and Balagtas 2014; Chalioti and Serfes 2024; Kim 2024; Janssen and Kasinger 2026). We similarly contribute to the literature quantifying the foreign cost share of domestic retail goods (Burstein et al. 2003; Devereux et al. 2003; Anderson and Van Wincoop 2004; Goldberg and Campa 2010; Berger et al. 2012), pass-through of exchange rate shocks to retail prices (Auer et al. 2021), and the combination of pass-through estimates along the distribution network (Nakamura and Zerom 2010; Koujianou Goldberg and Hellerstein 2013; Flaaen et al. 2025). By studying the endogenous response of retailers to tariffs, we also contribute to the literature documenting the relationships between foreign producers, domestic retailers, and trade shocks (Raff and Schmitt 2006, 2009; Eckel 2009; Basker and Van 2010; Basker et al. 2012; Raff and Schmitt 2012, 2016; Cole and Eckel 2018; Jaccard 2023).

Finally, we provide novel empirical evidence regarding pass-through by seller shares and buyer shares within the context of consumer retail goods (Atkeson and Burstein 2008; Amiti et al. 2016; Alviarez et al. 2023). We find that larger producers pass-through a higher share of the tariff cost shock, while larger retailers reduce the level of pass-through to consumer prices. Furthermore, the product replacement mechanism we document appears to be most relevant for larger producers, because the gap between unit value and product-specific price pass-through is largest for firms with some monopolistic power.

Section 2 describes our data. Section 3 provides our headline estimates of pass-through into retail unit values alongside estimates of the foreign cost share. Section 4 decomposes unit value pass-through into components associated with continuing barcodes, the full panel of barcodes, and compares within-barcode to within-product estimate. Section 5 estimates heterogeneous tariff effects according to buyer and seller power, and Section 6 concludes.

2 Data

We bring together several data sources for the empirical analysis of this paper: i) barcode-level price and expenditure from NielsenIQ Retail Scanner Data, ii) barcode-level country-of-origin data from both Label Insight and scraped from the website of a large U.S. retailer, iii) firm identifiers from GS1, iv) tariff data and import unit values at the eight-digit Harmonized System (HS) level from the U.S. Trade Representative (USTR), and v) a concordance between NielsenIQ product modules and HS-10 digit codes (constructed by the authors).[13]

2.1 Price and Expenditure

For information on prices and expenditures at the barcode-level, we make use of the NielsenIQ Retail Scanner Panel over the period 2015-20. Products from all NielsenIQ departments are included in the data, such as alcoholic beverages, dairy, deli, dry grocery, fresh produce, frozen foods, packaged meat, general merchandise, health and beauty care, and non-food grocery. Depending on the year, data are included from approximately 30,000-50,000 participating grocery, drug, mass merchandiser, and other stores.

The raw data is available at the barcode-store-week level. For the analysis, we build a panel dataset of prices and expenditures at the barcode-retailer-month level. To abstract away from changes in the set of reporting stores over time, we focus on the set of stores that report positive sales in every year of our sample period ($n = 29,054$). To compute sales and prices at the retailer level, we make use of the NielsenIQ retail chain identifiers.

2.2 Country of Origin

One limitation of the Retail Scanner Data when used in empirical studies of international trade is the lack of information regarding a product's country of origin. To overcome this, we supplement scanner data with two separate data sources that provide origin country information at the barcode level. First, we use a proprietary dataset purchased from Label Insight, a private firm specializing in the provision of detailed product attribute data. Second, we scrape data from the website of a large U.S. retailer, which further improves our product coverage.

[13] This is an extension of the concordance between NielsenIQ product modules and HS-6 digit codes used in Bai and Stumpner (2019).

2.3 Tariff Data

There were four waves of tariff changes associated with the 2018-2019 trade war, each implemented separately. Waves 1 and 2 (begining in July and August 2018 respectively) targeted mostly machinery and chemicals, and were therefore limited in scope.[14] By contrast, Waves 3 and 4a (beginning in September 2018 and September 2019) covered a much broader set of goods - $206bn worth of imports, across 5,756 HS-8 digit codes, were covered in Wave 3, while an additional $114bn, across 3,229 HS-8 digit codes, were covered in Wave 4a. Figure 2 plots the tariff rates by wave, as well as the average tariff rate across all imports from China. The targeted categories account for roughly two thirds of all US imports from China in 2017, while the remaining third (e.g. list 4b) were left as a threat point during the U.S.-China negotiations and were eventually left untaxed after the signing of the Phase One Deal in January 2020.

We use data from the official tariff lists published by the United States Trade Representative (USTR) for each wave. Figure 1 shows the weighted average tariff on imports from China between 2017-2021, using 2017 import volumes as weights. While the average tariff increase during 2017-19 for NielsenIQ consumer product modules was significantly lower than for other targeted categories, it was still over 7 percentage points, compared to a pre-2018 baseline average tariff rate of around 3%.

The main part of our analysis focuses on the comparison between Wave 3 and the never-taxed categories. Waves 1 and 2 focused largely on intermediate goods and have a poor representation in NielsenIQ scanner data. While our data contains a sufficient number of products across both Waves 3 and 4a, our main results focus on Wave 3 (the largest wave by trade volume), with Wave 4a results provided as a robustness. We do this for two reasons. First, as shown in Figure 2, the wave 3 tariffs experienced two increases: an increase from 2.5% to 12.5% in September 2018, and a further increase to 27.5% in May 2019. Wave 3 goods then maintained this heightened tariff rate through 2020. Wave 4a categories, on the other hand, experienced a tariff increase from 5% to 20% in August 2019 which was followed by a reduction to 12.5% in early 2020. The dynamic effects associated with a short-lived, dramatic, increase in tariff rates followed by a mediation of rates to 12.5% complicates effect

[14] Wave 1 tariffs covered $30bn worth of imports, while Wave 2 tariffs covered an additional $15bn.

Figure 1: Trade War Tariffs by Dataset

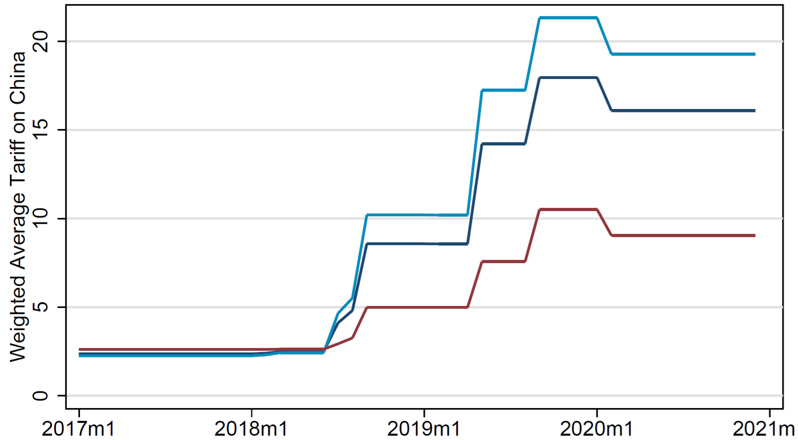

Figure 2: Trade War Tariffs by Wave within NielsenIQ Sample

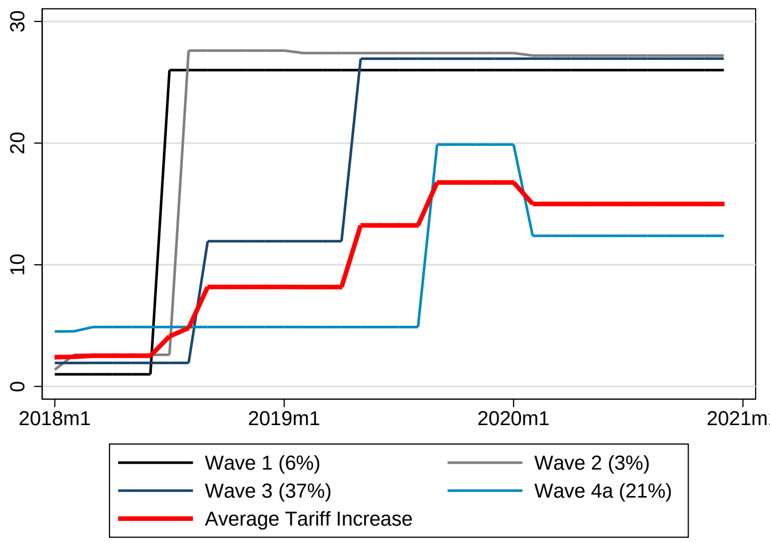

Notes: Figure 1 provides the average tariff rate applied by the U.S. to imports from China over time. All products are in dark blue, products outside of the NielsenIQ scanner data are in light blue, and products within NielsenIQ scanner data in dark red. Figure 2 provides the average tariff applied to Chinese imports over time within NielsenIQ categories, and denotes the total share of categories within the NielsenIQ data associated with each tariff wave during the 2018-2019 trade war. These shares do not sum to one as the remaining NielsenIQ categories had not tariff change, and therefore constitute a control group.

interpretation, particularly given the lagged effects of tariffs we document below. Second, since wave 4a tariff rates stabilize at the start of 2020, there is only a short period prior to the Covid-19 pandemic with which to estimate tariff effects. While we control for country-month and product category-month shocks, as well as seasonality controls specific to tariff-affected goods, it is still possible that the Covid-19 pandemic induced a distribution of demand

and supply shocks that are unobserved and potentially correlated with treatment, which we cannot account for here. Still, we show below that at least qualitatively, and to a large extent in magnitudes as well, estimates associated with the Wave 4a tariffs were similar.

One potential concern with the tariff data is that actually applied tariffs may differ from initially imposed tariff rates. Chor et al. (2025) document the relevance of tariff exemptions that companies could apply for. We evaluate whether exclusion requests are important in our context using the micro data on exclusion requests from regulations.gov. We find that out of roughly 52.4k exclusion requests spanning waves 1 to 4a, roughly 6.5k have been granted which corresponds to a success rate of 12.4%. However, most granted requests were concentrated in waves 1 and 2 (70% of all granted requests), even though those waves only corresponded to about 9% of US imports from China. While on average one out of three requests was granted in waves 1 and 2, only out of twenty was successful in waves 3 and 4a. The approval rate drops down further to about 4% once we focus on final products in waves 3 and 4a, which are the products that we study.[15] We therefore conclude that, while tariff exclusions were important overall, they most likely do not have a relevant effect on our estimates.

2.4 Concordance between NielsenIQ and tariff data

Barcode price and expenditure data in NielsenIQ are organized in roughly 1,200 different product modules. In order to apply tariff rates to barcodes we require a mapping between product modules and HS codes. We build on the concordance in Bai and Stumpner (2019) that maps NielsenIQ product modules to HS6 codes, forming 359 product categories. The number of product categories is lower than the number of modules due to many cases where several product modules and several HS6 codes have to be combined to form one consistent "category."[16] There are two reasons why this HS6-level concordance is unsuitable for our

[15]The much lower approval rates in waves 3 and 4a are likely related to the fact that the products in these waves were much less likely to meet official approval criteria. For instance, for the Wave 3 tariff exclusion procedure, the Federal Register Vol 83, N. 133 states that exclusion requests must specify whether the particular product is available only from China, whether the imposition of tariffs would cause severe economic harm to the requestor or other U.S. interests, and whether the product is related to any Chinese strategic industrial program such as "Made in China 2025."

[16]For example, the category "cameras and accessories" maps into four separate product modules (6057, 6060, 7410 and 7420), and five distinct HS6 codes (852540, 900211, 900653, 900691 and 900791).

setting. First, the original concordance matches product module codes to HS6 codes of the 2002 HS revision, while the tariffs relevant to our study were applied using HS codes of the 2017 revision. Second, the concordance links NielsenIQ modules to six-digit HS codes, while tariffs are typically applied at the eight-digit (HS8) or, in some cases, ten-digit (HS10) level.

To overcome this problem, we translate the concordance in Bai and Stumpner (2019) to the HS2017 revision and refine this concordance, wherever possible, to map to 8-digit or 10-digit HS codes. This results in a more granular concordance between modules and HS codes, creating 445 different categories. Since these categories may map to several HS8 codes, it is possible that the underlying HS8 codes of a given category belong to different tariff waves, creating measurement error via aggregation of heterogeneous tariff levels within a given category. We therefore restrict our attention to a set of "clean" categories, for which all underlying HS8 codes belong to the same tariff wave. Our final sample consists of 200 product categories, of which 39 have never been treated, and 161 belong to Wave 3. We then attribute unique tariff rates to the product modules that match to these clean categories, which provides us with a set of 476 product modules, 91 of which have never been treated and 385 that belong to Wave 3. We use these product modules to compute average prices and control for sectoral trends in our analysis.[17]

2.5 Sample Construction and Descriptive Statistics

After restricting to a balanced panel of stores and aggregating the sales data to the barcode-retailer-month level, our sample consists of 787k barcodes in 2017, sold by 68 retailers, and summing up to roughly $219B USD of expenditure in 2017. This is reduced to 206k barcodes in 2017 (out of which 8.4k are Chinese), or 26% of all barcodes in the data once we restrict attention to barcodes with: i) valid firm identifiers, and ii) country of origin information. Among others, this stage excludes retailer-owned private label barcodes, since these barcodes are masked by Nielsen and therefore can not be matched to GS1 firm identifiers nor to country of origin information. However, because Label Insight prioritizes products with the highest volume of sales, our matched barcodes account for 63% ($136B USD) of total expenditure

[17]Wave 4a consists of an additional 104 categories that map into 285 product modules.

in the retail panel.[18] As shown in Figure A.1, the distribution across NielsenIQ departments (both in terms of the number of UPCs and expenditure) looks similar for the matched and broader samples.

3 Tariff Pass-through into Average Prices

We begin by estimating the elasticity of U.S. retail prices with respect to changes in U.S. import tariffs during the 2018-2019 trade war. In order to allow for the broadest possible set of mechanisms associated with endogenous retail price changes, we begin by studying pass-through into aggregated unit values at the firm-module-country-retailer-month level. We denote firms as f, product modules as m, origin countries as j, retail chains as r, and periods as t. Throughout, we refer to the firm-module-country triplet fmj as k to ease notation, such that krt denotes our unit of observation in this section.

3.1 Empirical Model

Let y_{krt} denote some observed outcome associated with the triplet fmj sold at retailer r in month t. We estimate the effect of tariffs on unit values u_{krt} and quantities sold q_{krt}, where both measures are adjusted for the underlying size of barcodes which constitute each fmj triplet. That is, we construct $q_{krt} = \sum_{i \in \Omega_{krt}} q_{irt}$ where Ω_{krt} denotes the set of barcodes manufactured by firm f and sold at retailer r in month t within the module-country pair mj. q_{irt} denotes the total quantity sold of barcode i at retailer r in common units such as liters or kilograms. We compute this quantity using data on the raw number of packages sold, $\tilde{q}_{irt}$ and the quantity per package, w_i, such that $q_{irt} = \tilde{q}_{irt} w_i$. Moving forward all variables with a tilde denote variables which are unadjusted for product size.[19] Average prices, which we

[18] Flaaen et al. (2025) also study the pass-through of tariffs to consumer prices, focusing on very detailed data on wine imports. Our categories include wine sales, which accounts for \$3.9B USD, or 2.8% of all sales in our data.

[19] Not all barcodes within a given product module exhibit the same unit of measure. When possible, we convert units into a standard benchmark, such as "OZ" to "mL." Still, some recorded units, such as "Counts," cannot be readily converted to comparable units. For each product module, we calculate the most common unit of measure in our pre-tariff period, 2017, and drop all barcodes with units of measure which cannot be readily converted to this predominant unit of measure. In practice, this data restriction maintains over 95% of expenditure in our raw data. A benefit of using barcodes as our definition of a product is that any change to the size of a barcode, and therefore the product label, will lead to a new barcode such that w_i is constant

refer to as unit values, are simply expenditure v_{krt} divided by the quantity sold:

$$u_{krt} = \frac{v_{krt}}{q_{krt}} = \frac{\sum\limits_{i \in \Omega_{krt}} p_{irt} q_{irt}}{\sum\limits_{i \in \Omega_{krt}} q_{irt}} = \frac{\sum\limits_{i \in \Omega_{krt}} p_{irt} \tilde{q}_{irt} w_i}{\sum\limits_{i \in \Omega_{krt}} \tilde{q}_{irt} w_i}.$$

We denote the tariff level common to all barcodes within module m and produced in origin country j in period t as τ_{jmt}. By applying OLS to Equation 1, we recover the parameter of interest, β:

$$\ln y_{krt} = \beta \ln(1 + \tau_{jmt}) + \alpha_{krt} + \gamma_{k,q(t)} + \epsilon_{krt}. \tag{1}$$

When the dependent variable represents unit values, β captures the elasticity of retail unit values to import tariffs, which we would expect to be bounded below by zero and above by one. An estimate of one would imply that a 10% increase in tariffs generates a 10% increase in retail unit values. When the dependent variable represents total quantity sold, β captures the quantity elasticity of tariffs, which we expect to be negative. For ease of exposition, we collect all fixed effects within the vector α_{krt}, such that α_{krt} represents the following set of fixed effects:

$$\alpha_{krt} = \alpha_{kr} + \alpha_{rt} + \alpha_{c(k)t} + \alpha_{j(k)t}.$$

We include fixed effects at the $fcjr$, or kr, level as this is the unit of analysis, such that we control for any level differences in prices and quantities across firm-category-country-retailer units of observation. We then include fixed effects which account for shocks at the retailer-month level, category-month level, and origin country-month level. The category-month shocks account for any demand or supply shocks which are common to barcodes within a given product category. Country-period fixed effects isolate the role specifically of tariffs, by comparing the effect of tariff changes on Chinese goods to other imports from China for which tariffs did not change. Importantly, these fixed effects control for exchange rate changes which may be endogenous to the tariff change as well as any tariff effects which alter perceived future tariff rates associated with all goods produced in China. Finally, we include retailer-period fixed effects to account for potentially differential exposure across retail chains during the trade war.

---

across all observations of barcode i.

The vector $\gamma_{k,q(t)}$ constitutes seasonality controls, which we include in order to account for seasonal fluctuations which we observe in the tariff-targeted goods prior to the U.S.-China trade war. These controls constitute a dummy variable for each quarter of the calendar year interacted with the treatment dummy. Such seasonality is not captured by the fixed effects described above as these shocks are not specific to Chinese exports or specific categories, but the specific interaction of Chinese varieties within tariff-targeted categories. We therefore include a vector of seasonality controls, in which $\Gamma_k = 1$ is a treatment indicator equal to one if the combination of product category $c(k)$ and origin country $j(k)$ are ultimately targeted with a tariff by the U.S. We interact this indicator with indicator variables at the quarter level, such that $q(t) = 1$ denotes that period t is in the first quarter.

$$\gamma_{k,q(t)} = \sum_{q'=1}^{4} \gamma_{q'} \times \Gamma_k \times \mathbf{1}[q(t) = q']$$

The identification strategy employed here is effectively a triple-difference: we compare outcomes between Chinese goods and non-Chinese goods across tariff and non-tariff product categories before and after tariff implementation. As such, we can adjust our specification to an event study so as to check for the presence of pre-trends. To do so, we apply OLS to Equation 2 and plot the estimate vector β_γ.

$$\ln(y_{krt}) = \sum_{\tau=\underline{T}}^{\bar{T}} \beta_\tau \times \Gamma_k \times \mathbf{1}[t = T - \tau] + \alpha_{krt} + \gamma_{k,q(t)} + \epsilon_{krt} \tag{2}$$

T denotes the date of the first tariff change. While our primary focus in this paper are the wave 3 tariffs, we also generate event study figures for product categories targeted by the wave 4a tariffs. In all cases, the control group is the never-treated set of product categories.

In constructing variables at the firm-category-country-retailer level, we restrict ourselves to the set of firm-category-country-retailer quadruplets present in each period of our sample, such that at the krt level we estimate β using a balanced panel. Crucially, as we discuss in Section 4.2, we construct unit values and aggregate quantities using all underlying barcodes within each krt observation, such that while the estimating sample is balanced at the krt level, the underlying set of barcodes used to construct aggregate variables is unbalanced.

Table 1: Tariff Pass-through into Unit Values and Tariff Substitution

	(1) Log Unit Value	(2) Log Quantity	(3) Log Unit Value	(4) Log Quantity
$\log(1+\tau)$	0.293***	-0.457***	0.280***	-0.480***
	(0.069)	(0.148)	(0.068)	(0.124)
Exclude 2020			✓	✓
R^2	1.00	0.98	1.00	0.98
Observations	4,183,892	4,183,892	3,139,829	3,139,829

*Notes: This table provides estimates of β from applying OLS to Equation 1. Data are aggregated to the firm-category-country-retailer level. All models include fixed effects by firm-country-category-retailer, retailer-month, category-month, and country-month. Columns (1) and (2) estimate, respectively, pass-through into unit values and tariff substitution in quantities for the entire sample period. Columns (2) and (4) restrict the sample to only time periods prior to 2020, such that our sample is January 1, 2017, to December 31, 2019. Standard errors are clustered three-way at the category, country, and month level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.*

3.2 Tariff Elasticity of Retail Unit Values

Table 1 provides our baseline estimates of retail unit value and retail substitution tariff elasticities associated with the wave 3 tariffs applied to imports to the U.S. from China. Columns (1) and (2) represent our preferred estimates, in which the entire sample from January 1 2017 to December 31 2020 is included in our estimation. Column (1) provides the headline result of this paper: retail unit values exhibit a tariff elasticity of 0.293 (0.069). Taken at face value, this estimate implies tariff pass-through into retail prices was 29.3%. To our knowledge, this paper is the first to document a non-zero tariff elasticity of U.S. retail prices for consumer goods during the 2018-2019 U.S.-China trade war. We document significant substitution away from tariff-targeted goods, with a quantity elasticity estimate in Column (2) of -0.457 (0.148). The implied price elasticity of these two estimates is $-(-0.457)/0.293 = -1.56$, which is plausible given the relatively small estimates of price elasticities associated with consumer packaged goods (Faber and Fally 2021) and the relatively short time-horizon we consider here.

Columns (3) and (4) show that these estimates are broadly robust to excluding the Covid-19 pandemic from our sample. As discussed earlier, the onset of the pandemic in early 2020 could plausibly lead to unobserved demand and supply shocks which may be correlated with treatment. At the same time, the lack of sample periods post-tariffs but prior to the onset of the pandemic limits our ability to recover the dynamic effects of tariffs on unit values and quantities sold when 2020 is removed from our sample. Still, we recover almost identical estimates, with a unit value elasticity prior to 2020 of 0.280 (0.068) and a quantity tariff elasticity of -0.480 (0.124).

Figure 3 provides event study estimates for unit values (top) and quantities sold (bottom). Each figure includes two vertical lines which denote the two tariff changes associated with tariff wave 3: an initial increase from 2.5% to 12.5% in September 2018 and a further increase to 27.5% in May 2019. There are no appreciable pre-trends in either unit values or quantity units sold prior to the first wave 3 tariff change in September 2018. This is reassuring, as earlier tariff waves (1 and 2) had already occurred, which could have led to anticipatory effects on later tariff waves. We also do not observe any particularly striking behavior during the pandemic, suggesting that the fixed effects implemented in Equation 1 are able to difference out any Covid-related shocks.

For both unit values and quantities, the initial tariff change exhibits modest effects, with the four quarters after the second tariff increase highlighting the dynamic effect of these tariff changes on price and quantities. For unit values, this tariff-induced price differential stabilizes at a roughly 8% increase. For quantities, the relative decrease in tariff-affected quantities sold stabilizes at roughly -17%. We therefore take the wave 3 estimates in Table 1 associated with the entire sample as our headline estimates, such that the tariff elasticity of retail unit values is 0.293.

Robustness across tariff waves. As noted above, the wave 4a tariffs pose a challenge from an identification standpoint, as these tariffs were implemented and then reduced only months prior to the onset of the Covid-19 pandemic. Thus, we focus our analysis in this section on the Wave 3 tariffs, using categories which were never targeted by any wave as control. Still, the results derived here broadly hold for the wave 4a categories as well. Table A.1 confirms that unit value tariff elasticity associated with wave 4a was 0.380 (0.082). The wave 4a quantity elasticity estimate of -0.512 (0.295) is more noisily estimated than the wave 3 estimate, but implies a price elasticity of 1.35, which is similar to the implied price elasticity for wave 3. Figure A.2 provides the event study figures associated with wave 4a unit values and quantities, again confirming a lack of pre-trends in both cases.

Figure 3: Event Study Estimates for Unit Value Pass-through (Wave 3)

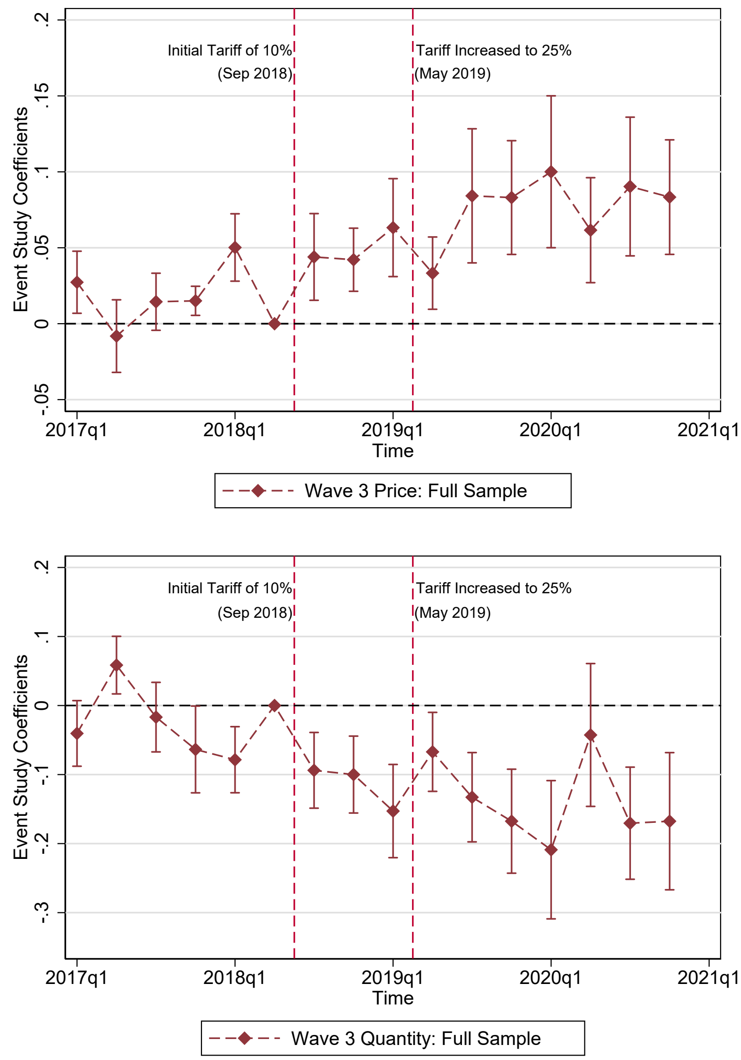

Notes: This figure provides estimates associated with the event study design at the aggregated firm-country-category-retailer level outlined in Equation 2. All models include fixed effects by firm-country-category-retailer, retailer-month, category-month, and country-month. 95% confidence intervals included, with standard errors clustered three-way at the category, country, and month level. Only wave 3 and never-treated categories included. Red vertical lines denote tariff changes, from 2.5% to 12.5% in September 2018, and from 12.5% to 27.5% in May 2019.

3.3 Pass-through and the Foreign Cost Share of U.S. Retail

While our unit value tariff elasticity of 0.293 is significant in its own right, as it suggests U.S. consumers bore roughly 30% of the incidence associated with these tariffs, our estimate is still well shy of the full pass-through into border prices recorded by numerous papers studying this same episode (Amiti et al. 2019; Fajgelbaum et al. 2020; Cavallo et al. 2021). This is surprising given that our estimate makes use of unit values, thereby offering a plausibly consistent comparison between the unit values recorded in customs data and those recorded in retail prices.[20]

As suggested by Cavallo et al. (2021), any differential in pass-through estimates at the border versus at the store could reflect the ability of U.S. intermediaries, such as retailers, to absorb this cost shock. Given estimates of full pass-through into border prices and null pass-through into store prices, it seems natural to assume that U.S. retailers bore the entire incidence of these import tariffs. However this result can only be inferred, as the distribution of incidence across intermediaries and domestic consumers is not directly observed. In this section, we instead consider the possibility that pass-through into retail prices is unlikely to fully reflect the same percentage change in prices as recorded in customs data, even in the presence of full pass-through into border and retail prices. Simply put, tariffs are only levied on the value share of a given imported good when that good crosses the border. All post-tariff costs associated with intra-national distribution and retailing occur after the tariff is levied, such that even full pass-through into border and retail prices would only manifest as the foreign cost share multiplied by the tariff change, when measured using retail prices.[21]

As a simple exercise to highlight this point, consider a retailer importing a good with some unit value u_k^f, as recorded at the border via customs data. This retailer then incurs additional costs associated with post-border services (e.g. transportation, warehousing, etc.) in order to sell this imported good to final consumers. Ignoring second-order substitution effects and assuming a constant pass-through elasticity across the underlying set of barcodes in each kr unit, the effect of a tariff change $d\ln\tau$ on the final price offered at retailer r is simply

$$\beta = \frac{d\ln u_{kr}}{d\ln\tau} = F_{kr}\frac{d\ln u_k^f}{d\ln\tau} = F_{kr}\beta^f.$$

That is, the effect of a tariff τ on the unit values associated with product k as observed

[20] Flaaen et al. (2020), Cavallo et al. (2021), and Flaaen et al. (2025) do not have quantity sales data for U.S. retailers, and are therefore unable to provide a direct comparison between retail unit values and customs data unit values. Instead, these papers estimate within-product estimates of the tariff elasticity of retail prices, rather than the elasticity of unit values aggregated over multiple products. We discuss this further in Section 4.

[21] See Sangani (2025) for a similar argument in the context of commodity price shocks and Flaaen et al. (2025) for evidence of this effect in wine tariffs.

at the border, u_k^f, is mediated by the cost share of purchasing this foreign good within the total cost of selling this product on the shelves of retailer r. In this case, F_{kr} denotes this foreign cost share and is defined as $F_{kr} \equiv u_k^f / u_{kr}$. An implication of this first-order cost shock approximation is that estimates of pass-through into retail prices which are lower than estimates of pass-through into border prices may in fact capture the exact same magnitude of tariff pass-through.[22]

Existing estimates of the foreign cost share of U.S. retail imports suggest a range of 30%-50% (Berger et al. 2012). For E.U. wine exports to the U.S., Flaaen et al. (2025) document a foreign cost share of approximately 38%, and provide evidence from the Bureau of Economic Analysis that across consumer goods, the average foreign cost share is just under 50%.[23] We make use of our correspondence between scanner and customs data to estimate the foreign cost share directly at the country-category level in the year 2017, prior to any tariff changes. We therefore impose the restriction $F_{kr} = F_{cj}$, and calculate unit values, or total import value by unit size, at the country-by-HS8 level using U.S. customs data. We then map these unit values into country-by-product module average import unit values, thereby calculating u_{cj}^f: the unit values associated with module c and country j as measured at the border. By calculating the same unit values with retail data, u_{cj}, we construct an estimate of the foreign cost share $\hat{F}_{cj} = u_{cj}^f / u_{cj}$.

We recover an average estimate of $\hat{F}_{cj} = 0.30$, with imports from China exhibiting an average of 0.32. Section B.1 provides further details on the construction of the foreign cost share, and Figure B.1 provides a histogram of foreign cost shares that we recover from the data. We manually link our product categories to those studied in Berger et al. (2012) and find that for categories consistent across the two datasets, Berger et al. (2012) recover foreign cost share estimates in the range of 30%-40%, suggesting that our lower estimates may be partly due to the composition of categories we study. It is also plausible that the share of domestic service content within imported retail goods increased between 2010 – when Berger et al. (2012) recorded their FCS estimates – and 2017, which is the reference year we use to

[22]If the underlying production function of retailing services is Cobb-Douglas in both the product being purchased and the retailing costs associated with providing this product to consumers, then this result is no long a first-order approximation, but holds exactly for a Cobb-Douglas cost parameter F_{kr} associated with the imported good as input.

[23]In a non-U.S. context, Errico (2026) estimates a foreign cost share of 43% within Chilean grocery goods.

uncover our FCS estimates.

As a validation of our FCS estimates, we consider heterogeneous unit value elasticities according to variation in the foreign cost share across categories. Specifically, we interact tariff changes in Equation 1 with alternative measures regarding the FCS of a given category, with Table A.2 providing these estimates. Column 1 re-states our headline estimate, with no interaction applied to tariffs. Column 2 provides an additional linear interaction between the FCS at the category level and tariff changes, with a recovered estimates of 0.529 ($p<0.01$) for the interaction term. If our estimates of the FCS contained no measurement error, and pass-through was full into retail prices, one would expect an estimate of one for this interaction. Our FCS estimates, likely contain considerable noise, given the aggregations required to calculate these estimates, and we therefore consider this exercise as a qualitative validation of our FCS estimates. Similarly, column 3 provides a binary high/low FCS comparison, with the high-FCS categories exhibiting a unit value elasticity significantly larger than, and distinguishable from, the unit value elasticity of low-FCS categories.

Combining our FCS estimate of 30% with our headline unit value tariff elasticity of 0.293, we infer a pass-through of this tariff shock to retail prices that is almost complete, since $\beta^f = \beta/F_{kr} = 0.293/0.30 \approx 1$. We therefore argue that our headline estimate of 0.293 is entirely compatible with the estimates of full pass-through into border prices recorded elsewhere, such that domestic intermediaries did not necessarily any share of this tariff cost shock. In fact, our estimates suggest that for consumer goods, U.S. consumers bore the entire incidence of this tariff shock. Crucially, however, this incidence is masked by the relatively low foreign cost share associated with U.S. retail goods.

Still, our estimates of pass-through into retail unit values are considerably larger than the within-product estimates associated with Cavallo et al. (2021), who document zero pass-through for this same tariff wave. Equipped with a unique combination of prices, quantities, and well-defined products, we dedicate the remainder of this paper to understanding this discrepancy. Ultimately, we provide evidence that the composition effect of tariffs on the set of products sold at U.S. retailers is responsible for the majority of tariff pass-through into unit values. We also provide evidence that these compositional effects are welfare relevant, in that they reflect product replacement bias, as in Nakamura and Steinsson (2012), via "shrink-flation": producers and retailers passed through tariffs to U.S. consumers primarily by reducing the average package size of barcodes offered. We therefore provide evidence that both the null estimates documented by Cavallo et al. (2021) and the complete pass-through estimates associated with unit values at the border (Amiti et al. 2019; Fajgelbaum et al. 2020; Cavallo et al. 2021) are not necessarily inconsistent, nor does this discrepancy in existing estimates necessarily imply that intermediaries absorbed any substantial share of the cost burden associated with these tariffs.

4 Margins of Adjustment

This section provides a decomposition of the unit value tariff elasticity estimates described above into various margins of adjustment. By incorporating price, quantity, and entry and exit data both within barcode-retailer pairs and within our aggregated product definition of a firm-category-country-retailer, we provide evidence that product replacement within the unbalanced panel of underlying barcodes is the key driver of the unit value estimates described earlier.

4.1 Decomposing Unit Value Pass-through

In order to organize the empirical estimates to follow, we begin by providing a decomposition of the tariff elasticity of unit values—total expenditure divided by total quantities sold—as observed at the firm-retailer-module-country-time level. The ultimate goal is to understand the conditions under which the within-barcode estimates of tariff effects, as documented in Cavallo et al. (2021), might differ from estimates associated with aggregated unit values, as measured in customs data.

We first decompose the aggregate unit value elasticity into two components: unit values associated with continuing barcodes and an adjustment factor which captures the role of entry and exit:

$$u_{krt} = u^c_{krt} \times \frac{u_{krt}}{u^c_{krt}}.$$

The superscript c denotes variables associated with a continuing set of barcodes, which we also refer to as the “balanced" panel of barcodes: barcode-retailer combinations observed in every month. In the absence of barcode turnover, the effect of tariffs on the unit value of continuing barcodes is closely related to the effect that is estimated using within-barcode variation for the same set of barcodes. Denoting by $\beta_{ir}^p \equiv \frac{d\ln p_{irt}}{d\ln\tau}$ and $\beta_{ir}^q \equiv \frac{d\ln q_{irt}}{d\ln\tau}$ the (potentially heterogeneous) effects of tariffs on barcode prices and quantities, we can write the effect on the unit value of continuing goods as follows.[24]

$$\frac{d\ln u_{krt}^c}{d\ln\tau} = \sum_i \Big(\frac{v_{irt}}{v_{krt}^c}\Big)\Big(\beta_{ir}^p + \beta_{ir}^q[1 - u_{krt}^c/p_{irt}]\Big).$$

If the effect of tariffs on barcode-level prices and quantities does not differ across barcodes within a firm-retailer-module-country cell, then the effect of tariffs on unit values of continuing barcodes equals the within-barcode effect: $\frac{d\ln u_{krt}^c}{d\ln\tau} = \beta^p$. In general, however, the two elasticities will likely differ due to heterogeneous responses across barcodes. For instance, the unit value elasticity would exceed the barcode-level price response if high market share barcodes have a higher price elasticity to tariffs, but would fall short of the barcode-level elasticity if initially more expensive barcodes yielded a larger quantity response. Thus, even in the continuing set of underlying barcodes, aggregated estimates via unit values might differ from within-barcode estimates due to correlation between heterogeneous effects, expenditure weights, and the initial distribution of prices across barcodes.

If this were the only point of departure between unit value and within-barcode estimates, it would be reasonable to consider the superiority of within-barcode estimates, in particular from the perspective of identification. Within-barcode estimates have the benefit of capturing a clear counterfactual price change within the exact same barcode by comparing an identical product before and after the tariff change. If the underlying correlation in heterogeneous tariff effects and market shares is weak, or non-existent, then for the continuing set of underlying barcodes, within-barcode estimates are arguably the correct approach to

[24]To derive this equation we note that $\frac{d\ln u_{krt}^c}{d\ln\tau} = \sum_i \frac{v_{irt}}{v_{krt}^c}\Big(\frac{d\ln p_{irt}}{d\ln\tau} - \frac{d\ln(q_{irt}/q_{krt}^c)}{d\ln\tau}\Big)$. Incorporating $\frac{d\ln q_{krt}^c}{d\ln\tau} = \sum_i \frac{q_{irt}}{q_{krt}^c}\frac{d\ln q_{irt}}{d\ln\tau}$, we derive:

$$\frac{d\ln u_{krt}^c}{d\ln\tau} = \sum_i \left(\frac{v_{irt}}{v_{krt}^c}\beta_{ir}^p + \left(\frac{v_{irt}}{v_{krt}^c} - \frac{q_{irt}}{q_{krt}^c}\right)\beta_{ir}^q\right) = \sum_i \frac{v_{irt}}{v_{krt}^c}\left(\beta_{ir}^p + \beta_{ir}^q[1 - u_{krt}^c/p_{irt}]\right).$$

measuring the price effects of tariffs.

However, within-barcode estimates are unable to account for the possibility of product replacement bias: the explicit use of barcode replacement by producers/retailers as a means of raising prices while mitigating the salience of this price change in the eyes of consumers (Nakamura and Steinsson 2012). To the extent that product replacement reflects a genuine change in the price of an otherwise comparable product, this change in price should be seen as relevant for welfare and therefore pass-through. To capture the role of product replacement, we consider the extent to which unit value tariff elasticities may differ between the continuing and unbalanced panel of barcodes, such that compositional changes in the set of barcodes generate the observed unit value tariff elasticities discussed above. In order to pass on tariffs into higher normalized prices, retailers can either shift the product mix to those products with higher price per unit weight or to products with identical sticker price but smaller packaging. To quantify these two different channels, we further decompose the product replacement channel into a channel that captures changes in average sticker prices, and a channel that captures changes in average package size.

We first write quantity as the product of the number of packages sold times package size: $q_{irt} = \tilde{q}_{irt} w_i$. The raw sticker price is then given by the sales volume divided by the number of packages sold, $\tilde{p}_{irt} = v_{irt}/\tilde{q}_{irt}$. Analogously, we therefore calculate the average sticker price as

$$\tilde{u}_{krt} = \frac{v_{krt}}{\tilde{q}_{krt}} = \sum_i \frac{\tilde{q}_{irt}}{\tilde{q}_{krt}} \tilde{p}_{irt}.$$

and average package size as $w_{krt} = \sum_i \frac{\tilde{q}_{irt}}{\tilde{q}_{krt}} w_i$. Unit values can then be expressed as

$$\begin{aligned} u_{krt} &= u^c_{krt} \times \frac{v_{krt}/\tilde{q}_{krt}}{v^c_{krt}/\tilde{q}^c_{krt}} \times \frac{\tilde{q}_{krt}/q_{krt}}{\tilde{q}^c_{krt}/q^c_{krt}} \\ &= u^c_{krt} \times \frac{\tilde{u}_{krt}}{\tilde{u}^c_{krt}} \times \frac{w^c_{krt}}{w_{krt}}. \end{aligned}$$

When differentiated with respect to some tariff change $d \ln \tau$, we derive

$$\frac{\partial \ln u_{krt}}{\partial \ln \tau} = \frac{\partial \ln u^c_{krt}}{\partial \ln \tau} + \left(\frac{\partial \ln \tilde{u}_{krt}}{\partial \ln \tau} - \frac{\partial \ln \tilde{u}^c_{krt}}{\partial \ln \tau} \right) + \left(\frac{\partial \ln w^c_{krt}}{\partial \ln \tau} - \frac{\partial \ln w_{krt}}{\partial \ln \tau} \right).$$

We turn now to providing tariff elasticity estimates associated with unit values and within-barcode prices for both the continuing panel of barcodes and the full set of underlying barcodes. In doing so, we show that roughly two thirds of our headline unit value estimate described in Section 3 is attributable to product replacement. We then highlight the role of "shrink-flation" in driving this differential effect, with a decrease in package size accounting for roughly half of the unit value tariff elasticity.

4.2 Tariff Elasticities by Aggregation and Barcode Tenure

This section provides the empirical estimates necessary to fully quantify the decomposition provided in the previous section. Specifically, we first make use of the same estimating equation from Section 3, Equation 1, and estimate the tariff elasticity of unit values and quantities sold at the firm-category-country-retailer level. While the estimates reported earlier in Table 1 construct unit values and quantities using the entire set of underlying barcodes, regardless of tenure, we now explicitly compare these estimates to those associated with only the continuing set of underlying barcodes. Finally, we provide the exact same comparison but for elasticity estimates associated with within-barcode changes over time, as opposed to changes in aggregate unit values and quantities.

Specifically, we re-create the triple-difference identification strategy associated with Equation 1 but at the barcode-retailer level of disaggregation, as illustrated in Equation 3:

$$\ln(y_{irt}) = \beta \ln(1 + \tau_{j(i)c(i)t}) + \alpha_{irt} + \gamma_{i,q(t)} + \epsilon_{irt}. \tag{3}$$

As in Equation 1, α_{irt} denotes a vector of fixed effects, with the sole exception in this case being the replacement of firm-category-country-retailer fixed effects α_{fcjr} with barcode-retailer fixed effects α_{ir}:

$$\alpha_{irt} = \alpha_{ir} + \alpha_{rt} + \alpha_{m(i),t} + \alpha_{j(i),t}.$$

Again, $\gamma_{i,q(t)}$ denote seasonality controls at the quarterly level, which are identical to those incorporated in Equation 1. As before, we can again estimate this model as an event study, in which barcode-retailer pairs within Wave 3 categories are defined as treated, whereas all

Table 2: Tariff Elasticity by Aggregation and Barcode Tenure

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Log Price		Log Unit Value		Log Quantity			
$\log(1+\tau)$	0.136***	0.001	0.114***	0.293***	-0.235***	-0.689***	-0.274**	-0.457***
	(0.034)	(0.045)	(0.024)	(0.069)	(0.081)	(0.184)	(0.108)	(0.148)
Within-Barcode	✓	✓			✓	✓		
Within-Product			✓	✓			✓	✓
Continuing	✓		✓		✓		✓	
All		✓		✓		✓		✓
R^2	1.00	1.00	1.00	1.00	1.00	1.00	0.98	0.98
Observations	15,393,284	43,819,632	4,183,892	4,183,892	15,393,284	43,819,632	4,183,892	4,183,892

*Notes: Columns 3-4 and 7-8 provide estimates of β from applying OLS to Equation 1, with data aggregated to the firm-category-country-retailer-month, whereas columns 1-2 and 5-6 provide estimates of β from applying OLS to Equation 3, with data disaggregated at the barcode-retailer-month level. All models include fixed effects by retailer-month, category-month, and country-month, as well as unit-specific fixed effects. Odd columns use only the continuing panel of barcodes, whereas even columns use the unbalanced panel of all barcodes. In all cases, standard errors are clustered three-way at the category, country, and month level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.*

never-affected categories act as control.

Table 2 provides these estimates, with columns 4 and 8 simply re-stating the unit value and quantity tariff elasticities associated with aggregating our data to the firm-category-country-retailer level and including all underlying barcodes, regardless of tenure, within this aggregation. As reported in Table 1, we find a unit value elasticity of 0.293 (column 4) and a quantity elasticity of -0.457 (column 8). Columns 3 and 7 provide tariff elasticities associated with the same outcomes, respectively, but aggregated using only the continuing panel of underlying barcodes. Interestingly, these estimates are considerably smaller in magnitude, with our unit value elasticity reduced to 0.114 — roughly a quarter of our headline estimate — and a quantity elasticity of -0.274.[25]

Yet this same gap in tariff effects across the continuing sample of barcodes and the unbalanced panel of barcodes does not persist when effects are estimated within barcodes. Specifically, columns 1 and 5 provide within-barcode estimates of price and quantity tariff effects via Equation 3, with the estimating sample subset to only the continuing panel of barcodes. Reassuringly, these estimates are similar to the aggregated unit value and quantity estimates, with a price elasticity of 0.136 and a quantity elasticity of -0.235, both significant

[25] When estimating tariff pass-through and quantity elasticities for the aggregate unbalanced panel, we restrict our observations to only those $fcjr$ quadruplets with non-zero balanced panel expenditure, such that the observations recorded below are identical for the "continuing" and "all" panel estimates. The terms "continuing" and "all" instead refer to the underlying set of barcodes used either as the sample itself, or as the set of barcodes with which we construct aggregated measures.

at $p < 0.01$. However the within-barcode price elasticity for the unbalanced panel of barcodes – reported in column 2 – is precisely zero, with a point estimate of 0.001.

The null effect of tariffs on within-barcode prices for the unbalanced panel of barcodes is striking for two reasons. First, it suggests that the gap in unit value tariff effects between the continuing panel and the unbalanced panel of barcodes is not simply driven by heterogeneous tariff elasticities across these two sets of barcodes. If anything, the larger price effect within-barcodes for the continuing panel is surprising, given that for unit values, the continuing panel exhibits much smaller tariff effects than the unbalanced panel. Second, these estimates form a bridge with those reported in Cavallo et al. (2021), who document effectively zero tariff pass-through into U.S. consumer gods when estimating these effects within narrow products.[26]

Finally, column 6 highlights an additional puzzle when comparing estimates associated with the continuing panel and the unbalanced panel of barcodes: despite a precisely null effect of tariffs on barcode prices for the unbalanced panel, the quantity elasticity for this same level of aggregation is -0.689, and significant at $p < 0.01$. The implied price elasticity for this unbalanced panel is therefore excessively large, suggesting that retailers do not allow consumers to substitute away from tariff-affected barcodes, but are instead removing tariff-affected products without altering their price. Incorporating all estimates of pass-through recorded up to this point, we recover a striking image of how the typical U.S. consumer interacted with these tariff-induced price changes. If a consumer had selected a random unbalanced panel of barcodes and recorded within-barcode price changes throughout this tariff episode, they would recover exactly zero tariff-induced price effects. If they instead focused their efforts on a continuing panel of barcodes, they would recover a tariff elasticity of retail prices of 14%. If they recorded the sales-weighted average price per barcode, such that they interpreted sticker prices as the object of interest, they would recover a tariff price elasticity of 11%. Finally, if they recorded the average price per unit weight of barcodes sold they would recover a tariff unit value elasticity of 29%. As discussed in Section 3.3, in order to then transform these elasticity estimates into estimates of pass-through, this consumer

[26]To the best of our knowledge, the null result in Cavallo et al. (2021) is similarly estimated using an unbalanced panel of products.

would need to then divide each price change by the foreign cost share of U.S. retail imports, which we estimate to as roughly 30%.

In the subsequent section we provide further evidence that product replacement was particularly associated with "shrink-flation". However, first we briefly consider some robustness exercises to ensure that the discrepancy in estimates reported in Table 2 is not driven by a flawed empirical design, but instead represents a true gap in tariff price effects across aggregations and tenures. First, Figure A.3 provides an event study design of unit value and quantity elasticities associated with both the continuing panel and the full panel of underlying barcodes, and confirms that the discrepancy in estimates between columns 3 and 4 of Table 2 is not attributable to differential pre-trends. Similarly, Figure A.4 finds a significant discrepancy between event study estimates for the unbalanced and balanced panel for categories associated with the wave 4a tariffs, with the same pattern emerging: continuing product unit values exhibit considerably smaller tariff elasticity estimates when compared to the aggregate unit values elasticities.

At the barcode-level, we confirm almost identical price and quantity elasticity estimates within the balanced panel of underlying barcodes when 2020 is excluded, thereby controlling for any Covid-19 shocks which may differentially affect the within-barcode estimates to a greater degree than the aggregated unit value estimates (Table A.3). Figure A.5 and Figure A.6 provide event study plots for prices and quantities and across both tariff waves for the within-barcode estimates, with again no evidence for differential pre-trends as the source of differing estimates across disaggregated and aggregated effects.

Finally, we consider the extent to which SUTVA violations might possibly bias our estimates of the within-barcode estimates downwards, potentially above and beyond the effect that such violations would have on our aggregated estimates. Table A.4 provides estimates of a modified model in which all Chinese goods are dropped and the China-specific tariff rates are applied to all products within a given category. Table A.4 therefore checks for whether or not non-tariff-targeted goods within tariff-affected categories exhibited a differential post-tariff price evolution than goods in non-tariff-affected categories. Notice that the identifying variation is now at the category-month level, such that we must drop fixed effects at this same level of variation. Columns (1) and (3) therefore provide no fixed effects to control for product classification shocks, whereas columns (2) and (4) provide department-time fixed effects, with departments defined as aggregations of similar product categories. In all cases, we find little to no evidence of spillover affects on the prices of non-treated products, either for aggregated unit values or within-barcodes.

Given that spillover effects would bias estimates downwards, and that our within-barcode estimates are significantly smaller than the unit value estimates, we provide an additional battery of robustness checks regarding potential SUTVA violations at the barcode level in Section B.2 in Appendix B. We consider the possibility of endogenous spillover effects to prices of non-targeted barcodes, tariff "smoothing" by retailers across products, and the effect of tariffs on intermediate goods leading to cost shocks for domestically produced goods. In all cases, we find no evidence of effects which would suggest that low barcode-level estimates described here are the result of downward bias via SUTVA violations.

4.3 Product Replacement and "Shrink-flation"

This section further decomposes tariff-induced price changes in unit values by barcode tenure into changes in the "sticker price", or unadjusted price per barcode purchased, and a change in the average size of barcodes within a given firm-category-country-retailer quadruplet. We replace the dependent variable in our triple-difference estimating equation—Equation 1—with the "sticker price" unit value $\tilde{u}_{krt}$ and the average package size w_{krt}. We then compare estimates of β across these two variables for the continuing set of barcodes and the unbalanced set of barcodes. As always, we also estimate this triple-difference identification strategy as an event study design, outlined in Equation 2. Figure 4 provides these event study figures, with the blue circle series denoting estimates associated with the continuing set of barcodes, and the red diamond series denoting estimates associated with the unbalanced set of all barcodes. The upper panel provides "sticker price" unit values, whereas the lower panel provides average weight estimates.

As shown in the upper panel of Figure 4, average sticker prices increased for the continuing set of barcodes, but only after the additional tariff change to 27.5% in May, 2019, and only by a modest amount (+2 ppt). The tariff-induced increase in average sticker prices associated with the unbalanced panel of barcodes, however, exhibits a clear positive response to both

Figure 4: Event Study Estimates for Unit Value Pass-through Decomposition

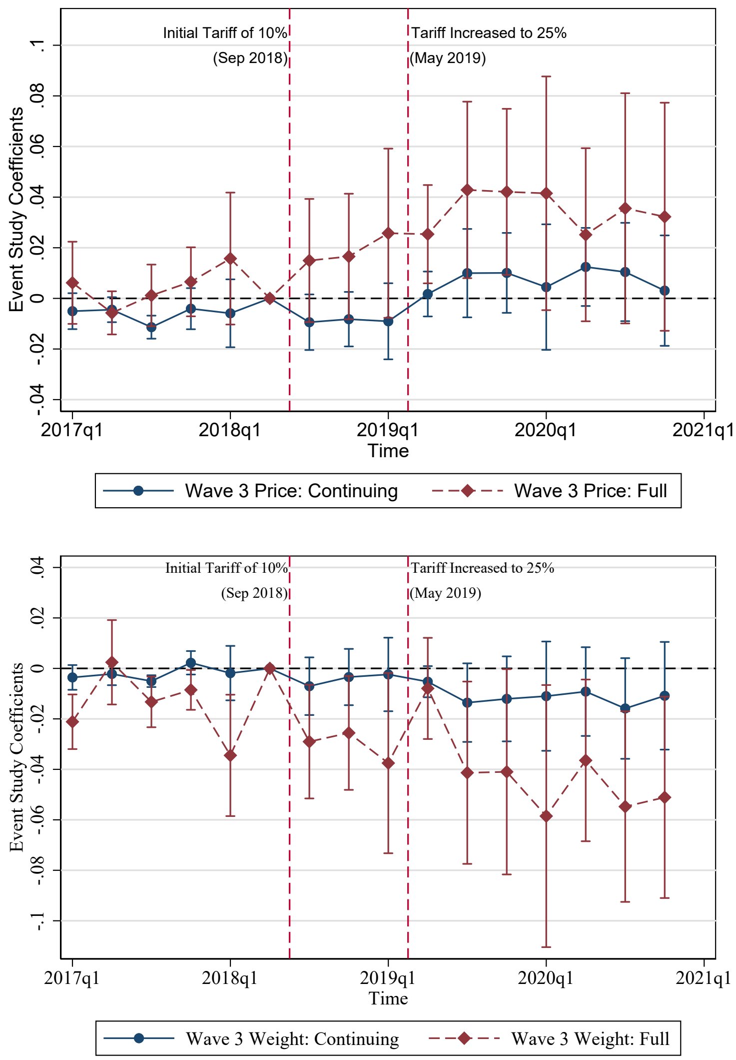

Notes: This figure provides estimates associated with an event study design at the aggregated firm-country-category-retailer level. All models include fixed effects by firm-country-category-retailer, retailer-month, category-month, and country-month. The dependent variable in the upper panel is $\ln \tilde{u}_{krt}$*, expenditure per barcode purchase counts, and in the lower panel the dependent variable is* $\ln w_{krt}$*, the average package size of barcodes purchased. All estimates are associated with wave 3 categories as treated and never-tariff categories as control. Red diamonds denote estimates using the entire sample of underlying barcodes; blue circles denote estimates associated with the continuing set of barcodes. Vertical lines denote tariff changes. 95% confidence intervals included, with standard errors are clustered three-way at the category, country, and month level..*

tariff changes, with relative sticker prices stabilizing around four percentage points higher than the control group after the second tariff change. Recall from Figure 3, however, that aggregated unit values in the unbalanced panel increased by as much as eight percentage points after the second tariff change. The bottom panel of Figure 4 highlights the role of changes in the denominator of unit values, or average weight, in explaining this discrepancy between unit values and sticker prices. As expected, the continuing set of barcodes shows no appreciable change in average package size over time, as this sample exhibits no entry or exit by construction.[27] However the unbalanced panel exhibits a striking negative trend which begins immediately following the first change in tariffs in September 2018, but intensifies after the second tariff increase in May 2019. By the first quarter of 2020, the average package size of barcodes sold in the unbalanced panel had decreased by roughly five percentage points when compared to Chinese barcodes in non-tariff categories. In all cases, and in particular for the estimates of average package size, pre-trends are non-existent, suggesting that these outcomes are attributable specifically to the import tariffs applied to these goods.

Table 3 provides the relevant point estimates associated with these event studies. Columns (1) and (2) simply re-iterate the unit value tariff elasticity estimates associated with, respectively, the balanced and unbalanced panels, as in Table 2. Columns (3) and (4) provide equivalent estimates but for the average sticker price, or price per barcode purchased, with the balanced panel exhibiting a sticker price tariff elasticity of 0.075 (0.037) and the unbalanced panel exhibiting a sticker price elasticity of 0.150 (0.060).

By construction, the sum of the sticker price elasticity and the average package elasticity equal the unit value elasticity, such that the estimates reported in Table 3 represent a true decomposition. Columns (5) and (6) provide the first clear evidence to date that just under half of the tariff pass-through into unit values was in fact achieved by an endogenous decrease in the average package size of goods sold in tariff-affected categories of imports. Specifically, the tariff elasticity of average package size for the continuing set of barcodes is indistinguishable from zero at -0.039 (0.031) while the tariff elasticity of average package size for the unbalanced set of barcodes is -0.143 (0.075), and distinguishable from zero at 90% confidence levels. To put these estimates in context, a 10% tariff increase generates a 1.4% decrease in average package size, while the full tariff increase for wave 3 products of 27.5% generated a decrease in package size of 3.9%, on average.

While the average package size elasticity estimate reported in column (6) is admittedly

[27] We confirm via the barcode provider GS1 that any change (increase or decrease) to the legally-required declared net content that is printed on the pack, requires assignment of a new GTIN, or barcode, such that there are no changes in package size within a continuing barcode over time.

Table 3: Decomposition of Unit Value Pass-through by Barcode Tenure

	(1)	(2)	(3)	(4)	(5)	(6)
	Log Unit Value		Log Sticker Price		Log Average Size	
$\log(1+\tau)$	0.114***	0.293***	0.075**	0.150**	-0.039	-0.143*
	(0.024)	(0.069)	(0.037)	(0.060)	(0.031)	(0.075)
Continuing	✓		✓		✓	
Full		✓		✓		✓
R^2	1.00	1.00	0.99	0.98	1.00	1.00
Observations	4,183,892	4,183,892	4,183,892	4,183,892	4,183,892	4,183,892

*Notes: This table provides estimates of β from applying OLS to Equation 1. Data are aggregated to the firm-category-country-retailer level. All models include fixed effects by firm-country-category-retailer, retailer-month, category-month, and country-month. Columns (1), (3), and (5) construct unit values and quantities sold using a balanced panel of underlying barcodes with non-zero expenditure in all periods between 2017 and 2020 for each $fcjr$ quadruplet. Columns (2), (4) and (6) construct measures using the raw unbalanced panel of all barcodes recording any sales in a given period. Only $fcjr$ quadruplets with a non-empty set of continuing barcodes are included in estimation, such that both samples have the same observations. Columns (1) and (2) estimate pass-through into unit values. Columns (3) and (4) estimate pass-through into "sticker prices", which is total expenditure divided by the un-adjusted number of units sold: $\tilde{u}_{krt} = v_{krt}/\tilde{q}_{krt}$. Columns (5) and (6) estimate tariff elasticities of average package size. Standard errors are clustered three-way at the category, country, and month level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.*

only marginally significant, we show below that this estimate masks considerable heterogeneity across categories, with those categories more amenable to the possibility of "shrinkflation" exhibiting large and statistically significant estimates associated with this mechanism. Finally, we show that these same mechanisms exist within the Wave 4 estimates. Figure A.7 provides estimates of the sticker price (left) and average package size (right) effects of tariffs around the introduction of each tariff change. Similar to the estimates for Wave 3, we find that a reduction in package size accounts for roughly half of the aggregate unit value tariff elasticity, with little evidence for similar outcomes in the continuing panel of barcodes.

4.4 Replacement Bias Versus Endogenous Quality

The previous results provide evidence that even within narrow definitions of product—specific firms selling narrow product modules from a given country and at a specific retail chain—product turnover led to a significant increase in the unit values associated with tariff-affected goods. Whether or not this tariff-induced change in unit values should be considered welfare-relevant, however, remains a first-order concern. A clear benefit to measuring tariff pass-through within continuing products—as in Flaaen et al. (2020), Cavallo et al. (2021), Cavallo et al. (2025), and Flaaen et al. (2025)—is that the researcher directly observes a pre-tariff price associated with each product. The fact that for entering barcodes we have no pre-tariff price is therefore worrisome, as our estimates may conflate endogenous changes in the set of products sold with genuine tariff pass-through via product replacement.[28] Specifically, if tariffs generate an endogenous exit of low-quality, or low-price, products, then the measured unit values may endogenously increase even in the absence of any true tariff pass-through.

From an identification perspective, this alternative hypothesis poses a problem, as we cannot directly observe the true "quality" of each barcode, and we cannot state with certainty that each entering barcode is equivalent to an alternative which existed prior to tariff implementation. Instead, we provide two complementary analyses which provide further evidence that product replacement is in fact the dominant mechanism through which tariff costs were passed through to U.S. consumers, and therefore the unit value tariff elasticities described earlier should be considered as welfare-relevant measures of pass-through. First, we consider the extent to which product categories may differ in their exogenous suitability to the use of "shrink-flation". Specifically, we consider the extent to which categories with large variation in package sizes prior to tariffs exhibit greater *ex post* use of "shrink-flation" when compared to categories with smaller size variation. The explicit argument being made is that categories with a high initial variance in sizes are likely to allow for greater use of "shrink-flation" when compared to categories which, in the limit, may have only a single size of barcode offered.

Specifically, we assign a product category as "high size variance" if the coefficient of variation across sizes offered is above the median size coefficient of variation across all categories, which we calculate at the category level using data across all months and retailers in 2017. This classification of categories as high and low size variance is therefore established prior to tariffs being implemented.[29]

[28]Of course firms may rely on product replacement as a means of increasing prices for the very same reason: it is difficult for consumers to compare the new, more expensive, product directly to the cheaper alternatives which were offered prior to the tariff change and subsequently dropped.

[29]We make use of the coefficient of variation, as opposed to variance or the standard of deviation, as neither of these alternative measures are invariant to the units of measure being used. Since categories may

Table 4: Unit Value Elasticity Decomposition by Shrink-flation Suitability

	Def.: Above Median Coeff. of Variation				
	(1)	(2)	(3)	(4)	(5)
	Unit Value	Sticker Price	Weight	Quantity	Units Sold
$\log(1+\tau)$	0.254***	0.215***	-0.039	-0.429*	-0.390*
	(0.079)	(0.076)	(0.068)	(0.221)	(0.203)
$\log(1+\tau)\times$ High Size Variance	0.060	-0.100*	-0.160***	-0.042	0.117
	(0.037)	(0.053)	(0.048)	(0.207)	(0.204)
All	✓	✓	✓	✓	✓
R^2	1.00	0.98	1.00	0.98	0.97
Observations	4,183,892	4,183,892	4,183,892	4,183,892	4,183,892

*Notes: This table provides estimates of β when applying OLS to equation 1, with tariffs interacted with a module-level measure of above- or below-median suitability to "shrink-flation": the coefficient of variation in sizes offered within a category. Standard errors are clustered three-way at the category, country, and month level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.*

Table 4 provides estimates associated with unit values, sticker prices, average package size, quantities sold, and units sold, calculated using the entire set of underlying barcodes, regardless of tenure. In each case, we provide a baseline tariff effect associated with categories which are below median, in terms of their size variation, and an added effect associated with categories being above median in their size variation. Thus, the total effect is simply the sum of both rows, with significance capturing the extent to which above- and below-median category elasticities differ. We find strong evidence that high size variance modules made greater use of "shrink-flation" as a means of passing through tariff costs to consumers, with this mechanism substituting for higher sticker prices – the dominant mechanism for low size variance modules. Specifically, low size variance modules exhibit a unit value tariff elasticity of 0.254, compared to 0.314 in high size variance modules, although the differential effect for high size variance modules is not distinguishable from zero at standard confidence levels. Crucially, changes in sticker prices account for roughly 85% of the unit value elasticity in low size variance modules, with a sticker price elasticity of 0.215, whereas high size variance categories exhibit a sticker price elasticity of only 0.115, with this differential elasticity significant at $p < 0.10$.

The average package size elasticity, on the other hand, is only -0.039, and indistinguishable from zero, for low size variance categories, but -0.199 for high size variance categories,

differ in their unit of measure, the coefficient of variation offers a unit-less measure which is comparable across categories.

with the differential effect between the two significant at $p < 0.01$. Thus, we provide evidence that the use of "shrink-flation" as a means of passing through a cost shock is highly concentrated in categories with greater variance in sizes offered pre-tariff shock. For low size variance modules, reductions in package size accounted for less than a fifth of the unit value elasticity, whereas in high size variance modules, "shrink-flation" accounted for just under two thirds of this same elasticity. Table 4 therefore provides evidence that a plausibly exogenous characteristic of product modules – the distribution of sizes offered pre-tariffs – generated differential use of "shrink-flation" as a mechanism to pass-through tariff costs, which stands in contrast to an endogenous quality composition mechanism, which would plausibly be independent of the initial size variation within a given module.

One potential caveat to the estimates provided in Table 4 is that our method of classifying modules into high and low size variance types relies on the aggregate size distribution across all barcodes within each module, whereas our estimates of "shrink-flation" occur within firm-module pairs. Table A.5 therefore provides the same estimates as Table 4 but with an alternative allocation method of placing modules in high and low size variance classes: the average within-firm coefficient of variation across sizes offered within a given firm-module pair. We again split modules by above- and below-median, and estimates are broadly similar to those found Table 4: reductions in package size account for roughly half of the unit value elasticity in high size variance modules, but less than a third in low size variance modules (with the low size variance elasticity not significant at standard confidence levels).

Columns (4) and (5) of Table 4 provide qualitative evidence regarding the potentially differential effects of sticker price and package size changes on sales. Column (4) provides estimates of the quantity sold tariff elasticity and finds only a marginal difference between the two module types, consistent with the fact that the unit value elasticity is also similar between the two. Column (5) instead provides a comparison of the units sold tariff elasticity, which plausibly captures consumer behaviour to a greater extent than total quantity sold, as this represents the raw number of instances in which a consumer purchased any barcode. While the estimate of the differential effect is imprecisely estimated, there is some evidence that the use of "shrink-flation" mitigated the loss in units sold. When taken at face value, the estimates in Table 4 suggest that high size variance modules exhibited a units sold to unit values elasticity of only -0.87, whereas for low size variance modules this elasticity was -1.54. This is true even though the implied true price elasticities for both high and lower size variance modules were, respectively, -1.50 and -1.69. While the comparison is not exact, these estimates align to some extent with existing estimates, which suggest that consumers are between 25%-50% less responsive to changes in package size when compared to sticker price demand sensitivity (Çakır and Balagtas 2014; Janssen and Kasinger 2026). Ultimately, such comparisons are at best qualitative, however, given the possibility that high and low size variance modules differ in other dimensions (including their price elasticity of demand), and the relative imprecise nature of the estimates reported here.

As a final result in favor of interpreting our unit value elasticities as indicative of welfare-relevant tariff pass-through, we show in Appendix B.3 that while tariffs generated a net loss of barcodes – consistent with a negative extensive margin elasticity of tariffs – the barcode entry elasticity of tariffs is in fact positive. While we relegate details to Appendix B.3, including the empirical model and results, this result is important as it suggests that the response of firms and retailers to tariffs is somewhat more nuanced than the quality composition hypothesis would suggest. Specifically, it is unclear why high quality barcodes would enter a tariff-affected market, and in doing so endogenously increase the unit value via composition, as those same high-quality barcodes would almost certainly exhibit positive profits prior to tariffs being implemented as well.

We therefore provide evidence that product turnover and replacement generated the headline unit value elasticity documented in Section 3, with a reduction in package size accounting for roughly half of this aggregate effect. This average package size effect was stronger in modules with an initially high variance in package sizes offered, and tariff-affected units exhibited both greater exit and entry than non-tariff-affected units, consistent with product replacement as the dominant mechanism of tariff pass-through.

5 Heterogeneity in Tariff Pass-through

As a final exercise, we study heterogeneity in pass-through across product and retailer characteristics. The key motivation for this exercise is the substantial literature both theoretically grounding, and empirically documenting, heterogeneous pass-through according to market structure and product-specific market power (Atkeson and Burstein 2008; Amiti et al. 2016; Alviarez et al. 2023). In general, these papers either confirm or predict that pass-through should be decreasing in market power, as these firms can reduce mark-ups in order to gain a greater share of the market, whereas small firms are constrained to charging a fixed mark-up over marginal costs. Alviarez et al. (2023), however, suggest that nuance is also needed when mapping pass-through across bilateral buyer-seller relationships, such that both buyer and seller power play a role in shaping the pass-through which occurs when the seller faces a cost shock. By combining scanner data with a well-defined cost shock, the tariffs implemented by the U.S. in 2018-2019, our setting provides a unique opportunity to estimate the relative effects of buyer and seller power on tariff elasticities.

In constructing market shares, we consider three separate definitions of buyer/seller power. In all cases, we use expenditure shares in 2017 to define each measure. First, we define barcode market shares as the within-category expenditure share of a given barcode. Second, we define the seller share, or firm share, of a given firm within a given retail chain as the share of revenue accruing to retailer r generated by sales associated with firm f. Finally, we consider the buyer share of a given retailer within a given firm, which we define as the share of all revenue accruing to a given firm which is generated within a given retail chain.[30]

Table A.6 provides estimates of tariff pass-through and substitution by barcode market share. We estimate pass-through elasticities which are increasing in market share, with each MS_i denoting elasticity estimates associated with barcodes in the i tercile of 2017 market share. From low to high market share, we recover pass-through estimates of 10.4%, 20.8%, and 23.2%, respectively. Interestingly, tariff quantity elasticities exhibit the reverse pattern, with tariff substitution largest for barcodes in the lowest market share tercile, despite these barcodes exhibiting the lowest level of tariff pass-through.

While tariff pass-through into prices depends on the super-elasticity of the demand system, the substitution effect on quantities depends largely on the initial level of the demand elasticity. The results therefore suggest that a larger substitution effect for low market share

[30] 87% of all expenditure on Chinese barcodes is concentrated in the lowest quintile of market share, which hosts 30% of all Chinese barcodes. At the high end of the distribution, 1.6% of all Chinese barcodes are in the highest market share quintile, but these barcodes account for 14% of all expenditure on Chinese barcodes.

goods is the result of a very high initial demand elasticity of low market share barcodes – so high that it overcompensates the smaller increase in prices. Taking our results on prices and quantities together, our results are therefore consistent with Atkeson and Burstein (2008) in the ranking of demand elasticities across market shares (low market share goods have higher elasticities), but inconsistent in the ranking of super-elasticities, as our price results suggest that the demand elasticity of high market share products is less elastic, leading to higher pass-through.

We then turn to understanding how pass-through varies by firm and retailer bargaining power via the buyer and seller shares described earlier. For example, Alviarez et al. (2023) build a model of two-sided bargaining according to which pass-through should be increasing in the firm share of retailer sales and decreasing in the retailer's share of firm sales. Figure A.8 plots the distribution of these two bilateral variables in the data in 2017. The blue bars highlight that the vast majority of firms are small relative to retailers: in roughly 80% of all firm-retailer relationships, the firm accounts for less than 0.1% of total retailer sales, and in almost all firm-retailer relationships the firm accounts for less than 1% of retailer sales. In contrast, many retailers account for a sizable share of firm sales: In roughly half of all firm-retailer relationships, the retailer accounts for more than 1% of firm sales, and in almost 15% of cases for more than 10%.

We test for the presence of buyer and seller power by interacting our tariff change in Equation 1 with two indicators: an indicator for whether a given firm exhibits an above-median seller share at a given retailer, and an indicator for whether a given retailer exhibits an above-median buyer share for a given firm. These two indicators are not mutually exclusive, and the interaction between them generates four possible combinations of buyer and seller power interactions. We then estimate the tariff elasticity of unit values, sticker prices, and average package size for the complete set of barcodes, while interacting tariff changes with these measures of buyer and seller power.

Table 5 provides these estimates. In all cases, the indicator for buyer/seller power is additive with respect to the baseline estimate of the tariff elasticity, provided in the first row. Column (2) provides our headline unit value estimates, and highlights the remarkable extent to which large and small firms/retailers differ in the extent to which unit values

Table 5: Unit Value Elasticity Decomposition by Buyer and Seller Shares

	Tariff Elasticity by Buyer/Seller Share		
	(1) Log Unit Value	(2) Log Sticker Price	(3) Log Average Weight
$\log(1+\tau)$	0.170***	0.018	-0.151**
	(0.061)	(0.055)	(0.074)
$\log(1+\tau)\times$ High retailer share in firm sales	-0.212***	-0.140***	0.073***
	(0.040)	(0.025)	(0.024)
$\log(1+\tau)\times$ High firm share in retailer sales	0.256***	0.220***	-0.035
	(0.033)	(0.047)	(0.034)
All	✓	✓	✓
R^2	1.00	0.98	1.00
Observations	4,183,892	4,183,892	4,183,892

*Notes: This table provides estimates of β when applying OLS to equation 1, with tariffs interacted with both an indicator for a high buyer share and an indicator for a high seller share within a given firm-retailer relationship. Standard errors are clustered three-way at the category, country, and month level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.*

respond to tariff changes. In particular, we find a baseline unit value elasticity of 0.170 for a small producer selling at a small retailer, such that neither monopsony or monopoly power are likely. For a large producer selling at a small retailer, however, such that some monopolistic power may exist, the unit value elasticity increases to 0.426, and yet the unit value elasticity in the opposite scenario – a large retailer stocking barcodes from a relatively small firm – is effectively zero (-0.042). Thus, monopoly power tends to increase the unit value tariff elasticity, whereas monopsony power decreases this elasticity, consistent with the theory and empirical evidence documented in Alviarez et al. (2023). When both sources of power are relevant, the unit value elasticity is 0.214.

The use of package size reductions as a means of passing through tariff costs, however, does not vary at all according to whether a producer exhibits monopolistic power: the 26 percentage point increase in the unit value elasticity associated with monopoly power occurs alongside a 22 percentage point increase in sticker prices, and only a marginal (and statistically insignificant) decrease in package size. Large retailers, however, seem to make use of their monopsony power by reducing the use of “shrink-flation", such that for a small producer, the average package size tariff elasticity is effectively reduced in magnitude by 50% when selling at a larger retail versus at a smaller retailer. Interestingly, large retailers also exhibit negative sticker price tariff elasticities when selling barcodes produced by small firms, such that there is some evidence of these larger retailers off-loading inventory of these barcodes. Thus, all producer seem willing, and able, to leverage "shrink-flation" as a means of passing through tariff costs to consumers, but large retailers seem to have an incentive to reduce the use of this mechanism, and indeed to pass-through as little of the tariff cost shock to consumer prices as possible.

A consequence of these estimates is that consumers purchasing tariff-affected barcodes associated with large firms, but sold at small retail outlets, experienced the largest implied tariff pass-through associated with these tariffs, whereas consumers purchasing barcodes associated with small firms, but purchased at large retail outlets, were generally unaffected. This remarkable variance in tariff-related price outcomes has yet to be fully explored in the literature, and suggests that there may be considerable variation in the underlying exposure to tariffs according to shopping patterns across firms and retailers of differential size, such as those documented in Faber and Fally (2021).

6 Conclusion

This paper provides novel evidence regarding the pass-through and substitution effects of U.S. import tariffs applied to Chinese consumer goods exports during the 2018-2019 U.S.-China trade war. Existing evidence suggests a puzzle, in that tariffs were fully passed through into border prices with little effect on retail prices. We show that in fact such a discrepancy in estimates does not necessarily imply zero pass-through to U.S. consumers. Instead, we show that product replacement bias—the replacement of products with higher price per unit weight alternatives—was the predominant mechanism whereby import tariffs were passed through to U.S. consumers, such that tariff price elasticity estimates within a given product might be zero while estimates for retail unit values are as high as 0.29. We provide evidence that post-tariff margins associated with mark-ups and/or intra-national distribution and retailing costs imply that a retail price elasticity of 0.29 is not inconsistent with full pass-through into border prices, and implies a pass-through into retail prices that was likely complete. Thus, we conclude that U.S. consumers bore the majority, if not the entirety, of tariff incidence associated with these import tariffs.

Finally, we shed light on the heterogeneous effects of tariffs across barcodes, firms, and retailers, according to buyer and seller power within these producer-retailer relationships. We document increasing pass-through for higher market share barcodes and higher pass-through for firms which constitute a relatively large share of expenditure within a given retailer. Retailers which constitute a relatively large share of expenditure within a given firm, however, dampen the level of pass-through, in some cases to effectively zero. While both retailers and firms made use of "shrink-flation" as a means of passing through tariff costs to consumers via product replacement, the use of this technology was particularly concentrated among producer-retailer pairs in which monopoly power was large, wheras in settings with high monopsony power – large retail chains – the use of this mechanism was dampened significantly.

While our estimates offer a solution to existing empirical puzzles within the tariff pass-through literature, the effects estimated here are ultimately limited to the study of consumer goods. A significant share of tariffs levied in the 2018-2019 trade war were applied to intermediate and capital goods, for which scanner data is not readily available. We therefore caution that our estimates apply only to those tariffs associated with consumer goods, and leave the concordance of our estimates with those for other product categories to future research.

References

Alviarez, Vanessa I, Michele Fioretti, Ken Kikkawa, and Monica Morlacco, “Two-sided market power in firm-to-firm trade,” Technical Report, National Bureau of Economic Research 2023.

Amiti, Mary, Chris Flanagan, Sebastian Heise, and David E Weinstein, “Who Is Paying for the 2025 US Tariffs?,” Technical Report, Federal Reserve Bank of New York 2026.

_ **, Oleg Itskhoki, and Jozef Konings**, “International Shocks, Variable Markups, and Domestic Prices,” *The Review of Economic Studies*, 2016, *86* (6), 2356–2402.

_ **, Stephen J Redding, and David E Weinstein**, “The Impact of the 2018 Tariffs on Prices and Welfare,” *Journal of Economic Perspectives*, 2019, *33* (4), 187–210.

Anderson, James E and Eric Van Wincoop, “Trade costs,” *Journal of Economic literature*, 2004, *42* (3), 691–751.

Atkeson, Andrew and Ariel Burstein, “Pricing-to-market, trade costs, and international relative prices,” *American Economic Review*, 2008, *98* (5), 1998–2031.

Auer, Raphael, Ariel Burstein, and Sarah M Lein, “Exchange rates and prices: evidence from the 2015 Swiss franc appreciation,” *American Economic Review*, 2021, *111* (2), 652–686.

Bai, Liang and Sebastian Stumpner, “Estimating US consumer gains from Chinese imports,” *American Economic Review: Insights*, 2019, *1* (2), 209–24.

Basker, Emek and Pham Hoang Van, “Imports "R" Us: Retail Chains as Platforms for Developing-Country Imports,” *American Economic Review*, 2010, *100* (2), 414–18.

_ **, Shawn Klimek, and Pham Hoang Van**, “Supersize it: The growth of retail chains and the rise of the "Big-Box" store,” *Journal of Economics & Management Strategy*, 2012, *21* (3), 541–582.

Benguria, Felipe, Jaerim Choi, Deborah L Swenson, and Mingzhi Jimmy Xu, "Anxiety or pain? The impact of tariffs and uncertainty on Chinese firms in the trade war," *Journal of International Economics*, 2022, *137*, 103608.

Berger, David, Jon Faust, John H Rogers, and Kai Steverson, "Border prices and retail prices," *Journal of international Economics*, 2012, *88* (1), 62–73.

Blanchard, Emily J, Chad P Bown, and Davin Chor, "Did Trump's Trade War Impact the 2018 Election?," *Journal of International Economics*, 2024, *148*, 103891.

Bown, Chad P, "The US–China trade war and Phase One agreement," *Journal of Policy Modeling*, 2021, *43* (4), 805–843.

Broda, Christian and David E Weinstein, "Globalization and the Gains from Variety," *The Quarterly journal of economics*, 2006, *121* (2), 541–585.

Burstein, Ariel T, Joao C Neves, and Sergio Rebelo, "Distribution costs and real exchange rate dynamics during exchange-rate-based stabilizations," *Journal of monetary Economics*, 2003, *50* (6), 1189–1214.

Çakır, Metin and Joseph V Balagtas, "Consumer response to package downsizing: Evidence from the Chicago ice cream market," *Journal of Retailing*, 2014, *90* (1), 1–12.

Cavallo, Alberto, Gita Gopinath, Brent Neiman, and Jenny Tang, "Tariff pass-through at the border and at the store: Evidence from us trade policy," *American Economic Review: Insights*, 2021, *3* (1), 19–34.

_ , Paola Llamas, and Franco M Vazquez, "Tracking the short-run price impact of US tariffs," Technical Report, National Bureau of Economic Research 2025.

Chalioti, Evangelia and Konstantinos Serfes, "Shrinkflation," *Economics Letters*, 2024, *244*, 111959.

Charbonneau, Karyne B and Anthony Landry, "The trade war in numbers," Technical Report, Bank of Canada 2018.

Chen, Natalie and Luciana Juvenal, "Quality, trade, and exchange rate pass-through," *Journal of International Economics*, 2016, *100*, 61–80.

Chor, Davin, Matthew Grant, and Bingjing Li, "Exclusions for Sale? Tariff Exclusions in the US-China Trade War," *mimeo*, 2025.

Cole, Matthew T and Carsten Eckel, "Tariffs and markups in retailing," *Journal of International Economics*, 2018, *113*, 139–153.

Devereux, Michael B, Charles Engel, and Cédric Tille, "Exchange rate Pass-Through and the welfare effects of the euro," *International Economic Review*, 2003, *44* (1), 223–242.

Eckel, Carsten, "International trade and retailing," 2009.

Errico, Marco, "Domestic wedges and the (in) sensitivity of CPI to exchange rates," *Journal of International Economics*, 2026, p. 104237.

Faber, Benjamin and Thibault Fally, "Firm heterogeneity in consumption baskets: Evidence from home and store scanner data," *Review of Economic Studies*, 2021.

Fajgelbaum, Pablo D and Amit K Khandelwal, "The economic impacts of the US–China trade war," *Annual Review of Economics*, 2022, *14* (1), 205–228.

_ , Pinelopi K Goldberg, Patrick J Kennedy, and Amit K Khandelwal, "The return to protectionism," *The Quarterly Journal of Economics*, 2020, *135* (1), 1–55.

Feenstra, Robert C, "New product varieties and the measurement of international prices," *The American Economic Review*, 1994, pp. 157–177.

Fetzer, Thiemo and Carlo Schwarz, "Tariffs and politics: Evidence from Trump's trade wars," *The Economic Journal*, 2021, *131* (636), 1717–1741.

Flaaen, Aaron, Ali Hortaçsu, and Felix Tintelnot, "The production relocation and price effects of US trade policy: the case of washing machines," *American Economic Review*, 2020, *110* (7), 2103–27.

Flaaen, Aaron B, Ali Hortaçsu, Felix Tintelnot, Nicolás Urdaneta, and Daniel Xu, "Who Pays for Tariffs Along the Supply Chain? Evidence from European Wine Tariffs," Technical Report, National Bureau of Economic Research 2025.

Freitag, Andreas and Sarah M Lein, "Endogenous product adjustment and exchange rate pass-through," *Journal of International Economics*, 2023, *140*, 103706.

Ganapati, Sharat and Colin Hottman, "Did Foreigners Pay America's Tariffs? Quantity Discounts, Scale Economies and Incomplete Pass-Through," Technical Report, CESifo 2026.

Goldberg, Linda S and José Manuel Campa, "The sensitivity of the CPI to exchange rates: Distribution margins, imported inputs, and trade exposure," *The Review of Economics and Statistics*, 2010, *92* (2), 392–407.

Goldberg, Pinelopi Koujianou and Rebecca Hellerstein, "A structural approach to identifying the sources of local currency price stability," *Review of Economic Studies*, 2013, *80* (1), 175–210.

Gopinath, Gita and Brent Neiman, "The Incidence of Tariffs: Rates and Reality," Technical Report, National Bureau of Economic Research 2026.

Handley, Kyle, Fariha Kamal, and Ryan Monarch, "Rising import tariffs, falling exports: When modern supply chains meet old-style protectionism," *American Economic Journal: Applied Economics*, 2025, *17* (1), 208–238.

Hankins, Kristine W, Morteza Momeni, and David Sovich, "Consumer credit and the incidence of tariffs: Evidence from the auto industry," *American Economic Review (Forthcoming)*, 2025.

Hinz, Julian, Aaron Lohmann, Hendrik Mahlkow, and Anna Vorwig, "America's Own Goal: Who Pays the Tariffs?," Technical Report, Kiel Policy Brief 2026.

Jaccard, Torsten, "Retailers and Import Penetration: Evidence from US Dollar Stores," 2023.

Janssen, Aljoscha and Johannes Kasinger, "Shrinkflation and consumer demand," *Marketing Science*, 2026, *45* (1), 142–158.

Kim, In Kyung, "Consumers' preference for downsizing over package price increases," *Journal of Economics & Management Strategy*, 2024, *33* (1), 25–52.

Ludema, Rodney D and Zhi Yu, "Tariff pass-through, firm heterogeneity and product quality," *Journal of International Economics*, 2016, *103*, 234–249.

Metzler, Lloyd A, "Tariffs, the terms of trade, and the distribution of national income," *Journal of Political Economy*, 1949, *57* (1), 1–29.

Nakamura, Emi and Dawit Zerom, "Accounting for incomplete pass-through," *The review of economic studies*, 2010, *77* (3), 1192–1230.

_ **and Jón Steinsson**, "Price setting in forward-looking customer markets," *Journal of Monetary Economics*, 2011, *58* (3), 220–233.

_ **and** _ , "Lost in transit: product replacement bias and pricing to market," *American Economic Review*, 2012, *102* (7), 3277–3316.

Raff, Horst and Nicolas Schmitt, "Exclusive dealing and common agency in international markets," *Journal of International Economics*, 2006, *68* (2), 485–503.

_ **and** _ , "Buyer power in international markets," *Journal of International Economics*, 2009, *79* (2), 222–229.

_ **and** _ , "Imports and the structure of retail markets," *Canadian Journal of Economics/Revue canadienne d'économique*, 2012, *45* (4), 1431–1455.

_ **and** _ , "Retailing and international trade," *Handbook on the Economics of Retailing and Distribution. Edwar Elgar Publishers, Cheltenham*, 2016, pp. 157–179.

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

Sangani, Kunal, "Complete Pass-Through in Levels," 2025.

A Appendix: Additional Tables and Figures

This section provides additional tables and figures, as described in the main text.

A.1 Additional Tables

Table A.1: Tariff Pass-through and Substitution by Wave

	(1) Log Unit Value	(2) Log Quantity	(3) Log Unit Value	(4) Log Quantity
$\log(1+\tau)$	0.293***	-0.457***	0.380***	-0.512*
	(0.069)	(0.148)	(0.082)	(0.295)
Wave 3	✓	✓		
Wave 4a			✓	✓
R^2	1.00	0.98	1.00	0.98
Observations	4,183,892	4,183,892	3,858,741	3,858,741

*Notes: This table provides estimates of β from applying OLS to Equation 1. Data are aggregated to the firm-category-country-retailer level. All models include fixed effects by firm-country-category-retailer, retailer-month, category-month, and country-month. All columns provide results estimated using the unbalanced panel of underlying barcodes. Columns (1) and (2) estimate the effect of Wave 3 tariffs on unit values and quantities, with all Wave 4a categories omitted. Columns (3) and (4) estimate the effect of Wave 4a tariffs on unit values and quantities with all Wave 3 categories omitted. Standard errors are clustered three-way at the category, country, and month level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.* *Return to Section 3.2.*

Table A.2: Unit Value Tariff Elasticity Interacted with FCS

	(1)	(2)	(3)
		Log Unit Value	
$\log(1+\tau)$	0.293***	0.060	0.102*
	(0.069)	(0.068)	(0.061)
$\log(1+\tau)\times$ FCS		0.529***	
		(0.143)	
$\log(1+\tau)\times$ High FCS			0.213***
			(0.055)
All	✓	✓	✓
R^2	1.00	1.00	1.00
Observations	4,183,892	4,114,606	4,114,606

*Notes: This table provides estimates of β from applying OLS to Equation 1. Data are aggregated to the firm-category-country-retailer level. All models include fixed effects by firm-country-category-retailer, retailer-month, category-month, and country-month. All columns provide results estimated using the full panel of underlying barcodes. Column 2 interacts tariffs with the category-level foreign cost share estimate, and column 3 interacts tariffs with an indicator for a given category exhibiting an above-median foreign cost share. Standard errors are clustered three-way at the category, country, and month level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.* *Return to Section 3.3.*

Table A.3: Estimates of Tariff Pass-through and Substitution at the Barcode Level

	(1) Log Price	(2) Log Quantity	(3) Log Price	(4) Log Quantity
$\log(1+\tau)$	0.136***	-0.235***	0.112***	-0.361***
	(0.034)	(0.081)	(0.040)	(0.043)
Sample			<2020	<2020
Fixed Effects	✓	✓	✓	✓
R^2	1.00	0.97	1.00	0.98
Observations	1.54e+07	1.54e+07	1.16e+07	1.16e+07

Notes: This table provides estimates of β from applying OLS to Equation 3. Columns (1) and (2) use the entire sample of observations, with column (1) documenting price pass-through and column (2) the tariff substitution elasticity within barcodes. Columns (3) and (4) provide the same estimates, respectively, but with 2020 excluded from the sample. All estimates are associatd with the continuing panel of balanced barcodes. Standard errors are clustered three-way at the category, country, and month level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively. Back to Section 4.2.

Table A.4: SUTVA Violation Check by Aggregation

	(1)	(2)	(3)	(4)
	Log Unit Value		Log Price	
$\log(1+\tau)$	0.041*	-0.009	0.030	0.001
	(0.021)	(0.026)	(0.019)	(0.025)
Within-Barcode			✓	✓
Within-Firm-Category-Country	✓	✓		
Department-Month FE		✓		✓
Observations	4,063,120	4,063,120	15,093,623	15,093,620

Notes: all models exclude any products from China. Tariff variation occurs at the category-month level, with $\tau_{ct} = \tau_{ct,CN}$*. Columns (1) and (3) therefore drop the fixed effects at the category-month level, and columns (2) and (4) provide additional fixed effects at the department-month level. Columns (1) and (2) provide estimates for aggregated unit value spillover effects of tariffs on non-Chinese products, and columns (3) and (4) provide the same estimates but for within-barcode prices. Standard errors are clustered three-way at the category, country, and month level.* *, **, *and* *** *denote statistical significance at the 10%, 5%, and 1% levels, respectively.* *Back to Section 4.2.*

Table A.5: Unit Value Elasticity Decomposition by Shrink-flation Suitability (Alternative Classification)

	Def.: Above Median Mean Within-Firm Coeff. of Variation				
	(1)	(2)	(3)	(4)	(5)
	Unit Value	Sticker Price	Weight	Quantity	Units Sold
$\log(1+\tau)$	0.250***	0.167*	-0.083	-0.439*	-0.356*
	(0.089)	(0.087)	(0.068)	(0.233)	(0.208)
$\log(1+\tau)\times$ High Size Variance	0.058	-0.024	-0.082*	-0.025	0.057
	(0.053)	(0.070)	(0.041)	(0.221)	(0.234)
All	✓	✓	✓	✓	✓
R^2	1.00	0.98	1.00	0.98	0.97
Observations	4,183,892	4,183,892	4,183,892	4,183,892	4,183,892

*Notes: This table provides estimates of β when applying OLS to equation 1, with tariffs interacted with a module-level measure of above- or below-median suitability to "shrink-flation": the mean coefficient of variation in sizes offered across firms within a category. Standard errors are clustered three-way at the category, country, and month level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively. Back to Section 4.4.*

Table A.6: Pass-through and Substitution by Market Share Bins

	(1)	(2)	(3)	(4)
	Log Price		Log Quantity	
$\log(1+\tau) \times MS_1$	0.104**	0.081*	-0.282	-0.397**
	(0.039)	(0.044)	(0.219)	(0.191)
$\log(1+\tau) \times MS_2$	0.208***	0.174***	-0.090	-0.280
	(0.050)	(0.054)	(0.214)	(0.196)
$\log(1+\tau) \times MS_3$	0.232***	0.217***	-0.197	-0.253
	(0.065)	(0.058)	(0.310)	(0.256)
Sample		<2020		<2020
Fixed Effects	✓	✓	✓	✓
R^2	1.00	1.00	0.97	0.98
Observations	15393284	11556925	15393284	11556925

*Notes: tariff pass-through by barcode market share, from low market share (s_1) to high market share (s_3). Market shares are computed within product modules based on 2017 sales volume. Columns (1) and (3) estimate price and quantity elasticities for the entire sample, with columns (2) and (4) providing the same respective estimates with 2020 excluded. Standard errors are clustered three-way at the category, country, and month level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.* *Back to Section 5.*

A.2 Additional Figures

Figure A.1: Distribution of Barcodes by Department

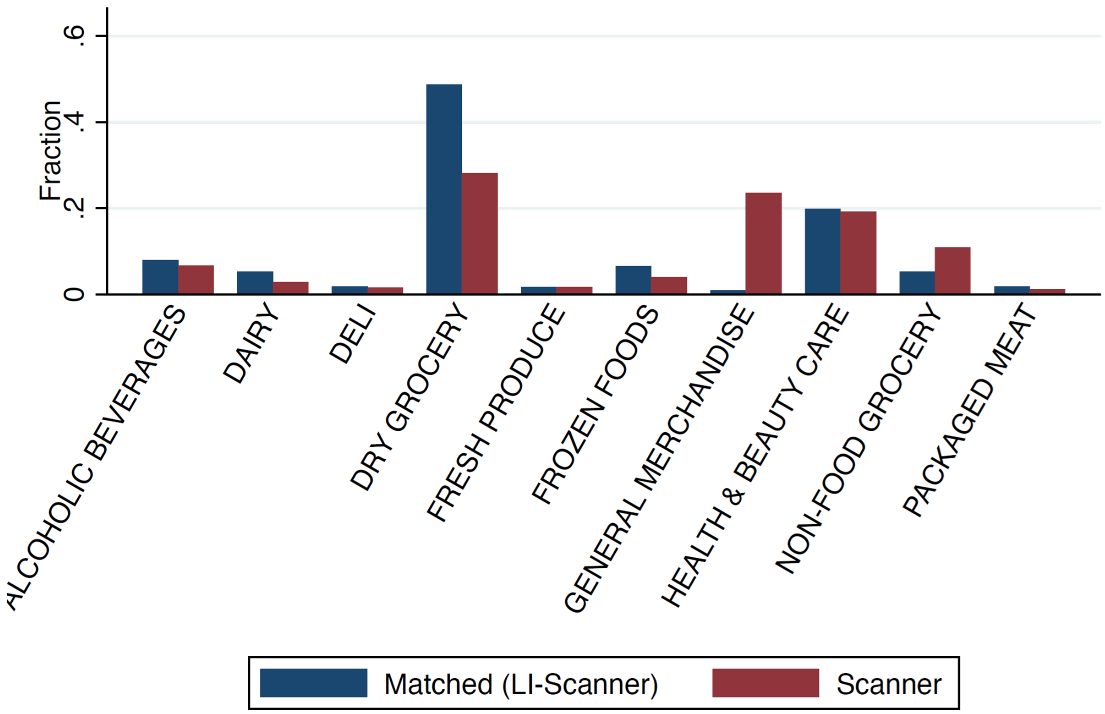

Notes: Figure 1 provides the distribution of barcodes across NielsenIQ product departments according to the raw scanner data (red) and the merged Label Insight-NielsenIQ dataset used in this paper (blue).

Figure A.2: Event Study Estimates for Unit Value Pass-through (Wave 4a)

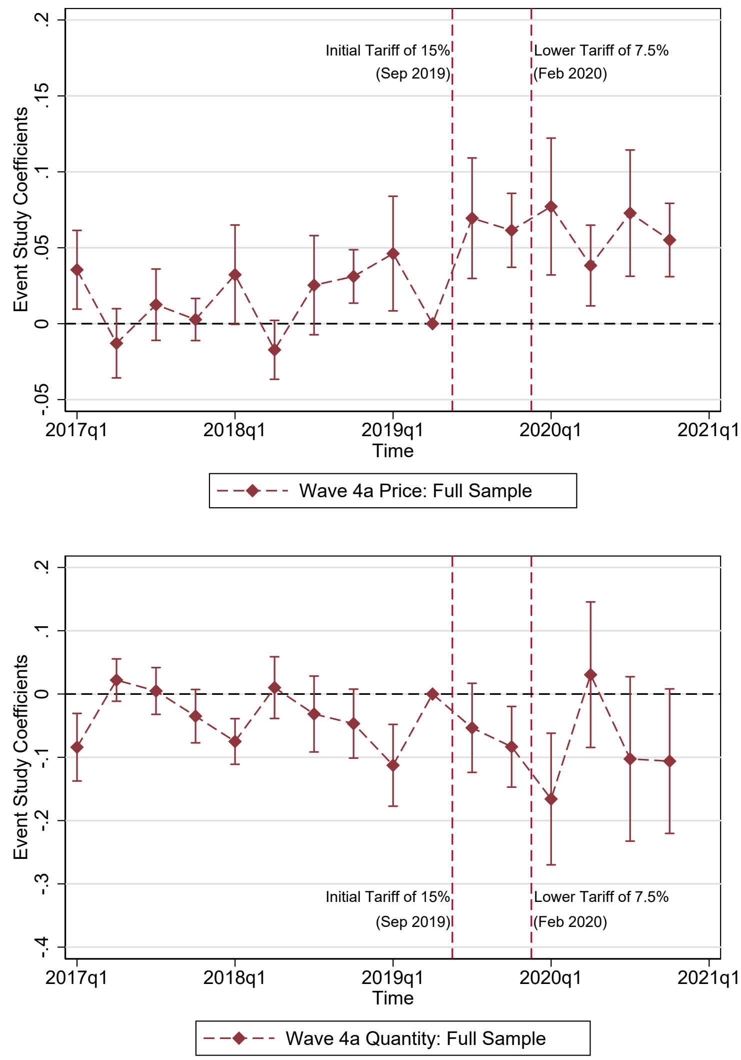

Notes: This figure provides estimates associated with the event study design at the aggregated country-category-retailer level outlined in Equation 2. All models include fixed effects by firm-country-category-retailer, retailer-month, category-month, and country-month. 95% confidence intervals included, with standard errors clustered three-way at the category, country, and month level. Only wave 4a and never-treated categories included. Red vertical lines denote tariff changes, from 5.0% to 20.0% in August 2019, and from 20.0% to 12.5% in January 2020. Back to Section 3.2.

Figure A.3: Event Study Estimates for Unit Value Pass-through by Barcode Tenure

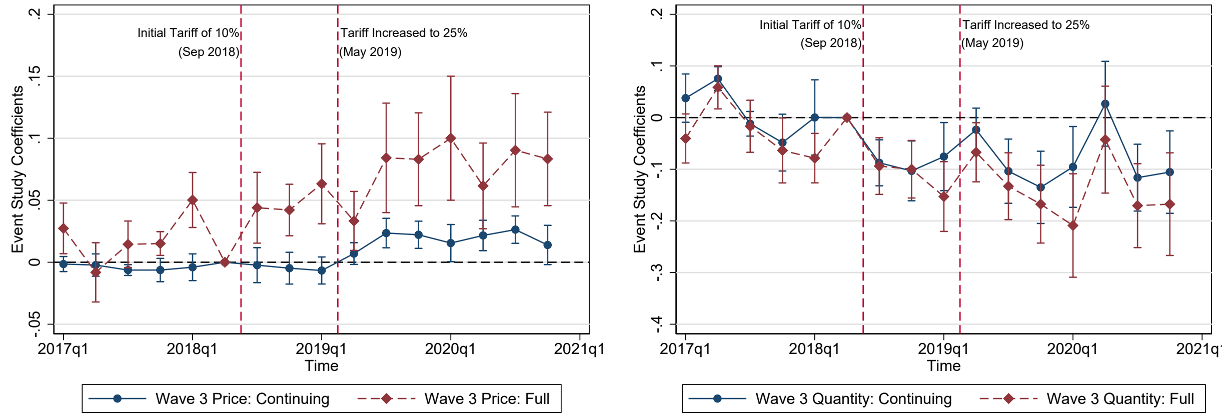

Notes: This figure provides estimates associated with the event study design at the aggregated country-category-retailer level outlined in Equation 2. All models include fixed effects by country-category-retailer, retailer-month, category-month, and country-month. 95% confidence intervals included, with Standard errors clustered three-way at the category, country, and month level. Back to Section 4.2.

Figure A.4: Event Study Estimates for Unit Value Pass-through by Barcode Tenure (Wave 4a)

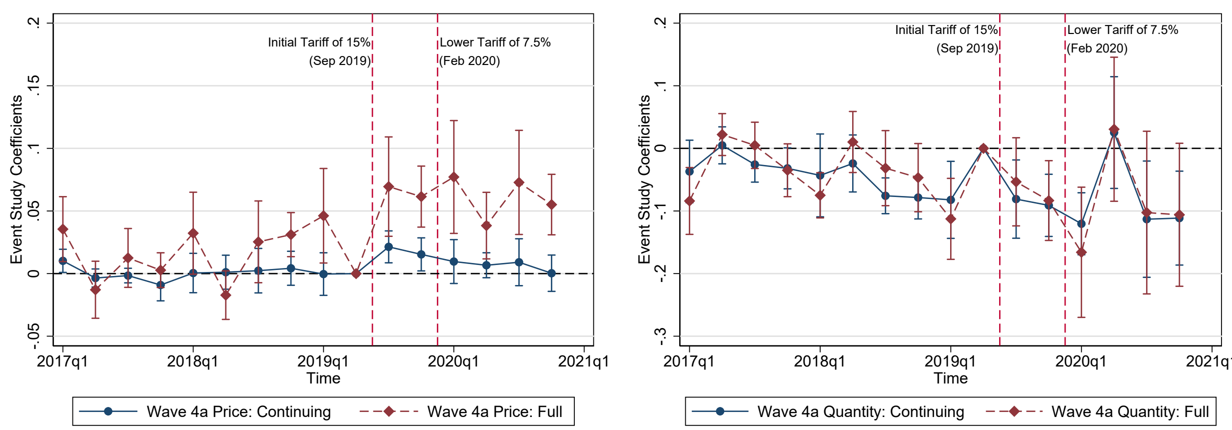

Notes: This figure provides estimates associated with the event study design at the aggregated firm-country-category-retailer level outlined in Equation 2. All models include fixed effects by firm-country-category-retailer, retailer-month, category-month, and country-month. 95% confidence intervals included, with Standard errors clustered three-way at the category, country, and month level. Back to Section 4.2.

Figure A.5: Event Study Estimates at Barcode Level (wave 3)

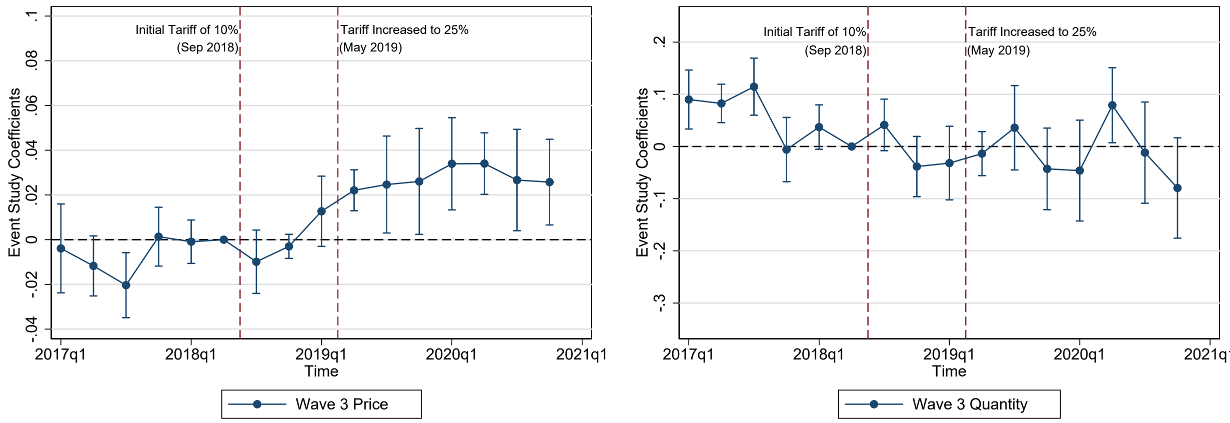

Notes: This figure provides estimates associated with a barcode-retailer-level event study. The dependent variable is p_{irt} (left) and q_{irt} (right). All never-tariff categories are included as the control group, whereas the only treated group included are those categories targeted by tariff wave 3. Barcode-retailer, retailer-month, category-month, and country-month fixed effects included. 95% confidence intervals included, with Standard errors clustered three-way at the category, country, and month level. Back to Section 4.2.

Figure A.6: Event Study Estimates at Barcode Level (wave 4a)

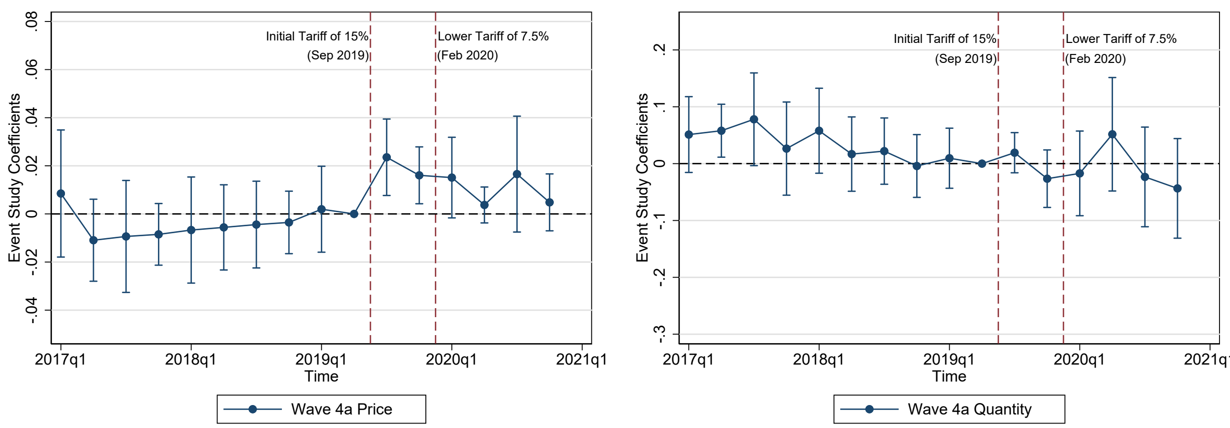

Notes: This figure provides estimates associated with estimating Equation 3. The dependent variable is p_{irt} (left) and q_{irt} (right). All never-tariff categories are included as the control group, whereas the only treated group included are those categories targeted by tariff wave 4a. Barcode-retailer, retailer-month, category-month, and country-month fixed effects included. 95% confidence intervals included, with Standard errors clustered three-way at the category, country, and month level. Back to Section 4.2.

Figure A.7: Event Study Estimates for Decomposition of Unit Value Pass-through (Wave 4a)

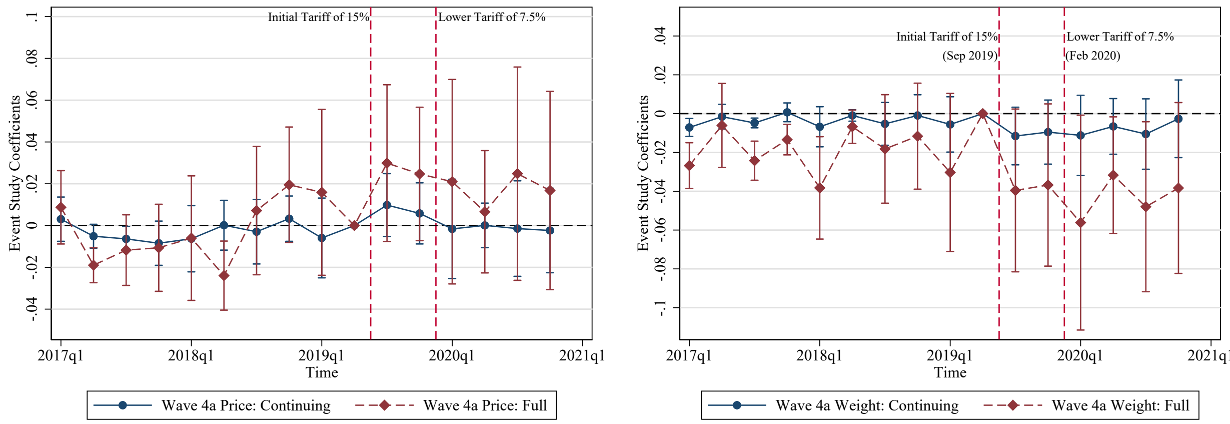

Notes: This figure provides estimates associated with the event study design at the aggregated firm-country-category-retailer level outlined in Equation 2. All models include fixed effects by firm-country-category-retailer, retailer-month, category-month, and country-month. 95% confidence intervals included, with Standard errors clustered three-way at the category, country, and month level.

Figure A.8: Firm and Retailer Bilateral Dependence

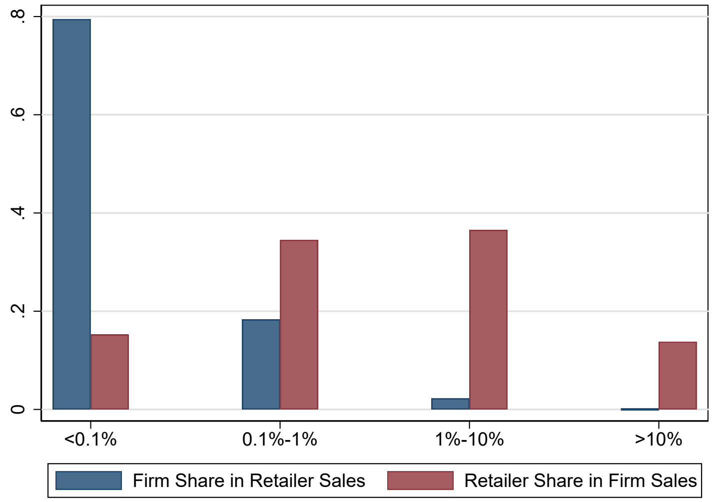

Notes: Distribution of the firm share in retailer sales and the retailer share in firm sales in 2017. Both variables are computed for 2017 data. Firm share in retailer sales is the total sales by products of firm f in retailer r, divided by total sales by retailer r. Retailer share in firm sales is the total sales by products of firm f in retailer r, divided by total sales by firm f.

B Empirical Appendix

This section provides further detail regarding data construction, empirical methodology, and results.

B.1 Constructing Foreign Cost Share Estimates

While we do not have data on the border prices of individual barcodes, we make progress by comparing category-by-origin unit values between retail and customs data. Since the U.S. computes tariff duties using free-on-board (FOB) prices, we use the HS10 customs valuation unit values by country of origin, as opposed to cost-in-freight (CIF) unit values which include international transport and insurance. We drop observations with invalid or missing quantity units and aggregate the import data to the level of product category-by-origin-by-unit of measurement, and compute unit values at that level.

We then try to harmonize as much as possible quantity units of measurement in both the scanner and the customs data. Some unit values cannot be compared across the two datasets, for instance if the unit in the scanner data is simply a "count", but it is measured in kilograms in the customs data. For the pairs for which we can identify a common quantity unit, we compute the foreign cost share as the ratio of the customs unit value and the scanner unit value. In rare cases, there are multiple quantity units of measurement in a given category-by-origin country cell. In all of these cases, one of the units is a product count while the other one is a more cleanly defined unit, such as volume or weight. In these cases we prefer the unit value measured from the non-count data. This procedure is arguably imperfect, since we cannot be sure that we compare the same barcodes in the customs and scanner data. Still, within the roughly two thousand country-category pairs we merge between datasets, 95% of all recovered estimates of the foreign cost share are less than one, which is reassuring as the theoretical upper bound of the FCS is one. Reassuringly, we recover an unweighted correlation between unit values in scanner data and customs data of 0.70. Figure B.1 provides a histogram of the foreign cost share estimates we recover across all imported country-category pairs.

We find no clear relationship with the foreign cost share estimates we recover and import

Figure B.1: Distribution of Foreign Cost Share Estimates

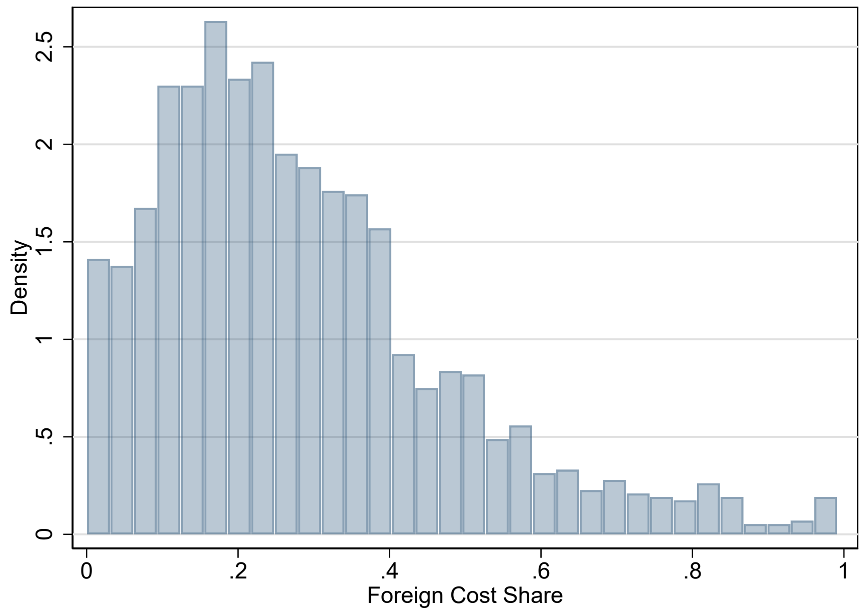

Notes: This figure provides a histogram of our foreign cost share estimates for all estimates which we reliably link between customs data and scanner data. $N = 1852$.

origin characteristics, such as distance to to the U.S. or GDP per capita, conditional on the product category in question. We do find some evidence that foreign cost shares associated with specifically with imports from China may have within-category foreign cost shares which are, on average, 3.6 percentage points lower than average. Still, this evidence remains qualitative at best, given the limited number of observations.

B.2 Robustness of Barcode-Level Estimates

This section describes a number of robustness checks associated with spillover effects of tariffs to non-treated barcodes. First, any model of strategic pricing would suggest that tariffs on Chinese varieties might induce an increase in markups for competing varieties from other origin countries. We test this effect directly using a differences-in-differences identification strategy, and comparing the effect of tariffs on the price of non-targeted products within tariff categories to products in non-tariff categories. Specifically, we estimate the following

Table B.1: Mark-ups of Non-Tariff Goods

	(1)	(2)	(3)	(4)
	log Price			
$\log(1+\tau_{CN})$	0.001	-0.001	0.019	0.011
	(0.021)	(0.019)	(0.026)	(0.023)
$\log(1+\tau_{CN}) \times (Share_{CN} > 0)$			-0.032	-0.020
			(0.025)	(0.022)
Sample		<2020		<2020
Fixed Effects	✓	✓	✓	✓
R^2	1.00	1.00	1.00	1.00
Observations	1.51e+07	1.13e+07	1.51e+07	1.13e+07

*Notes: This table provides estimates of β from applying OLS to Equation B.1. In columns (1) and (3) we estimate our specification on the whole sample, while in columns (2) and (4) we restrict attention to the comparison between wave 3 and the never treated, and stop the sample pre-Covid. Standard errors are clustered two-way at the category-country and month level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.*

empirical model:

$$\ln(p_{irt}) = \beta \ln(1 + \tau^{cn}_{c(i)t}) + \alpha_{ir} + \alpha_{rt} + \alpha_{j(i)t} + \alpha_{g(i)t} + \gamma_{w(i)q(t)} + \epsilon_{irt} \tag{B.1}$$

where $\tau^{cn}_{c(i)t}$ denotes the tariff applied to Chinese varieties within category c, which is now the explanatory variable for all non-Chinese varieties within that same category. All other fixed effects remain the same, other than the module-time fixed effects. Since we now exclude Chinese products from the estimation, these fixed effects would entirely absorb the tariff variation across categories. To still control for differential sector-level price trends we include fixed effects at a higher level of category classification, $\alpha_{g(c)t}$, where $g(c)$ denotes roughly 100 NielsenIQ product groups. Results are in Table B.1 and do not provide evidence for an endogenous mark-up response of non-tariff-affected goods. Estimated coefficients are all close to zero with relatively tight confidence intervals. It should be noted that in a similar exercise, Cavallo et al. (2021) also find little evidence for anti-competitive behavior within consumer goods. However, within consumer packaged goods, China's expenditure share in the U.S. is smaller than in many other export markets, with shares that are typically less than 5% of all sales. Thus, it is not immediately clear that the combination of low market shares and (relatively) small tariff changes would lead to observable anti-competitive behavior.

Second, we consider the possibility that firms or retailers may choose to spread price increases across many goods instead of concentrating them on the taxed items. While there is some anecdotal evidence for this type of strategy, particularly by retailers, we are not aware of any study testing whether this is a common pattern.[31] We test this mechanism as follows: do retailers (or firms) more exposed to a particular tariff wave tend to increase their prices of non-taxed goods by more than other retailers? To avoid retailers (or firms) being treated multiple times through different waves, we again restrict attention to the treatment through wave 3, and stop the sample in December 2019. We then define a set of treated firms and retailers using information on the firm and retailer product portfolio in 2017. Among the roughly 4,200 firms in our sample, we define a firm as treated if in 2017 it sold some Chinese products that were later taxed through Wave 3. This applies to 188 firms. Since all 67 retailers have some level of exposure to Chinese Wave 3 goods, we divide retailers by the median 2017 exposure. We then estimate the following specification on only non-Chinese barcodes,

$$\ln(p_{irt}) = \beta \text{Treated Retailer}_r \times \text{Post}_t + \alpha_{irt} + \epsilon_{irt}, \tag{B.2}$$

where $textPost_t$ is a dummy that equals one after the implementation of the initial Wave 3 tariffs in September 2018.

Since we want to identify the effect using variation across retailers, we now need to suppress the retailer-by-time FE. In column (1) of Table B.2 we estimate equation B.2 using barcode-by-retailer, country-by-time, and module-by-time effects, and we do the same in column (2) when we look at the firm effect. Column (3) adds both effects in the same specification. Column (4) then implements a more demanding specification to estimate the firm effect, where we suppress all cross-retailer variation. Finally, specification (5) implements a more demanding specification to estimate the smoothing effect for retailers, using only variation for the same barcode sold at the same point in time, across different retailers, which we do by adding a barcode-by-time effect.[32]

[31] As an example: https://www.oliverwyman.com/our-expertise/insights/2025/jan/retail-strategies-to-adapt-to-new-us-tariffs.html

[32] Notice that a barcode-by-time effect wipes out the firm smoothing effect, since a barcode always remains in the same firm.

Table B.2: Smoothing of Price Increases

	(1)	(2)	(3)	(4)
	Log Price			
Treated Retailer × Post	0.004**		0.004**	
	(0.002)		(0.002)	
Treated Firm × Post		0.003	0.003	0.002
		(0.004)	(0.004)	(0.004)
Sample	<2020	<2020	<2020	<2020
Fixed Effects	ir,mt,jt	ir,mt,jt	ir,mt,jt	ir,mt,rt,jt
R^2	1.00	1.00	1.00	1.00
Observations	1.51e+07	1.51e+07	1.51e+07	1.51e+07

Notes: In column (1) we estimate equation B.2 using barcode-by-retailer, country-by-time, and module-by-time effects, and we do the same in column (2) when we look at the firm effect. Column (3) adds both effects in the same specification. In column (4) we estimate a more demanding specificaiton with retailer-by-time FEs. Standard errors are clustered two-way at the category-country and month level. *, **, *and* *** *denote statistical significance at the 10%, 5%, and 1% levels, respectively.*

Table B.2 provides some evidence that, among non-taxed goods only, retailers that were initially more exposed to tariffs raised their prices by more, suggesting that retailers did smooth tariff increases. However, the magnitude is quite small, with prices of unaffected goods increasing by only 0.8-0.9% after implementation of Wave 3. While this channel is of interest *per se*, it is certainly not strong enough to explain the relatively low barcode-level pass-through estimates we recover of 10%.

Finally, while many consumer goods were taxed in waves 3 and 4 of the 2018/19 trade war, earlier waves (1 and 2) focused on intermediate goods. This might lead to price increases of domestically produced consumer goods to the extent that those goods use taxed imported intermediates as inputs in production. Since we cannot directly observe imported intermediates at the firm-level in our data, we instead use the fact that this concern only applies to U.S.-produced goods in the control group, but not to foreign-produced goods. In Table B.3, we re-estimate Equation 3, excluding U.S.-produced goods entirely. The coefficients stay very close in magnitude to our main estimates, suggesting that imported intermediates are not a first-order concern and do not bring us closer to estimates of pass-through in the range of 40%, which we recovered in Section 3.

Table B.3: Comparison to other Imported Products

	(1)	(2)
	Log Price	
$\log(1+\tau)$	0.101^{**}	0.093^{**}
	(0.044)	(0.042)
Sample		<2020
Fixed Effects	✓	✓
R^2	1.00	1.00
Observations	2.76e+06	2.07e+06

Notes: This table provides estimates of β when applying OLS to equation 3, excluding US-made goods from the estimation. Standard errors are clustered two-way at the category-country and month level. *, **, *and* *** *denote statistical significance at the 10%, 5%, and 1% levels, respectively.*

B.3 Barcode Survival, Entry, and Exit

This section provides the empirical model and estimation results regarding survival, entry, and exit of barcode-retailer pairs. We therefore make use of our barcode-retailer level dataset, such that all observations are at the irt level for barcode i, retailer r, and month t. For our main outcome, we simply indicate a value of one for barcode-retailer pair ir if period t exhibits non-zero expenditure recorded for ir, and zero otherwise. Thus, equation B.3 represents a linear probability model, with $\mathbf{1}[q_{irt} > 0]$ the binary indicator of survival described above:

$$\mathbf{1}[q_{irt} > 0] = \beta^{LPM} \ln(1 + \tau_{j(i)c(i)t}) + \alpha_{ir} + \alpha_{rt} + \alpha_{jt} + \alpha_{ct} + \gamma_{i,q(t)} + \epsilon_{irt}. \tag{B.3}$$

This panel is completely balanced across ir units and time periods t, such that that β^{LPM} represents the net entry/exit effect of tariffs. As always, we also estimate the model as an event study, with wave 3 denoting the treated categories and the never-tariff categories as control. As shown in Figure B.2, there is evidence of a downward effect of tariffs on the extensive margin, with the post-tariff window exhibiting an average probability of non-zero expenditure that is roughly five percentage points less than the stable pre-period level of non-zero expenditure. Thus, there is some evidence that tariffs generated a net exit of barcodes from tariff-affected units, which is expected. Table B.4 confirms this visual result: we estimate a negative and significant effect of tariffs on barcode-retailer non-zero

Table B.4: Survival, Entry, and Exit

	(1) Expdt.>0	(2) Exit	(3) Entry
$\log(1+\tau)$	-0.267***	0.292***	0.552***
	(0.082)	(0.082)	(0.045)
Observations	88,308,050	76,009,220	118,347,884

*Notes: Estimates associated with applying OLS to equation B.3. All models represent linear probability outcomes associated with a binary dependent variable, and all models include fixed effects at the barcode-retailer, retailer-month, module-month, and country-month level. Standard errors are clustered three-way at the category, country, and month level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.*

expenditure, with an estimate of -0.267 (0.082) in column (1).

In order to interpret our estimates in Section 3 as representative of product turnover, however, it is important to understand whether the negative effect on survival of tariffs documented above also allows for the possibility of barcode entry. We therefore re-estimate Equation B.3 but with two alternative dependent variables: an indicator for whether a given time period is after the last recorded sale of a given barcode-retailer pair (exit), and an indicator for whether a given time period is after a barcode-retailer pair recorded their first sale (entry). In each case, however, we opt for an unbalanced panel, such that when estimating the exit effects of tariffs, we exclude all barcode-retailer-months prior to entry, and when estimating the entry effects of tariffs, we exclude all barcode-retailer-months after exit. This is important, as these observations would receive a value of zero despite being unrepresentative of, for example, exit, when the ir pair in question had not yet entered.

These alternative samples, however, lead to baseline probabilities which differ across the groups, such that the aggregate effect of summing across the entry and exit effects of tariffs is not necessarily a useful exercise. We therefore continue with our "survive" estimates as the net entry effect of tariffs, with the entry and exit analyses instead capturing the sign and significance of tariff effects. Column (2) of Table B.4 exhibits a strong negative effect of tariffs on exit, with a point estimate of 0.292 (0.082), while column (3) indicates a strong effect of tariffs on entry, with a point estimate of 0.552 (0.045). Figure B.3 provides the relevant event study for this same entry variable, again indicating a large and significant effect of each tariff wave on entry and therefore product churn, with little evidence for pre-trends.

Figure B.2: Event Study of Barcode Survival

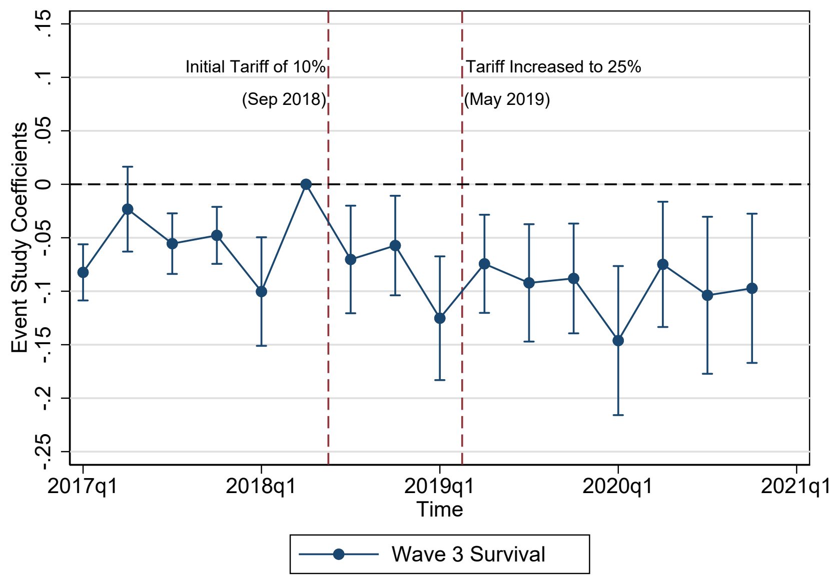

Notes: Event study estimates associated with an expenditure indicator at the barcode-retailer pair, as outlined in equation B.3. 95% confidence intervals provided. Standard errors are clustered three-way at the category, country, and month level.

Figure B.3: Event Study of Barcode Entry

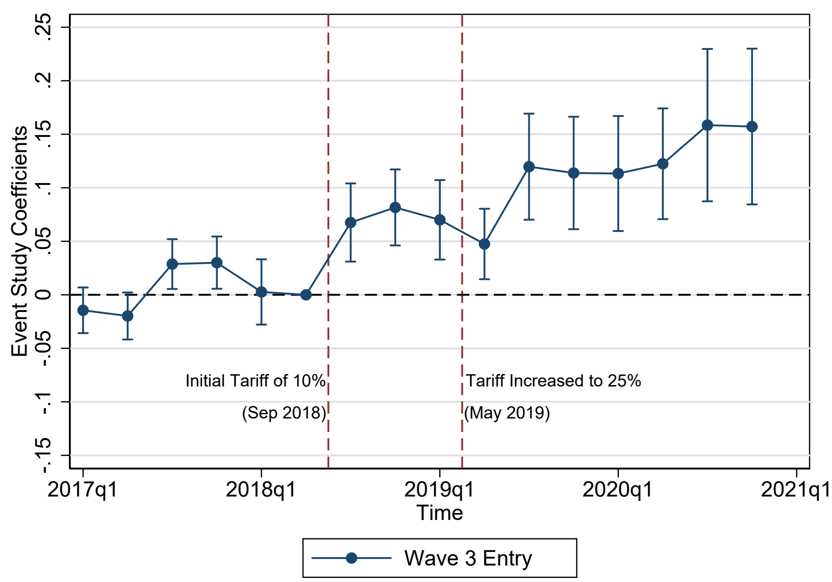

Notes: Event study estimates associated with an entry indicator at the barcode-retailer pair, as outlined in equation B.3. 95% confidence intervals provided. Standard errors are clustered three-way at the category, country, and month level.