Markups and the Distributional Impact of Tariffs
Evidence from Scotch Whisky

Minuk Kim*

Bryn Mawr College

mkim2@brynmawr.edu

August 2026

Abstract

When tariff pass-through is stronger for low-markup products, a uniform tariff falls hardest on consumers who buy low-markup products. Using product-level data on the 2019 U.S. tariff on single malt Scotch, with untariffed blended Scotch as control, I find a 14% retail price increase for cheap single malts versus 4–5% for premium products. Combining this pass-through with a demand model shows that low-income consumers lose a 1.6 times larger share of consumer surplus than high-income consumers do. Despite a uniform statutory tariff rate, income differences in price sensitivity generate regressivity, which heterogeneous pass-through amplifies by a quarter.

*An earlier version of this paper was circulated under the title "The Differential Impact of Tariffs by Quality." I am indebted to Thomas Holmes, Amil Petrin, and Joel Waldfogel for their advice and guidance. I thank seminar and conference discussants at the Midwest Trade Conference, IIOC, the Philly Trade Group at Drexel, the WIFO–WU Vienna Sanctions Workshop, the Econometric Society World Congress, Swarthmore College, and Haverford College for helpful comments. Researcher(s)' own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researcher(s) and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

1 Introduction

Aggregate evidence from the 2018–2019 trade war shows that tariffs pass through close to one-for-one into US import prices on average. But product-level variation in pass-through is substantial, and the average masks distributional consequences that depend on which consumers buy which products within a tariffed product category. If low-markup, commodity-tier products pass through the cost shock more than high-markup, premium-tier products, a uniform tariff more adversely affects low-income consumers who concentrate their purchases at the low-markup end of the quality ladder. The size of this channel is an empirical question with direct implications for trade policy design.

This paper measures that channel using the United States' 25% tariff on single malt Scotch (October 2019 to March 2021). Comparing tariffed single malts to untariffed blended Scotch, I find that a 7% average retail price increase conceals a 14% increase for cheap single malts against 4–5% for mid- and expensive-tier products, and that the price response declines steeply with the product's markup. The incidence is correspondingly skewed: low-income consumers bear roughly 1.6 times the proportional welfare loss of high-income consumers, and their probability of buying single malt falls by a third, against a quarter. Furthermore, substitution flows nearly entirely into lower-priced whiskey (blended Scotch and Irish), suggesting that the tariff's main behavioral effect is a downgrade along the quality ladder rather than a shift toward bourbon and other whiskeys.

The 25% tariff was authorized by the World Trade Organization in October 2019 as part of the United States' response to European Union subsidies of Airbus in a long-running dispute over civilian aircraft. The dispute began in 2004 and was unrelated to single malt Scotch, which makes the tariff plausibly exogenous to the industry. The tariff applied to single malt Scotch only, leaving blended Scotch (a near substitute facing nearly identical supply and demand shocks) as a credible control. I study the Pennsylvania market in particular, a "control" state where spirits are sold only through a state regulator, which suits the research question for two reasons. First, the regulator, the Pennsylvania Liquor Control Board (PLCB), reports a census of product-level sales, which allows for measuring the tariff effects on every product and thus provides a complete picture of the consumer substitution behavior. Second, the PLCB priced non-strategically during the period of interest, i.e., during the tariff window, the PLCB initiated no retail price changes of its own. As it only passed supplier-initiated cost changes through pre-existing product-specific markups, the retail response to the tariff reflects the upstream cost shock rather than a retailer's own markup adjustment we would observe in other states. Section 2 discusses the tariff, the industry, and the PLCB's pricing in detail.

The PLCB publishes weekly product-level sales in dollars and units for every store in Pennsylvania along with quarterly retail price catalogs. I use these to construct a balanced quarterly panel of 105 single malt and blended Scotch products from January 2019 through March 2021. Wholesale prices come from the New York State Liquor Authority (NYSLA) monthly distributor reports. The structural demand model also draws on the NielsenIQ Consumer Panel for income-conditional purchase patterns.

Section 3 develops the reduced-form estimates. The price response builds gradually over the tariff window, consistent with the staggered renegotiation of PLCB supplier pricing contracts. I use pre-tariff prices to split the products into three terciles, and use a difference-in-differences design by comparing the changes in retail prices relative to the control group of blended Scotches during the tariff window. I also estimate a saturated event-study triple difference-in-differences model with a continuous interacted variable to formalize the cheap-versus-expensive gap as a single coefficient. By the last quarter of the tariff window (Q1 2021), the tariff's retail price effect is sharply declining in that variable, whether it is pre-tariff retail price, Whisky Advocate rating, or structurally derived Lerner markup. The same pattern appears upstream in New York wholesale prices, which are posted to a competitive market in a neighboring state.

Pre-tariff retail price is a composite of cost, brand premium, vertical quality, and market power, and thus may be a proxy for markups and, under a similar logic, so are expert product ratings. The markup-based specification, in turn, is the structural counterpart to the observable-proxy results. I find that at the tariff's mature peak (Q1 2021), the pass-through heterogeneity implies a single-malt retail price increase of 20% for products in the 10th percentile of the markup distribution and 2% in the 90th percentile.

Section 4 estimates a random coefficients nested logit (RCNL) model on the same panel to recover consumer heterogeneity in price sensitivity by income. Consumers choose each pricing period to purchase a bottle of Scotch or Irish whiskey or to consume an outside good, with utility determined by price, alcohol content, bottle size, and brand effects. Substitution patterns are governed by an unobserved consumer-specific shifter on price and a three-nest structure (single malt Scotch, blended Scotch, and Irish whiskey, plus the outside good). I model consumers as belonging to one of three income bins (under \$45K, \$45K–\$100K, over \$100K) with distinct price coefficients estimated by matching purchase patterns observed in the NielsenIQ Consumer Panel. Prices are instrumented using the tariff indicator variable, New York wholesale prices as a cost shifter, and within- and cross-nest differentiation instruments.

Section 5 estimates the pass-through schedule in the model-derived markups and shows that, to a first-order approximation valid in any demand system, the within-category regressivity of a uniform tariff equals the product of two sufficient statistics: the rate at which pass-through declines with markups and the difference in the average markup of the products that low- and high-income consumers buy. It then combines the schedule with the demand model to compute counterfactual prices and consumer welfare under the heterogeneous tariff. The relevant welfare object is compensating variation, which integrates over consumer substitution via the inclusive-value identity for nested logit. A two-by-two decomposition separates this regressivity into a demand-side channel (income heterogeneity in price sensitivity) and a supply-side channel (heterogeneous pass-through). Income heterogeneity alone produces the standard regressive incidence, pass-through heterogeneity alone produces none, but together, the tariff is about a quarter more regressive than a uniform-pass-through welfare exercise would find.

Two qualifications apply. First, blended Scotch is itself a substitute for the tariffed products, so substitution away from single malt could raise control-group prices. Any such spillover attenuates the single-malt-versus-blended differential, and the reported pass-through magnitudes are therefore lower bounds on the partial-equilibrium response. Second, the welfare exercise holds the product set and the prices of non-tariffed goods fixed, so it captures the intensive-margin response to the estimated price changes and abstracts from product exit and any associated loss of variety.

This paper relates to several literatures. The empirical literature on tariff pass-through ranges from industry-level studies (Feenstra 1989, Irwin 2019, Flaaen et al. 2020) to the recent aggregate analyses of the 2018–2019 US-China trade war (Fajgelbaum et al. 2019, Amiti et al. 2019b, Amiti et al. 2020, Cavallo et al. 2021), which find pass-through close to one-for-one on average. The closest related work documents the same mechanism on the supply side. Chen and Juvenal (2022) document cross-destination heterogeneous markup adjustment for Argentinean wine exports, with high-quality varieties absorbing more of the variable trade cost in their FOB prices. Sangani (2026) shows that across retail gasoline, food products, and U.S. manufacturing, log pass-through of common cost shocks is systematically lower for products with larger gaps between price and marginal cost, so a uniform cost shock generates heterogeneous percentage burdens across products. Relative to these papers, this paper estimates the response to a single discrete policy shock, corroborates the markup mechanism with markups estimated from a structural demand model, and uses the same model to quantify the welfare incidence across the income distribution. Measuring the full chain in a single market requires a product-specific exogenous cost shock, a matched untreated control, census sales data, and measurable markups, which the Pennsylvania Scotch market provides.

A second strand examines how trade shocks affect households through their consumption baskets. Fajgelbaum and Khandelwal (2016) study this mechanism for import exposure, while Auer et al. (2024) examine exchange-rate movements. Closest to this paper, Cravino and Levchenko (2017) decompose the distributional effects of the 1994 Mexican devaluation into across-category and within-category components. They find that poorer households tend to buy cheaper varieties within categories, and that those varieties, which carry lower retail distribution margins, experienced larger price increases after the devaluation. They also note that variable markups across varieties could generate the same pattern but are not separately measurable in their data. In this paper, the Pennsylvania setting where the state retailer held product-specific markups fixed through the tariff window largely removes the retail distribution-margin channel, so I can examine the supplier markup channel in relative isolation. Acosta and Cox (2024) identify another source of regressivity. They find that tariff rates are systematically higher on lower-end variants of goods than on higher-end variants, a pattern that originated in twentieth-century trade negotiations and persists today. This paper studies a complementary channel. A uniform tariff rate on a single product category (25% across all single malt Scotch) is also regressive in incidence because pass-through itself varies with firm market power. Thus, a tariff schedule that is structurally regressive and that also impacts cheap, low-markup products more, even at uniform rates (this paper), is more regressive than either mechanism alone implies.

Finally, the paper draws on the literature on demand curvature and market power in cost pass-through. Weyl and Fabinger (2013) characterize the relationship between demand curvature and pass-through under imperfect competition, and Griffith et al. (2018), Birchall et al. (2024), and Miravete et al. (2023a,b) examine the role of functional form restrictions on demand curvature in discrete-choice demand models. Empirically, markup adjustment as a source of incomplete pass-through has been documented for import costs, exchange rates, and tariffs in international settings (Nakamura and Zerom 2010, Goldberg and Hellerstein 2013, Amiti et al. 2014, Ludema and Yu 2016, Amiti et al. 2019a) and for tax incidence in domestic industry studies (Kim and Cotterill 2008, Allcott et al. 2019, Miravete et al. 2018, 2020). This paper's contribution is the integration of a clean reduced-form quasi-experimental design with a structural demand model on the same data, allowing the markup-mediated heterogeneity in pass-through to be both estimated reduced-form and microfounded structurally.

The remainder of the paper is organized as follows: Section 2 describes the setting and data, Sections 3–5 present the reduced-form, demand, and welfare analyses, and Section 6 concludes.

2 Institutional Context and Data

Scotch Whisky Industry

Due to the distinctive brand and image associated with Scotch, the Scotch industry is tightly regulated by the government of the United Kingdom. In 2009, the United Kingdom adopted *The Scotch Whisky Regulations 2009* (SI 2009/2890), which lay out nine conditions for a distilled spirit to be defined as Scotch, from the minimum years of maturation to the capacity of the casks. As a protected brand name under EU law, Scotch must be distilled and aged in Scotland, and single malt Scotch in particular must be bottled before it is exported. These regulations ensure that Scotch is produced in the UK, and thus production relocation of Scotch in response to the tariffs is limited. Single malt Scotch and blended Scotch constitute the majority of the market, accounting for 32% and 60% of Scotch exports by value in 2022, respectively. Single malt Scotch is a Scotch whisky produced from only water and malted barley at a single distillery by batch distillation in pot stills, with popular brands such as Glenlivet, Glenfiddich, and Macallan. Blended Scotch is defined as a combination of one or more single malt Scotch whiskies with one or more single grain Scotch whiskies, with popular brands such as Johnnie Walker, Ballantine's, Chivas Regal, and Dewar's. The U.S. is the largest export market for Scotch, with around 1 billion GBP in exports in 2022, constituting around one-sixth of total exports by volume.

The largest players in the industry are Diageo, Pernod Ricard, William Grant & Sons, with around 41%, 22%, and 8% of the Scotch market, by sales. These firms often produce both single malt Scotch whiskies and blended Scotch whiskies. Diageo, for example, produces a line of single malt Scotch whiskies from its Talisker Distillery but also uses whisky from the same distillery to blend with whiskies from its other distilleries to produce blended Scotches such as Johnnie Walker Black Label. Distilling, maturing, bottling, storing, and transportation costs are shared between single malt and blended Scotch, which suggests that any potential cost shock to single malt Scotch arising from Brexit or the COVID-19 pandemic that propagates to retail prices would be shared with blended Scotch as well. Thus, blended Scotch serves as a natural control group for studying the impact of the tariffs on single malt Scotch.

Single malt Scotch whiskies are known for each product's idiosyncratic taste and aroma coming from the particular distillery, aging, cask, and use of peat, amongst many other factors used in production. On the other hand, blended Scotches are made from whiskies from different sources and often "average out" these idiosyncrasies to create a consistent flavor. Scotch, but single malt Scotch in particular, is differentiated along two main dimensions. The age of the whisky is the main method of vertical differentiation. Age statements are generally found on single malt Scotch whiskies, and is determined by the youngest spirit in the blend[1]. Age is often associated with higher quality due to the complex flavors drawn out of the spirit due to the longer maturation time in the oak casks (although there is a point of decreasing returns at the extreme upper end). Thus, an older age bottle is generally considered a higher quality bottle, both within a brand and across brands (e.g., a Glenlivet 18 would generally be considered higher quality than a Glenlivet 12 *and* a Macallan 12).

Branding is another method for product differentiation, as a mix of both horizontal and vertical. For single malt Scotches, the brand is often the distillery from which the Scotch was produced, e.g., Macallan branded Scotches are distilled at the Macallan distillery in the Speyside region of Scotland. Distilleries are known for their distinctive characteristics; for example, the Laphroaig, Lagavulin, and Ardbeg distilleries are known for their smoky tastes coming from the peat used in the production process. Preference for this smoky taste, also known as "peatiness," is highly idiosyncratic, and thus creates horizontal differentiation. On the other hand, some brands, such as Macallan, are generally considered higher quality than others and act as a form of vertical differentiation. As the Scotch Whisky Association, an industry group, notes, "many consumers buy these products because of their provenance and are unlikely to shift to products produced elsewhere," suggesting that the branding is indeed a strong form of product differentiation.

These two dimensions of differentiation (age and brand) determine the overall (idiosyncratic or common) perceived "quality" of a bottle, and both command large price premia. Within a brand, longer age statements sell at substantially higher prices, and within an age statement, premium distillery brands sell at substantially higher prices.

Using Pennsylvania data, Table 1 shows that the number of products in each age statement is generally decreasing while the median price is increasing, suggesting that more vertically differentiated products do enjoy less immediate competition, which may lead to higher markups.

United States Section 301 Tariffs on the European Union

The Section 301 tariffs were imposed by the United States on the European Union (EU) in response to what the U.S. considered illegal subsidies for Airbus. The dispute dates back to 2004 when the U.S. first filed a complaint with the World Trade Organization (WTO) alleging that the EU provided prohibited subsidies to Airbus, negatively impacting Boeing's competitiveness. After years of lawsuits and counterclaims, in April 2019, the U.S.

[1]Single malt Scotch is technically a blend as well, just that the component spirits are from the same distillery

Age Statement	# Products	Median Price
12	46	$49.99
15	14	$67.49
18	23	$114.99
21	10	$209.99
25	10	$599.99

Table 1: Product Space by Age Statement

announced a list of products to be tariffed pending WTO approval, which was authorized in October 2019. This led to the U.S. imposing a series of tariffs on various European imports.

As tariffs are not assigned randomly, there is concern of selection bias in identifying the tariff effects. If the products to be tariffed were correlated with demand or cost shocks specific to those products, or motivated by protection of a domestic competitor, we would conflate the tariff's causal effect with these confounding forces. No direct evidence rules out such selection, but two features of the list weigh against this concern. First, the composition of the list is more consistent with political-symbolic targeting than with industry-driven selection (besides for the Airbus/Boeing aircrafts in question). In addition to single malt Scotch whisky, the U.S. tariffs covered Irish butter and cream, French and Italian cheese, Spanish olive oil, and English wool coats. The European Union's parallel retaliation list, already in force at the time, included Harley-Davidson motorcycles, bourbon whiskey, and Levi's jeans. Both lists target nationally iconic exports of the offending parties. On both sides the objective appears to have been political rather than economic harm. Second, a domestic-protection rationale would have been inconsistent with the product chosen. Had the goal been to shield U.S. whiskey producers from European competition, blended Scotch would have been the more appropriate target, since blends compete with American whiskey on price and character, while single malts occupy a premium segment in which domestic producers have limited presence. Taken together, these features support treating the tariff on single malt Scotch as plausibly exogenous to demand and supply conditions in the U.S. spirits market, and as a credible natural experiment for the analysis that follows.

In March 2021, the U.S. and the EU decided to suspend the tariffs for four months, and in June 2021, both sides reached an agreement to suspend them for five years, putting a temporary end to this long-standing trade dispute over aircraft subsidies. Ahead of the suspension's June 2026 expiry, the U.S. announced the permanent removal of the tariffs on Scotch whisky, in connection with King Charles III's 2026 state visit to the United States; like its imposition, the tariff's removal also appears to have been driven by diplomatic considerations rather than by conditions in the whisky market.

Data Sources

The main data source used for both evaluating the impact of the tariffs and for estimating the demand model comes from the Pennsylvania Liquor Control Board (PLCB), acquired under a Right to Know Law request. The data contains the sales amount in dollars and sales quantity in units for every product sold at each liquor store in Pennsylvania, for each week from January 2019 to March 2021. I use a combination of a catalogue spreadsheet file from the PLCB archives and the quarterly product listing published by the PLCB to identify the Scotch products in the sales data. I aggregate each weekly sales record to the pricing-period level (27 periods, each roughly a month) and each store to the state level. The structural demand model in Section 4 uses these pricing periods as its markets; the reduced-form analysis in Section 3 further aggregates them to nine quarters (Q1 2019–Q1 2021). The PLCB also maintains constant prices throughout most of the state[2].

Until the enactment of Act 39 of 2016, the PLCB priced every product by a statutory formula applying a uniform 30 percent markup over the supplier acquisition cost. But since 2016, while the statutory default remains that retail prices "shall be proportional with prices paid by the board to its suppliers," plus a per-bottle handling fee, there is a carve-out permitting the Board to price its best-selling items (the 150 most-sold brands and product types of liquor and of wine, measured by units sold and recalculated semi-annually) and one-time limited-purchase buys "in a manner that maximizes the return on the sale of those items."[3] In practice the Board runs flexible pricing through acquisition-cost negotiations with suppliers, whose wholesale prices are confidential, with occasional Board-initiated retail adjustment rounds announced by supplier notice. While retail-side markups can be a relevant channel for tariff incidence in general (Cole and Eckel 2018), the Board's statutorily required annual pricing reports document its conduct during the tariff window.[4] "The PLCB did not pursue any retail price increases in 2020," with the listed

[2]With a couple exceptions (customer-relations-management offers and clearance of delisted items), which are small and the analysis aggregates stores to the state level.

[3]47 P.S. §2-207(b). The carve-out is capped at the 150 best-selling brands and product types of each of liquor and wine, though the PLCB's annual pricing reports describe all listed products as subject to flexible pricing. The same subsection requires a single uniform retail price throughout the state.

[4]Pennsylvania Liquor Control Board, *Annual Report on Method and Rationale for Product Pricing*, April 1 of 2020, 2021, and 2022.

portfolio instead recording 373 items with supplier-initiated retail price increases alongside 372 supplier cost increases, and the Board "did not unilaterally initiate retail price increases for a second consecutive year" in 2021, when 553 supplier cost increases drove 534 retail increases. The one Board-initiated adjustment round planned during the window, announced to suppliers in early March 2020 for implementation that August, was halted when the pandemic's economic toll became apparent. The retail price changes observed in Pennsylvania during the tariff window therefore reflect supplier-driven changes in acquisition cost, passed through under the PLCB's pre-existing product-specific markups.

However, to better understand supplier markup adjustments, I also collect wholesale price data from New York state. While not a state-owned retailer as in Pennsylvania, New York operates a post-and-hold system under which wholesalers must publicly post the monthly prices they charge retailers. These distributor list prices sit one layer below PLCB retail. They exclude the retail markup but include the wholesaler's margin, so they isolate upstream cost pass-through better than retail prices, though they do not observe the supplier's FOB price directly. The NYSLA wholesale prices support a robustness check on the role of the PLCB's monopoly and monopsony position in driving the results. As the sole licensed buyer and seller of distilled spirits in Pennsylvania, the PLCB faces every supplier as a single buyer and every consumer as a single seller. A monopsony and monopoly might respond to cost shocks differently than buyers and sellers in competitive wholesale markets. New York instead licenses many wholesalers and retailers, so the NYSLA data provides a more liberalized market to compare against for any tariff-induced pricing patterns observed in Pennsylvania. The wholesale prices are also used to identify exogenous price shifts in the demand model.

Consumer demographic and purchase data come from the NielsenIQ Consumer Panel. I restrict the sample to Pennsylvania households, classify each into one of three income bins, and compute income-conditional purchase statistics that link household income to price, product choice, and substitution patterns in the demand model below.

To proxy product quality in the reduced-form heterogeneity analysis (Section 3), I use external Whisky Advocate (WA) ratings, which are scored on a 50–100 scale by independent expert reviewers, with most ratings predating the tariff.

3 Heterogeneous Pass-Through of the Tariff

Tariff Impact on Prices

Figure 1 plots quarterly average prices of single malt and blended Scotch from Q1 2019 to Q1 2021, with the shaded region marking the tariff window (Q4 2019 onward). Average prices are computed on a balanced panel to remove product entry and exit from the comparison. Single malt prices increase noticeably from Q2 2020 onward, lagging the October 2019 tariff onset[5]. The tariff period overlaps with some macro-level shocks (e.g. Brexit and the COVID-19 pandemic), but blended Scotch shares the supply chain and was not subject to the tariff, making it a natural control for these shocks. Blended Scotch prices stay flat throughout the tariff-window, so the increase in single malt prices suggests a pricing response to the tariffs.

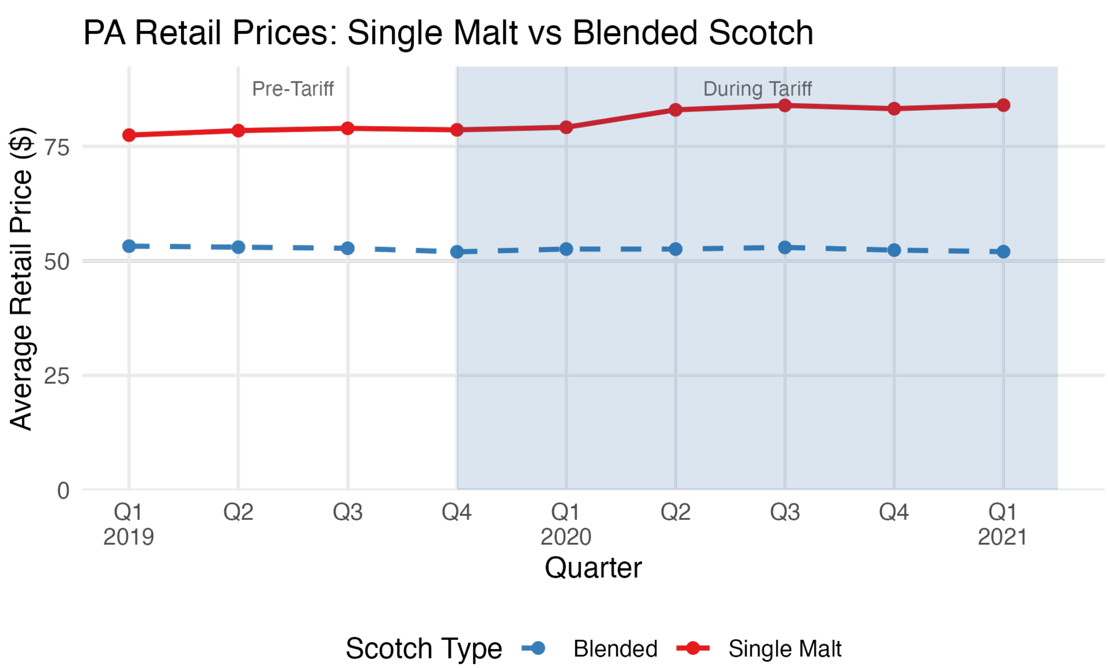

Figure 1: Average Price of Single Malt and Blended Scotches in Pennsylvania

To quantify the causal effect of the tariffs on prices, I estimate the following two-way fixed effects model,

$$ln(p_{it}) = \gamma_i + \lambda_t + \sum_{t=1}^{2} \beta_t D_{it} + \sum_{t=4}^{9} \beta_t D_{it} + \varepsilon_{it} \tag{1}$$

[5] The lag is consistent with PLCB pricing contracts that fix retail prices several quarters in advance.

where p_{it} is the price of product i in quarter t, γ_i is a product fixed effect, λ_t is a quarter fixed effect, and D_{it} is an interaction of the quarter fixed effect with a dummy variable indicating that product i is a single malt Scotch. The third quarter of 2019, the last full quarter before the tariff took effect on October 18, 2019, is the omitted reference period. The coefficients on the interaction terms (β_t's) measure the average price difference between single malt and blended Scotch relative to that reference, with β_1 and β_2 testing for pre-trends and β_4 through β_9 tracing the dynamic treatment effect. The data span Q1 2019 through Q1 2021, ending in March 2021 with the repeal of the tariff.

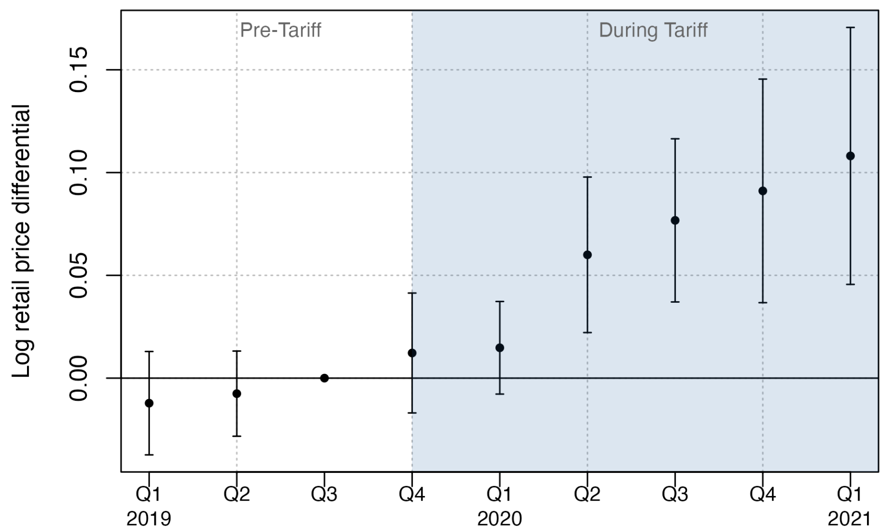

Figure 2: Single Malt versus Blended Price Difference

The corresponding event-study plot is in Figure 2 (full per-quarter coefficients are reported in Appendix Table 7). The pre-tariff coefficients β_1 and β_2 are statistically indistinguishable from zero, and the treatment coefficients build gradually. Single malt prices begin to diverge from blended in Q2 2020, and the gap widens through Q1 2021, where it reaches roughly 11%. Averaged over the during-tariff window (Q4 2019–Q1 2021), the pooled single-malt retail price increase is 6.9%. The trajectory has not visibly flattened by the end of the sample, so the steady-state response may be larger than the Q1 2021 estimate. The gradual buildup is consistent with PLCB pricing contracts that fix retail prices several quarters in advance and renegotiate on staggered schedules.

The pooled estimate may mask substantial heterogeneity in pass-through by the product-level markups. The remainder of this section presents reduced-form evidence that pass-through varies systematically with product-level markup, with high-markup products showing smaller price increases than low-markup products. But since markups are not directly observed prior to the structural demand estimation in Section 4, the analysis here relies on two observable measures: pre-tariff price tiers and Whisky Advocate (WA) quality ratings. The version using structurally-estimated markups is deferred to Section 5.

The single malt (SM) sample is partitioned into three bins by pre-tariff price (cheap, mid, and expensive),[6] with blended Scotch products in the same SM-defined dollar ranges serving as the within-tier control group. Using same-priced blended Scotch as the control makes each within-tier comparison closer to apples-to-apples, with cheap single malts matched against cheap blended Scotches, mid-tier against mid-tier, and so on. Price-correlated supply and demand shocks then affect both groups symmetrically and net out of the DiD. Estimating Equation 1 separately in each bin yields the per-tier event-study coefficients plotted in Figure 3 (full per-quarter coefficients are reported in Appendix Table 8).

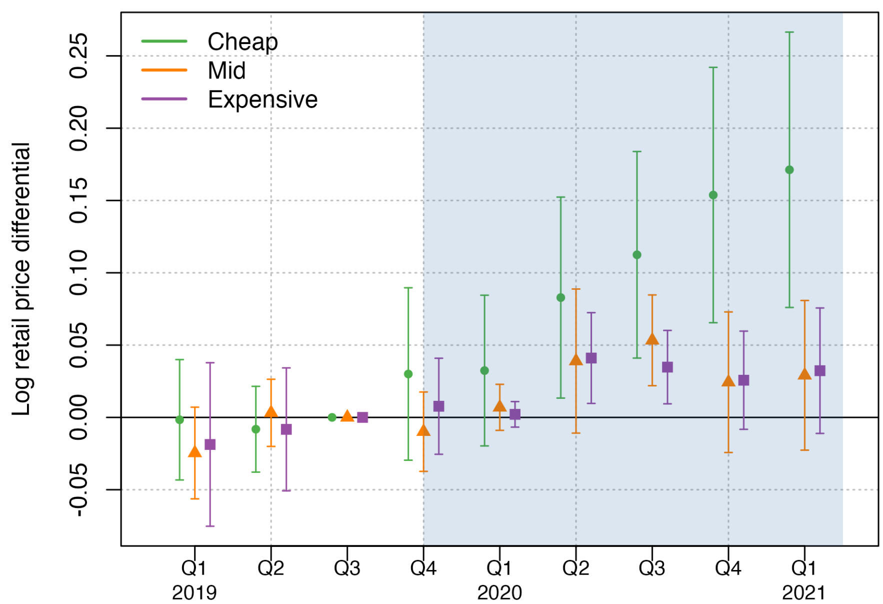

Figure 3: Event Study by Pre-Tariff Price Tercile

Pre-tariff coefficients are statistically indistinguishable from zero across all three terciles, consistent with no pre-trends within each price band. But treatment coefficients then differ sharply by tercile. The cheapest tercile shows the largest price increases by Q1 2021, while the middle and most expensive terciles show much smaller increases, statistically indistinguishable from blended Scotch in most quarters. The pooled estimate

[6] Khandelwal (2010) finds that retail prices are highly correlated with quality (defined as market shares conditional on price) in product categories with wide quality variation, supporting the use of pre-tariff retail price as a vertical-differentiation proxy.

therefore masks substantial heterogeneity, with the cheapest products exhibiting a much larger price response to the tariffs than the most expensive products.

A formal cross-tercile test is reported in Appendix Table 9, where a fully interacted specification that pools the mature treatment window (Q6–Q9) against the pre-tariff quarters puts the cheap-tier mature price increase at +14.3% (two-sided $p < 0.001$), with mid- and expensive-tier differentials of -9.8 and -9.9 percentage points relative to the cheap tier (two-sided $p = 0.030$ and 0.031). The heterogeneity is best summarized as a cheap-versus-other-tiers gap rather than a smooth quality gradient at the tercile level. The tercile heterogeneity also replicates upstream in New York wholesale prices, presented below.

Continuous Heterogeneity Gradient

In addition to the tercile contrasts that summarize the heterogeneity through a few bin comparisons, estimating the pass-through heterogeneity as a single coefficient in a continuous specification may provide a sharper test of the hypothesis and use the full cross-product variation. I apply that specification to the product's Whisky Advocate (WA). The rating, set by expert reviewers years before the tariff, isolates quality, and unlike price, it is not an equilibrium market outcome, so it carries no cost or markup component and cannot itself respond to the tariff.

The sample shrinks slightly to 102 products as three balanced-panel products are not covered by Whisky Advocate. The specification is as follows:

$$\begin{aligned}\ln(p_{it}) = \gamma_i &+ \lambda_t \\ &+ \sum_{q\neq 3} \theta_q \, \mathbf{1}\{t = q\} \, \mathrm{SM}_i \\ &+ \sum_{q\neq 3} \phi_q \, \mathbf{1}\{t = q\} \, \tilde{r}_i \\ &+ \sum_{q\neq 3} \kappa_q \, \mathbf{1}\{t = q\} \, \mathrm{SM}_i \, \tilde{r}_i + \varepsilon_{it}\end{aligned} \tag{2}$$

where $\tilde{r}_i = \mathrm{rating}_i - \overline{\mathrm{rating}}_{\mathrm{SM}}$ is the WA rating centered on the mean rating across single malts, and Q3 2019 is the omitted reference (so $\theta_3 = \phi_3 = \kappa_3 = 0$). The coefficient κ_q measures how the SM-Blended price gap varies with quality in quarter q. A negative κ_q means the tariff-period price increase is smaller, in percentage terms, for higher-rated single malts than for lower-rated ones.

Equation 2 is a triple difference-in-differences. The first difference is over time. Quarter fixed effects absorb shocks common to all products, with Q3 2019 the omitted reference. The second is single malt versus blended. Product fixed effects absorb time-invariant product levels, so the θ_q terms recover the standard SM-versus-blended DiD. The third is across quality. The ϕ_q terms absorb any rating-correlated price trend common to both categories (e.g. quality-tilted inflation or pandemic-era premiumization), so κ_q is identified from the within-single-malt rating gradient of price changes net of its within-blended counterpart. Because $\tilde{r}_i$ enters continuously, the third difference is taken as a derivative rather than as a contrast between two bins (discretizing $\tilde{r}_i$ into high and low groups would recover the more-familiar triple-DiD coefficient). The identifying assumption is correspondingly a parallel-gradients condition: absent the tariff, the quality gradient of price changes would have evolved identically for single malt and blended Scotch. The specification is saturated in event time (every non-reference quarter receives its own θ_q, ϕ_q, and κ_q), so the gradual contractual phase-in documented above requires no choice of pooling window, and the pre-tariff coefficients κ_1 and κ_2 serve as direct placebo tests of the parallel-gradients condition.

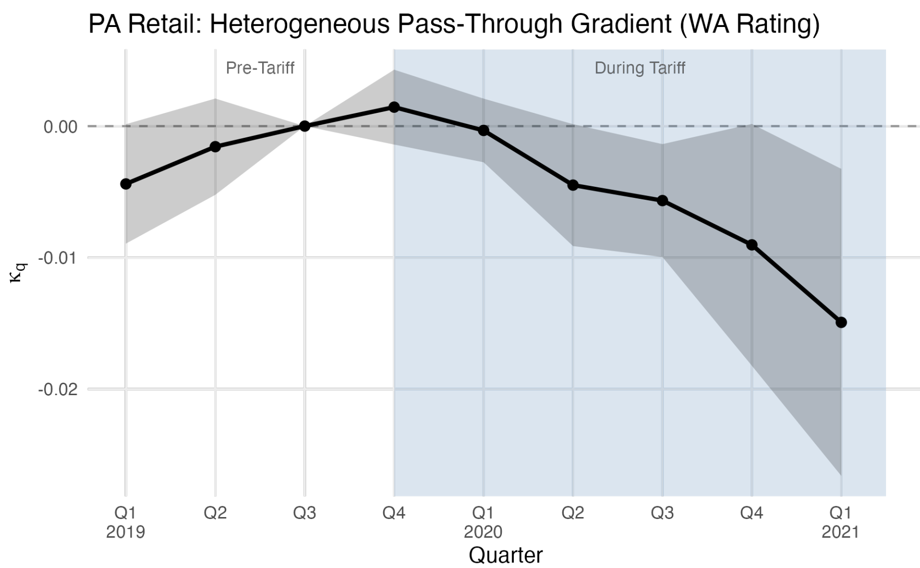

Figure 4: Per-Quarter SM-Blended Pass-Through Slope (κ_q) in the WA Rating. Plotted in raw κ_q units — the SM-specific log-price slope per one-point change in WA rating, with WA rating centered at the SM mean.

Figure 4 reports the per-quarter κ_q estimates. The series stays near zero through the early-tariff quarters and deepens monotonically through the mature window, reaching $\kappa_9 = -0.0149$ in Q1 2021 (SE 0.006, two-sided $p = 0.014$). In implied pass-through terms, products at the 10th percentile of WA rating (rating 83) face an implied retail price increase of +18.4%, while products at the 90th percentile (rating 93) face only +1.3%, a 17 percentage-point spread. The pre-trend evidence is mixed. The cluster-robust Wald test on $\kappa_1 = \kappa_2 = 0$ rejects ($F = 3.97$, $p = 0.019$), driven primarily by κ_1 (two-sided $p = 0.061$), while κ_2 is insignificant ($p = 0.40$). Reassuringly, the slope is null in Q4–Q5 (the early-tariff window before PLCB contracts had renegotiated) and only deepens monotonically from Q6 onward, consistent with PLCB's contractual phase-in.

The rating slope, the tercile contrasts, and the price-based versions (Pennsylvania retail in Appendix A.3, New York wholesale below) all point the same way. Higher-quality single malts show smaller price increases than lower-quality single malts, with the gap widening through the mature window. The Pennsylvania retail price slope is the weakest of these ($\kappa_9 = -0.098$, two-sided $p = 0.094$), consistent with price's threshold structure, but it agrees in sign and timing. The structural test, in which model-estimated markups from the demand model in Section 4 are used directly as the interacted variable, appears in Section 5.

Replication in New York Wholesale Prices

New York provides an independent check on both the Pennsylvania-specificity and the monopsony and monopoly concerns. NYSLA wholesale list prices are posted monthly to a competitive market with many licensed wholesalers and retailers rather than to a single state buyer, and they sit one layer upstream of retail. The balanced Q1–Q9 wholesale panel contains 675 products (537 single malts and 138 blends), six times the Pennsylvania panel.

The wholesale data reproduce both forms of heterogeneity. The tercile contrasts show the same cheap-tier-loaded pattern, with mid- and expensive-tier mature increases 1.4 percentage points below the cheap tier (two-sided $p = 0.062$ and 0.033; Appendix Table 10). Because Whisky Advocate coverage is too thin on the wholesale panel, the continuous specification uses log pre-tariff wholesale price as the interacted variable. Figure 5 shows the per-quarter slope deepening monotonically to $\kappa_9 = -0.0081$ in Q1 2021 (SE 0.003, two-sided $p = 0.007$), with pre-trends individually small and jointly insignificant ($F = 0.87$, $p = 0.42$). Wholesale magnitudes are smaller than retail, as expected for one component of the retail price, but the cross-product pattern is the same. More expensive single malts show smaller proportional wholesale price increases.

Scope of the Reduced-Form Estimates

The estimates above identify the intensive-margin price response, the change in retail price for products that remain on PLCB shelves throughout the panel. The balanced-panel design omits products withdrawn or never reintroduced after the tariff, which likely concentrate at the higher end of the quality distribution. The welfare analysis later in Section 5 also holds the product set fixed at its pre-tariff composition, so it captures the intensive-

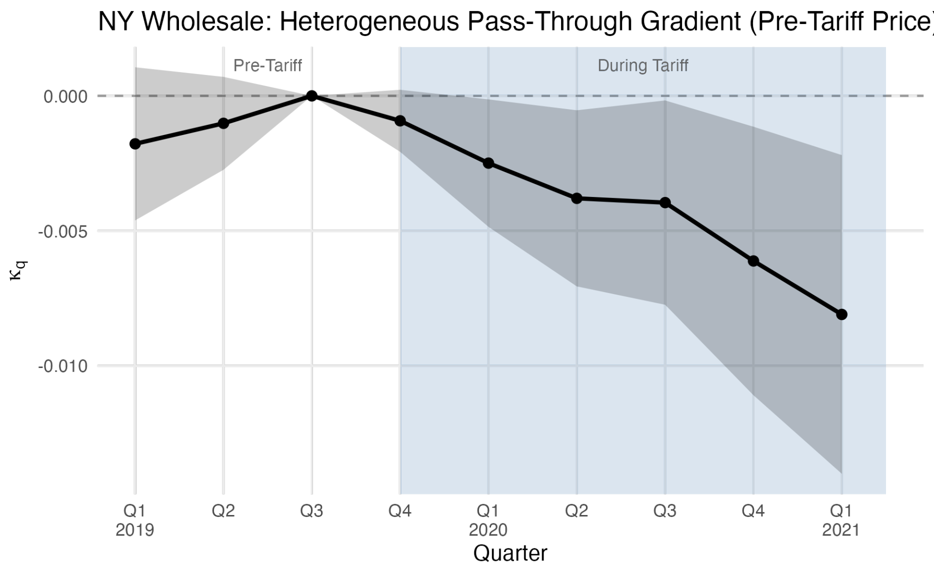

Figure 5: Per-Quarter SM Wholesale Pass-Through Slope (κ_q). Saturated event-study triple-DiD with log pre-tariff wholesale price as the continuous interacted variable. $\kappa_q < 0$ means more expensive single malts show smaller proportional wholesale price increases than cheaper ones in quarter q. Reference: Q3 2019. Pre-trends in κ_q are individually small and jointly insignificant ($F = 0.87$, $p = 0.42$).

margin price response and the resulting substitution, not the value of any variety the tariff removed.

Another potential concern regards the stable-unit-treatment-value assumption (SUTVA). Blended Scotch is the natural control group because it shares the supply chain with single malt and is exposed to the same common shocks (COVID, Brexit, distribution costs), but blended and single malt are also consumer substitutes. A tariff that raises single malt prices therefore raises blended Scotch demand via cross-elasticity, which in turn raises blended Scotch equilibrium prices.[7] The single-malt-versus-blended differential observed in the DiD is the net of the direct tariff response on single malt and any substitution-induced price response on blended Scotch. In principle the spillover could bias the differential either way. The cross-elasticity channel just described would compress it, while the offsetting multiproduct channel in the footnote above would widen it.

Directly, the control group barely moves. Blended Scotch average prices are nearly flat across the tariff window, falling 1.4% in levels from Q3 2019 to Q1 2021 (the blue series in Figure 1), and against their own pre-tariff baseline they show no statistically significant

[7] An offsetting channel involves multiproduct producers. Berry et al. (1999) note that a firm producing both treated and substitute products may lower the substitute's price to retain consumers substituting away from the cost-shocked product. The two channels operate in opposite directions, and the realized blended-Scotch response is the net of both.

deviation in any quarter or jointly (cluster-robust Wald $F = 1.21$, $p = 0.30$), with a clean pre-trend ($F = 0.48$, $p = 0.62$). This is suggestive evidence of a substitution-induced spillover too small to detect. Furthermore, the continuous specification is more robust to this concern than the pooled specification. The ϕ_q term in the same equation absorbs any quality-correlated movement in blended Scotch prices, so the single-malt-specific slope κ_q is identified from within-single-malt variation in quality after netting out the corresponding within-blended-Scotch quality variation. To the extent that substitution generates a uniform quality-correlated shift in blended Scotch prices, that shift would affect ϕ_q rather than κ_q, and the heterogeneity claim is identified more cleanly than the average pass-through specification.

4 A Structural Model of Whiskey Demand

The previous section established that the tariff's price response varies sharply across the quality ladder. Identifying the mechanism requires firm markups, which are unobserved and recoverable only from a demand system. Furthermore, quantifying the incidence of the tariffs (i.e. how those heterogeneous price changes map into welfare across income groups) requires consumer substitution and income-specific price sensitivities that the reduced form does not estimate. I therefore estimate a random coefficients (BLP) nested logit demand model[8] for Scotch and Irish whiskey in Pennsylvania, which recovers product-level markups and accommodates flexible substitution and consumer heterogeneity in income.

Model of Demand

In each pricing period, a consumer decides to purchase a bottle of Scotch or Irish whiskey (inside good), or opts for some other whiskey (outside good), of sizes 375ml, 750ml, 1000ml, or 1750ml. Consumer i's indirect utility of purchasing product j in pricing period (market) t is given as:

$$U_{ijt} = x_j\beta - \alpha_i p_{jt} + \xi_{jt} + \tilde{\varepsilon}_{ijt} \tag{3}$$

[8] Previous studies have used the random coefficients nested logit (RCNL) model in estimating demand for alcohol, such as Miller and Weinberg (2017), Miravete et al. (2018), and Conlon and Rao (forthcoming).

x_j is a vector of product characteristics, which includes bottle size and proof, and p_{jt} is the price of product j. The unobserved mean valuation for product j in time t is defined as

$$\xi_{jt} = \xi_{\tilde{j}} + \xi_t + \Delta\xi_{jt}, \tag{4}$$

where $\xi_{\tilde{j}}$ is a fixed effect for the brand of a bottle (e.g., Johnnie Walker) and ξ_t is a fixed effect for the pricing period. Including the pricing period fixed effects will account for market-wide trends; for example, the online alcohol retailer Drizly has reported that Scotch market shares increase in the holiday season as consumers look to buy more expensive bottles. Finally, $\Delta\xi_{jt}$ represents product-specific unobserved mean valuation that is not explained by brand and pricing period effects and will be the structural error that will form the moment conditions for estimation.

I allow consumer heterogeneity in the price coefficient, and model the bottle-size and proof coefficients in β as constant across consumers. Consumer i's price sensitivity α_i is log-normal,

$$\log \alpha_i = \alpha + \Pi D_i + \Sigma \nu_i, \quad \nu_i \sim N(0,1),$$

where D_i represents consumer i's income-bin indicators, Π the demographic interactions, and $\Sigma\nu_i$ the unobserved random-coefficient shifter. The log-normal form keeps α_i strictly positive, so the price coefficient $(-\alpha_i)$ in the indirect utility is negative for every consumer.

Individuals are split into three bins of household income: less than \$45,000, between \$45,000 and \$100,000, and greater than \$100,000. The lowest bin is the omitted reference, leaving two free demographic interaction parameters. The income classification draws on the bins available in the NielsenIQ Consumer Panel. The bins above and below \$100,000 capture the relevant variation in price sensitivity for this market without producing thin cells in any bin.

The $\tilde{\varepsilon}_{ijt}$ forms the stochastic term, and can be further decomposed into

$$\tilde{\varepsilon}_{ijt} = \zeta_{ijt} + (1 - \rho_g)\varepsilon_{ijt} \tag{5}$$

for a product j in nest g. ε_{ijt} is the independent and identically distributed extreme value error term and ζ_{ijt} is the nest-level component, drawn so that $\tilde{\varepsilon}_{ijt}$ follows the extreme value distribution. The nesting parameter ρ_g governs within-nest substitution for nest g, bounded between zero and one. Nests are defined as single malt Scotch (SM), blended Scotch (BL), Irish whiskey (IR), and the outside good. The Irish single malt and blended Irish products are combined into a single Irish nest, as treating them as separate nests yields an unstable estimate because the Irish single malt sample is small. I allow each nest to have its own ρ_g, so that ρ_{SM}, ρ_{BL}, and ρ_{IR} are estimated separately. Larger ρ_g values indicate stronger within-nest substitution. At $\rho_g = 1$ consumers substitute only to other products within nest g, and at $\rho_g = 0$ for all nests the model collapses to the mixed logit model.

I define the potential market as the market for whiskey and measure shares in liters of ethanol equivalent (an 80-proof 750 ml bottle contributes 300 ml of ethanol). The advantage of this definition is that the PLCB data fix the market size directly, without the auxiliary assumptions a broader "all spirits" market would require. As in any discrete-choice model, the outside option (here, all non-Scotch, non-Irish whiskey) carries a normalized mean utility and is not repriced, so the model abstracts from strategic price responses by, for instance, bourbon producers. This abstraction does not bind on the exercises below. The tariff moves only single malt Scotch prices, and substitution out of Scotch toward the rest of the whiskey market is still captured through the outside-good margin. No counterfactual in the paper perturbs the price of an outside product.

I decompose the indirect utility into a mean utility δ_{jt} that is common across consumers and an individual-specific deviation μ_{ijt},

$$U_{ijt} = \delta_{jt} + \mu_{ijt} + \tilde{\varepsilon}_{ijt} \tag{6}$$

where

$$\delta_{jt} = x_j\beta + \xi_{\tilde{j}} + \xi_t + \Delta\xi_{jt} \tag{7}$$

$$\mu_{ijt} = -\alpha_i\, p_{jt}. \tag{8}$$

The bottle-size and proof characteristics carry coefficients that are common across consumers and so enter the mean utility δ_{jt}; price is the only individual-specific coefficient and is therefore the sole term in μ_{ijt}. For a product j in nest g with nest-specific parameter ρ_g, the conditional probability that consumer i chooses product j in market t is then

$$s_{ijt} = \frac{\exp((\delta_{jt} + \mu_{ijt})/(1-\rho_g))}{\exp(I_{ig}/(1-\rho_g))} \frac{\exp(I_{ig})}{\exp(I_i)}$$

with the McFadden (1978) inclusive values defined as

$$I_{ig} = (1-\rho_g) \ln \sum_{j=1}^{J_g} \exp((\delta_{jt} + \mu_{ijt})/(1-\rho_g))$$

$$I_i = \ln \left(1 + \sum_{g=1}^{G} \exp(I_{ig})\right)$$

for product set J_g of each nest. The outside good is the degenerate nest normalized to $\delta_{0t} = 0$, which is the source of the 1 in I_i. Thus product j's shares in market t are then simulated as

$$\hat{s}_{jt} = \frac{1}{N_t} \sum_{i=1}^{N_t} \frac{\exp((\delta_{jt} + \mu_{ijt})/(1-\rho_g))}{\exp(I_{ig}/(1-\rho_g))} \frac{\exp(I_{ig})}{\exp(I_i)} \tag{9}$$

Model of Supply

The firms in the supply model are the upstream suppliers (e.g., Diageo, Pernod Ricard), competing Bertrand-Nash and pricing each product by the inverse-elasticity rule. The retail layer requires care. The PLCB is nominally the shelf-price setter, and Act 39 of 2016 permits it to price best-selling and limited-purchase items to maximize the return on their sale, authority the Board interprets as covering its entire listed portfolio (Section 2). In principle, then, the PLCB could add its own strategic markup on top of each supplier's wholesale price in response to demand conditions. If the PLCB were exercising that flexibility during the tariff window, observed retail price changes would mix supplier pricing decisions and the PLCB's own response to those decisions, and the Bertrand-Nash specification below would describe supplier wholesale prices but not the observed retail prices. The Board's documented conduct during the window, however, is not that of an active monopoly pricer. As Section 2 records, it initiated no retail price increases in 2020 or 2021 and cancelled the one adjustment round it had planned, and when pandemic-era shortages left some products scarce, it imposed purchase limits rather than raising prices, declining the scarcity rents a return-maximizing monopolist would have taken.[9] Supplier pricing decisions during this period are therefore made against final consumer demand, with the PLCB collapsing to a passive intermediary that applies fixed markups, and the Bertrand-Nash multi-product specification below models the relevant strategic interaction (competition among suppliers).

There are F firms, each producing a subset $\mathcal{J}_{ft}$ of products in each period t. Firm f's profit function in market t (ignoring t subscripts), assuming constant marginal costs, and normalizing the market size to one, is

$$\Pi_f = \sum_{j \in \mathcal{J}_f} (p_j - mc_j) s_j(\mathbf{p}) - FC_f \tag{10}$$

where $\mathcal{J}_f$ is the firm's product set, mc_j is product j's marginal costs, $s_j(\mathbf{p})$ is the market

[9] Pennsylvania Liquor Control Board, *Annual Report on Method and Rationale for Product Pricing*, April 1, 2022: "While flexible pricing afforded the PLCB the authority to increase the retail prices on these supply-constrained products to offset demand, the agency instead chose to implement purchase limits."

share of product j as a function of market prices $\mathbf{p}$, and finally FC_f is the fixed cost of firm f.

The ownership matrix Ω is defined as in Nevo (2001), where

$$\Omega_{jr} = \begin{cases} -\partial s_r / \partial p_j & \text{if } \exists f : \{r, j\} \subset \mathcal{J}_f \\ 0 & \text{otherwise} \end{cases} \tag{11}$$

I do not separately specify and estimate a model of marginal costs, but use the estimates from the demand model and the supply model specification to back out marginal costs. The first-order conditions give the markup equation

$$p - mc = \Omega^{-1} s(\mathbf{p}), \tag{12}$$

which I solve for each product's marginal cost. The corresponding Lerner index,

$$\mu_j = \frac{p_j - mc_j}{p_j}, \tag{13}$$

the markup as a fraction of price (bounded in $[0, 1]$), will be used in Section 5 in assessing the role of markups in determining pass-through.

Identification and Estimation

The arguments for identification and the estimation procedure are based on the procedure outlined in Berry (1994) and Berry et al. (1995). I use the NielsenIQ Consumer Panel Data to form micro moments that identify the demographic interactions with prices and the nesting parameters. I draw the income distribution from the NielsenIQ Consumer Panel for 2018-2021, assigning each panelist to one of the three income bins above (less than \$45K, \$45K-\$100K, greater than \$100K). I then construct three sets of micro moments. The first is the average purchase price of 750mL Scotch by income bin, conditional on a Scotch purchase (Petrin 2002), which disciplines the demographic interactions on price. Conditioning on 750mL eliminates the size composition confound (the cheapest 1750mL bottles are mostly blended). The second is the share of Scotch purchases that are single malt by income bin, which identifies how the within-Scotch quality split varies with income. The third is the within-nest repurchase rate for each of the three estimation nests, following Conlon and Rao (forthcoming): for households that buy at least two distinct products, what fraction of repeat purchases stay within the same nest? This third group disciplines the nest-specific ρ_g parameters directly through observed within-nest substitution patterns. The lowest income bin's coefficient is excluded for normalization.

The unobserved heterogeneity in price sensitivities is identified through the exogenous shifts in prices and the corresponding changes in shares, which the distribution of consumer heterogeneity must rationalize. To identify the exogenous changes in prices, I instrument price with three instruments: the first is an indicator variable for the tariff, which is credibly uncorrelated with the structural error term $\Delta\xi_{jt}$ due to the nature of the tariffs, but correlated with the price. The tariff serves as an ideal cost shock for "tracing out" the demand curve. For my second instrument, I match each product to the data of wholesale liquor prices in New York state and use the corresponding wholesale price in New York (the prices that suppliers charge to wholesalers in New York) as an instrument for retail prices in Pennsylvania. The identifying assumption behind this instrument is that wholesale prices in New York would reflect cost shocks that affect Pennsylvania as well, while being uncorrelated with any demand shocks in Pennsylvania. This is similar to the argument for "Hausman instruments," and as such, is susceptible to the same threats to identification as discussed in Nevo (2001). For example, if there is an unobserved cross-state demand shock that leads manufacturers to charge higher prices to both Pennsylvania and neighboring New York, retail prices in Pennsylvania and wholesale prices in New York would be correlated. However the inclusion of period fixed effects and brand fixed effects helps mitigate such concerns as identification now relies on wholesale prices in New York being uncorrelated with $\Delta\xi_{jt}$, the product-specific deviation from brand and period fixed effects. Unless the correlated demand shocks in both markets are brand-period specific, these instruments would satisfy the exclusion restriction. Finally, the third instrument is a dummy variable indicating whether a product is offered on sale or not. Most of the price variation is due to temporary price reductions and promotions, which are scheduled by the manufacturer or PLCB several months in advance, and thus, the identification holds as long as the sales are not correlated with $\Delta\xi_{jt}$, the product specific deviation from brand and period fixed effects. Using promotions as an instrument does risk attributing stock-piling behavior for elastic demand, but Seim and Waldfogel (2013) also study the Pennsylvania liquor market and find no correlation between consumers' distances to store and changes in sales.

The unobserved heterogeneity in the price coefficient (the random coefficient Σ) is identified by the exogenous changes in the product choice sets across markets and the correlation between shares and the characteristics of the remaining products. And beyond the second-choice micro moments, the nest-specific parameters ρ_g are also identified by the change in each nest's share of the total inside market as products enter and exit the choice set, with the number of products in each nest entering as an instrument for the corresponding ρ_g.

I also utilize differentiation instruments from Gandhi and Houde (2019). Namely, I utilize the local and interacted form of the differentiation IVs, using exogenously predicted prices, proof, and bottle categories. These IVs capture the crowdedness of the product space for each product, for example, "how many similarly priced own-firm and rival-firm bottles are there near mine?".

Defining the set of instruments as $Z = [z_1, \ldots, z_M]$, I define the population moment conditions as $E[Z'\omega(\theta)]$ where $\theta = [\beta, \Pi, \Sigma, \rho]$ are the parameters to be estimated and $\omega(\theta)$ is an error term defined as the unobserved mean valuation $\Delta\xi_{jt}$. Specifically, using Equation 7, I derive

$$\omega(\theta) = \delta_{jt} - x_j\beta - \xi_{\tilde{j}} - \xi_t$$

The δ_{jt} comes from matching the predicted shares to the market shares, $\hat{s}_{jt}(\delta_{jt}, \theta) = s_{jt}$, and can be inverted numerically with the dampening term from Grigolon and Verboven (2014). Then, the GMM estimate is

$$\hat{\theta} = \arg\min \ \big(Z'\omega(\theta)\big)' A^{-1} \big(Z'\omega(\theta)\big)$$

for the appropriate weighting matrix A. The micro-moment residuals constructed from the NielsenIQ targets are stacked with the demand-error moments $Z'\omega(\theta)$ and enter the same quadratic form.

I simulate consumers using 250 agents per income bin per market (750 total per market), with each bin's agents carrying the integration weight that reproduces the NielsenIQ target income proportions (30.5%, 34.2%, 35.3%) in every market. Each agent receives a single one-dimensional Halton draw for the unobserved component of the price coefficient. I estimate the parameters by one-step GMM with a robust weighting matrix and report robust standard errors. Estimation is implemented in PyBLP (Conlon and Gortmaker 2020), with the micro moments incorporated following the framework of Conlon and Gortmaker (2025).

Estimation Results

The results of the estimation are presented in Table 2, with the corresponding micro moment fits in Table 3. The mean price coefficient α is not separately identified from the constant term in the demographic projection (π_0) and is absorbed into the latter. The demographic interactions are in the expected direction. Relative to the omitted lowest-income bin, the middle-income coefficient is -0.33 and the high-income coefficient is -1.57, both implying lower price sensitivity at higher incomes. The middle-income coefficient is within one standard error of zero and is not statistically distinguishable from the low-income reference; the high-income coefficient is significant at $t \approx -3.4$. The implied median price sensitivity $|\alpha_i|$ ranges from 2.57 at the lowest income to 0.54 at the highest, a roughly five-fold ratio (the price coefficient itself, $-\alpha_i$, is correspondingly negative). The random coefficient $\sigma = 1.74$ generates substantial within-bin heterogeneity that overlaps across income bins.

The bottle-size coefficients show consumers preferring 1750 mL bottles to the 750 mL reference and disliking the 375 mL half-bottle. The 1000 mL and proof coefficients are small and statistically indistinguishable from zero. The nest-specific ρ parameters are around 0.22 for blended Scotch and 0.26 for single malt Scotch, indicating moderate within-nest correlation in unobserved tastes. The Irish nest's ρ sits at the lower boundary, reflecting the small Irish product set. The Irish nest contains few products, leaving little within-nest substitution variation to discipline ρ_{IR}, so it is weakly identified and is driven to its lower bound. Because Irish whiskey is only a peripheral substitution margin for the single malt Scotch tariff at the center of the paper, this imprecision does not materially affect the results of interest.

Table 3 compares the moment targets to the model fits. The within-nest cross-purchase moments fit nearly exactly. The price moments fit closely for the middle bin, over-predict the low-income bin's average price by \$4, and under-predict the high-income bin's by \$2. The SM share moment for the highest income bin under-fits by 12 percentage points (predicting 0.15 vs. observed 0.27). The model captures the qualitative ordering of all three groups but understates the strength of within-Scotch quality sorting at higher incomes. I return to this limitation in the welfare discussion in Section 5.

Figure 6 plots the implied distribution of the absolute price coefficient $|\alpha|$ for each income bin. Median price sensitivity is roughly $5\times$ higher for low-income consumers ($|\alpha| = 2.57$) than for high-income consumers ($|\alpha| = 0.54$), but the within-bin heterogeneity generated by the random coefficient is comparable in scale to the between-bin shift. Income classification therefore captures a meaningful share of cross-consumer variation in price sensitivity but not all of it. The substantial overlap is the demand-side reason why the welfare incidence in Section 5 is regressive but moderately so.

Figure 7 plots the implied diversion ratios from single malt Scotch when prices rise, averaged within SM quality tier. The substitution patterns differ across tiers in an economically meaningful way. When cheap single malts rise in price, displaced demand splits roughly equally between other single malts (34%), blended Scotch (31%), and Irish whiskey (31%), with a small fraction (3%) going to the outside good, other whiskeys.

	Mean Utility	Random Coeff	Demographic Interactions (Π)		
	(β)	(Σ)	<$45K	$45K - $100K	>$100K
Price	0.946 (0.433)	1.740 (0.091)	-	-0.334 (0.629)	-1.567 (0.464)
Proof	0.003 (0.003)				
375 ML	-4.05 (0.23)				
750 ML	-				
1000 ML	0.028 (0.107)				
1750 ML	3.24 (0.11)				
		Nesting Parameters (ρ)			
Single Malt Scotch			0.257 (0.011)		
Blended Scotch			0.216 (0.011)		
Irish whiskey			0.010 (0.012)		

Table 2: Demand Model Estimates. Robust standard errors in parentheses. Brand and market fixed effects absorbed. One-step GMM with 14 excluded instruments and 9 micro moments.

Moment	Observed	Fitted
Group 1: $E[\,p\,|\,\textit{income},\ 750mL,\ j>0\,]$ (*PA Scotch*)		
<$45K	$34.88	$39.05
$45K–$100K	$40.55	$40.57
>$100K	$50.44	$48.14
Group 2: $\Pr(\textit{SM Scotch}\,|\,\textit{income},\ j>0)$ (*national*)		
<$45K	0.116	0.108
$45K–$100K	0.159	0.115
>$100K	0.272	0.149
Group 3: $\Pr(\textit{same nest}\,|\,\textit{first choice in nest})$ (*national*)		
Single Malt Scotch	0.402	0.411
Blended Scotch	0.496	0.498
Irish whiskey	0.346	0.355

Table 3: Micro Moment Targets and Model Fits. All moments come from the NielsenIQ Consumer Panel and weight purchase occasions equally; sample scope for each group as labeled.

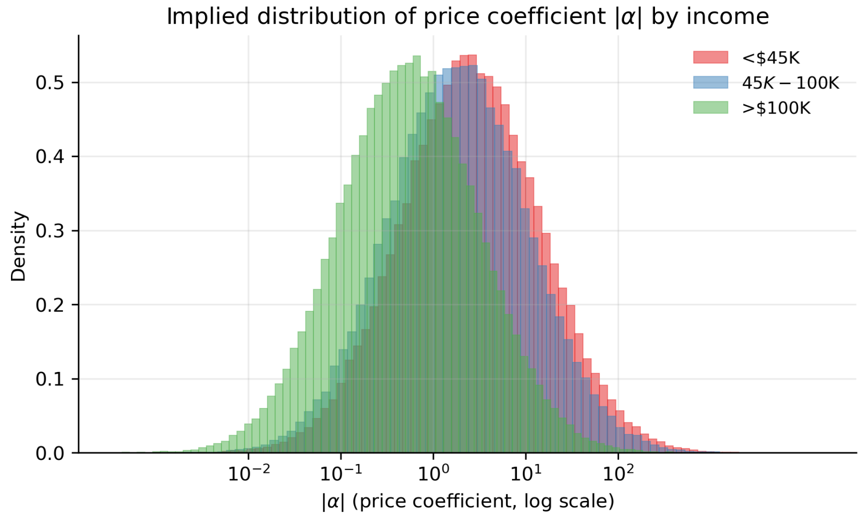

Figure 6: Implied distribution of the absolute price coefficient $|\alpha|$ by income bin, from the log-normal random coefficient. Medians: 2.57 (low income), 1.84 (middle), 0.54 (high).

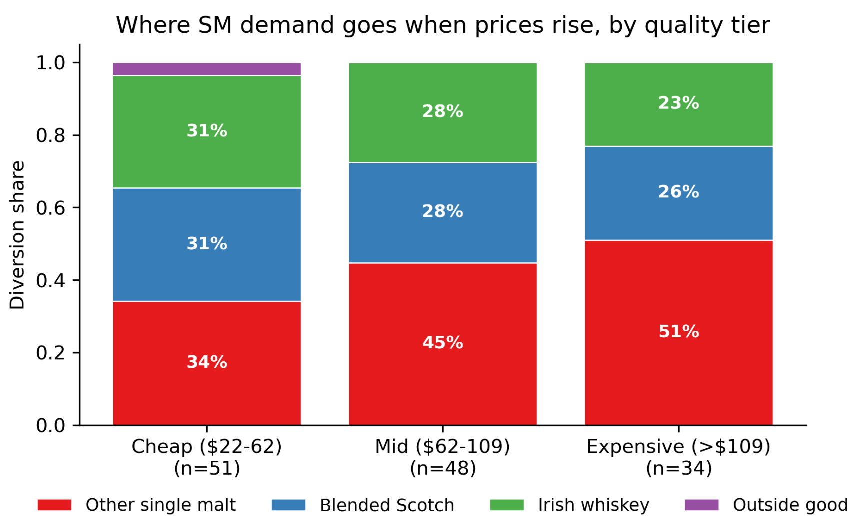

Figure 7: Diversion ratios from single malt Scotch when prices rise, by SM quality tier. Bars show the share of displaced demand going to other single malts, blended Scotch, Irish whiskey, or the outside good.

When expensive single malts rise in price, more than half of displaced demand stays within the SM nest (51%); the rest splits between blended Scotch (26%) and Irish whiskey (23%); and essentially nothing leaves Scotch and Irish. The 31% diversion from cheap SM to blended Scotch is the demand-side counterpart of the quality-downgrade channel that Section 5 measures on the extensive margin.

5 Markups and the Distribution of the Tariff Burden

Section 3 documented that the tariff passed through more strongly to cheap single malt Scotch than to expensive products. This section asks two related questions. First, does the heterogeneity operate through firm markups? Price and WA rating both correlate with markup but each conflates several dimensions (cost, brand premium, vertical quality), and the structural demand model from Section 4 gives a direct measure. Second, what does the heterogeneity imply for welfare incidence by income? A uniform tariff on a category whose pass-through varies systematically across the product ladder will fall hardest on consumers who concentrate at the high-pass-through, low-markup end. I estimate the size of that channel using the BLP demand model.

5.1 Markup as the Mechanism

I re-run the saturated event-study triple-DiD from Section 3 with the BLP-estimated Lerner markup (Equation 13) substituted for the observable proxies:

$$\begin{aligned} \ln(p_{it}) &= \gamma_i + \lambda_t \\ &\quad + \sum_{q \neq 3} \theta_q \, \mathbf{1}\{t = q\} \, \mathrm{SM}_i \\ &\quad + \sum_{q \neq 3} \phi_q \, \mathbf{1}\{t = q\} \, \tilde{\mu}_i \\ &\quad + \sum_{q \neq 3} \kappa_q \, \mathbf{1}\{t = q\} \, \mathrm{SM}_i \, \tilde{\mu}_i + \varepsilon_{it} \end{aligned} \tag{14}$$

where $\widetilde{\mu}_i = \mu_i - \overline{\mu}^{SM}$ is the SM-mean-centered Lerner markup recovered from the demand model, averaged over the pre-tariff markets (Q1–Q3 2019). Quarter $q = 3$ (Q3 2019) is the reference. The main coefficient of interest is κ_9, the markup slope of the SM price response at the mature peak of the tariff window.

Table 4 reports the per-quarter κ_q. The coefficient is small and insignificant during the first two post-tariff quarters, when PLCB supplier contracts had not yet renegotiated, and grows in magnitude through the mature window. At Q1 2021, $\kappa_9 = -0.465$ (SE 0.195, two-sided $p = 0.019$). Figure 8a plots the per-quarter estimates with 95% confidence intervals.

Combined with $\theta_9 = 0.099$, the implied SM retail price increase at Q1 2021 ranges from +20% for products in the 10th percentile of the SM markup distribution to +2% in the 90th percentile. Figure 8b plots the implied per-product pass-through curve. The model-estimated markups reproduce the heterogeneous pattern observed in the reduced form, and its agreement with the WA-rating slope in Section 3 and the price-based slopes in Appendix A.3 and Figure 5 confirms the pattern is not driven by how any single markup measure is constructed.

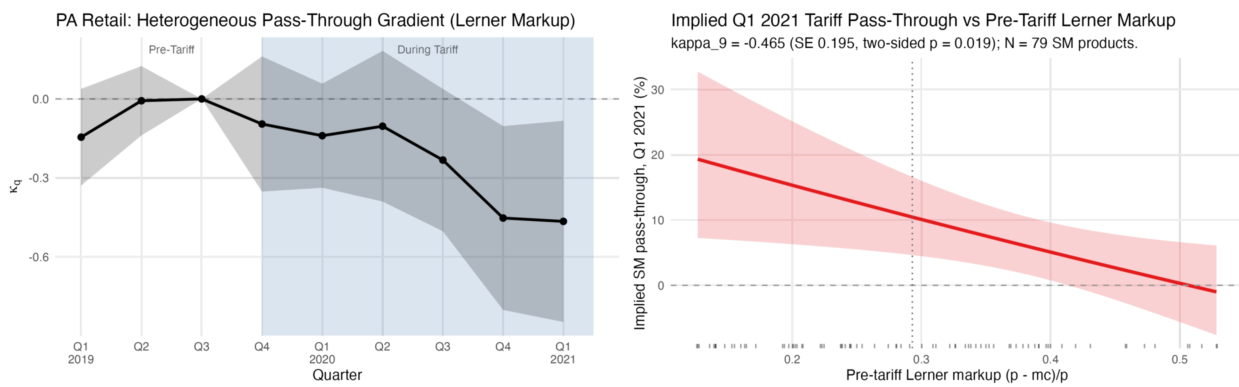

(a) Per-quarter κ_q

(b) Implied Q1 2021 pass-through

Figure 8: Markup interaction: per-quarter κ_q and implied per-product pass-through

A negative κ is consistent with two distinct supply-side mechanisms. One, incomplete pass-through with imperfect competition means high-markup firms have the margin buffer to absorb cost shocks while low-markup firms transmit them. Two, under complete pass-through in levels (Sangani 2026), an ad-valorem tariff raises each product's price by a fixed dollar amount roughly proportional to its cost, which is a larger *percent* increase for low-markup products because they carry a higher cost share. Both accounts predict $\kappa_9 < 0$.[10] I do not attempt to identify precisely which mechanism is driving the results, and the welfare analysis below does not depend on which one operates.

[10] A companion regression replacing the log price with the dollar price change yields a markup slope in levels of $+15.2$ (SE 13.6, two-sided $p = 0.27$). Dollar price increases are statistically flat in markup, if anything rising. Active margin absorption predicts a negative slope in dollars, while complete pass-through in levels predicts a weakly positive one, so the data lean toward the cost-share account without sharply distinguishing the two. None of the welfare results depend on the distinction.

Dependent Variable:		$\log p_{it}$
SM-by-markup interaction κ_q:		
κ_1	(Q1 2019)	-0.1453
		(0.0937)
κ_2	(Q2 2019)	-0.0070
		(0.0672)
κ_4	(Q4 2019)	-0.0954
		(0.1310)
κ_5	(Q1 2020)	-0.1398
		(0.1010)
κ_6	(Q2 2020)	-0.1037
		(0.1461)
κ_7	(Q3 2020)	-0.2323*
		(0.1380)
κ_8	(Q4 2020)	-0.4526**
		(0.1782)
κ_9	(Q1 2021)	-0.4653**
		(0.1951)
Fixed effects:	product, quarter	Yes
Observations		945
R^2		0.99

Cluster-robust standard errors (by product) in parentheses.

Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$ (two-sided).

Table 4: Saturated event-study triple-DiD in the BLP-derived Lerner markup. Each row reports the per-quarter SM-by-markup interaction coefficient κ_q, with Q3 2019 as the omitted reference. Sample: balanced Q1–Q9 panel of $n = 105$ Scotch products (79 SM, 26 blended).

5.2 A First-Order Incidence Framework

With the price-change schedule in hand, it is useful to state generally when a uniform tariff is regressive within the tariffed category. Indexing products by j and income groups by g, let w_{jg} denote group g's expenditure share on product j within the category at pre-tariff prices, with $\sum_j w_{jg} = 1$. Let $\Delta \ln p_j$ denote the tariff-induced price change. To a first order approximation, the tariff's cost to group g, expressed as a share of the group's category spending, is

$$R_g = \sum_j w_{jg} \, \Delta \ln p_j, \tag{15}$$

the expenditure-weighted average price increase. This is the standard envelope approximation. Behavioral responses affect welfare only at second order, so Equation 15 holds in any demand system and requires only observed baskets and price changes.

The distributional content of the tariff is the gap in R_g across groups. Because weights sum to one within each group, any common component of the price increases nets out of the comparison,

$$R_\ell - R_h = \sum_j (w_{j\ell} - w_{jh}) \, \Delta \ln p_j, \tag{16}$$

where ℓ and h index the low- and high-income groups. Two implications follow. First, under uniform pass-through ($\Delta \ln p_j$ constant across products), the right-hand side is zero regardless of how differently the groups shop, so a uniformly passed-through tariff cannot be regressive within the category. Second, heterogeneous pass-through creates regressivity exactly when price increases covary with the products the low-income group favors.

Proposition 1. *Suppose price increases follow the linear schedule $\Delta \ln p_j = \bar{\theta} + \bar{\kappa}(\mu_j - \bar{\mu})$, where μ_j is product j's markup. Then, to first order,*

$$R_\ell - R_h = \bar{\kappa} \left(\bar{\mu}_\ell - \bar{\mu}_h \right), \qquad \bar{\mu}_g \equiv \sum_j w_{jg} \, \mu_j, \tag{17}$$

where $\bar{\mu}_g$ is group g's expenditure-weighted basket markup. The tariff is regressive within the category if and only if $\bar{\kappa}$ and $\bar{\mu}_\ell - \bar{\mu}_h$ share the same sign, which here means pass-through declining in markups combined with a low-income consumption basket tilted toward low-markup products.

The proposition reduces within-category incidence to two sufficient statistics. One, the slope of pass-through in markups, and two, the income gap in the markup content of consumption baskets. Neither is specific to this setting — pass-through of cost shocks declines with markups across a broad range of industries (Sangani 2026), and richer households systematically source their consumption from higher-markup firms and premium varieties (Faber and Fally 2022, Jaravel 2019). Any tariffed category sharing these two features inherits the regressivity, and Equation 17 prices it with a pass-through slope and basket data alone. Appendix B derives Equation 15, proves Proposition 1, covers the discrete-choice case, and reports the income-level extension.

Implementing Equation 15 is a two-step calculation. Each single malt receives the price increase its markup implies under the mature-window schedule estimated above ($\bar{\theta} = 0.080$, $\bar{\kappa} = -0.313$, averaged over Q2 2020–Q1 2021), and each income group averages those increases over its own basket. The result is an effective tariff rate on the group's single-malt spending. Of each dollar the low-income group spends on single malt, the tariff claims 9.15 cents, against 8.24 cents for the high-income group, a 0.90 percentage point gap.[11] The uniform counterfactual shows where the gap comes from. Repeating the calculation with every product marked up by the average ($\Delta \ln p_j = \bar{\theta}$ for all j) yields an effective rate of 8.29% for every income group and a gap of exactly zero, although the groups' baskets differ throughout. Shopping patterns alone cannot generate within-category regressivity; the entire 0.90 point gap arises from heterogeneous pass-through interacting with who buys what.

The proposition also works as a formula. In place of the product-by-product calculation above, Equation 17 needs only two numbers, the pass-through slope $\bar{\kappa}$ and the gap in basket markups. The latter is measured directly. The expenditure-weighted basket markup rises monotonically in income, 0.269 for the low-income group against 0.296 for the high-income group, so the low-income basket carries 2.7 Lerner points less markup. Multiplying this gap by $\bar{\kappa}$ predicts a 0.84 percentage point effective-rate difference, against the 0.90 computed from the full baskets. The first-order identity holds to within a tenth of a percentage point, and the small residual reflects the second-order terms the approximation drops.

5.3 Welfare Incidence: A Two-by-Two Decomposition

The previous framework analyzes the welfare loss of the tariff at fixed baskets, as it uses just a first order approximation. However, substitution away from the hardest-hit products attenuates the loss at second order. This subsection repeats the exercise above using the demand model to let consumers substitute toward less-affected single malts, into blended Scotch or Irish whiskey, or out of the category entirely.

[11] The proposition requires only observed group-level baskets. Product-by-income expenditure weights for Pennsylvania are too thin to measure directly in the NielsenIQ panel, so the implementation takes w_{jg} from the choice probabilities of the estimated demand model, which are disciplined by the panel's income-conditional purchase moments.

I use the markup-κ_q schedule to construct counterfactual prices in the pre-tariff baseline market (period 9, Q3 2019). For each SM product j the predicted log price change under the heterogeneous tariff is

$$\Delta \ln p_j = \overline{\theta}_{Q6\text{–}Q9} + \overline{\kappa}_{Q6\text{–}Q9} \cdot (\mu_j - \overline{\mu}^{SM}) \tag{18}$$

with the coefficients averaged over the mature window (Q2 2020 through Q1 2021). The mature-window average is used because the actual tariff was in place from Q4 2019 through Q1 2021. The mature window approximates the steady-state pass-through that prevailed once supplier contracts had renegotiated. Equation 14 is fit on the balanced panel of 79 SM products plus 26 blended Scotch controls (105 in total). The welfare counterfactual concerns single malts only. The period-9 baseline market contains 133 SM products, of which 79 are in the balanced panel and 54 entered late or had data gaps that excluded them. I apply the fitted schedule to all 133; restricting to the 79 panel SM produces qualitatively identical results.

The relevant welfare object is compensating variation: the dollar amount that would compensate the consumer for the price change. The inclusive-value identity for nested logit with an outside good gives a closed form,

$$CV_i = \frac{1}{|\alpha_i|} \left(\log \frac{1}{P_{0i}^{\text{base}}} - \log \frac{1}{P_{0i}^{\text{cf}}} \right), \tag{19}$$

where P_{0i} is agent i's probability of choosing the outside good (so the inclusive value is $I_i = -\log P_{0i}$), evaluated at the pre-tariff baseline prices (base) and at the counterfactual prices with the estimated tariff increases imposed (cf). The BLP demand model integrates over all of agent i's substitution responses to the price change implicitly through the choice probabilities. The relative welfare metric is CV_i/CS_i^{base}[12], the loss in consumer surplus from the inside-good market expressed as a fraction of pre-tariff Scotch and Irish CS.[13]

The regressivity in the welfare metric is the joint product of a demand-side heterogeneity (heterogeneity in α across income groups) and a supply-side one (heterogeneity in Δp across products). To isolate each contribution, I run the welfare exercise across a two-by-two comparison of α regimes and Δp schedules:

(1) Homogeneous α + uniform Δp at $\exp(\overline{\theta}) - 1 = 8.3\%$.

[12] Dollar compensating variation deflates each agent's utility loss by the marginal utility of money $1/|\alpha_i|$, which differs mechanically across income groups, so it is not a valid cross-income comparison. Baseline consumer surplus carries the same $1/|\alpha_i|$ factor, so the ratio CV_i/CS_i^{base} cancels it and the comparison is scale-free.

[13] The CS object is the welfare from the inside-good (Scotch + Irish whiskey) market against the outside good (non-Scotch spirits), not total welfare across all consumption.

(2) Homogeneous α + heterogeneous Δp from the markup-κ_q schedule.

(3) Heterogeneous α + uniform Δp.

(4) Heterogeneous α + heterogeneous Δp, the actual world.

"Homogeneous α" replaces each agent's income demographics with the population shares, so α no longer varies systematically by income. The unobserved draws $\sigma\nu_i$ are kept (agents still differ in price sensitivity) and are held identical across the four cells, so the cell-to-cell differences isolate the income and price channels rather than picking up differences in the draws. Cells (1)–(3) correspond to approaches that build in one form of heterogeneity but not the other. Cell (2), heterogeneous pass-through with a representative consumer, is the focus of the aggregate trade-war pass-through literature (e.g., Cavallo et al. 2021, Amiti et al. 2019b); cell (3), income-heterogeneous demand facing a uniform price change, is what a standard BLP welfare exercise delivers; cell (1) is the no-heterogeneity benchmark. Cell (4) combines both.

Scenario	<\$45K	\$45K–\$100K	>\$100K	Spread
(1) Hom α + Uniform	0.364%	0.364%	0.364%	0.000 pp
(2) Hom α + Hetero	0.369%	0.369%	0.369%	0.000 pp
(3) Het α + Uniform	0.417%	0.391%	0.289%	0.128 pp
(4) Het α + Hetero	**0.436%**	0.403%	**0.277%**	**0.159 pp**

Table 5: Loss in pre-tariff Scotch and Irish consumer surplus from the heterogeneous tariff, by income bin. Each cell crosses an α regime with a Δp schedule. Cell (4) is the actual-world counterfactual. The spread is the low-income minus high-income difference. CV computed via Equation 19; % CS loss reported.

The four cells separate the tariff's regressivity into its demand-side source (income heterogeneity in α) and its supply-side source (heterogeneous pass-through), and show that the actual tariff is more regressive than either source alone reveals. The Spread column of Table 5 traces this regressivity across the four cells. Pass-through heterogeneity by itself does nothing. With homogeneous α (cells (1)–(2)), every income bin has the same price sensitivity, so the percentage loss is uniform whatever the Δp schedule. Income heterogeneity by itself opens a 0.13-point gap (cell (3)). Even under a uniform tariff, low-income agents lose a larger *share* of their surplus, because their substitution options are limited (cheap single malt stays tariffed) and their baseline Scotch-and-Irish surplus is small (\$12.60 per period against \$90.10 for high-income). The two together (cell (4)) widen the gap to 0.16 points: the extra 0.03 is the interaction, heterogeneous pass-through concentrating the largest increases on the cheap products low-income consumers buy and amplifying the income-driven regressivity by about a quarter (24%). A uniform-tariff welfare exercise (cell (3)) thus understates the regressivity of the actual, heterogeneous tariff (cell (4)). Measuring that gap is the contribution of the decomposition.

These levels are small in absolute terms. The tariff falls on a narrow product and the loss is measured against the full Scotch-and-Irish surplus, so the magnitude reflects the size of the tariffed category rather than the importance of the mechanism. The object of interest is the distributional structure, i.e., low-income consumers lose about 1.6 times the proportional surplus of high-income consumers, and heterogeneous pass-through amplifies that gap by a quarter. While these magnitudes are specific to Scotch and to Pennsylvania, pass-through declines with markups across a broad range of industries, and any tariffed category with that pattern faces the same incidence logic.

5.4 Extensive Margin and Quality Downgrade

The percentage-CS metric of Section 5.3 integrates every substitution response into a single welfare number. On the extensive margin (the change in single-malt purchase probability), the regressivity is far more pronounced. Of the lost SM purchases across all income groups, roughly 95% substitute to non-SM inside goods, split about equally between blended Scotch and Irish whiskey (Figure 7), and only 5% leave Scotch and Irish whiskey for the outside good, other whiskeys. The behavioral response to the tariff is overwhelmingly a downgrade within Scotch and Irish whiskey rather than a shift toward bourbon and other whiskeys.

Table 6 reports the same two-by-two decomposition for the extensive margin. The structure mirrors the welfare decomposition above. Pass-through heterogeneity on its own ((2)−(1)) again gives zero spread. Under homogeneous α all agents respond identically. Income heterogeneity on its own ((3)−(1)) produces a 5.7 percentage point spread, larger than on the welfare margin because the response is purely behavioral and is not normalized by baseline CS. The interaction adds another 2.2 percentage points, for a total cell (4) spread of 7.9 percentage points — a 38% amplification of the standard income-driven response.

The extensive-margin spread is large in absolute terms relative to the welfare spread for two reasons. First, the welfare metric divides the compensating variation by total inside-good consumer surplus (Scotch and Irish), so the single-malt tariff registers as a small fraction of a broad base. The extensive-margin metric is the raw percentage change in

Scenario	<\$45K	\$45K–\$100K	>\$100K	Spread
(1) Hom α + Uniform	-24.46%	-24.46%	-24.46%	0.00 pp
(2) Hom α + Hetero	-29.43%	-29.43%	-29.43%	0.00 pp
(3) Het α + Uniform	-26.69%	-25.66%	-21.00%	-5.69 pp
(4) Het α + Hetero	**-32.53%**	-31.10%	**-24.66%**	**-7.88 pp**

Table 6: Change in $\Pr(SM)$ per period by income bin, under each cell of the two-by-two decomposition. Counterfactual probabilities computed from the BLP demand model at the corresponding cell's prices and α regime.

single-malt purchase probability, undiluted by that denominator. Second, the metric isolates a behavioral response on a small base. Low-income consumers buy SM in only 1.42% of periods at baseline, so a 33% reduction is a large percentage change off a small starting probability. The welfare cost of the lost purchases is captured in the CV-based metric in Section 5.3. The extensive-margin metric records the size of the behavioral response without converting it to dollars.

The within-Scotch quality downgrade is the substitution channel that the BLP model attributes most of the welfare loss to. Whether the downgrade carries first-order welfare cost depends on the marginal rate of substitution between single malt and blended Scotch in consumer preferences. The BLP demand model assigns positive utility to product-specific differentiation that is partially preserved when the consumer downgrades. The welfare cost is smaller than in a counterfactual where substitution flowed to the outside good, but is not zero, and the CV-based metric in Section 5.3 integrates over this substitution.

6 Conclusion

This paper documents that the welfare incidence of a uniform tariff depends on how pass-through varies with firm markups. Using product-level data from Pennsylvania during the United States' 25% tariff on single malt Scotch, I find that pass-through declines with pre-tariff Lerner markups: implied retail price increases run from 20% at the 10th percentile of the markup distribution to 2% at the 90th, and in the coarser price-tercile view cheap products rise 14% while mid- and expensive-tier prices rise only 4–5%. The decline is estimated using the BLP-derived markup distribution and reproduced by independent specifications using pre-tariff retail price and Whisky Advocate ratings as observable proxies.

The heterogeneous pass-through translates into regressive welfare incidence on two margins. The percentage loss in pre-tariff Scotch and Irish consumer surplus is 0.44% for low-income consumers, against 0.28% for high-income, and the per-period probability of purchasing single malt Scotch falls 33% for low-income consumers, against 25% for high-income, with substitution flowing nearly entirely to lower-priced whiskey (blended Scotch and Irish), a downgrade along the quality ladder rather than a shift toward other whiskeys. A two-by-two decomposition that crosses the α regime with the Δp schedule shows that conditional on income heterogeneity in price sensitivity already producing regressive incidence, heterogeneous pass-through amplifies the regressivity by approximately 24%; pass-through heterogeneity alone, without income heterogeneity, generates no welfare regressivity. The amplification operates through the covariance of consumer income with supply-side pass-through across products. The welfare burden concentrates at the bottom of the income distribution.

Several limitations qualify the analysis. First, the welfare exercise is partial equilibrium. The demand model holds fixed both the product set and the prices of non-tariffed goods. Cross-elasticity from single malt to blended Scotch may attenuate the single-malt-versus-blended differential toward zero, so the reported pass-through magnitudes are plausibly lower bounds on the partial-equilibrium response, though the blended-Scotch price path is close to flat and the attenuation is small either way. Second, the BLP model under-predicts high-income consumption of single malt (15% predicted versus 27% in the NielsenIQ Consumer Panel). To the extent high-income single-malt demand concentrates on expensive products more than the model captures, the welfare regressivity reported here is likely conservative. Third, the Lerner markups in Section 5 are recovered from the same demand system used for the welfare analysis, with the tariff serving as a price instrument in that estimation and under a maintained conduct assumption of Bertrand-Nash supplier competition through a passive retail layer, so the markup result is best read as an internally consistent structural description rather than independent evidence of a supplier-markup channel. The independent weight is carried by the observable-proxy estimates of Section 3, pre-tariff price and Whisky Advocate rating, which require neither the demand model nor the tariff instrument. The markup specification shows that the same pattern survives translation into the structural object. Finally, inference rests on a balanced panel of 105 products with standard errors clustered one-way by product, and the welfare results are point estimates that do not propagate estimation uncertainty.

These findings nonetheless carry a direct policy implication. Because pass-through declines with markup, a uniform tariff concentrates its burden on the cheap, low-markup products that lower-income consumers disproportionately buy. A quality-graduated schedule (lower rates on cheap, low-markup varieties and higher rates on expensive, high-markup ones) would raise the same revenue with less regressive incidence, and could in principle be implemented at existing tariff-line resolution using observable price tiers as a proxy for markup. Characterizing the welfare-maximizing schedule, however, requires extrapolating pass-through beyond the rate range observed here together with a model of firm pricing under non-uniform tariffs, which I leave to future work.

Appendix

A Additional Results and Robustness

A.1 Full Coefficient Tables for §3 Event Studies

Tables 7 and 8 report the per-quarter coefficients underlying the pooled and tercile event-study figures in §3 (Figures 2 and 3).

	Dependent variable:	
	log(price)	
Q1 2019 (T–2)	-0.012	(0.013)
Q2 2019 (T–1)	-0.008	(0.010)
Q3 2019 (T 0, ref)	—	
Q4 2019 (T+1)	0.012	(0.015)
Q1 2020 (T+2)	0.015	(0.011)
Q2 2020 (T+3)	0.060***	(0.019)
Q3 2020 (T+4)	0.077***	(0.020)
Q4 2020 (T+5)	0.091***	(0.027)
Q1 2021 (T+6)	0.108***	(0.032)
Observations	945	
R^2	0.992	
	*p<0.1; **p<0.05; ***p<0.01	

Table 7: Pooled Pass-Through of the Single Malt Tariff. Single malt $\times$ quarter interactions from Equation 1; product and quarter fixed effects absorbed. Standard errors clustered by product in parentheses. Q3 2019 is the reference period.

A.2 Tercile Contrasts and Anticipation Check

To formalize the cross-tercile comparison in §3, Equation 20 estimates a fully interacted DiD specification, which pools across the mature treatment window (Q6–Q9) and uses the cheap tercile as the omitted reference:

	Dependent variable: log(price)		
	Cheap	Mid	Expensive
Q1 2019 (T–2)	-0.002 (0.021)	-0.025 (0.016)	-0.019 (0.028)
Q2 2019 (T–1)	-0.008 (0.015)	0.003 (0.011)	-0.008 (0.021)
Q3 2019 (T 0, ref)	—	—	—
Q4 2019 (T+1)	0.030 (0.030)	-0.010 (0.013)	0.008 (0.016)
Q1 2020 (T+2)	0.032 (0.026)	0.007 (0.008)	0.002 (0.004)
Q2 2020 (T+3)	0.083** (0.035)	0.039 (0.024)	0.041** (0.015)
Q3 2020 (T+4)	0.113*** (0.035)	0.053*** (0.015)	0.035*** (0.012)
Q4 2020 (T+5)	0.154*** (0.044)	0.024 (0.024)	0.026 (0.017)
Q1 2021 (T+6)	0.171*** (0.047)	0.029 (0.025)	0.032 (0.021)
Observations	405	279	261
R^2	0.973	0.917	0.987

*p<0.1; **p<0.05; ***p<0.01

Table 8: Single Malt Scotch Price Changes Relative to Blended Scotch, by Price Tercile. Three separate event-study estimations of Equation 1, one per pre-tariff price tercile (Spec A: SM-defined dollar cutpoints applied to both treated and control products). Product and quarter fixed effects absorbed. Standard errors clustered by product in parentheses. Q3 2019 is the reference period.

$$
\begin{aligned}
\ln(p_{it}) = \gamma_i &+ \lambda_t \\
&+ \beta_M \, \text{Post}_t \, \text{Mid}_i + \beta_E \, \text{Post}_t \, \text{Exp}_i \\
&+ \delta_0 \, \text{Post}_t \, \text{SM}_i \\
&+ \delta_M \, \text{Post}_t \, \text{SM}_i \, \text{Mid}_i + \delta_E \, \text{Post}_t \, \text{SM}_i \, \text{Exp}_i + \varepsilon_{it}
\end{aligned}
\tag{20}
$$

where $\text{Post}_t = 1$ if quarter $t \in \{6,7,8,9\}$ and 0 if $t \in \{1,2,3\}$ (sample restricted to these seven quarters); SM_i is the single malt indicator; and Mid_i and Exp_i are the mid- and expensive-tercile indicators (cheap is omitted as reference). Q4 and Q5 are dropped from this regression to account for the pricing contract lag in the early-tariff quarters. The coefficient δ_0 measures the cheap-tier price increases directly. The coefficients δ_M and δ_E measure how the mid- and expensive-tier mature price increases differ from the cheap-tier baseline. Standard errors are clustered by product.

	Dependent variable:
	log(price)
During (Q6–Q9) × Mid tercile	0.055* (0.032)
During (Q6–Q9) × Expensive tercile	0.043 (0.033)
During (Q6–Q9) × SM (Cheap baseline)	0.133*** (0.037)
During (Q6–Q9) × SM × Mid tercile	-0.090** (0.041)
During (Q6–Q9) × SM × Expensive tercile	-0.091** (0.041)
Observations	735
R^2	0.992

*p<0.1; **p<0.05; ***p<0.01 (two-sided)

Table 9: Fully Interacted Tercile DiD: Q1–Q3 (pre) versus Q6–Q9 (mature). Cheap tercile is the omitted reference. Triple interactions measure pass-through differences relative to the cheap-tier baseline. Tier samples: 27/26/26 single malts and 18/5/3 blended controls in the cheap/mid/expensive tiers.

Table 9 reports the results. The cheap-tier mature price increase is +14.3% (two-sided

$p < 0.001$). The Mid-versus-Cheap differential is -9.8 percentage points (two-sided $p = 0.030$), and the Expensive-versus-Cheap differential is -9.9 percentage points (two-sided $p = 0.031$). Both differentials reject the null of equal pass-through across terciles, with smaller price increases at the higher quality tiers. The Mid-versus-Expensive comparison, by contrast, is essentially zero. The heterogeneity is best summarized as a cheap-versus-other-tiers gap rather than a smooth decline across the three tiers.

Anticipation Check. One concern with the difference-in-differences design is that suppliers may have anticipated the tariff and adjusted single malt prices upward before October 2019, which would raise the Q3 2019 baseline and attenuate the estimated post-tariff coefficients. I check for this by regressing log retail price on the interaction of single malt status with quarter over the pre-tariff window (Q1–Q3 2019), with product and quarter fixed effects and standard errors clustered by product. The estimated single-malt-versus-blended linear drift is +0.006 log points per quarter (SE 0.006, $p = 0.34$), indistinguishable from zero.

A.3 Continuous Price Gradient (Pennsylvania Retail)

Substituting log pre-tariff retail price, centered on the single malt mean ($\tilde{p}_i$), for the WA rating in Equation 2 gives a price-based version of the §3 specification on the full 105-product balanced panel. Pre-tariff price is a composite of cost, quality, and market power, and the tercile evidence indicates a threshold structure (the cheap tier behaves distinctively while the mid and expensive tiers look alike), so a linear term in price is a coarse summary. It is reported here for completeness and for comparability with the New York wholesale slope below, which uses the same interacted variable.

Figure 9 plots the per-quarter κ_q estimates. The series stays near zero through Q1 2020 and deepens monotonically through the rest of the sample, peaking at $\kappa_9 = -0.098$ in Q1 2021 (SE 0.058, two-sided $p = 0.094$). The corresponding pre-trend test is borderline at essentially the same threshold ($F = 2.36$, $p = 0.095$ for the cluster-robust Wald test of $\kappa_1 = \kappa_2 = 0$). Both are marginal, consistent with the threshold structure above, and the direction matches the WA-rating slope in §3.

A.4 New York Supplier Prices

The NYSLA wholesale data has two functions in the paper. The first is a check on whether the PLCB's monopsony position is what generates the heterogeneous pass-through in §3. The PLCB is the only licensed buyer and retailer of distilled spirits in Pennsylvania, so

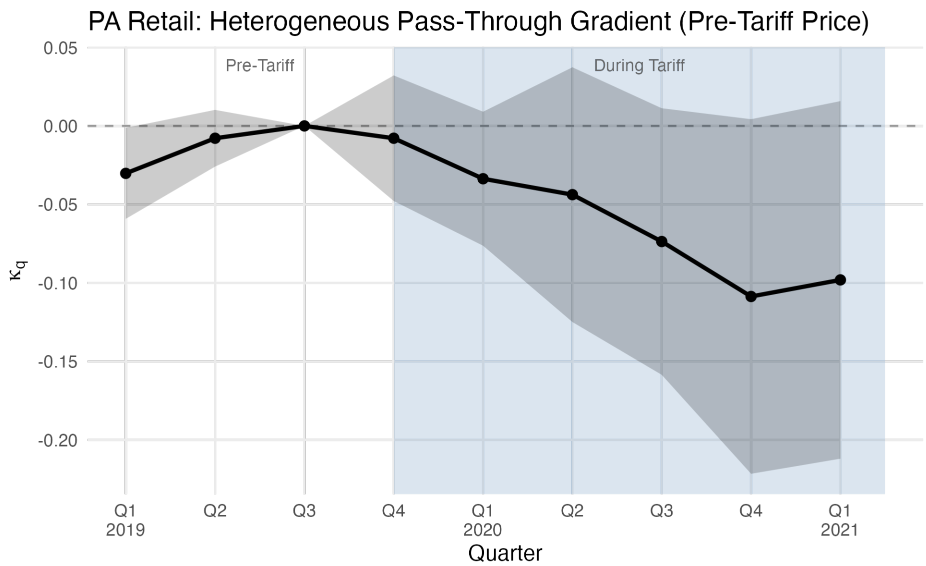

Figure 9: Per-Quarter SM-Blended Pass-Through Slope (κ_q) in Log Pre-Tariff Retail Price. Shaded ribbon is the 95% two-sided confidence interval.

every supplier faces it as a single buyer. A monopsony might absorb cost shocks asymmetrically across product tiers for reasons that do not generalize to competitive wholesale markets. New York instead licenses many wholesalers and retailers, and suppliers there face standard competitive bargaining. If PLCB monopsony were the explanation, the same heterogeneity by quality should not appear in NYSLA wholesale prices.

The second function is to look at supplier prices without PLCB markup decisions in the way. NYSLA wholesale list prices are public, uniform across wholesalers, and updated monthly under the post-and-hold system discussed in Conlon and Rao (forthcoming). They have fewer contractual lags and fewer nominal rigidities than the PLCB retail prices used in the main text.

I aggregate the monthly NYSLA list prices to quarterly observations so the panel matches the PA retail window (Q1 2019 – Q1 2021). I classify products as single malt or blended Scotch, restrict the sample to 750 ml bottles, drop products whose quarterly prices have a coefficient of variation above 15% (which mechanically removes posting errors), and review all remaining classifications by hand. The balanced Q1–Q9 panel contains 675 products (537 single malts, 138 blends).

Figure 10 plots quarterly geometric-mean wholesale prices by type. SM and Blended trajectories are roughly flat through 2019, with mild SM divergence beginning Q4 2019.

Tercile Event Study. I replicate the price-tercile heterogeneity test from §3.2 on the wholesale panel using the same Spec A construction. Tercile cutpoints are computed from single

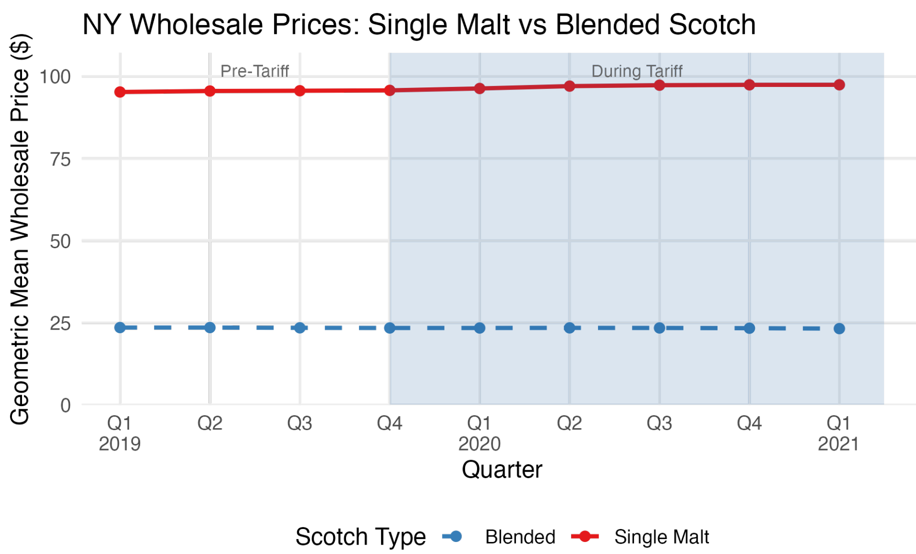

Figure 10: Geometric mean wholesale prices, single malt vs blended. The geometric mean is used because the NYSLA panel includes a heavy right tail of ultra-luxury bottlings (for example, a Macallan Fine & Rare 1940 listed at $22,815 per 750 ml). These products have stable listed prices across the panel and pass the CV filter despite negligible sales. They pull the SM arithmetic mean to roughly four times the SM geometric mean. The PLCB does not stock products of comparable price, so the corresponding figure for PA retail uses the arithmetic mean.

malt pre-tariff prices, and Blended Scotch products are assigned to the same dollar ranges. Figure 11 plots the per-tercile event-study coefficients.

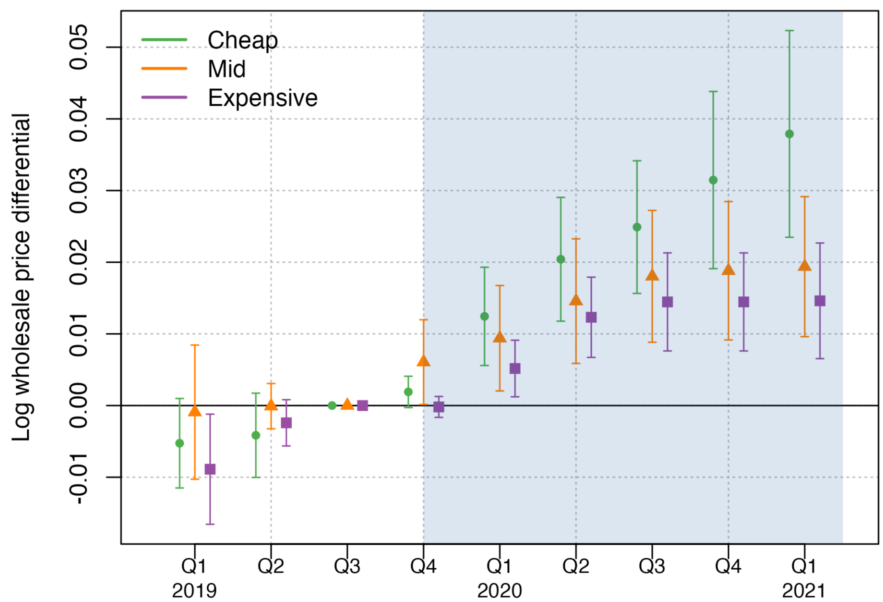

Figure 11: Wholesale Event Study by Price Tercile (Spec A: SM-Defined Bins). Each line is a separate event study estimated within the corresponding tercile, using the SM-defined dollar cutpoints applied to both treated and control products. Reference: Q3 2019. Standard errors clustered by product. Within-tercile pre-trends are individually insignificant (joint Wald $p = 0.24$ for cheap, 0.98 for mid, 0.07 for expensive).

The qualitative pattern matches PA retail. The cheap tercile has the largest mature-window response, the expensive tercile the smallest. Wholesale magnitudes are roughly a quarter to a third of the PA retail magnitudes. The NYSLA list price sits one layer upstream of retail and includes the wholesaler's margin, and posted wholesale prices are strikingly sticky (the median posted price change across the tariff window is zero, with the response concentrated among the minority of products that reprice), so the levels of the wholesale and retail responses need not match. The informative comparison is the cross-tier pattern, which matches.

Continuous Heterogeneity Gradient. The continuous wholesale specification and its headline estimate are reported in §3 (Figure 5; $\kappa_9 = -0.0081$, two-sided $p = 0.007$). The unrated products that preclude a WA-based wholesale specification are concentrated among independent bottlers and private labels. Relative to the Pennsylvania retail price slope above ($\kappa_9 = -0.098$, Figure 9), the wholesale slope is roughly an order of magnitude smaller but goes in the same direction, with more expensive single malts showing smaller proportional price increases at the wholesale stage as well.

Pre-Trend Caveat on Pooled and Tier-Level Levels. The pooled SM-vs-Blended event study on the wholesale panel has a same-direction pre-trend in the level differential (joint Wald $F = 3.17$, $p = 0.04$). The SM-Blended log-price gap is roughly 0.75 percentage points smaller in Q1 2019 than in Q3 2019, leaning in the same direction as the post-tariff effect. The pre-trend likely reflects NYSLA price-list updates concentrated in late 2018 / early 2019 alongside the public progression of the WTO Boeing-Airbus dispute through 2019. I therefore do not report a pooled NYSLA pass-through estimate. The within-tercile pre-trends in Figure 11 pass formally, but the Q1 coefficients in Cheap (-0.005) and Expensive (-0.009) are similar in magnitude to the pooled and lean the same way.

The cross-tercile differentials and the continuous κ_q slope avoid this issue. The differentials measure how SM tier responses differ from one another rather than how SM moves relative to Blended at the level. The slope measures how SM responds to within-quarter cross-sectional variation in pre-tariff price. Both are immune to common SM-Blended level drift, so the headline at wholesale is the cross-sectional heterogeneity in Figure 5 and Table 10, not the pooled level.

Tercile Contrasts. Table 10 reports the fully interacted tercile DiD on the strict pre-vs-mature panel (Q1–Q3 versus Q6–Q9, with Q4–Q5 dropped to mirror the PA retail spec in §3.2).

The cross-tier differentials net out the level pre-trend. The Mid–Cheap differential is -1.4 percentage points (two-sided $p = 0.062$), and the Expensive–Cheap differential is -1.4 percentage points (two-sided $p = 0.033$). The Expensive–Cheap differential rejects the null of equal pass-through across terciles at the 5% level and the Mid–Cheap differential at the 10% level, with smaller price increases at the higher quality tiers.[14] The cheap-tier baseline level $\hat{\delta}_0 = +0.032$ inherits the same pre-trend as the pooled regression and is an upper bound on the partial-equilibrium cheap-tier effect, not a clean treatment estimate.

[14] Percent-change differentials are computed as differences in tier-level percent changes via $\exp(\hat{\delta}_0 + \hat{\delta}_k) - \exp(\hat{\delta}_0)$, expressed in percentage points; the underlying log coefficients in Table 10 are $\hat{\delta}_M = -0.014$ and $\hat{\delta}_E = -0.014$.

	Dependent variable:
	log(price)
During (Q6–Q9) × Mid tercile	0.006
	(0.004)
During (Q6–Q9) × Expensive tercile	0.008***
	(0.003)
During (Q6–Q9) × SM (Cheap baseline)	0.032***
	(0.005)
During (Q6–Q9) × SM × Mid tercile	-0.014*
	(0.007)
During (Q6–Q9) × SM × Expensive tercile	-0.014**
	(0.007)
Observations	4,725
R^2	0.999

*p<0.1; **p<0.05; ***p<0.01 (two-sided)

Table 10: NYSLA Wholesale: Fully Interacted Tercile DiD, Q1–Q3 (pre) vs Q6–Q9 (mature). Cheap tercile is the omitted reference. Standard errors clustered by product (*prod_item*). Stars and the p-values reported in the prose below are two-sided.

B The First-Order Incidence Framework: Derivation and Extensions

B.1 Derivation and Proof of Proposition 1

Let $e(p, u)$ denote a consumer's expenditure function and let $\Delta p_j = p_j^1 - p_j^0$ be the tariff-induced price changes. The compensating variation is $e(p^1, u^0) - e(p^0, u^0)$. By Shephard's lemma, $\partial e/\partial p_j$ equals compensated demand, which coincides with the observed basket at the initial point, so the first-order term of the expansion is the old basket priced at the new prices, $\sum_j x_j \Delta p_j$. Because e is concave in prices, this term is a global upper bound on the true compensating variation, not merely a local approximation, and the second-order correction $\frac{1}{2}\Delta p' S \Delta p$ is non-positive because the Slutsky matrix S is negative semidefinite: substitution can only attenuate the loss. Dividing by initial category expenditure and using the first-order equivalence of percent and log changes gives Equation 15; group-level R_g follows by expenditure-weighted aggregation, which is exact because first-order losses are linear in the weights. For Proposition 1, substituting $\Delta \ln p_j = \bar{\theta} + \bar{\kappa}(\mu_j - \bar{\mu})$ into Equation 15 and using $\sum_j w_{jg} = 1$ gives $R_g = \bar{\theta} + \bar{\kappa}(\bar{\mu}_g - \bar{\mu})$; differencing across groups eliminates the common terms. □

B.2 Discrete Choice

Demand in this paper is discrete choice, where the textbook derivation's interior-demand assumption does not literally apply. The same result follows from the expected-maximum-utility function: its derivative with respect to p_j is $-\alpha_i s_{ij}$ (McFadden 1978), so the money-metric first-order loss is $\sum_j s_{ij}\Delta p_j$, with choice probabilities playing the role of quantities. This is exactly the object the implementation in Section 5.2 computes.

B.3 Income-Level Extension

Proposition 1 scales the burden by category spending. Scaling instead by income multiplies R_g by the group's category expenditure share of income, adding a participation margin that the within-category statement holds fixed. Empirically the participation margin does not overturn the result: the annualized tariff burden is 0.0036% of income for low-income consumers against 0.0023% for high-income consumers, a ratio of roughly 1.5, with the middle bin bearing the least (0.0016%). The magnitudes appear small because single malt Scotch is a narrow category.

References

Acosta, Miguel and Lydia Cox (2024), "The Regressive Nature of the U.S. Tariff Code: Origins and Implications."

Allcott, Hunt, Benjamin B Lockwood, and Dmitry Taubinsky (2019), "Regressive Sin Taxes, with an Application to the Optimal Soda Tax." *The Quarterly Journal of Economics*, 134, 1557–1626.

Amiti, Mary, Oleg Itskhoki, and Jozef Konings (2014), "Importers, Exporters, and Exchange Rate Disconnect." *American Economic Review*, 104, 1942–1978.

Amiti, Mary, Oleg Itskhoki, and Jozef Konings (2019a), "International Shocks, Variable Markups, and Domestic Prices." *The Review of Economic Studies*, 86, 2356–2402.

Amiti, Mary, Stephen J Redding, and David E Weinstein (2019b), "The Impact of the 2018 Tariffs on Prices and Welfare." *Journal of Economic Perspectives*, 33, 187–210.

Amiti, Mary, Stephen J Redding, and David E Weinstein (2020), "Who's Paying for the US Tariffs? A Longer-Term Perspective." *AEA Papers and Proceedings*, 110, 541–546.

Auer, Raphael, Ariel Burstein, Sarah Lein, and Jonathan Vogel (2024), "Unequal Expenditure Switching: Evidence from Switzerland." *The Review of Economic Studies*, 91, 2572–2603.

Berry, Steven, James Levinsohn, and Ariel Pakes (1995), "Automobile Prices in Market Equilibrium." *Econometrica*, 63, 841–890.

Berry, Steven, James Levinsohn, and Ariel Pakes (1999), "Voluntary Export Restraints on Automobiles: Evaluating a Trade Policy." *American Economic Review*, 89, 400–430.

Berry, Steven T (1994), "Estimating Discrete-Choice Models of Product Differentiation." *The RAND Journal of Economics*, 25, 242–262.

Birchall, Cameron, Debashrita Mohapatra, and Frank Verboven (2024), "Estimating Substitution Patterns and Demand Curvature in Discrete-Choice Models of Product Differentiation." *The Review of Economics and Statistics*, 1–40.

Cavallo, Alberto, Gita Gopinath, Brent Neiman, and Jenny Tang (2021), "Tariff Pass-Through at the Border and at the Store: Evidence from US Trade Policy." *American Economic Review: Insights*, 3, 19–34.

Chen, Natalie and Luciana Juvenal (2022), "Markups, quality, and trade costs." *Journal of International Economics*, 137, 103627.

Cole, Matthew T. and Carsten Eckel (2018), "Tariffs and Markups in Retailing." *Journal of International Economics*, 113, 139–153.

Conlon, Christopher and Jeff Gortmaker (2020), "Best practices for differentiated products demand estimation with PyBLP." *The RAND Journal of Economics*, 51, 1108–1161.

Conlon, Christopher and Jeff Gortmaker (2025), "Incorporating Micro Data into Differentiated Products Demand Estimation with PyBLP." *Journal of Econometrics*, 105926.

Conlon, Christopher and Nirupama L. Rao (forthcoming), "The cost of curbing externalities with market power: Alcohol regulations and tax alternatives." *Journal of Political Economy*.

Cravino, Javier and Andrei A. Levchenko (2017), "The Distributional Consequences of Large Devaluations." *American Economic Review*, 107, 3477–3509.

Faber, Benjamin and Thibault Fally (2022), "Firm Heterogeneity in Consumption Baskets: Evidence from Home and Store Scanner Data." *The Review of Economic Studies*, 89, 1420–1459.

Fajgelbaum, Pablo D, Pinelopi K Goldberg, Patrick J Kennedy, and Amit K Khandelwal (2019), "The Return to Protectionism." *The Quarterly Journal of Economics*, 135, 1–55.

Fajgelbaum, Pablo D. and Amit K. Khandelwal (2016), "Measuring the Unequal Gains from Trade." *The Quarterly Journal of Economics*, 131, 1113–1180.

Feenstra, Robert C (1989), "Symmetric Pass-Through of Tariffs and Exchange Rates Under Imperfect Competition: An Empirical Test." *Journal of International Economics*, 27, 25–45.

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

Gandhi, Amit and Jean-Francois Houde (2019), "Measuring Substitution Patterns in Differentiated-Products Industries." *NBER Working Paper Series*, w26375.

Goldberg, Pinelopi Koujianou and Rebecca Hellerstein (2013), "A Structural Approach to Identifying the Sources of Local Currency Price Stability." *The Review of Economic Studies*, 80, 175–210.

Griffith, Rachel, Lars Nesheim, and Martin O'Connell (2018), "Income Effects and the Welfare Consequences of Tax in Differentiated Product Oligopoly." *Quantitative Economics*, 9, 305–341.

Grigolon, Laura and Frank Verboven (2014), "Nested Logit or Random Coefficients Logit? A Comparison of Alternative Discrete Choice Models of Product Differentiation." *The Review of Economics and Statistics*, 96, 916–935.

Irwin, Douglas A (2019), "Tariff Incidence: Evidence From U.S. Sugar Duties, 1890 - 1914." *National Tax Journal*, 72, 599–616.

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

Khandelwal, Amit (2010), "The Long and Short (of) Quality Ladders." *The Review of Economic Studies*, 77, 1450–1476.

Kim, Donghun and Ronald W Cotterill (2008), "Cost Pass-Through in Differentiated Product Markets: The Case of U.S. Processed Cheese." *The Journal of Industrial Economics*, 56, 32–48.

Ludema, Rodney D and Zhi Yu (2016), "Tariff Pass-Through, Firm Heterogeneity and Product Quality." *Journal of International Economics*, 103, 234–249.

McFadden, Daniel (1978), "Modelling the Choice of Residential Location."

Miller, Nathan H and Matthew C Weinberg (2017), "Understanding the Price Effects of the MillerCoors Joint Venture." *Econometrica*, 85, 1763–1791.

Miravete, Eugenio J, Katja Seim, and Jeff Thurk (2018), "Market Power and the Laffer Curve." *Econometrica*, 86, 1651–1687.

Miravete, Eugenio J, Katja Seim, and Jeff Thurk (2020), "One Markup to Rule Them All: Taxation by Liquor Pricing Regulation." *American Economic Journal: Microeconomics*, 12, 1–41.

Miravete, Eugenio J., Katja Seim, and Jeff Thurk (2023a), "Elasticity and Curvature of Discrete Choice Demand Models."

Miravete, Eugenio J., Katja Seim, and Jeff Thurk (2023b), "Pass-Through and Tax Incidence in Differentiated Product Markets." *International Journal of Industrial Organization*, 90, 102985.

Nakamura, Emi and Dawit Zerom (2010), "Accounting for Incomplete Pass-Through." *The Review of Economic Studies*, 77, 1192–1230.

Nevo, Aviv (2001), "Measuring Market Power in the Ready-to-Eat Cereal Industry." *Econometrica*, 69, 307–342.

Petrin, Amil (2002), "Quantifying the Benefits of New Products: The Case of the Minivan." *Journal of Political Economy*, 110, 705–729.

Sangani, Kunal (2026), "Complete Pass-Through in Levels." *The Quarterly Journal of Economics*, 141, 1077–1135.

Seim, Katja and Joel Waldfogel (2013), "Public Monopoly and Economic Efficiency: Evidence from the Pennsylvania Liquor Control Board's Entry Decisions." *American Economic Review*, 103, 831–62.

Weyl, E Glen and Michal Fabinger (2013), "Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition." *Journal of Political Economy*, 121, 528–583.