The Pink Tax: (Why) Do Women Pay More?*

Kayleigh Barnes† Jakob Brounstein‡

June 22, 2026

Abstract

We evaluate the existence of the pink tax: the hypothesized price premium on women's consumer goods. Using retail scanner data, we find women pay 5.1% more for consumer packaged goods (CPGs) in the same product-by-location market. When we observe wholesale prices, we find women's goods feature lower markups. Estimating a constant elasticity of substitution model of demand on the near universe of CPGs and a differentiated products demand model within the disposable razors market yields similar conclusions: the pink tax is not sustained by higher markups, but by women sorting into goods with higher marginal costs and implied quality valuations.

JEL codes: J16, D12, L11

Keywords: Gender inequality, Pink tax

*We would like to thank Ben Handel, Jon Kolstad, Ben Faber and Nano Barahona for their mentorship and guidance. We thank Matt Backus, Mathilde Bombardini, Nina Buchmann, Bryan Chu, Giovanni Compiani, Stefano DellaVigna, Thibault Fally, Jonathan Holmes, Xavier Jaravel, Leticia Juarez, Kei Kawai, Peter Levell, Sarah Moshary, Mathilde Muñoz, Nina Rousille, and Damian Vergara, as well as participants in the UC Berkeley IO and Trade lunch seminars and University of Washington PHEnOM seminar for helpful comments. Our own analyses and calculations are based in part on NielsenIQ Retail Measurement Service and Consumer Panel Service data for all categories for the 16-year period ending in 2019, for the total US. The conclusions drawn from the NielsenIQ data provided by Nielsen Consumer LLC are those of our own and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein or in developing, reviewing or confirming the research approaches used in connection with this report. Researchers own analyses calculated (or derived) based on data from National Promotion Reports, LLC and marketing databases provided by the Kilts Center for Marketing Data Center at the University of Chicago School of Business. The views expressed in this work are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System.

†Corresponding Author. Federal Reserve Board of Governors, 20th Street and Constitution Avenue N.W., Washington, D.C. 20551; email: kayleigh.n.barnes@frb.gov

‡Revelio Labs; The Institute for Fiscal Studies; Skatteforsk - Centre for Tax Research; email: jakob.brounstein@gmail.com

1 Introduction

Is it more expensive to be a woman? The notion that there exists a price premium on women's consumer goods relative to those of men is colloquially referred to as the "pink tax". The concept has received considerable attention in popular media and has spurred recent legislation in US states to prohibit gender based price discrimination on comparable goods and services. Existing studies of the pink tax focus on narrow sets of goods that explicitly feature gender-based marketing, like health and beauty products, and compare prices between men's and women's products that are otherwise identical, finding mixed evidence for the existence of a pink tax (Moshary, Tuchman, and Vajravelu 2023; Guittar et al. 2022; NYCDCA 2015; Duesterhaus et al. 2011; Manzano-Antón, Martinez-Navarro, and Gavilan-Bouzas 2018). We build on this work by studying gender differences in prices and markups across the near-universe of consumer packaged goods (CPG). This broad approach allows us to capture the role of differential sorting by men and women into products that differ in cost and quality in generating price and markup disparities between men and women across their entire consumption bundles rather than just health and beauty products.

This paper explores the existence and underlying mechanisms of the pink tax. We start by describing consumption baskets for men and women and analyze how they vary by quantity, price, and diversity of products consumed. We then decompose observed price differences into markups and marginal costs. We consider a broad definition of the pink tax based on any channel through which women may face higher prices in the retail consumer packaged goods (CPG) space. This definition expands upon the legislative and common definition of the pink tax, which applies to differential pricing of otherwise identical products from the same manufacturer on the basis of gender (e.g. pink razor versus blue razor).[1] Our broader definition is consistent with Stigler's (1987) definition of second degree price discrimination as markup differences between groups due to versioning and product differentiation. With this definition, we capture the role of differential sorting by men and women—a channel that the legislative definition does not address, and one that prior work has not fully characterized—in

[1] In some instances, the pink tax refers to the luxury, sales, or value added taxes statutorily placed on women's hygienic products. Our analysis focuses on the more general case of price differences between men's and women's consumer goods.

generating price and markup disparities between men and women.

We find that, averaged across the entire grocery consumption basket, women pay 5.1% higher per unit prices than do men for products in the same product-by-location market. This price difference is sustained not just by purchases of explicitly gendered products, like men's and women's razors, but also by differences in purchasing habits between men and women for food and household items. This finding could be driven by three distinct economic mechanisms that contribute to either the markup or marginal cost component of a product's price: (i) women could exhibit less elasticity of demand than men leading to higher markups on products disproportionately bought by women, (ii) women could consume products with greater market power or from less competitive markets than men resulting in women paying higher markups, or (iii) women could consume products with higher marginal costs. Disentangling these mechanisms will allow us to attribute the observed price premium women pay to price discrimination or to differences in preferences between men and women.

To characterize the pink tax and gender differences in consumption habits, we employ several data sets that contain detailed information on individuals and their purchases, store-level product offerings, and retail prices. The NielsenIQ Consumer Panel Survey features a rotating panel of households and the near-universe of their purchases at big box retailers and grocery stores. Importantly, the data includes rich household demographic information as well as highly detailed product and purchase characteristics—including deal or sale usage, prices paid and quantities consumed, and a hierarchy that aggregates products into tractable market definitions. In order to assign purchases to a specific gender, we restrict the bulk of our analysis to single-member households, an approach we discuss and justify below. We augment the Consumer Panel with the NielsenIQ Retailer Scanner data which contains store level data on prices and quantities sold in any given week. Lastly, we incorporate data from PriceTrak to inform the wholesale costs for a subset of goods in our data, from which we can directly compute retailer markups.

We begin by establishing the existence of systematic gender differences in consumption and pricing along two margins: consumer behavior and the product space. To document consumer behavior, we describe consumption bundles for men and women, documenting

differences in their unit prices and composition. We find that women spend about 6.7% more annually than men do on retail CPG consumption and that their consumption bundles are larger and more diverse. The products that women purchase are on average 5.1% more expensive per unit than those purchased by men in the same product-by-location market.

In the product space, we document a significant share of products that are exclusively bought by one gender, with the majority of these products gendered towards women. These products are particularly common in markets with explicit gender differentiation in marketing and product design, such as in beauty and personal care goods. We categorize products bought at least 90% of the time by one gender (after applying the NielsenIQ probability weights to reweight purchases) as "gendered" products, categorizing all other products as "ungendered" (using alternate cutoffs to demonstrate the robustness of our results). We then decompose the average 5.1% price premium paid by women into a contribution from differential sorting into ungendered products and from purchases of highly gender-segregated products, finding that women pay on average 4.8% higher price-per-unit on *ungendered* products relative to men and that women pay an average of 21.2% higher prices on gendered products relative to those of men. While gendered items carry large price premiums for women, they make up a small share of actual purchases; the bulk of the price premium is driven by women buying more expensive ungendered items than men.

We then turn our attention to understanding the demand and supply mechanisms that give rise to women paying higher prices. Profit maximizing firms set prices as a function of own-price elasticities, market shares, cross-price elasticities of products owned by the same parent company, and marginal costs. To assess the relative importance of these different channels, we model and estimate demand and supply under a variety of assumptions, attributing differences in pricing and product choice to markups and marginal costs.

We estimate demand elasticity differences between men and women across the entire retail grocery consumption basket. We develop a simple, tractable model following Faber and Fally (2022), assuming a constant elasticity of substitution within each market $\times$ gender category that allows us to estimate demand by gender in the aggregate population. We aggregate individual-level purchase data to the gender-product-location-market level and we find that,

on average, women exhibit a price-elasticity of demand between 11% and 26% greater in magnitude than do men. This finding implies that, in a standard inverse-elasticity markup setting, women face lower rather than higher markups on average. This elasticity difference varies across CPG departments (e.g. dry grocery, health and beauty goods, etc.), but we find no department in which women exhibit significantly less price elasticity than men.

We further adapt the model of Faber and Fally (2022) to our setting, allowing us to estimate by-gender subjective product quality, intrinsic quality, and quality-adjusted price across goods. We find that women's consumer goods are assigned 25% greater intrinsic quality than goods consumed at gender parity, whereas those consumed by men are assigned only 10% greater quality. Individuals with the greatest female-goods concentration of their consumption baskets feature an average quality-consumption level about 10% higher than the mean-product quality, whereas those with the greatest male-goods concentration of their consumption baskets see a 5% lower average quality consumption level. We also document that women's goods do not operate in less-competitive markets than do those of men's, implying that a competition-channel is unlikely to generate greater markups for women.

We corroborate this central finding by implementing several additional designs that leverage complementary data and identification techniques. We combine the scanner data with data on wholesaler prices paid by retailers from PriceTrak. Wholesale prices represent the cost of the product charged to the retailer. We construct retail markups and observe that conditioning on wholesaler costs largely eliminates the observed pink tax; we also find no significant difference in retailer markups paid by men and women. Studying the goods themselves, we find higher markups placed on men's goods than on women's goods: within market, the most-male goods face a 10pp higher markup than the most-female goods; when weighting by expense, this difference expands to 30pp.

Finally, we estimate overall markups and marginal costs of production for disposable razors using a differentiated products demand model (Berry, Levinsohn, and Pakes (1995)). Thus far, our work speaks to average differences in prices paid, demand elasticities, product quality and markups across the entire consumption basket. By focusing on a narrow product market, we can incorporate additional model complexity, like heterogenous tastes, as well as

product ownership information, which will allow us to test for price discrimination within a brand. We allow consumers to have heterogeneous tastes over both prices and product gender, which generates flexible estimates of own- and cross-price elasticities. We combine these elasticity estimates with observed prices and data on product ownership in a simple supply side model to construct estimates of markups and marginal costs. We find women's razors have higher marginal costs of production and that women's razors have 11pp lower markups than men's razors produced by the same firm.

We find that women *do* pay higher prices than do men for goods in the same product category, but that the pink tax is not driven by price discrimination, but rather marginal costs and quality. While our primary aim is to comprehensively document and understand the mechanisms underlying gender differences in prices and consumption, our results carry natural implications for ongoing policy debates. Recently ratified and proposed legislation across different countries and subnational jurisdictions[2] aim to ban differential pricing of products that differ only in gendered marketing. Like Moshary, Tuchman, and Vajravelu (2023), our results suggest that these laws are likely to be ineffective at addressing price disparities between men and women, as the majority of our observed pink tax can be explained by men and women sorting into products that differ by more than just gender. Moreover, because we attribute gendered price differences to differences in marginal cost rather than markups, such legislation may risk inducing product exit. However, a full welfare assessment would require explicitly modeling and estimating product exit and weighing these costs against any redistributive benefits of regulation.

1.1 Related literature

Our paper contributes to the literature on gender disparities, consumption inequality and price discrimination. In spite of its prevalence in popular discourse and policy, there are few studies that rigorously substantiate the pink tax. Much of the direct evidence on the pink

[2]The state of New York has banned pricing on the basis of gender through bill S2679 which took effect in 2020. A similar bill, AB 1287, was signed into law in California by Governor Gavin Newsom on Sept. 27, 2022. The Pink Tax Repeal Act has been presented in Congress four times and aims to put national law in place similar to the New York and California policy.

tax comes from government reports or academic articles that consider either a limited set of products that gender matched in a subjective manner and document differences in list prices rather than actual prices paid. (NYCDCA 2015; Duesterhaus et al. 2011; Manzano-Antón, Martinez-Navarro, and Gavilan-Bouzas 2018; Manatis-Lornell et al. 2019) These studies find that women's goods have about a 5-7% price premium but do not attempt to allocate this price difference to differences in markups or differences in the cost of production. Recently, Moshary, Tuchman, and Vajravelu 2023 examine the existence of the pink tax in health and beauty product markets. Consistent with our findings, they document higher prices for women's products; however, after controlling for ingredients, they find little evidence of a pink tax. Their results suggest limited scope for price discrimination in the form of different prices for otherwise identical products. We build on these findings by examining the extent to which differential sorting into products contributes to gender differences in prices paid across all CPG product markets, not just health and beauty. We then analyze whether these price premiums are driven by markups—consistent with price discrimination—or by marginal costs arising from differences in consumer preferences.

There is a large literature on the gender wage gap and in savings behavior (Blau and Kahn 2017). Taking into account differences in the cost of consumption prompts us to re-frame the widely-studied difference in wages between men and women as a *nominal* wage gap. Moretti (2013) has shown that population specific price indices have important implications for wage inequality in real terms. In line with this, the presence of a gender gap in consumption prices augments these wage inequalities by reducing women's purchasing power in nominal terms, which also has implications for aggregate savings and wealth accumulation differences between men and women. We note, however, that to the extent these price differences reflect women sorting into higher-quality goods rather than price discrimination, a revealed preference argument would suggest women derive commensurately higher utility from their consumption bundle, tempering the welfare implication of this nominal gap.

Our work is closely related to research on inequality in consumption and product offerings. The consumption literature has documented that inflation, price indices and product offerings exacerbate inequality between rich and poor households (Jaravel 2019; Argente

and Lee 2017; Faber and Fally 2022). Our work on gender explores a new angle through which price index inequality may shape wealth inequality at large and our findings suggest that women may experience inflation and product innovation in different ways from men. Our finding that women and men sort into inherently different products suggests that their preferences are systematically different, which may be a result of differences in social norms or market experience. Bronnenberg et al. (2015) demonstrate that market and professional experience affect product choice, where pharmacists and chefs less frequently purchase more expensive brand name items (as opposed to generic-brand equivalents) in CPGs.

Finally, we contribute to the literature on price discrimination and optimal pricing. The work on gender based price discrimination focuses on first degree discrimination in bargaining contexts, finding that women often pay higher markups (Ayres and Siegelman 1995; Goldberg 1996; Trégouët 2015; Castillo et al. 2013). Our work investigates the existence of second degree price discrimination (also known as *versioning*) against women in CPG markets. Product differentiation and second degree price discrimination are sometimes thought of as separate phenomena but Stigler (1987) defines price discrimination as any markup difference between consumer groups. Relatedly, we speak to the body of work studying discrimination among CPG retailers: Hendel and Nevo (2013) finds that grocery store chains utilize promotional sales as a way to intertemporally price discriminate against consumers. However, other work finds that retail chains do not necessarily engage in optimal pricing decisions: DellaVigna and Gentzkow (2019) find substantial price mis-optimization for retail chains, where stores typically implement uniform prices across locations irrespective of local demand and cost factors. Our work examines how pricing of differentiated products could be a form of second degree price discrimination with certain consumer groups being charged systematically higher markups, which we conclude is *not* the case along the gender dimension.

2 Data

We combine data from three main sources to conduct our analysis. Our main analyses rely on data from NielsenIQ including the HomeScan Panel (HMS) and the Retailer Scanner Data (RMS). The HMS data contains purchase histories for a rotating panel of households from 2004 to 2019. In brief, we use the HMS to assign gender to products (detailed in Section 3.2) and to directly study differences in consumption patterns by gender. The RMS data contains anonymized purchases of products aggregated to the Universal Product Code (UPC) $\times$ Store $\times$ Week level from 2004 to 2018. We use the RMS to more accurately observe product prices and study differential demand sensitivity along the UPC-gender measures we construct with the HMS data. Lastly, we incorporate data from National Promotion Reports' PriceTrak database (PromoData), which features data on wholesaler prices charged to retailers for certain products from 2006-2011. While we discuss these data in turn, see Bronnenberg et al. (2015) and Allcott, Lockwood, and Taubinsky (2019) for further discussion of the NielsenIQ data.

HomeScan Panel ("HMS"). The entire HMS features data on the shopping trips and transactions of approximately 60,000 households per year. Households remain in the panel for on average 55 months, with approximately 200,000 distinct households rotating through the HMS in total. The data report purchases made by households on the 20 million shopping trips from 2004 to 2019 made by the panelists. For each individual item purchase, we observe transaction metadata such as date, store/retailer-info, and panelist identifier, as well as granular data on product and transaction details including prices paid, amounts and units of quantities purchased, deal or sale usage, and detailed nests of product identifiers.

We primarily use the HMS data to document differences in the purchasing behavior of men and women and understand how product markets differ for men and women. To confidently assign product purchases to consumer gender demographics, we restrict our consumer panel to single-individual household-years, which eliminates 74.6% of panelist-years in the HMS. This restriction is crucial for our approach in assigning product gender to goods, as purchases in multi-member households cannot be confidently attributed to a specific individual and their gender. The relevant concern with this restriction is whether the product gender

classifications derived from singles generalize externally to the consumption of multi-member households. We find strong evidence that they do, where 95.4% of UPC purchase volume in the unrestricted sample is accounted for by products that receive a gender classification from our singles sample. This high degree of overlap suggests that the consumption sets of single- and multi-member households coincide substantially, and that our product gender classifications are neither driven by goods idiosyncratic to singles nor systematically censoring goods idiosyncratic to multi-member households. This said, it is not obvious ex-ante how the gender price differential and its causes would differ between these two sample specifications, particularly given that single- and multi-member households almost certainly differ in their demand characteristics and price/substitution elasticities, and that it is not clear how shopping behavior varies when a principal purchases on behalf of another household member. Ultimately, we view concerns with this restriction as pertaining largely to external validity rather than internal validity, and the near-complete overlap in brand coverage between our sample and the unrestricted panel suggests this concern is limited.[3]

Our final panelist sample consists of 49,256 households which we use to study differences in consumer behavior. We report summary statistics for the sample in Table A.1. Our main specifications also control for demographic differences that might otherwise present compositional concerns in comparing single men and single women.

Table A.1 shows that the single men and women in our sample appear similar in many respects, but do exhibit relatively small but statistically significant differences in observable characteristics. The men in our sample report on average 20% higher income[4] and slightly higher levels of education; we frequently control for these demographic characteristics in our specifications in order to avoid comparing men and women that differ significantly along these characteristics. Also note that women are overrepresented in our sample as they are in the general HMS, comprising about 70% of our panelists. However, applying the HMS

[3]We also restrict to households-years with at least 12 shopping trips, which eliminates only 15 out of nearly 162 million panelist-years for both family and single-person panelist households. Lastly, we drop single-person households that report multiple genders compositions (31 individuals in the singles sample) for the reason that we cannot distinguish coding errors from panelists that undergo gender transition.

[4]The HMS reports panelist household income in discrete buckets. All results referring to HMS panelist income make use of the midpoint of each discrete income buckets used for the household income field, following Allcott, Lockwood, and Taubinsky (2019).

projection weights yields a near 50-50 gender balance. We employ these weights for the entirety of our analysis that uses the HMS.

The second component of our analysis focuses on how the product market space varies by gender. For this analysis, we restrict our data to products to which we can confidently assign a UPC gender. We describe our methodology in detail in Section 3.2. Purchases in our main sample comprise approximately 1.87 million UPCs. We are able to confidently assign gender to 1.17 million UPCs, comprising 98.7% of the total expenditure volume of our singles panel.

Retailer Scanner Data ("RMS"). The RMS data contain UPC-store-week level prices and volumes of products purchased by consumers from 2004 to 2018. The strength of the RMS data lies in its comprehensiveness of US sales. We use the RMS data to model demand in select markets that have a high level of gendered products (as identified in the HMS data).[5] While the HMS tracks all retail purchases for a household from any store, the RMS contains a select set of stores. For our main analysis, we keep only stores that are part of a larger retail chain rather than independent stores which is important for constructing our price instruments described in Section 4. The NielsenIQ data has limited coverage of independent stores, so this constraint has limited impact on differential selection *within* the NielsenIQ data.

Both components of the NielsenIQ data feature a highly detailed product hierarchy classification that organizes all goods into smaller nests with increasing degrees of specificity. Products in the NielsenIQ are identified with their Universal Product Code (UPC) which corresponds to a unique barcode. We define products at the UPC level. All UPCs fit into one of ten *departments* (the broadest category, e.g. "Health and Beauty" and "Dry Grocery"). From here, products in a department are allocated to *product groups*—of which there are 120 total—such as "Shaving Needs". Finally, UPCs in the same Product Group are assigned to *product modules*—the most granular grouping of multiple products—e.g. "Disposable Razors". The NielsenIQ data identifies over 1,300 distinct product modules. Brand description represents an alternate grouping that features the brand name for a given set of UPCs,

[5]The RMS is also used to generated the prices paid as observed in the HMS.

not strictly contained in any single product module or group contained, such as "Venus" (a division of razors marketed to women by Gillette), for the brand of razors. We consider product modules as constituting a self-contained goods market; however for certain reduced-form analyses, we further divide product modules into module-units to compare quantities and quantity-weighted prices across purchases. For example, the coffee product module contains bagged coffee measured in weight (ounces) and Keurig cup coffee measured as a count (number of K-cups).

We construct transaction-level UPC unit prices as the total amount paid for a UPC divided by the observed quantity. We construct observed quantity as product of 1) the quantity of units contained in a single purchase-unit (e.g. 5 ounces) 2) "multipack" (a variable indicating the number of units contained in the package; for non-multipacks, this variable is equal to one), and 3) the number of overall units purchased. For example, we calculate the unit price corresponding to a \$10 purchase of a 5 count multipack of razors, where each multipack contains 2 units, and each unit contains 4 razors as
$\$10/(5\ multipacks\ \cdot\ 2\ units\ per\ multipack\ \cdot\ 4\ razors\ per\ unit) = \$0.25\ per\ razor$.[6]

Pricetrak PromoData ("Pricetrak"). Lastly, the PriceTrak PromoData report wholesaler prices charged to retailers and therefore allow us to impute retailer markups. The PriceTrak data feature retailer cost-data of individual UPCs for a variety of time- and geographic-denominations from 2006 to 2011, with geographic disaggregations covering 55 markets (coinciding with the metropolitan areas around large US cities). The match rate of UPCs in the Promodata to the NielsenIQ datasets is relatively low. Only about 10% of the 430,000 distinct UPCs in the RMS data match to PromoData (largely because the PriceTrak data report on a smaller subset of UPCs); however, these UPCs account for 40% of purchase volume observed in the HMS. We combine the data from PriceTrak on wholesaler prices with NielsenIQ data on post-deal consumer prices to compute retailer markups relative to wholesaler prices.

[6]We find that women consume multipack items about 9% less frequently than do men and that multipack purchases constitute around 5% of transactions. We also show that differential bulk and multipack purchases between men and women do little to explain our price gap estimates in Section 3.2.

3 Price Disparities Across the Consumption Bundle

3.1 Consumer Behavior by Gender

Aggregated consumption baskets. We begin by analyzing differences in how men and women consume and purchase products. Figure A.1 establishes a baseline finding that, controlling for income and demographic characteristics, women's yearly spending on CPGs exceeds that of men's by 6.2%.[7] Additionally, women's product diversity measured by unique UPC purchases is greater than men's by about 31.3% and their consumption baskets are larger than men's in terms of items purchased by about 9.4%. This pattern is primarily driven by differences in behavior in consumption of Health and Beauty products, where women spend 70% more than men, consume 73% more unique products, and consume 66.5% more items. However, we continue to observe significantly estimated, if mildly attenuated differences in spending and product purchase patterns for women relative to men after excluding Health and Beauty products. Among these products, women spend 1.8% more, have 27% greater product diversity and 7% more items than men.[8]

This aggregate spending difference between men and women could arise from differences in prices paid for similar goods, differences in quantities purchased, or differences in the set of goods purchased. As a clarifying example, consider consumption habits for hair products. Both men and women buy shampoo, but women tend to have longer hair than men, which may lead them to buy more bottles of shampoo over the course of a year increasing their total spending through quantity. It is also possible that women have preferences for higher-priced shampoos, increasing their total spending by paying higher prices. Finally, it could be that women buy leave in conditioner or mousse that is not typically bought by men, increasing their total spending by purchasing a different set of products than men.

[7] We approximate coefficient estimates $\hat{\beta}$ from log price regressions throughout the paper as percentage differences as $e^{\hat{\beta}} - 1$.

[8] Table A.2 displays corresponding tabular estimates for these objects in means and with both univariate and multivariate regressions results to estimate female-male differences. Beyond corroborating the differences in expenditure volume and purchase volume/diversity, the table also shows that female consumers make significantly greater use of coupons and deals.

Differences in unit prices. Popular discussion of the pink tax is often focused on differences in prices paid between men and women for similar goods in the same market, so we focus primarily on the price channel. We compare per unit prices paid by men and women for products in the same market with the following specification:

$$log(P_{ijct}) = \phi_{m(j)ct} + \beta \mathbb{1}\{Woman_i = 1\} + \gamma X_i + \epsilon_{ijct}, \tag{1}$$

where i denotes the individual, j denotes the product purchased, $m(j)$ is the module of product j, and c denotes geography, and t denotes time. Table 1 Panel (a) presents the results. Column (1) regresses log unit UPC price on a woman indicator and includes fixed effects for the interaction of product module, units the good is sold in and the year of purchase; Column (2) estimates this same regression on the most-restrictive sample specification as Column (7), for the purposes of comparison across a constant sample. We interpret the 2.9% estimate as the raw difference in prices paid between single men and women in the USA, not accounting for other demographic factors or location and retail chain sorting. Columns (3)-(7) sequentially add permutations of fixed effects and additional covariates. Note that sample size changes across columns reflect singleton purchases that cannot be estimated under different fixed effect specifications.

Column (3) runs the same specification but adds in controls for age, income, and race. The increase in the coefficient estimate from 2.9% to 5.1% highlights the importance of demographic differences between single men and single women because older and lower income consumers tend to buy lower priced products. Columns (3) and (4) subsequently add in county and retailer fixed effects. Column (4) can be interpreted as the contribution of women sorting into more or less expensive locations, because the coefficient change is small, the contribution is minimal. Similarly, Column (5) can be thought of as the contribution of sorting into more or less expensive retail chains, i.e. Whole Foods vs. Walmart. Controlling for the retail chain lowers our price premium estimate to 4.2%, suggesting that retail chain sorting plays a small but significant role. Finally, in Column (7) we add in fixed effects for month rather than year. The results indicate that women spend about 5% more than do men per unit of goods in the same product market, bought in the same retail chain, county, and

month. We consider this our preferred specification because it controls for a wide variety of potential differences that could arise between the two groups other than gender.[9] We estimate this preferred specification for each department in Table A.3, finding that women also pay higher unit prices in all departments except in Alcohol and General Merchandise.

Table 1 Panel (b) estimates Equation (1) while including UPC-level fixed effects instead of module-level fixed effects. In this specification, the interpretation of the coefficient becomes the difference in prices paid between men and women for the *exact same* product. Differences in prices paid for the same good can be attributed to differences in price shopping behavior, like coupon usage and sale shopping, suggestive of greater price elasticity. We sequentially add in fixed effects in the same manner as Table 1 Panel (a), so the coefficients can be interpreted as a raw difference between men and women in column (1) and then iteratively making comparisons between demographics, location, retail chain and month.

While we find that women, on average, buy more expensive products than do men, we find that they consistently spend less than men on the *same* product. In Column (3) we find that controlling for demographics attenuates this gap, likely driven by differences in use of coupons by age and income. Controlling for retail chain in column (4) increases the gap, which is consistent with women sorting into higher price chains. Column (7) shows that women pay 1.4% less for the same product than do men that are demographically similar shopping in the same retail chain-location-month market. Combining our results from Panels A and B suggests that women purchase higher-priced goods while also exhibiting behaviors associated with higher-elasticity shopping patterns. This finding is further corroborated by Table A.2, which also shows that women make greater usage of deals and coupons. Hendel and Nevo (2013) study promotional sales as a form of intertemporal price discrimination, our results would indicate that women are likely to comprise a larger share of the consumer base that benefits from this type of price discrimination (also substantiating previous related findings that women engage in price shopping to a greater degree than do men, e.g. Aguiar and Hurst (2005)).

[9]Table A.4 estimates a set of regressions analogous to the module-level results in our main specification of Table 1 Panel (a) while including the log number of units contained in a purchase (where non-multi-packs have unit value one). The table shows that our descriptive price differences are unaffected by multi-pack purchases.

3.2 Gender in the Product Space

We now shift our focus from consumer behavior to understanding how the product space varies by gender. Our descriptive evidence shows that women buy more expensive and larger consumption bundles and that the products they buy are more expensive relative to similar products bought by men. However, these observations could be driven by differences in purchase intensity of otherwise "ungendered" products or by purchases of products that are primarily bought by one gender. To more comprehensively characterize the pink tax, we develop a continuous measure of "product gender" on the UPC level, and decompose the observed descriptive pink tax of 5% into its respective contributions from gendered products and differential purchasing of ungendered products.

First, we assign values of gender-stratification to each UPC. We begin by calculating a woman purchase share for each UPC in our data as the fraction of overall purchase volume by women within our panel of single individuals. We define the estimator for the time-invariant woman purchase share of UPC j (the "UPC-gender") as

$$\hat{w}_j = \frac{\sum_{i \in \mathcal{I}} Purchase_{ij} \mathbb{1}\{woman_i = 1\} \cdot NielsenWeight_i}{\sum_{i \in \mathcal{I}} Purchase_{ij} \cdot NielsenWeight_i}. \tag{2}$$

This fraction assigns to each product a value $\hat{w}_j \in [0, 1]$ where 0 denotes a good that is only bought by men and 1 denotes a good that is exclusively bought by women. $NielsenWeight_i$ refers to the NielsenIQ projection weight assigned to individual i in the HMS data. We assign goods with $\hat{w}_j \leq .1$ as men's products, and those with $\hat{w}_j \geq .9$ as women's products.[10] Because employing these weights yield a sample representative of US consumers, we define UPC gender parity as $\hat{w}_j = 0.5$, as opposed to the initial gender balance of sample when unweighted. We use the continuous measure of UPC gender for most of our analysis, but we also employ these binary categories to more intuitively capture comparisons involving "highly gendered goods".

To reduce measurement error in our measure of UPC-gender, we only assign an observed women purchase share to products that are observed to be bought with sufficient frequency.

[10]The choice of cutoff is inherently arbitrary, but is guided by the empirical distribution of UPC gender, as displayed in Figure 1. We assess sensitivity to this threshold below.

To illustrate the necessity of this decision, consider that approximately two-thirds of the UPCs purchased by NielsenIQ panelists are only ever observed to be purchased once (although these UPCs represent less than one percent of the overall purchase volume reported in the HMS). By merit of only observing a single purchase, these UPCs would always be assigned to having an explicit gender of 0 or 1. Conceptually, each UPC in our data has a *true* UPC-gender, w_j, that we do not observe. We observe an empirical UPC-gender $\hat{w}_j$ as well as the UPC's number of unique purchases, n_j. An observed UPC-gender represents a draw from a binomial distribution. We can express the probability that the value $\hat{w}_j$ lies more than 5 percentage points away from its true value w_j as

$$P(w_j \notin (\hat{w}_j - .05, \hat{w}_j + .05)) = \int_{x \geq |\hat{w}_j \pm .05|} \binom{n_j}{\lceil \hat{w}_j n_j \rceil} x^{\lceil \hat{w}_j n_j \rceil} (1-x)^{n_j - \lceil \hat{w}_j n_j \rceil} f(x) dx, \qquad (3)$$

where $f(x)$ is the empirical probability density function of woman purchase share. We use Equation (3) to calculate thresholds for discrete bins of UPC-gender of radius 0.00025 from 0 to 1, $n^*_{jb(\hat{w}_j)}$, such that the probability that an individual UPC's observed UPC-gender deviates from its true UPC-gender by a value less than 0.05 is 95%. We map each observed UPC-gender to its nearest bin and only keep UPC-gender observations when the underlying number of unique purchases n_j exceeds $n^*_{jb(\hat{w}_j)}$.[11] Table A.5 replicates the preferred specifications from Table 1 using only the set of goods with identifiable UPC gender, yielding quantitatively and qualitatively similar results.

Figure 1 plots the distribution of woman purchase share for all products, Health and Beauty products, and all products excluding Health and Beauty. We observe significant excess mass at the right tail of the distribution where goods are bought exclusively by women, but only mild excess mass at the left tail of the distribution where goods are bought

[11] While we are only able to confidently assign UPC-gender to 700,000 out of 1.8 million unique UPCs, Figure A.2 shows those we are able to assign gender to account for greater than 95% expenses in the data. We find that gendered products make up a small share of purchases, 3.6% for men and 4.6% for women. Within Health and Beauty products however, gendered products make up 20% of women's purchases and 10% of men's purchases. Figure A.3 displays the gender-composition of UPCs by within each NielsenIQ department.

only by men.

UPC gender and UPC unit prices. We now describe how prices vary along our measure of UPC-gender. We map each UPC-gender $\hat{w}_j$ to a 10pp-width bin centered at values $b = 10 \cdot (\lfloor 10 \cdot (\hat{w}_j + .05) \rfloor) = \{0, 10, 20, \ldots, 90, 100\}$ and estimate the following regression:

$$\log P_{jct} = \theta_{m_j ct} + \sum_{b \in \mathcal{B}} \gamma_b \mathbb{1}\{w_j \in Bin_b\} + \varepsilon_{jct}, \tag{4}$$

where $\mathcal{B}$ is the set of 10pp-width bins that partition the interval $[0, 100]$, allowing for different price estimates along the UPC gender distribution. $\theta_{n_j ct}$ is a set of indicators for module m_j, county c, and half-year t.

Figure 2 plots the coefficients from estimating this equation, taking the 50th percentile bin as the reference point. The coefficients can be interpreted as averages across comparisons made of products in the same market and bought in the same location and time frame relative to products bought equally by men and women.

We observe significant price premiums of between $10 - 40\%$ for goods purchased exclusively by women relative to goods in the same product market purchased at gender parity. We observe no price premium for goods purchased exclusively by men. However, beyond the left and right tails of the graph, a striking pattern emerges: prices tend to monotonically increase in woman purchase share beginning around goods with 30% female purchase share. This increase in prices *along* woman purchase share suggests that our overall price premium of 4% from Table 1 is not only driven by explicitly gendered products (i.e. pink products and blue products) but also by differences in preferences for non-overtly gendered products. This finding is consistent with women having preferences for higher (perceived) quality items such as, for example, organic products, which we explore in Section 4.3 and in Section B (Ureña, Bernabéu, and Olmeda (2008)).

UPC and Consumer gender. To explore the interaction of consumer and UPC-gender, we run the same regression specification as in Table 1 column (5), but now include an indicator for whether a good is gendered and an interaction between the woman indicator

and the product gender indicator:

$$\begin{aligned} log(P_{ijct}) = \phi_{m(j)ct} + \beta_1 \mathbb{1}\{Woman_i = 1\} + \beta_2 \mathbb{1}\{GenderedProduct_j = 1\} \\ + \beta_3 \mathbb{1}\{Woman_i = 1\} \cdot \mathbb{1}\{GenderedProduct_j = 1\} + \gamma X_i + \epsilon_{ijct}. \end{aligned} \tag{5}$$

This specification has the desirable quality that by interacting consumer gender with a gendered product indicator, we allow for differential price estimates for male-gendered and female-gendered goods. Table 2 presents the results of these regressions. We find that women pay a price premium of 4.8% on ungendered products relative to ungendered products bought by men. Across all departments, men pay lower prices on their gendered products than they do for ungendered products by about 5%. The interaction coefficient shows that women pay about 14.5% more for gendered goods relative to ungendered goods. Overall, we find that women pay approximately 21.2% higher prices on gendered products than do men (calculated by exponentiating the difference of sums of coefficients for gendered product purchases between women and men). Columns (2) and (3) of Table 2 demonstrate that these findings are not driven by health and beauty products.[12]

However, while the magnitude of the price difference for gendered goods purchased by women is large, its contribution to the overall price premium is small. Figure A.2 gives some intuition for this small contribution, showing that highly gendered products make up an overall small share of consumers' consumption baskets. Indeed, the 5% price premium from Table 1 closely aligns with the purchase-weighted price averages that women pay on ungendered items. While we find evidence that female-gendered products see significantly higher unit prices, the vast majority of our observed pink tax is in fact driven by differential sorting between men and women on less-overtly-gendered products.

While our descriptive results show that on average, women do pay greater unit prices for goods in the same module × market than do men, they do not speak to the underlying mechanisms of this phenomenon. We now turn to estimating differences in supply, demand,

[12]Table A.6 replicates this finding using a more lenient threshold of UPC gender, where we assign goods below 25% female purchase share as "male" and those above 75% female purchase share as "female". We obtain nearly identical results with only very mild attenuation.

and competition between men and women and their respective goods in explaining the forces that generate the pink tax.

4 Gender differences in demand

To estimate demand differences between men and women, we augment the constant elasticity of substitution (CES) model used in Faber and Fally (2022). A CES model offers several desirable characteristics in our setting. We allow men and women to differ both in their price elasticity and in their subjective product quality evaluations. Additionally, the nesting structure of the model also features flexibility in allowing for heterogeneous price sensitivity across product modules or departments. Furthermore, beyond the model structure, the estimating equation we derive also features the intuitive interpretation of capturing the average difference in price elasticity of women versus men over product modules.

This approach ultimately allows us to make aggregated comparisons of the purchasing habits of men and women across a wide range of products to make inference on their average price elasticities and average product quality consumption. If the per-unit price premium observed on women's goods and on goods purchased by women more broadly is attributable to differences in markups, we should observe that women exhibit lower price-sensitivity on average as a consumer demographic. However, if the observed price-premium on goods purchased by women is attributable to differences in marginal costs, we should observe that women consume higher-quality goods than do men.

The model characterizes a representative consumer of gender g. The consumer allocates their income in an additively separable manner between a vector of retail goods G and consumption of the outside option:

$$U = U(V_G(g), C_g).$$

We assume that the basket of goods comprising the outside option C_g is consumed as a normal good.

The model aggregates products in two tiers. The consumer allocates consumption *across product modules* with Cobb-Douglas elasticity and *substitutes between goods within modules* with module×gender-specific constant elasticity of substitution. We index product modules as $n \in \mathcal{N}$, which perfectly partition the set of UPCs. Each module n contains goods G_n. The consumer maximizes their utility subject to their budget constraint by choosing a vector of quantities, G, that represents their consumption bundle across all goods:

$$V_G(g) = \prod_{n \in \mathcal{N}} \left[\sum_{j \in G_n} \left(q_j(g) \varphi_j(g) \right)^{\frac{\sigma_n(g)-1}{\sigma_n(g)}} \right]^{\alpha_n(g) \frac{\sigma_n(g)}{\sigma_n(g)-1}}. \tag{6}$$

Here, q_j refers to quantity consumed of UPC j, $\varphi_j(g)$ refers to the subjective product quality (product appeal) of a product j as evaluated by a consumer of gender g. $\sigma_n(g)$ represents the elasticity of substitution for gender g within module n, and $\alpha_n(g)$ denotes the share of expenditures by gender g allocated to a module $n \in \mathcal{N}$.[13]

Specifying the upper tier as Cobb-Douglas implies that comparisons of consumption amounts between products within the same module depend on their relative quality-adjusted prices:

$$\frac{b_j(g)}{b_k(g)} = \left(\frac{p_j / \varphi_j(g)}{p_k / \varphi_k(g)} \right)^{1-\sigma_n(g)}, \tag{7}$$

where $b_j(g)$ is the budget share spent on product j. From Equation (7), expressing gender in subscripts for notational ease, we derive our estimating equation while also allowing for consumption and prices to vary over geography and time:

$$\Delta log(b_{jgct}) = (1 - \sigma_{ng}) \Delta log(p_{jct}) + \eta_{ngct} + \varepsilon_{jgct}, \tag{8}$$

where differences are taken between two consecutive periods and η_{ngct} represents module × gender × county × half-year fixed effects to capture the local CES price index. Note that while Equation (7) is written in relative terms using the reference good k, in a CES system, it is the aggregate price of all other goods in the nest that matters, so that after taking logs

[13] Under a Cobb-Douglas upper nest it is the case that $\sum_{n \in \mathcal{N}.} \alpha_n(g) = 1$.

and differences, the CES price index is common across products and is absorbed by the fixed effects. We derive this estimating equation from a CES demand model, but this estimating equation can also be interpreted as an average price response within a market. Additionally, pooling all module markets together allows us to estimate a module-market average price responsiveness for each gender: $1 - \sigma_g$.

We face the standard issues of simultaneity in demand estimation where price changes may be correlated with demand shocks. To address this issue, we rely on two identifying assumptions employed frequently in empirical works estimating product demand. First, we assume that local demand shocks are uncorrelated and idiosyncratic across localities while supply shocks are correlated across space and retailers (e.g. Hausman (1999)). Second, we assume that retail chains set prices at the national or regional level and that these prices are set independent of local demand shocks following evidence presented in DellaVigna and Gentzkow (2019). From these assumptions, we estimate $(1 - \sigma_g)$, the gender-specific average substitution elasticity, using two instruments for prices. The first are Hausman instruments, which we construct as national leave-out means in price changes at the county or Nielsen Designated Market Area (DMA) level, $\Delta H_{jct} := \frac{1}{N^c_{jt}-1} \sum_{c' \neq c} \Delta \log(p_{jc't})$ where N^c_{jt} refers to the number of observations of UPC/Brand j across counties in half-year t.[14] The second are instruments that follow DellaVigna and Gentzkow (2019) and further developed by Allcott, Lockwood, and Taubinsky (2019) which are constructed as national leave-out means of price changes at the county-retailer chain level, $\Delta DV_{jcry} := \frac{1}{N^{cr}_{jy}-1} \sum_{c' \neq c, r' \neq r} \Delta \log(p_{jc'r't})$ where N^{cr}_{jy} refers to the number of observations of UPC/Brand j across retailer-counties in the half-year.[15] Both instruments may see exclusion restriction violations from national demand shocks, for example from national advertising campaigns or national holidays that drive certain purchases (e.g. hot dogs on the fourth of July). To help control for this possibility, we include market-by-time fixed effects in our estimation.

[14] County c level prices of good j in half-year t are constructed as simple means over observations of good j across retailers in the half-year within each county.

[15] County-retailer level prices of good j in half-year t are constructed as simple means over observations of good j in the half-year within each retailer-county. Much of the variation in the DellaVigna-Gentzkow instrument is driven by variation in how often a product is placed on a promotional sale. The timing of these sales is driven by a bargaining process between the retailer and the manufacturer and typically only one manufacturer is put on promotional sale at a time.

In our estimation procedure, we aggregate observations to the Brand × Nielsen Designated Market Area (DMA) × consumer gender × semester level. We aggregate UPCs to the brand-level in order to overcome sparsity issues where individual UPCs are not necessarily observed in consecutive semesters within a specific market.[16] This decision also allows us to compare our estimates with those of Faber and Fally (2022) to gauge the validity of our approach. We also include an additional distinction for the retail chain r depending whether we use instruments from DellaVigna and Gentzkow (2019). We estimate our model at the half-year level because many product categories are prone to stockpiling, which when observed in shorter time intervals, would bias our demand estimates towards greater price elasticity; this bias may further confound our estimates of $\sigma_{nf} - \sigma_{nm}$ if men and women exhibit differential stockpiling behavior. To address auto-correlation in the error term, we cluster standard errors at the brand × Nielsen DMA level.

Under the assumption of single-product firms, we derive own-price elasticity of demand as a linear interpolation of the elasticity of substitution between perfect monopoly and perfect competition:

$$\varepsilon_{jg} = \sigma_g - (\sigma_g - 1) \cdot s_{jg} \tag{9}$$

Where s_{jg} is the module-level market share of product j for gender g. Thus, we can calculate ε_{jg} as a function of known and estimated parameters. In the case of monopolistic competition, all market shares are approximately zero and ε_{jg} coincides with the elasticity of substitution, σ_g. Mapping elasticities to markups, we assume firms compete on prices and maximize firm profits given the demand that they face.[17]

4.1 CES Model Results

We begin by estimating differences in the elasticity of substitution, $\sigma_n(g)$, between men and women. Table 3 presents results of estimating an analogue of Equation (8) that fully

[16] We aggregate prices from UPCs to the brand level using a Törnquist price index. We use the procedure outlined in Section 3.2 to confidently assign continuous values of brand genders to 70.6% of 240,000 brands that appear in our data, comprising 99.85% of aggregate purchase volume.

[17] We do not assume single-product firms in computing markups. Section 5 calculates markups using retailer costs and sale prices directly and Section 6 assumes multi-product firms and constructs an ownership matrix to compute markups.

interacts price changes with an female consumer indicator and pooling the elasticities across all departments.[18] Note that the coefficient on $\log \Delta p_{jct}$ estimates $\beta_m := 1 - \sigma_m$. Therefore, when the female indicator turns on, the coefficient on the interaction of the female indicator and $\log \Delta p_{jct}$ estimates $\beta_{m-f} := (\sigma_m - \sigma_w)$ so that the net substitution elasticity for women is indeed given as $(1 - \sigma_f)$. Succinctly, assuming $\beta_m,\ \beta_{m-f} < 0$, one can map the estimated coefficients to gendered substitution elasticities as $\sigma_m = 1 + |\beta_m|$ and $\sigma_f = 1 + |\beta_m| + |\beta_{m-f}|$.

Accordingly, our main coefficient of interest is that on the full interaction term $\mathbb{1}\{g(i) = female\} \cdot \log \Delta p_{jct}$, given by $\hat{\sigma}_m - \hat{\sigma}_w$, the average difference in elasticity of substitution between men and women. The second row displays estimates of $\hat{\beta}_m := 1 - \hat{\sigma}_m$. In odd-numbered columns we include gender-module-market-semester fixed effects and even-number columns include gender-module-market-retailer-semester fixed effects. These specifications aim to capture the differences in demand elasticities between men and women for the same price change for the same product in the same market and purchasing setting, for which reason our preferred specification includes the retailer fixed effects. We use equal weights across UPCs, as expenditure weighting would mechanically upweight high-price observations and distort the price variation driving identification, although using expenditure weights produces consistent results.

Columns (1) and (2) display coefficients from estimation by Ordinary Least Squares. We obtain positive price coefficients, suggesting the presence of substantial simultaneity bias. However, in interpreting the coefficient on the gender differential price sensitivity term, if it is the case that demand shocks affect men and women identically, this regression does not need to be instrumented since the endogenous portion is differenced-out. In both specifications, this difference term indicates that women exhibit greater substitution elasticity do men, however only significantly so when controlling for retailer. However, we view this assumption as unnecessarily restrictive, and we proceed by instrumenting accordingly.

Columns (3)-(8) display different IV specifications. First, note that the sign of log price coefficient changes from positive to negative going from OLS to IV specifications. This pattern exactly mirrors findings in Faber and Fally (2022), who—using NielsenIQ scanner

[18] Table A.7 presents the first stage and reduced form results respectively. The table shows strong first stages and consistently negative reduced form estimates across specifications.

and Homescan data—also document this sign reversal in estimates of $(1 - \hat{\sigma})$ when moving from OLS to IV specifications. We interpret this similarity as consistent with the presence of meaningful bias in OLS estimates.

Across instrumental variables specifications, we find that women consistently demonstrate significantly *greater* price sensitivity than do men. We find that for within the same market, women are between about 11 and 26 percent more price-elastic than men. Our preferred specification, Column (8), uses both Hausman and DellaVigna-Gentzkow retailer instruments as well as retailer fixed effects and gives an elasticity of substitution for men of 1.51 and for women of 1.73.[19] This differential is markedly larger than its OLS counterpart. Had demand shocks affected men and women identically, the endogenous component would have cancelled-out in the difference, leaving the gender gap unbiased and therefore unchanged by instrumenting; our observation that it instead grows substantially under IV shows that the endogeneity does not perfectly difference out, and that it falls more heavily on the women's side. This finding is consistent with women's prices co-moving more strongly with demand, although we do not identify this channel directly.

Interestingly, we also estimate elasticities of substitution slightly lower in magnitude than estimates from Faber and Fally (2022) ($\hat{\sigma} = 1.66$ in column (6) versus $\hat{\sigma} = 2.20$ in the analogous specification in Faber and Fally (2022)). However, our sample differs from Faber and Fally (2022), as our subsample of the HMS data only consists of single individuals, whereas Faber and Fally (2022) do not make this restriction. Nonetheless, our estimates imply positive, finite markups, consistent with a standard monopolistic competition framework. Moreover, the estimated implied average markup ranges between 150% and 300%, consistent with the higher end of recent estimates such as in Döpper et al. (2025).

We now turn our focus to how elasticities of substitution vary across product departments. Table 4 presents elasticities of substitution results pooled to the department level

[19]Our instruments may violate the exclusion restriction if demand shocks are national in scope. As a robustness check, Table A.8 presents an uninstrumented analogue of Table 3 that includes brand or UPC fixed effects, so that identification comes from men and women facing the same price change for the same product in the same market. With the product held fixed, the assumption that demand shocks affect both genders equally is considerably more plausible, and the gender differential in the elasticity of substitution is consistently identified without instrumenting. These results also indicate that women are more price-elastic than men when purchasing the same brand or product.

in all of these specifications, we use both instruments as well as retailer fixed effects (in addition to our other market definition fixed effects).[20] We find that women are either more elastic than men are or are not significantly different than men in terms of elasticity across all departments. However, for some departments, namely Health & Beauty, Alcohol, and General Merchandise, we estimate less elastic demand with substantial noise. These departments have substantially less coverage in our data, as many purchases of these products can be made at stores that are not included in the NielsenIQ data. Further, these products tend to be bought less frequently, which can exacerbate sparsity issues. In contrast, for departments that see reliable coverage in the NielsenIQ data and do not face sparsity issues due to frequent purchases, like Dry Grocery, we estimate reasonable demand elasticities with precision.

The table shows that across almost all food products and general non-food merchandise women engage in substitution with significantly greater price elasticity than do men. We find that $\hat{\sigma}_w - \hat{\sigma}_m \in [-1.23, -0.49]$ among precisely estimated specifications. For half of the ten departments, including Health and Beauty, dry groceries, deli products, non-food groceries, and alcohol, we find no significant difference in elasticity of substitution between men and women.[21] The table also shows consistent results for price elasticity differentials, computed following Equation (9) and averaged within departments using expense weighting. The vast majority of purchases that constitute the retail consumption basket in the NielsenIQ data are food purchases, so our finding that women substitute more elastically applies to the bulk of the consumption basket. However, the majority of gendered products exist in non-food purchases, particularly Health and Beauty products. We interpret this finding as evidence that women demonstrate greater substitution elasticity across markets even with little explicit gendering. But, we cannot reject that women are *less* elastic in markets with significant gendering.

[20] Table A.9 displays this disaggregation without retailer fixed effects, yielding similar results.

[21] We find that Health and Beauty, Alcohol and General merchandise products tend to exhibit less price elastic demand than other departments. The finding that consumption of Health and Beauty products is more inelastic than that for other types of products is consistent with findings in Faber and Fally (2022) as well as with our findings in Section 6.

4.2 The competition channel

Up to this point, we have estimated elasticities of substitution, $\sigma(g)$. However, price elasticities of demand are given by Equation (9) and are a function of the elasticity of substitution and market shares. Under this derivation, price elasticities of demand will range from σ_g (in case of monopolistic competition where market shares approach zero) to 1 (in case of monopoly where the market share of the single good is 1). Because we have found that women generally substitute more elastically than men, the primary remaining channel for women to be less elastic on average as a consumer demographic is through lower market competitiveness for women's markets than for those of men. From Figure A.1 in the descriptive analysis, we know that women purchase a greater number of distinct products than do men by about 27%. This observation suggests *prima facie* that the markets women consume from may feature greater product diversity and competition than do those of men.

To further assess the competition channel, we study how brand and UPC market shares vary through the product-gender distribution as well as in men's and women's consumption baskets. First, we construct the market share of brand/UPC j in semester t as the aggregate expense observed in the HMS on j as a share of the overall national expense on the module of brand/UPC j in the same semester (weighted using NielsenIQ's projection weights):[22]

$$s_{jt} = \frac{\sum_{i \in \mathcal{I}} \sum_{l \in \mathcal{J}_{ij}} Expense_l \cdot NielsenWeight_i}{\sum_{i \in \mathcal{I}} \sum_{k \in G_j} \sum_{l \in \mathcal{J}_{ik}} Expense_l \cdot NielsenWeight_i}, \tag{10}$$

where $\mathcal{J}_{ik}$ is the set of transactions of individual i purchasing product k, indexed by l. We correlate this object with measures of product and consumer measures of gender.

To speak to the market share composition of individuals' consumption baskets, we assign each individual $\times$ half-year an expense-weighted average market share and assign each individual the mean of these values over half-years. We also assign each individual in the a gender-consumption rank based on their time-invariant percentile of expense-weighted UPC-

[22]We construct these measures in the HMS. Constructing these shares in the RMS could yield greater precision, but 1) our results demonstrate that the HMS weights yield considerable precision and 2) computing national-level market shares for all products using the RMS would be computationally infeasible due to the size of the RMS.

gender average.[23] Figure 3 plots the conditional distribution of expense-weighted average market share of goods consumed along gendered-consumption rank. In both panels, the left-most values correspond with individuals with the highest consumption share of "male" goods, whereas those on the right correspond with individuals with the highest consumption share of "female".

First, Panel (a) shows an inverse U-shape of raw expense-weighted average market share consumed along gendered consumption rank, which is expected since ungendered goods mechanically exhibit consumption from both men and women and would exhibit greater market shares. Second, observe that the female-end of the gendered consumption distribution features slightly lower average market share than the male end, suggesting the presence of greater levels of competition among female-gendered goods than ungendered goods and, to a lesser extent, male goods, indicating that female-male price differentials are not augmented by differences in competition environments. Panel (b) elaborates on this result, plotting average market shares demeaned against semester and module fixed effects, which controls for gendered sorting across modules as well as potential changes in competition over time. Interestingly, the graph demonstrates a negative relationship between market share and gendered consumption rank, albeit of significantly smaller magnitude than Panel (a). We interpret this graph to indicate that when controlling for the competitive environment across modules and over time, individuals consuming female goods (who are increasingly likely to be women by definition) consume products with lower market shares and therefore less monopolistic pricing power.[24]

We observe that women tend to consume goods with *lower* market shares and that female-gendered goods themselves also exhibit *lower* market shares. Abstracting from the role of pricing in the context of multiproduct firms, these results rule out a differential competition

[23] We construct the expense-weighted average UPC-gender consumed by individual i as $\frac{\sum_{j \in \mathcal{J}_i} Expense_j \cdot \hat{w}_j}{\sum_{j \in \mathcal{J}_i} Expense_j}$ for all transactions $\mathcal{J}_i$ of household i.

[24] Figure A.4 displays an analogous set of results, using brand gender itself as the variable on the x-axis. Panels (a) and (b) show that across the bulk of the distribution of brands, market shares are similar for men's and women's products. Panels (c) and (d) replicate Panels (a) and (b), weighting each brand × year observation by its respective module log expenditure amount. Notably, across all specifications, brands consumed nearly entirely by women exhibit the lowest market shares, reinforcing that the competition channel does not augment gender price differences.

channel in driving differences in price paid by men and women. We can conclude that women are more price elastic consumers than are men and that our average documented price differences paid by men and women—the pink tax—do not reflect differences in markups paid, but rather differences in marginal costs on average.

4.3 Differences in product quality along the gender distribution

Higher marginal costs could reflect greater product quality among female-targeted goods—a hypothesis we now investigate. Our model prescribes a method for estimating subjective and intrinsic product quality. Applying a logarithm to and rearranging Equation (7) yields an expression for subjective product quality of good j to good k in the same module n:

$$\log \frac{\varphi_j(g)}{\varphi_k(g)} = \left(\frac{1}{\sigma(g) - 1} \right) \cdot \log \frac{x_j(g)}{x_k(g)} - \log \frac{p_j(g)}{p_k(g)}. \tag{11}$$

As in Faber and Fally (2022), we let household quality evaluations depend on an intrinsic quality term $\log \phi_j$ and multiplicative term $\gamma_n(g)$ such that $\log \varphi_j(g) = \gamma_n(g) \log \phi_j$. Using a normalization such that within each module the multiplicative term $\gamma_n(g)$ averaged over genders equals one, the intrinsic quality term corresponds with the democratic average quality evaluation across households:

$$\log \phi_j = s_{female} \cdot \log \varphi_j(female) + (1 - s_{female}) \cdot \log \varphi_j(male) \approx \bar{\varphi}_j \tag{12}$$

for s_{female} corresponding with the share of female consumers in the entire market (typically around 0.5; i.e. not only consumers of the good).[25] In the case that a good is only purchased by a single gender, the intrinsic quality equals the subjective quality of the purchasing gender, as implicitly, goods that are never purchased by a given consumer gender correspond with a subjective quality valuation of $-\infty$ for that consumer gender. Following Faber and Fally (2022), we view the consumer gender that does not purchase a good at all as *not* comprising part of the consumer base that constitutes the democratic average.

[25] Note that in this setup, subjective quality-adjusted price can be written as $\log \frac{\tilde{p}_i(g)}{\tilde{p}_j(g)} := \log \frac{p_i}{p_j} - \log \frac{\varphi_i(g)}{\varphi_j(g)} = \frac{1}{1-\sigma(g)} \cdot \log \frac{x_i(g)}{x_j(g)}$ and intrinsic quality-adjusted price as $\log \frac{\tilde{p}_i}{\tilde{p}_j} := \log \frac{p_i}{p_j} - \log \frac{\phi_i}{\phi_j}$.

Interpretation of product quality. We view our quality measures as capturing components of both vertical quality and horizontal match value. In the CES framework, $\varphi_j(female)$ is a residual capturing all product appeal not explained by price, gender, and elasticity of substitution, and is therefore agnostic between vertical quality (an objective, universally-agreed-upon characteristic) and horizontal match value (idiosyncratic appeal that differs across consumer types). However, heterogeneous patterns in subjective quality across men and women would also be consistent with a horizontal match component of quality. This said, our framework only requires that products with higher intrinsic appeal command higher willingness-to-pay, which would be satisfied under either interpretation. As corroborating evidence that our estimated quality component capture at least a partially vertical component, Table A.10 documents a positive correlation between estimated intrinsic quality and USDA organic certification (which we observe in the NielsenIQ data), which we interpret as serving as an independently verifiable quality attribute with broad consensus across consumer types.[26]

Constructing brand quality estimates. We follow this structure to estimate brand-level intrinsic quality and quality-adjusted price for brands relative to their module-averages in the half year, as well as subjective quality for men and women separately. This process features several steps. First, we construct budget shares for each individual $\times$ brand $\times$ half-year, which we express relative to their own module-average budget shares, and apply a logarithm (per Equation (11)). We then average these values within each gender $\times$ brand $\times$ half-year (applying the appropriate NielsenIQ proprietary weights). Next, we assign each gender-department an elasticity of substitution from Table 4. Our main specification uses only departments with elasticities that are precisely estimated as greater than one in absolute value—a restriction we later relax to demonstrate robustness. We measure brand- and module-level unit prices as the ratio of aggregate expense to aggregate amounts purchased in a half-year. From here, we use Equation (11) to compute gender $\times$ brand $\times$ half-year

[26]Section B further explores gendered consumption of organic products, finding that the concentration of organic purchases increases sharply in women purchase share. Table B.1 shows that women are between 2.5 and 3 times as likely as men to purchase a USDA-certified organic product. Figure B.1 shows that within a module $\times$ market, the probability that a good is organic increases nearly monotonically in female purchase share intensity.

levels of subjective quality and related objects and Equation (12) to compute brand × half-year measures of "intrinsic" quality and related measures. We residualize these measures relative to half-year and module averages.[27]

Results. To speak to the quality-composition of individuals' consumption bundles, we compute for each individual × half-year their residualized expense-weighted quality consumed and quality-adjusted price consumed and assign each individual their mean over half-years (as in Section 4.2). We then plot values of expense-weighted average intrinsic quality and quality-adjusted price consumed along household gendered-consumption rank. Figure 4 plots the evolution of average quality consumed and quality-adjusted unit price faced along ranks of gendered consumption. The first panel shows a flat relationship between quality and gendered-consumption rank until around the 55th percentile of the gendered-consumption rank, at around 5% lower quality than module × half-year mean quality. Starting at this point, the relationship broadly monotonically increases to around 10% higher quality at the highest levels of gendered-consumption. We take this result to corroborate our finding that the average gender price differential is driven by marginal costs and quality rather than markups. Panel (b) elaborates on this result, plotting quality-adjusted price. The figure shows that, after accounting for product quality, individuals with the highest expense-concentration of female goods actually pay around 1% lower per unit relative to the average quality-unit × unit price within the module × half-year. We also observe slightly lower quality-adjusted prices paid by individuals with high expense-concentration of male goods.[28]

Figure 5 disaggregates this result by plotting female- and male-subjective quality valuations of goods along brand-gender. The result demonstrates two interesting asymmetries which likely drive our intrinsic quality results. First, women assess the quality of female

[27]We validate our quality measures by replicating Figure 2 in Faber and Fally (2022) in our setting. Figure A.5 shows a monotonically increasing relationship between average quality consumed and total individual expense (proxying for income, which is imprecisely measured in the NielsenIQ survey) and a decreasing relationship between quality-adjusted price and total individual expense. The magnitudes of our estimates are slightly smaller than those documented in Faber and Fally (2022), which we attribute to differences in our sample construction and our methodological approach (namely, we study individuals where Faber and Fally (2022) studies households, and we leverage the continuity of the HMS panel over time whereas Faber and Fally (2022) effectively treat their data as half-year cross sections.

[28]Figure A.6 replicates Figure 4 using brand-gender along the x-axis, demonstrating similar results.

goods higher than ungendered goods and male goods, both of which are assigned approximately equal quality valuations. Men, however, assess male goods most highly, followed by female goods, with ungendered goods valued substantially lower than both.[29] Second, the magnitude of the subjective quality valuations differ between men and women. Namely, Panel (a) shows that women assess the quality of women's goods at around 30% higher than ungendered and male goods, but men assess the quality of men's goods around 20% higher than ungendered goods and around 10% higher than female goods. These asymmetries allocate relatively greater quality-mass to female goods than do men's goods.[30]

To the extent that quality proxies for marginal cost (e.g., as in the supply-side model of Faber and Fally (2022)), these results indicate that women's consumer goods exhibit higher marginal costs. Combined with our findings that women demonstrate greater price sensitivity than men and do not consume from less competitive markets, this result implies that the pink tax is not driven by markups or demand-side factors, but rather reflects quality differences and gender-differential sorting into goods along these dimensions.

5 Wholesale Prices and Retail Markups

To speak to the markup channel more directly, we link our scanner data environment to data on wholesale prices from distributors to retailers from PriceTrak. These data consist of wholesale price information on the UPC-geography-year level from 2006-2011, from which we construct retailer markups.[31]

Although these data do not represent direct manufacturing or production costs, they allow us to directly observe a portion of the markup-setting process. Consider the double marginalization problem of a two-stage production-to-consumer setting in which an inte-

[29] We interpret men and women as most-highly valuing their respective goods as further validating our approach.

[30] To assess robustness to this criterion, Figure A.7 presents results including all departments with elasticities bounded at least one standard error above one (adding Meat), and Figure A.8 further includes departments with point estimates above one (adding Dairy and Non-food Groceries). Both sets of results are consistent with our main estimates.

[31] We align PriceTrak markets with the ScanTrack market codes used in NielsenIQ based on market name. We link 55 out of 67 PriceTrak markets; for the remaining 12 markets in PriceTrak that do not correspond with a ScanTrak market code, we use national-level wholesale prices (also reported by PriceTrak).

grated manufacturer-wholesaler produces and sells a product to a retailer that sells to a final consumer.[32] Let c represent the per-unit manufacturing cost of a good. The manufacturer sets a manufacturing markup μ^m and the retailer adds a markup μ^r so that the consumer pays a final unit price of $p = \mu^r \mu^m c$, which we observe in the NielsenIQ data. In this setting, the PriceTrak data specifically allow us to observe retailer cost $c^r = \mu^m c$ and impute retailer markups μ^r.

Our inference on gender differences in *retailer* markups $\mu^r_f - \mu^r_m$ will yield unbiased inference on gender differences in *overall* markups $\mu_f - \mu_m$ under the following condition:

$$\mathbb{E}[\log(\mu_f) - \log(\mu_m) \mid \log(\mu^r_f) - \log(\mu^r_m)] \approx \mathbb{E}\Big[\%\Delta c + \%\Delta \mu^m \mid \%\Delta \mu^r\Big] = 0 \qquad (13)$$

for a locally approximate proportion difference between women and men $\%\Delta x := \log(x_f) - \log(x_m)$. This assumption requires that, conditional on the observed gender difference in retailer markups, the gender differences in manufacturer cost $\%\Delta c$ and manufacturer markup $\%\Delta \mu^m$, which we cannot observe, introduce no additional difference in overall markups—i.e., all information content in the overall gender markup difference is captured by the retailer markup difference. We view imposing this assumption as largely innocuous: since the manufacturer's demand base aligns with the retailer's, the two margins of the double marginalization problem are unlikely to move in opposite directions in sufficient magnitude.

There are additional important caveats to using the PriceTrak data. First, these data only cover a subset of the NielsenIQ data. Within the 2006-2011 timeframe, only 9% of unique UPCs observed in the NielsenIQ data map to a PriceTrak wholesale price observation.[33] This matched subsample accounts for 37% of purchase volume we observe in the HMS panel during this timeframe.[34] Several UPCs have multiple observations on the UPC-geography-year level

[32] Our discussion is largely unimpacted by having distinct wholesaling and distribution entities. Section D explores the double marginalization problem in our context and further characterizes conditions under which inference on retailer markups generalizes to overall markups from manufacturer to final sale.

[33] We report the following match rate by department: 1) Health and beauty (5.5%), 2) Dry grocery (12.5%), 3) Frozen foods (14.2%), 4) Dairy (12.8%), 5) Deli (8.3%), 6) Packaged meats (16.5%), 7) Produce (3.3%), 8) Non-food grocery (10.3%), 9) Alcohol (0.2%), 10) General merchandise (2.7%). We exclude alcohol from this part of the analysis due to the low match rate.

[34] Table A.11 demonstrates balance of consumer gender based on whether UPC×market×year observations merge to the PriceTrak data. Around 14% of transactions map to the PriceTrak Data. Women are about 3% more likely than men to make a purchase that maps to the PriceTrak data (+0.4pp relative to a 14.2% baseline). Restricting to purchases made within the same year × module × county × retailer combination, observations that merge to the PriceTrak data are between 8 and 9.4% more expensive per unit than are

featuring multiple unique wholesale price values; in this case we use the lowest-observed per-unit wholesale price. Within this sample of NielsenIQ purchases that successfully matches to a wholesale price, around 8% of transactions exhibit a negative markup (i.e. observed wholesale prices greater than unit prices). We discard these goods with observed negative markups in this part of the analysis because we cannot apply a logarithm to these negative values; moreover, we cannot easily distinguish whether such observations reflect measurement error or loss-leader pricing.[35] Because we find little evidence of gender bias in PriceTrak coverage and because these data give us a unique insight into a component of the markup-setting process, we view these data as acutely informative in understanding the mechanisms underlying the pink tax.[36]

Results. Table 5 Panel (a) estimates average unit price differences paid by women and men analogous to Table 1 using the PriceTrak data. Columns (1) and (4) reproduce the least- and most-saturated specifications from Table 1 (Columns (1) and (6)). Columns (2) and (5) run these same specifications on the PriceTrak sample. Lastly, Columns (3) and (6) control for log wholesale price as observed in the PriceTrak data. The female coefficient in Column (3) is significant and negative, indicating that comparing goods in the same product module (and purchased in the same year), women actually pay a *lower* unit price than do men. I.e., conditioning on this measure of wholesale price, there is *no* pink tax. The coefficient in Column (6) is positive and statistically significant but very close to zero. The estimated coefficient represents an approximately 77% reduction in the gender-differential unit prices paid relative to as reported in Column (5)—after conditioning on location, age, race, retailer, and income demographic. The remaining gender price differential after conditioning on

those that do not merge. These relationships do not threaten the internal validity of our approach to using the PriceTrak data, but may have important external validity implications.

[35]We find no evidence of a differential presence of negative markups based on consumer gender. A transaction-level regression of a binary variable for the presence of a negative markup on a binary indicator for female purchaser gender yields a coefficient of 0.0007 (standard error 0.001, p-value 0.488) off of a male baseline of 0.079. Including module fixed effects yields a female dummy coefficient of 0.0016 (standard error 0.001, p-value 0.104).

[36]Another central limitation of the PriceTrak data is that they inform retailer costs only for the set of goods purchased from wholesalers. It is possible that other goods see other vertical production structures, including goods sold directly to consumers by manufacturers or goods sold from a manufacturer directly to a retailer. Goods produced in either of these cases would not be covered by PriceTrak data.

wholesale cost is small, but to assess whether it reflects markup differences, we turn to examining markups directly.

Figure 6 displays the coefficients of regressions analogous to Figure 2 projecting markups onto decile-bins of female purchase share (relative to the decile of UPCs with near gender parity in purchase share). Panel (a) shows the results of this regression with no differential weighting across UPCs. The figure follows a U-shaped pattern in female purchase share, where goods purchased at near-gender-parity see the lowest markups, and highly gendered goods see greater markups. More striking however, is that the markups exhibited by male-goods are significantly higher than female goods. At the extreme ends, goods purchased nearly exclusively by men see 40% higher markups than goods at near gender parity, whereas goods purchased nearly exclusively by women see only 30% higher markups. Panel (b) estimates this same regression while including analytic weights on the amount UPC expense recorded among HMS panelists (including NielsenIQ sample adjustments); the graph illustrates a similar shape in markup evolution, with even greater relative markups for male goods than female goods: weighting on UPC purchase volume, male goods see on average 60% higher markups than the gender-parity good, visibly increasing with male purchase intensity, whereas female goods see only a 30% greater markup and a much more shallow increase in female purchase intensity.[37]

Lastly, Table 5 Panel (b) displays the coefficients from several transaction-level regressions of imputed retail markups on indicators for female purchaser gender. The coefficients illustrate minimal difference in average markup faced by women and men.[38][39] This finding holds across departments, though with notable heterogeneity: perhaps surprisingly, women face significantly lower markups in Health and Beauty, the department most commonly as-

[37]Figure A.9 presents analogous results for retailer costs, showing that wholesale costs broadly increase with female purchase intensity, confirming that the price premium on women's goods reflects higher costs rather than higher markups.

[38]Table A.12 estimates an analogous set of regressions with analytic weighting for each observation equal to the consumer budget share represented by the specific good transaction. Alternating between significant negative coefficients and insignificant coefficients on the female dummy, these results imply that women either spend a lower proportion of their budgets on markups or that there is little difference in relative budget share allocation to markups per transaction.

[39]The table documents average markups of around 90%, smaller than those implied by Table 3; this observation can be at least partially reconciled by the fact that these estimates capture only the retailer markup relative to wholesale cost.

sociated with the pink tax (Table A.13). If markups do not account for the gender price differential, costs must. Indeed, Table A.14 shows that women face 3–5% higher wholesale costs across specifications under both equal transaction weighting (Panel (a)) and budget-share weighting (Panel (b)), with the gap widening in our most saturated specification. The pink tax, as observed in these data, reflects differences in costs rather than in markups.

6 Differentiated Products: Markups and Marginal Costs

Lastly, we turn to investigating marginal costs and markups in a more stylized market. Up to this point, our analysis has focused on capturing gender differences in the consumption and characteristics of goods across the consumption basket, but at the cost of model complexity in terms of flexible substitution patterns and market structure. Additionally, our CES demand analysis for Health and Beauty products showed that women may be less elastic than men for this specific subset of products, but the analysis was underpowered due to sparsity issues from infrequent purchases in the individual level data. To complement these prior analyses and address some of their limitations, we estimate markups and marginal costs using weekly store level data, allowing for significantly more model complexity and incorporating product ownership data at the cost of narrowing our focus to a single product market. Our weekly store level data solves our sparsity issues, but comes at the cost of no longer being able to attribute purchases to a specific gender. To overcome this limitation, we rely on our observed woman purchase share, $\hat{w}_j$ that we calculate using the individual level purchase data and map to the products in the weekly store level data. In addition to allowing for flexible substitution patterns, product ownership data will allow us to include firm level fixed effects that will allow us to compare men's and women's product markups for products made by the same manufacturer, directly testing for price discrimination.

We model demand in the market for disposable razors. We choose disposable razors because they have a high level of dispersion of $\hat{w}_j$ across their product space and because they are commonly referred to as the canonical pink tax item. Disposable razors exhibit near

complete gender segregation, and they also feature observable product characteristics.[40]

6.1 Differentiated Products Demand Model and Estimation

We follow the standard differentiated products market demand model presented in Berry, Levinsohn, and Pakes (1995) (BLP). Our main departure is that instead of typical product characteristics, we include our measure of the woman purchase share of a product, $\hat{w}_j$ and allow for heterogeneity in preferences for how gendered a product is. For the disposable razor product module, consider $t = 1, ..., T$ markets defined as a retail store $\times$ month combination each with $i = 1, ..., I_T$ customers. The indirect utility that customer i receives from choosing product j in market t is:

$$u_{ijt} = \beta_i \mathbf{x}_{jt} + \xi_{jt} + \epsilon_{ijt}, \tag{14}$$

where x_{jt} is a three dimensional vector of a constant term, the woman purchase share of product j and the price of product j in market t. $\xi_{jt} = \xi_{jr(t)} + \xi_{m(t)} + \Delta\xi_{jt}$ are product-retail chain fixed effects, month fixed effects, product fixed effects that capture unobservable product and brand characteristics, and ϵ_{ijt} is a mean-zero idiosyncratic error term that assumes a Type I Extreme Value distribution. The key deviation from our CES model or a logit demand system is that the three dimensional vector β_i of coefficients on prices, woman purchase share, and the outside option, are individual-specific coefficients.[41] We can parameterize these individual coefficients as a population mean preference parameter that is absorbed by the fixed effects and an individual random taste shock that captures unobserved heterogeneity in preference for the outside option and the woman purchase share

[40] Figure A.10 plots the histogram of woman purchase share for the sales of razors included in our sample for demand analysis. The distribution is bi-modal with the vast majority of razors either bought mostly by women or mostly by men.

[41] We define the outside option to include no purchases of any razor or shaving product as well as purchases of non-disposable razors and depilatories. In our data, men and women that have ever purchased a shaving product are similarly likely to have purchased a disposable razor (53% of men vs 61% of women). Men and women are equally likely to purchase a non-disposable razor (18%) while women are more likely to purchase a depilatory (20% vs 5%). Despite this observation, the men and women in our sample appear equally likely to substitute or switch product type purchases. The average man in our sample that purchases a disposable razor purchases from 1.85 product modules that are related to shaving and the average women purchases from 1.8 shaving-related product modules.

of the product:

$$\beta_i = \mathbf{\Sigma} \cdot \mathbf{v}_i, \qquad \mathbf{v}_i \sim N(0, \mathbf{I_3})$$

Heterogeneity in preferences for product gender will generate more flexible substitution patterns than our CES demand model. Under CES demand, price increases on a woman's razor will lead to equal levels of substitution from the women's razor into other women's razors and men's razors. Now, the random coefficient on woman purchase share will generate substitution patterns that feature men's razors substituting to men's razors and women's razors substituting to women's razors. Allowing for heterogeneity in preferences for the outside option is important as the value of the outside option likely differs between men and women in many of these markets. For example, the value of the outside option for disposable razors depends on the social stigma attached to shaving for men versus women.[42]

The resulting market share for product j in market t can be written as:

$$s_{jt} = \int \frac{\exp(\beta_i \mathbf{x}_{jt} + \xi_{jt})}{1 + \sum_k \exp(\beta_i \mathbf{x}_{kt} + \xi_{kt})} d\beta_i \tag{15}$$

We instrument for prices with the same instruments we use for our constant elasticity of substitution analysis: Hausman instruments that are a national level leave out mean of prices and DellaVigna-Gentzkow instruments that are a retail level leave out mean of prices. The Hausman instruments rely on the assumption that demand shocks are uncorrelated across markets while supply shocks are correlated across space and time. The DellaVigna-Gentzkow instrument's validity relies on retail chain level pricing being largely exogenous from local demand shocks. In addition to price instruments, we identify substitution patterns across products with quadratic differentiation instruments developed by Gandhi and Houde (2019). The instruments take the form $Z_{jt}^{diff} = \sum_k d_{jkt}^2$, where $d_{jkt} = x_{kt} - x_{jt}$ and x_{jt} is the woman purchase share of product j. We utilize two versions of this instrument: one with differences summed over products that are rivals; that is, products that are owned by other firms, and

[42]Many papers that estimate differentiated products demand models include demographic moments as in Nevo (2001), here we do not because our product characteristic is effectively a demographic moment and will be mechanically correlated.

one for products produced by the same firm. The instrument captures proximity in the product space in terms of woman purchase share and is rooted in the idea that substitution likely occurs among products that are similar in gender.

We estimate the model using the Python package, *pyBLP* (Conlon and Gortmaker (2020)), which solves for the parameters of interest using a two-step generalized method of moments. We estimate the model using the store level RMS data, which contains weekly observations of prices and quantities sold of each product offered. Because disposable razors constitute infrequent and storable purchases, we aggregate the data to the quarter level. We implement a series of restrictions to our sample to aid in estimation. We use data from 2013 and 2014[43] and keep only larger retail chains: those that have above median count of store-quarter observations and require that each store has observed razor purchases in each quarter or week of the data. In order to ameliorate sparsity issues, we drop stores that within the bottom quarter of unique product offerings as well as products within the bottom quartile of store coverage. We make a final sample restriction to ensure that our instruments have enough power for identification: we calculate the coefficient of variation of prices within the retail chain level and keep chains that have an above median coefficient of variation. These retail chains are those that put items on sale at higher frequency, which give our price instruments greater power.

We fit the supply side of the model by assuming firms, f, maximize their profits across the set of products they produce, $\mathcal{J}_f$ given the demand they face.[44]

$$\pi_{ft} = \sum_{j \in \mathcal{J}_f} (p_{jt} - mc_{jt}) s_{jt},$$

We construct a square ownership matrix of J products, Ω, that maps each product in our data to a common owner so that element jk is equal to 1 if product j and product k are owned by the same firm and 0 otherwise.[45] Let J be the matrix of estimated demand

[43] We restrict the data to two post-recession years for computational tractability purposes.

[44] We do not formally model the double marginalization that takes place in Consumer Packaged Goods. The reason for this is twofold: 1) many large CPG retailers self-distribute, so their products are only marked up once and 2) given that both wholesalers and retailers ultimately share the same demand base, it is unlikely that wholesale markups differ in direction from retailer markups.

[45] We construct the ownership matrix through manual search, Capital IQ, and newspaper articles. Although

derivatives, so that element jk is $\frac{\partial s_j}{\partial p_k}$. The price-cost markup is then given by:

$$\frac{p^* - mc}{p^*} = -(\Omega J)^{-1} \frac{s(p^*)}{p^*}, \tag{16}$$

Because we observe price, identified markups also identify marginal costs.[46]

6.2 Differentiated Products Demand Model Results

Table 6 shows average differences in prices, own-price elasticities, markups, and marginal costs between men's and women's disposable razors across products produced within the same firm. We define a razor as a women's razor if its woman purchase share is above 0.6, otherwise we classify the razor as a men's (or non-gendered) razor; this definition fits the bimodal distribution of razor female purchase share according to Figure A.10. We present results from a series of regressions on the product-store-quarter level of the outcome of interest on an indicator for whether the razor is a women's razor or a men's razor and include fixed effects for the store-quarter and firm; these regressions are weighted by the total sales volume for each product-store-quarter observation. We interpret the coefficient as the difference between the average men's razor and a woman's razor produced by the same firm purchased at the same store in the same time period.

The firm fixed effects allow us to directly test for second degree price discrimination by a firm through versioning. The majority of our work thus far has looked at differences in prices, elasticities, and markups between men and women across a product market. However, price discrimination is often thought of as a way for a firm to segment their market and profit maximize. In our CES analysis, we could not explore multi-product firm offerings, as the NielsenIQ data does not include parent company data for products. Restricting our analysis to a smaller product market, we can feasibly construct the product ownership matrix for disposable razors. This parent company data is contained within the ownership matrix, Ω, that captures multi-product firm's pricing incentives.

the NielsenIQ data contain information about brands, firms in the data frequently own multiple brands.

[46] Table A.15 presents the estimated parameters.

Within a firm, women's razors exhibit higher prices and marginal costs, have larger own-price elasticities and lower markups. The first column shows the difference in prices, which we observe in the data. Within a firm, women's razors are, on average, priced 35 cents higher per unit than men's in the data. The average per unit price of a men's razor in our data is 1.63, so women's razors are associated with about a 21% price premium. The second column presents the results for our estimated own-price elasticities which are given by $\varepsilon_{jt} = \frac{\partial s_{jt}}{\partial p_{jt}} \frac{p_{jt}}{s_{jt}}$. We find that women's razors have own price elasticities that are about 35 pp larger than men's razors produced by the same firm, which suggests that women are consuming razors more elastically. We find that the average estimated men's price elasticity in our sample is about -2, which is in line with estimates in the literature and yields reasonable estimates of marginal costs and markups. (Döpper et al. 2025) In the third column we present the difference in estimated marginal costs for men's and women's razors produced by the same firm, which are identified by solving the firm's profit maximization problem. We find that women's razors are associated with 35 cent higher marginal cost than men's on average. In column (4) we present our estimated markups, which are given $\frac{p-mc}{p}$. We find that women's razors have lower markups by about 12 percentage points. These results directly refute the common narrative around the pink tax in the media, that women are charged unfair markups on nearly identical products.[47]

Table A.16 shows the results excluding firm fixed effects. We find that across all razors, women's razors have higher prices, no significant differences in own price elasticities, higher marginal costs but also higher markups than men's razors. This result may seem counter-intuitive, as on average the difference in marginal costs fully offsets the difference in price. The finding that women's razors have higher markups is due to the non-linear relationship between markups and price along with the correlation structure of prices and marginal costs in the men's razors versus women's. The difference between the these results and the firm fixed effect results in Table 6 highlight the important role of sorting across firms. Overall, we take our results as evidence that although women's razors have higher markups on average, gendered price differences in razors do not seem to be driven by systemic price discrimination,

[47]Section C studies the product-level relationship between observable razor quality characteristics and "razor gender". Table C.1 shows that women's razors have more blades and are over 50% more likely to have an ergonomic handle, which we interpret to imply that women's razors are more costly to produce.

as women appear to be no differently elastic than men are.

7 Conclusion

We evaluate the existence of a "pink tax" on women's consumer goods relative to men's. We document a robust price premium on women's goods compared to men's goods in the same product category of 5.1% on average. Further corroborating this descriptive result, we find that within markets of similar goods, unit price increases nearly monotonically in women purchase share, relative to a gender-parity baseline. Simultaneously, we observe similar prices for men's goods as for goods purchased at gender-parity. Not only do we observe a consistent women's price premium of 21.2% on overtly gendered goods, but we also observe that women sort into purchasing less-overtly-gendered products with prices that are 4.8% higher, such as organic foods.

We proceed by studying the causal components of this pink tax. We distinguish three broad potential mechanisms at play: price elasticity of demand, competitive structure, and marginal costs. We estimate a CES model of demand and find that women as a consumer demographic are consistently *more* price elastic than are men. On average, women are approximately 15% more price sensitive than are men. We further adapt the model of Faber and Fally (2022) to our setting allowing us to estimate by-gender subjective product quality, intrinsic quality, and quality-adjusted price across goods. We find that women's consumer goods are assigned 25% greater intrinsic quality than goods consumed by both genders, whereas those consumed by men are assigned only 10% greater quality. Individuals with the greatest female-goods concentration of their consumption baskets feature an average quality-consumption level about 10% higher than the mean-product quality, whereas those with the greatest male-goods concentration of their consumption baskets see 5% lower average quality-consumption level than the mean-product quality. We also document that women's goods do not operate in less-competitive markets than those of men's, implying that a competition-channel is unlikely to generate greater markups for women.

We corroborate this central finding by implementing several additional designs that leverage complementary data and identification techniques. We combine our scanner data with

data from PriceTrak for a subset of our goods on wholesaler prices paid by retailers. Wholesale prices represent the cost of the product charged to the retailer. We construct retailer markups and observe that conditioning on wholesaler costs largely eliminates the observed pink tax; we also find no significant difference in retailer markups paid by men and women. Studying the goods themselves, we find slightly higher markups placed on men's goods than on women's goods: within market, the most-male goods face a 10pp higher markup than the most-female goods; when weighting by expense, this difference expands to 30pp.

Finally, we estimate overall markups and marginal costs of production for disposable razors using a differentiated products demand model (Berry, Levinsohn, and Pakes (1995)). We view this part of our analysis as complementary insofar as we can focus on one of the canonical examples of "pink tax goods", incorporate additional model complexity, and include firm fixed effects to directly test for price discrimination. We use this model to estimate product-market level demand elasticities, marginal costs, and markups. We find that women's razors have more elastic demand, higher marginal costs of production and lower markups when compared to men's razors produced by the same firm.

We conclude from our analysis a novel set of facts to frame the discussion of the pink tax: women pay around 5% more per unit than do men for goods in the same product market; when we study overtly gendered goods, this price difference rises to 21.2%. Taking consumption habits as fixed, the pink tax represents a *real* cost of living difference that exacerbates measures of the *nominal* gender wage gap by around 15-20% (Blau and Kahn (2017)). Contrary to popular discussion that attributes the pink tax to price discrimination, we find the pink tax is driven by women sorting into goods of higher marginal cost. However, it is almost certainly the case that preferences are *not* exogenous to gender; it is likely that the sorting processes we identify reflect societal expectations of women's and men's consumption behaviors in addition to personal taste. Nonetheless, our results suggest that recent and ongoing legislation aiming to prohibit price differences for gendered products is likely to prove ineffective at addressing price disparities between men and women—and that the forces generating those disparities are more deeply rooted in gendered consumption patterns and marginal costs than in gender- differential price discrimination.

References

Aguiar, Mark and Erik Hurst (2005). "Consumption versus Expenditure". In: *Journal of Political Economy* 113.5, pp. 919–948.

Allcott, Hunt, Benjamin B Lockwood, and Dmitry Taubinsky (2019). "Regressive sin taxes, with an application to the optimal soda tax". In: *The Quarterly Journal of Economics* 134.3, pp. 1557–1626.

Argente, David and Munseob Lee (2017). "Cost of living inequality during the great recession". In: *Kilts Center for Marketing at Chicago Booth–Nielsen Dataset Paper Series*, pp. 1–032.

Ayres, Ian and Peter Siegelman (1995). "Race and gender discrimination in bargaining for a new car". In: *The American Economic Review*, pp. 304–321.

Berry, Steven, James Levinsohn, and Ariel Pakes (1995). "Automobile prices in market equilibrium". In: *Econometrica: Journal of the Econometric Society*, pp. 841–890.

Blau, Francine D. and Lawrence M. Kahn (Sept. 2017). "The Gender Wage Gap: Extent, Trends, and Explanations". In: *Journal of Economic Literature* 55.3, pp. 789–865.

Bronnenberg, Bart J et al. (2015). "Do pharmacists buy Bayer? Informed shoppers and the brand premium". In: *The Quarterly Journal of Economics* 130.4, pp. 1669–1726.

Castillo, Marco et al. (2013). "Gender differences in bargaining outcomes: A field experiment on discrimination". In: *Journal of Public Economics* 99, pp. 35–48.

Conlon, Christopher and Jeff Gortmaker (2020). "Best practices for differentiated products demand estimation with pyblp". In: *The RAND Journal of Economics* 51.4, pp. 1108–1161.

DellaVigna, Stefano and Matthew Gentzkow (2019). "Uniform pricing in us retail chains". In: *The Quarterly Journal of Economics* 134.4, pp. 2011–2084.

Döpper, Hendrik et al. (2025). "Rising Markups and the Role of Consumer Preferences". In: *Journal of Political Economy* 133.8, pp. 2462–2505. DOI: 10.1086/735510.

Duesterhaus, Megan et al. (2011). "The cost of doing femininity: Gendered disparities in pricing of personal care products and services". In: *Gender Issues* 28.4, pp. 175–191.

Faber, Benjamin and Thibault Fally (2022). "Firm heterogeneity in consumption baskets: Evidence from home and store scanner data". In: *The Review of Economic Studies* 89.3, pp. 1420–1459.

Gandhi, Amit and Jean-François Houde (2019). "Measuring substitution patterns in differentiated-products industries". In: *NBER Working Paper* w26375.

Goldberg, Pinelopi Koujianou (1996). "Dealer price discrimination in new car purchases: Evidence from the consumer expenditure survey". In: *Journal of Political Economy* 104.3, pp. 622–654.

Guittar, Stephanie Gonzalez et al. (2022). "Beyond the pink tax: gender-based pricing and differentiation of personal care products". In: *Gender Issues* 39.1, pp. 1–23.

Hausman, Jerry (1999). "Cellular telephone, new products, and the CPI". In: *Journal of business & economic statistics* 17.2, pp. 188–194.

Hendel, Igal and Aviv Nevo (2013). "Intertemporal price discrimination in storable goods markets". In: *American Economic Review* 103.7, pp. 2722–51.

Jaravel, Xavier (2019). "The unequal gains from product innovations: Evidence from the us retail sector". In: *The Quarterly Journal of Economics* 134.2, pp. 715–783.

Manatis-Lornell, Athena J et al. (2019). "Gender-related cost discrepancies in a cohort of 110 facial moisturizers". In: *Journal of Cosmetic Dermatology* 18.6, pp. 1765–1766.

Manzano-Antón, Roberto, Gema Martinez-Navarro, and Diana Gavilan-Bouzas (2018). "Gender Identity, Consumption and Price Discrimination". In: *Revista Latina de Comunicación Social* 73, pp. 385–400.

Moretti, Enrico (2013). "Real Wage Inequality". In: *American Economic Journal: Applied Economics* 5.1, pp. 65–103.

Moshary, Sarah, Anna Tuchman, and Natasha Vajravelu (2023). "Gender-Based Pricing in Consumer Packaged Goods: A Pink Tax?" In: *Marketing Science* 0.0, null. eprint: `https://doi.org/10.1287/mksc.2023.1452`.

Nevo, Aviv (2001). "Measuring market power in the ready-to-eat cereal industry". In: *Econometrica* 69.2, pp. 307–342.

NYCDCA (2015). *From Cradle to Cane: The Cost of Being a Female Consumer*. Tech. rep.

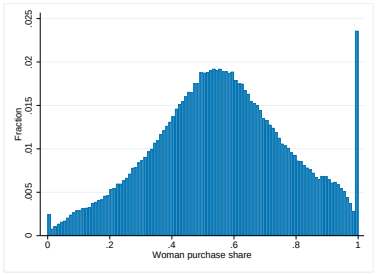
Stigler, George J. (1987). *The Theory of Price*. Macmillan Publishing Co., Inc.

Trégouët, Thomas (2015). "Gender-based price discrimination in matching markets". In: *International journal of industrial organization* 42, pp. 34–45.

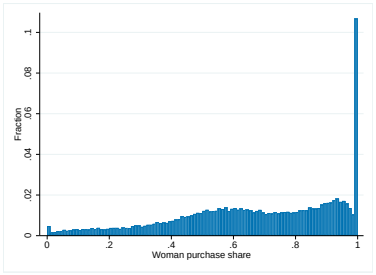
Ureña, Félix, Rodolfo Bernabéu, and Miguel Olmeda (2008). "Women, men and organic food: differences in their attitudes and willingness to pay. A Spanish case study". In: *international Journal of consumer Studies* 32.1, pp. 18–26.

Main exhibits

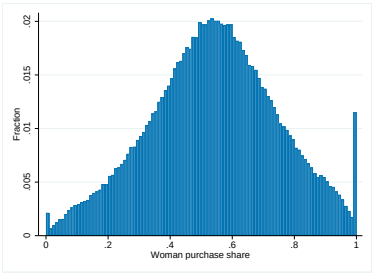
Figure 1: Distribution of Woman Purchase Share (UPC-gender) Across UPCs



(a): All Departments



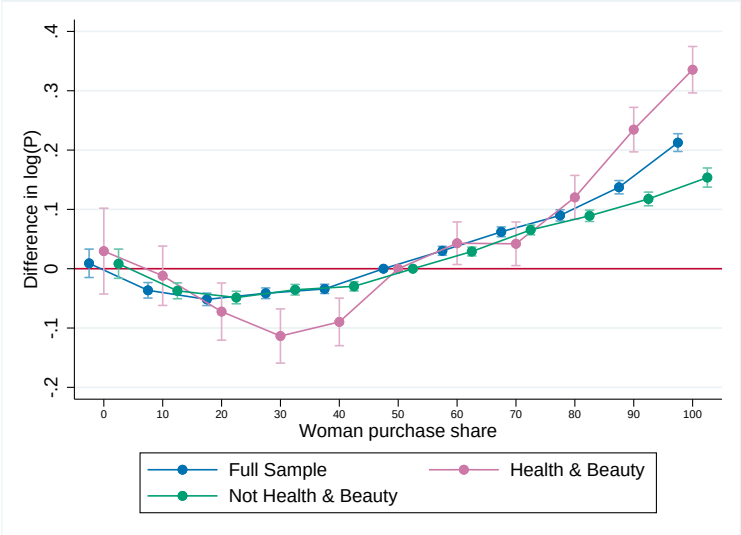
(b): Health and Beauty



(c): Non-Health and Beauty

Note: This figure plots a histogram of the share of times a UPC is bought by women. We restrict to UPCs that have above a varying cutoff number of purchases by unique individuals over the panel, with the cutoff number corresponding to 95% confidence that a product's true purchase share is within a 10 percentile-width bin centered around its observed share. This figure is referred to in Section 3.2.

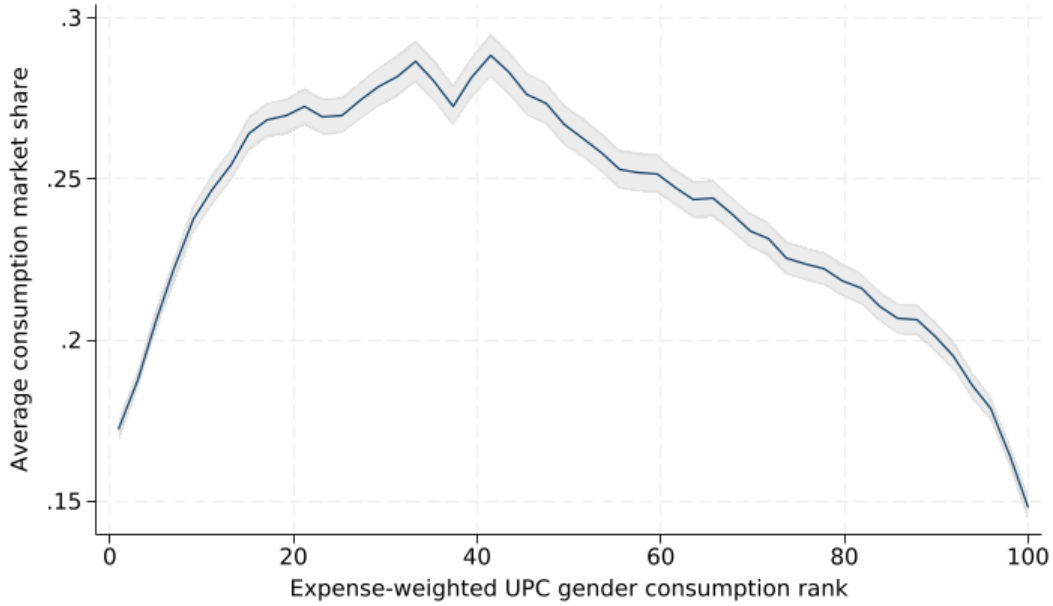
Figure 2: Prices of UPCs by Woman Purchase Share



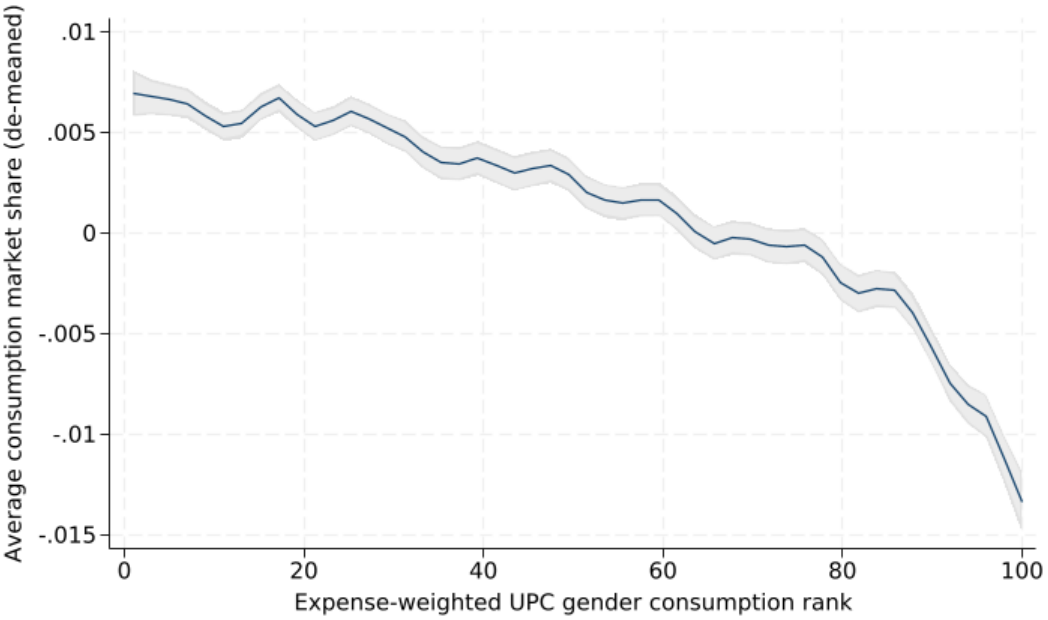
Note: This figure presents plots of the results of the regression $\log P_{jct} = \theta_{n_j ct} + \sum_{b \in \mathcal{B}} \gamma_b \mathbb{1}\{w_j \in Bin_b\} + \varepsilon_{jct}$. Bins $b \in \mathcal{B}$ include five-percentile-radius bins centered at values $5 + 10k$ and two bins for pure gender stratification at the tails partitioning the interval $[0, 1]$. The regression includes fixed effects for product module m_j, county c, and half-year t. Results are presented for the whole sample, also separating out Health and Beauty and Dry Grocery. Standard errors are clustered at the UPC-county level. This figure is referred to in Section 3.2.

Figure 3: Average market shares consumed by households ranked by gendered consumption

(a) Raw average market share



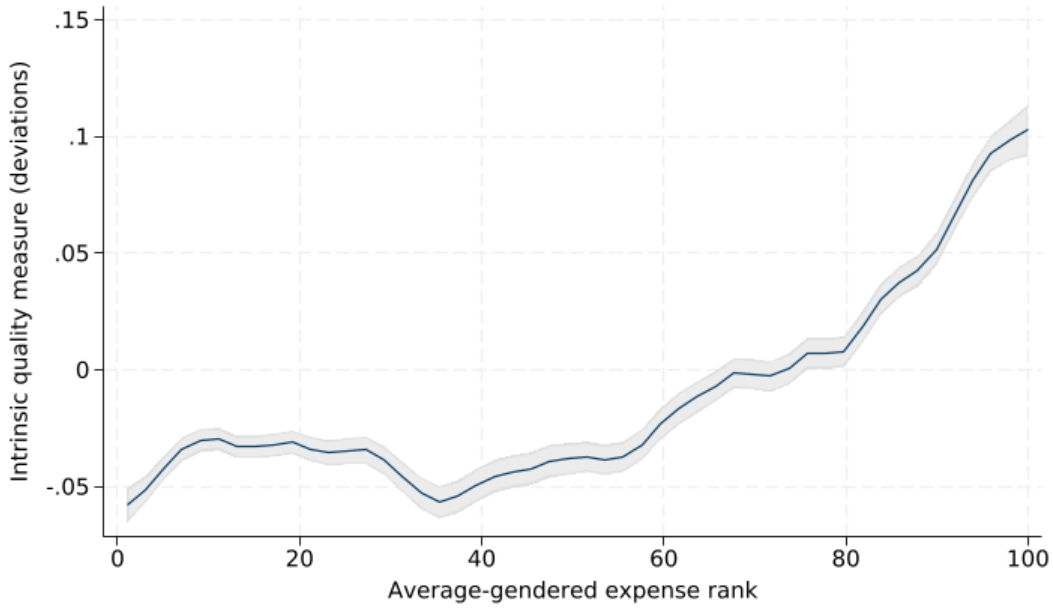
(b) Market shares residualized against module-average and half-year



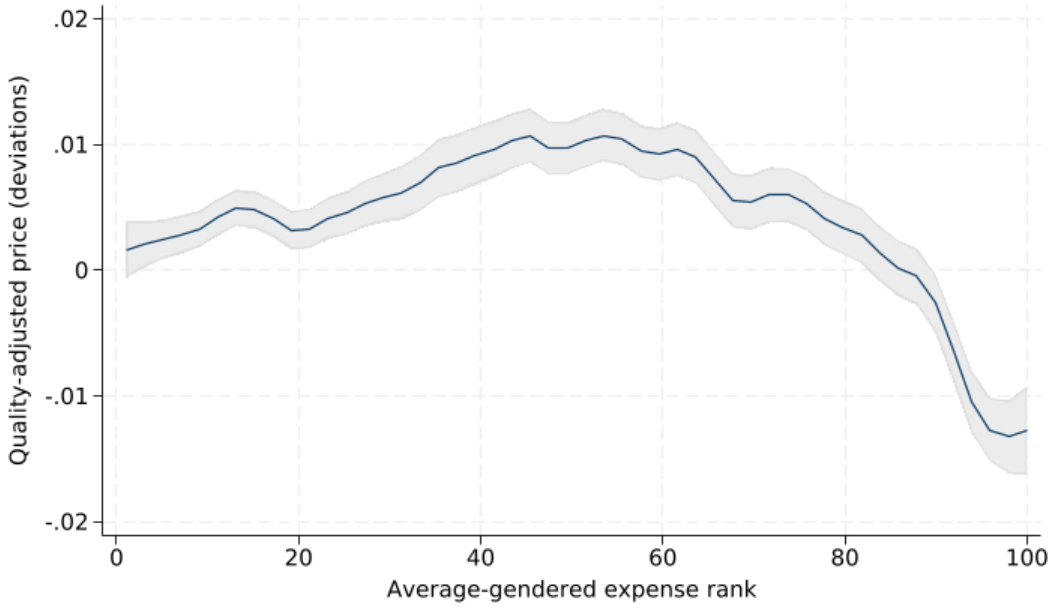
Note: These figures display average market shares consumed along percentile ranks of the distribution of gendered consumption. On the x-axis in both panels, the left-most values corresponds with individuals consuming the greatest share of most "male products" and the right-most values corresponds with individuals consuming the greatest share of most "female products". Panel (a) plots the average expense-weighted market share consumed as the dependent variable; Panel (b) plots the average of these market shares residualized against half-year and module fixed effects. Each observation is a household. The fitted relationships correspond to local polynomial regressions, and the shading represents 95% confidence intervals. This figure is referred to in Section 4.2.

Figure 4: Average intrinsic quality and quality-adjusted price
Households ranked by gendered consumption (male-to-female)

(a) Intrinsic quality



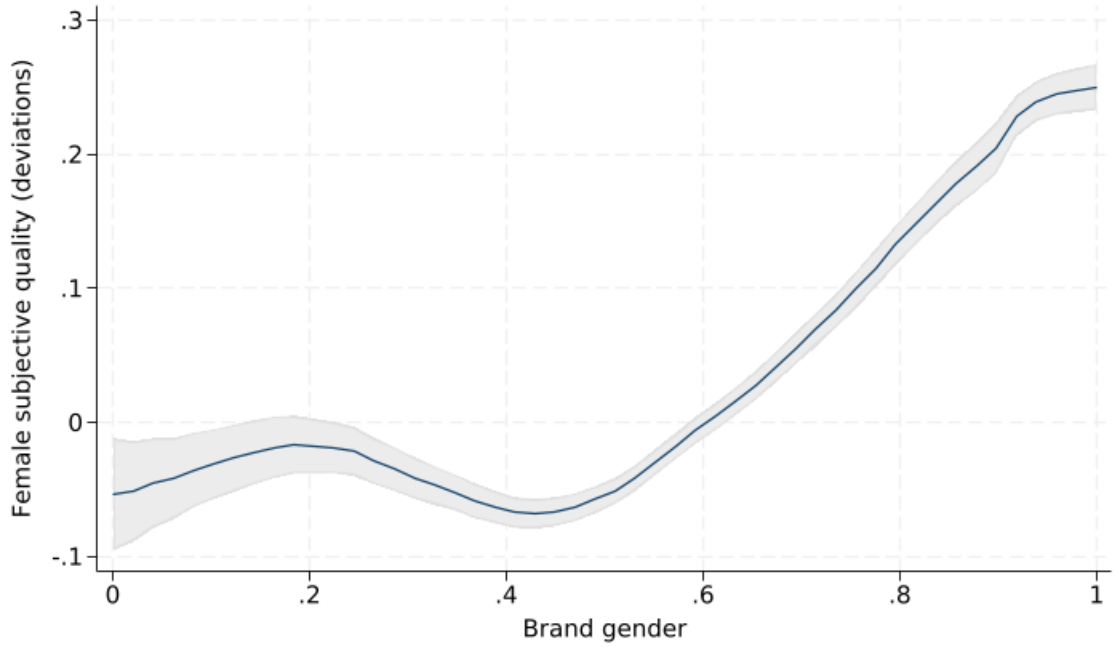
(b) Quality-adjusted price



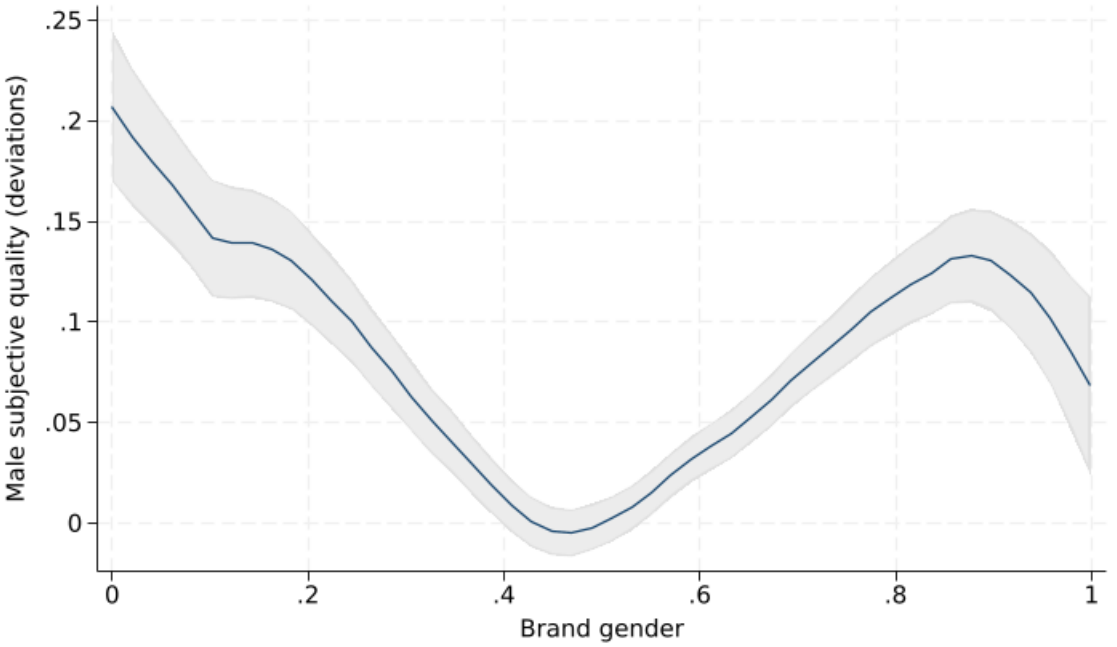
Note: These graphs plot intrinsic quality (Panel (a)) and quality-adjusted price (Panel (b)) along individual-ranks of gendered share of consumption. Each observation is a unique individual in the NielsenIQ HomeScan Panel. Individuals on the left have the highest expense-concentration of "male goods" in their consumption baskets; individuals on the right have the highest expense-concentration of "female goods" in their consumption baskets. Individual observations are weighted using NielsenIQ proprietary probability weights. Quality and quality-adjusted price are measured in units percent deviation relative to the mean good in each module × half-year. The fitted relationships correspond to local polynomial regressions, and the displayed confidence intervals are at the 95% level. This figure is referred to in Section 4.3.

Figure 5: Subjective quality valuations by brand-gender

(a) Female subjective quality valuations



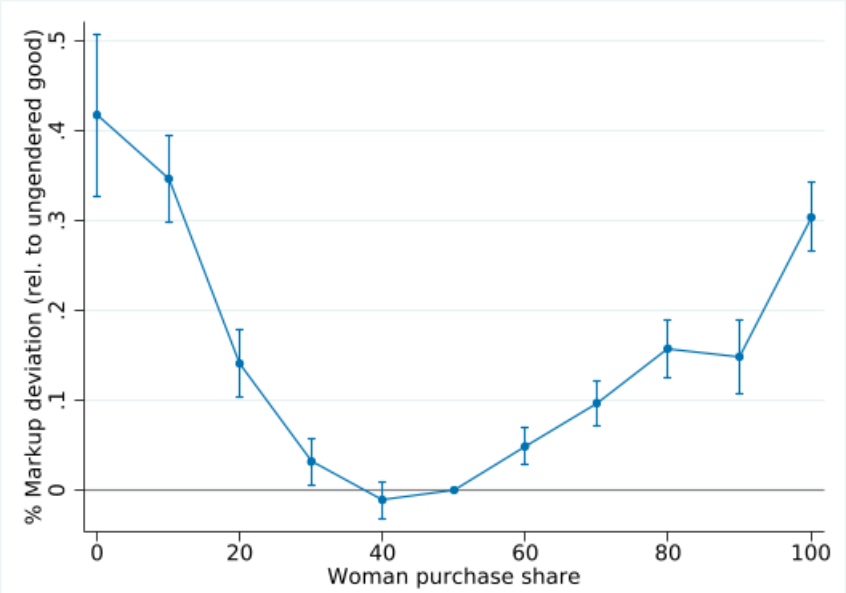
(b) Male subjective quality valuations



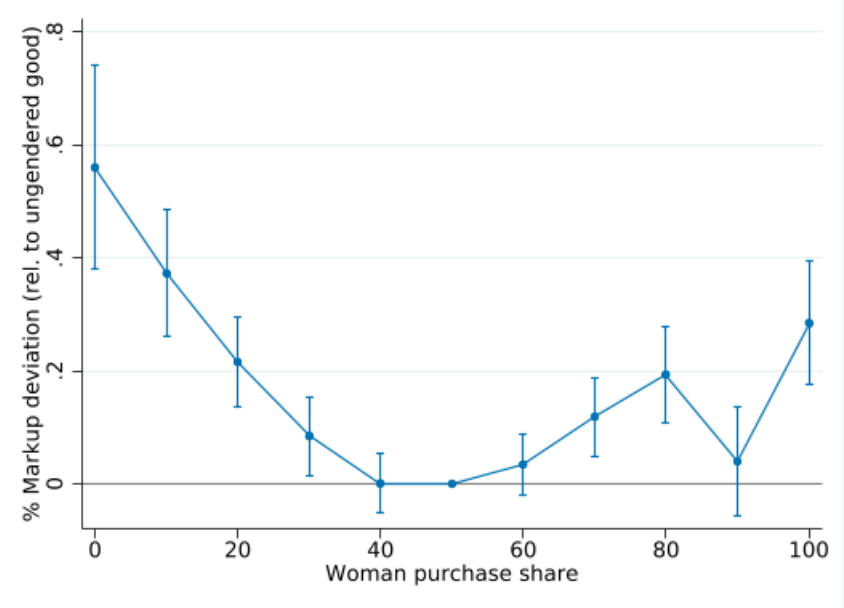
Note: These graphs plot subjective quality valuations by women and men (Panel (a) and Panel (b), respectively) along brand-gender. Brand gender is constructed as the expense-weighted mean UPC-gender of UPCs contained within the brand. The x-axis can be interpreted thus as the share of the consumption purchased by women (i.e. increasing in "female-ness" from left to right). Subjective quality is measured in units percent deviation relative to the mean good in each module × half-year, and constructed for men and women separately. The fitted relationships correspond to local polynomial regressions, and the displayed confidence intervals are at the 95% level. This figure is referred to in Section 4.3.

Figure 6: Log markup by woman purchase intensity

(a) Unweighted



(b) Weighted by expense



Note: These figures display the coefficients estimated from the following regression on the UPC-year level: $\log \mu_{u,t} = \alpha + \sum_{b \in \mathcal{B}} \gamma_b \mathbb{1}\{g_u \in Bin_b\} + \theta_{m,l,t} + \varepsilon_{u,m,c,l,t}$. Markup μ is constructed using PriceTrak data on wholesale prices and NielsenIQ final sale prices. Bins $b \in \mathcal{B}$ represent ten-percentile-width bins centered at multiples of 10 (truncated at 0 and 100) partitioning the interval $[0, 1]$; these bins reflect the aggregate amount of a UPC purchased by single women (as opposed to single men). The regression includes fixed effects for product module, county and half-year. Coefficients γ_b are estimated relative to goods in the same product module purchased at approximate gender-parity (between 45 and 55%). Panel (a) estimates this regression with equal weighting for all observations. Panel (b) presents the coefficients estimated from an analogous regression with analytic weights on UPC-year expenditure as recorded in the NielsenIQ HMS data. Standard errors are clustered at the UPC level. This figure is referred to in Section 5.

Table 1: Gender differences in log unit prices

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel (a): Unit prices in same product module							
Women	0.0230***	0.0299***	0.0514***	0.0545***	0.0489***	0.0381***	0.0501***
	(0.0035)	(0.0038)	(0.0037)	(0.0031)	(0.0025)	(0.0027)	(0.0026)
Men's Average	$1.329	$1.276	$1.276	$1.276	$1.276	$1.276	$1.276
Mod. X Units FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R-squared	0.83	0.84	0.84	0.88	0.91	0.92	0.92
N	153,333,383	97,715,726	97,715,726	97,715,726	97,715,726	97,715,726	97,715,726
Number of clusters	49,256	49,248	49,248	49,248	49,248	49,248	49,248
Panel (b): Unit prices for same product							
Women	-0.0089***	-0.0104***	-0.0030	-0.0060	-0.0120***	-0.0190***	-0.0139***
	(0.0017)	(0.0031)	(0.0031)	(0.0045)	(0.0034)	(0.0040)	(0.0039)
Men's Average	$1.293	$1.709	$1.709	$1.709	$1.709	$1.709	$1.709
UPC FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R-squared	0.95	0.95	0.95	0.96	0.97	0.97	0.97
N	151,184,396	44,114,794	44,114,794	44,114,794	44,114,794	44,114,794	44,114,794
Number of clusters	49,256	49,202	49,202	49,202	49,202	49,202	49,202
Common fixed effects sample	No	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	No	No
Month × Year FE	No	No	No	No	No	Yes	Yes
County FE	No	No	No	Yes	Yes	Yes	Yes
Retailer FE	No	No	No	No	Yes	Yes	Yes
Demographic FE	No	No	Yes	Yes	Yes	No	Yes

Individual level clustered standard errors in parentheses

* $p < .10$, ** $p < .05$, *** $p < .01$

Note: Panel A of this table presents estimates from the regression: $log(P_{ijt}) = \phi_{t(j)} + \beta Woman_i + \gamma X_i + \epsilon_{ijt}$ where P_{ijt} is the per-unit price of a UPC. $Woman_i$ indicates whether individual i is a woman, ϕ_t is a market-time fixed effect and X_i is a vector of demographic controls including income, county, age, race and education which we add in sequentially. Panel B of this table presents estimates from the regression: $log(P_{ijt}) = \phi_{jt} + \beta Woman_i + \gamma X_i + \epsilon_{ijt}$ where P_{ijt} is the per-unit price of a UPC. $Woman_i$ is an indicator for whether the individual is a woman, ϕ_t is a UPC-market-time fixed effect and X_i is a vector of demographic controls including income, county, age, race and education which we add in sequentially. Column 1 can be thought of as a raw gap between single men and single women within the same year, and each subsequent column demonstrates the contribution of controlling for an additional market or demographic factor. Note that the sample size in column (1) is larger than in the other columns, which use the sample estimated under the most restrictive fixed effect specification in column (7), where many singletons are dropped. All standard errors are clustered at the individual-level. This table is referred to in Section 3.1.

Table 2: Log unit prices by gender of product and consumer

	(1) All	(2) Health & Beauty	(3) Non-Health & Beauty
Women	0.0473***	0.0654***	0.0468***
	(0.0026)	(0.0078)	(0.0026)
Gendered Product	-0.0523***	-0.0027	-0.0571***
	(0.0065)	(0.0195)	(0.0068)
Women × Gendered Product	0.1452***	0.1220***	0.1424***
	(0.0076)	(0.0201)	(0.0085)
Men's Ungendered Average	$1.242	$4.012	$1.153
Module × unit × retailer × county × month FE	Yes	Yes	Yes
Demographic FE	Yes	Yes	Yes
Adj. R-squared	0.92	0.88	0.92
N	92,291,891	3,867,947	88,423,944
Number of clusters	49,247	47,492	49,243

Note: This table presents estimates from the regression: $log(P_{ijct}) = \phi_{jct} + \beta_1 Woman_i + \beta_2 GenderedProduct_j + \beta_3 Woman_i \cdot GenderedProduct_j + \gamma X_i + \epsilon_{ijct}$. ϕ_{jct} is a vector of fixed effects for the interaction of product module, units denomination, retailer chain, county, and month defined by product-location-time tuple (j, c, t). X_i includes with demographic controls for income, age, race and education. Gendered products are defined as UPCs purchased exclusively 90% or more (by amount) by one gender. Columns 2 and 3 separate out Health and Beauty products. Standard errors are clustered on the individual-level. This table is referred to in Section 3.2.

Table 3: Elasticities of Substitution

	OLS		Hausman IV		Retailer IV		Both IVs	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Substitution elasticity difference ($\sigma_m - \sigma_w$)	-0.00679	-0.0760***	-0.143***	-0.205**	-0.236***	-0.332***	-0.159***	-0.221***
	(0.00529)	(0.00941)	(0.0455)	(0.0922)	(0.0707)	(0.112)	(0.0417)	(0.0792)
Price coefficient ($1 - \sigma_m$)	0.289***	0.680***	-0.272***	-0.623***	-0.162**	-0.275***	-0.243***	-0.511***
	(0.00638)	(0.00828)	(0.0381)	(0.0814)	(0.0635)	(0.101)	(0.0356)	(0.0709)
% Elasticity dif. $(100 \cdot (\lvert\frac{\sigma_w}{\sigma_m}\rvert - 1))$	0.95	23.75	11.24	12.63	20.31	26.04	12.79	14.63
1st Stage F-Stat	-	-	1121	839.6	788.2	552.2	3173	2560
Module× gender×DMA× half-year FE	Yes	No	Yes	No	Yes	No	Yes	No
Mod.× gen. × retailer× DMA× HY FE	No	Yes	No	Yes	No	Yes	No	Yes
Observations	13,398,189	7,928,011	9,578,130	5,429,742	13,398,189	7,928,011	9,578,130	5,429,742
Number of clusters	1.02E+06	756,534	763,079	573,630	1.02E+06	756,534	763,079	573,630

* $p < .10$, ** $p < .05$, *** $p < .01$

Note: This table presents the results OLS and IV regressions on the product-consumer gender-retailer-Nielsen DMA-semester level of estimating average log budget share on log price changes, controlling for different fixed effects: $\Delta log(b_{gjt}) = (1 - \sigma_t(g))\Delta log(\bar{P}_{gjt}) + \eta_{gt} + \varepsilon_{gjt}$. Standard errors are clustered on the Brand-Nielsen DMA level. This table is referred to in Section 4.1.

Table 4: Elasticities of Substitution by Department

	(1) H&B	(2) Dry Groc.	(3) Frozen	(4) Dairy	(5) Deli	(6) Meat	(7) Produce	(8) Non-food Groc.	(9) Alcohol	(10) Gen. Merch.
Substitution elasticity difference ($\sigma_m - \sigma_w$)	0.0752	-0.105	-0.637**	-0.489**	0.638	-0.862**	-0.274*	-0.164	-0.612	-1.232**
	(0.458)	(0.126)	(0.284)	(0.190)	(0.749)	(0.370)	(0.145)	(0.470)	(2.796)	(0.598)
Price coefficient ($1 - \sigma_m$)	0.189	-0.661***	-0.901***	-0.0828	-1.766**	-0.418	-0.361**	-0.0412	1.668	0.795
	(0.395)	(0.109)	(0.237)	(0.169)	(0.693)	(0.312)	(0.152)	(0.446)	(2.763)	(0.568)
% Elasticity dif. ($100 \cdot (\lvert\frac{\sigma_w}{\sigma_m}\rvert - 1)$)	-9.27	6.32	33.51	45.16	-23.07	60.79	20.13	15.75	-91.62	600.98
Expense-weighted male price elasticity ($\lvert\varepsilon_m\rvert$)	0.866	1.548	1.623	1.057	1.593	1.335	1.196	1.035	-0.54	0.586
Expense-weighted female price elasticity ($\lvert\varepsilon_f\rvert$)	0.8	1.642	2.049	1.408	1.388	2.027	1.348	1.175	0.02	1.228
Observations	326,867	2,737,720	588,628	538,573	103,133	204,639	283,830	521,080	37,210	87,925
Number of clusters	58,072	268,531	62,751	38,296	11,819	17,077	27,068	62,713	8,527	18,661
MGDRT FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Hausman DMA IV	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Retailer IV	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

* $p < .10$, ** $p < .05$, *** $p < .01$

Note: This table presents the results OLS and IV regressions on the product-consumer gender-retailer-Nielsen DMA-semester level of estimating average log budget share on log price changes, controlling for different fixed effects: $\Delta log(b_{gjt}) = (1 - \sigma_t(g))\Delta log(\bar{P}_{gjt}) + \eta_{gt} + \varepsilon_{gjt}$. Each column stratifies the data by department; every column uses both Hausman and Retailer leave-out instruments in an instrumental variables regression. "MGDRT FE" refers to module×gender×Nielsen DMA×retailer×half-year fixed effects. Standard errors are clustered on the Brand-DMA level. This table is referred to in Section 4.1.

Table 5: PriceTrak prices, costs, and markups

Panel A: Log prices, controlling for retailer cost

	(1)	(2)	(3)	(4)	(5)	(6)
Women	0.0230***	0.0143***	-0.0071**	0.0501***	0.0394***	0.0092***
	(0.0035)	(0.0045)	(0.0034)	(0.0026)	(0.0039)	(0.0033)
Log wholesale cost			0.7600***			0.7041***
			(0.0018)			(0.0036)
Men's mean (levels USD)	$1.32	$0.43	$0.43	$1.28	$0.37	$0.37
PriceTrak sample	No	Yes	Yes	No	Yes	Yes
Module × Units	Yes	Yes	Yes	Yes	Yes	Yes
Demographics	No	No	No	Yes	Yes	Yes
County	No	No	No	Yes	Yes	Yes
Retailer	No	No	No	Yes	Yes	Yes
Year	Yes	Yes	Yes	No	No	No
Month	No	No	No	Yes	Yes	Yes
Adj. R-squared	0.83	0.75	0.86	0.92	0.88	0.92
N	153,333,383	17,901,305	17,901,305	97,715,726	9,690,298	9,690,298
Number of clusters	49,256	28,412	28,412	49,248	28,356	28,356

Panel B: Log markups, controlling for retailer cost

	(1)	(2)	(3)	(4)	(5)	(6)
Women	0.0086	0.0035	0.0090	0.0060*	0.0098***	0.0044
	(0.0073)	(0.0061)	(0.0061)	(0.0032)	(0.0029)	(0.0041)
Men's mean (percent markup)	89%	89%	89%	90%	91%	93%
Demographics	No	Yes	Yes	Yes	Yes	Yes
Module × Units	No	No	Yes	Yes	Yes	Yes
County	No	No	No	Yes	Yes	Yes
Retailer	No	No	No	No	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	No
Month × Year FE	No	No	No	No	No	Yes
Adj. R-squared	0.00	0.47	0.47	0.63	0.71	0.75
N	18,076,261	18,076,147	18,076,147	17,258,171	15,310,301	9,834,314
Number of clusters	28,412	28,412	28,412	28,406	28,393	28,357

This table presents estimates from transaction-level regressions uniquely identified by an individual i purchasing product j at time t. Panel (a) estimates the form: $log(P_{ijt}) = \phi_{t(j)} + \beta \mathbb{1}\{Woman_i = 1\} + \gamma C_{jt} + \Gamma X_i + \epsilon_{ijt}$ where P_{ijt} is the per-unit price of a UPC. $\mathbb{1}\{Woman_i = 1\}$ is an indicator for whether the individual is a woman, ϕ_t is a market-time fixed effect, C_{jt} is the wholesale price of UPC j in year t as observed in PriceTrak (included only in columns (3) and (6)), and X_i is a vector of demographic controls including income, county, age, race and education. Panel (b) estimates a similar set of regressions, however with log markup as the dependent variable. Each column restricts to the set of UPCs matching to the PriceTrak data and varies the level of fixed effects. Standard errors are clustered at the individual-level. This table is referred to in Section 5.

Table 6: Differentiated Products Model Estimates

	(1) Prices	(2) Own Elasticities	(3) Marginal Costs	(4) Markups
Women's razor indicator	0.3536***	-0.3526***	0.3510***	-0.1159***
	(0.0026)	(0.0027)	(0.0183)	(0.0033)
Observations	1,071,486	1,071,486	1,071,486	1,071,486
Adjusted R^2	0.094	0.195	-0.040	-0.028
Men's Mean	1.6317	-2.0207	0.3763	0.7340
Store×Quarter FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes

Standard errors in parentheses

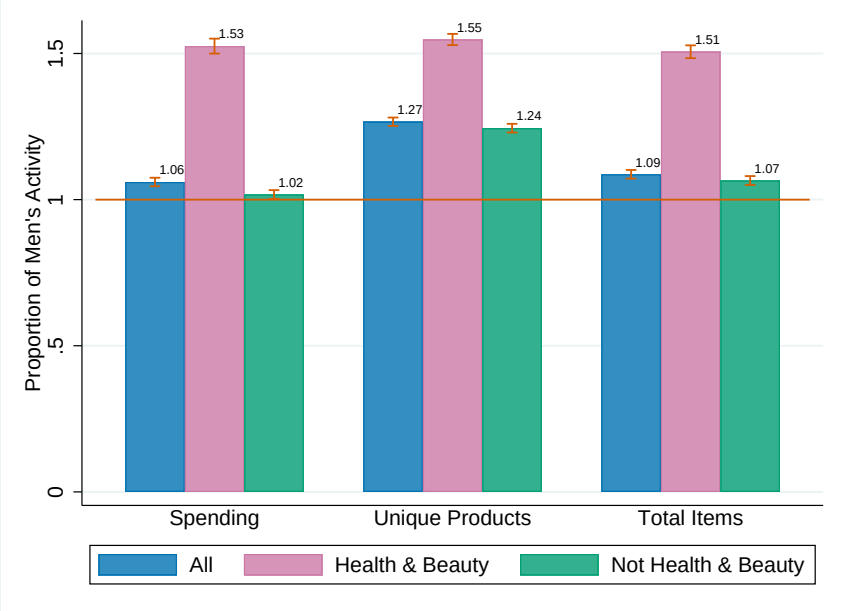
* $p < .10$, ** $p < .05$, *** $p < .01$

Note: This table displays regressions of prices and estimated price elasticities, marginal costs, and markups of disposable razors on a female-razor indicator. Female razors are defined as razors purchased at least 60% of the time by women, although 99% of razors see female purchase share above 80% or below 42%. Prices are directly observed in our data. Price elasticities, marginal costs, and markups are estimated from our differentiated products demand model described in Section 6. Standard errors take structural estimates as given. This table is referred to in Section 6.

Appendix A Additional exhibits

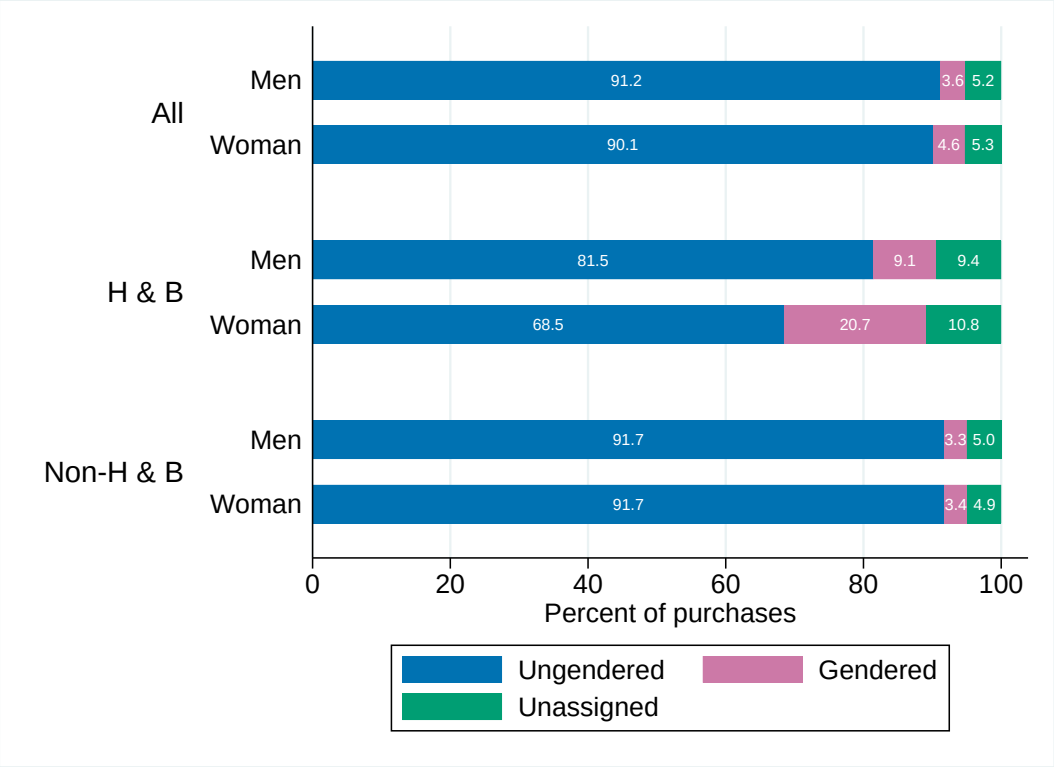
A.1 Additional descriptive results

Figure A.1: Women's yearly retail consumption spending relative to men's



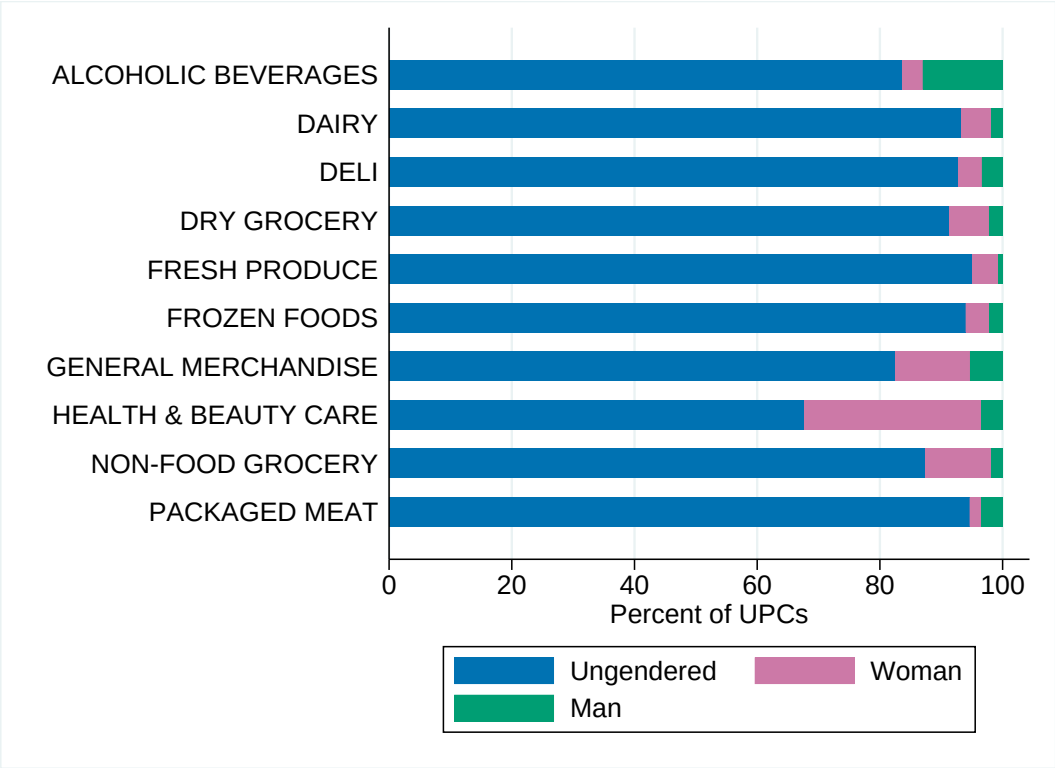
Note: This figure plots the coefficients estimated from a regression of log expenditure on an indicator for the individual being a woman and demographic controls: $\log y_{it} = \alpha + \beta \cdot \mathbb{1}\{Woman_i = 1\} + \Gamma X_{it} + \varepsilon_{it}$, for dependent variables including yearly spending, unique products purchased, and total items purchased. $\mathbb{1}\{Woman_i = 1\}$ is an indicator for whether the individual is a woman, and X_{it} is a vector of time- and time-id-varying controls including income, county, age, race and education. The height and standard errors of the bars correspond with the value $(1+\hat{\beta})$ estimated from each respective regression to emphasize the comparison to a gender parity standard at unity. Standard errors are clustered at the individual-level. This figure is referred to in Section 3.1.

Figure A.2: Consumption Basket Composition by Product Gender



Note: This figure presents plots the decomposition of purchases made by men and women into gendered, ungendered and unassigned products. The first rows show this for all product departments while the next two separate out health and beauty products. This figure is referred to in Section 3.2.

Figure A.3: Assigned UPC-gender Across Departments



Note: This figure plots the percentage distribution of UPCs assigned to Ungendered, Female, and Male within each department. We restrict to UPCs that are observed with great enough purchase frequency to be assigned a UPC-gender with false positive probability of 5%. This figure is referred to in Section 3.2.

Table A.1: Demographics of HMS panelists sample of single-member households

	Total	Women	Men	Difference
Income	44249	39178	50231	-11053.8**
	(36755.08)	(33639.89)	(39288.29)	(509.3439)
Age	53.39	53.12	53.71	-.593**
	(16.5003)	(17.2805)	(15.5241)	(.2221)
High school	0.600	0.636	0.558	.078**
	(.4899)	(.4812)	(.4966)	(.0063)
College	0.240	0.209	0.276	-.068**
	(.4269)	(.4064)	(.4472)	(.0054)
Post-grad	0.120	0.113	0.129	-.017**
	(.3253)	(.3163)	(.3355)	(.0041)
White	0.781	0.762	0.802	-.04**
	(.4138)	(.4258)	(.3982)	(.0055)
Black	0.135	0.158	0.107	.051**
	(.3417)	(.3651)	(.3097)	(.0044)
Asian	0.0260	0.0230	0.0280	-.005*
	(.1577)	(.1499)	(.1663)	(.0023)
Hispanic	0.0690	0.0700	0.0670	0.00200
	(.2531)	(.2549)	(.2508)	(.0038)
No. households	49,256	35,428	13,828	21,600

* $p < .10$, ** $p < .05$, *** $p < .01$

This table displays demographic data of men and women constituting single-member households as well as their differences. These figures and their corresponding gender-differences were computed using the proprietary analytic household weights included in the NielsenIQ Consumer Panel Survey. Dollar amounts are expressed in USD 2016. This table is referred to in Section 2.

Table A.2: Gender differences in shopping behavior

	Total	Women	Men	Difference	Difference (covariate-adj.)
Agg. per-year differences					
Log trips	4.473	4.447	4.503	-.057**	-.055**
	(.6531)	(.6333)	(.6744)	(.0086)	(.0085)
Log spending	7.742	7.749	7.733	.016**	.062**
	(.6311)	(.6201)	(.6436)	(.0083)	(.0082)
Spending inc. share	0.130	0.148	0.109	.039**	.008**
	(.2173)	(.2448)	(.1775)	(.0029)	(.0021)
Log purchases	6.229	6.271	6.180	.091**	.104**
	(.6455)	(.6213)	(.6695)	(.0086)	(.0085)
Log unique products	5.516	5.638	5.372	.266**	.269**
	(.6119)	(.5634)	(.635)	(.008)	(.008)
Log unique modules	4.346	4.480	4.188	.292**	.255**
	(.6882)	(.6397)	(.7094)	(.0089)	(.0082)
Log unique groups	3.257	3.355	3.141	.213**	.149**
	(.8171)	(.7999)	(.822)	(.0104)	(.009)
Log coupon value	3.060	3.269	2.798	.472**	.579**
	(1.9135)	(1.8137)	(2.0008)	(.0269)	(.0268)
Log coupon use	2.524	2.730	2.266	.464**	.551**
	(1.8058)	(1.7266)	(1.8686)	(.025)	(.0251)
Log deal use	3.859	4.070	3.604	.466**	.553**
	(1.89)	(1.7545)	(2.013)	(.0255)	(.0251)
Additional agg. differences					
Years in Panel	4.548	4.291	4.850	-.558**	-.124**
	(4.226)	(4.0326)	(4.4237)	(.0531)	(.0439)
Coupon value of spending	0.0600	0.0660	0.0530	.013**	.017**
	(.099)	(.1008)	(.0962)	(.0013)	(.0013)
Coupons per purchase	0.212	0.225	0.197	.027**	.039**
	(.2244)	(.2191)	(.2296)	(.0029)	(.0029)
Deals per purchase	0.0240	0.0270	0.0200	.007**	.008**
	(.0429)	(.0454)	(.0394)	(.0006)	(.0006)

* $p < .10$, ** $p < .05$, *** $p < .01$

Note: This table summarizes shopping behaviors of single-individual household NielsenIQ panelists per month and select aggregated characteristics. Columns "Total", "Women", and "Men" give unconditional (projection-weighted) averages for each respective group unconditional differences between genders. Columns "Difference" and "Difference (covariate-adjusted)" regress each respective row variable on an indicator for female consumer gender, with standard errors clustered on the individual level. The column "Difference (covariate-adjusted)" also controls for income and demographic characteristics. Monetary values are expressed in 2016 USD. This table is referred to in Section 3.1.

Table A.3: Prices paid across departments

	(1) H&B	(2) Dry Groc.	(3) Frozen	(4) Dairy	(5) Deli	(6) Pack. Meat	(7) Produce	(8) Non-food Groc.	(9) Alcohol	(10) Gen. Merch.
Panel A: Per unit prices within product module										
Women	0.0764***	0.0649***	0.0494***	0.0494***	0.0311***	0.0534***	0.0168***	0.0414***	0.0302***	-0.0094
	(0.0073)	(0.0029)	(0.0041)	(0.0033)	(0.0087)	(0.0052)	(0.0055)	(0.0044)	(0.0103)	(0.0171)
Men's Average	$3.86	$0.28	$1.12	$0.36	$3.56	$0.71	$1.57	$0.96	$1.93	$12.99
MURLM FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Demographic FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R-squared	0.88	0.87	0.92	0.94	0.90	0.85	0.78	0.91	0.99	0.86
N	4,443,135	44,889,271	10,084,043	12,196,328	4402975	2,742,710	8,416,509	5,688,481	1,299,644	3,546,033
Number of clusters	47,790	49,221	48,443	48,765	45,445	45,626	45,803	48,212	26,642	46,953
Panel B: Per unit price for same UPC										
Women	-0.0201	-0.0016	-0.0101***	0.0009	-0.0350***	-0.0156***	-0.0229***	-0.0110***	0.0053*	0.0300
	(0.0140)	(0.0016)	(0.0029)	(0.0015)	(0.0103)	(0.0038)	(0.0044)	(0.0021)	(0.0030)	(0.0207)
Men's Average	$3.95	$0.28	$1.21	$0.38	$3.79	$0.74	$1.45	$0.76	$2.36	$15.18
URLY FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Demographic FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R-squared	0.95	0.97	0.97	0.98	0.92	0.95	0.90	0.99	1.00	0.94
N	4,253,347	40,792,894	9,247,164	12,428,193	4,412,234	2,623,437	8,947,522	5,968,132	1,261,528	2,657,264
Number of clusters	46,854	49,208	48,089	48,718	44,770	44,809	46,444	47,919	23,278	44,421

* $p < .10$, ** $p < .05$, *** $p < .01$

Note: This table estimates $log(P_{ijt}) = \phi_t + \beta \mathbf{1}_{w(i)} + \gamma X_i + \epsilon_{ijt}$, stratifying by department across columns. P_{ijt} is the per-unit price of a UPC. $\mathbb{1}\{Woman_i = 1\}$ is an indicator for whether the individual is a woman and X_i is a vector of demographic controls including income, county, age, race. In panel A, ϕ_t is a vector of fixed effects for the interaction of product module, units, retailer chain, county, and half-year. In Panel B ϕ_t is a vector of fixed effects for the interaction of product (UPC), retailer chain, county, and half-year. "MURLM FE" refers to Module × Unit × Retailer × County × Month fixed effects; "URLY" refers to UPC × Retailer × County × Year. Standard errors are clustered at the household-level. This table is referred to in Section 3.1.

Table A.4: Differences in log unit prices paid within module
Including multi-pack controls

	(1)	(2)	(3)	(4)	(5)	(6)
Female	0.0234***	0.0473***	0.0515***	0.0469***	0.0383***	0.0503***
	(0.0034)	(0.0033)	(0.0028)	(0.0022)	(0.0027)	(0.0026)
Log multipack count	-0.0892***	-0.0920***	-0.0899***	-0.0598***	-0.0481***	-0.0487***
	(0.0018)	(0.0018)	(0.0014)	(0.0015)	(0.0021)	(0.0021)
Men's Average	$1.33	$1.33	$1.29	$1.23	$1.28	$1.28
Module × Units FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	No	No
Month X Year FE	No	No	No	No	Yes	Yes
County FE	No	No	Yes	Yes	Yes	Yes
Retailer FE	No	No	No	Yes	Yes	Yes
Demographic FE	No	Yes	Yes	Yes	No	Yes
Adj. R-squared	0.83	0.83	0.87	0.91	0.92	0.92
N	153,333,383	153,333,383	150,043,380	137,137,029	97,715,726	97,715,726
Number of clusters	49,256	49,256	49,256	49,252	49,248	49,248

Individual level clustered standard errors in parentheses

* $p < .10$, ** $p < .05$, *** $p < .01$

Note: This table presents estimates from the regression: $log(P_{ijkt}) = \phi_{t(j)} + \beta Woman_i + \delta \log Multi_k + \gamma X_i + \epsilon_{ijt}$ where P_{ijt} is the per-unit price of a UPC. $Woman_i$ indicates whether individual i is a woman, $Multi_k$ counts the number of units included (where non-multipacks are assigned a value of log(1)), ϕ_t is a market-time fixed effect and X_i is a vector of demographic controls including income, county, age, race and education which we add in sequentially. Standard errors are clustered on the household-level. This table is referred to in Section 3.1.

Table A.5: Differences in prices paid between men and women for products with identifiable UPC gender

	(1) Price per Unit	(2) Same UPC
Female	0.0487***	-0.0139***
	(0.0026)	(0.0039)
Men's Average	$1.24	$1.74
ModuleXUnits FE	Yes	No
UPC FE	No	Yes
Month FE	Yes	Yes
County FE	Yes	Yes
Retailer FE	Yes	Yes
Demographic FE	Yes	Yes
Adj. R-squared	0.92	0.97
N	92,291,891	42,702,383
Number of clusters	49,247	49,194

Individual level clustered standard errors in parentheses

* $p < .10$, ** $p < .05$, *** $p < .01$

Note: The table replicates the results from Table 1 Column (6) of Panels (a) and (b) using all products for which we can identify UPC gender following the procedure developed in Section 3.2. This table presents estimates from the regression $log(P_{ijt}) = \phi_{t(j)} + \beta Woman_i + \gamma X_i + \epsilon_{ijt}$ where P_{ijt} is the per-unit price of a UPC. $Woman_i$ indicates whether individual i is a woman, ϕ_t is a market-time fixed effect and X_i is a vector of demographic controls including income, county, age, race and education which we add in sequentially. Panel B of this table presents estimates from the regression: $log(P_{ijt}) = \phi_{jt} + \beta Woman_i + \gamma X_i + \epsilon_{ijt}$ where P_{ijt} is the per-unit price of a UPC. $Woman_i$ is an indicator for whether the individual is a woman, ϕ_t is a UPC-market-time fixed effect and X_i is a vector of demographic controls including income, county, age, race and education which we add in sequentially. All standard errors are clustered at the individual-level. This table is referred to in Section 3.2.

Table A.6: Unit prices in same product module by UPC and consumer gender, 75-25 Cutoff

	(1) All	(2) Health & Beauty	(3) Non-Health & Beauty
Female	0.0408***	0.0407***	0.0408***
	(0.0026)	(0.0083)	(0.0026)
Gendered UPC	-0.0186***	0.0455***	-0.0212***
	(0.0029)	(0.0090)	(0.0031)
Female × Gendered UPC	0.1177***	0.1402***	0.1144***
	(0.0034)	(0.0104)	(0.0036)
Men's Ungendered Average	$1.24	$4.01	$1.15
MURLM FE	Yes	Yes	Yes
Demographic FE	Yes	Yes	Yes
Adj. R-squared	0.92	0.88	0.92
N	92,291,891	3,867,947	88,423,944
Number of clusters	49,247	47,492	49,243

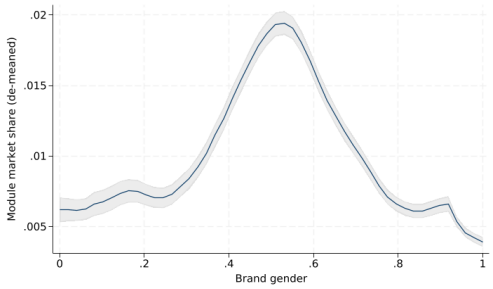
Individual level clustered standard errors in parentheses

* $p < .10$, ** $p < .05$, *** $p < .01$

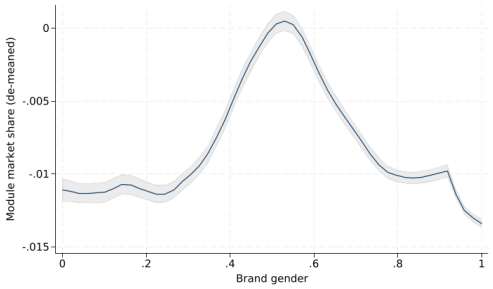
Note: This table presents estimates from the regression: $log(P_{ijt}) = \phi_{t(j)} + \beta_1 \mathbf{1}_{w(i)} + \beta_2 \mathbf{1}_{g(j)} + \beta_3 \mathbf{1}_{w(i)} \cdot \mathbf{1}_{g(j)} + \gamma X_i + \epsilon_{ijt}$. $\phi_{t(j)}$ is a vector of fixed effects for the interaction of product module, units denomination, retailer chain, county, and half-year. X_i includes with demographic controls for income, age, race and education. Columns 2 and 3 separate out Health and Beauty products. This table corresponds to table 6 in the paper but with the gendered product cutoff at 25-75 rather than 10-90. "MURLM FE" refers to module×unit×retailer×county×month fixed effects. This table is referred to in Section 3.2.

A.2 Additional CES exhibits

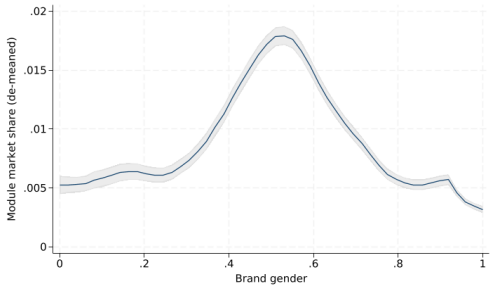
Figure A.4: Market shares by brand gender



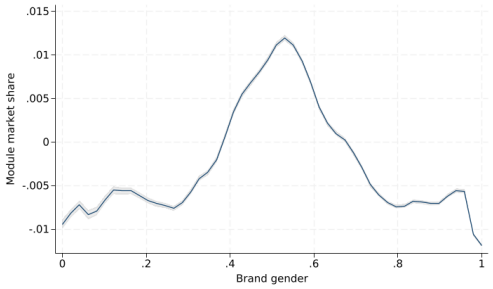
(a) Raw average market share



(b) Market shares residualized against module-averages and half-years



(c) Raw average market share: Log module expense weight



(d) Market shares residualized against module-averages and half-years: Log module expense weight

Note: these graphs display the average national market shares of brands within modules based on their observed brand gender. Each observation is a brand-semester between 2004 and 2019. Panel (a) shows raw average market shares; Panel (b) plots average market shares residualized within each half-year and module. Panels (c) and (d) replicate Panels (a) and (b), however weighting each brand × year observation by its respective module log expenditure amount. The fitted relationships correspond to local polynomial regressions, and the shading represents 95% confidence intervals. This figure is referred to in Section 4.2.

Table A.7: CES First stage and reduced form results of price change instruments

Panel A: First stage

	(1)	(2)	(3)	(4)	(5)	(6)
Hausman IV	0.130***	0.139***			0.112***	0.113***
	(0.00261)	(0.00219)			(0.00253)	(0.00203)
Retailer IV			0.233***	0.237***	0.252***	0.244***
			(0.00350)	(0.00425)	(0.00506)	(0.00617)
MGDT FE	Yes	No	Yes	No	Yes	No
MGDRT FE	No	Yes	No	Yes	No	Yes
Observations	9,578,130	5,429,742	13,398,189	7,928,011	9,578,130	5,429,742
F-Statistic	2,469	4,019	4,452	3,108	3,173	2,560
Adjusted R^2	0.272	0.473	0.24	0.447	0.274	0.476
Number of clusters	76,3079	573,630	1.02E+06	756,534	763,079	573,630

Panel B: Reduced form

	(1)	(2)	(3)	(4)	(5)	(6)
Hausman IV	-0.0326***	-0.0782***			-0.0304***	-0.0740***
	(0.00451)	(0.0100)			(0.00455)	(0.0103)
Female×Hausman IV	-0.0242***	-0.0437***			-0.0190***	-0.0340***
	(0.00565)	(0.0120)			(0.00573)	(0.0124)
Retailer IV			-0.0328**	-0.0555***	-0.0337*	-0.0436
			(0.0128)	(0.0203)	(0.0176)	(0.0290)
Female×Retailer IV			-0.0692***	-0.102***	-0.0636***	-0.0841**
			(0.0152)	(0.0238)	(0.0211)	(0.0338)
MGDT FE	Yes	No	Yes	No	Yes	No
MGDRT FE	No	Yes	No	Yes	No	Yes
Observations	9,578,130	5,429,742	13,398,189	7,928,011	9,578,130	5,429,742
Adjusted R^2	0.182	0.314	0.153	0.299	0.182	0.314
Number of clusters	763,079	573,630	1,015,000	756,534	763,079	573,630
F-Statistic	120.1	136.4	55.09	59.06	67.53	74.58

Brand-DMA level clustered standard errors in parentheses

* $p < .10$, ** $p < .05$, *** $p < .01$

Note: This table presents the first stage and reduced form results of the instrumental variables design using Hausman and retailer leave-out instruments to substitution elasticities of men and women. The first stage displayed in Panel A regresses within-market log price changes between half years on within-market changes in the values of the instruments between half years. "MGDT FE" refers to module×gender×Nielsen DMA×half-year fixed effects. "MGDRT" refers to module×gender×Nielsen DMA×retailer×half-year. This table is referred to in Section 4.1.

Table A.8: CES with Brand and UPC level fixed effects

	(1) Brands	(2) UPCs
$\sigma_m - \sigma_w$	-0.0414***	-0.2468***
	(0.0032)	(0.0069)
Observations	4,562,152	3,279,384
Adjusted R^2	0.026	0.037
BrandXTimeXDMAXRet FE	Yes	No
UPCXTimeXDMAXRet FE	No	Yes

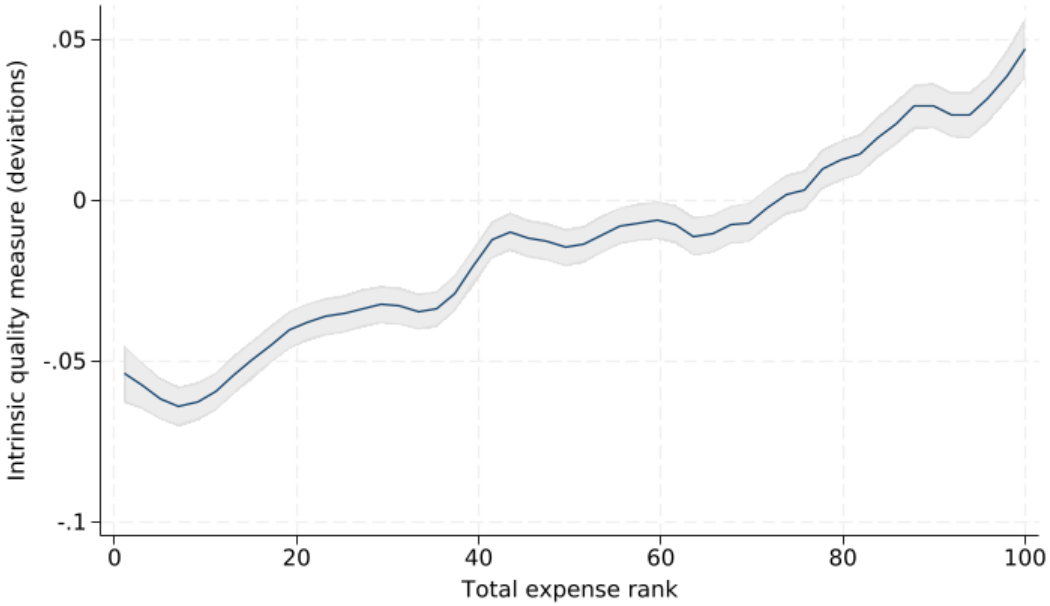
Brand-DMA and UPC-DMA level clustered standard errors in parentheses

* $p < .10$, ** $p < .05$, *** $p < .01$

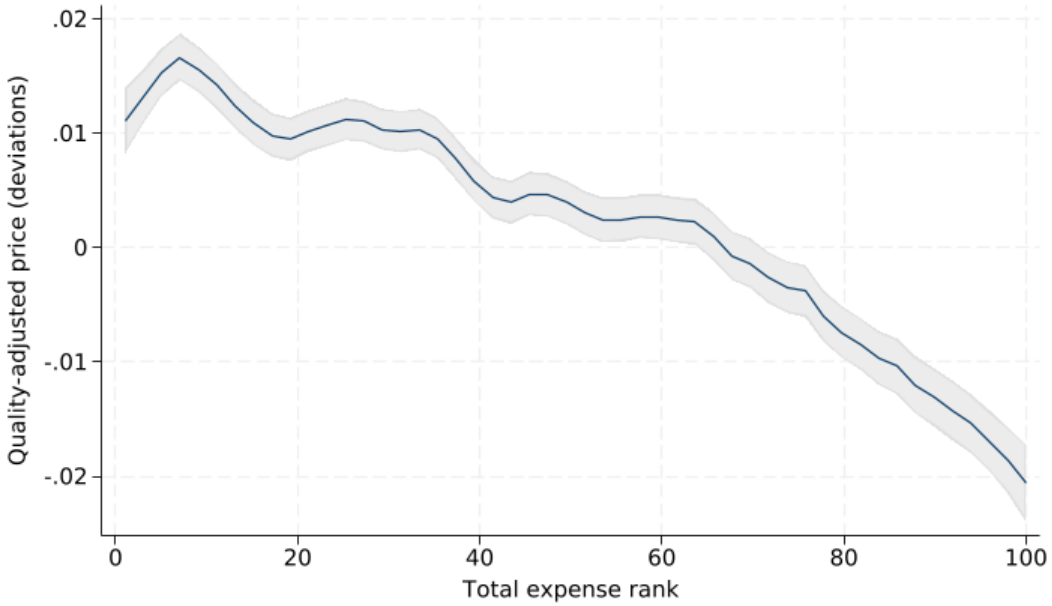
Note: This table presents results of univariate regressions similar to our main CES specification in Table 6 but includes fixed effects at the brand and UPC level rather than the module level. This specification identifies the difference in price responsiveness between men and women for purchases of the same brand or UPC in the same Nielsen DMA, retail chain and half year. This table is referred to in Section 4.1.

Figure A.5: Average intrinsic quality and quality-adjusted price
By total expense rank of individuals
Faber and Fally (2022) replication - validation of our approach

(a) Intrinsic quality



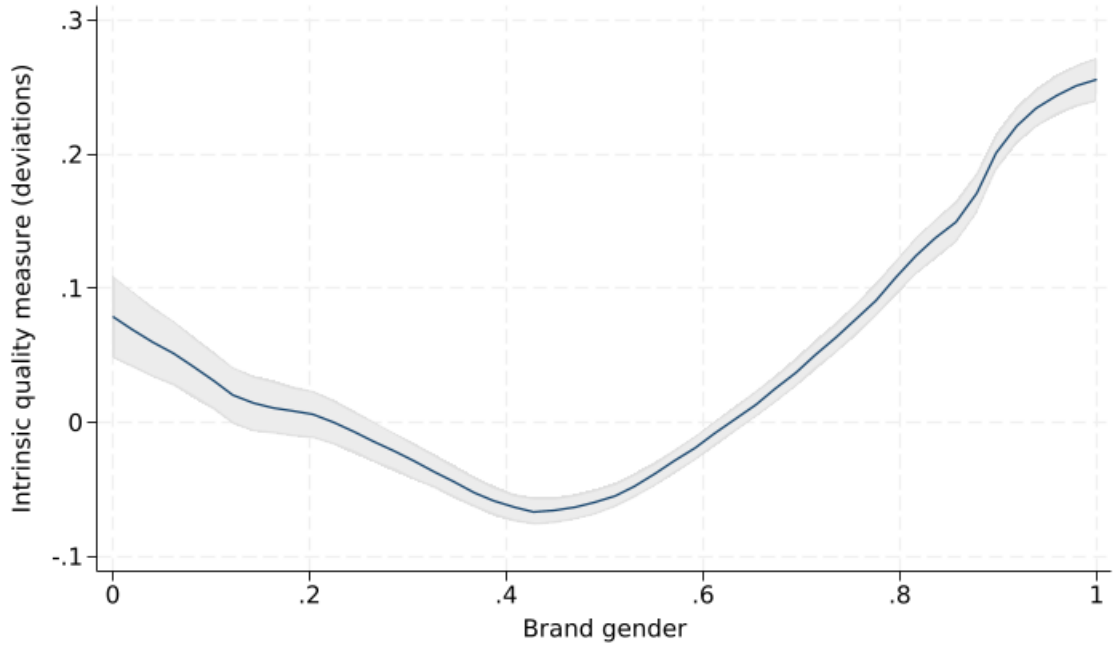
(b) Intrinsic quality-adjusted price



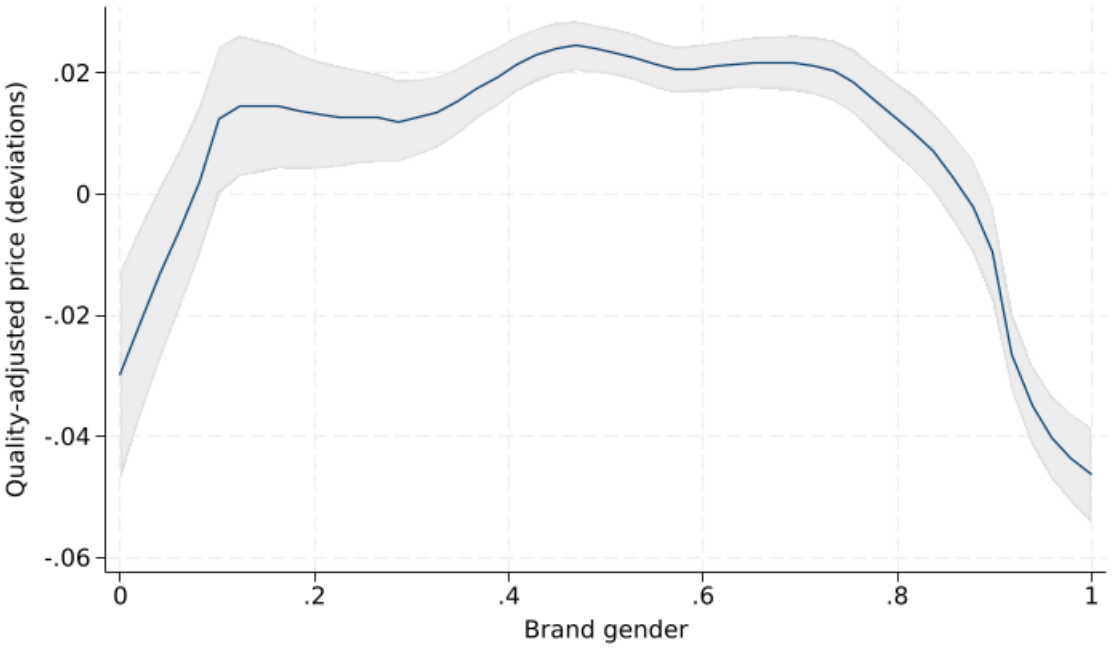
Note: These graphs replicate Figure 2 from Faber and Fally (2022) in our setting, validating the construction of our quality-measures. The graphs plot intrinsic quality (Panel (a)) and quality-adjusted price (Panel (b)) along household-expenditure ranks (weighted by NielsenIQ proprietary probability weights). Brand gender is constructed as the expense-weighted mean UPC-gender of UPCs contained within the brand. Quality and quality-adjusted price are measured in units percent deviation relative to the mean good in each module × half-year. The fitted relationships correspond to local polynomial regressions, and the shading represents 95% confidence intervals. This figure is referred to in Section 4.3.

Figure A.6: Average intrinsic quality and quality-adjusted price by brand-gender

(a) Intrinsic quality

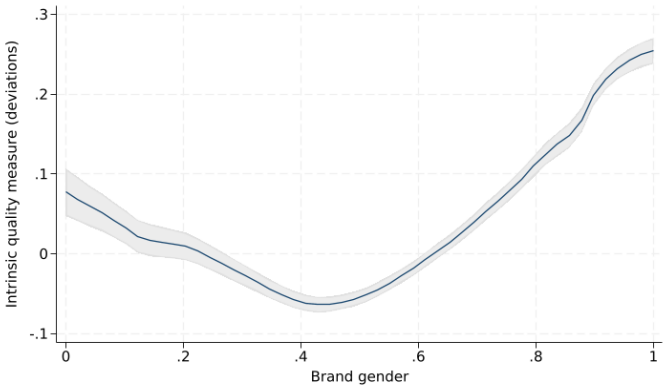


(b) Quality-adjusted price

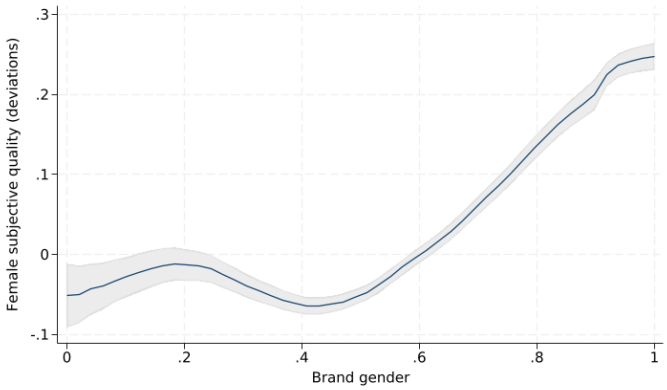


Note: These graphs plot intrinsic quality (Panel (a)) and quality-adjusted price (Panel (b)) along brand-gender. Each observation is a brand-semester. Brand gender is constructed as the expense-weighted mean UPC-gender of UPCs contained within the brand. The x-axis can be interpreted thus as the share of the consumption purchased by women (i.e. increasing in "female-ness" from left to right). Quality and quality-adjusted price are measured in units percent deviation relative to the mean good in each module × half-year. The fitted relationships correspond to local polynomial regressions, and the shading represents 95% confidence intervals. This figure is referred to in Section 4.3.

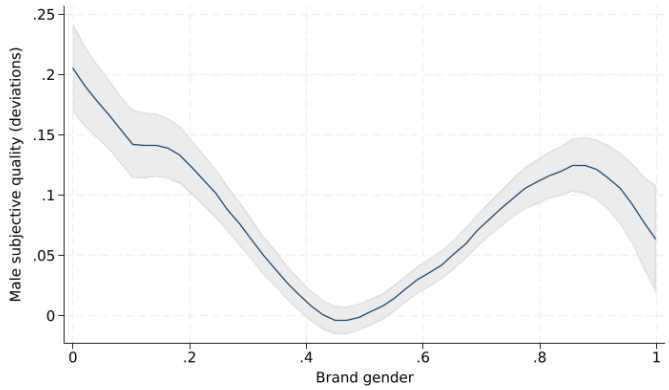
Figure A.7: Intrinsic and subjective quality by brand-gender
Robustness: Departments with elasticities of substitution bounded above one by at least one standard error



Panel (a) Intrinsic quality



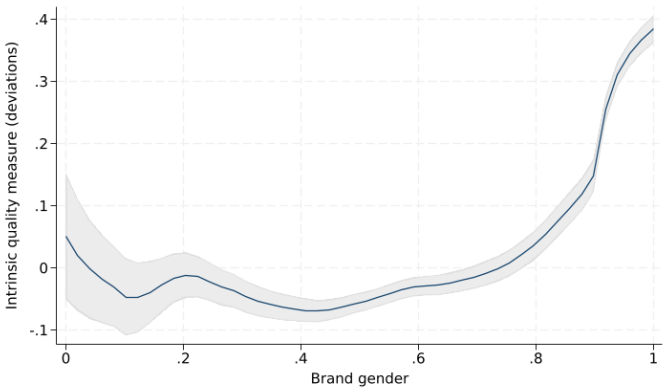
Panel (b) Female subjective quality valuations



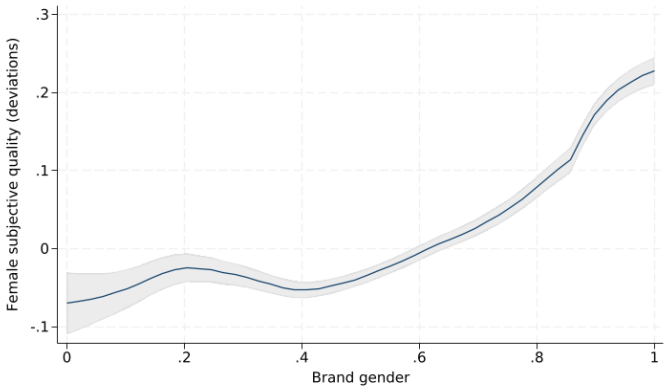
Panel (c) Male subjective quality valuations

Note: These graphs plot intrinsic quality (Panel (a)), and subjective quality valuations by women and men (Panel (b) and Panel (c), respectively) along brand-gender. Brand gender is constructed as the expense-weighted mean UPC-gender of UPCs contained within the brand. This graph replicates results from Section 4.3 however using all departments with point estimates for elasticity of substitution at least one standard error above one (see Table 4). The x-axis can be interpreted thus as the share of the consumption purchased by women (i.e. increasing in "female-ness" from left to right). Quality is measured in units percent deviation relative to the mean good in each module × half-year. The fitted relationships correspond to local polynomial regressions, and the shading represents 95% confidence intervals. This figure is referred to in Section 4.3.

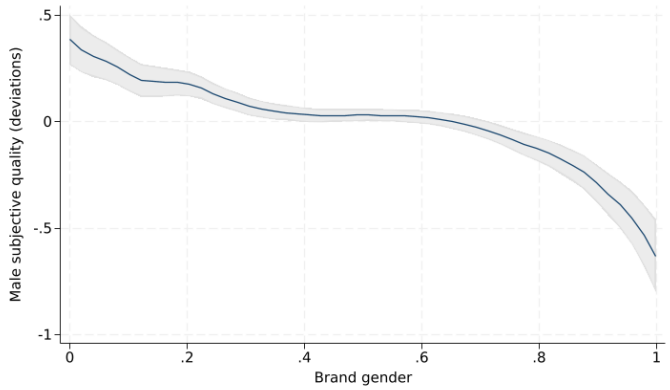
Figure A.8: Intrinsic and subjective quality by brand-gender
Robustness: Departments with elasticities of substitution greater than one



Panel (a) Intrinsic quality



Panel (b) Female subjective quality valuations



Panel (c) Male subjective quality valuations

Note: These graphs plot intrinsic quality (Panel (a)), and subjective quality valuations by women and men (Panel (b) and Panel (c), respectively) along brand-gender. Brand gender is constructed as the expense-weighted mean UPC-gender of UPCs contained within the brand. This graph replicates results from Section 4.3 however using all departments with point estimates for elasticity of substitution greater than one (see Table 4). The x-axis can be interpreted thus as the share of the consumption purchased by women (i.e. increasing in "female-ness" from left to right). Quality is measured in units percent deviation relative to the mean good in each module × half-year. The fitted relationships correspond to local polynomial regressions, and the shading represents 95% confidence intervals. This figure is referred to in Section 4.3.

Table A.9: Elasticities of Substitution by Department
No retailer fixed effects

	(1) H&B	(2) Dry Groc.	(3) Frozen	(4) Dairy	(5) Deli	(6) Meat	(7) Produce	(8) Non-food Groc.	(9) Alcohol	(10) Gen. Merch.
Substitution elasticity difference ($\sigma_m - \sigma_w$)	0.233	-0.141**	-0.446***	-0.185**	0.0148	-0.401*	-0.187*	0.0484	-0.0823	-0.682*
	(0.203)	(0.0715)	(0.168)	(0.0771)	(0.248)	(0.231)	(0.0953)	(0.162)	(0.67)	(0.389)
Price coefficient ($1 - \sigma_m$)	-0.11	-0.303***	-0.497***	-0.0766	-0.575***	-0.429**	-0.293***	-0.0982	0.565	0.471
	(0.127)	(0.0592)	(0.141)	(0.0649)	(0.218)	(0.207)	(0.0934)	(0.145)	(0.566)	(0.382)
% Elasticity dif. ($100 \cdot (\lvert\frac{\sigma_w}{\sigma_m}\rvert - 1)$)	-20.99	10.82	29.79	17.18	-0.94	28.06	14.46	-4.41	18.92	128.92
Observations	633,365	4,626,095	964,377	947,026	255,957	317,418	574,360	961,507	60,436	236,745
Number of clusters	82,060	351,177	78,704	48,601	15,561	21,475	33,362	90,969	11,303	29,481
MGDR FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Hausman DMA IV	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Retailer IV	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

* $p < .10$, ** $p < .05$, *** $p < .01$

Note: This table presents the results of estimating elasticities of substitution by regressing changes in the log budget share of a product on changes in log price for men and women controlling for the location, retail chain, and half-year corresponding to the following regression: $\Delta log(b_{gjt}) = (1-\sigma_t(g))\Delta log(\bar{P}_{gjt}) + \eta_{gt} + \varepsilon_{gjt}$. Each column stratifies the data by department; every column uses both Hausman and Retailer leave-out instruments in an instrumental variables regression. "MGDR FE" refers to module×gender×Nielsen DMA×half-year fixed effects. This table is referred to in Section 4.1.

Table A.10: Regression of estimated brand quality on organic indicator

	(1) Intrinsic quality	(2) Male-subjective quality valuation	(3) Female-subjective quality valuation
Organic	0.0796***	0.0627*	0.0715**
	(0.0301)	(0.0377)	(0.0289)
Fixed effects	Module	Module	Module
Adj. R-squared	0.04	0.01	0.04
N	2333	2333	2333

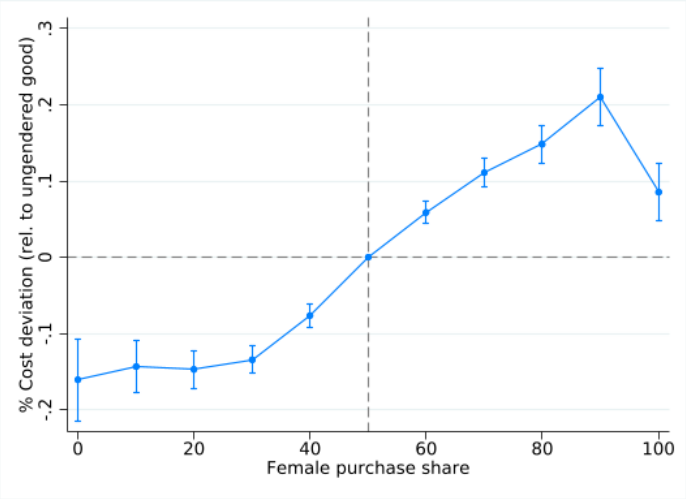
* $p < .10$, ** $p < .05$, *** $p < .01$

Note: This table presents estimates of univariate regressions of estimated product quality on indicators on the brand-level for organic certification from the US Department of Agriculture, which we observe in our data. These regressions consider brands within the Meat, Dairy, Deli, and Fresh Produce departments where organic indicators are best populated. The dependent variables, corresponding with log intrinsic and subjective quality measures, are estimated via the CES model procedure in Section 4 and is de-meaned with respect to product module average averages, so that a coefficient β yields the interpretation $(100 \cdot \beta)\%$ greater (or lesser) quality valuation than the module average. Heteroskedasticity-robust clustered standard errors are included in parentheses. This table is referred to in Section 4.3.

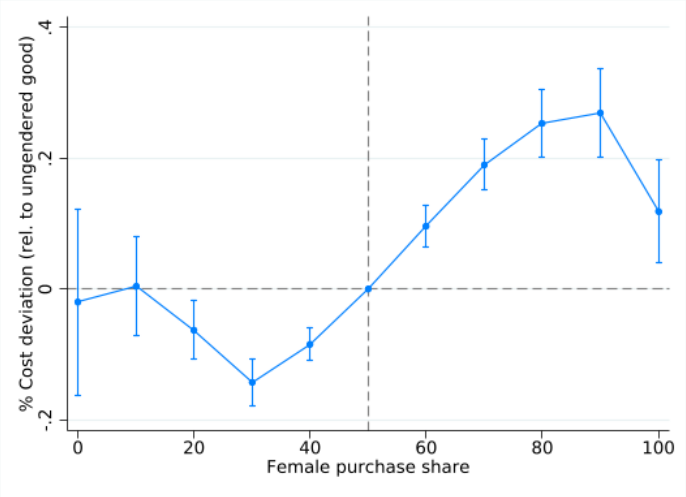
A.3 Additional results on retailer costs and markups

Figure A.9: Average retailer costs, within market, by female purchase intensity

(a) Equal weighting



(b) Expense-weighting



Note: These figures display the coefficients estimated from the following regression on the UPC-year level: $y_{u,t} = \alpha + \sum_{b \in \mathcal{B}} \gamma_b \mathbb{1}\{g_u \in Bin_b\} + \theta_{m,l,t} + \varepsilon_{u,m,c,l,t}$. Bins $b \in \mathcal{B}$ represent ten-percentile-width bins centered at multiples of 10 (truncated at 0 and 100) partitioning the interval $[0, 1]$; these bins reflect the aggregate amount of a UPC purchased by single women (as opposed to single men). The regression includes year fixed effects. Coefficients γ_b are estimated relative to goods in the same product module purchased at approximate gender-parity (between 45 and 55%). Panel (a) estimates uses log markup as the dependent variable; Panel (b) uses directly observed log retailer cost (distributor price) as the dependent variable. Both specifications feature equal weighting for all observations. Standard errors are clustered at the UPC level. This figure is referred to in Section 5.

Table A.11: Balance on PriceTrak merge

Panel (a): Regressing consumer gender on the merge indicator

	(1)	(2)	(3)	(4)	(5)	(6)
Female	0.0068***	0.0045***	0.0044***	0.0065***	0.0053***	0.0034***
	(0.0007)	(0.0003)	(0.0003)	(0.0015)	(0.0005)	(0.0003)
Constant	0.1402***	0.1414***	0.1421***	0.1112***	0.1118***	0.1154***
	(0.0026)	(0.0008)	(0.0004)	(0.0055)	(0.0007)	(0.0003)
Weight	Nielsen	Nielsen	Nielsen	Expense	Expense	Expense
Year	No	Yes	Yes	No	Yes	Yes
Module	No	Yes	Yes	No	Yes	Yes
County	No	No	Yes	No	No	Yes
Retailer	No	No	Yes	No	No	Yes
Adj. R-squared	0.000	0.503	0.721	0.000	0.515	0.777
N	1.55e+08	1.55e+08	1.38e+08	1.53e+08	1.53e+08	1.37e+08
Number of clusters	1,869,584	1,869,373	1,661,846	1,858,083	1,857,874	1,650,113

Panel (b): Regressing the merge indicator on transaction log unit price

	(1)	(2)	(3)	(4)	(5)	(6)
Merges to PriceTrak	-0.2121***	0.1352***	0.0894***	-0.7880***	0.1054***	0.0768***
	(0.0589)	(0.0050)	(0.0037)	(0.2029)	(0.0056)	(0.0035)
Constant	-1.4507***	-1.5006***	-1.5151***	-0.6812***	-0.7838***	-0.8536***
	(0.0581)	(0.0040)	(0.0034)	(0.2027)	(0.0030)	(0.0016)
Weight	Nielsen	Nielsen	Nielsen	Expense	Expense	Expense
Year	No	Yes	Yes	No	Yes	Yes
Module	No	Yes	Yes	No	Yes	Yes
County	No	No	Yes	No	No	Yes
Retailer	No	No	Yes	No	No	Yes
Adj. R-squared	0.002	0.829	0.907	0.014	0.897	0.959
N	1.53e+08	1.53e+08	1.37e+08	1.53e+08	1.53e+08	1.37e+08
Number of clusters	1,858,045	1,857,874	1,650,113	1,858,045	1,857,874	1,650,113

UPC-clustered standard errors in parentheses

* $p < .10$, ** $p < .05$, *** $p < .01$

Note: This table displays regressions pertaining to the balance of the PriceTrak merge. Panel (a) regresses an indicator for whether a UPC-market-year from the PriceTrak data merges onto the NielsenIQ data onto a female-consumer indicator. Panel (b) regresses log unit transaction price on the PriceTrak merge indicator. Each column varies the fixed effect and weighting specification. This table is referred to in Section 5.

Table A.12: Log markup by purchaser gender, budgetshare-weighted

	(1)	(2)	(3)	(4)	(5)	(6)
Female	-0.0294**	-0.0175	-0.0182	-0.0054	0.0014	-0.0012
	(0.0142)	(0.0121)	(0.0118)	(0.0045)	(0.0044)	(0.0057)
Male mean (levels)	89.3%	89.3%	89.3%	90.0%	91.5%	93.9%
Demographics	No	Yes	Yes	Yes	Yes	Yes
Module	No	No	Yes	Yes	Yes	Yes
County	No	No	No	Yes	Yes	Yes
Retailer	No	No	No	No	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
Month	No	No	No	No	No	Yes
Adj. R-squared	0.00	0.46	0.46	0.69	0.77	0.82
N	17,901,420	17,901,305	17,901,305	17,084,000	15,142,190	9,690,298
Number of clusters	28,412	28,412	28,412	28,406	28,393	28,356

Individual-level clustered standard errors in parentheses

* $p < .10$, ** $p < .05$, *** $p < .01$

Note: This table presents estimates from the transaction-level regression: $log(\mu_{ijt}) = \phi_{t(j)} + \beta \mathbf{1}_{w(i)} + \Gamma X_i + \epsilon_{ijt}$ where μ_{ijt} is the retailer markup UPC inferred from the PriceTrak and NielsenIQ data. $\mathbb{1}\{Woman_i = 1\}$ is an indicator for whether the individual is a woman, ϕ_t is a market-time fixed effect, and X_i is a vector of demographic controls including income, county, age, race and education. Observations are weighted by the transaction expense as a share of the individual's annual income. Standard errors are clustered at the individual-level. This table is referred to in Section 5.

Table A.13: Log markup by purchaser gender × department

	(1) H&B	(2) H&B	(3) Dry Groc.	(4) Dry Groc.	(5) Frozen	(6) Frozen	(7) Dairy	(8) Dairy	(9) Deli	(10) Deli
Female	-0.2652***	-0.0875***	-0.0054	0.0176***	0.0482***	0.0070	0.0087	-0.0202	0.0293	-0.0305***
	(0.0146)	(0.0178)	(0.0096)	(0.0052)	(0.0134)	(0.0095)	(0.0192)	(0.0155)	(0.0198)	(0.0111)
Male mean (levels)	73.9%	67.3%	71.35%	77.2%	31.7%	30.7%	352.1%	358.1%	44.5%	43.6%
Module	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Demographics	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
County	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Retailer	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Adj. R-squared	0.01	0.75	0.00	0.75	0.00	0.58	0.00	0.82	0.01	0.68
N	755,352	225,292	9,405,824	5,310,685	2,380,817	1,491,274	1,788,708	1,059,547	612,284	439,534
Number of clusters	27,572	21,102	28,395	28,257	27,794	26,300	27,933	26,398	25,155	21,780

Individual level clustered standard errors in parentheses

* $p < .10$, ** $p < .05$, *** $p < .01$

	(11) Pack. Meat	(12) Pack. Meat	(13) Produce	(14) Produce	(15) Non-food Groc.	(16) Non-food Groc.	(17) Gen. Merch.	(18) Gen. Merch.
Female	0.0203*	0.0051	0.0237	0.0239	-0.0437***	-0.0065	0.0493***	-0.0498**
	(0.0117)	(0.0106)	(0.0250)	(0.0342)	(0.0110)	(0.0117)	(0.0128)	(0.0242)
Male mean (levels)	39.1%	37.4%	272.4%	287.4%	47.1%	43.2%	102.8%	102.3%
Module	No	Yes	No	Yes	No	Yes	No	Yes
Demographics	No	Yes	No	Yes	No	Yes	No	Yes
County	No	Yes	No	Yes	No	Yes	No	Yes
Retailer	No	Yes	No	Yes	No	Yes	No	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month	No	Yes	No	Yes	No	Yes	No	Yes
Adj. R-squared	0.01	0.50	0.02	0.83	0.00	0.68	0.01	0.63
N	752,874	453,095	259,734	108,374	1,759,695	635,798	351,650	109,466
Number of clusters	26,190	23,200	22,113	14,437	28,147	25,031	26,403	17,796

Individual level clustered standard errors in parentheses

* $p < .10$, ** $p < .05$, *** $p < .01$

Note: This table presents estimates from the transaction-level regression: $log(\mu_{ijt}) = \phi_{t(j)} + \beta \mathbf{1}_{w(i)} + \Gamma X_i + \epsilon_{ijt}$ where μ_{ijt} is the retailer markup UPC inferred from the PriceTrak and NielsenIQ data. $\mathbb{1}\{Woman_i = 1\}$ is an indicator for whether the individual is a woman, ϕ_t is a market-time fixed effect, and X_i is a vector of demographic controls including income, county, age, race and education. Regressions are stratified by product department. Alcoholic products are largely not covered in the Pricetrak Data and so are excluded from this table. All transaction-observations are given equal weight. Standard errors are clustered at the individual-level. This table is referred to in Section 5.

Table A.14: Log retailer cost by purchaser gender

Panel (a): By transaction

	(1)	(2)	(3)	(4)	(5)	(6)
Female	0.0262***	0.0283***	0.0434***	0.0484***	0.0404***	0.0429***
	(0.0048)	(0.0028)	(0.0028)	(0.0021)	(0.0017)	(0.0022)
Male mean (levels)	$0.29	$0.29	$0.29	$0.28	$0.27	$0.26
Module	No	No	Yes	Yes	Yes	Yes
Demographics	No	Yes	Yes	Yes	Yes	Yes
County	No	No	No	Yes	Yes	Yes
Retailer	No	No	No	No	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
Month	No	No	No	No	No	Yes
Adj. R-squared	0.00	0.76	0.76	0.84	0.88	0.89
N	18,076,261	18,076,147	18,076,147	17,258,171	15,310,301	9,834,314
Number of clusters	28,412	28,412	28,412	28,406	28,393	28,357

Panel (b): Budgetshare-weighted

	(1)	(2)	(3)	(4)	(5)	(6)
Female	0.0361***	0.0274***	0.0344***	0.0461***	0.0433***	0.0471***
	(0.0081)	(0.0040)	(0.0041)	(0.0026)	(0.0022)	(0.0027)
Male mean (levels)	$0.43	$0.43	$0.43	$0.38	$0.34	$0.31
Module	No	No	Yes	Yes	Yes	Yes
Demographics	No	Yes	Yes	Yes	Yes	Yes
County	No	No	No	Yes	Yes	Yes
Retailer	No	No	No	No	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
Month	No	No	No	No	No	Yes
Adj. R-squared	0.00	0.79	0.79	0.88	0.91	0.93
N	17,901,420	17,901,305	17,901,305	17,084,000	15,142,190	9,690,298
Number of clusters	28,412	28,412	28,412	28,406	28,393	28,356

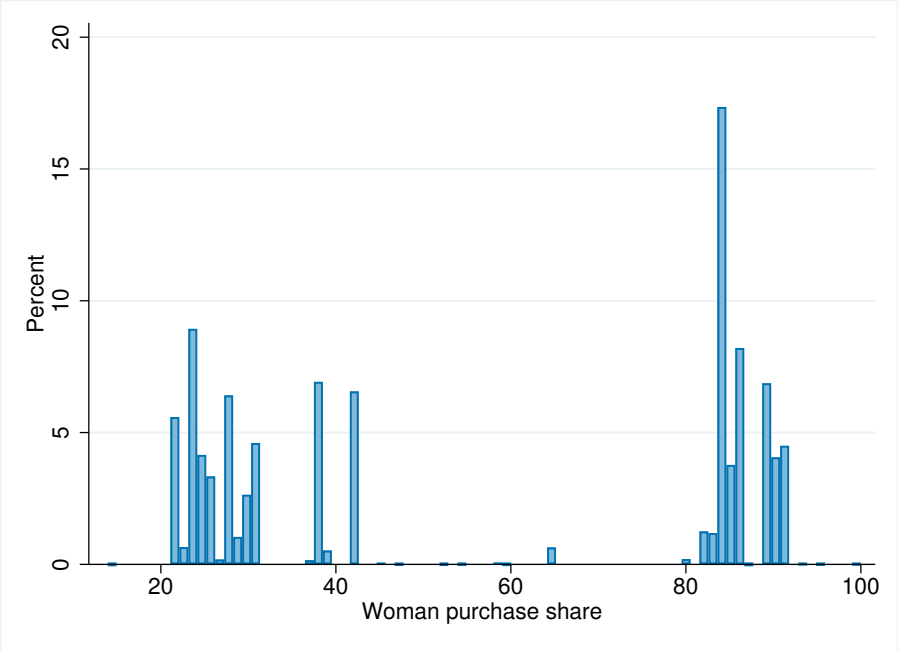
Individual-level clustered standard errors in parentheses

* $p < .10$, ** $p < .05$, *** $p < .01$

Note: This table presents estimates from the transaction-level regression: $log(C_{ijt}) = \phi_{t(j)} + \beta \mathbf{1}_{w(i)} + \Gamma X_i + \epsilon_{ijt}$ where C_{jt} is the wholesale price of UPC j in year t as observed in PriceTrak. $\mathbb{1}\{Woman_i = 1\}$ is an indicator for whether the individual is a woman, ϕ_t is a market-time fixed effect, and X_i is a vector of demographic controls including income, county, age, race and education. Panel (a) estimates this regression with equal weighting for all transaction-observations; Panel (b) weights each transaction-observations by the transaction expense as a share of the individual's annual income. Standard errors are clustered at the individual-level. This table is referred to in Section 5.

A.4 Additional differentiated products demand results

Figure A.10: Distribution of brand-gender for disposable razors



Note: This figure presents the distribution of brand gender, or woman purchase share, $\hat{w}_j$, for disposable razors. Brand gender is defined as the time-invariant expense-weighted average share of a product brand purchased by women. This figure is referred to in Section 6.

Table A.15: Differentiated Products Model Parameter Estimates

	Disposable Razors
Linear Coefficients	
Price (α)	-2.29***
	(0.22)
Nonlinear Coefficients	
Price (Σ_p)	0.81***
	(0.08)
Outside Option (Σ_1)	0.13
	(10.27)
Woman Purchase Share (Σ_w)	7.54E-05
	(70.22)

Market level clustered standard errors in parentheses

* $p < .10$, ** $p < .05$, *** $p < .01$

Note: This table presents estimated parameters from our differentiated products model using PyBLP. We allow for heterogeneity in taste for the outside option, price and the woman purchase share of the razor. We include product, store and week fixed effects. This table is referred to in Section 6.

Table A.16: Differentiated Products Model Estimates - No Firm Fixed Effects

	(1) Prices	(2) Own Elasticities	(3) Marginal Costs	(4) Markups
Women's razor indicator	0.0940***	-0.0033	0.0995***	0.1433***
	(0.0028)	(0.0035)	(0.0095)	(0.0026)
Observations	1071486	1071486	1071486	1071486
Adjusted R^2	-0.005	0.030	-0.040	-0.036
Men's Mean	1.6317	-2.0207	0.3763	0.7340
StoreXQuarter FE	Yes	Yes	Yes	Yes
Firm FE	No	No	No	No

Standard errors in parentheses

* $p < .10$, ** $p < .05$, *** $p < .01$

Note: This table displays regressions of prices and estimated price elasticities, marginal costs, and markups of disposable razors on a female-razor indicator. Female razors are defined as razors purchased at least 60% of the time by women, although 99% of razors see female purchase share above 80% or below 42%. Prices are directly observed in our data. Price elasticities, marginal costs, and markups are estimated from our differentiated products demand model described in Section 6. Standard errors take structural estimates as given. This table is referred to in Section 6.

Appendix B Purchase of Organic Products

Our analysis of prices and marginal costs paid by women suggests that women are sorting into products that have higher prices, higher marginal costs and lower markups. In this section we look at how one product attribute associated with higher costs of production, organic products, vary with gender. Organic products typically have higher costs of production because they cannot be grown with lower cost pesticides and require certification with the US government. For this analysis we restrict to consider only food products, as these departments have more reliable information on organic status. We present two main analyses that capture the difference between men and women in the purchasing habits of organic products. First, we estimate the difference in the share of organic purchases within a given market for men and women in Table B.1. Second, we plot the difference in share of products that are organic by woman purchase share relative to products bought equally by men and women in the same market in Figure B.1.

Table B.1 shows that women are about 0.2pp more likely to purchase an organic product than men are for products in the same product market. While this number is small, it is highly significant and reflects the overall low level of organic purchases. On average, the men in our sample have an average organic purchase rate of about 0.8% meaning that we estimate that women are between 25% and 32% more likely to buy organic products, a notable difference in propensity.

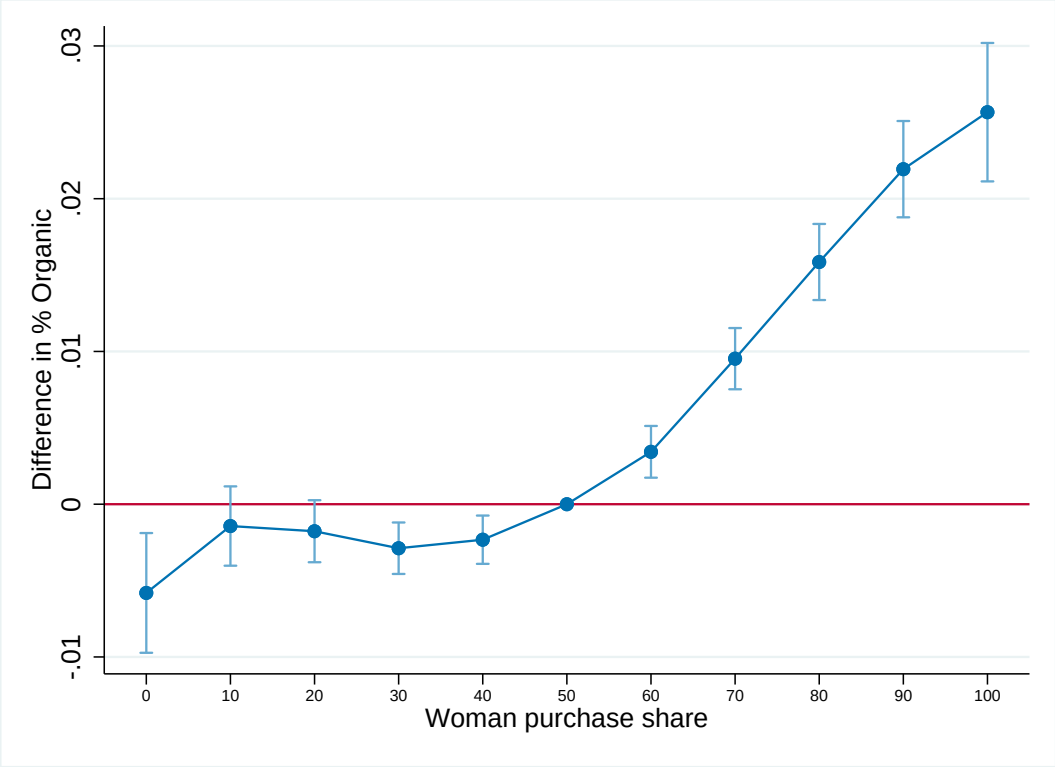
Figure B.1 plots the coefficients from a regression of organic status of a product on decile of woman purchase share normalized to products that are bought equally by men and women. The graph shows that products primarily bought by men are slightly less likely to be organic while products more often bought by women are significantly more likely to be organic. The orders of magnitude are comparable to those found in Table B.1, with women's products being about 0.2pp more likely to be organic.

Table B.1: Purchases of Organic Products

	(1)	(2)	(3)	(4)
Women	0.0023***	0.0026***	0.0022***	0.0021***
	(0.0003)	(0.0003)	(0.0002)	(0.0002)
Men's average	0.008	0.008	0.008	0.008
Module X Units FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	No
Month FE	No	No	No	Yes
County FE	No	Yes	Yes	Yes
Retailer FE	No	No	Yes	Yes
Demographic FE	Yes	Yes	Yes	Yes
Adj. R-squared	0	0	0	0
N	122,260,336	120,476,974	114,028,782	108,583,777
Number of clusters	49,252	49,250	49,245	49,249

Note: This table presents estimates from the regression: $\mathbb{1}_{O(ijt)} = \phi_{t(j)} + \beta \mathbb{1}_{w(i)} + \gamma X_i + \epsilon_{ijt}$ where $\mathbb{1}_{O(ijt)}$ is an indicator turned on if the purchase is an organic product. $\mathbb{1}\{Woman_i = 1\}$ is an indicator for whether the individual is a woman, ϕ_t is a market-time fixed effect and X_i is a vector of demographic controls including income, county, age, race and education which we add in sequentially. Standard errors are clustered at the individual-level. Column 1 can be thought of as a raw gap between single men and single women, each subsequent column demonstrates the contribution of controlling for an additional market or demographic factor. This table is referred to in Section 4.3.

Figure B.1: Share of Organic UPCs by Women's Purchase Share



Note: This figure presents plots of the results of the regression $\mathbb{1}\{IsOrganic_l\} = \alpha + \sum_{b \in \mathcal{B}} \gamma_b \mathbb{1}_{g(j_l)} \in Bin_b\} + \theta_{t_l} + \varepsilon_{j_l t_l}$ for a transaction l. Bins $b \in \mathcal{B}$ include ten-percentile-width bins centered at and two bins for pure gender stratification at the tails partitioning the interval $[0, 1]$. The regression includes fixed effects for product module, county and half-year. Results are presented for the whole sample and also separating out Health and Beauty and Dry Grocery. Standard errors are clustered at the UPC-county level. This figure is referred to in Section 4.3.

Appendix C Marginal Cost Validation: Razors Case Study

We validate our marginal cost findings for disposable razors with information about product characteristics that are likely correlated with the costs of production. Specifically, we scrape information on a razor's number of blades, the existence of a moisture strip, and the shape and contents of the handle. We create an indicator for whether the handle is ergonomic based on it having a shape that requires more plastic in comparison to a straight handle or whether it has additional rubber grip in the handle. We are able to gather information on product attributes for 90 out of the 176 razor product lines in our data (226 UPCs), however we capture those products that have the largest market share and are able to capture information for 73% of purchases that are made on private label disposable razors.

We present purchase-weighted comparisons of the product characteristics of the average women's razor to the average men's razor in Table C.1. In Panel A, we find that women's razor purchases have 0.3 more blades than men's, with the average razor purchase having between two and three blades. Women's razor purchases are slightly less likely to have a moisture strip, by about 5pp. Finally, women's razor purchases are about 17pp more likely to have an ergonomic handle. We take this as evidence that the razors that women purchase have characteristics associated with having higher costs of production as they require more materials to produce than men's razor purchases. Panel B presents UPC level results. We do not find significant differences in the average number of blades or moisture strips between men's and women's product offerings but do find that women's product offerings are significantly more likely to have ergonomic handles. The difference in findings between Panels A and B highlights the important role that sorting plays when considering differences in how men and women consume products.

Overall, we find that the product attribute data support our finding that women's razors have higher marginal costs of production. This should give confidence that while our marginal cost estimates may be biased downwards due to using a static model or other competitive factors, the trend lines and elasticity estimates are capturing meaningful differences in firm's pricing and production of products.

Table C.1: Women's and Men's Razor Attributes

	(1) Number of blades	(2) Has moisture strip	(3) Has ergonomic handle
Panel A: Purchase Level			
Women's Razor	0.3050***	-0.0477***	0.1663***
	(0.0001)	(0.0000)	(0.0000)
Men's Razor Average	2.239	0.758	0.310
Adj. R-squared	0.04	0.00	0.03
N	664,347,126	661,511,494	661,511,494
Panel B: UPC Level			
Women's Razor	0.0829	0.0504	0.2127***
	(0.1108)	(0.0521)	(0.0718)
Men's Razor Average	2.484	0.815	0.369
Adj. R-squared	-0.00	-0.00	0.03
N	226	224	224

* $p < .10$, ** $p < .05$, *** $p < .01$

This table plots coefficients from regressions of a given product characteristic on whether or not the product is a women's razor. Data on gender and characteristics were created by searching product and brand descriptions. Panel A presents results weighting each razor by the number of purchases observed in the RMS data. Panel B weights observations equally. Heteroskedasticity-robust standard errors are presented in parentheses. This table is referred to in Section 6.

Appendix D Retailer markups: model and mathematical appendix

How do our results on retailer markups inform the mechanisms underlying the pink tax? We find that gendered differences in unit prices nearly disappear and possibly become negative after conditioning on retailer costs. Additionally, we find no average difference in retailer markups paid by men and women and even greater retail markups on male-gendered goods than on female-gendered goods. Setting aside concerns on external validity to goods that do not match to PriceTrak and goods under alternative vertical integration settings, under what conditions would these results imply that all of the pink tax is attributable to differences manufacturing cost?

Consider a decomposition of average difference in prices paid for a female and male good in a setting with an integrated manufacturer-distributor that sells to a retailer who sells to the final consumer:

$$\begin{aligned}
\mathbb{E}[\Delta\%p] &= \mathbb{E}[\Delta\%c + \Delta\%\mu^m + \Delta\%\mu^r] > 0 \\
&\implies \mathbb{E}[\Delta\%c] + \mathbb{E}[\Delta\%\mu^m] + \underbrace{\mathbb{E}[\Delta\%\mu^r]}_{\leq 0} > 0 \\
&\implies \mathbb{E}[\Delta\%c] > -\Big(\underbrace{\mathbb{E}[\Delta\%\mu^m]}_{\text{Not observed}} + \underbrace{\mathbb{E}[\Delta\%\mu^r]}_{\leq 0}\Big).
\end{aligned}$$

In rationalizing the observed pink tax in this setting, it is necessarily the case that $\mathbb{E}[\Delta\%c] > 0$ if $\mathbb{E}[\Delta\%\mu^m] < 0$. Therefore, a natural question to ask is: given our observation that $\mathbb{E}[\Delta\%\mu^r] \leq 0$, is it reasonable to suppose also that $\mathbb{E}[\Delta\%\mu^m] < 0$? This condition is implied by the sufficient but not necessary condition in Equation (13), but this condition is less restrictive and amounts to a bounding condition: that on average, differences in markups by UPC-gender are of identical sign along the vertical integration chain. Absent gender-differential competitive environments between layers of manufacturing, distribution, and final sale, markups at each production/supply layer are set according to ultimate consumer

demand, so that it is likely that $\mathbb{E}[\Delta\%\mu^m] \leq 0$ given our observation that $\mathbb{E}[\Delta\%\mu^r] < 0$; i.e. it is highly likely the case that $\mathbb{E}[\Delta\%c] > 0$.

D.1 Markups in a vertically integrated setting

Consider the environment of vertical integration as in Section 5. To explore the bounding problem of the average gender difference in manufacturer-distributor markup, we want to explore under what conditions they differ in sign from the average gender difference in retailer markups.

We suppose an integrated manufacturer-distributor that produces a female and male good, distributing this good to a retailer that sells to a final consumer. We maintain this structure of vertical integration in order to align with the PriceTrak data as well as our discussion in Section 5, but this discussion generalizes to other vertical integration structures as well, such as one with a distinct wholesaler and distributor.

A single manufacture produces two goods to respective gender demand bases $g \in \{f, m\}$ separately at marginal costs c_m and c_f. The final consumer demand bases for these products exhibit iso-elastic price-sensitivity ε_m, and ε_f respectively in a manner independent of consumption of the other good.

The manufacturer sells both products to a single retailer, and the retailer resells these products as final goods to the ultimate consumers.

The manufacturer's problem is

$$\max_{p_m^m, p_f^m, Q_m, Q_f} (p_m^m - c_m)Q_m + (p_f^m - c_f)Q_f,$$

consisting of the price and quantity combination that maximizes rents from the retailer. Let superscripts refer to stages of the production process ($k \in m$ for manufacturer-distributor and r for retailer) and subscripts refer to UPC-gender.

Finally, the retailer sets prices in selling to the final consumer with $c_m^r = p_m^m$ and $c_f^r = p_f^m$:

$$\max_{p_m, p_f, Q_m, Q_f} (p_m - c_m^r)Q_m(p_m) + (p_f - c_f^r)Q_f(p_f),$$

facing their differentiated iso-elastic consumer bases and taking wholesaler prices as exogenous. We define the final price as the retail price $p_h := p_h^r$.

The setup yields a standard multi-marginalization problem with the each stage setting prices according to a standard Lerner markup rule:[48]

$$p_h^k = c_h^k(1 - \frac{1}{|\varepsilon_h|})^{-1},$$

for each stage of the production process $k \in \{m, r\}$. This price-setting process results in a final price to consumers of

$$p_h = c_h(1 - \frac{1}{|\varepsilon_h|})^{-2}.$$

We are interested in knowing whether it is possible to observe the following simultaneously:

$$p_f^r := p_f > p_h =: p_h^r,$$
$$c_f^r = p_f^m > p_m^m = c_m^r,$$
$$\frac{p_f^r}{c_f^r} - 1 = \mu_f^r < \mu_m^r = \frac{p_m^r}{c_m^r} - 1,$$
$$|\varepsilon_f| > |\varepsilon_m|,$$

and

$$\mu_f^m < \mu_m^m + .$$

I.e. given our observation that 1) female prices at retail exceed male prices at retail, 2)

[48] Quantities are ultimately set by the consumer. Because firms linearly maximize profit, each stage internalizes consumers' down-the-line demand response to prices.

female retail costs (i.e. manufacturer-distributor prices) exceed male retail costs, 3) male retail markups exceed female retail markups, and 4) elasticity of demand on the female goods exceeds that of the male good in absolute value, can it be the case that the manufacturer-distributor markup for female goods exceeds that of male goods? In this simplified environment, the answer is no. Without alternate assumptions on the structure of competition within and between layers, lower retailer markup and greater elasticity on part of women likely imply a lower overall markup between manufacturing to final sale to consumer, which determines overall demand.