Gender-Based Pricing in Online Markets

Diego Aparicio, *IESE Business School*[1]
Maxime Cohen, *McGill University*
Amber Liu, *Santa Clara University*

Abstract

We study gender-based pricing in online retail markets, commonly referred to as the "pink tax," using product-level data from Amazon in beauty and personal care categories. We infer product gender at scale using large language models and construct matched pairs of highly similar products based on textual similarity. We document that male-targeted products are priced lower than comparable female-targeted products, while female-targeted products are more likely to be promoted, indicating that firms use both posted prices and temporary discounts to implement gender-based price differences. Extending beyond a binary perspective, we show that gender-neutral products are typically priced between male- and female-targeted products. Finally, we find that algorithmic pricing is associated with a compression of gender-based price differences. Overall, we provide new evidence on gender-based pricing in online markets by jointly examining prices, promotions, and algorithmic pricing in a large-scale retail setting.

Keywords: Gender Pricing; Pink Tax; Economics of Gender; Artificial Intelligence; Online Platforms
JEL CODES: J16; D12; L11; L81

[1]Correspondence to: DAparicio@iese.edu. Manuscript version: March 2026.

1 Introduction

The "pink tax" refers to the phenomenon whereby products marketed toward female consumers are priced higher than comparable products targeted at men. Evidence of such price differences has been documented across a range of consumer goods and services, raising concerns about systematic disparities in pricing across gender-targeted products. These differences are often discussed in the context of broader gender inequalities, including well-documented gaps in earnings (Blau and Kahn 2017, Goldin et al. 2017). Reflecting these concerns, policy discussions have increasingly focused on the role of gender in retail pricing. For example, when introducing legislation aimed at addressing such disparities, Congresswoman Jackie Speier noted that women "make less for doing the same work and pay more for the same product or service because it's for women" (The Hill 2018).

Concerns about gender-based pricing have existed for decades, but the term "pink tax" gained broader public and policy attention in the mid-1990s. Early investigations documented systematic price differences across a variety of consumer goods and services.[2]

One of the earliest policy responses emerged in California. In 1994, the California Assembly Office of Research—under the leadership of then–Assemblywoman Jackie Speier—conducted a statewide investigation into gender-based price disparities. These findings contributed to the passage of the Gender Tax Repeal Act of 1995, which prohibited businesses from charging different prices for comparable *services* based solely on gender. The law states that "no business establishment may discriminate, with respect to the price charged for services of a similar or like kind, because of the person's gender," and remains in force today.

Interest in gender-based pricing intensified in the 2010s following several high-profile investigations. In particular, a study commissioned by the New York City Department of Consumer Affairs (DCA) in 2015 found that products marketed to women were priced higher than comparable men's products in 42% of cases and by 7% on average (New York City Department of Consumer Affairs 2015). The report received widespread media attention and contributed to renewed policy attention.

Recent years have seen additional legislative efforts to regulate gender-based pricing. In 2020, New York enacted legislation stating that businesses may not "charge a price for any two goods that are substantially similar, if such goods are priced differently based on the gender of the individuals for whom the goods are marketed and intended" (N.Y. GBS Section 391-u) (New York State Senate 2020). California subsequently extended its earlier service-based prohibition to consumer goods through Assembly Bill 1287, signed in 2022 and effective in 2023, introducing penalties of up to 100,000 per violation (California State Legislature 2022). As of 2025, several additional states—including New Jersey, Connecticut, Massachusetts, Washington, and Minnesota—have proposed or enacted similar measures.

At the federal level, the proposed *Pink Tax Repeal Act*, first introduced in 2016 and reintroduced multiple times since, would prohibit gender-based pricing nationwide, although it has not been

[2]Examples include a 1993 California Assembly Office of Research survey of five major Californian cities and a 1996 South Dakota Legislative Research Council issue memorandum summarizing multi-state evidence.

enacted.[3]

Beyond directly regulated products, gender-neutral alternatives have also emerged in public discussions as a practical way for consumers to avoid gender-based price discrimination. Consumer advocacy campaigns and media coverage often recommend comparing products across gender lines or choosing gender-neutral versions of personal care goods as a strategy to bypass higher prices associated with female-targeted products (Crockett 2016, Time 2016). For instance, a commissioner of the New York City DCA noted in an interview with Time: "when I buy shampoo for our house, I buy one kind of shampoo, and I buy a gender-neutral shampoo. I will not buy a product that I know has been marked up falsely" (Time 2016). This evolution suggests that analyses focused solely on male–female comparisons may overlook an increasingly relevant segment of the market.

More broadly, the growing presence of gender-neutral products reflects a shift in product design and marketing beyond binary gender segmentation. Despite this shift, policy discussions and empirical analyses have primarily focused on comparisons between male- and female-targeted goods, leaving little evidence on gender-neutral products. This raises a natural question: how are gender-neutral products priced relative to male- and female-targeted products?

We contribute to this debate by examining gender-based pricing in a major online marketplace. Our analysis focuses on Amazon, which accounts for roughly 40% of U.S. e-commerce sales.[4] We assemble a large-scale dataset covering more than 15,000 products and over 18 million product-day price observations across more than 20 beauty and personal care categories. This scale substantially expands the scope of prior empirical studies, which typically examine a smaller number of categories or products (e.g., Duesterhaus et al. 2011, New York City Department of Consumer Affairs 2015, Gonzalez Guittar et al. 2021, Moshary et al. 2023).

The high-frequency nature of our data allows us to examine not only price levels but also promotional activity. Using daily price observations, we study whether male- and female-targeted products differ in the likelihood of receiving price promotions—an aspect that has received limited attention in the context of gender-based pricing. The online retail environment, combined with rich time-series data, also provides a natural setting to study the interaction between gender-based pricing and algorithmic pricing. As one of the earliest and largest adopters of algorithmic pricing at scale (Chen et al. 2016, Aparicio and Misra 2022, Aparicio et al. 2024), Amazon offers a setting in which to examine whether the use of artificial intelligence in pricing attenuates or amplifies gender-based price differences.

These issues are increasingly relevant from a policy perspective. Policymakers have begun to scrutinize the role of algorithmic and data-driven pricing in retail markets. For example, New York recently introduced the Preventing Algorithmic Pricing Discrimination Act, which would require retailers to disclose when prices are generated using algorithms that rely on consumer data (New

[3]Gender-based pricing has also received attention in Europe (Sommen 2025). Several jurisdictions have reclassified menstrual products as essential items: Germany reduced VAT from 19% to 7% in 2020 (New York Times 2019); the United Kingdom abolished the "tampon tax" in 2021 (AP News Wire 2021); and by 2023, more than 20 U.S. states had eliminated sales taxes on such products (New York Times 2023).

[4]E-commerce retail sales totaled approximately $1.1 trillion in the United States in 2023 (Federal Reserve Bank of St. Louis 2025).

York State Senate 2025). Similarly, the Federal Trade Commission has launched inquiries into personalized pricing practices (Federal Trade Commission 2025)—often referred to as "surveillance pricing." Although these efforts are not explicitly framed around gender-based pricing, they reflect broader concerns about how AI-driven pricing technologies may generate or amplify disparities, particularly when they rely on consumer characteristics.

We develop a scalable empirical strategy to identify gender targeting and product comparability across large online assortments. First, we infer product gender using information contained in product titles and descriptions. Because gender targeting is often embedded in marketing language rather than structured product attributes, we employ large language models (LLMs) to process textual product information and classify items as male-targeted, female-targeted, or gender-neutral. In addition, we leverage Amazon's internal categorization, which in some cases explicitly separates products by gender (e.g., perfumes and fragrances listed under distinct categories for men and women).

Second, we identify highly comparable products within each category using a combination of text-based similarity measures. Specifically, we leverage structured product attributes and textual descriptions to construct matched pairs of products that are closely comparable in their characteristics, are produced by the same manufacturer, but differ in their target gender. This approach allows us to construct *matched pairs* of products that are marketed toward different genders but are otherwise closely comparable in their descriptions and features. By identifying the closest comparable products across genders, we can directly evaluate price differences while minimizing confounding variation in product characteristics. Together, these methods enable a scalable and replicable approach to studying gender-based pricing disparities in large digital retail environments.

We document several empirical findings. First, we find that male-targeted products are systematically priced lower per unit than comparable female-targeted products, even after controlling for high-dimensional fixed effects. This pattern is consistent with the presence of a "pink tax" in the beauty and personal care categories we study. Second, we show that female-targeted products are more likely to be promoted, indicating that firms can use two pricing levers, namely posted prices and temporary discounts, to implement gender-based price discrimination.

Extending beyond the binary comparison emphasized in the existing literature, we find that gender-neutral products tend to be priced between male- and female-targeted products, suggesting that they occupy an intermediate position in firms' segmentation strategies and may provide an outside option for consumers seeking to avoid gender-based price differences. Finally, we provide evidence that algorithmic pricing is associated with a compression of gender-based price differences, particularly when comparing highly similar products, suggesting that increased pricing flexibility may limit the persistence of such differences.

The remainder of the paper is structured as follows. We review the related literature in Section 2. Section 3 describes the data. Sections 4 and 5 report evidence on gender-based pricing and promotions, respectively. Section 6 extends the analysis to gender-neutral products. Section 7 discusses the role of algorithmic pricing. Section 8 reports robustness checks and Section 9 concludes.

2 Literature Review

The "pink tax" sits at the intersection of two strands of literature: price discrimination and gender-based inequality. While both strands are well developed, relatively little research has directly examined gender-based pricing in consumer product markets. Conceptually, the pink tax can be understood through the lens of classic models of price discrimination. Third-degree price discrimination occurs when sellers are "charging different prices to different markets or classes of customers" (Schmalensee 1981), implying that the seller distinguishes customer segments and charges each segment a different price. Meanwhile, second-degree price discrimination occurs when the seller "presents a schedule which diverts individuals of different types towards price-quantity combinations" (Leland and Meyer 1976), whereby consumers self-select among available product options. Examples of the latter include offering products of varying quality or discounts (Anderson and Dana Jr 2009).

In practice, however, gender-based pricing in retail markets may not map cleanly into either form of price discrimination. On one hand, products targeted toward men and women frequently differ in packaging, ingredients, or branding, allowing consumers to self-select into gendered variants even when the products serve similar purposes. On the other hand, product categorization in retail environments can segment the choice set that consumers face. Many online marketplaces explicitly organize assortments into gendered categories such as "Men" and "Women." When products are displayed in separate sections, consumers browsing one category may never encounter comparable products listed in another. In this sense, digital categorization can resemble a retail environment in which consumers effectively enter different aisles depending on gender, limiting the likelihood that they directly compare alternatives marketed to another group and making the outcome closer to third-degree price discrimination. This feature contrasts with traditional brick-and-mortar retail environments, where consumers typically observe the full assortment on the same shelf. Consistent with this view, Moshary et al. (2023) show that gender segmentation is pervasive in personal care markets and that men's and women's products often differ substantially in formulation and attributes, although price differences largely disappear when comparisons are restricted to highly similar products.

Our work is not the first to examine gender-based price differences. Early work on gender inequality primarily focuses on the labor market, documenting persistent gender pay gaps (Goldin 1990, Bertrand et al. 2010, Blau and Kahn 2017, Goldin et al. 2017). In this line of research, there is broad consensus that "everywhere in the world, women earn less than men," as noted in Goldin (2024)'s Nobel lecture.

More recently, studies have documented gender-based price differences across a variety of markets, including automobile sales and repairs (Goldberg 1996, Harless and Hoffer 2002, Morton et al. 2003, Busse et al. 2017, Ozturk et al. 2022), real estate transactions (Harding et al. 2003, Goldsmith-Pinkham and Shue 2023, Andersen et al. 2021), credit markets (Alesina et al. 2013), and consumer packaged goods (New York City Department of Consumer Affairs 2015, Moshary et al. 2023, Barnes and Brounstein 2025). Much of the evidence from these settings comes from

markets with negotiated prices or substantial seller discretion, where gender differences may arise through bargaining dynamics or informational asymmetries.

By contrast, consumer packaged goods are typically sold in posted-price retail environments where consumers can directly observe and compare alternatives and negotiation is generally not feasible. Closer to our work, using Nielsen scanner data, Barnes and Brounstein (2025) show that women pay more than men on average, although this premium is mainly driven by the purchase of lower-markup but higher-quality products, again suggesting that gender-based price differences can be explained by differences in purchasing behavior and product selection (Moshary et al. 2023). We contribute to this literature by documenting gender-based pricing and promotion strategies in online marketplaces, focusing on beauty and personal care categories. In addition, we extend the analysis to gender-neutral products, which have received little attention in the existing literature.

Our empirical setting of online marketplaces connects to the growing literature on algorithmic pricing in digital markets. Advances in data analytics allow firms to adjust prices rapidly and condition them on a wide range of market signals, including demand conditions, competition, and consumer browsing or purchasing patterns (Elmaghraby and Keskinocak 2003, Aparicio and Misra 2022). While such technologies can improve pricing efficiency, they may also raise concerns about fairness and transparency. A large literature shows that consumers perceive price differences as unfair when they appear arbitrary or difficult to justify, which can lead to negative emotional responses, lower purchase intentions, and reputational costs for firms (Huppertz et al. 1978, Kahneman et al. 1986, Campbell 1999, Xia et al. 2004). In response, recent research has begun to explore pricing strategies that incorporate fairness considerations or regulatory constraints (Cohen et al. 2022). We contribute to this literature by examining the interplay of algorithmic pricing in gender-based pricing.

Even when posted prices are identical, effective prices may still differ systematically through discounts and promotions (Anderson et al. 2023). An advantage of our high-frequency online data is that it allows us to measure promotional activity and discount frequency precisely. While there is an extensive marketing literature on promotions, relatively little work has examined whether discounting patterns differ across gender-targeted products. Existing evidence suggests that women tend to be more price-sensitive and more responsive to coupons and promotional offers (Harmon and Hill 2003, NielsenIQ 2024, Barnes and Brounstein 2025). We contribute to this literature by examining whether promotion frequency differ across male-, female-, and gender-neutral products in online marketplaces.

3 Data Description

Our analysis uses product-level price and attribute data for items sold on Amazon.com. We obtain these data from `Keepa.com`, a commercial price-tracking platform that records historical prices and product characteristics for items listed on Amazon.

We collect data for the top-selling 2,000 items in each of 21 categories within the personal care

department over the period March 2011 to November 2025. Focusing on top-selling products allows us to capture the portion of the market that accounts for the majority of consumer demand, while abstracting from long-tail products with limited sales and infrequent price updates. As a result, our sample reflects pricing behavior for the most economically relevant products in the marketplace.

Each item (ASIN) is classified according to Amazon's hierarchical product taxonomy. For example, the product *Dove Beauty Bar Soap* (Figure 1) is categorized as: Personal Care → Skin Care → Body Skin Care → Body Cleansers → Body Soap. We define the highest level ("Personal Care") as the department and the most granular level ("Body Soap") as the product category.

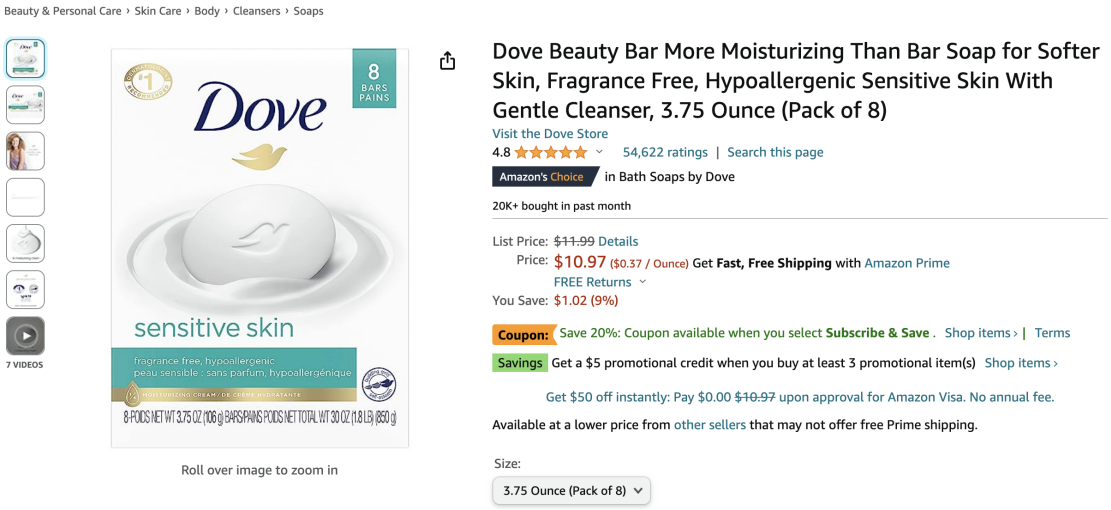


Figure 1: Product Page on Amazon

We construct a panel dataset at the product × date level. When multiple price observations are recorded within a day, we use the median daily price. We exclude observations in which products are listed as out of stock. Furthermore, the dataset contains a rich set of product attributes, including manufacturer, brand, category, title, description, quantity, weight, and size. Product titles provide concise descriptions, while product descriptions offer more detailed information that we later use for gender inference. Additional details on data collection and variable construction are provided in Section A1 in the Appendix.

The resulting dataset includes over 18 million product × date observations, from March 2011 to November 2025, and covers 15,748 items in 21 categories in the personal care department. Table 1 presents summary statistics. The average price of a product is $25.54 and the average shelf time we observe is 55.46 months (roughly 4.6 years).

Table 1: Summary Statistics

Category	# Observations	Avg. price	# Products	Avg. month span
BathShowerGels	1,045,468	19.04	992	47.96
BathSoaps	673,075	13.55	608	60.45
BodyLotions	984,453	20.21	949	59.66
BodyMoisturizers	1,156,898	16.56	899	67.98
BodySprays	535,180	15.13	622	42.43
CCCream	39,170	25.23	44	43.66
Cologne	1,220,166	43.33	981	62.68
ConcealersAndNeutralizers	594,939	17.01	591	46.17
Deodorant	745,673	16.18	799	48.90
EaudeParfum	1,666,283	58.89	2,003	37.79
FaceMoisturizers	1,162,846	32.64	735	63.24
FacialCleansingWashes	1,046,389	21.12	813	56.42
HairColor	709,544	12.44	641	54.96
HairConditioner	1,874,733	20.71	1,334	60.36
HairLossProducts	348,050	34.12	308	54.93
HairShampoo	1,661,763	22.65	1,179	58.86
HandCreamsAndLotions	567,997	18.53	507	52.12
HandWash	809,274	24.60	711	55.06
Masks	453,161	27.76	356	54.89
ShavingCreamsLotionsGels	834,336	14.94	751	58.74
ShavingLotions	23,779	24.74	41	55.71
Total (All Beauty)	**18,152,816**	**25.54**	**15,748**	**55.46**

3.1 Strategies to Infer Gender

We classify products into three gender-targeting categories: *female-targeted*, *male-targeted*, and *gender-neutral*. We adopt an inclusive definition of gender as self-identified rather than biologically assigned sex. Gender is not restricted to a binary classification and may lie along a continuum. Accordingly, we classify products as gender-neutral when they are not explicitly targeted toward male or female consumers.

To ensure robustness, we implement three complementary classification strategies:

- **Amazon Internal Classification.** We use Amazon's internal product taxonomy, which in some cases explicitly encodes gender targeting. When this classification exists, male-targeted products are listed in a separate category than female-targeted products. Products classified as "male" or "female" are assigned accordingly, while products without an explicit gender designation are classified as gender-neutral.

- **Keyword-Based Classification.** We infer gender targeting from keywords in product titles and descriptions. Terms such as "men", "male", "women", or "female" provide direct signals of gender targeting. Products containing both or neither sets of keywords are classified as gender-neutral.

- **Generative AI Classification.** We use ChatGPT to classify each product based on its title and description. The model is prompted to assign each item to one of three categories: male, female, or unisex.[5] Compared to keyword-based approaches, this method captures more subtle contextual cues in product descriptions and branding, providing a more flexible and comprehensive classification.

3.2 Matched Products

To further isolate the role of gender targeting, we exploit the fact that Amazon frequently offers highly similar products in male- and female-targeted versions. We construct matched pairs of such products and compare their prices within pairs. We identify these pairs using a text-similarity approach that combines fuzzy string matching on product titles with additional filtering based on product attributes (e.g., size, quantity, and category). For each product, we select the closest counterpart of the opposite gender within the same category based on a similarity score ranging from 0 to 100. We retain unique product pairs with similarity scores above 50, ensuring that matched items are comparable in observable characteristics.

Figure 2 illustrates examples of matched product pairs. These pairs consist of products that are highly similar in description and positioning but differ in their gender targeting, enabling a narrow comparison of prices across male- and female-targeted products.

Table 2 reports the distribution of products by gender using the AI-based classification. In the full sample, gender-neutral products account for the majority of observations (66.6%), while female- and male-targeted products represent 17.4% and 16.0%, respectively. The matched sample, by construction, excludes gender-neutral products and is approximately balanced between female (50.1%) and male (49.9%) items. The reduction in sample size reflects the requirement that each product has a sufficiently comparable counterpart of the opposite gender.

[5]Prompt: *Based on the following Amazon product title and description, is the following Amazon product for men, women, or unisex? Please answer by reporting one of the three options and nothing else.* The model returns a valid classification for 99.97% of products; the remaining cases are manually classified.



(a) Suave Men 2 in 1
Hair and Body Wash
15 Fl Oz



(b) Suave Essentials
Body Wash
15 Fl Oz



(c) Tom's of Maine
Deodorant for Men
2.6 Oz



(d) Tom's of Maine
Deodorant for Women
2.25 Oz

Figure 2: Examples of Gen-AI Matching Pairs

Table 2: Products by Target Gender

	Full sample		Matching sample	
AI Gender	#	%	#	%
Neutral	10,193	66.63	–	–
Female	2,654	17.35	1,753	50.19
Male	2,450	16.02	1,740	49.81
Total	15,297	100.00	3,493	100.00

4 Evidence on Gender Pricing

We begin our empirical analysis by estimating price differences across products targeted to different genders. We first focus on products targeted to male and female consumers, and subsequently

extend the analysis to include gender-neutral products.

Let $i \in I$ index products, $j \in J$ product categories, $k \in K$ firm (manufacturer), and $t \in T$ date. Each product i belongs to a single category and a single firm, although multiple products may share the same category or firm. We estimate the following fixed-effects specification:

$$\log Price_{it} = \alpha + \beta Men_i + \gamma_{jkt} + \epsilon_{it}, \tag{4.1}$$

where $\log Price_{it}$ denotes the log of price per unit (price divided by total ounces w_i), and Men_i is an indicator equal to one if product i is targeted to male consumers. The fixed effects γ_{jkt} absorb all time-varying shocks at the firm $\times$ category $\times$ date level, thereby ensuring that identification comes from within-firm-category-day comparisons. To ensure comparability across products of different sizes, we define our main outcome as price per unit of weight.

The coefficient β captures the price difference between male- and female-targeted products. Thus, a negative value of β indicates that male-targeted products are priced lower, on a per-unit basis, than female-targeted products. We estimate Eq. (4.1) separately using three alternative gender classification strategies described in Section 3.1: (i) Amazon's internal categorization, (ii) classification based on keywords in product titles and descriptions, and (iii) classification using generative AI. In each case, the definition of Men_i is adjusted accordingly. Additionally, we estimate the model on a restricted sample that includes only products for which Amazon and AI classifications agree. This provides a conservative benchmark that mitigates concerns about misclassifications.

Finally, we implement a matched-sample approach that compares highly similar male- and female-targeted products using structured similarity measures applied to product titles and descriptions (Section 3.1). Specifically, we estimate:

$$\log Price_{it} = \alpha + \beta Men_i + \delta_{ii'} + \gamma_j + \gamma_t + \epsilon_{it}, \tag{4.2}$$

where $\delta_{ii'}$ denotes fixed effects for matched product pairs, and γ_j and γ_t denote category and date fixed effects. This specification restricts comparisons to products that are closely matched in observable characteristics.

Table 3 reports the results. Across all specifications, male-targeted beauty and personal care products are consistently and significantly priced lower per unit than female-targeted products. This pattern is robust to alternative gender classification methods, to restricting the sample to products with consistent classifications, and to comparisons within matched product pairs.

The estimated effects are economically meaningful. Specifically, the coefficients imply that male-targeted products are priced between approximately 3.7% and 14.4% lower than comparable female-targeted products (corresponding to $e^{-0.038}-1$ and $e^{-0.155}-1$, respectively). Taken together, these findings provide evidence consistent with systematic gender-based price differences. Moreover, the estimates are qualitatively in line with the 7-10% Pink Tax price gaps reported in the literature (New York City Department of Consumer Affairs 2015, Moshary et al. 2023).

Table 3: Price Effects

	(1) $Price_{it}$	(2) $Price_{it}$	(3) $Price_{it}$	(4) $Price_{it}$	(5) $Price_{it}$
Men Targeted	-0.038***	-0.075***	-0.046***	-0.155***	-0.045***
	(0.001)	(0.002)	(0.002)	(0.001)	(0.002)
Constant	1.983***	1.696***	2.005***	1.635***	1.664***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	3,280,916	3,851,326	2,507,813	3,132,039	1,170,490
R-squared	0.791	0.828	0.795	0.872	0.876
Gender source	Amazon	AI	Intersection	Keywords	Matching AI
Firm × Categ. × Date FE	✓	✓	✓	✓	×
Match Pair FE	×	×	×	×	✓
Date FE	×	×	×	×	✓

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

5 Price Promotions

In this section, we examine whether there are systematic differences in promotion patterns between male- and female-targeted products. Specifically, we study how the likelihood of a product being discounted varies with gender targeting.

We define the reference price of a product as the maximum price observed over the previous 60 days. A price promotion occurs when the current price is at least 20% lower than this reference price. (We consider alternative definitions in Section 8.) Formally, we define the discount depth variable Δ_{it}^{f} as the log deviation of the current price from the reference price, where $f \in \{60\}$ denotes the look-back window:

$$\Delta_{it}^{f} = \log(P_{it}) - \log\left(\max_{r\in\{t,t-1,\ldots,t-f-1\}} P_{ir}\right).$$

Using this measure, we construct a binary promotion indicator $Y_{it}^{f,d}$ that equals 1 if product i on date t has a price discount of at least $d \in \{0.20\}$ and 0 otherwise. This serves as our main dependent variable. Further, our baseline specification is a conditional logit model, which allows us to compare products within the same firm × category × date group. We also report results from alternative specifications, including a linear probability model and different promotion definitions, in Section 8.

The probability of observing a promotion is given by:

$$\mathbf{P}(Y_{it}^{f,d} = 1|X_{it}) = \frac{\exp(\beta Men_i + \gamma_{jkt})}{1 + \exp(\beta Men_i + \gamma_{jkt})} \tag{5.1}$$

where γ_{jkt} denotes firm × category × date fixed effects, and Men_i is an indicator for male-targeted products. The coefficient β captures the difference in log-odds of receiving a promotion for male-targeted products relative to female-targeted products. A negative coefficient implies that male-targeted products are less likely to be promoted.

The results are reported in Table 4. Starting with Column (1), which uses Amazon's gender classification, the coefficient on male-targeted products implies that the odds of a promotion for male-targeted products are approximately 17% lower ($e^{-0.181}$) than for comparable female-targeted products within the same firm-category-date group. This pattern is consistent across all specifications: male-targeted products are systematically less likely to exhibit a price promotion than female-targeted products. The magnitude of the estimates ranges from approximately 9% to 17% lower odds of promotion.

Overall, these findings indicate that promotional activity differs systematically across gender-targeted segments. Combined with the price gap results from Section 4, the evidence suggests that female-targeted products are both priced higher but promoted more frequently, pointing to a joint role of prices and promotions in shaping gender-based price differences.

Table 4: Frequency of Price Promotions

	(1) $Prom_{it}^{60,20}$	(2) $Prom_{it}^{60,20}$	(3) $Prom_{it}^{60,20}$	(4) $Prom_{it}^{60,20}$	(5) $Prom_{it}^{60,20}$
Men Targeted	-0.181***	-0.135***	-0.189***	-0.142***	-0.090***
	(0.005)	(0.005)	(0.006)	(0.005)	(0.005)
Observations	1,861,247	1,985,369	1,336,103	1,639,495	326,616
Gender Inference	Amazon	AI	Intersection	Keywords	Matching AI
Firm × Categ. × Date FE	✓	✓	✓	✓	×
Match Pair × Date FE	×	×	×	×	✓

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Finally, we also examine whether male- and female-targeted products differ in the size of price discounts, conditional on a promotion taking place. Restricting the sample to observations where $Y_{it}^{f,d} = 1$, we estimate:

$$|\Delta_{it}^{f}| = \alpha + \delta Men_i + \gamma_{jkt} + \varepsilon_{it}, \tag{5.2}$$

where $|\Delta_{it}^{f}|$ denotes the absolute log discount depth relative to a reference price computed over the

previous $f \in \{60\}$ days. To mitigate sensitivity to isolated price spikes, we define the reference price as the 90th percentile of prices over the previous 60 days.

We find some evidence that male-targeted products exhibit slightly larger discounts, although the magnitudes are relatively small. For example, a coefficient of 0.005 in log discount depth corresponds to approximately a 0.5% difference in discount size. The results are reported in Appendix A2.5.

6 Gender-Neutral Products

In this section, we extend the analysis by incorporating products marketed as gender-neutral or unisex. The growing availability of such products reflects broader shifts in product positioning and consumer preferences, where firms increasingly move beyond strictly gendered marketing. While the existing literature on gender-based pricing has primarily focused on differences between male- and female-targeted products, comparatively less attention has been paid to the role of gender-neutral offerings.

Studying the pricing of these products is economically relevant because they may interact with gender-based price discrimination mechanisms documented in the "pink tax" literature. If gender-based pricing reflects firms' ability to segment consumers into male and female markets, the introduction of gender-neutral products creates an additional segment that firms may target, potentially reshaping pricing strategies and substitution patterns across segments.

Interestingly, several mechanisms may govern the pricing of gender-neutral products. For instance, gender-neutral products may act as a competitive outside option, particularly for consumers facing higher prices in gender-targeted segments. In this case, substitution toward neutral products can discipline pricing, especially for female-targeted goods, leading to lower relative prices for neutral offerings. Alternatively, gender-neutral products may reflect an intermediate positioning strategy, targeting consumers who are less responsive to gendered differentiation, leading to prices that fall between those of male- and female-targeted products. Overall, these mechanisms yield ambiguous theoretical predictions.

To examine these patterns, we estimate the following fixed-effects specification:

$$\log Price_{it} = \alpha + \beta_1 Neutral_i + \beta_2 Men_i + \gamma_{kjt} + \epsilon_{it}, \tag{6.1}$$

where $\log Price_{it}$ denotes the log price per unit, and γ_{kjt} are firm $\times$ category $\times$ date fixed effects. Further, β_1 and β_2 capture price differences for gender-neutral and male-targeted products, respectively, relative to female-targeted products (baseline). In this analysis, gender is inferred using generative AI and keyword-based methods.

Table 5: Price and Promotion Effects — Gender Neutral

	(1) $Price_{it}$	(2) $Price_{it}$	(3) $Prom_{it}^{60,20}$	(4) $Prom_{it}^{60,20}$
Gender-Neutral Targeted	-0.042***	-0.014***	-0.058***	-0.100***
	(0.001)	(0.001)	(0.004)	(0.004)
Men Targeted	-0.068***	-0.071***	-0.133***	-0.120***
	(0.001)	(0.001)	(0.005)	(0.005)
Constant	1.045***	1.025***		
	(0.001)	(0.001)		
Observations	12,787,511	12,787,511	6,154,977	6,154,977
R-squared	0.832	0.832		
Gender Inference	AI	Keywords	AI	Keywords
Firm × Categ. × Date FE	✓	✓	✓	✓

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 5 reports the results. Focusing on Column (1), which uses AI-based gender classification, gender-neutral products are priced slightly below female-targeted products, while male-targeted products continue to exhibit larger price discounts. The coefficient on gender-neutral products implies that prices are roughly -4.1% lower ($e^{-0.042} - 1$) than female-targeted products, and the coefficient on male-targeted products corresponds to a price difference of about -6.6%.

These results indicate that gender-neutral products occupy an intermediate position in the price distribution, while male-targeted products remain consistently cheaper on a per-unit basis. This suggests that while gender-neutral products do not command a premium, they may serve as a partial substitute for female-targeted products, potentially constraining the extent of gender-based price differentiation.

We next examine whether promotion patterns differs across gender targeting. Specifically, we estimate the probability that a product is on promotion using a conditional logit model with the same fixed effects structure:

$$\mathbf{P}(Y_{it}^{f,d} = 1|X_i) = \alpha + \beta_1 Neutral_i + \beta_2 Men_i + \beta_3 \log(w_i) + \gamma_{kjt} + \varepsilon_{it}, \tag{6.2}$$

The results are shown in Table 5, Columns 3 and 4. We find that both gender-neutral and male-targeted products are less likely to be promoted than female-targeted products. Focusing on Column (3), and interpreting coefficients as odds ratios, the estimate for gender-neutral products implies that their odds of having a price promotion are approximately 6% lower ($e^{-0.058}$) than those of comparable female-targeted products, and for male-targeted products it corresponds to

12% lower odds of a price promotion.

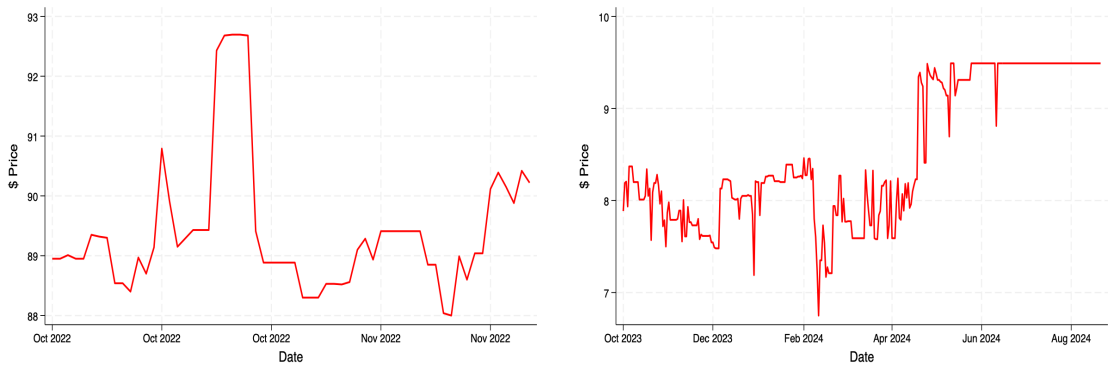
Taken together, these findings indicate that female-targeted products are both priced higher and promoted more frequently, while gender-neutral products occupy an intermediate position in both price levels and promotional intensity. This joint pattern is consistent with firms using promotions as an additional pricing instrument to further segment customers *within* a given gender. More broadly, these results highlight that price discrimination in this setting operates not only through posted prices but also through promotional activity. Incorporating promotions into the analysis therefore provides a more complete view of gender-based pricing, as promotions can further segment consumers based on their price sensitivity and purchasing timing.

7 The Role of Algorithmic Pricing

The increasing adoption of algorithmic pricing tools has transformed how firms and platforms update prices, particularly in online marketplaces where prices can be adjusted frequently and at low cost. In this environment, pricing rules may become more flexible and responsive to demand and competitive conditions, potentially affecting the extent to which firms sustain or compress price differences across gender-targeted products. We therefore examine whether the use of algorithmic pricing amplifies or attenuates gender-based price differences.

A direct empirical challenge is that we do not observe when Amazon uses an algorithmic pricing system. Following the existing literature, we therefore rely on the frequency of price changes as a proxy for algorithmic pricing intensity (Aparicio and Misra 2022, Musolff 2026, Aparicio et al. 2026, Assad et al. 2024). The underlying idea is that products whose prices adjust more frequently are more likely to be managed using automated or algorithmic pricing tools, whereas products with infrequent price changes are more likely to follow more rigid or manually determined pricing rules. We interpret this measure cautiously as capturing pricing dynamics consistent with algorithmic pricing, rather than as a direct measure of algorithm adoption.

To build intuition about the presence of algorithmic pricing, it is useful to consider a few illustrative examples. Figure 3 presents two such cases. A noticeable feature is that, during certain periods, prices change frequently from day to day. In some instances, these adjustments are small, with prices fluctuating from \$8.20 to \$8.01, \$8.05, \$8.34, \$8.05, and \$8.13 over a short interval. Such high-frequency price adjustments are consistent with the stylized patterns associated with algorithmic pricing (Aparicio et al. 2026).



(a) Eau de Parfum **(b)** Hair Shampoo

Figure 3: Examples of High-Frequency Price Changes

The implications of algorithmic pricing for gender-based price differences are theoretically ambiguous. On the one hand, if algorithmic pricing enables firms to better exploit differences in willingness to pay across gender-targeted segments, more dynamic pricing may widen the gender price gap. On the other hand, if algorithmic pricing makes prices more responsive to competition and substitution across similar products, greater pricing dynamics may compress the gap.

To operationalize this idea, we construct a product × week measure of algorithmic pricing using daily price changes. For each product, we count the number of day-to-day price changes within a given week.[6] We classify a product-week as *algorithmic* if the number of price changes exceeds two thresholds simultaneously: it must be greater than three price changes within the week, and it must exceed the product-specific median number of weekly price changes observed over the sample period. This definition captures periods in which pricing activity is high both in absolute terms and relative to the product's historical pattern. Section 8 reports robustness specifications.

We then estimate whether the gender price gap differs between low- and high-dynamic pricing environments by interacting the male-targeting indicator with the algorithmic pricing indicator:

$$\log Price_{it} = \alpha + \beta_1 Men_i + \beta_2 Algorithmic_{it} + \beta_3 \left(Men_i \times Algorithmic_{it}\right) + \gamma_{jkt} + \epsilon_{it}, \qquad (7.1)$$

where $Algorithmic_{it}$ equals one when product i exhibits high pricing dynamics at time t. The coefficient β_1 captures the gender price gap among low-dynamic products, while β_3 measures how that gap changes with algorithmic pricing.

[6] We exclude price changes associated with promotions, as these mechanically lower prices by construction.

Table 6: Price Effects and Algorithmic Pricing

	(1) $Price_{it}$	(2) $Price_{it}$	(3) $Price_{it}$	(4) $Price_{it}$	(5) $Price_{it}$
Men Targeted	-0.039***	-0.079***	-0.047***	-0.157***	-0.049***
	(0.001)	(0.002)	(0.002)	(0.001)	(0.002)
Algorithmic	-0.021***	-0.035***	-0.029***	-0.021***	-0.045***
	(0.002)	(0.003)	(0.002)	(0.002)	(0.005)
Men Target. × Algorithmic	0.023***	0.061***	0.008	0.028***	0.065***
	(0.005)	(0.005)	(0.006)	(0.004)	(0.006)
Constant	1.984***	1.699***	2.007***	1.636***	1.574***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	3,280,916	3,851,326	2,507,813	3,132,039	1,337,999
R-squared	0.791	0.828	0.795	0.872	0.882
Gender Inference	Amazon	AI	Intersection	Keywords	Matching AI
Firm × Categ. × Date FE	✓	✓	✓	✓	×
Match Pair FE	×	×	×	×	✓
Date FE	×	×	×	×	✓

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 6 reports the results. Across specifications, we again find that male-targeted products are priced lower than female-targeted products. In addition, products identified as using algorithmic pricing exhibit lower prices on average, consistent with more responsive pricing to competitive conditions. The interaction term between male targeting and algorithmic pricing is consistently positive and statistically significant, indicating that the gender price gap is smaller in environments with greater use of algorithmic pricing.

To illustrate magnitudes, consider Column (2), which uses AI-based gender classification. The coefficient on male-targeted products (-0.079) implies that, among products with low use of algorithmic pricing, male-targeted items are approximately 7.6% cheaper than female-targeted products. The interaction term (0.061) implies that this gap is substantially reduced among products with higher use of algorithmic pricing, decreasing to approximately $e^{-0.079+0.061} - 1 \approx -1.7\%$.

While the previous empirical strategy captures periods of high pricing activity, it does not explicitly account for changes in pricing regimes. Following Assad et al. (2024), we therefore complement our analysis by identifying structural breaks in pricing behavior at the product level. Intuitively, algorithmic pricing may reflect a shift in the pricing regime that alters how frequently prices are updated over time.

To implement this approach, we first construct a product × week measure of pricing activity

by counting the number of day-to-day price changes for each product. We then apply a structural break test to this series, allowing for changes in both the level and the trend of pricing activity. Specifically, for each product and each candidate break week, we estimate:

$$n_{it} = \alpha + \beta_1 t + \beta_2 Post_t + \beta_3 (Post_t \times t) + \epsilon_{it}, \tag{7.2}$$

where n_{it} denotes the number of price changes for product i in week t, and $Post_t$ is an indicator for periods after a candidate break. We evaluate all possible break points within the sample and select the week that maximizes the corresponding F-statistic (Assad et al. 2024). This yields a product-specific estimate of the timing at which pricing behavior changes most sharply.

We then map the estimated break week into daily data and define a post-break indicator, $Break_{it}$, equal to one for dates after the detected break and zero otherwise. To focus on local price adjustments around the change in pricing regime, we restrict the estimation sample to a symmetric window of 45 days before and 45 days after the estimated break date. This approach reduces the influence of longer-run price trends and isolates pricing behavior in the immediate vicinity of the regime change. Then we estimate a modified version of Eq. (7.1):

$$\log Price_{it} = \alpha + \beta_2 Break_{it} + \beta_3 \left(Men_i \times Break_{it}\right) + \gamma_{jkt} + \gamma_i + \epsilon_{it}, \tag{7.3}$$

where $Break_{it}$ equals one when product i is observed in the post-break period. Because we include product fixed effects, γ_i, the main effect of Men_i is absorbed. The coefficient β_2 captures the average price change after the estimated break, while β_3 measures how the gender price gap changes in the post-break period. A positive estimate of β_3 implies that the price gap between female- and male-targeted products narrows after the onset of algorithmic pricing.

Table 7 reports the results. Across all five specifications, the coefficient on the post-break indicator is negative and statistically significant, indicating that prices decline on average after the detected change in pricing behavior. More importantly, the interaction between male targeting and the post-break indicator is positive and statistically significant in every column. This pattern implies that the pre-existing price discount for male-targeted products becomes smaller after the structural break, consistent with a compression of the gender price gap. The magnitude of the effect is economically meaningful: depending on the gender-classification strategy, the interaction term ranges from 0.008 to 0.020 log points in the baseline specifications, and remains positive in the matched-sample specification.

Table 7: Price Effects and Algorithmic Pricing

	(1) $Price_{it}$	(2) $Price_{it}$	(3) $Price_{it}$	(4) $Price_{it}$	(5) $Price_{it}$
Algorithmic	-0.011***	-0.009***	-0.013***	-0.011***	-0.007***
	(0.003)	(0.003)	(0.004)	(0.003)	(0.003)
Men Target. × Algorithmic	0.020***	0.008***	0.017***	0.015***	0.011***
	(0.003)	(0.003)	(0.004)	(0.003)	(0.002)
Constant	2.293***	1.602***	2.186***	1.473***	1.665***
	(0.001)	(0.001)	(0.002)	(0.001)	(0.001)
Observations	56,404	53,652	31,970	44,965	46,848
R-squared	0.996	0.998	0.997	0.998	0.996
Gender Inference	Amazon	AI	Intersection	Keywords	Matching AI
Firm × Categ. × Date FE	✓	✓	✓	✓	×
Item FE	✓	✓	✓	✓	✓
Match Pair FE	×	×	×	×	✓
Date FE	×	×	×	×	✓

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Overall, these results indicate that algorithmic pricing is associated with a compression of gender-based price differences. This pattern suggests that more flexible and responsive pricing systems may limit the extent to which firms sustain persistent price differences across gender-targeted products.

8 Robustness Checks

We estimate several alternative specifications to confirm the robustness of our main analyses.

- We re-estimate the matching models using pairs of products sold at the same time period. The results are unchanged and are shown in Appendix A2.1.
- We re-estimate the price gap models in Eq. (4.1) using prices per unit without the log-transformation. The results are similar and are shown in Appendix A2.2.
- We re-estimate the price promotion models in Eq. (5.1) with an alternative discount definition. We use 30 days instead of 60 days, and reduce the promotion size to 10% instead of 20%. The results are similar and are reported in Appendix A2.3.

- We re-estimate the price promotion models in Eq. 5.1 using a linear probability model. The results are reported in Appendix A2.4. More specifically, we estimate the following model:

$$Y_{it}^{f,d} = \alpha + \beta Male_i + \gamma_{kjt} + \varepsilon_{it}, \tag{8.1}$$

  where $Y_{it}^{f,d}$ denotes the promotion indicator, and γ_{kjt} denotes firm $\times$ category $\times$ date fixed effects. We estimate this model for several price promotion definitions as well as for the specifications including gender-neutral targeted products.

- We re-define the indicator of algorithmic pricing at the product $\times$ month level using the number of daily price changes within each month. Specifically, we classify a product-month as algorithmic if the number of price changes exceeds both the product-specific median and a threshold of 10 price changes per month. We then re-estimate Eq. 7.1. The results, reported in Appendix A2.6, are qualitatively similar.

9 Conclusion

This paper documents novel evidence of systematic gender-based price differences in a large online marketplace using high-frequency data on beauty and personal care products sold on Amazon. We find that male-targeted products are priced lower than comparable female-targeted products, and that female-targeted products are more likely to be promoted. Extending the analysis beyond the binary comparison emphasized in most of the existing literature, we show that gender-neutral products tend to occupy an intermediate position in both prices and promotions. Finally, we provide evidence that algorithmic pricing is associated with a compression of gender-based price differences, suggesting that more flexible pricing systems may reduce the persistence of these gaps in digital retail environments.

These findings have several implications for policy and for the broader study of pricing in digital markets. First, they suggest that assessments of the Pink Tax should account not only for posted prices but also for promotional activity, since effective price differences may be shaped by both regular pricing and temporary discounts. Second, the results on gender-neutral products indicate that a more inclusive perspective on gender pricing is increasingly important in markets where firms move beyond strictly binary product positioning. Third, the evidence on algorithmic pricing highlights that new pricing technologies may affect not only average prices but also how prices differ across consumer segments. As policymakers continue to evaluate gender-based pricing and the growing use of automated pricing tools, understanding the interaction between forms of price discrimination, promotions, and pricing technology will be essential for designing effective and forward-looking regulation.

Appendix

A1 Data Description

Our analysis uses product and price data for items sold on Amazon.com obtained from *Keepa.com*, a price-tracking platform that continuously monitors and records historical information for products listed on Amazon. Keepa collects data directly from Amazon product pages and stores time-stamped observations whenever key variables change. As a result, price observations correspond to recorded price changes and therefore occur at irregular intervals over time.

For each product, the dataset includes a unique product identifier (ASIN), brand name, category classification, historical prices, sales rank, and product attributes such as package size, quantity, and weight. To construct a consistent panel dataset, we aggregate the event-based price history to the daily level. Thus, the unit of analysis is: ASIN $\times$ Date. If multiple price observations occur within the same day, we use the median price as the representative daily price. For days without a recorded price change, we carry forward the most recent observed price. This procedure produces a balanced product-day panel while preserving the underlying timing of price adjustments.

A2 Robustness Specifications

A2.1 Robustness: Matched Products

We re-estimate Eq. (4.2) and (5.1) for the matching sample, restricting it to include only those pairs that are listed during the same time period. Table A1 reports unchanged results.

Table A1: Robustness: Matched Products and Same Time Period

	(1) $Price_{it}$	(2) $Prom_{it}^{60,20}$
Men Targeted	-0.031***	-0.013***
	(0.001)	(0.000)
Constant	1.626***	0.138***
	(0.001)	(0.000)
Observations	3,692,641	3,692,641
R-squared	0.876	0.111
Gender Inference	Matching AI	Matching AI
Match Pair FE	✓	✓
Categ. × Date FE	✓	✓

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

A2.2 Robustness: Price Effects

We re-estimate Eq. (4.1) and (4.2) using prices per unit without a log-transformation. Table A2 reports similar results.

Table A2: Robustness: Price Effects

	(1) $Price_{it}$	(2) $Price_{it}$	(3) $Price_{it}$	(4) $Price_{it}$	(5) $Price_{it}$
Men Targeted	-2.305***	-1.754***	-2.282***	-1.839***	-2.413***
	(0.035)	(0.025)	(0.037)	(0.029)	(0.025)
Constant	18.006***	17.843***	17.509***	16.107***	16.056***
	(0.018)	(0.017)	(0.021)	(0.019)	(0.013)
Observations	3,280,916	3,851,326	2,507,813	3,132,039	1,383,955
R-squared	0.670	0.750	0.650	0.735	0.835
Gender Inference	Amazon	AI	Intersection	Keywords	Matching AI
Firm × Categ. × Date FE	✓	✓	✓	✓	×
Match Pair FE	×	×	×	×	✓
Categ. × Date FE	×	×	×	×	✓

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

A2.3 Robustness: Price Promotions using Alternative Definitions

We re-define a price promotion using the last 30 days and a 10% discount threshold. Thus, a promotion indicator $Y_{it}^{30,10}$ equals 1 if the current price is at least 10% lower than this reference price, and 0 otherwise. Table A3 reports the results.

Table A3: Robustness: Price Promotions (30 Days and 10%)

	(1) $Y_{it}^{30,10}$	(2) $\mathrm{d}Y_{it}^{30,10}$	(3) $Y_{it}^{30,10}$	(4) $Y_{it}^{30,10}$	(5) $Y_{it}^{30,10}$
Men Targeted	-0.079***	-0.069***	-0.029***	-0.055***	-0.051***
	(0.005)	(0.005)	(0.006)	(0.006)	(0.004)
Observations	2,038,292	2,334,507	1,491,011	1,843,458	419,138
Gender Inference	Amazon	AI	Intersection	Alternative	Matching AI
Firm × Categ. × Date FE	✓	✓	✓	✓	×
Match Pair FE	×	×	×	×	✓

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

A2.4 Robustness: Price Promotions using Linear Probability Model

As robustness, we estimate the likelihood of a price promotion following Eq. (8.1) using a linear probability model.

First, we report the results for male- versus female-targeted products. Table A4 reports estimates when discounts are defined using the maximum price in last 60 days and a discount of at least 20%, Table A5 when discounts are defined using the maximum price in last 30 days and a discount of at least 10%, and Table A6 when discounts are defined using the 90th percentile price in last 60 days and a discount of at least 10%.

Table A4: Robustness: Price Promotions using LPM (60 Days and 20%)

	(1) $Y_{it}^{60,20}$	(2) $Y_{it}^{60,20}$	(3) $Y_{it}^{60,20}$	(4) $Y_{it}^{60,20}$	(5) $Y_{it}^{60,20}$
Men Targeted	-0.181***	-0.135***	-0.189***	-0.142***	-0.142***
	(0.005)	(0.005)	(0.006)	(0.005)	(0.005)
Observations	3,280,916	3,851,326	2,507,813	3,132,039	3,982,248
R-squared	0.268	0.293	0.288	0.291	0.110
Gender Inference	Amazon	AI	Intersection	Keywords	Matching AI
Firm × Categ. × Date FE	✓	✓	✓	✓	✓
Match Pair FE	×	×	×	×	✓

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table A5: Robustness: Price Promotions using LPM (30 Days and 10%)

	(1) $Y_{it}^{30,10}$	(2) $Y_{it}^{30,10}$	(3) $Y_{it}^{30,10}$	(4) $Y_{it}^{30,10}$	(5) $Y_{it}^{30,10}$
Men Targeted	-0.018***	-0.010***	-0.016***	-0.012***	-0.009***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.000)
Constant	0.199***	0.189***	0.200***	0.195***	0.184***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Observations	3,280,916	3,851,326	2,507,813	3,132,039	3,982,248
R-squared	0.270	0.296	0.289	0.293	0.104
Gender Inference	Amazon	AI	Intersection	Keywords	Matching AI
Firm × Categ. × Date FE	✓	✓	✓	✓	✓
Match Pair FE	×	×	×	×	✓

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table A6: Robustness: Price Promotions using LPM (90th, 60 Days and 10%)

	(1) $Y_{it}^{60,10}$	(2) $Y_{it}^{60,10}$	(3) $Y_{it}^{60,10}$	(4) $Y_{it}^{60,10}$	(5) $Y_{it}^{60,10}$
Men Targeted	-0.020***	-0.015***	-0.019***	-0.018***	-0.011***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.000)
Constant	0.217***	0.205***	0.218***	0.212***	0.199***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Observations	3,280,916	3,851,326	2,507,813	3,132,039	3,982,248
R-squared	0.267	0.295	0.286	0.293	0.102
Gender Inference	Amazon	AI	Intersection	Keywords	Matching AI
Firm × Categ. × Date FE	✓	✓	✓	✓	✓
Match Pair FE	×	×	×	×	✓

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Finally, we report the results including gender-neutral products. Once again, female-targeted products are the baseline segment. The estimates are shown in Table A7.

Table A7: Robustness: Price Promotions using LPM (60 Days and 20%)

	(1) $Y_{it}^{60,20}$	(2) $Y_{it}^{60,20}$
Gender-Neutral Targeted	-0.006***	-0.011***
	(0.000)	(0.000)
Men Targeted	-0.015***	-0.013***
	(0.001)	(0.001)
Constant	0.126***	0.130***
	(0.000)	(0.000)
Observations	12,787,511	12,787,511
R-squared	0.295	0.295
Gender Inference	AI	Keywords
Firm × Categ. × Date FE	✓	✓

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

A2.5 Price Promotion Depth

For the analysis of promotion depth (size of a discount), we estimate the model in Eq. (5.2). The baseline results using the 90th percentile over the last 60 days and a minimum discount of 10% are shown in Table A8.

Table A8: Promotion Depth

	(1) Δ_{it}^{60}	(2) Δ_{it}^{60}	(3) Δ_{it}^{60}	(4) Δ_{it}^{60}	(5) Δ_{it}^{60}
Men Targeted	-0.003***	0.005***	0.003***	0.006***	0.007***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Constant	0.251***	0.249***	0.248***	0.242***	0.247***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Observations	506,380	520,633	374,143	438,102	281,343
R-squared	0.395	0.424	0.431	0.402	0.236
Gender Inference	Amazon	AI	Intersection	Keywords	Matching AI
Firm × Cat. × Date FE	✓	✓	✓	✓	×
Match Pair FE	×	×	×	×	✓
Date FE	×	×	×	×	✓

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

As robustness, we re-estimate the same models using a minimum discount of 15%. The results are shown in Table A9.

Table A9: Promotion Depth

	(1) Δ_{it}^{60}	(2) Δ_{it}^{60}	(3) Δ_{it}^{60}	(4) Δ_{it}^{60}	(5) Δ_{it}^{60}
Men Targeted	-0.004***	0.003***	0.001	0.006***	0.007***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Constant	0.295***	0.293***	0.293***	0.285***	0.290***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Observations	355,995	357,812	258,155	298,593	210,921
R-squared	0.419	0.449	0.459	0.428	0.261
Gender Inference	Amazon	AI	Intersection	Keywords	Matching AI
Firm × Cat. × Date FE	✓	✓	✓	✓	×
Match Pair FE	×	×	×	×	✓
Date FE	×	×	×	×	✓

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

A2.6 Robustness: Algorithmic Pricing

As robustness, we re-estimate Eq. (7.1) when inferring the use of algorithmic pricing using a high frequency of daily price changes at the product × month level. The results are similar and are reported in Table A10.

Table A10: Price Effects and Algorithmic Pricing

	(1) $Price_{it}$	(2) $Price_{it}$	(3) $Price_{it}$	(4) $Price_{it}$	(5) $Price_{it}$
Men Targeted	-0.040***	-0.080***	-0.046***	-0.158***	-0.052***
	(0.001)	(0.002)	(0.002)	(0.001)	(0.002)
Algorithmic	-0.022***	-0.038***	-0.030***	-0.016***	-0.057***
	(0.002)	(0.003)	(0.002)	(0.002)	(0.005)
Men Target. × Algorithmic	0.020***	0.064***	-0.005	0.026***	0.079***
	(0.004)	(0.005)	(0.005)	(0.004)	(0.007)
Constant	1.985***	1.700***	2.008***	1.636***	1.669***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	3,280,916	3,851,326	2,507,813	3,132,039	1,170,490
R-squared	0.791	0.828	0.795	0.872	0.876
Gender Inference	AI	AI	Intersection	Keywords	Matching AI
Firm × Categ. × Date FE	✓	✓	✓	✓	×
Match Pair FE	×	×	×	×	✓
Date FE	×	×	×	×	✓

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

References

Alberto F. Alesina, Francesca Lotti, and Paolo Emilio Mistrulli. Do women pay more for credit? evidence from italy. *Journal of the European Economic Association*, 11(suppl_1):45–66, 2013.

Steffen Andersen, Julie Marx, Kasper Meisner Nielsen, and Lise Vesterlund. Gender differences in negotiation: Evidence from real estate transactions. *The Economic Journal*, 131(638): 2304–2332, 2021. doi: 10.1093/ej/ueab027. URL `https://www.nber.org/papers/w27318`.

Eric T Anderson and James D Dana Jr. When is price discrimination profitable? *Management Science*, 55(6):980–989, 2009.

Simon Anderson, Alicia Baik, and Nathan Larson. Price discrimination in the information age:

Prices, poaching, and privacy with personalized targeted discounts. *The Review of Economic Studies*, 90(5):2085–2115, 2023. doi: 10.1093/restud/rdac073.

AP News Wire. Uk abolishes 'sexist' tax on women's sanitary products. *The Independent*, 2021. URL https://www.independent.co.uk/news/uk-abolishes-sexist-tax-on-womens-sanitary-products-ireland-tax-tampon-tax-women-women-b1html.

Diego Aparicio and Kanishka Misra. Artificial intelligence and pricing. In K. Sudhir and O. Toubia, editors, *Artificial Intelligence in Marketing*, pages 103–124. Emerald Publishing Limited, Bingley, 2022.

Diego Aparicio, Zachary Metzman, and Roberto Rigobon. The pricing strategies of online grocery retailers. *Quantitative Marketing and Economics*, 22(1):1–21, 2024.

Diego Aparicio, Dean Eckles, Madhav Kumar, and Catherine Tucker. Algorithmic pricing and consumer price sensitivity. *Working Paper*, 2026.

Stephanie Assad, Robert Clark, Daniel Ershov, and Lei Xu. Algorithmic pricing and competition: Empirical evidence from the german retail gasoline market. *Journal of Political Economy*, 132 (3):723–771, 2024.

Kayleigh Barnes and Jakob Brounstein. The pink tax: Why do women pay more? *Kilts Center at Chicago Booth Marketing Data Center Paper*, 2025.

Marianne Bertrand, Claudia Goldin, and Lawrence F. Katz. Dynamics of the gender gap for young professionals in the financial and corporate sectors. *American Economic Journal: Applied Economics*, 2(3):228–255, 2010.

Francine D. Blau and Lawrence M. Kahn. The gender wage gap: Extent, trends, and explanations. *Journal of Economic Literature*, 55(3):789–865, September 2017. doi: 10.1257/jel.20160995. URL https://www.aeaweb.org/articles?id=10.1257/jel.20160995.

Meghan R. Busse, Ayelet Israeli, and Florian Zettelmeyer. Repairing the damage: The effect of gender and price knowledge on auto-repair price quotes. *Journal of Marketing Research*, 54 (1):75–95, 2017.

California State Legislature. Cal. Civ. Code § 51.14 (2022). https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=202120220AB1287, 2022. Assembly Bill No. 1287, Ch. 555, approved Sept. 27, 2022.

Margaret C Campbell. Perceptions of price unfairness: antecedents and consequences. *Journal of Marketing Research*, 36(2):187–199, 1999.

Le Chen, Alan Mislove, and Christo Wilson. An empirical analysis of algorithmic pricing on amazon marketplace. In *Proceedings of the 25th International Conference on World Wide Web (WWW '16)*, pages 1339–1349, Montreal, Canada, 2016. International World Wide Web Conferences Steering Committee. doi: 10.1145/2872427.2883089.

Maxime C. Cohen, Adam N. Elmachtoub, and Xiao Lei. Price discrimination with fairness constraints. *Management Science*, 68(11):7914–7938, 2022. doi: 10.1287/mnsc.2022.4317.

Emily Crockett. Women's products cost more. here's how to avoid the "pink tax.". *Vox*, March 2016. URL https://www.vox.com/2016/3/1/11139280/avoid-pink-tax-women. Accessed: 2025-08-22.

Megan Duesterhaus, Liz Grauerholz, Rebecca Weichsel, and Nicholas A. Guittar. The cost of doing femininity: Gendered disparities in pricing of personal care products and services. *Gender Issues*, 28(4):175–191, 2011. doi: 10.1007/s12147-011-9106-3. URL https://doi.org/10.1007/s12147-011-9106-3.

Wedad Elmaghraby and Pınar Keskinocak. Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions. *Management science*, 49 (10):1287–1309, 2003.

Federal Reserve Bank of St. Louis. E-commerce retail sales (ecomsa). https://fred.stlouisfed.org/series/ECOMSA, 2025. Units: Millions of Dollars, Seasonally Adjusted. Accessed: September 14, 2025.

Federal Trade Commission. Ftc surveillance pricing study indicates wide range of personal data used to set individualized consumer prices, 2025.

Pinelopi Koujianou Goldberg. Dealer price discrimination in new car purchases: Evidence from the consumer expenditure survey. *Journal of Political Economy*, 104(3):622–654, 1996.

Claudia Goldin. *Understanding the gender gap: An economic history of American women.* New York, 1990.

Claudia Goldin. Nobel lecture: An evolving economic force. *American Economic Review*, 114(6): 1515–1539, 2024.

Claudia Goldin, Sari Pekkala Kerr, Claudia Olivetti, and Erling Barth. The expanding gender earnings gap: Evidence from the lehd-2000 census. *American Economic Review*, 107(5):110–14, May 2017. doi: 10.1257/aer.p20171065. URL https://www.aeaweb.org/articles?id=10.1257/aer.p20171065.

Paul Goldsmith-Pinkham and Kelly Shue. The gender gap in housing returns. *Journal of Finance*, 78(2):1097–1145, 2023. doi: 10.3386/w26914. URL https://www.nber.org/papers/w26914.

Stephanie Gonzalez Guittar, Liz Grauerholz, Shameika D. Daye, and Megan McLaughlin. Beyond the pink tax: Gender-based pricing and differentiation of personal care products. *Gender Issues*, 39(3):284–305, 2021. doi: 10.1007/s12147-021-09285-x.

John P Harding, Stuart S Rosenthal, and C. F. Sirmans. Estimating bargaining power in the market for existing homes. *Review of Economics and Statistics*, 85(1):178–188, 2003. doi: 10.1162/003465303762687741.

David W Harless and George E Hoffer. Do women pay more for new vehicles? evidence from transaction price data. *American Economic Review*, 92(1):270–279, 2002.

Susan K. Harmon and C. Jeanne Hill. Gender and coupon use. *Journal of Product & Brand Management*, 12(3):166–179, 2003. doi: 10.1108/10610420310476924.

John W Huppertz, Sidney J Arenson, and Richard H Evans. An application of equity theory to buyer-seller exchange situations. *Journal of marketing research*, 15(2):250–260, 1978.

Daniel Kahneman, Jack L Knetsch, and Richard Thaler. Fairness as a constraint on profit seeking: Entitlements in the market. *The American economic review*, pages 728–741, 1986.

Hayne E Leland and Robert A Meyer. Monopoly pricing structures with imperfect discrimination. *The Bell Journal of Economics*, pages 449–462, 1976.

Fiona Scott Morton, Florian Zettelmeyer, and Jorge Silva-Risso. Consumer information and discrimination: Does the internet affect the pricing of new cars to women and minorities? *Quantitative marketing and Economics*, 1(1):65–92, 2003.

Sarah Moshary, Anna Tuchman, and Natasha Vajravelu. Gender-based pricing in consumer packaged goods: A pink tax? *Marketing Science*, 07 2023. doi: 10.1287/mksc.2023.1452.

Leon Musolff. Algorithmic pricing, price wars and tacit collusion: Evidence from e-commerce. *Working Paper*, 2026.

New York City Department of Consumer Affairs. From cradle to cane: The cost of being a female consumer, 2015. URL `https://www1.nyc.gov/assets/dca/downloads/pdf/partners/Study-of-Gender-Pricing-in-NYC.pdf`.

New York State Senate. §391-u. pricing goods and services on the basis of gender prohibited, 2020. URL `https://www.nysenate.gov/legislation/laws/GBS/391-U`.

New York State Senate. Senate bill s7033: Preventing algorithmic pricing discrimination act. `https://www.nysenate.gov/legislation/bills/2025/S7033`, 2025. 2025–2026 Legislative Session, In Committee (Consumer Protection).

New York Times. Tampons to be taxed as essential, not luxury, items in germany. *The New York Times*, 2019. URL `https://www.nytimes.com/2019/11/12/world/europe/tampon-tax-germany.html`. Accessed July 21, 2025.

New York Times. Texas is the latest state to drop the 'tampon tax'. *The New York Times*, 2023. URL `https://www.nytimes.com/2023/09/01/well/live/tampon-tax-texas.html`. Accessed 2025-07-21.

NielsenIQ. Shaping success: A deep dive into women's impact on the cpg landscape. `https://nielseniq.com/global/en/insights/analysis/2024/shaping-success-a-deep-dive-into-womens-impact-on-the-cpg-landscape/#niq-tab-1d8c3a6c-37fb-4f98-9061-2c7c7b475d99`, 2024. Accessed: 2025-09-12.

O. Cem Ozturk, Cheng He, and Pradeep K. Chintagunta. Inequalities in dealers' interest rate markups? a gender and race-based analysis. Available at SSRN: `https://ssrn.com/abstract=4224944`, September 2022. URL `http://dx.doi.org/10.2139/ssrn.4224944`.

Richard Schmalensee. Output and welfare implications of monopolistic third-degree price discrimination. *The American Economic Review*, 71(1):242–247, 1981.

Liesbet Sommen. Question for written answer e-000379/2025 to the commission: Gender-based

price discrimination (pink tax), 2025. URL https://www.europarl.europa.eu/doceo/document/E-10-2025-000379_EN.pdf. Parliamentary question under Rule 144.

The Hill. Speier introduces bill to end the 'pink tax', 2018.

Time. Meet the woman who's calling out companies for charging women more. *TIME*, 2016. URL https://time.com/4245619/pink-tax-study/. Accessed: 2025-09-01.

Lan Xia, Kent B Monroe, and Jennifer L Cox. The price is unfair! a conceptual framework of price fairness perceptions. *Journal of marketing*, 68(4):1–15, 2004.