What's In the Box? The Effect Of Self-Preferencing On Amazon

Florian Dendorfer & Regina Seibel*

April 11, 2025

Abstract

We examine product recommendations in Amazon's "Similar items to consider" box in the US and Canada, finding evidence of self-preferencing in Canada. In our dataset, alternatives to Amazon Basics (AB) products are sometimes recommended in the US but never in Canada, while non-AB products are sometimes recommended in Canada but never in the US. By comparing sales across domains, we find that non-AB products not recommended in Canada due to self-preferencing experience a 22% sales decrease compared to those that are not exposed to self-preferencing in the same way.

1 Introduction

Economists and policymakers, particularly in Europe, warn of the potential harm caused by self-preferencing, which is broadly defined as the preferential treatment of a gatekeeper platform's own products and services over third-party offerings, often through higher visibility or better placement. Cabral, Haucap, Parker, Petropoulos, Valletti and Van Alstyne (2021) argue that "self-preferencing [should be] deemed anti-competitive and 'per se' disallowed" (p.13), a position supported to varying degrees by the theoretical literature on self-preferencing (De Corniere and Taylor, 2019; Dendorfer, 2024; Hagiu, Teh and Wright, 2022; Padilla, Perkins and Piccolo, 2022; Tirole and Bisceglia, 2023). However, identifying self-preferencing and quantifying

*Dendorfer: Department of Economics, University of Toronto. Seibel: Rotman School of Management, University of Toronto. Email: florian.dendorfer@utoronto.ca, and regina.seibel@rotman.utoronto.ca.

its effects is notoriously difficult. As Peitz (2023) notes, "it can be challenging to detect self-preferencing bias as opposed to legitimate differential treatment" (p.315). In Amazon's case, first-party products may be shown to consumers due to self-preferencing or simply because consumers prefer them over third-party alternatives. A nascent literature attempts to disentangle these two causes. Notably, Chen and Tsai (2023) have made significant inroads by leveraging quasi-random variation in first-party product availability.

We contribute to this literature by describing the differential treatment of products sold under Amazon's largest brand, Amazon Basics (AB), and non-AB products in the "Similar items to consider" recommendation box across Amazon's US and Canadian domains. Moreover, we provide a conservative estimate for the effect this differential treatment has on directly affected products.

To this end, we scraped product pages for 14,610 products on Amazon, from both the US and Canadian domains. We intentionally selected product groups that include Amazon Basics items. We find that average ratings and prices are very similar across domains, while products have fewer reviews and lower rankings in Canada which is consistent with the smaller size of the Canadian market. In both domains, sales rank is correlated with whether the product is labeled Amazon's Choice, whether it is sold by Amazon, and its price. We find a notably higher degree of self-preferencing in Canada: while both AB and non-AB products appear in the recommendation box in the US, non-AB products never do in Canada. At the same time, US consumers browsing an AB product page may be offered an alternative product in this box, whereas Canadian consumers will not encounter such a box. Our regression analysis confirms that the single most important factor for being in the box or having the box is whether a product is an AB product or not.

We next assess whether this form of self-preferencing effectively steers consumer behavior. Assuming that, absent self-preferencing, the relative sales difference for non-AB products across markets is unrelated to whether a product is recommended in the US, we estimate the effect of self-preferencing on non-AB products that are no longer featured in the recommendation box, as well as on AB products that lose the box on their page. Our results indicate that non-AB products subject to self-preferencing, being recommended in the US but not in Canada, experience a 22% decrease in sales relative to non-AB products that are not recommended in either domain. This estimate can be interpreted as a lower bound of the effect since also products that are not recommended in both domains might be negatively affected by Amazon unfairly promoting its own products. We do not find a significant sales difference for AB products that stand to benefit from self-preferencing — where an alternative product is recommended in the US but not in Canada

— compared to AB products for which no alternative is recommended in either domain.

Lastly, we discuss potential reasons for this type of self-preferencing in Canada. First, AB products may be generally less popular in Canada which may increase Amazon's incentive to engage in self-preferencing. Second, regulatory pressure may appear higher in the US compared to Canada in the time period studied. The higher degree of self-preferencing leads to potentially large distortions in consumer behavior: Simple back-of-the-envelope calculations suggest that non-AB sellers in the US would have sold 120 million fewer units if they were exposed to the same degree of self-preferencing as in Canada.

The remainder of this paper is organized as follows. Section 2 reviews the related literature. Section 3 describes the data. Section 4 outlines the estimation strategy and presents the results. Section 5 discusses the validity of the results and Section 6 concludes.

2 Literature

The incentives for platforms to engage in self-preferencing, as well as its impact on the platform-run marketplace, have been extensively explored in the theoretical literature (De Corniere and Taylor, 2019; Dendorfer, 2024; Hagiu et al., 2022; Padilla et al., 2022; Tirole and Bisceglia, 2023; Zennyo, 2022). This literature generally finds that self-preferencing is anti-competitive and harms consumers (with Zennyo (2022) being the exception), but offers diverging views on whether self-preferencing is ultimately profitable for the platform.

A fast-growing empirical literature seeks to identify evidence of self-preferencing and document its effects on Amazon Marketplace (Chen and Tsai, 2023; Farronato, Fradkin and MacKay, 2023; Lee and Musolff, 2021; Raval, 2022; Waldfogel, 2024). Much of this work is surveyed in Etro (2024). Farronato et al. (2023) use consumer search data from Amazon Marketplace and find that Amazon-branded products rank higher in search results than comparable non-AB products, although it is unclear whether this is due to self-preferencing or to consumers' inherent preference for Amazon's own brands. Waldfogel (2024) finds that Amazon-branded products have lower sales ranks and argues that self-preferencing must account for at least part of this advantage. He notes that their average sales rank worsened significantly after Amazon was designated a "gatekeeper" under the Digital Markets Act, interpreting this shift as a sign of increased regulatory compliance by Amazon.

In contrast to these studies, we document differential treatment of the *same* product across geographic domains, with respect to product recommendations. We show that Amazon-branded products are treated differently from products of other brands in the Canadian domain compared to the US domain of Amazon Marketplace, a difference that cannot be explained by observable factors such as the price or the average rating. Similar to Waldfogel (2024), we suspect that this difference is due, at least in part, to varying levels of antitrust scrutiny between countries.

Lee and Musolff (2021) and Raval (2022) study products sold both by Amazon and by third-party vendors, finding that Amazon's offers are disproportionally selected into the "Buy Box". Lee and Musolff (2021) estimate a structural model of Amazon's marketplace and, through one of their main counterfactuals, determine that the bias toward Amazon largely reflects consumers' legitimate preferences for Amazon as a seller, rather than unfair steering. In contrast, we document that Amazon Basics products receive more favorable treatment in product recommendations in the Canadian domain than in the US domain which cannot easily be attributed to consumer preferences.

Lastly, Chen and Tsai (2023) examine "Frequently Bought Together" product recommendations on Amazon Marketplace. They find that a product is, on average, eight percentage points less likely to be recommended when Amazon stocks out of a complementary product, thus losing the incentive to promote the product. Relatedly, we find that Amazon favors its own products over non-AB products in the "Similar items to consider" recommendations. Unlike Chen and Tsai (2023), we estimate the sales changes sellers experience as a result of self-preferencing.

3 Data

This section is structured as follows. Section 3.1 describes our data collection process and the contents of our data sample. In Section 3.2, we showcase the data by replicating findings from the existing literature. Section 3.3 documents evidence of self-preferencing in the "Similar items to consider" box.

3.1 Data Collection

Between January 24 and March 11, 2025, we scraped product pages for 14,610 Amazon Standard Identification Numbers (ASINs) from both the US (Amazon.com) and Canadian (Amazon.ca) domains of the Amazon Marketplace

website. We obtained the ASINs from the first search result page on Amazon Marketplace for each of 200 search terms (e.g., "Clothing Steamers") that we suspect includes at least one product sold under Amazon's largest private label, Amazon Basics (AB). Although Amazon sells under other brands as well, virtually all private label products in our data sample are AB products.[1]

We collected the Buy Box price (in USD and CAD, respectively) excluding delivery fees, the number of reviews, and the average rating on a 1-to-5 star scale. Additionally, we recorded whether a product was designated as "Amazon's Choice", whether it was an AB product, and its sales rank.

A single product can have multiple sales ranks, one for each category it belongs to. Typically, we observe a broad category (e.g., "Electronics") along with at least one more specific category (e.g., "Home Cinema Cables"). In total, we identified 495 categories, many of which include AB products. To give the reader a sense of the categories included in our dataset, Table A.1 in the Supplemental Appendix lists the twenty largest categories, along with the total number of products and the number of AB products in each.

Apart from sales ranks, we collected sales data. Product pages occasionally provide a lower bound on recent sales, such as "700+ bought in the past month" or "2k+ bought in the past month." There are 36 different such sales indicators: x00+, xK+, x0K+, and x00K+ for $x \in \{1, 2, 3, 4, 5, 6, 7, 8, 9\}$. We observe sales information of this kind for 55% of observations in our data sample. We acknowledge that, due to censoring, the reported sales data are systematically understated, and we view them as conservative estimates of actual sales. Nonetheless, we believe th quality of the sales data is sufficiently high to use them in our analysis in Section 4.

Lastly, we record whether a product page includes a "Similar item to consider" box, sometimes labeled "Consider a similar item" or "Competitively priced item", along with the ASIN of the product featured in the box. Figure 1 illustrates how this box appears on a product page. The feature is widespread: 31% of Canadian product pages in our sample contain the box, compared to 53% of US pages. Among products in our sample, 4% appear in the box in Canada, while this share is 9% for the US.

For many products, we can identify both the (Buy Box-winning) seller and the shipper. All products for which we register the shipper are shipped or "fulfilled" by Amazon (FBA).[2] Similarly, Gutierrez (2021) reports that, in his dataset, the vast majority of products (81%) are shipped by Amazon. In

[1] A small number of products are sold under Amazon's Rivet and Amazon Commercial brands. We treat them as Amazon Basics products in our analysis.

[2] For the Canadian domain, we successfully scraped the shipper information for 78% of product pages. Although we were unable to scrape the shipper information for the US domain, we believe that most, if not all of the products are FBA too.

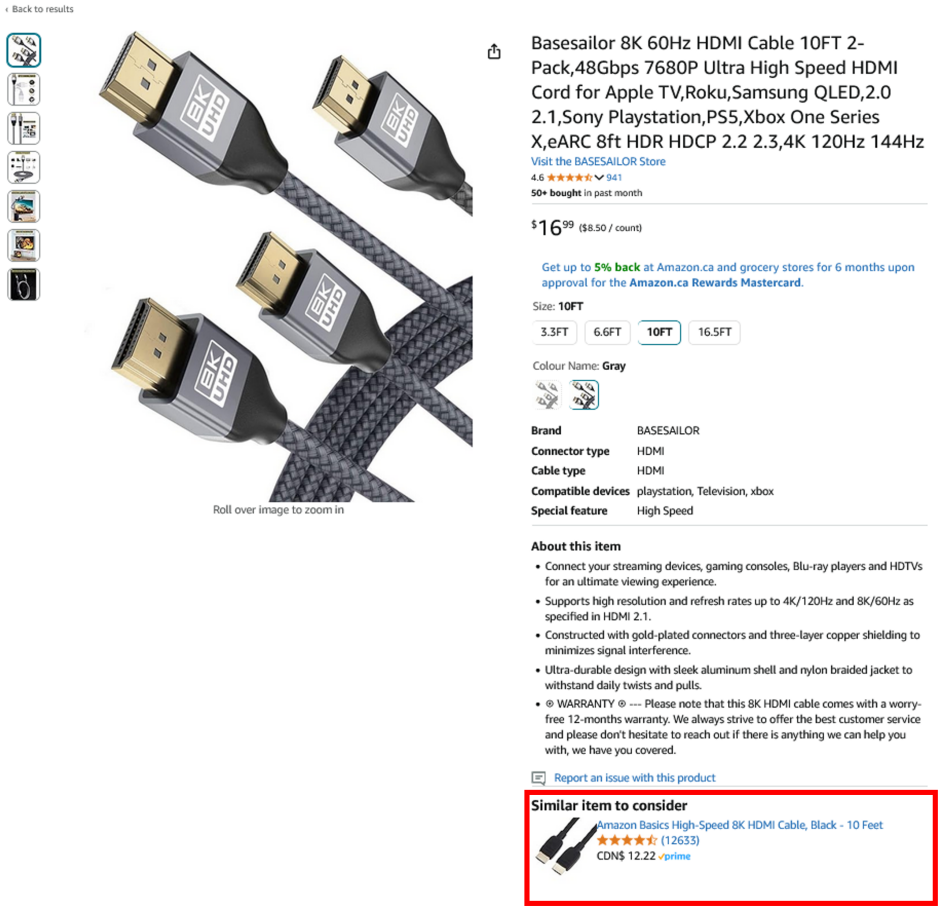

Figure 1: "Similar item to consider" box (red frame added).

our dataset, the seller with the largest number of products is by far Amazon itself. Third-party sellers tend to be small; half of them sell two or fewer products. In our sample, 5% of products are AB-branded, which coincides with the 5% share of first-party private label sales reported in Gutierrez (2021), who considers a broader range of products. With 32%, the share of third-party products sold by Amazon in our sample is similar to the one found by Gutierrez (2021) too.

3.2 Descriptives

Table 1 presents sample averages for AB and non-AB products in the Canadian and US domains for various variables. AB products differ from non-AB products. On average, they are rated higher and are approximately 30% cheaper than non-AB products in both domains. This aligns with findings

by Farronato et al. (2023), who report similar statistics for their data sample. While Farronato et al. (2023) find that Amazon-branded products have about four times more reviews than other products, our sample shows that AB products have, on average, around eight times more reviews compared to non-AB products.

	CA		US	
	non-AB	AB	non-AB	AB
Avg. rating	4.45 stars	4.47 stars	4.49 stars	4.55 stars
Review count	5,434	44,697	6,256	52,529
Avg price	62.48 CAD	41.84 CAD	40.75 USD	27.38 USD
Rank	26,460	9,817	38,922	17,411
Rank ("smallest" category)	10,265	2,548	15,772	3,544
No. products	13,978	632	13,978	632

Table 1: Summary statistics by domain and product type (AB vs non-AB)

Compared to non-AB products, AB products have on average 63% and 55% lower sales ranks than non-AB products, in Canada and the US respectively. This gap is smaller in Farronato et al. (2023), where the reported difference is 24%. Notably, sales ranks in our data are generally much larger, likely because our sample includes broad categories such as "Clothing" or "Tools & Home Improvement". When focusing on the category associated with the smallest rank for each product, the average rank decreases in both domains, while the difference between AB and non-AB products widens. All of this suggest that consumers *could* plausibly prefer AB to non-AB products. As AB products may enjoy greater prominence for this reason, it is difficult to establish self-preferencing and isolate its effects. Table 1 also highlights that the Canadian and US domains are reasonably similar. Difference in the average prices between domains comes down to the exchange rate. Unsurprisingly, the Canadian market is relatively smaller. The average sales rank and the average review count are respectively 32% and 15% lower than in the US.

Next, we investigate whether the average sales rank of AB products differs from that of non-AB products, conditional on observables. For this purpose as well as all regressions that follow, we account for differences across. Since a product may be associated with multiple categories, we define observations on the ASIN-category-domain level, but weigh each observation in accordance with the number of observations belonging to the same ASIN and domain.[3] Accordingly, we cluster the standard errors of our estimates on

[3]Let n_{am} be the observations with ASIN a in domain m. We weigh each observation with n_{am}^{-1} in all regressions.

the category-domain level. Note that our dataset consists of 57,620 observations which cover 14,610 products in each of two domains, totaling 29,220 product-domain observations. Table 2 presents the estimated coefficients from a regression of the logged sales rank on the following variables: an indicator for whether the product is AB-branded, an indicator for whether the product is labeled Amazon's Choice, an indicator for whether it is sold by Amazon (1P), the logged price, the average rating, the logged review count, their interaction, and category-domain fixed effects. We do not find evidence that AB products have a lower sales rank than non-AB products, all else equal.[4] The other coefficients are largely as expected. We estimate that, *ceteris paribus*, a 1% increase in price is associated with an 17% higher sales rank. The Amazon's Choice label is linked to a substantial 58% reduction in the sales rank. Higher ratings coincide with a lower sales rank if accompanied by a sufficient number of reviews. More reviews tend to coincide with a lower sales rank, likely because both the sales rank and the review count reflect past sales. Somewhat surprisingly, first-party selling is associated with a higher sales rank. We emphasize that none of these coefficients should be interpreted causally. The estimated AB coefficient, in particular, does not provide evidence for or against self-preferencing. It may reflect consumer preferences as well as potential self-preferential behavior by Amazon. The statistically non-significant AB coefficient could, for example, mask that Amazon employs self-preferencing to boost poor-selling AB products.

3.3 Self-Preferencing

Table 3 demonstrates the extent to which the "Similar items to consider" box is assigned to product pages and products are assigned to the box in each domain in our data sample, depending on whether they are AB or non-AB products.

None of the non-AB products in Canada are promoted in the "Similar items to consider" box (referred to as "being in the box" for brevity). In contrast, about 6% of non-AB products in the US domain are in at least one box. At this point, we want to highlight that figures relating to the number of products in the box are likely understated because we only consider boxes

[4]This contrasts with Waldfogel (2024), where a similar estimation suggests that AB products have significantly lower sales ranks on average. We believe this discrepancy stems from differences in both sample and specification. Our data is limited to first-page results, while Waldfogel (2024) includes the first three pages. Additionally, our specification includes a dummy for whether a product is Amazon's Choice, a variable not used in Waldfogel (2024). Since AB products are more likely to be Amazon's Choice, this may partly explain the difference.

Dependent Variable:	log(rank)
Variables	
Amazon Basics	-0.09
	(0.09)
Amazon's Choice	-0.87***
	(0.04)
1P	0.17***
	(0.04)
log(price)	0.16***
	(0.03)
rating	0.20
	(0.16)
log(review count)	-0.06
	(0.14)
rating × log(review count)	-0.07**
	(0.03)
Fixed-effects	
category-domain	Yes
Fit statistics	
Observations	57,620
R^2	0.697
Within R^2	0.221

Clustered (category-domain) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 2: Determines of the sales rank.

on the product pages of items within our sample. Among AB products, 82% and 77% are in the box in Canada and the US, respectively.

At the same time, none of the AB product pages in the Canadian domain feature a "Similar items to consider" box (referred to as "having the box" for brevity) as is evident in the second-to-last column in the table. By contrast, in the US, 29% of AB product pages have the box. For non-AB products, these numbers are 33% in Canada and 54% in the US.

To test whether these differences between domains and product types are statistically significant, we estimate the effect of being an AB product on the likelihood of being in the box (having the box) in Canada, conditional on being in the box (having in the box) in the US. Column (1) in Table 4 shows that AB products have an 80% chance of featuring in the box in Canada, regardless of whether they feature in the box in the US, while non-AB products have zero chance. Importantly, the AB product is singificantly

	Number of products	Share of products			
		In Box		**Has Box**	
		CA	**US**	**CA**	**US**
non-AB	13,978	0.0%	5.6%	32.6%	54.2%
AB	632	82.4%	76.6%	0.0%	29.3%
Total	14,610	3.6%	8.9%	31.2%	53.2%

Table 3: Share of products in box / having box by domain and product type (AB vs non-AB)

more likely to feature in the box even after accounting for whether a product has the Amazon's Choice label, whether it is sold by Amazon, the logged price, the average rating, the logged review count, and their interaction (all for the Canadian domain) as well as category fixed effects. Column (3) reports that, all else equal, non-AB product pages have a 48% chance of having the box in Canada if they have the box in the US, and a 15% chance if they do not. In contrast, AB products have zero probability having the box in Canada, irrespective of whether they have the box in the US or not. The finding that AB products are significantly less likely to have the box in Canada is robust to the inclusion of control variables as well.

Overall, Table 4 suggests that Amazon treats the *same* product that is assigned the box in the US, whether by being shown in the box or by having the box displayed on its page, differently in Canada, even after accounting for observable attributes such as the price.

Importantly, we do not claim that the "Similar items to consider" box is the only or most significant form of self-preferencing in either domain.[5] Nor do we claim that self-preferencing does not occur in the US domain, as we cannot determine exactly how and why Amazon assigns boxes to products and vice versa. To better understand Amazon's algorithm, we regress the likelihood of a product being in the box or having the box on various observables using only US observations. Our results show that, all else equal, AB products are more likely to be in the box but less likely to have their page feature the box (see Table A.2 in the Supplemental Appendix) though we cannot say if this is self-preferencing. Although the inner workings of the algorithm are opaque to us, we argue that the differing treatment of AB and non-AB products across domains in relation to the box constitutes self-preferencing.

We expect self-preferencing to impact sales as follows. By not displaying

[5]For the sake of brevity, we will refer to this specific instance of self-preferencing simply as self-preferencing in the following.

Dependent Variables:	In Box (CA)		Has Box (CA)	
Model:	(1)	(2)	(3)	(4)
Variables				
Amazon Basics	0.80***	0.82***	-0.15***	-0.31***
	(0.03)	(0.02)	(0.01)	(0.02)
Has Box (US)			0.33***	0.28***
			(0.02)	(0.02)
Has Box (US) × Amazon Basics			-0.33***	-0.27***
			(0.02)	(0.03)
In Box (US)	0.00	0.0008		
	(0.00)	(0.0006)		
In Box (US) × Amazon Basics	0.04	0.02		
	(0.03)	(0.03)		
Amazon's Choice		0.008***		-0.04***
		(0.002)		(0.01)
1P		0.00		0.02**
		(0.0005)		(0.01)
log(price)		-0.002***		0.007
		(0.0007)		(0.01)
log(review count)		-0.005		-0.04
		(0.005)		(0.03)
rating		-0.002		-0.03
		(0.004)		(0.03)
rating × log(review count)		0.0009		0.01
		(0.001)		(0.007)
Constant	0.00***		0.15***	
	(0.00)		(0.01)	
Fixed-effects				
category	No	Yes	No	Yes
Fit statistics				
Observations	28,810	28,810	28,810	28,810
R^2	0.818	0.828	0.141	0.263
Within R^2		0.819		0.135

Clustered (category) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 4: Regression of being in / having the box in CA on being in / having the box in US

boxes on AB product pages, Amazon shields its own products from competition, making it more difficult for consumers to compare them with viable alternatives, thereby increasing AB product sales. Conversely, by not showing non-AB products in the box, consumers are unable to compare these products with the one they are currently viewing, reducing the likelihood of purchase and lowering average non-AB sales. However, the sales effect is not uniform, as products are differently exposed to self-preferencing, an aspect we exploit to estimate the sales effect.

4 Estimation

In this section, we first discuss our estimation strategy and identifying assumption (Section 4.1) and then present our estimation results (Section 4.2).

4.1 Estimation Approach

For identification, we exploit the fact that, among AB products and among non-AB products, some products are directly exposed to different degrees of self-preferencing while others are not. Recall from Section 3 that non-AB products are sometimes shown in the box in the US but never in Canada. If a product has the box in the US, it is directly exposed to self-preferencing in Canada because it "loses" the box, and we expect it to have relatively lower sales in Canada. Conversely, if a product is not in the box in the US, its sales should not be directly affected by self-preferencing. Similarly, recall that AB product pages do not have the box in Canada, whereas in the US they might. If an AB product already does not have the box in the US, the product's sales difference between domains is not directly affected by self-preferencing in Canada. However, if the product has the box in the US, we expect relatively higher sales in Canada because there AB products do not have the box.

We take advantage of the fact that each product appears in both domains, allowing us to compare the *difference* in sales across domains. While a product being in the box or having a box may be influenced by factors like the product's quality, our identifying assumption is that relative sales differences between Canada and the US are independent of these factors, conditional on whether a product is in or has a box in the US. This lets us attribute any additional sales differences to *direct* exposure to self-preferencing in Canada (through Amazon's manipulation of the product's *own* box assignment), rather than to *indirect* exposure (through manipulation of the box assignments of *other* products). One concern is that some categories

generally display more boxes than others, and products in these categories may sell fewer units in Canada than in the US. For example, products in the category "Beach Towels" may face lower demand in Canada than in the US. To account for this, we add category-domain fixed effect.[6]

Equation (1) and Equation (2) outline our specifications.

$$\begin{aligned}\ln(y_{im}) &= \alpha + \gamma \text{In Box (US)}_i + \delta \text{CA}_m \times \text{In Box (US)}_i \\ &\quad + \theta X_{im} + \eta_{cm} + u_{im} \qquad (1) \\ \ln(y_{im}) &= \alpha + \gamma \text{Has Box (US)}_i + \delta \text{CA}_m \times \text{Has Box (US)}_i \\ &\quad + \theta X_{im} + \eta_{cm} + u_{im} \qquad (2)\end{aligned}$$

y_{im} is the dependent variable, either the sales rank or sales, for product i in domain m.[7] CA_m is an indicator for whether the observation is from the Canadian domain. In Box (US)_i and Has Box (US)_i indicate whether the ASIN features in the box and if its product page includes the box in the US domain, respectively. We also add additional control variables (X_{im}) as well as category-domain fixed effects (η_{cm}). Importantly, we estimate Equation (1) using non-AB products only, whereas we estimate Equation (2) using AB products only. Hence, the δ coefficient in Equation (1) captures how the difference in the dependent variable between domains varies for non-AB products that feature in the box in the US but not in Canada, compared to non-AB products that do not feature in the box in either domain. Conversely, δ in Equation (2) measures the y difference for AB products whose page in the US has the box but the one in Canada does not, compared to AB products whose page never features the box.

Note that, while the above specification is reminiscent of a difference-in-differences estimation, our identifying assumption does not require that products for which the box assignment is unchanged between the US and Canada are unaffected by the change in the degree of self-preferencing. We interpret δ as a conservative estimate of the effect of being *directly* exposed to self-preferencing, by no longer being in the box (not having the box), on the *sales (dis-)advantage* of non-AB (AB) products compared to non-AB (AB) products that are only *indirectly* exposed. This is because all products in

[6]Note that our results are qualitatively robust to the inclusion of category-domain fixed effects.

[7]Following Chevalier and Goolsbee (2003), we use the observed sales together with the sales rank to predict sales for the entire sample. This approach has been applied to books (Chevalier and Mayzlin, 2006; Reimers and Waldfogel, 2021) as well as a broad range of products (Gutierrez, 2021; He, Hollenbeck and Proserpio, 2022) sold on Amazon. See A.1 in the Supplemental Appendix for details.

Canada face intensified preferential treatment of AB products compared to the US.

A concern could be that non-AB products not featured in the box in the US may benefit if their competitors are no longer shown in the box in Canada. We explore this possibility in more detail in Table A.3 in the Supplemental Appendix. Specifically, we identify products that are closer substitutes within the same category to those featured in a box in the US.[8] We find that substitute products experience a similar difference in sales rank compared to less close substitutes, suggesting that the measured sales difference is primarily driven by the sales difference of directly exposed non-AB products.

4.2 Estimation Results

Table 5 present our estimation results. Column (1) in Table 5 shows that the sales rank for non-AB products that are not in a box in the US is similar in the US and Canada, while Column (5) in the same table finds sales are lower by 60%, presumably because the Canadian market is relatively smaller. Non-AB products recommended through the box in the US significantly outperform those that are not in the box, but only in the US. According to our preferred specification, Column (2) in Table 5, which includes controls and category-domain fixed effects, the former group of products has a 36% lower sales rank (recall that a lower rank indicates higher sales) in the US, whereas in Canada, their sales rank is, on average, only 5% lower. According to Column (5), in the US domain, non-AB products in the box have approximately 25% higher sales compared to non-AB products that are not in the box. However, this gap closes in Canada, where the former group's sales are, on average, 3% lower. Overall, self-preferencing substantially reduces the sales rank of non-AB products by 48% while decreasing their sales by 22%.

For AB products, by contrast, we do not find any significant difference in sales ranks (Columns (3) and (4)) or sales (Columns (7) and (8)), once we include fixed effects and controls, between products whose page contains the box in the US and those whose page does not. One possible explanation is limited statistical power given the relatively small number of AB products.

For reference to the results on AB products, we also run a similar analysis for non-AB products. Non-AB products whose US product page features a box with another non-AB product in it are expected to perform relatively better in Canada, where such a box does not appear. To test this, we re-estimate

[8]We identify substitute products as follows: For each product in a US box, we collect the products listed in the "Compare with similar items" section, a set Amazon deems to be similar.

Dependent Variable:	log(rank)				log(sales)			
	non-AB only		*AB only*		*non-AB only*		*AB only*	
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Variables								
In Box (US)	-1.3***	-0.44***			0.58***	0.22***		
	(0.24)	(0.07)			(0.06)	(0.05)		
Has Box (US)			0.09	-0.28*			0.13	0.05
			(0.45)	(0.16)			(0.12)	(0.04)
CA	-0.19		0.19		-0.92***		-1.5***	
	(0.68)		(0.74)		(0.04)		(0.10)	
In Box (US) × CA	0.69**	0.39***			-0.27***	-0.25***		
	(0.35)	(0.12)			(0.08)	(0.06)		
Has Box (US) × CA			-0.07	0.03			-0.15	-0.03
			(0.58)	(0.24)			(0.15)	(0.07)
Amazon's Choice		-0.85***		-0.57***		0.26***		0.12***
		(0.04)		(0.11)		(0.02)		(0.04)
1P		0.16***				-0.04***		
		(0.04)				(0.01)		
log(price)		0.16***		0.33***		-0.02***		-0.10***
		(0.03)		(0.07)		(0.006)		(0.02)
rating		0.09		6.0***		-0.15*		-0.64*
		(0.14)		(1.7)		(0.07)		(0.35)
log(review count)		-0.17		3.8***		0.09		-0.19
		(0.12)		(1.0)		(0.06)		(0.19)
rating × log(review count)		-0.04		-0.92***		0.05***		0.11**
		(0.03)		(0.23)		(0.01)		(0.04)
Constant	6.6***		5.2***		5.5***		6.5***	
	(0.49)		(0.55)		(0.03)		(0.08)	
Fixed-effects								
category-domain	No	Yes	No	Yes	No	Yes	No	Yes
Fit statistics								
Observations	54,644	54,644	2,976	2,976	44,982	44,982	2,400	2,400
R^2	0.007	0.701	0.0008	0.668	0.157	0.587	0.404	0.862
Within R^2		0.217		0.184		0.420		0.577

Clustered (category-domain) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 5: Effect of self-preferencing on the sales rank

Equation (2), restricting the sample to non-AB products and distinguishing between those whose US page includes a box with a non-AB product and those without such a box. Consistent with this prediction, we find that the former group has, on average, a 42% lower sales rank and 19% higher sales in Canada compared to the latter group, all else equal (see Table A.4 in the Supplemental Appendix).

5 Discussion

In this section, we briefly discuss potential motivations for self-preferencing in Canada and the implications of our findings. While there is no official explanation for cross-jurisdictional differences in self-preferencing, it is noteworthy that AB products experience lower sales relative to non-AB products in Canada. Columns (5) and (7) in Table 5 show that, holding the box assignment constant, AB products exhibit a larger sales decline in Canada than non-AB products. We estimate that non-AB product sales are 60% lower in Canada compared to the US, while this number is 78% for AB products. This may give Amazon a stronger incentive to engage in self-preferencing in the Canadian domain.

What is more, Amazon may also have perceived more regulatory pressure to scale down or discontinue the practice in the US. In 2019, NBC News reported that Amazon reduced private label promotions, including "Similar items to consider" links, many of which were removed after some policymakers in the US advocated for breaking up large technology companies.[9] Competition enforcement in the US has been more active in recent years, particularly concerning digital markets, which may help explain why self-preferencing practices are more limited in the US domain. Compared to Canada, US competition law is also sometimes viewed as having stronger enforcement mechanisms.[10]

A simple back-of-the-envelope calculation suggests that the sales loss for third-party sellers in the US, if the "Similar item to consider" box were reserved for AB products only, as is the case in Canada, would be substantial. In our data, non-AB products in the box account for roughly 20% of US sales. Assuming this 20% share is representative of the full product population, and using our estimates from Table 5, removing featured positions

[9]See https://www.nbcnews.com/tech/tech-news/amazon-quietly-removes-promotions-its-own-products-calls-tech-regulation-n990666.

[10]For example, unlike in the US, firms in Canada cannot be compelled to provide information for market studies (see https://www.cbc.ca/news/business/competition-laws-review-1.6855634).

for non-AB products would reduce Amazon sales of all non-AB products by approximately 4.4%. Based on 2023 Amazon sales data and the facts that third-party sellers were responsible for about 62% of all units sold on Amazon in the fourth quarter of 2024,[11] a 4.4% drop corresponds to more than 120 million items.[12] Note that this decrease is a conservative estimate for the reduction in non-AB product sales, as it excludes the indirect effect on non-AB products of AB products benefiting from more prominent placement, but it would have to be weight against the increase in sales by AB products to get to the total effect of self-preferencing.

6 Conclusion

In this article, we mainly make two contributions. First, we document that Amazon treats Amazon Basics (AB) products differently from products that are not sold under this brand across two geographic domains. In the Canadian domain, AB products do not have a "Similar item to consider" box displayed on their product page, while non-AB products are systematically excluded from being shown in this box. Amazon's asymmetric treatment of its own and other products constitutes a form of self-preferencing.

Second, we estimated the extent to which this self-preferencing steers consumers' purchase decisions and distorts product sales. By comparing the sales difference between the Canadian domain (where there is self-preferencing with respect to the box assignment) and the US domain (where there is no such self-preferencing) for products exposed to self-preferencing versus those that are not, we show that self-preferencing significantly reduces non-AB product sales, both statistically and economically.

[11] See https://www.marketplacepulse.com/stats/amazon-percent-of-units-by-third-party-sellers

[12] Reportedly, Amazon sold more than 4.5 billion items in 2023 in the US alone, see https://sell.amazon.com/blog/amazon-stats.

References

Cabral, Luis, Justus Haucap, Geoffrey Parker, Georgios Petropoulos, Tommaso M Valletti, and Marshall W Van Alstyne, "The EU digital markets act: a report from a panel of economic experts," *Cabral, L., Haucap, J., Parker, G., Petropoulos, G., Valletti, T., and Van Alstyne, M., The EU Digital Markets Act, Publications Office of the European Union, Luxembourg*, 2021.

Chen, Nan and Hsin-Tien Tsai, "Steering via algorithmic recommendations," *The RAND Journal of Economics*, 2023.

Chevalier, Judith A and Dina Mayzlin, "The effect of word of mouth on sales: Online book reviews," *Journal of marketing research*, 2006, *43* (3), 345–354.

Chevalier, Judith and Austan Goolsbee, "Measuring prices and price competition online: Amazon. com and BarnesandNoble. com," *Quantitative marketing and Economics*, 2003, *1*, 203–222.

Corniere, Alexandre De and Greg Taylor, "A model of biased intermediation," *The RAND Journal of Economics*, 2019, *50* (4), 854–882.

Dendorfer, Florian, "First-party selling and self-preferencing," *International Journal of Industrial Organization*, 2024, *97*, 103098.

Etro, Federico, "e-Commerce platforms and self-preferencing," *Journal of Economic Surveys*, 2024, *38* (4), 1516–1543.

Farronato, Chiara, Andrey Fradkin, and Alexander MacKay, "Self-preferencing at Amazon: evidence from search rankings," in "AEA Papers and Proceedings," Vol. 113 American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203 2023, pp. 239–243.

Gutierrez, German, "The welfare consequences of regulating Amazon," *Job Market Paper, New York University*, 2021.

Hagiu, Andrei, Tat-How Teh, and Julian Wright, "Should platforms be allowed to sell on their own marketplaces?," *The RAND Journal of Economics*, 2022, *53* (2), 297–327.

He, Sherry, Brett Hollenbeck, and Davide Proserpio, "The market for fake reviews," *Marketing Science*, 2022, *41* (5), 896–921.

Lee, Kwok Hao and Leon Musolff, "Entry into two-sided markets shaped by platform-guided search," *Job Market Paper, Princeton University*, 2021.

Padilla, Jorge, Joe Perkins, and Salvatore Piccolo, "Self-Preferencing in Markets with Vertically Integrated Gatekeeper Platforms," *The Journal of Industrial Economics*, 2022, *70* (2), 371–395.

Peitz, Martin, "How to apply the self-preferencing prohibition in the DMA," *Journal of European Competition Law & Practice*, 2023, *14* (5), 310–315.

Raval, Devesh, "Steering in one click: Platform self-preferencing in the amazon buy box," *Unpublished manuscript*, 2022.

Reimers, Imke and Joel Waldfogel, "Digitization and pre-purchase information: the causal and welfare impacts of reviews and crowd ratings," *American Economic Review*, 2021, *111* (6), 1944–1971.

Tirole, Jean and Michele Bisceglia, "Fair gatekeeping in digital ecosystems," 2023.

Waldfogel, Joel, "Amazon Self-preferencing in the Shadow of the Digital Markets Act," Technical Report, National Bureau of Economic Research 2024.

Zennyo, Yusuke, "Platform encroachment and own-content bias," *The Journal of Industrial Economics*, 2022, *70* (3), 684–710.