A Framework for Measuring Platform Self-Preferencing: Evidence from Kindle Daily Deals

Imke Reimers
Northeastern University

Joel Waldfogel
University of Minnesota, NBER, ZEW

April 20, 2023

PRELIMINARY DRAFT - DO NOT CITE OR CIRCULATE

Abstract

Platforms selling their own products have faced scrutiny for possible self-preferencing, a practice that will soon be illegal in some jurisdictions. Yet, it is not really clear what self-preferencing is, or how it might be measured. We propose a theoretical definition of self-preferencing and a test for self-preferencing as well as a framework for assessing its welfare effects. We illustrate our approach by examining possible self-preferencing in Amazon's daily rank-ordered ebook promotion Kindle Daily Deal.

Introduction

Platforms have come to play significant roles in many retail sectors, and platforms with interests in the products they sell are facing scrutiny for giving their own products preferential treatment relative to those of other suppliers. For example, the EU's Digital Markets Act forbids gatekeepers from giving preferential treatment to their own products. Under the proposed American Innovation and Choice Online Act, it would be unlawful for a platform to "preference the products, services, or lines of business of the covered platform operator over those of another business user on the covered platform in a manner that would materially harm competition."[1] Even before the imposition of the DMA as an ex ante remedy, the European Union imposed a $2 billion fine on Google as an ex post penalty for giving its own products preferential ranking.

The paradigmatic example of the behavior that regulators seek to control is a platform's display ranking of products. Self-preferencing arises when a platform selling its own products alongside those of competitors gives its own products better rankings than those products' attributes warrant. While prohibiting self-preferencing sounds simple and appealing, is definition and measurement are not straightforward, even as it has now become urgent. A determination that a product has received a better ranking than it warrants begs the question of what ranking a product warrants and requires a coherent method for defining appropriate ranking.

One natural approach is to "round up the usual suspects," i.e. to regress platform rankings on plausibly relevant product characteristics, asking whether the platform's own products receive better rankings. The inherent challenge dogging this *conditioning on observables* (COO) approach, however, is the determination of what factors are relevant, as well as the question of whether all relevant factors have been included so that the platform product

[1] See `https://www.congress.gov/bill/117th-congress/senate-bill/2992/text`.

coefficient can be interpreted as bias, rather than correlation with some unobservable but legitimate determinant of rankings.

An alternative approach, pursued in this paper, is to articulate a model of the platform's objective in ranking the products. Then, if one can observe both the platform's ranking as well as the outcome the platform seeks to control (e.g. sales or revenue), then one can both define and measure bias. Theoretically, self-preferencing is the ranking of a platform's products better than, say, sales maximization would require. It is then possible to construct a test for bias that avoids the inherent challenges of the COO approach. All that's needed to demonstrate self-preferencing is to show that, conditional on platform ranks, the ranking's impact on the platform's own products is lower than its impact on other products. We term this the "outcome-based self-preferencing test."

If one can isolate the causal effect of platform rankings on sales, then it's possible to measure the degree of self-preferencing in rank terms. This gives rise to an outcome-based measure of rank bias. Finally, this approach can easily be embedded in an explicit model of platform rankings for welfare analysis. This paper presents such an approach, along with an illustration using Amazon's Kindle Daily Deal ebook promotions.

In our application, we first look for evidence of self-preferencing on Amazon's Kindle Daily deals, asking whether Amazon books receive better promotional ranks than their sales prospects warrant. We find that conditional on books' promotional ranks, Amazon books sell fewer copies than non-Amazon books during the promotion. The sales differential is equivalent to Amazon giving its own books promotional ranks that are an average of about ten times better than their sales prospects warrant. We then turn to measuring the welfare cost of Amazon's self-preferencing. We compare status quo rankings, which incorporate the self-preferencing we document, with debiased rankings. We find that the self-preferencing in current practice forgoes 12.9 percent of possible consumer surplus and 5.4 percent of possible revenue in debiased rankings (and more than half of the welfare at stake in the ranking choice among currently included products).

The paper proceeds in five sections after the introduction. Section 1 provides background on regulatory developments necessitating ways of quantifying self-preferencing. Section 2 presents a simple model of platform rankings choices from which we derive both a coherent definition of, and a simple test for, platform self-preferencing in the determination of rankings. We also outline how to implement the model empirically in order to estimate the welfare impact of potentially biased rankings. Section 3 describes the data used in the application. Section 4 presents empirical evidence on the effects of Amazon's promotion on sales, as well as estimates of the degree of Amazon's self-preferencing in its promotional rankings. Section 5 uses the simple model to develop estimates of the effects of Amazon's ranking decisions on publisher revenue and consumer surplus, compared to counterfactual scenarios without self-preferencing and without any Amazon books in the promotions.

1 Background

1.1 Policy context

Antitrust authorities around the world are now implementing or contemplating restrictions on retail platforms that would prevent them from giving preference to their own products. For example, under the EU's Digital Markets Act, implemented in 2022, "the gatekeeper should not engage in any form of differentiated or preferential treatment in ranking on the core platform service... ...in favour of products or services it offers itself." Moreover, the determinants of its rankings should be "generally fair and transparent."[2] Under the proposed US American Innovation and Choice Online Act, it would be unlawful for a platform to "preference the products, services, or lines of business of the covered platform operator over those of another business user on the covered platform in a manner that would materially

[2] See https://www.consilium.europa.eu/media/56086/st08722-xx22.pdf.

harm competition."[3] Competition authorities in other countries are also concerned about self-preferencing among online platforms.[4]

The canonical problem that these policies seek to address is a platform ranking decision. This arises when a platform chooses a display ranking, or an algorithm for ranking search results. The self-preferencing question is the question of whether the platform's own products (or some other group of products the platform is suspected of favoring) obtain better ranking or page position than is appropriate for those products. While policy now seeks to proscribe such behavior, a challenge is that it is not clear what constitutes self-preferencing. Given that it is, or will be, the law in many places, there is a pressing need for both a definition and a way to measure self preferencing.

1.2 Relevant literature

There is a substantial theoretical literature on platform bias and a growing empirical literature.[5] Edelman (2011) provides evidence of bias in search results, while Chen and Tsai (2019) document self-preferencing at Amazon through the choice of items that the platform refers to as "frequently bought together." Aguiar et al. (2021) test for bias in Spotify playlists. See also Zhu and Liu (2018); Hunold et al. (2018, 2020); Cure et al. (2022); Gutierrez (2021); He et al. (2021). Concern about self-preferencing is not universal: Dubé (2022) argues that Amazon's promotion of its own products is moderate compared to the way that offline retailers treat their store brands.

[3]See `https://www.congress.gov/bill/117th-congress/senate-bill/2992/text`.

[4]See `https://www.outlookindia.com/business/explained-why-is-competition-commission-of-india-probing-amazon-news-194362`.

[5]Theoretical contributions include Armstrong and Zhou (2011), Hagiu and Jullien (2014), Parker et al. (2020), De Corniere and Taylor (2019), Bourreau and Gaudin (2022).

2 Theory: rankings and self-preferencing

In order to measure self-preferencing, we first need to define it. This section presents a theoretical framework in which a platform chooses product rankings to maximize some objective (e.g. sales quantities or revenues), along with possible self-preferencing. We put the model to two uses: We develop a test for self-preferencing, and we outline how to quantify the degree of self-prefencing, and measure the welfare consequences of self-preferencing.

Specifically, we model the platform's current ranking decisions, which incorporate possible self-preferencing, as well as a way to model a counterfactual environment with self-preferencing removed. One benefit of this approach is that we don't need to know all factors determining the platform's status quo rankings. Rather, we understand that the platform may observe things that we do not observe, and we assume that the platform makes privately optimal decisions that might include self-preferencing.

Measuring welfare effects of self-preferencing requires three steps. We first need a model that can deliver the status quo as the solution to the platform's optimization problem. Second, we need a way to ascertain the degree of self-preferencing in current rankings. Third, we need a way to remove the self-preferencing while leaving the platform's pursuit of other objectives intact.

2.1 The model and self-preferencing

2.1.1 Consumer demand

Consumers choose how much of each product j to purchase based on the products' attributes, as well as the promotional ranks to which the products are assigned (r). Rankings may convey information about quality to consumers, and it is well-documented that rankings affects sales (Ursu, 2018). We think of the realized quantity that consumers purchase as being separable into two multiplicative components. These are the rank to which the product

j is assigned (r) and the underlying saleability of the product (perhaps captured by prior sales), which we term q_j^0. Hence, realized sales are given by $q_j(r) = q_j^0 f(r)$, with $f^{'} < 0$; and $rev_j(r) = p_j q_j^0 f(r)$.

2.1.2 Platform ranking choices

For illustration, suppose the platform's objective function includes revenue. If better rankings have larger multiplicative effects on revenue, maximization of total revenue would be achieved by ranking the products in descending order of rank-independent revenue $rev_j^0 = p_j q_j^0$.[6] In general, we observe the platform's estimate of rank-independent sales with error. That is, there is a distinction between the rank-independent revenues the platform expects rev_j^0 and the rank-independent revenue observed by the researcher, $r\hat{e}v_j^0$. For clarity, however, we begin by assuming that we can observe rank-independent revenue in the same way as the platform. In that case, $r\hat{e}v_j^0 = rev_j^0$.

If the platform were not self-preferencing, then the platform would take a constant share of revenue (c_0) as its commission, and the platform's payoff would be proportional to revenue. Then, the platform would rank products in descending order of the observed rank-independent revenue. That is, $r\hat{e}v_j^0$ would decline monotonically in ranks.

Bias arises if the platform derives a differential payoff from selling its own products, relative to non-platform products. With a constant differential payoff for selling its own products, the platform's rank-independent payoff for product j would be

$$\phi_j \equiv c_0 rev_j^0 + c_1 \delta_j, \tag{1}$$

where c_1 is the additional platform payoff for selling its own product, and δ_j is an indicator for

[6] Our chosen functional form, which we introduce later, guarantees that placing products with greater saleability at better ranks raises revenue. We provide evidence for this assumption in Section 4.3 below.

platform products.[7] If the platform ranked products in descending order of ϕ_j, if $c_1 > 0$, then for the relationship between rev_j^0 and rank, the rev_j^0 values would lie along one monotonic function for non-platform products and on another function for platform products, which lies below by c_1 units. This is depicted in the left panel of Figure 1, which shows the relationship between rank-independent revenue and the rank assigned by a platform that is maximizing its payoff and is engaged in self-preferencing. This setup suggests a straightforward notion of self-preferencing: Self-preferencing is present when $c_1 > 0$ or, equivalently, when the platform's products receive better ranks, conditional on rank-independent revenue.[8]

We now relax the assumption that we observe the platform's rank-independent revenue (rev_j^0) in the same way as the platform. Instead, we observe it with error. Our estimate of the platform's payoff from product j is $\hat{\phi}_j \equiv c_0 r\hat{e}v_j^0 + c_1\delta_j$, while the payoff observed by the platform is $\phi_j = \hat{\phi}_j + \epsilon_j$, where ϵ_j is the part of the platform's payoff from product j that we do not observe.

We illustrate this in the right panel of Figure 1. The dots correspond to $r\hat{e}v_j^0$, while the lines represent $rev^0(r)$ functions.[9] Given that we do not observe all that the platform observes, our observed analogues of rank-independent revenue $r\hat{e}v_j^0$ (the dots) will not decline monotonically in rank, even within type, even if the rev^0 functions (the lines) do. The platform's rank ordering reflects the condition for platform payoff maximization, that the platform ranks products in descending order by rev_j^0, or that $\hat{\phi}_j + \epsilon_j > \hat{\phi}_k + \epsilon_k$ for $j < k$.

We expect a monotonically negative relationship between $r\hat{e}v_j^0$ and ranks *on average* rather than observation-by-observation. Given a downward-sloping relationship, we can test for bias by asking whether the $r\hat{e}v_j^0(r)$ function for platform products lies below the function

[7]The term c_1 could literally be a different commission rate, or a notional payoff derived from, say, an expectation of future returns from promoting platform products.

[8]Aguiar et al. (2021) implement an analogous approach for measuring bias at Spotify's New Music playlists.

[9]Note that under the assumption that the $rev^0(r)$ functions are linear, this suggests that if we fit a line through the dots, we obtain the true $rev^0(r)$ functions.

for non-platform products.

2.1.3 Quantifying the degree of self-preferencing

Quantifying the potential bias requires a way to calculate $r\hat{e}v_j^0$ as well as a choice of functional form for the $rev^0(r)$ function. To these ends, we first derive estimates of $r\hat{e}v_j^0$. Observed revenue for product j at rank position r is given by $rev_j(r) = rev_j^0 f(r)$. Assuming that $f(r) = r^\gamma$, this implies $rev_j^0 = {}^{rev_j(r)}/_{r^\gamma}$. We discuss the implementation of this in Section 4.3. Then, using our estimates of rev_j^0, we specify a linear function relating $r\hat{e}v_j^0$ to rank and product type:

$$\ln(r\hat{e}v_j^0) = \pi_0 + \pi_1 \ln(r_j) + \pi_2\delta_j + \nu_j. \tag{2}$$

The parameterized component of this estimate, $\pi_0 + \pi_1 r_j + \pi_2\delta_j$, can satisfy conditions needed for a (downward-sloping) $rev_j^0(r)$ function that delivers status quo rankings; and in our empirical implementation we use $\hat{\pi}_0 + \hat{\pi}_1 \ln(r_j) + \hat{\pi}_2\delta_j$ as our empirical estimate of the true $\ln(rev_j^0(r))$ function.

We put this estimated function to multiple uses. First, the estimated function provides some evidence about platform motives. If the platform prefers greater seller revenue to less, then rank-independent revenue will decline in ranks ($\pi_1 < 0$), all else equal.

Second, this equation provides both a test for self-preferencing and a quantification, in rank terms, of the magnitude of the self-preferencing. Self-preferencing is present if $\pi_2 < 0$. That is, self-preferencing is present if the platform's products have lower rank-independent revenue, conditional on the rank to which they were assigned. This is equivalent to $c_1 > 0$ in Equation (1). It arises if the platform assigns its own products to higher ranks than their rank-independent revenues warrant. The term π_2/π_1 provides a measure of the self-preferencing: e^{π_2/π_1} is the multiplicative factor by which the platform preferences its own products' ranks.

We can also create a debiased index – and a ranking – that removes self-preferencing, using $\ln(rev_j^{0,db}(r)) = \ln(rev_j^0(r)) - \pi_2\delta_j$. This approach preserves the status quo rank ordering of products within type (platform-owned vs not) but changes the rank-orderings across types. Using $\ln(rev_j^0(r))$ and $\ln(rev_j^{db}(r))$ and their associated rank orderings, we can make a comparison between actual rankings, which solve the platform's status quo maximizing problem, and counterfactual rankings that remove self-preferencing while leaving other rank determinants, that we as researchers cannot observe, intact.

2.2 Welfare comparisons

The consumer surplus (CS) associated with product j ranked at r is given by the definite integral of the demand function for product j from 0 to its observed sales under the promotion q_j. Given our chosen functional form, $\ln(q_j) = X_j\beta + \gamma\ln(r_j) + \alpha p_j + \varepsilon_j$, the CS from a promotion day with a ranking of products R is:

$$CS(R) = \sum_j CS_j(r_j) = \sum_j \frac{1}{\alpha} q_j \left(\ln(q_j) - 1 - q_j e^{-\alpha p_j} - p_j\right), \tag{3}$$

where q_j is understood to be the q_j arising when j is ranked at r, i.e., $q_j = q_j^0 r^\gamma$. The revenue associated with product j when ranked at r is given by

$$rev(R) = \sum_j rev_j(r) = \sum_j p_j q_j. \tag{4}$$

To determine the welfare cost of any self preferencing in actual rankings, we need a benchmark. Our main comparison will be between actual and debiased rankings, allowing us to ask how much CS or revenue the bias in actual rankings forgoes. These alternatives take the set of ranked products as given and measure the welfare consequences of changing the location of platform-owned products in the ranking

3 Data

The data for the application consist of two things. First, for each date between April 4 and July 12, 2022, we observe the roughly 50 ebook titles promoted by Amazon in their Kindle Daily Deal, as well as the order of the promoted titles. This is a total of 76 daily promotions of 3,738 promotional listings. These promotional listings include 2,892 distinct titles because some titles are promoted on more than one day. Second, we observe daily Amazon sales and prices for each of these ebook titles both before and after promotion. Our main analyses make use of data for 20 days prior to promotion as well as the day of the promotion and the following day. For each title we also observe its genre, the publisher type, and its sales during 2021.

The Kindle Daily deal data are drawn from Amazon's page listing the ebook titles promoted that day.[10] The remaining data are from Bookstat, a vendor providing data on book titles as well as daily data on price and estimates of quantities sold at Amazon.

Table 1 summarizes the data. Of the promotional listings, 17.9 percent are either self-published through Amazon's Kindle Direct Publishing or are published by one of Amazon's imprints. The mean daily sales prior to the promotion are 8.5 for non-Amazon and 21.6 for Amazon titles. For the day of the promotion and the following day, sales average 138.9 for non-Amazon and 63.7 for Amazon books. Prior to the promotion, prices average \$8.83 for non-Amazon titles and \$4.73 for Amazon titles. On the promotion day, the price averages \$4.17 for non-Amazon books and \$2.26 for Amazon.

[10] See `https://www.amazon.com/b?node=6165851011` and `https://www.amazon.com/s?rh=n%3A6165851011&fs=true` for the current list.

4 Descriptive evidence

This section presents descriptive facts and results to motivate the empirical modeling. We first present evidence on the the impacts of the promotions on prices, quantities sold, and revenue. Second, we document that Amazon books get better ranks. Third, we offer evidence for the model assumption employed above, that higher ranks deliver larger proportional effects on revenues, and more so for books that are more saleable prior to the promotion. Finally, we present evidence of bias.

4.1 Promotion effects

We first seek to document effects of the promotions on prices and sales. For these analyses, we include 10 days before, and 10 days after, the promotion day. Some books are promoted more than once; and in some instances a book is promoted fewer than 10 days before, or fewer than 10 days after, the focal promotion. These books appear more than once in these analyses, and therefore we also include an indicator for whether the book is included in another day's promotion.

In particular, we estimate the following models:

$$y_{jt} = \psi_\tau + \mu_j + \rho\mathbb{1}_{jt} + \epsilon_{jt}, \tag{5}$$

where y is price, quantity, or revenue for title j on day t, depending on the specification, ψ_τ is a set of flexible effects showing the level of the dependent variable τ days before or after the last day untreated by the focal promotion, μ_j is a title-by-promotion fixed effect, and $\mathbb{1}_{jt}$ is an indicator that is 1 if the book appears in another day's promotion on day t.[11]

Figure 2 shows the ψ terms reflecting the evolution of the prices, separately for Amazon

[11] Day 1 of these graphs is the promotion day. We do not know what time of day the data are collected, so day 0 may be partially treated, whereas day -1 is fully untreated.

and non-Amazon books, around the day of the promotion. Prices are essentially stable, then fall by \$3 for Amazon books and by almost \$4.5 for non-Amazon books on day 1. Prices return to their pre-promotion level the day after the promotion (day 2).

Figure 3 shows quantity effects, and the quantity sold increases substantially on the day of the promotion. The daily sales jump 50 units for Amazon books and by over 150 units for non-Amazon books. Sales remain high the second day, then return to pre-promotion levels by the third day. Figure 3 does not control for price. If we control for price, the day 1 quantity effect becomes smaller but remains statistically significant. It reaches 20 for Amazon books, and about 110 for non-Amazon books.

Finally, Figure 4 shows the effect on revenue. For non-Amazon books, revenue rises substantially in day 1 and remains high – around \$400 per promoted book – on day 2. For Amazon books, the higher quantities sold on day 1 are offset by lower prices, and revenues do not rise on day 1. Revenues rise on day 2, however, reaching nearly \$100 per title.

These figures show clear effects of the promotions. Moreover, the effects are much larger for non-Amazon than for Amazon titles; and prices, quantities and revenues return to their pre-promotion levels two days after the promotion. Larger promotional effects on sales for non-Amazon books provide a hint of pro-Amazon bias, although more evidence – provided below – is needed for firmer conclusions.

4.2 Do Amazon books get better ranks?

Of the books promoted each day, about 18 percent are Amazon titles; and the average promotional rank for Amazon's titles is 17.3, while the average rank for other titles is 26.3. As Figure 5 shows, Amazon's share of the top 10 is between 30 and 40 percent, while its share of promoted titles falls as the promotional rank gets worse, to about ten percent among ranks between 30 and 50. In conjunction with larger promotional effects for non-Amazon books, these patterns provide suggestive evidence that Amazon's rankings reflect self-preferencing.

4.3 Evidence of self-preferencing

Figure 6 shows the log quantities $\ln(q_j(r))$ sold under the promotion in the left panel, and log revenues $\ln(rev_j(r))$ in the right, as functions of the ranks to which the books are assigned. Realized sales and revenues decline in rank, and the schedules for Amazon books lie below the non-Amazon schedules. These figures are suggestive of both platform motivations to generate sales and revenue, as well as self-preferencing. However, it is important to note that the relationship between sales (or revenue) and ranks reflects both the possibility that more saleable books are ranked higher as well as the causal effect of rank on realized sales under the promotion.

Figure 7 provides some direct evidence that more saleable books are ranked higher. Both pre-promotion quantities sold and revenues decline in rank: Books ranked in the top 5 have average pre-promotion sales of roughly 50 per day, while books ranked 10^{th} have average pre-promotion sales of nearly 10; and books ranked worse have lower pre-promotion sales still. The relationship for pre-promotion revenue is similar in shape. To the extent that past sales reflect books' saleability, the figure provides evidence that the platform's rankings are broadly consistent with concern about both quantity and revenue.

Recall that $rev_j = rev_j^0 r^\gamma$.[12] To measure bias in rank assignments, we need to identify the relationship between the books' inherent saleabilities q_j^0 (or rev_j^0) and the multiplicative effects of the rankings to which they are assigned, r_j^γ. This, in turn, requires us to extract a plausibly causal γ reflecting the effect of rank. The general concern is that books at different ranks have different observable and unobservable attributes related to saleability, so that an estimated relationship between, say, $\ln(q_j(r))$ and $\ln(r)$ – even with book characteristics

[12] By using this functional form, we impose that better ranks have a larger proportional impact on revenue for books that are more saleable. To test this assumption, we regress both product j's quantity during the promotion as well as the revenue that it receives, as a function of its rank r and its sales before the promotion, along with an interaction of the rank with the pre-promotion sales of the product. The interaction term is consistently significantly negative, indicating that the effect of better ranks is bigger for more saleable products.

included – might reflect correlations with book types and not the causal effect of rank.

Table 2 shows a sequence of regressions relating logs of realized sales and revenue (logs of $q_j(r)$ and $rev_j(r)$) to log ranks. The first two columns show the raw relationships. The estimates of the rank parameter γ are -0.516 (standard error = 0.0179) for sales and -0.497 (0.0195) for revenue.

Some of the books j appear in more than one promotion, at different ranks. This allows us to difference out the time-constant title-specific unobservable characteristics for j. We can identify the effect of rank using the relationship between the cross-promotion sales differential and the difference in assigned ranks. Columns (3) and (4) of Table 2 report estimates with title fixed effects, and the γ estimates are smaller in absolute value than in columns (1) and (2). They are -0.267 (0.106) for sales and -0.264 (0.111) for revenue. Columns (5) and (6) replace the fixed effects with controls such as pre-promotion sales, and the estimates are similar to those in columns (1) and (2). The smaller FE estimates support our expectation that the overall relationship between realized sales and ranks is due partly, but not completely, to the causal effect of ranks on sales.

To explore bias, we proceed with the fixed effects estimates of γ.[13] Given our assumed functional form and an estimate of γ, we construct rank-independent sales for each title j as

$$\hat{q}_j^0 = q_j(r)/r_j^\gamma, \tag{6}$$

and rank-independent revenue $r\hat{e}v_j^0 = rev_j(r)/r_j^\gamma$. Figure 8 shows how rank-independent log revenue $(p_j\hat{q}_j^0)$ varies across ranks, separately for Amazon and non-Amazon products. A few things are clear. First, the platform ranks "better" books higher (at ranks closer to unity), especially among the best rankings. Second, on average, Amazon books have lower rank-independent sales at any rank. This indicates self-preferencing in the assignment of

[13] In Section 5.3 we explore implications of different values of γ for our results.

promotional ranks and also suggests that most Amazon books should be ranked last among the promoted books. Figure 8 also indicates the magnitude of the self-preferencing in rank terms. This is the horizontal distance between the functions implied for the Amazon and non-Amazon relationships between $\ln(\hat{q}_j^0)$ and log rank $\ln(r_j)$.

One simple way to measure the rank bias is to linearize the relationships via our functional specification of $\ln(r\hat{e}v^0)$ from Equation (2). The degree of bias in rank terms is then e^{π_2/π_1}. Finding values of $\pi_2 = -0.835$ (se $= 0.046$) and $\pi_1 = -0.497$ (se $= 0.020$), we estimate this to be 12.23 (1.86).

The bias implied by this test is substantial. Amazon books' actual ranks average 17.47. Correcting the bias would result in most Amazon books being ranked at the end of the lists. The substantial degree of implied bias also raises a question of whether too many Amazon books are included in the rankings, not just whether they are ranked too high.

5 Welfare estimates

This section presents estimates of the welfare outcomes associated with actual and debiased rankings, as well as rankings that include all, and no, Amazon products.

5.1 Welfare foregone from ranking choices

We calculate status quo and debiased CS and revenue using the observed and debiased rankings, R^a and R^{db}, respectively. Table 3 reports the percentages by which actual welfare values fall short of the debiased values. For example, we calculate CS foregone as

$$\frac{CS(R^{db}) - CS(R^a)}{CS(R^{db})}. \tag{7}$$

The self-preferencing in actual rankings chosen by the platform forgo 12.89 percent of possible CS and 5.36 percent of possible revenue. This is our main result.

While the share of welfare foregone relative to debiased rankings is of interest, this kind of measure understates the impact of self-preferencing on what is at stake through the ranking of promoted titles.

The range of welfare outcomes available from changing Amazon books' positions in the ranking runs from the effects of putting all Amazon books at the top, and at the bottom, of the rankings. It is of interest to quantify the share of the change in welfare at stake from the rankings is forgone by actual rankings.

Putting Amazon books at the bottom delivers maximal CS and revenue. Putting Amazon books at the top of the ranking, by contrast, produces the minimum values of CS and revenue achievable from different Amazon book rankings of the chosen set of promoted books. Define these minimal values as $\widetilde{CS}$ and $\widetilde{Rev}$. Of the CS at stake in the platform's ranking of Amazon relative to non-Amazon books, the share that the actual ranking forgo is given by

$$\frac{CS(R^{db}) - CS(R^{a})}{CS(R^{db}) - \widetilde{CS}},$$

and we define the analogous measure for revenue similarly.

As Table 3 shows, actual rankings forgo more than half of the CS and revenue at stake from rankings: 56.87 percent of the CS at stake in rankings and 62.62 percent of the revenue at stake in rankings.

5.2 The number of Amazon books

The welfare losses calculated above take as given the number of Amazon books included in the daily promotions. Figure 8 shows that most Amazon books have rank-independent revenue below the lowest non-Amazon books' rank-independent revenue. Hence, it is possible that status quo rankings, in addition to giving Amazon books preferential ranks, also include too many Amazon titles. We explore the welfare consequence of the number of Amazon books included by creating rankings that include no, and that include only, Amazon books.

Define R^{NA} (for "no Amazon") as a ranking that includes no Amazon books. We create this by replacing the $r\hat{e}v_j^0$ of Amazon books with the average values of $r\hat{e}v_j^0$ for non-Amazon books appearing at the rank positions where Amazon books had appeared. This gives rise to welfare measures $CS(R^{NA})$ and $Rev(R^{NA})$ and welfare losses from actual rankings measured by $\frac{CS(R^{NA})-CS(R^a)}{CS(R^{NA})}$ or, for revenue, by $\frac{Rev(R^{NA})-Rev(R^a)}{Rev(R^{NA})}$.

As row 2 of Table 3 shows, the actual ranking forgoes 5.81 percent of CS and 20.12 percent of revenue relative to rankings that include no Amazon books.

Keeping Amazon's rank ordering constant, we can similarly evaluate the welfare gains from promoting non-Amazon books. To this end, we calculate the share of welfare at stake in the number of Amazon books included in the ranking, using

$$\frac{CS^{NA} - CS(R^a)}{CS^{NA} - CS^{allA}},$$

and an analogous equation for revenue.

Because the all-Amazon ranking delivers very low welfare, the self-preferencing in current rankings forgo just slightly more of the difference between no-Amazon and all-Amazon than they forgo of the no-Amazon levels. See row 2 of Table 3.

5.3 Robustness to the ranking parameter

The ranking parameter γ affects our estimates through two conflicting channels. First, the higher is γ, the larger the effect of placing products with higher (resulting) rank-independent revenues at better ranks. Second, however, because $rev_j = rev_j^0 r_j^{\gamma}$, the higher is γ, the smaller the resulting relationship between rev^0 and rev. Hence, increases in γ can increase the calculated welfare effects, but only up to a point.

To explore the sensitivity of our results to its value, we consider a range of γ values between 0 and the OLS values in Table 3. For each possible value, we recalculate the implied rank-independent revenue rev_j^0 and the associated status quo and debiased product rankings.

Figure 9 shows the percentages of revenue and CS debiased CS and revenue forgone by actual rankings, for a range of γ values. The gray, vertical line on the left depicts the estimate of γ from the fixed effects model, and the black, vertical line shows the estimate of γ from the model with controls. As the ranking parameter increases in absolute value, the welfare losses from self-preferencing rise and reach their largest negative values between values of γ between -0.2 and -0.3 and then decline again.

6 Conclusion

Platforms have come to assume large roles in retail markets, and in many cases platforms make recommendations among both their own – and other producers' – products. This has, in turn, attracted the interest of regulators around the world.

This paper documents the existence – and the extent – of self-preferencing by Amazon in its daily ebook promotions. We find substantial self-preferencing in the ranks that Amazon chooses for the 50 books it promotes each day. Amazon books receive ranks that are substantially better than their characteristics warrant. Rather than averaging a rank of 17.5, most Amazon books should appear at the back of the rankings, if they are included at all. In addition to documenting the existence of self-preferencing, we also measure its welfare cost.

Relative to debiased rankings, the current rankings sacrifice 12.89 percent of CS and 5.36 percent of revenue. Of the welfare at stake through the ranking of included products, the current ranking foregoes over half of both CS and revenue. Finally, relative to a ranking that replaces all Amazon products with the average sorts of non-Amazon products placed at the same ranks, the current rankings sacrifice 5.89 percent of CS and 20.29 percent of revenue. These results, on self-preferencing in an environment that is easy to observe, suggest that possible self-preferencing by major platforms may warrant further scrutiny from regulators.

References

Aguiar, L., J. Waldfogel, and S. Waldfogel (2021): "Playlisting favorites: Measuring platform bias in the music industry," *International Journal of Industrial Organization*, 78, 102765.

Armstrong, M. and J. Zhou (2011): "Paying for prominence," *The Economic Journal*, 121, F368–F395.

Bourreau, M. and G. Gaudin (2022): "Streaming platform and strategic recommendation bias," *Journal of Economics & Management Strategy*, 31, 25–47.

Chen, N. and H.-T. Tsai (2019): "Steering via algorithmic recommendations," *Available at SSRN 3500407.*

Cure, M., M. Hunold, R. Kesler, U. Laitenberger, and T. Larrieu (2022): "Vertical integration of platforms and product prominence," *Quantitative Marketing and Economics*, 1–43.

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

Dubé, J.-P. (2022): "Amazon Private Brands: Self-Preferencing vs Traditional Retailing," *Available at SSRN 4205988.*

Edelman, B. (2011): "Bias in search results: Diagnosis and response," *Indian JL & Tech.*, 7, 16.

Gutierrez, G. (2021): "The welfare consequences of regulating Amazon," Tech. rep., mimeo, New York University.

Hagiu, A. and B. Jullien (2014): "Search diversion and platform competition," *International Journal of Industrial Organization*, 33, 48–60.

He, L., I. Reimers, and B. Shiller (2021): "Does Amazon exercise its market power? Evidence from Toys R Us," *Evidence from Toys R Us (August 24, 2021).*

Hunold, M., R. Kesler, and U. Laitenberger (2020): "Rankings of online travel agents, channel pricing, and consumer protection," *Marketing Science*, 39, 92–116.

Hunold, M., R. Kesler, U. Laitenberger, and F. Schlütter (2018): "Evaluation of best price clauses in online hotel bookings," *International Journal of Industrial Organization*, 61, 542–571.

Parker, G., G. Petropoulos, and M. W. Van Alstyne (2020): "Digital platforms and antitrust," *Available at SSRN 3608397.*

Ursu, R. M. (2018): "The power of rankings: Quantifying the effect of rankings on online consumer search and purchase decisions," *Marketing Science*, 37, 530–552.

Zhu, F. and Q. Liu (2018): "Competing with complementors: An empirical look at Amazon. com," *Strategic management journal*, 39, 2618–2642.

7 Figures and Tables

Figure 1: Platform rankings of products by rank-independent revenue

Panel A: no error

Panel B: error

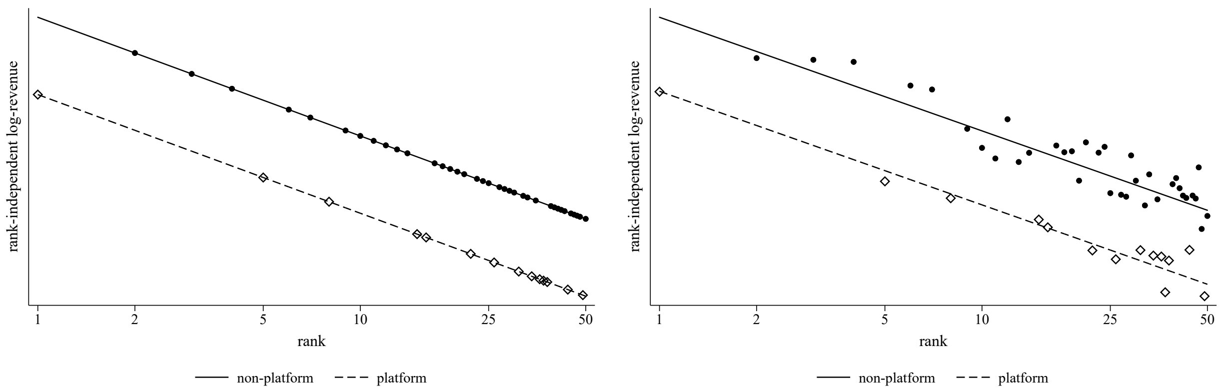

Figure 2: Price effect of promotion

Panel A: non-Amazon

Panel B: Amazon

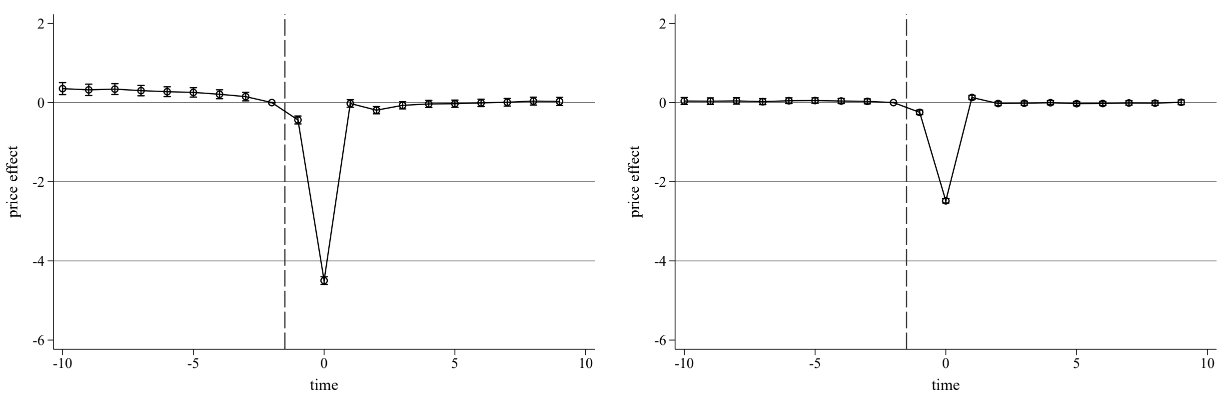

Figure 3: Quantity effect of promotion

Panel A: non-Amazon

Panel B: Amazon

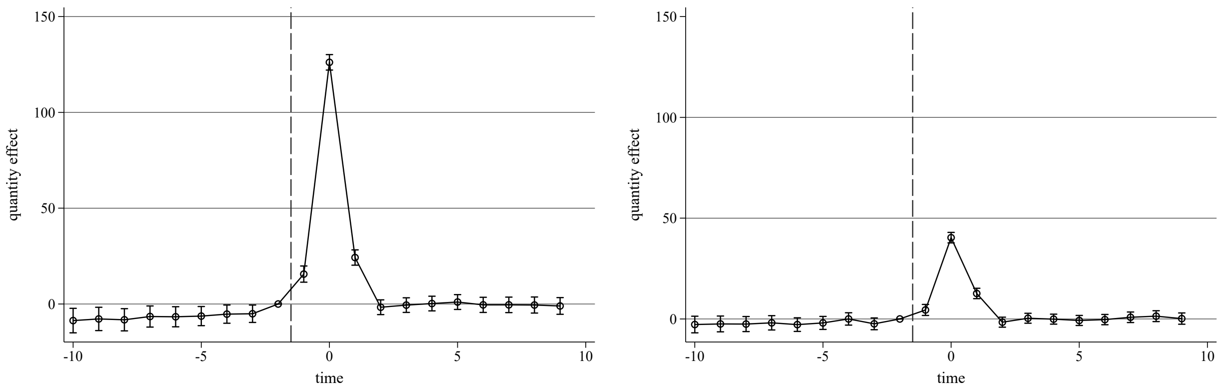

Figure 4: Revenue effect of promotion

Panel A: non-Amazon

Panel B: Amazon

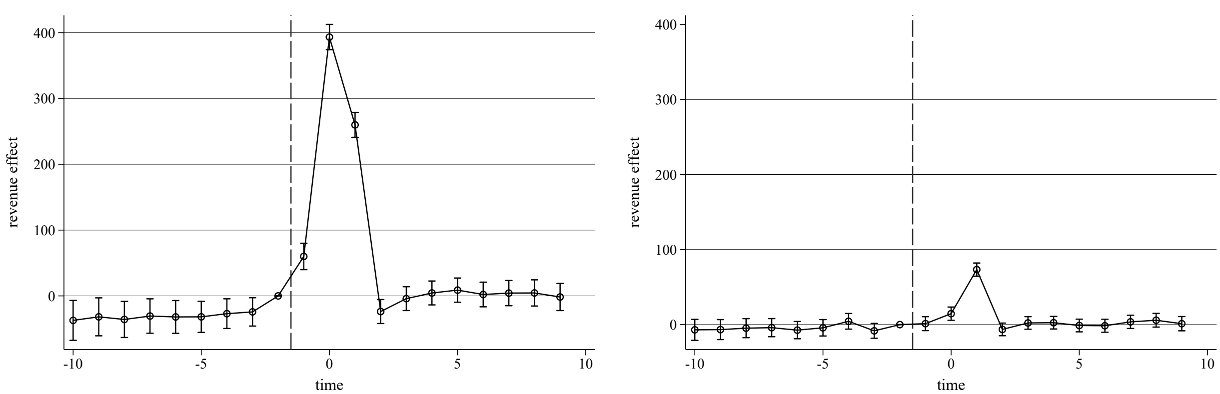

Figure 5: Share of Amazon books by promotional rank

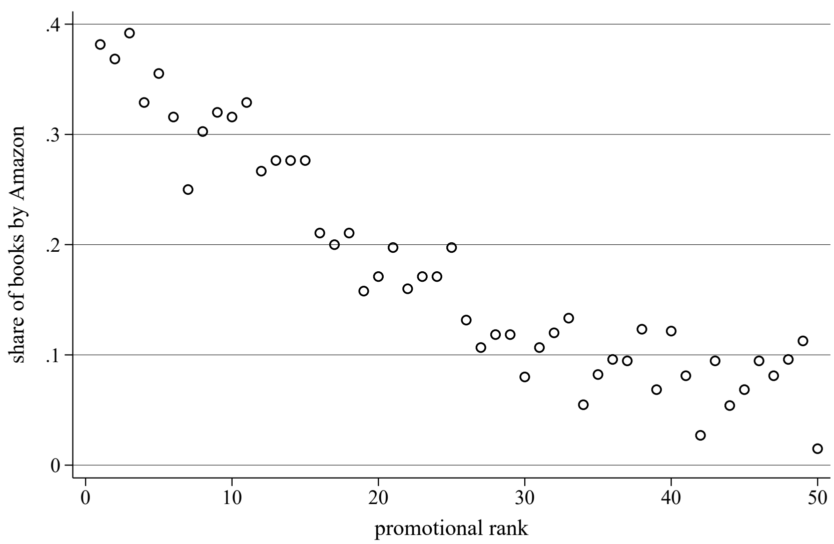

Figure 6: Post-promotion outcomes by promotional rank and type (log-rank scale)

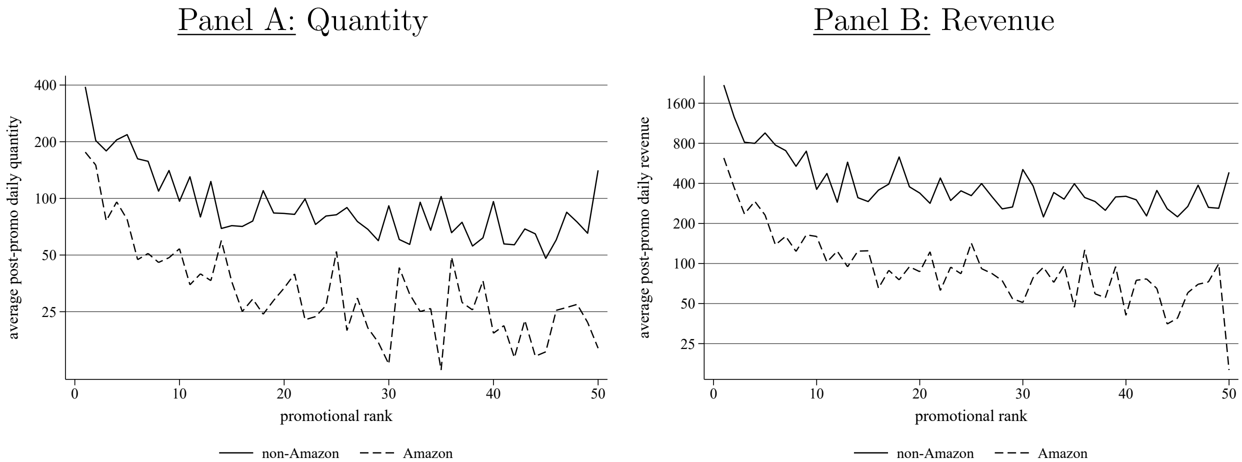

Figure 7: Pre-promotion sales by promotional rank

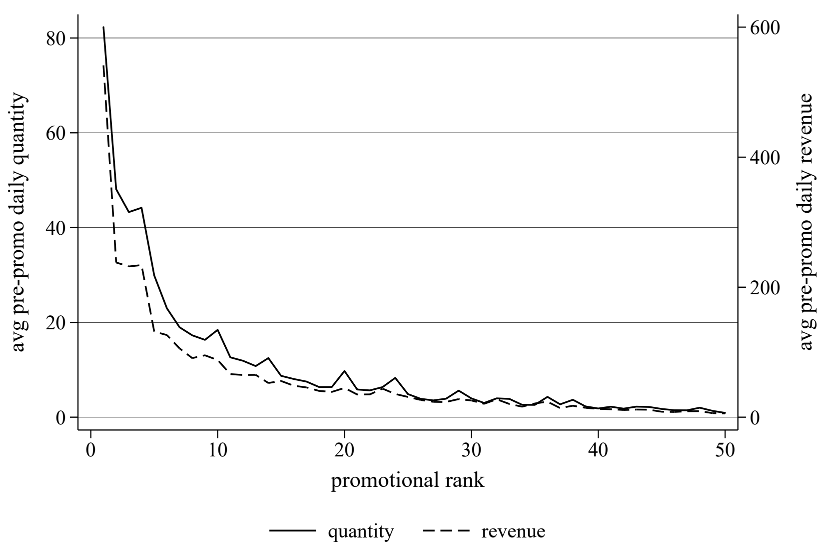

Figure 8: Rank-independent revenue for Amazon and non-Amazon

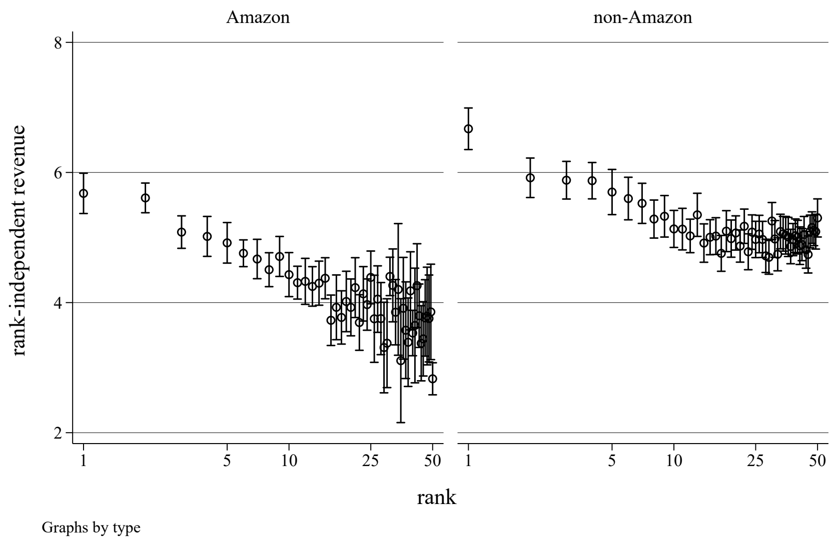

Figure 9: Welfare results and rank parameter (γ)

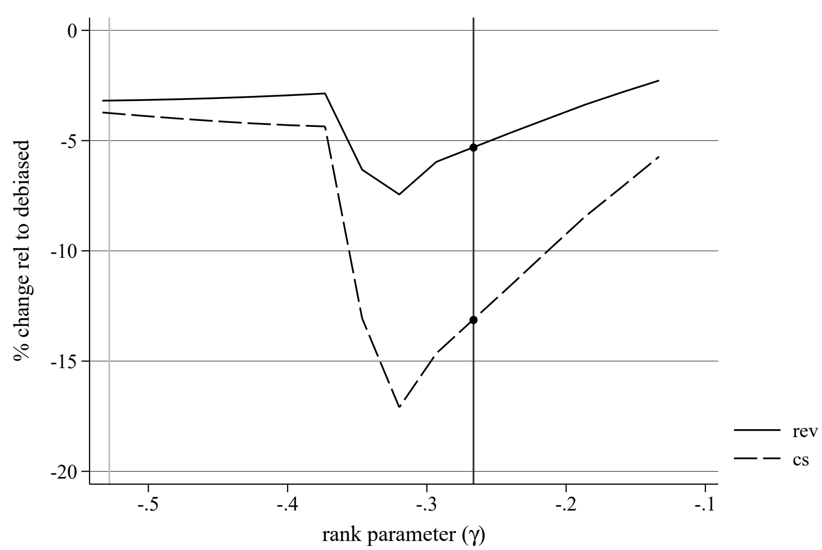

Table 1: Summary Statistics

	Non-Amazon		Amazon	
	Mean	Std. Dev.	Mean	Std. Dev.
daily avg. price up to 20 days prior	8.83	3.99	4.73	1.21
price on promotion day	4.17	3.87	2.26	1.72
daily avg. sales up to 20 days prior	8.50	22.55	21.59	50.75
sales on promotion day	138.90	266.76	63.66	109.94
Amazon promotion rank on first day	26.95	14.08	17.60	13.17
Observations	3,069		669	

Table 2: Outcome-based test for self-preferencing

	Raw		Asin FE		Controls	
	(1)	(2)	(3)	(4)	(5)	(6)
	ln(quantity)	ln(revenue)	ln(quantity)	ln(revenue)	ln(quantity)	ln(revenue)
Amazon book	-0.528***	-0.835***			-0.523***	-0.854***
	(0.0428)	(0.0456)			(0.0440)	(0.0459)
ln rank	-0.516***	-0.497***	-0.267**	-0.264**	-0.528***	-0.411***
	(0.0179)	(0.0195)	(0.106)	(0.111)	(0.0312)	(0.0335)
price	-0.0286***		-0.0856***		-0.0290***	
	(0.00433)		(0.00585)		(0.00446)	
ln pre-promo sales					-0.0133	
					(0.0289)	
day after promo	-0.505***	0.245***	-0.230***	0.303***	-0.503***	0.250***
	(0.0342)	(0.0326)	(0.0359)	(0.0295)	(0.0345)	(0.0326)
ln pre-promo rev						0.0839***
						(0.0241)
Observations	6819	6819	6824	6824	6819	6763
$\overline{R^2}$	0.278	0.222	-0.495	-0.681	0.278	0.227

Notes: .

Table 3: Welfare effects of the self-preferencing in current rankings

	overall		relative to range	
	% ΔCS	% Δrev	% ΔCS	% Δrev
actual rel to debiased	-12.89%	-5.36%	-28.70%	-29.68%
actual rel to no Amazon	-5.89%	-20.29%	-6.12%	-23.33%

Notes: The first two columns show the welfare forgone to self-prefencing in actual rankings relative to the welfare available from debiased rankings. The last two columns shows the shares of the ranges foregone. For the first first row, the relevant range runs from rankings the debias and place Amazon products at the back vs rankings that place Amazon products uniformly at the head of the rankings. For the second row, the range runs from all Amazon products to no Amazon products included in the rankings.