INSTITUTE FOR INNOVATION
AND PUBLIC PURPOSE

# Behind the Clicks: Can Amazon allocate user attention as it pleases?

**Rufus Rock**
Researcher
| UCL Institute for Innovation and Public Purpose

**Ilan Strauss**
Senior Research Associate
| UCL Institute for Innovation and Public Purpose

**Tim O'Reilly**
Founder, CEO, and Chairman of O'Reilly Media
| Visiting Professor of Practice at UCL Institute for Innovation and Public Purpose

**Mariana Mazzucato**
Director and Professor in the Economics of Innovation and Public Value
| UCL Institute for Innovation and Public Purpose

Institute for Innovation
and Public Purpose

WORKING PAPER
2023/11

**About the Institute for Innovation and Public Purpose**

The Institute for Innovation and Public Purpose (IIPP) at University College London (UCL) aims to develop a new framework for creating, nurturing and evaluating public value in order to achieve economic growth that is more innovation-led, inclusive and sustainable. This requires rethinking the underlying economics that has informed the education of global civil servants and the design of government policies. Our work feeds into innovation and industrial policy, financial reform, institutional change and sustainable development. A key pillar of IIPP's research is its understanding of markets as outcomes of the interactions between different actors. In this context, public policy should not be seen as simply fixing market failures, but also as actively shaping and co-creating markets. Re-focusing and designing public organisations around mission-led, public purpose aims will help tackle the grand challenges facing the 21st century.

IIPP is a department within UCL - and part of The Bartlett, which consistently ranks in the top two faculties for architecture and the built environment in the world.

**ISSN 2635-0122**

**This working paper can be referenced as follows:**

Rock, R., Strauss, I., O'Reilly, T. and Mazzucato, M. (2023). Behind the Clicks: Can Amazon allocate user attention as it pleases? UCL Institute for Innovation and Public Purpose, Working Paper Series (IIPP WP 2023-11). Available at: https://www.ucl.ac.uk/bartlett/public-purpose/wp2023-11

# Behind the Clicks: Can Amazon allocate user attention as it pleases?

Rufus Rock, Ilan Strauss, Tim O'Reilly, and Mariana Mazzucato*

December 2023

**Abstract**

We investigate Amazon's power to extract economic rents through the algorithmic arrangement of search results. We analyze 2023 Amazon Marketplace search result data, focusing on product features and the top-3 most clicked products, to estimate what drives user clicks. Our econometric results show that increased visual prominence and positioning ("attention share") correlates strongly with more clicks, even when products have a higher price or worse ratings. Among the top five search results, where typically four are advertisements, we find that neither decreased relevancy nor increased price significantly sways user decisions. This suggests that users tend to satisfice, accepting the products displayed prominently by Amazon's algorithms, rather than exhaustively searching for optimal choices. Advertising on Amazon exploits this behavioural dynamic by transforming product prominence into a mechanism for rent-extraction.

**JEL Codes**: B52; D30; D46; D83; D91.
**Keywords**: Attention Markets; Amazon; Bounded Rationality; Decision Science; Digital Advertising.

*Rufus Rock is a Researcher, Ilan Strauss is a Senior Research Associate, Tim O'Reilly is a Visiting Professor of Practise, and Mariana Mazzucato is a Professor in the Economics of Innovation and Public Value at University College London, *Institute for Innovation and Public Purpose*. Research funded by The Omidyar Network. Corresponding Author: Rufus Rock (rufus.rock.21@ucl.ac.uk), 11 Montague St, London WC1B 5BP, United Kingdom. We acknowledge the helpful comments from Gordon Wei (UK CMA). All errors are our own. Declaration of Interests: None.

# Contents

# 1 Introduction

Theories of platform market power remain hobbled by the fact that platforms generally provide users with a free or subsidized service in order to take advantage of network or scale effects (Rochet et al. 2003; Jin et al. 2021), while monopolistic harm tends to be measured by consumers being charged higher pecuniary prices or their equivalent, i.e. lower quality or less choice (Areeda et al. 2023).[1] As we have argued elsewhere (Strauss et al. 2023), a dominant platform can use advertising as a price substitute (E. Hovenkamp 2018) to get users to pay with more of their attention (i.e., time) than what would prevail in a competitive market, without information and behavioural imperfections.[2] We call this "attention rent" (O'Reilly et al. 2023; Mazzucato, Ryan-Collins, et al. 2023). For users, excessive advertising may lead to inferior product matching (Roth 2015), higher search costs (Areeda et al. 2023),[3] and less variety due to considerable duplication of products in search results (Section 2.4).

Advertising output above a certain level may also unfairly exploit the platform's third-party ecosystem of advertisers and/or suppliers. This follows from each side of a platform competing for the user's attention and the platform's algorithmic attention allocations deciding how value is allocated between these competing sides (O'Reilly et al. 2023; Strauss et al. 2023). When advertising crowds out organic results, it can create a dynamic in which suppliers are strongly incentivized to pay to receive user attention, rather than to earn it competitively. This is precisely what we see on Amazon today: the first four search results are more likely to be advertisements than organic results. As a result, almost one-third (31.8%) of the top-3 most clicked products in the most popular search results on Amazon's third-party Marketplace are sponsored results (Section 2).[4]

A platform's successful exploitation of its users through advertising relies on fairly persistent user click behaviour on advertising results. Otherwise more advertising output will not be profitable and instead will lead to a decline in total profits as advertising prices and/or product sales fall.

This paper examines whether Amazon possesses the algorithmic capability to effectively direct user attention toward specific search results and explores the extent of this influence. More specifically, we estimate the extent to which Amazon is able to exploit users' "position-bias" heuristic (clicks driven by screen position) through the successful algorithmic arrangement and formatting of search results in

[1] Areeda et al. (2023, Section 501. Market Power Defined): "Raise price above the competitive level without loss of sales".

[2] More advertising amounts to a higher effective shadow price paid by users (Baye et al. 2020).

[3] Areeda et al. (2023, Section 2023, Agreements Pertaining to Advertising and Related Dissemination of Product Information): "increased consumer search costs can yield higher pricing".

[4] See also: Fowler (2022), which finds up to half of all Amazon.com search results are advertisements or Amazon products. 71% of third-party sellers on Amazon now advertise to gain user attention (Jungle Scout 2022).

favour of advertising products, which may be of lesser quality (i.e. relevance).

To test this hypothesis, our empirical strategy is to combine scraped[5] search result data from Amazon's third-party Marketplace on product characteristics and positions, with data on the top-3 most clicked products for the same 2,250 product search queries.[6] Our final dataset includes 154,172 product listings collected over 8 days in June 2023. We run a logistic regression to test what drives user clicks (probability of being a "top-3 most clicked" product): Amazon's algorithmic placement of a product in a relatively more favourable screen position, or competitive variables such as product price and relevancy (proxied by organic rank).[7] To estimate a product's relatively more favourable screen position, we create a variable called "attention share", which measures the share of total user attention a product captures relative to all the other products returned by a search query. This accounts for a product's relative size and screen position, adjusted for user cognition (Janiszewski 1998; Luck et al. 1997; Ahn et al. 2018). Our key findings are that:

1. Advertising (i.e. sponsored product results) creates potential harms in several ways when it crowds out organic results and creates incentives for merchants to pay to achieve visibility. Almost one-third (31.8%) of the top-3 most clicked products in the most popular search results on Amazon's third-party Marketplace are sponsored results. Harms from advertising occur, firstly, through directing users to more expensive, less relevant, products. The top-3 most clicked advertised products are around 17% more expensive than organic ones ($19.3 vs. $16.5) and one-third less relevant (organic rank of 4 vs. 3), Table 1. Secondly, advertising harms consumers through creating considerable duplication, thereby greatly reducing product variety facing the consumer (Table 2). One-quarter of product search results on the first page are adverts. This leads to 48.3% of advertised results having at least one duplicate organic result on the first page, and 93.6% of top-3 most clicked ads being duplicated.

2. Econometrically, we find that the relevancy of a product (its organic rank) and the product's relative visual prominence ("attention share") both strongly influence its probability of being a top-3 most clicked product. The "Amazon's Choice" badge heuristic also significantly increases click probability. Product price has relatively little impact once we account for the above factors.

[5] Data scraping is the practice of automated information extraction from a web-page into a database or spreadsheet.

[6] We scrape just the first page of search results, since more than 70% of Amazon shoppers do not click past it (FTC 2023). We only keep products which have both an organic listing and advertised result.

[7] We used a product's organic rank, as determined by Amazon's search algorithms, to proxy for the product's overall relevance, encompassing all algorithmically-considered characteristics, including behavioural factors (Sorokina et al. 2016).

3. We estimate how far Amazon can push a user to click on inferior (less organically relevant) results. We find that increasing a product's attention share can offset the reduced clicks that a less relevant or more expensive product would otherwise receive. Once a product is highly prominent visually, negative product attributes such as high price or low relevancy fade in importance to users' click decisions. This indicates that users may less carefully evaluate highly prominent search results. Among the approximately 69 product search results shown to the user in the first page, an irrelevant[8] but highly visually prominent[9] product is as likely to be clicked on as a super-relevant (top 5 organically ranked) product that is visually obscured in the middle of the page.[10]

These findings have broader implications for the question of whether rents can persist in digital markets. Observed user behaviour online, as predicted by Neoclassical economics, should be highly complex (Simon 1956, 1978). If users are Neoclassical optimizers, clicks should be driven entirely by a product's characteristics and relevancy, not by screen position, unless position perfectly captures relevancy. Unconstrained by search and switching costs, cognitive limitations, or trust and positional biases, users should frequently compare products within and between platforms, instead of clicking on inferior advertised products (H. Hovenkamp 2023). Advertisement sales (or prices) on Amazon would fall over time. If users instead satisfice when making decisions, constrained by time and cognition (Simon et al. 1971), then they would be expected to act "fast" (Kahneman 2011) in a frictionless online environment (Sundin 2021). This means making decisions on the basis of heuristics (Simon 1977; Park et al. 2015), influenced by the screen's "choice architecture" (UK CMA 2022).

This paper contributes to the existing literature on how Amazon abuses its market power, which has thus far focused mainly on self-preferencing and traditional forms of anti-competitive behaviour (Jeffries et al. 2021; Etro 2022; Farronato et al. 2023). We show that platform power in a multi-sided context is not simply power over price (directly), but a power to allocate user attention to inferior content on a fairly persistent basis, in order to exploit a captive supplier base (FTC 2023; Jin et al. 2021). Amazon's third-party Marketplace is particularly susceptible to digital rent extraction as its merchants and advertisers are one and the same group.

Our attention share measure contributes to studies on user click behaviour online and positional bias (Joachims et al. 2017; Wang et al. 2018; Shi et al. 2021), by accounting not just for the relative visual prominence of a search result (Janiszewski 1998; Hong et al. 2004) – but also cognitive aspects

[8] Bottom 10 organic ranked.
[9] Top 1% for "attention share".
[10] Ranked 35 overall.

of attention (working memory) (W. Zhang et al. 2009; Ahn et al. 2018). The influence of a search result page's (SERPs) visual formatting on user attention is often overlooked (Google 2014), despite evidence of its importance (Reutskaja et al. 2011; Floh et al. 2013; K. Z. Zhang et al. 2018). This follows directly from Herbert Simon's insights on the role of information abundance in making user attention increasingly valuable (Simon 1998, 2013).

Our findings can be viewed in terms of the system 1 ('fast') and system 2 ('slow') thinking of Kahneman (2011), which share the control of attention. Algorithmic curation, screen design, and information formatting activate a user's system 1 thinking by framing familiar problems in a familiar way. This effectively circumvents the need for more deliberate system 2 ('slow') thinking, prompting users to rely on algorithmically supported system 1 heuristics for quick and easy decision-making, but prone to biases which can be exploited.

One limitation of our study is that we do not scrape individual product's pages due to time limitations. This means our price variable is not adjusted for the number of units in the product's search page listing. Another limitation is that according to the Amazon's search team, ranking is driven not just by the query or the characteristics of the product but by behavioural factors (Sorokina et al. 2016). But we can only capture these factors indirectly through organic rank.

A major policy implication of our work is that advertising can be exploited excessively by a platform in a native informational environment – and so might require regulation (Strauss et al. 2023; FTC 2023).[11] The key challenge is to establish a reasonable competitive benchmark for this behaviour – or to highlight direct evidence of market power as we have done. The empirical context presented in this paper serves to support the theoretical framework outlined in O'Reilly et al. (2023) and Strauss et al. (2023), aiming to ground regulatory and policy efforts in determining acceptable levels of such algorithmic influence. Regulating undue algorithmic influence, and in turn enhancing platform governance, requires better disclosures of a platform's operating metrics (Mazzucato, Strauss, et al. 2023).

Section 2 details our dataset and defines our key quantities of interest. Section 3 specifies our logistic regression model and core hypotheses. Section 3.3 discusses the findings from our regression, with some discussion of robustness. Section 4 extends our model to see how much Amazon can degrade results quality while still garnering user clicks. Section 5 provides a summary and conclusion. Our Appendix contains further information on the concepts of "attention competition" (A.3) and "attention decay" (A.4), data collection and cleaning process (A.1) and model validation (A.7).

[11] On native advertising see Wikipedia. "Native advertising". Online: https://en.wikipedia.org/wiki/Native_advertising.

# 2 Sample & Variables

This section provides an overview of our dataset and variables. We then discuss how we construct our measure of a product's relative visual prominence on a page - which we call its "attention share" - and combines attention competition and attention decay theories. We conclude by providing key descriptive statistics on our sample focusing on advertising results, duplication, and position driven user click behaviour.

## 2.1 Dataset Construction

Our product-level dataset combines data on product characteristics with data on user clicks from Amazon search results. First, we scraped product search results between 11-19$^{\text{th}}$ June 2023 from Amazon.com's third-party Marketplace covering: product price, number of reviews, mean product rating, search rank, "Amazon's Choice" status, "Best Seller" status, "Prime" status, product title, organic rank, size, page position, and whether the product is advertised.[12] We obtained data on the three most-clicked products from Amazon's Seller Central Search Term Report, focusing on 2,250 of the most searched queries in Amazon's third-party Marketplace. Each query was scraped once.

Only products with complete data were included. We remove products that are labelled as "out-of-stock" as they do not have price information. We also remove advertised products that are not duplicated as organic results in the search results page, as this ensures the availability of the "organic rank" variable (Section 2.2). This means that only results which feature as both an ad and an organic result are included in our regression sample. Our regression sample ($n$) is 112,727. This is because we run the regression on a sub-sample of the 154,172 observations and reserve 41,445 observations for validation. Our regression sample only contains product results that are in a matrix format page, i.e., a page where products are arranged in a grid with multiple products per row. We exclude from our sample any results provided in a list format. i.e, one product per row. This exclusion is due to the inability to directly compare the relative visual prominence of individual products across different layouts. For robustness, we ensure that our model's performance is maintained across both formats (Section 4.2).

We only scrape the first search engine results page (SERP) since more than 70% of Amazon shoppers do not click past it (FTC 2023). We define "page" as the complete webpage, including all the search

[12] To scrape our dataset we developed a custom Python(3.8) program utilizing Selenium WebDriver - an open source Python web-browser automation module: https://www.selenium.dev/documentation/webdriver/. Our program and dataset are available in the online appendix.

results on the first page returned for a specific Amazon query. This is distinct from a "screen", which refers to a portion of the page that is simultaneously viewable in the browser window. For example, the first screen contains the content that appears before any scrolling, and scrolling down to the bottom constitutes viewing the entire page. We do not scrape specific product pages, which would require us clicking on each individual product in the search result, because this would cost significant time and computational resources. This introduces some noise into the variables scraped, most importantly price, because prices will not be adjusted for the number of units included in the product listing (i.e., they are not normalized by quantity). This also means we can't scrape delivery time or the distribution of product reviews by star. But review data performed poorly in initial regressions and so no attempt was made to get more fine-grained review data. We ignore banners and video ads from the scraped results, including when calculating positions, focusing only on traditional search results as well as those included in side-scrolling carousels.

## 2.2 Variables

**Organic Rank**. Organic rank is an index telling us where a product appears in a list of Amazon's organic search results for a specific search query, when advertising is not included. The ordering of this organic product results list is driven by Amazon's organic search algorithm, which takes into account factors such as a product's relevancy to the user's query, its price, popularity (Sorokina et al. 2016), delivery time, and other factors that Amazon has determined will make the results most likely to satisfy their users and result in sales. Consequently, we use organic rank to define a product's relevancy across all of these dimensions, rather than a purely textual measure of how closely a product's title matches the user's search query.

**Search Rank**. Search rank is the indexed result on the screen of a product amongst all search results shown, including both organic and advertised results. Due to the proliferation of advertising results, a product's organic rank can deviate greatly from its search rank. A product might have an organic rank of 3, for example, indicating that Amazon's organic algorithm believes the product is the $3^{rd}$ best choice for the user's search, but if there are 10 advertising results placed ahead of it, then its search result rank will be 13. Advertising may indirectly influence a product's organic ranking to the extent that sales volume is used by Amazon as a ranking factor (Turton 2023). We accept this limitation given that Amazon's estimation of relevancy is based on a far more comprehensive set of variables than is available to external evaluators. We do not use search rank as a predictor.

**Price**. The price variable captures each product's $USD total price as listed on the main search result page. A limitation of our dataset is the omission of unit-price data, which tends to be found only on the the product's specific page. As a result we cannot account for instances in which the price listed on the results page includes multiple units of a product. We mitigate some of this noise by normalizing our price variable to between 0 and 1 for products within each search term (1 is the most expensive, 0 the least). This ensures that the price variable is search context specific.

**Attention Share**. Attention share captures the relative visual prominence of a product on the page, in terms of the attention it is likely to capture from the user (Section 2.3 below). A product's relative visual prominence does not track the product's organic rank that closely because of the noise generated by advertising in search rank positions (Section 4.2). This allows us to use both organic rank and relative visual prominence ("attention share") as separate predictors in the same regression. To compute attention share, we scrape various aspects of the SERP's information format: the dimensions and coordinates of each product, the dimensions of the browser window, and the locations of visual interruptions on the page such as banner/video ads or carousels.

**Amazon's Choice**: Amazon's Choice is a binary variable indicating whether a product has an "Amazon's Choice" label.

**Top-3 Most Clicked Products**: Top-3 most clicked products is a binary dependent variable indicating whether or not a product is reported by Amazon in the "Seller Central Search Term Report" as being one of the 3 most clicked products in a search result for a given search query. The Report does not detail the distribution of clicks between each of the product's cloned listings, when the listing appears multiple times on the search result page (Appendix A.2). Therefore, when a top-3 most clicked product appears multiple times in a SERP, we label each of the product's listings as top-3 most clicked product, adjusting the attention share metric for product duplication, as described in Section 2.3.

## 2.3 Attention Share

In this subsection we explain the rationale behind a key predictor of our regression called "attention share". This measure proxies for Amazon's ability to allocate value, through allocating more user attention to a product, by changing the relative prominence of a product in search results.

We combine the dual theories of attention decay (Luck et al. 1997; Ahn et al. 2018) and attention competition (Janiszewski 1998; Hong et al. 2004) to estimate the "attention share" that a product search result is likely to capture from the user. A larger attention share value indicates that more attention from the user is likely to be allotted to that product. For example, a product with an attention share of 1% implies that it will likely receive 1% of the total attention that a user dedicates to all the products on the search results page (across screens). Attention share combines, through multiplication, a product's relative visual prominence (attention competition) with a user's working memory capacity (attention decay). To find the attention share for each product, $j$, for a given Amazon search result query, $q$, we multiply the product's attention decay and competition values:

$$\text{Attention Share}_{jq} = \text{Attention Decay}_{jq} \times \text{Attention Competition}_{jq} \tag{1}$$

The final attention share for each product is normalized by dividing it by the sum of all attention share values on the page. Product $j$'s attention share is always query specific $q$, because it is a function of the layout of the products on the page. Any visual interruptions are also relative to position within that specific results page. Amazon's algorithm dynamically determines each page layout and can therefore shape product attention competition and decay dynamics.

Based on the formulae, we can see that a product will receive more attention if a small amount of the user's working memory has decayed away (W. Zhang et al. 2009; Ahn et al. 2018; Lorigo et al. 2008), and/or the product is more visually prominent than its peers (Janiszewski 1998; Hong et al. 2004).[13] Conversely, a small attention share value indicates that the product is less prominent, existing in a more visually crowded part of the page, and/or a large part of a user's attention has decayed away.

**Attention Decay**. A user's attention devoted to sequentially presented information shown on the screen is found to decay exponentially in online search environments (W. Zhang et al. 2009; Ahn et al. 2018; Luck et al. 1997). We apply a parameter ($\lambda_1 = -0.809$) from Ahn et al. (2018) to model the rate of decay.

But attention can also be reignited. This is called "attention renewal" (Berti et al. 2003; Lépine et al. 2005; Ahn et al. 2018), such "interruptions" in the form of video ads, banners, or other changes in how search results are presented, can refresh the visual working memory buffer of the user. Therefore, the decay coefficients are not fixed across all products (see Appendix A.4): they are query specific, represented by subscript $q$, since visual interruptions appear in different locations for each query. Attention decay for product $j$, with search result position or index $i$, is defined as:

[13] We invert the attention competition metric to align with the decay scale, enabling coherent multiplication.

$$\text{Attention-Decay}_{qj} = e^{-\lambda i}$$

**Attention Competition**. The amount of visual competition a product, $j$, receives from all of the simultaneously visible products on a screen is a function of three variables (Janiszewski 1998; Hong et al. 2004):

$$\text{Attention-Competition} = C_i = \sum_{j \neq i}^{\text{Competing}} \frac{\sqrt{\text{Size}_j}}{\text{Dist}_{ij}} \tag{2}$$

- $\text{Size}_j$: The size (area) of a competing search result $j$.
- $\text{Dist}_{ij}$: The distance between $i$ and $j$.
- $n$: The number of competing results viewable simultaneously as $i$.

Attention competition and attention decay provide complementary levels of analysis. Attention competition provides a local indication of how visually prominent a search result is amongst its peers on the screen at any given time, whilst decay provides a global analysis of how a user's attention is likely to be allocated between all the products on the page. Again, by "page" we refer to the entire webpage returned by the Amazon search whilst a "screen" refers to a portion of the webpage which is simultaneously viewable in the browser window without scrolling. Decay theory describes how much of the user's visual working memory remains available for processing new product information, whilst competition theory reports how likely a result is, amongst those that are visible at a given time, to attract a user's remaining quantum of attention. Attention competition, unaided by decay, cannot provide an explanation for user's scrolling behaviour because it can only describe a product's relative prominence amongst simultaneously viewable products with a screen, not between products on the whole page.

We show how attention share is distributed empirically among products on the search result pages below. More details on attention decay and competition theories can be found in Appendices A.4 and A.3, respectively. Lastly, empirical justification for our metric is in Appendix A.5.

## 2.4 Descriptive Statistics

Table 1 details key features of our dataset, focusing on our quantity of interest, which is the products that are amongst the top-3 most clicked for a given query by users. Firstly, particularly striking is that the median attention share of a top-3 most clicked product is 5.5% - meaning that it captures around

5.5% of the total attention that the user allocates to all the products in the search results page (across all screens). This underscores the importance of a product having a large relative visual prominence in order to capture a user's clicks.

Table 1. Top-3 Most Clicked Products: Ads vs. Organic Characteristics

| | Median Top-3 Clicked | MAD Top-3 Clicked | Median Top-3 Ads | Median Top-3 Organic |
|---|---|---|---|---|
| Attention Share (%) | 5.5% | 4.5% | 12.1% | 3.4% |
| Organic Rank | 3.0 | 4.2 | 4.0 | 3.0 |
| Price ($) | $17.2 | $26.4 | $19.3 | $16.5 |

*Note*: MAD is the median of the absolute deviations from the data's median, a measure of dispersion. A smaller organic rank is better, implying a higher rank on the page. Attention share values represent the portion of the attention that a user allocates to the entire page (across all screens) that is directed towards a specific product. A value of 10% would indicate that a specific product is likely to attract 10% of the users attention. The median indicates that a typical product captures only 0.5% of the total user attention dedicated to a webpage.

Secondly, the table highlights that the top-3 most clicked products overall have highly competitive organic ranks, suggesting that consumers do have a preference for very relevant products. However, this varies by ad vs. organic top-3 product result. Top-3 product listings that are ads have an organic rank of 4, compared with top-3 product listings that are organic results which have an organic rank of 3 – a 33% difference. This may seem small visually but in practise can lead to a significantly different screen position for the product. However, the median absolute deviation (MAD) of a top-3 product's organic rank is 4.2. This is larger than the median and indicates that the data is significantly skewed and/or has outliers. Around half of the top-3 most clicked products have an organic rank lying between 1 and 7.2.[14] The skew is caused by the expansive range of organic ranks that exist within the top-3 ads category, suggesting that advertising might be capable of stretching this preference for relevancy.

Thirdly, the table shows that top-clicked advertising results are more expensive than top-clicked organic results ($19.3 vs $16.5) or 17% higher. This may be because advertisements appear in more expensive product categories, or because advertising fees are passed along to consumers in the form of higher prices. However, the large median absolute deviation for price is indicative of significant variability in the prices of the most clicked products. Moreover, our price variable is noisy since it does not adjust for the number of units in the listing.

[14] About half of the data points lie within the range (median - MAD) to (median + MAD), constrained by the minimum organic rank being 1.

Next, Table 2 below provides some key descriptive statistics on the composition of the dataset as a whole. Of particular interest is the high proportion of top-3 clicked product listings that are advertisements (31.8%) - compared to 24.6% in our sample as a whole.

Table 2. Key Descriptive Findings from our Sample

| | |
|---|---|
| *Top-3 most clicked Products* | |
| % of Products that are Top-3 most clicked | 6.4% |
| % of Top-3 most clicked Products that are Ads* | 31.8% |
| *Amazon's Choice* | |
| % of Products that are "Amazon's Choice" | 1.9% |
| % of "Amazon's Choice" Products that are Top-3 most clicked Products | 73.9% |
| *Advertising* | |
| % of Products that are Ads* | 24.6% |
| % of Ads that are Duplicated* | 48.3% |
| % of Top-3 most clicked Ads that are Duplicated* | 93.6% |

*Note*: (*) Denotes statistics taken from a larger sample, before filtering out advertised products which do not have an organic counterpart (or duplicate).

Of the advertisements which are for top-3 most clicked products, over 90% are duplicated as an organic result, despite less than half of all ads being duplicated (48.3%). This means that flooding the screen with your product multiple times is the clearest path to click-based success on Amazon. The median advertisement appears 17 search ranks above its organic duplicate. However, top-3 most clicked advertisements tend to receive smaller boosts from sponsorship. Further research is required to properly investigate duplication behaviour on Amazon. This contrasts with Google's policy of "unfair Advantage", whereby one cannot advertise two listings for the same keyword. Google prevents any one advertiser controlling all the real estate (Meyerson 2023). Finally, 73.4% of "Amazon's Choice" products are amongst the top-3 most clicked on the page. This label's sway over consumer decision making raises questions about the undisclosed nature of the algorithms that assign it (Mazzucato, Strauss, et al. 2023; Shifflett et al. 2019).

Lastly, we provide heatmaps below to visualize our ranked search results data at a more fine-grained level. The heatmaps illustrate where user clicks for the top-3 most clicked products (Figure 1) and for product advertisements (Figure 2) are located on Amazon's search results pages. The top left box represents the first product search result, the second box to the right represents the second product search result, and so on. Figure 2 on the right demonstrates the degree to which organic results are being crowded out by advertised results. The first results slot has an 80% chance of being occupied by an advertising result, declining to 63% for the second slot.

Figure 1. User Clicks for Top-3 Clicked Products are Highly Concentrated

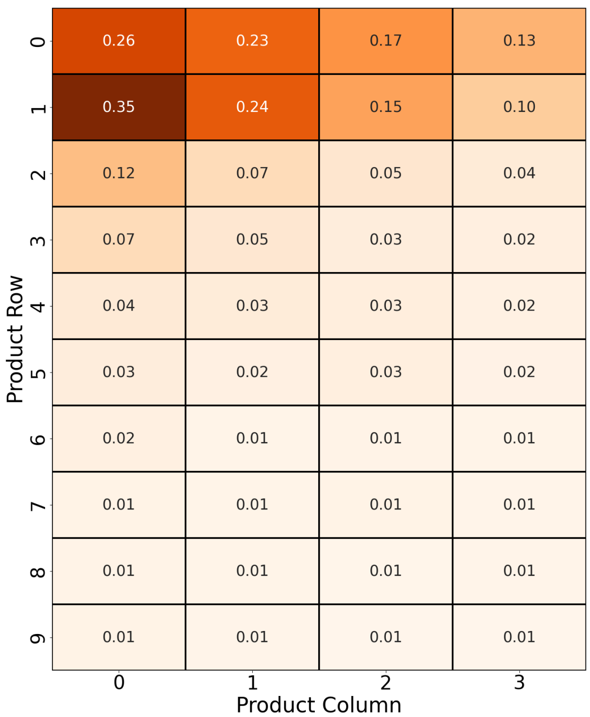


*Note*: This heatmap shows where top-3 most clicked products tend to be located on Amazon SERPs, illustrating the large influence of position-bias. Each square in the heatmap represents a screen position on an Amazon search page, with the relative shading indicating the probability of each position being occupied by one of the 3 most clicked products on the page. The density of top-3 products is strongly indicative of consumers' "position-bias".

Figure 2. The First Search Result has an 80% Chance of being an Ad

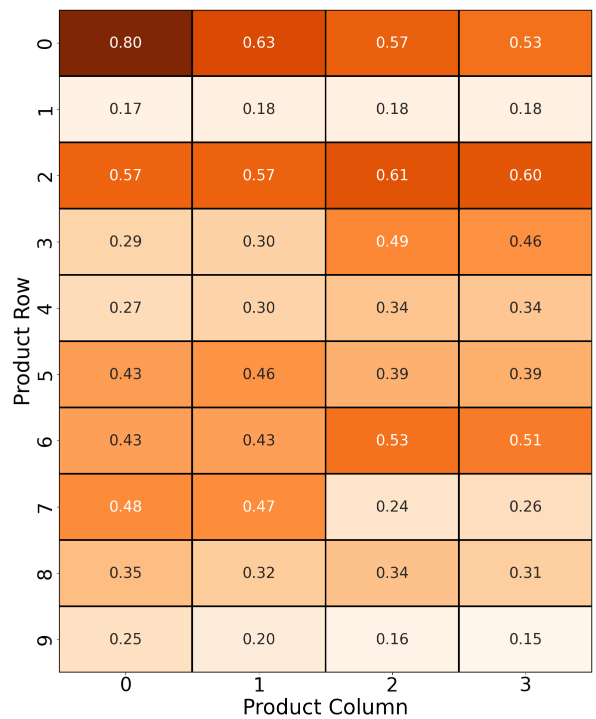


*Note*: This heatmap shows the probability of a search result slot on an Amazon results page being occupied by an Ad. This figure illustrates the degree of advertising saturation, and crowding out of organic results, present on Amazon SERPs.

The fifth result has a 17% chance of being occupied by an advertising result. The observed volume of advertised products in search results, whose position does not depend on their organic relevancy, substantially weakens the correlation between the attention share and the organic rank of products (and reduces regression endogeneity – Section 4.2).

Figure 1 on the left shows the influence of "position-bias" on clicks, revealing that the product in the top product position (the left-hand slot on the first row) has 0.26 probability (26% chance) of being one of the top-3 most clicked products on the page. The darker squares indicate a greater probability of being a top-3 most clicked product; the figure visualizes user position-bias by revealing that clicks are condensed in the top left hand quadrant of the page. The $5^{\text{th}}$ search result position holds the highest probability, at 35%, of featuring a top-most clicked product. The $5^{\text{th}}$ position is situated on the second row, beneath the screen fold (on matrix-format search result pages). This finding is notable because it contradicts prior research (Fessenden 2018), which emphasizes that consumer attention is predominantly allocated to products that are "above the fold", meaning that ones that are visible without scrolling.

Below-the-fold positions, such as those on the second or third row, may owe their higher than previously indicated probability of being occupied by a top-3 most clicked product to organic results being pushed further down the screen. The first position which has a higher chance of being an organic result than an advertised result is now the fifth result slot (having only a 17% chance of being an ad), as shown in 2. The fact that the $5^{\text{th}}$ slot is the most likely to contain a top-clicked product and the first organic result suggests that users are becoming more willing to spend their time scrolling in order to locate the most relevant, organically provided, search results. Finally, Figure 1 demonstrates that the left-hand column of product positions retains a slight advantage in receiving user attention over the second column, and is indicative of the frequently observed "F-shaped" routine undertaken by users to inspect search results (Pernice 2017; Shrestha et al. 2007).

# 3 Model and Results

This section details our logistic regression model, which predicts the probability of a product being a top-3 most clicked product for a given search query. We then estimate the model using maximum likelihood and present the results for discussion.

## 3.1 Baseline Logistic Regression Model

Our baseline logit model (Gelman et al. 2006) predicting a product's chances of being a top-3 most clicked product is estimated up until $\beta_4$ (i.e. with no interaction effects):

$$P(Y_j = 1) = \frac{e^{\beta_0 + \beta_1 X_{1j} + \beta_2 X_{2j} + \beta_3 X_{3j} + \beta_4 X_{4j} + \beta_5 X_{1j} * X_{2j} + \beta_6 X_{1j} * X_{3j}}}{1 + e^{\beta_0 + \beta_1 X_{1j} + \beta_2 X_{2j} + \beta_3 X_{3j} + \beta_4 X_{4j} + \beta_5 X_{1j} * X_{2j} + \beta_6 X_{1j} * X_{3j}}} \tag{3}$$

This runs from $j = 1 ... 112,732$ (instead of our full sample which is 154,172), since we set aside data points for model evaluation. The last two interaction terms of the model, $\beta_5$ and $\beta_6$, are excluded from our benchmark model. We later extend this model by including interaction effects, $\beta_5$ and $\beta_6$ (Section

4). The dependent variable $Y_j$ is the top-3 most clicked products in the search results pages for a given query. $P(Y_j = 1)$ is the probability that a product is a top-3 most clicked product.

The predictors for a given product $j$ are: $X_{1j}$ a product's attention share, its relative prominence in the search result for the user compared to all other products, ranges from near 0% to 30%. $X_{2j}$ is the product's organic rank, which proxies for its relevancy, and ranges from 1$^{\text{st}}$ being the most relevant (highest ranked) to around 50$^{\text{th}}$ being the least relevant (lowest rank) depending on the number of products on the page. $X_{3j}$ is the product's $ total price unadjusted for units sold. $X_{4j}$ is a dummy variable indicating whether the product is an Amazon's Choice. $\beta_0$ is a mean-centered constant term. Noisy variables that were not economically significant were removed.[15] A dummy variable representing whether or not a product is an advertisement was not included because when included it generated significant noise in the model. This is likely because top-3 most clicked products are often duplicated, appearing once as an organic result and once as an advertisement on the same page (see A.2). Furthermore, our dataset cannot identify how clicks are allocated between duplicated listings (Section 2).

We log and then mean-centre both organic rank and attention share variables. This significantly improves the stability of our model given how skewed several of the distributions are, while also providing the benefit of being able to meaningfully interpret the intercept coefficient and helping with interpreting the interaction effects later on (Gelman et al. 2006).

One way of conceptualizing what theoretical quantity our logit regression, Equation (3), is that it estimates a platform's direct market power to allocate user clicks (demand) as it degrades quality (O'Reilly et al. 2023). Our approach applies an institutional decision sciences perspective to the question of Amazon's market power and its estimation in a context where prices are not determinative of user welfare (Simon 2013; Strauss et al. 2023). For a firm with market power, the probability of a user clicking on a search result will remain high even on less relevant (lower quality) information (Newman 2015; Calvano et al. 2021; FTC 2023). Our regression, therefore, asks: after accounting for information quality or relevancy (i.e. the result's organic rank) can Amazon allocate clicks by raising the relative prominence of the result in the user's attention sphere, after accounting for quality differences? This is captured by the "attention share" variable. So for a firm with market power, after accounting for product quality (organic rank), we would expect to see a large and positive effect from "attention share' on user clicks,

[15] We initially incorporated two quality control variables into the regression, including the product's star rating and number of reviews. However, they were excluded from the final model. The product's star rating was found to be statistically insignificant, failing to improve the fit or performance of the model. Number of reviews, on the other hand, was omitted because it exhibited very high multicollinearity with the organic rank variable, destabilising the regression.

such that as Amazon increases the prominence it gives to a product, it can receive an increase in click probability. This is similar yet distinct from a regression that estimates a traditional elasticity, which looks at the percentage change in demand (dependent variable) arising from a percentage change in price or quality (independent variable). However, both are ways of testing to see if a firm has market power.

### 3.2 Hypotheses

In effect, our model estimates two competing hypotheses regarding the drivers of user clicks, one neoclassical and one institutional (Strauss et al. 2023):

- **Institutional Satisficing**: *Room for rents.* Following Simon (1978), users satisfice, making them reliant on algorithmic heuristics, such as product screen placement, for decision making under informational complexity. The "Attention Share" coefficient is expected to be + and large, indicating that as a product increases in relative prominence for the user compared to other search results, the likelihood of the product being clicked on increases. This would indicate that Amazon has considerable discretion in allocating value as it sees fit, since it can algorithmically increase the relative prominence of a product to increase click probability, irrespective of the product's intrinsic characteristics.

- **Neoclassical Optimizing**: *No room for rents.* Users optimize by engaging in self-directed search independent of algorithmic ordering and positioning of results (H. Hovenkamp 2023). Attention Share is expected to be small (economically insignificant) if users optimize, while relevancy (Organic Rank) would be large, negative, and economically significant, as users always click on the most relevant result. Organic Rank would be negative since a smaller organic rank means a more relevant rank, and so a greater effect on top-3 click probability for a product, implying an inverse (negative) relationship.

  Prices are largely uncorrelated with rank or organic rank, so under optimization we would also expect the Price coefficient to have a meaningful impact on outcomes, after estimating other effects. Users inspect and compare the intrinsic characteristics of every product and choose the optimal one, eliminating the impact of algorithmic positional heuristics on clicks (when the heuristic deviates from the optimal). According to this hypothesis, Amazon would be highly limited in its ability to extract rents from its suppliers by charging for screen visibility, since user behaviour is not driven by position and is therefore difficult to influence through advertising.

The robustness Section 4.2 shows that organic rank and attention share have only a mild correlation (given the ubiquity of advertising highlighted by Figure 2), ensuring that each variable identifies a separate hypothesis.

### 3.3 Baseline Model Results

Table 3 the results from our logit model estimated using maximum likelihood. All variables are significant. The "pseudo $R^2$" value is high at 0.46, even without fixed effects for product category and search query. Including these variables overfits the model – especially when we include interaction effects in the next section – and so are omitted. Noisy variables with little economic or statistical significance were excluded from the regression, including a product's star rating and number of reviews.

Overall, our findings suggest that Amazon is able to extract "attention rents" from its users since the "attention share" coefficient $\beta_2$ (relative visual prominence) is large and positive at 0.56, indicating a large positive impact of a product's relative position on the probability of a product being one of the top-3 most clicked on the page, after accounting for quality and price. However, the regression shows that users also demonstrate a preference for the most relevant products, since a product's organic rank $\beta_1$ has a large negative coefficient at $-1.44$, meaning that as organic rank becomes smaller (i.e. more relevant), click probability increases. This indicates that Amazon's power to allocate user clicks to advertising products may be circumscribed – something we explore in the next Section 4.1. We explore these coefficients in more detail below.

The log-odds of an "average"[16] product being one of the 3 most clicked on the page are given by our intercept, -4.20, as our predictors are mean centered or dummies. To convert log-odds into a probability, we use the inverse logistic function,[17] resulting in the average product having a probability of 0.015 (1.5%) of being a top-3 most clicked product.

**Attention Share,** $\beta_2$. The coefficient for attention share is large and positive ($\beta_2 = 0.56$). This indicates a large positive impact of relative visual prominence on the probability of a product being amongst the top-3 most clicked on the page. The impact, measured in log-odds, of a one unit increase

[16] A product with mean values of each of our predictors and no Amazon's Choice label.

[17] The inverse logistic function is given by:

$$P(X) = \frac{1}{1 + e^{-x}}$$

Table 3. Regression Results Predicting Top-3 Click Probability

| | Coef. $\beta$ (Standard Error) |
|---|---|
| Intercept $\beta_0$ | -4.20 (0.03) |
| Log(Organic Rank) $\beta_1$ | -1.44 (0.02) |
| Log(Attention Share) $\beta_2$ | 0.56 (0.01) |
| Amazons Choice $\beta_3$ | 1.87 (0.06) |
| Price $\beta_4$ | -0.54 (0.08) |
| df Residuals | 112,727 |
| df Model | 4 |
| Pseudo $R^2$ | 0.46 |
| Log-Likelihood | -15,784 |
| LL-Null | -29,629 |

*Note*: Estimated logit results from regression Equation 3 using maximum likelihood. No of observations ($n$) = 112,727. Dependent variable is a dummy variable indicating whether the product is a top-3 most clicked product in the search result query. All pseudo p-values are significant at the 0% level. LLR p-value = 0. $\textsf{Attention Share}_j$, $\textsf{Organic Rank}_j$, and $\textsf{Price}_j$ are mean centered.

in logged attention share (holding our other predictors constant at their mean value) is equal to the log-odds falling from -4.20 to -3.64 ($\beta_0 + \beta_2 \times 1$). We can then convert this to a probabilistic difference using the inverse logistic function described above. A one logged unit increase in attention share leads to an increase from 1.0% (for an average attention share product) to a 2.5% chance of being a top-3 most clicked product. Following the same interpretation, a product in the top decile by attention share (i.e., one of the first 5 most prominent products on a search results page) has a 4.5% chance of being a top-3 most clicked product (an increase of +3% from its position alone), holding all of our other predictors constant at their mean. The same product, given a bottom 10% attention share position receives a -0.9% decreased chance compared to an average attention share product.

Figure 3 plots the the predicted probability of a product being amongst the top-3 most clicked on the page (y-axis) as a function of attention share (x-axis), while holding all other predictors constant at 0 (their mean values). The curve is generated by applying the inverse logistic function to the sum of the intercept ($\beta_0$) and the product of the attention share coefficient ($\beta_2$) with attention share values. The x-axis spans the range of attention share observed in the regression dataset, from the minimum to maximum observed value.

The figure visualizes the impact of attention share on the predicted probability of being a top-3 most clicked product. One can see that, initially, as attention share increases, the probability of being a top-3 clicked product also gradually increases. Between the minimum attention share value (at the origin) and the 73rd percentile of attention share values (around half way along the x-axis) the probability of being a top-3 most clicked product increases to around 0.02 (2%). As the attention share value increases into the upper quartile, the rate of increase in click probability also increases. Between the 90th percentile and maximum values of attention share, the probability increases by +9% (from 5% to 14%).

Therefore, our attention share coefficient indicates that the relative visual prominence of a product (measured by attention share) significantly influences its chances of being one of the top-3 most clicked on the page. This is especially true of the elite screen positions, in the top 10% of those on the page, indicative of position-biased user behaviour.

Figure 3. Probability of Being a Top-3 Product ("Click Probability") Increases with Attention Share

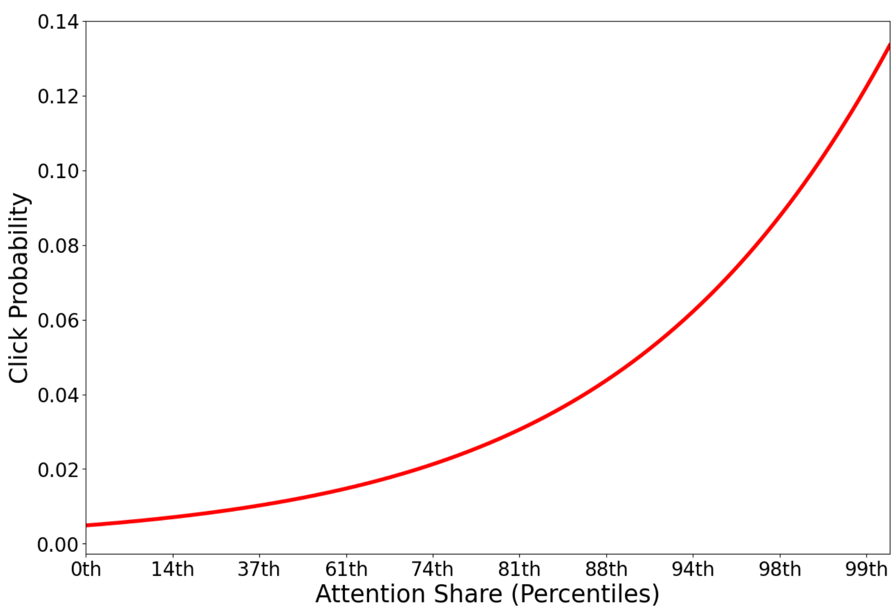


*Note*: This graph shows the impact of attention share on predicted top-3 probability, according its regression coefficient. The x-axis represents for the range of attention share values we observe in our dataset. This graph holds all other predictors constant at 0 and is given by: $Y = \frac{1}{1+e^{\beta_0+\beta_2\times X_{2j}}}$ where $\beta_0 = -4.20$ and $\beta_2 = 0.56$.

The significant role of relative visual prominence in driving clicks supports Simon's notion that heuristics drive decision making in information dense environments. Evidence is also provided by the large positive coefficient for the Amazon's Choice dummy ($\beta_3 = 1.87$). Interpreting this coefficient, we find that when a product (holding our other predictors constant at their mean) has an Amazon's Choice label, the log-odds of it being a top-3 most clicked product increase by +1.87, translating to a probabilistic increase of +8% for a product with average prominence and relevancy. Therefore, our coefficient signals that the "Amazon's Choice" label significantly drives clicks, a clear example of heuristic led consumer behaviour.

**Organic Rank,** $\beta_1$. The coefficient for organic rank is large and negative ($\beta_1 = -1.44$). This suggests a large negative impact of decreasing relevancy on a product's chances of being one of the top-3 most clicked on the page. A product's organic rank increasing by one logged unit from the mean is equivalent to it falling from the 25$^{\text{th}}$ organically ranked product to the 69$^{\text{th}}$ and results in its probability of being among the top-3 most clicked products on the page dropping from 1.5% to 0.4%. Conversely, going up the organic rankings by one logged unit (25$^{\text{th}}$ to 9$^{\text{th}}$ rank) leads to a probabilistic increase of 4.2% (from 1.5% to 5.9%).

The first organic product, the most relevant on the page according to Amazon, has a 51% chance of being a top-3 most clicked product on that page.[18] However, these positive effects from high relevancy rapidly diminish. This is visualized in Figure 4 which shows the relationship between organic rank (x-axis) and the predicted top-3 product probability (y-axis). It is calculated similarly to 3, using the organic rank coefficient (and holding attention share constant at 0).[19] The graph shows that very relevant products, in the top 5 by organic rank (1-5 on the x-axis), have a large, albeit quickly diminishing, predicted probability of being a top-3 most clicked product.

The size of the organic rank coefficient explains the density of clicks we observe in the second row of results, below the fold, where the initial organic results are located. Our organic rank coefficient points to a strong user preference for first few organically ranked products, the most relevant on the page according to Amazon's algorithm. Thus, our coefficient indicates that users tend to be willing to spend their time scrolling, past increasing numbers of advertisements, in order to find the most relevant organic results.

**Price,** $\beta_4$. Thirdly, our small and negative price coefficient indicates a slight consumer preference for cheaper products ($\beta_4 = -0.54$). The relative size of the coefficient (compared to our other predictors) is misleading due to the smaller scale of the price variable, caused by it being normalized to lie between 0 and 1 (within each search query), resulting in reduced economic significance. The relative price difference between the most and least expensive products on a page affects the probability of being amongst the 3

[18] The minimum value of the underlying organic rank variable is 1, representing the top organically ranked product on the page. After transformations the minimum becomes -2.94. Holding our other variables constant, we can determine the log-odds of this product being top-3 most clicked $\beta_0 + \beta_1 \times -2.94$, giving us 0.075. Using the inverse logistic function on 0.075, gives a probability of 0.51 or a 51% chance of being a top-3 most clicked product.

[19] It should be noted that an organic rank of 40 (the upper limit of the x-axis) is not the absolute maximum observed in our dataset (82).

Figure 4. Worse Organic Ranking Decreases Probability of a Product Being Top-3 Most Clicked

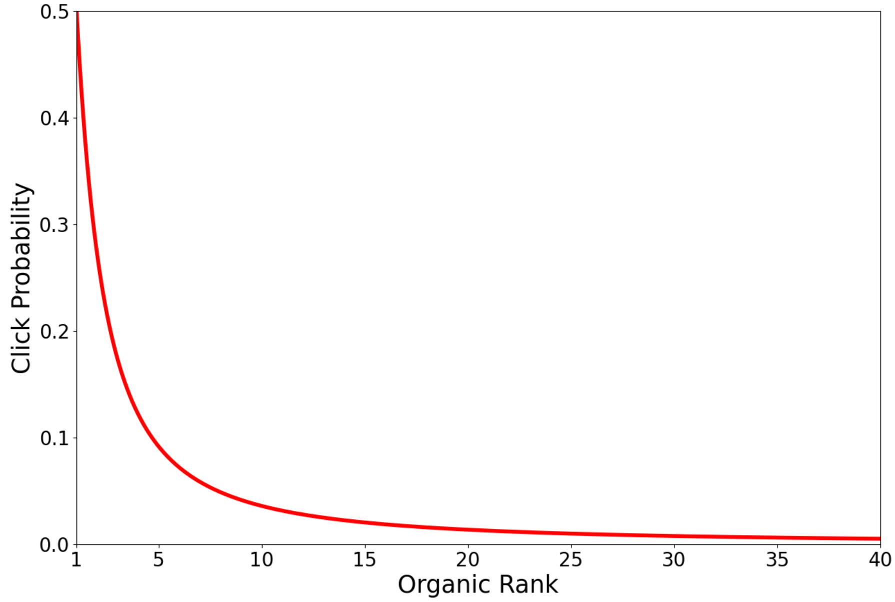


*Note*: This graph shows the impact of organic rank on predicted top-3 probability, according to its regression coefficient. To calculate we hold all other predictors constant at 0. The graph is given by: $Y = \frac{1}{1+e^{\beta_0+\beta_1 \times X_{1j}}}$

most clicked by only $\pm0.35\%$, holding our other variables constant. Given the scale of Amazon's platform, a $\pm0.35\%$ impact on the probability of being a top-3 most clicked product is not inconsequential. However, compared to our other predictors, price is by far the least influential that we include. The interpretation of this finding is complicated by a product's price presumably being one of the many factors considered by Amazon's organic ranking algorithms. The small coefficient may be caused by price information being contained within the organic rank variable. However, we observe a negligible variance inflation factor (VIF) and correlation coefficient between the price and organic rank, revealing that if price does influence organic rank, it is not a straightforward relationship.

# 4 Model Extension: How much market power does Amazon have?

This section extends the baseline model (Equation 3) by including the previously omitted interaction effect terms, $\beta_5$ and $\beta_6$, to assess how much market power Amazon has to degrade product result quality (organic rank relevancy) while still garnering user clicks. The extended model with interaction effects allows us to properly assess the extent to which Amazon can trade-off relevancy (quality) and still receive user clicks by increasing the share of user attention devoted to the product.

## 4.1 To What Extent Can Amazon Trade off Product Relevancy for Prominence?

We extend our baseline econometric model by adding two interaction terms to our regression:

$$\text{Interaction 1 } (\beta_5): \quad \textsf{Attention Share} \times \textsf{Organic Rank}$$
$$\text{Interaction 2 } (\beta_6): \quad \textsf{Attention Share} \times \textsf{Price}$$

These interaction effects estimate how the impact of a product's "attention share" on user click probability is mediated by a product's relevancy and its price. In essence, we test to what extent a product is able to use its relative visual prominence to trade off decreased relevancy or increased price. Interaction terms allow for the possibility that the effect of one predictor on the outcome is modified by the level of another (multiplicative effects). This is equivalent to estimating multiple random coefficients which vary by sub-groups, calculating not just average but marginal effects (Gelman et al. 2006). As a result, a much larger sample size is required to estimate interaction effects.

Table 4 shows the results from our extended logit regression model with interaction effects. P-values remain significant and the previous coefficients do not change much after adding in interactions. Interaction effects slightly improves model fit: the log-likelihood is reduced and the pseudo-$R^2$ increases marginally.

Table 4. Interaction Effects From Extended Logistic Regression

| | Coef. $\beta$ (Standard Error) |
|---|---|
| Attention Share $\times$ Organic Rank $\beta_5$ | 0.21 (0.01) |
| Attention Share $\times$ Price $\beta_6$ | 0.29 (0.07) |
| df Residuals | 112,732 |
| df Model | 6 |
| Pseudo $R^2$ | 0.47 |
| Log-Likelihood | -15,652 |
| LL-Null | -29,629 |

*Note*: Estimated logit results from extended regression Equation 3 including interaction effects $\beta_5$ and ($\beta_6$). Number of observations ($n$) = 112,732. Dependent variable is a dummy variable indicating if the product is top-3 most clicked product in the search result query. All pseudo p-values are significant at the 0% level. LLR p-value = 0.

Both our interaction effects are positive and similar in coefficient size: our attention share and organic rank interaction is $\beta_5 = 0.21$ and our price interaction is $\beta_6 = 0.29$. However, $\beta_5$ (attention share and organic rank) has a much more significant economic impact due to the smaller scale of our price variable. Price also has a much larger uncertainty than organic rank with SE(.07).[20]

The interaction effects indicate that the impact of price or relevancy on the predicted probability of user clicks is reduced for products that achieve high visual prominence. To interpret the interaction effect between organic rank and attention share for a unit increase in attention share, the interaction effect coefficient, $\beta_5$, is added to our coefficient for organic rank, $\beta_1$. This adjustment represents the additional impact of attention share on the relationship between organic rank and the dependent variable (probability of being a top-3 most clicked product). For a product with a one (logged) unit increase in attention share, the organic rank coefficient would be equal to $\beta_1 + \beta_5 = -1.76 + 0.21 = -1.55$. For a two unit increase: $-1.76 + 0.21 \times 2 = -1.34$. Therefore, increasing attention share reduces the magnitude of the organic rank coefficient, pushing it closer to 0 and diminishing its influence on the probability of a product being one of the top-3 most clicked.

To better contextualize the impact of these interaction effects, consider a product organically ranked amongst the worst 20% of products by relevancy. This leads to a four-fold decrease in the likelihood of it being amongst the 3 most clicked product on the page, compared to the average product (assuming an average attention share).[21]

As discussed in the previous section, a low-relevancy product could receive a considerable boost in probability of being a top-3 most clicked product from the attention-share coefficient alone, if it was located in a very prominent position. Calculating this boost from visual prominence (the top 5% of positions by attention share), excluding the interaction effect between attention share and organic rank, we see that probability of being a top-3 most clicked product increases to from 0.3% to 2.4%.[22]

Besides this boost purely from visual prominence, the interaction effect between attention share (prominence) and organic rank (relevancy) suggests that the deleterious effect of low relevancy is mitigated for products that are very prominent. Adding the interaction effect increases the predicted top-3 clicked product probability by about 1.5x (from 2.4% to 3.7%).[23] We can understand the interaction ef-

[20] Compared with our baseline regression, coefficients in the extended regression increase in magnitude: our organic rank coefficient increases slightly from $-1.44$ to $-1.76$, our attention share coefficient from 0.58 to 0.77, and our intercept from $-4.20$ to $-4.47$. Though these coefficients are not of core interest here.

[21] We can calculate this as follows: The 20$^{\text{th}}$ percentile organic rank (after logging and mean centering) is 0.75. The average product's chances of being a top-3 most clicked product (as shown above) are $\approx 1.5\%$. Holding attention share constant at 0, our new probability of being a top-3 most clicked product is given by $P(\text{Top3 Most Clicked}) = \frac{1}{1+e^{\beta_0+\beta_1\times 0.75}} = \frac{1}{1+e^{4.47+(-1.76)\times 0.75}} = 0.003$. Thus, the average product is 5 times more likely to be a top-3 most clicked product than one in the bottom 20% of the organic rankings.

[22] Holding price constant at the mean.

[23] The following calculation can be used to calculate the predicted probability of the product being a top-3 most clicked including the interaction effect:

$$P(\text{Top3 Most Clicked}) = \frac{1}{1+e^{\beta_0+\beta_1\times 0.75+\beta_2\times 2.73+\beta_5\times 2.73\times 0.75}} = 0.048$$

or 3.7%. Top 5% of attention share after logging and mean centering is 2.74 and bottom 20% of organic ranks begins at 0.75.

fect between organic rank and attention share as saying that the importance of organic rank (relevancy) as a predictor decreases for products that are more visually prominent.

Alternatively, we can look at the interaction effect the other way around: For each unit increase in (logged) organic rank we add the interaction effect ($\beta_5 = 0.21$) to our attention share coefficient ($\beta_2 = 0.77$), thus the size of the attention share coefficient increases as relevancy decreases. From this perspective we can understand the interaction as saying that as organic rank increases (i.e., as relevancy decreases) the importance of visual prominence as a predictor of clicks increases.

Either way, our $\beta_5 = 0.21$ coefficient suggests that users subject the most prominent products to a significantly less thorough inspection than those products in more visually obscure positions. This affords Amazon the capability to position less relevant, advertised results in prominent locations in order to allocate them demand (clicks) within certain limits of product relevancy.

Similarly our price and attention share interaction effect ($\beta_6 = 0.29$) shows that a product's increased relative visual prominence (attention share) is able to negate the negative impacts of increased product price on click probability. However, since the maximum influence of $\mathsf{Price}_j$ before being offset (for the most expensive product on a search page) is -0.35%, see discussion of price coefficient in previous section, this is not a highly economically significant finding.

Figures 5 and 6 below demonstrate the impact of our attention share and organic rank interaction effect by descriptively showing that organic results rely less on relative prominence and more on relevancy for clicks, whilst advertisements lean more on their attention share. The contour plots reveal how top-3 most clicked products are distributed in our dataset by attention share (x-axis) and organic rank (y-axis). The areas that are more brightly coloured indicate that there is a greater density of products within that range of attention share and organic rank values.

Comparing Figure 6 to Figure 5 one can see that all of the contours in the advertising only figure are stretched to the right towards the higher organic ranks (less relevant products). This stretching shows that amongst advertisements, top-3 most clicked products can be less relevant, indicating that advertising stretches consumers' tolerance for less relevant products.

We also see that screen position is less important to organic products. More precisely, in the advertising only Figure 5 the brightest contours exist almost exclusively between the 99$^{\text{th}}$ and 93$^{\text{rd}}$ percentile of attention share positions (y-axis), whereas in this organic-only contour plot the brightest areas exist between 65$^{\text{th}}$ percentile and the 94$^{\text{th}}$. Once again, this indicates that organic results do not depend as much on their relative visual prominence to receive clicks.

Figure 5. Advertising Increases the Range of Organic Ranks that Users are Willing to Click

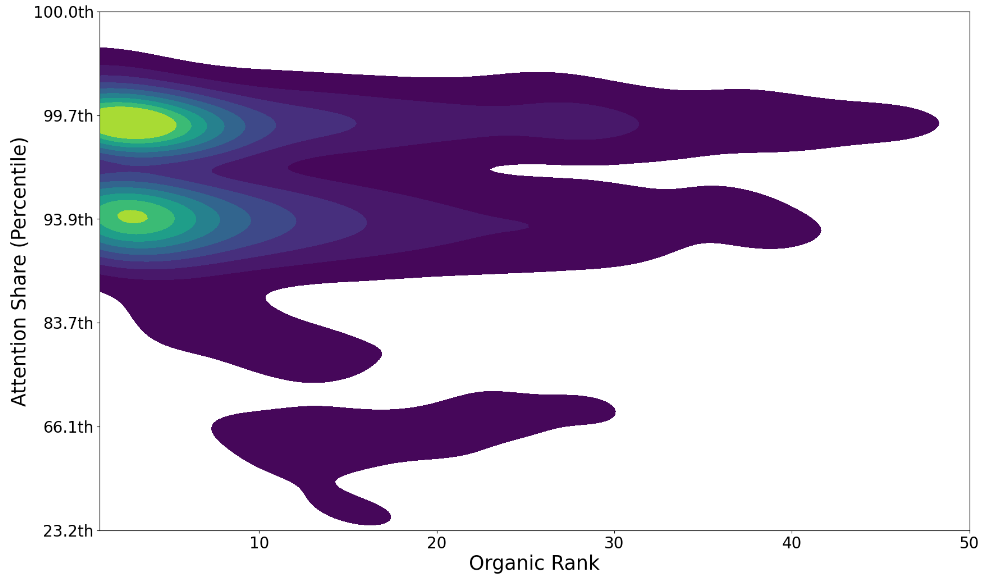


*Note*: Contour plot of top-3 most clicked products that are ads. The brighter the colour of the contour, the more probable it is that top-3 most clicked products will have attention share and organic values within that region of the graph. Note the narrow range of attention share values (top of the y-axis) in which the most brightly coloured areas are located, showing that ads require very favourable positions to receive clicks. The horizontal stretching of the contours along the x-axis, compared with 6, reveals how consumers preference for relevancy is stretched toward high organic rank products by advertising.

Figure 6. Screen Prominence is Less Important for Organic Results

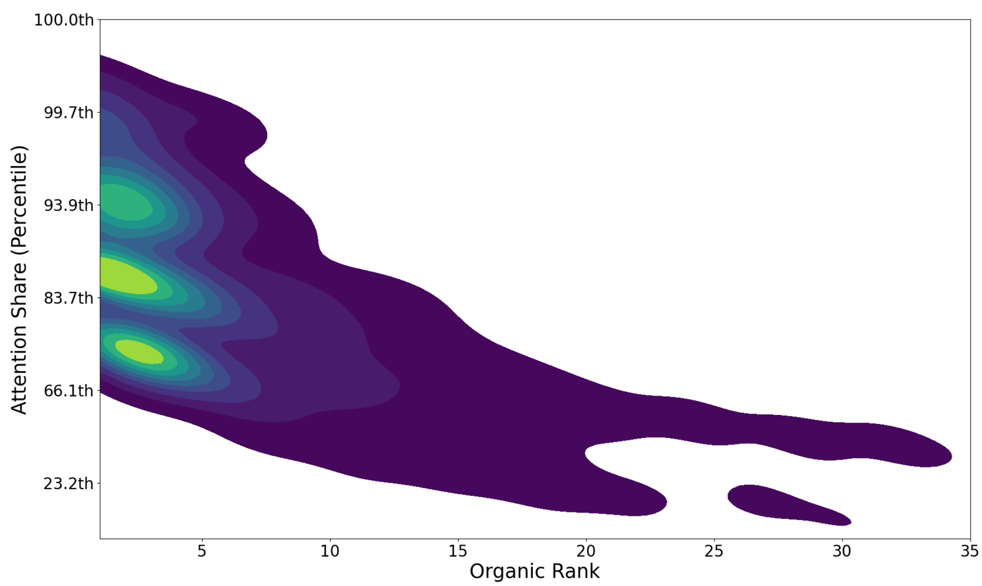


*Note*: Contour Plot of Top-3 most clicked Products that are Organic showing that organic products depend less on screen position for clicks, as the brightest contours exist in a greater vertical range than the advertised top-3 products.

## 4.2 Robustness

To verify the robustness of our model, we used it to predict top-3 most clicked results in the unseen list-format SERP data, as previously described. This experiment aimed to validate the predictive performance of both our model and attention share measure across SERP formats. We found that the predictive performance remained consistent, even improving slightly, for list-format data. The binned residual plot of our model revealed no systemic patterns, with all but two of the points lying between 95% error bounds, encouraging confidence (Gelman et al. 2006).

Fitting a new model to the list format SERP data, we compare the relative sizes and signs of our coefficients between our original model and the list-fitted model. We find that the relative coefficient sizes and signs remain constant (i.e., attention share remained large and positive, smaller than organic rank) with the exception of our price coefficient. We note a slight decrease in our $\mathsf{AttentionShare}_j$ coefficient, a finding which is likely explained by the increased simplicity of the visual information format on list style SERPs (Djamasbi et al. 2013; Guo et al. 2020; Xie et al. 2019). Our price coefficient decreases such that it loses statistical significance altogether in our list-format regression, becoming less than double its standard error. We suggest that may be the result of the much smaller sample size on which this model was fitted, as list-format SERPs generally contain less than half the number of products of their matrix counterparts and are deployed less often by Amazon (around 1/3 of pages were list-format in our scrape).

We also test the robustness of the decay parameter borrowed from Ahn et al. (2018). For our main regression we use a value selected from their findings: -0.809. To test for robustness we run two alternative regressions: The first with a decay value of -0.5 and the second -1.0. The first model ($\lambda = -0.5$) achieves very similar performance to our original ($\lambda = -0.809$), with a slight increase in precision and a slight decrease in recall, maintaining a constant F1-score on the positive class to two decimal places. The second, using a decay coefficient of -1.0 ($\lambda = -1.0$), performs slightly worse than our original model, losing two-hundredths of its F1-score. The marginal performance reduction indicates that our results are not highly sensitive to a narrow range ($\pm 0.3$) of decay coefficient values.

We performed a series of regressions including and excluding several control variables. Notably, adding a dummy representing whether a product has a "Best Seller" badge, a feature which has been found to be highly endogenous with the rank of products in search (Jungle Scout 2023b). In our "Best Seller" regression, we find that the the deviance of our model remains around the same, whilst the $\mathsf{OrganicRank}_j$ coefficient shrinks slightly, likely in light of the noted endogeneity. Removing our price coefficient and interaction term, we find that the deviance of the model only increases by 25, a marginal gain reflecting the slight economic significance reported by this variable. As previously reported, including SERP level product review data[24] does not improve the fit of the model, with these predictors found to be highly endogeneous with organic rank (number of reviews), or statistically insignificant (mean star rating). Further discussion of robustness and verification, including on the endogeneity between organic rank and attention share is located in Section 4.2.

Whilst there is a slight degree of collinearity between attention share and organic rank, demonstrated by a variable inflation factor of 1.7, it is not a significant concern in our regression analysis for the following reasons. Firstly, the prevalence of advertising on Amazon's SERPs, as illustrated in Figure 2, introduces considerable noise into organic rankings. The advertising factor contributes to the moderate VIF that we observe, which suggests that the collinearity is not problematic.

The Pearson's correlation coefficient of -0.34 between organic rank and attention share suggests a moderate statistical relationship but one that lacks practical significance in our regressions(Cohen 1987). This is caused by the highly uneven distribution of attention on Amazon's SERPs, shown in Figure 1, in which majority of user attention (clicks) is directed into products on the first couple of rows, which are highly advertisement dense. This is reflected in the distribution of attention share (see Section 2.3, which is highly skewed. Therefore, the region of attention share where it is significantly informative regarding click-activity shows a substantially weakened correlation with organic rank. This point is accentuated by our logistic regression analysis, highlighted in Figure 3. Here, the influence of attention share on clicks is most pronounced in the top decile of attention-heavy positions. When the dataset is narrowed to this decile, the Pearson coefficient reverses from -0.34 to 0.11. This reduction, even inversion, demonstrates that the correlation is minimized in the critical range of attention share where the majority of clicks occur. In other words, the correlation between organic rank and attention share manifests in page regions with minimal clicks, attention and advertising.

# 5 Conclusion

Using data from Amazon Marketplace on product features and the top-3 most clicked products, we have shown that Amazon has significant market power which it exploits to increase its profits, defined here as an ability to degrade output quality and still receive user clicks. This market power is reliant on user behaviour being strongly influenced by Amazon's algorithmic placement of screen results. This heuristic helps direct user clicks to inferior results.

[24] Mean star rating and number of reviews, not star rating breakdown or review text analysis as this data is located exclusively on product pages.

Despite consumers preferring highly relevant products, our econometric analysis indicates that Amazon is able to exploit users' attention and biases, maintaining a form of market power atypical to traditional (non-digital) monopolistic frameworks. We found that among the top five search results, where typically four are advertisements, neither decreased relevancy nor increased price significantly sways user decisions.

These empirical findings challenge Chicago School type legal arguments that platforms cannot extract rents from their ecosystems because users can costlessly optimize between platforms and search results, such that information quality can never harm (H. Hovenkamp 2023; Strauss et al. 2023). Instead, advertising on Amazon transforms product prominence into a mechanism for rent-extraction and shows that rents in digital markets can persist (O'Reilly et al. 2023; Mazzucato, Ryan-Collins, et al. 2023). This highlights the importance of not just user data, but a platform's algorithmic allocations of user attention to more profitable information outputs, as a way for it to exploit its ecosystem (Calvano et al. 2021).

External assessments of the process of value creation and allocation by dominant platforms are complicated by algorithmic opacity and marketplace dynamism. Thus, the findings in this paper reinforce our previous recommendations, including that technology companies should be required to disclose a "monetization narrative" regarding how exactly attention is monetized on their platforms (Mazzucato, Strauss, et al. 2023). Furthermore, regulating these currently intangible assets entails quantification: companies should be required to disclose their internally-utilized operating metrics (Strauss et al. 2023; O'Reilly et al. 2023). Access to this information would clarify to regulators the exact mechanisms that link the companies' revenues and costs to its free-to-use products. This would enable regulators to make informed decisions when delineating acceptable levels of algorithmic power.

Further research might aim to refine the measurement of attention share and incorporate a wider array of behavioural factors to further understand the interplay between algorithmic influence and consumer choice (Sorokina et al. 2016). Future work might also aim to extend this type of study temporally, capturing time-series data to better understand the evolution of Amazon and other platforms across the attention economy, including how user click behaviour dynamically adapts to algorithmic choices. Lastly, our simple (implicit) decision theoretic framework could in the future be extended to account for user clicks or purchasing stopping in response to more advertising (Sun et al. 2023). However, such data is very difficult to obtain or infer.

**Declaration of Generative AI and AI-assisted technologies in the writing process**: During the preparation of this work the author(s) used ChatGPT 4.0 to assist with the stylistic rewriting, grammatical, and spell checking of several sentences. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the contents of the publication.

# References


Ahn, Jae-Hyeon et al. (2018). “Attention Adjustment, Renewal, and Equilibrium Seeking in Online Search: An Eye-Tracking Approach.” *Journal of Management Information Systems* 35.4, pp. 1218–1250.

Alvarez, George A and Patrick Cavanagh (2004). “The Capacity of Visual Short-Term Memory is Set Both by Visual Information Load and by Number of Objects.” *Psychological science* 15.2, pp. 106–111.

Anstis, Stuart M (1974). “Chart Demonstrating Variations in Acuity with Retinal Position.” *Vision Research* 14.7, pp. 589–592.

Areeda, Phillip and Herbert Hovenkamp (2023). *Fundamentals of Antitrust Law*. Wolters Kluwer.

Baye, Michael R and Jeffrey Prince (2020). “The Economics of Digital Platforms: A guide for regulators.” *The Global Antitrust Institute Report on the Digital Economy* 34.

Berti, Stefan and Erich Schröger (2003). “Working Memory Controls Involuntary Attention Switching: Evidence from an auditory distraction paradigm.” *European Journal of Neuroscience* 17.5, pp. 1119–1122.

Calvano, Emilio and Michele Polo (2021). “Market Power, Competition and Innovation in Digital Markets: A survey.” *Information Economics and Policy* 54.

Cohen, Jacob (1987). *Statistical Power Analysis for the Behavioral Sciences*. Laurence Erlbaum Associates: Hillsdale, NJ, USA.

Djamasbi, Soussan, Adrienne Hall-Phillips, and Ruijiao Yang (2013). “Search Results Pages and Competition for Attention Theory: An exploratory eye-tracking study.” *Human Interface and the Management of Information. Information and Interaction Design*. Ed. by Sakae Yamamoto.

Etro, Federico (2022). “The Economics of Amazon.” *Available at SSRN 4307213*. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4307213.

Farronato, Chiara, Andrey Fradkin, and Alexander MacKay (2023). “Self-Preferencing at Amazon: evidence from search rankings.” *AEA Papers and Proceedings*.

Fessenden, Therese (2018). “Scrolling and Attention.” *Nielsen Norman Group* 15.4. URL: https://www.nngroup.com/articles/scrolling-and-attention/.

Floh, Arne and Maria Madlberger (2013). “The Role of Atmospheric Cues in Online Impulse-Buying Behavior.” *Electronic Commerce Research and Applications* 12.6, pp. 425–439.

Fowler, Geoffrey (Nov. 2022). *It’s not your imagination: Shopping on Amazon has gotten worse*. URL: https://www.washingtonpost.com/technology/interactive/2022/amazon-shopping-ads/.

FTC (2023). “FTC complaint Against Amazon.” Tech. rep. Filed 09/26/23. CASE NO.: 2:23-cv-01495. URL: https://www.ftc.gov/legal-library/browse/cases-proceedings/1910129-1910130-amazoncom-inc-amazon-ecommerce.

Fukuda, Keisuke, Edward Awh, and Edward K Vogel (2010). “Discrete Capacity Limits in Visual Working Memory.” *Current Opinion in Neurobiology* 20.2, pp. 177–182.

Gelman, Andrew and Jennifer Hill (2006). *Data Analysis Using Regression and Multilevel/Hierarchical Models*. Cambridge University Press.

Google (Nov. 2014). “The Importance of Being Seen.” *Think with Google*. URL: https://think.storage.googleapis.com/docs/the-importance-of-being-seen_study.pdf.

Guo, Ruocheng et al. (2020). “Debiasing Grid-based Product Search in E-commerce.” *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 2852–2860.

Hong, Weiyin, James Thong, and Kar Tam (Jan. 2004). “The Effects of Information Format and Shopping Task on Consumers’ Online Shopping Behavior: A cognitive fit perspective.” *Journal of Management Information Systems* 21, pp. 149–184.

Hovenkamp, Erik (2018). “Platform Antitrust.” *Journal of Corporation Law* 44, p. 713.

Hovenkamp, Herbert (2023). "Gatekeeper Competition Policy." *University of Penn, Institute for Law & Economics Research Paper* 23-08. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4347768#:~:text=The%20%E2%80%9CGatekeeper%E2%80%9D%20approach%20to%20competition,Congress%20would%20have%20done%20this..

Janiszewski, Chris (1998). "The Influence of Display Characteristics on Visual Exploratory Search Behavior." *Journal of Consumer Research* 25.3, pp. 290–301.

Jeffries, Adrianne and Leon Yin (Oct. 2021). *Amazon Puts its Own "Brands" First Above Better-Rated Products*. URL: https://themarkup.org/amazons-advantage/2021/10/14/amazon-puts-its-own-brands-first-above-better-rated-products.

Jin, Ginger Zhe and Liad Wagman (2021). "Big Data at the Crossroads of Antitrust and Consumer Protection." *Information Economics and Policy* 54, p. 100865.

Joachims, Thorsten, Adith Swaminathan, and Tobias Schnabel (2017). "Unbiased Learning-to-Rank with Biased Feedback." *Proceedings of the tenth ACM international conference on web search and data mining*, pp. 781–789.

Jungle Scout (2022). "The State of the Amazon Seller 2022." URL: https://www.junglescout.com/amazon-seller-report/.

— (2023a). "Amazon Advertising Report." URL: https://www.junglescout.com/amazon-advertising-report/.

— (2023b). "Amazon Best Sellers Rank (BSR): Everything you need to know in 2023." URL: https://www.junglescout.com/blog/amazon-best-sellers-rank/.

Kahneman, Daniel (2011). *Thinking, Fast and Slow*. Macmillan.

Lépine, Raphaëlle, Sophie Bernardin, and Pierre Barrouillet (2005). "Attention Switching and Working Memory Spans." *European Journal of Cognitive Psychology* 17.3, pp. 329–345.

Lorigo, Lori et al. (May 2008). "Eye Tracking and Online Search: Lessons learned and challenges ahead." 59, pp. 1041–1052.

Luck, Steven J and Edward K Vogel (1997). "The Capacity of Visual Working Memory for Features and Conjunctions." *Nature* 390.6657, pp. 279–281.

Mazzucato, Mariana, Josh Ryan-Collins, and Giorgos Gouzoulis (2023). "Mapping Modern Economic Rents: The good, the bad, and the grey areas." *Cambridge Journal of Economics* 47.3, pp. 507–534.

Mazzucato, Mariana, Ilan Strauss, et al. (2023). "Regulating Big Tech: The role of enhanced disclosures." *Oxford Review of Economic Policy* 39.1, pp. 47–69.

Meyerson, Mark (Oct. 2023). *5 Reasons Amazon Ads is better than Google Ads for Ecommerce*. URL: https://searchengineland.com/amazon-ads-vs-google-ads-ecommerce-433186.

Newman, John M (2015). "Antitrust in Zero-Price Markets: Foundations." *University of Pennsylvania Law Review*, pp. 149–206.

O'Reilly, Tim, Ilan Strauss, and Mariana Mazzucato (2023). *Algorithmic Attention Rents: A theory of digital platform market power*. UCL IIPP Working Paper 2023-10. URL: https://www.ucl.ac.uk/bartlett/public-purpose/sites/bartlett_public_purpose/files/oreilly_t._strauss_i._mazzucato_m._2023._algorithmic_attention_rents-_a_theory_of_digital_platform_market_power.pdf.

Park, Sangwon and Juan L Nicolau (2015). "Asymmetric Effects of Online Consumer Reviews." *Annals of Tourism Research* 50, pp. 67–83.

Pernice, Kara (2017). "F-shaped Pattern of Reading on the Web: Misunderstood, but still relevant (even on mobile)." *Nielsen Norman Group*. URL: https://www.nngroup.com/articles/f-shaped-pattern-reading-web-content/.

Reutskaja, Elena et al. (Apr. 2011). "Search Dynamics in Consumer Choice under Time Pressure: An eye-tracking study." *American Economic Review* 101.2, pp. 900–926.

Rochet, Jean-Charles and Jean Tirole (2003). "Platform Competition in Two-Sided Markets." *Journal of the European Economic Association* 1.4, pp. 990–1029.

Roth, Alvin E (2015). *Who Gets What – And Why: The new economics of matchmaking and market design*. Houghton Mifflin Harcourt.
Shi, Savannah Wei and Michael Trusov (2021). "The Path to Click: Are you on it?" *Marketing Science* 40.2, pp. 344–365.
Shifflett, Shane, Alexandra Berzon, and Dana Mattioli (Dec. 2019). *"Amazon's choice" isn't the endorsement it appears*. URL: https://www.wsj.com/articles/amazons-choice-isnt-the-endorsement-it-appears-11577035151.
Shrestha, Sav et al. (2007). ""F"' Pattern Scanning of Text and Images in Web Pages." *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*. Vol. 51. 18, pp. 1200–1204.
Simon, Herbert A (1956). "Rational Choice and the Structure of the Environment." 63.2, pp. 129–138.
Simon, Herbert A et al. (1971). "Designing Organizations for an Information-Rich World." *Computers, communications, and the public interest* 72, p. 37.
Simon, Herbert A (1977). "The Logic of Heuristic Decision Making." *Models of Discovery: and Other Topics in the Methods of Science*, pp. 154–175.
— (1978). "Rationality as Process and as Product of Thought." *The American economic review* 68.2, pp. 1–16.
— (1998). "Information 101: It's not what you know, it's how you know it." *The Journal for Quality and Participation* 21.4, p. 30.
— (2013). *Administrative Behavior*. Simon and Schuster.
Sorokina, Daria and Erick Cantu-Paz (2016). "Amazon Search: The joy of ranking products." *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*, pp. 459–460.
Strauss, Ilan, Tim O'Reilly, and Mariana Mazzucato (2023). *Amazon's Algorithmic Rents: The economic of information on Amazon*. UCL IIPP Working Paper 2023-12. URL: https://www.ucl.ac.uk/bartlett/public-purpose/wp2023-12.
Sun, Tianshu et al. (2023). "The Value of Personal Data in Internet Commerce: A High-Stakes Field Experiment on Data Regulation Policy." *Management Science*.
Sundin, Emma (2021). "Nudging and Design Friction: The impact on our decision making process." *Conference in Interaction Technology and Design*, p. 39.
Turton, Chris (2023). "Does Amazon Ads Impact On Amazon Organic SEO/Ranking?" *Ecommerce Intelligence*. URL: https://www.christurtonecommerce.com/does-amazon-ads-impact-on-amazon-organic-seo/.
UK CMA (2022). "Research and Analysis Evidence Review of Online Choice Architecture and Consumer and Competition Harm." URL: https://www.gov.uk/government/publications/online-choice-architecture-how-digital-design-can-harm-competition-and-consumers/evidence-review-of-online-choice-architecture-and-consumer-and-competition-harm.
Vogel, Edward K, Geoffrey F Woodman, and Steven J Luck (2001). "Storage of Features, Conjunctions, and Objects in Visual Working Memory." *Journal of experimental psychology: human perception and performance* 27.1, p. 92.
Wang, Xuanhui et al. (2018). "Position Bias Estimation for Unbiased Learning to Rank in Personal Search." *Proceedings of the eleventh ACM international conference on web search and data mining*, pp. 610–618.
Xie, Xiaohui et al. (2019). "Grid-based Evaluation Metrics for Web Image Search." *The World Wide Web Conference*, pp. 2103–2114.
Zhang, Kem ZK et al. (2018). "Online Reviews and Impulse Buying Behavior: The role of browsing and impulsiveness." *Internet Research* 28.3, pp. 522–543.
Zhang, Weiwei and Steven J Luck (2009). "Sudden Death and Gradual Decay in Visual Working Memory." *Psychological Science* 20.4, pp. 423–428.

# Appendices

For online publication.

## A Data Appendix

### A.1 Data Collection Technical Notes & Limitations

Our scrape was performed on a single computer and routed through a virtual private network (VPN) using rotating IP addresses. This approach served two purposes, firstly, helping to anonymise our traffic; secondly, geolocating us within the United States, as this study exclusively investigates Amazon's US Marketplace. We chose Amazon.com specifically because it is the single largest advertising market amongst Amazon's service regions (Jungle Scout 2023a). A potential drawback of utilizing VPN is the possibility that the anonymization of user traffic alters the delivery of search results and advertising due to personalization algorithms.

### A.2 Duplication & Churn

We found a high level of product duplication within SERPs. "Duplicated products" are products which have exactly matching "Amazon Standard Identification Numbers" (ASINs) that appear more than once in a single search results page. Having identical ASINs entails that it is an exact product match (same colour, skew, etc).

Amongst these duplicated products, we observed that the the most common form of duplication were pairs of products consisting of one appearance of the product as an organic result and one sponsored. We observed that around 40% of all ads on the first page of Amazon search results had an organic clone on the same page. However, this was not the only form of duplication, with as many as 4 duplicate appearances of one product on a single page as well as varying numbers of organic and sponsored appearances.

It is possible that Amazon opts to show duplicated listings in situations where they have a limited subset of highly relevant products, choosing to repeat those which are highly relevant rather than show the user less relevant products. However, the pervasiveness of duplication across search terms (at least one duplicate is observed in 94% of the search terms we scraped) indicates that product duplication is likely to be more systemic, as one might expect situations in which there is a highly limited pool of very relevant products to be the minority, given the scope Amazon's third party Marketplace.

In our scraped Amazon SERPs, we found a high level of product "churn" amongst the reported top-3 most clicked products. That is, we found that of the 19,518 products reported by Amazon to be amongst the top-3 most clicked, 7752 were not present in our scrape of their respective search terms. This degree of churn is remarkable, as nearly 40% of the products reported to be the best-performing amongst our search terms did not appear in the first page of search results.

The causes of this churn are uncertain, it is possible that, since our traffic was anonymized, Amazon considered the chances of purchasing a product in our searches as relatively low, thus altering the results that were provided through advertising. Amazon maintains the ability to do this through their "dynamic bidding" system, a mechanism which automatically increases or decreases an advertiser's bid price (the cost they are willing to incur for a click on their product) depending on Amazon's algorithmic conversion chance estimation.

Alternatively, churn may be caused simply by popular products being more likely to become out of stock. However, we observe that Amazon often leaves out of stock products within search results, adding an "Out-of-stock" label in place of their price. This demonstrates to searchers that Amazon does carry the product, and that it will be available in the future, rather than leaving them out of search and possibly leaving users with the impression that they are not available on Amazon at all. It is important to note that our churn evaluation took place before filtering out "Out-of-stock" marked products (as they do not have price information). Results pages without complete top-3 information were excluded from regressions.

### A.3 Attention Competition

Competition-for-attention theory asserts that all items within one's visual field are engaged in a constant "tug-of-war" for a consumer's attention (Janiszewski 1998). This theory relies on the neuroscientific idea, from Helmholtz, that "seeing", for the visual system, is equivalent to "recognizing differences in the intensity of signals generated by the boundaries of objects located in the visual field" (ibid.). Visual receptor cells are responsible for measuring the strength of these signals within the visual field and are densest at central point of the retina, the fovea. Items generating larger signals would compete more strongly for a user's attention.

The strength of the signal generated by an item can be influenced by several factors. These factors include (i) its distance from the object immediately in focus; and (ii) its size.

Firstly, as a given item is moved closer to the object that is currently in focus, its projection on the retina moves closer to the fovea. Therefore, the image of this competing object on the retina will be projected onto a greater density of visual receptor cells. This increases the strength of its signal to the perceptual system. That item can be considered to exert a greater 'tug,' or compete more strongly, for the viewer's attention. Secondly, as an item in the visual field becomes larger, its projection on the retina increases in area and thus is dispersed over a greater number of visual receptor cells. The item generates a larger signal and thus competes more strongly for a user's attention. Notably, in order to maintain a constant signal magnitude, an object which moves one degree further from the current object in focus must increase in size by 0.2 degrees, a linear relationship demonstrated by Anstis' infamous chart of Visual Acuity (Anstis 1974).

According to Janiszewski (1998), we can quantify the amount of competition that a given item in the visual field, Item X, faces for viewers' attention. We can do so by calculating and summing the signal strength of each of Item X's competitors in the visual field - in other words, the other items viewable simultaneously to Item X. The signal strength of each competitor can be calculated by taking the square root of its area and dividing this by its distance to Item X (ibid.). The sum of these competitors' signal strengths is equivalent to the amount of competition for attention faced by Item X.

Attention competition theory can be applied to Amazon SERPs to illustrate the influence of screen position on a search result's ability to capture consumers' attention. Products belonging to a Section of the results page that has (i) more competing products, (ii) competing products which are nearby and/or (iii) larger competing products, thus faces greater competition for the users' attention.

### A.4 Attention Decay

Attention decay theory is based on a neuroscientific understanding of human visual working memory (W. Zhang et al. 2009; Alvarez et al. 2004). It rests on the now widely accepted contention that during eye movements from one point to another, previously observed contents are not instantly forgotten by temporarily retained in working memory (Ahn et al. 2018; W. Zhang et al. 2009). Numerous studies, (Luck et al. 1997; Vogel et al. 2001; Fukuda et al. 2010), have reported that the upper limit of this visual working memory buffer contains around 6 items of visual information. Decay theory hypothesizes that as our remaining visual working memory capacity decreases, as we scroll through search results, processing new results requires more attentional resources. Decay theory models how users' cognitive capacity is diminished as they process information sequentially presented in a list.

In addition to modelling the rate of attention decay, the theory posits "attention renewal". Attention renewal theory arises from experimentation on attention switching (Berti et al. 2003; Lépine et al. 2005), which suggests that "interruptions" in the presentation of visual information can, in-part, refresh the visual working memory buffer. These "interruptions" in our case (Amazon search results pages) are defined as either large non-result HTML elements (banner or video advertisements) or changes in the format of the search results (such as a switch from a vertically scrolling list to a side scrolling carousel). Renewal theory hypothesizes that after one of these interruptions, there is a brief period of almost completely renewed attention, therefore, results placed immediately following an interruption are likely to receive more attention than those immediately prior. More technically, such disturbances to otherwise visually monotonous search results may "interrupt actions such as information search [and] lead to involuntary attention switching, subsequently, reversing attention-decaying behaviours" (Ahn et al. 2018).

Therefore, in order to calculate the decay after renewal, we capture the positions of banner advertisements, video advertisements and changes in the format of the search results. Our decay parameter, $\lambda_1$ is lifted from Ahn et al. (ibid.)'s paper on attention decay in online search environments. We test the robustness of this selection by altering the coefficients we use by $\pm(0.3)$ in Section 4.2.

One limitation of our approach is the unchanged attention decay coefficient post-advertising, despite Ahn et al. (ibid.) finding an increased rate of decay post users being exposed to advertising (attention renewal). The complexity of Amazon's SERPs (e.g., having numerous visual interruptions) limits our ability to apply these findings in our context. Future research could explore integrating increases in rates of decay in order to further refine the attention share measure.

### A.5 Empirical Validation of Attention Share

To justify the inclusion of attention share empirically over other measures of positional bias, such as the search result rank of a product or the distance to the top left hand corner of the screen we performed a series of three simplified regressions. These regressions were simplified out of necessity because of the endogeneity exhibited by our competing measures with our covariates. We include just the postional bias variable, our organic rank variable, and the Amazon's Choice dummy variable (as discussed in Section 2.2). In Table 5 we find that the model performance and fit are best with the attention share metric, although this gains are small.

Table 5. Model Performance with Different Measures of Visual Prominence

| Positional Bias Metric | F1-Score (Model Performance) | Deviance (Model Fit) | VIF |
|---|---|---|---|
| Attention Share | 0.48 | -74,282 | 1.70 |
| Distance to Top Left Corner | 0.47 | -75,848 | 2.92 |
| Search Result Rank | 0.46 | -76,106 | 3.05 |

*Note*: Empirical Justification For Attention Share vs Other Measures of Positional Bias. We see that the model demonstrating the best predictive performance used attention share, however, the observed increase in F1-score was relatively marginal. We also see that the model including attention share had the smallest deviance, and therefore the best fit of the 3 predictors. However, it was once again a relatively marginal advantage. These findings might be explained by comparing VIF scores, indicating that "attention share" was the least closely related predictor to the other variables test.

To explain model performance and fit findings in Table 5, we computed the variable inflation factor (VIF) for each of the models. VIF is a measure of multicollinearity between variables, where collinearity describes a situation in which two or more variables in the model are closely related, leading to distortion in the final model. Whilst it should be noted that none of our observed VIF scores were above the threshold for major concern (10), the relative increases may explain the reduced fit and predictive performance of the other models.

The observed VIF results alludes to the most important factor in favor of the inclusion of attention share - interaction effects. Whilst the previous analyses indicated only slight improvements in predictive power when considering attention share, those made possible through interaction effects were robust and significant as demonstrated in 4.1. Our other measures of positional bias contained a degree of multicollinearity which caused equivalent interaction effects to be unstable and/or not statistically significant hence our use of a simplified model for testing. We are able to achieve the best overall predictive performance and fit using attention share as our measure of position-bias.

## A.6 Choice of F1-Score for Model Evaluation

This metric captures the harmonic mean of the precision and recall of the model when predicting products that are amongst the top-3 most clicked. Precision quantifies the proportion of accurate predictions our model makes from all the products it identifies as within the top-3 most clicked product category. In contrast, recall measures how well the model detects the actual top-3 clicked products from the test dataset. Our choice of the F1-score relates to our dataset being intrinsically unbalanced - top-3 most clicked products are, by definition, outnumbered by the remaining products on the page. In such situations, traditional measures of predictive performance, such as accuracy, can be misleading. The null model, which only ever predicts the majority class, will achieve above 92% accuracy given our dataset's imbalance. In contrast, F1-score emphasises the performance of the model specifically on the crucial minority class.

## A.7 Model Evaluation

The error rate of a logistic regression model is often used to determine the predictive performance of the model and is defined as the proportion of cases for which the prediction made are wrong. The threshold value for prediction classification is usually set to 0.5 (Gelman et al. 2006), thus if the model's predicted probability of a product being in the top-3 most clicked class is $>= 0.5$ we classify the prediction as 1, and if the predicted probability $< 0.5$ we classify the prediction as 0. However, in our case, due to the major class imbalance present in our dataset we adjust this threshold down to a value of 0.337 which we find to optimize our F1-score on the positive class, the advantages of which are explained in Section A.6.

Our model has an error rate of 5.67%, this is compared to the null model, which only predicts the majority class, that has an error rate of 7.98%. Whilst, this may seem like a marginal increase in performance, this is, again, due to our class imbalance. In fact, this performance increase takes us from a null-model, which is incapable of correctly predicting the status of any products which are top-3 most clicked products, to our model which is capable of predicting 64% of all top-3 most clicked products correctly in our test data. Whilst this model performance is not radically impressive, it is indicative of a model which performs adequately for the discussion which follows, especially near the extremes of organic rank and attention share. Finally, the binned residual plot of our model displayed no systemic pattern with all but two of the residuals lying within the error bounds, inspiring confidence in the design and composition of our model.

UCL Institute for Innovation and Public Purpose
11 Montague Street, London, WC1B 5BP

@IIPP_UCL
ucl.ac.uk/IIPP