

Antitrust, Amazon, and Algorithmic Auditing

by

Abhisek Dash, Abhijnan Chakraborty, Saptarshi Ghosh, Animesh Mukherjee,
Jens Frankenreiter, Stefan Bechtold, and Krishna P. Gummadi*

In digital markets, antitrust law and special regulations aim to ensure that markets remain competitive despite the dominating role that digital platforms play today in everyone's life. Unlike traditional markets, market participant behavior is easily observable in these markets. We present a series of empirical investigations into the extent to which Amazon engages in practices that are typically described as self-preferencing. We discuss how the computer science tools used in this paper can be used in a regulatory environment that is based on algorithmic auditing and requires regulating digital markets at scale.

Keywords: Amazon, antitrust, self-preferencing, Digital Markets Act, regulating at scale

JEL classification code: K21, L40, L42

1 Introduction

In digital markets, a small number of digital platforms often mediate transactions between a large number of firms on the one side and millions of consumers on the other side. The platforms typically benefit from network effects that enable them to navigate in two-sided markets in a profit-maximizing way. Over the last years, policy makers around the world have discussed the extent to which such platforms should be regulated in order to preserve competition in digital markets, maintain a free discourse on the Internet, protect the privacy interests of consumers, and secure the intellectual property rights of authors and inventors.

* Abhisek Dash and Krishna P. Gummadi (corresponding authors): Max Planck Institute for Software Systems, Saarbrücken, Germany; Abhijnan Chakraborty: Indian Institute of Technology Delhi, India; Saptarshi Ghosh and Animesh Mukherjee: Indian Institute of Technology Kharagpur, India; Jens Frankenreiter: Washington University in St. Louis, USA; Stefan Bechtold: ETH Zurich, Switzerland. This research is supported in part by a European Research Council (ERC) Advanced Grant for the project "Foundations for Fair Social Computing," funded under the European Union's Horizon 2020 Framework Program (grant agreement 789373), and by a grant from the Max Planck Society through a Max Planck Partner Group at IIT Kharagpur. A. Dash was supported by a fellowship from Tata Consultancy Services for the major duration of this research. Excellent research assistance by Emma Deering and Simona Ramensperger is gratefully acknowledged.

Regulating digital platforms can be challenging. One reason is the fact that the Internet has torn down national borders and facilitated communication and trade across the globe. While, in principle, the transnational scope of platform activities seems to require a global response, international coordination mechanisms have often proven ineffective. Against this background, individual countries and regions have increasingly enacted legal regimes to govern the digital economy, even though such regulatory activity raises challenging questions about regulatory “spillovers” (Bradford, 2012, 2020; Frankenreiter, 2022; Peukert et al., 2022; Davis and Marotta-Wurgler, 2024).

Maybe even more importantly, regulators face a problem of scale. Digital platforms facilitate billions of transactions between firms and consumers. Google, for example, processes 1 billion searches and 1 billion advertisement auctions every day; Facebook connects over 3 billion people every month (Statista, 2023a); and users stream almost 1 billion hours of video on YouTube every day (CloudIncome, 2023). There are about 3.5 million mobile apps on the Google Play store, about 1.6 million apps on the Apple App store, and about 200 million active websites (Tambe and Jain, 2024; Statista, 2023b). In this environment, keeping track of the actions of digital platforms can be challenging. While digitization has led to a drastic decrease in transaction costs and a drastic increase in scale efficiencies for companies, public authorities have not benefited from similar efficiency gains when enforcing digital regulations against these companies.

Besides, various characteristic features of digital platforms pose particular challenges for regulators endeavoring to prevent harmful business practices. Platforms do not only establish the rules governing user interactions. They also design the underlying architecture facilitating these interactions. As Lawrence Lessig famously proclaimed about 25 years ago: “Code Is Law” (Lessig, 1999, chapter 3). On digital platforms, the “code” governing the platforms is written by private companies, not by the legislator. If a regulator wants to influence digital platforms, it has to deal with these privately created “codes” and often resort to indirect ways of regulation.

From the perspective of individual users, digital platforms serve as both gateways to information and architects of their interaction spaces. Digital platforms wield considerable influence over what is commonly termed *choice architectures* (Thaler and Sunstein, 2008), the environments in which users navigate decision-making processes. Extensive research in psychology and related fields has demonstrated that the ability to manipulate choice architectures can profoundly shape user behavior and outcomes (Thaler and Sunstein, 2008). Digital platforms may use technical choice architectures expressed in algorithms to shape consumer preferences, search behavior, and purchase patterns.

These features make it challenging for regulators to tame digital platforms. For one, designing choice architectures lies at the heart of platforms’ core functions. Choice architectures exist primarily to facilitate interactions among platform users, and a platform’s commercial success often hinges on the effectiveness of these choice architectures in engaging and retaining users. Regulators must therefore carefully consider interventions in platform design and operation, as heavy-handed

approaches could impede innovation and hinder platforms' agility in responding to shifting user preferences and market dynamics.

For another, assessing whether a platform utilizes a specific choice architecture to advance its own interests in improper ways poses a considerable challenge in its own right, a topic we explore in detail throughout this article. As we will demonstrate, comparing user interactions within existing choice architectures to hypothetical "neutral" alternatives is often impractical, as such alternatives might simply not exist. In the context of digital marketplaces, for instance, what is the benchmark against which to assess the platform's decision to include, among the information displayed to consumers, the total number of transactions a seller has completed – a feature that may favor larger sellers, particularly those associated with the platform provider? And even when a comparison between an existing choice architecture and a "neutral" alternative is feasible, distinguishing between design features serving legitimate business objectives and those potentially serving improper goals is often challenging. For example, it is not always obvious whether a digital marketplace's choice architecture that appears to nudge customers towards products affiliated with the platform provider is mainly the result of a discrimination of independent competitors or whether it is a response to preferences among consumers for such products, thereby enhancing the overall customer experience.

A significant challenge in this context is that users' preferences are not fixed. Rather, they can be shaped by and change in response to the choice architectures they encounter. This fluidity complicates the task of ascertaining consumers' true preferences. As a consequence of this endogeneity, assessing whether a design feature aligns with these evolving preferences can be elusive.

A related aspect is the pervasive role of algorithms in mediating interactions on online platforms. These algorithms dictate which information is presented to users and can significantly influence their action space at any given moment. Digital platforms employ a variety of algorithmic systems to facilitate interactions among their diverse user groups. Posts visible to users on social media platforms, the price offered to drivers accepting rides on a ride-sharing platform, or the products displayed to shoppers on Amazon are all determined by such complex computer programs. Usually, digital platforms keep the precise nature of the factors that influence the algorithms' decisions, as well as the way they are combined, a closely guarded secret. This opacity renders it challenging for regulators to understand the factors that shape user experiences and behaviors on these platforms, hindering their ability to effectively assess and address potential risks or harms.

Despite these challenges, the digital nature of the transactions mediated by digital platforms also offers new avenues for regulatory enforcement. Compared to traditional markets, market participant behavior is often easier to observe in these markets. Many digital marketplaces can be accessed by everyone, and legislators are increasingly establishing data access rights for regulators, researchers, and the wider public that offer opportunities to generate insights into the business models of these platforms.

That said, obtaining information about business models is usually not trivial. As a first step, a regulator or researcher seeking to assess the operations of a digital platform with the help of digital tools would typically try to access the platform through a regular web interface or engage in automated scraping of the platform's content. However, interfaces of digital platforms have become more interactive, oftentimes rendering simple web access insufficient to capture the inner workings of a digital platform. Digital platforms and markets enabled by these platforms are growing more complex, making it harder to gain a proper understanding of platforms' business practices.

In response, empirically oriented computer science research has developed various approaches to audit digital platforms over the last few years. Sandvig et al. (2014, p. 16) distinguish between "code audit[s]" (where researchers directly inspect a copy of the relevant algorithm), "noninvasive user audit[s]" (where users interacting with the platform answer questions about what they did online), "scraping audit[s]" (where researchers query a platform repeatedly, observing results), "sock puppet audit[s]" (where researchers use computer programs to impersonate users), and "crowdsourced audit[s]" (where researchers employ users as testers). This development points towards a future of algorithmic auditing, where potential adverse effects of algorithmic systems are investigated and evaluated using algorithmic auditing tools (Sandvig et al., 2014; Bandy, 2021; Metaxa et al., 2021). "The Answer to the Machine is in the Machine," as Charles Clark pronounced in the 1990s (Clark, 1996): As digital platforms are increasingly relying of algorithmic systems to govern their digital platforms, regulators will have to resort increasingly to algorithmic systems to audit compliance on these platforms.

This article explores some of the opportunities and challenges of algorithmic enforcement, using Amazon as an example. Recently, Amazon's business policies and deployed algorithmic systems have come under intense scrutiny due to Amazon's vertical integration, or in other words, its involvement in multiple stages of the production and distribution process for its products and services. In this article, we present a series of empirical investigations into the extent to which Amazon engages in practices that are typically described as *preferential treatment* or *self-preferencing*. In these investigations, we analyze different features of various choice architectures that Amazon presents to its customers. We investigate how the service decides the offer that is featured in its Buy Box (now called "Featured Offer"); the design feature of its page listing various offers for the same product; how it selects products offered to customers that use the conversational product search technology available through its virtual assistant technology Alexa; and how it displays results in its related item recommendation system.

The purpose of this analysis is not to determine whether all or some of the practices we observe constitute violations of antitrust laws or specialized regimes such as the European Union's Digital Markets Act. Conducting a full legal analysis of Amazon's business practices is beyond the scope of this paper. Our aim is to demonstrate how empirical methods originating in computer science can empower policy makers, regulators, and researchers to track digital platform behavior at a

granular level, thereby generating empirical evidence that could contribute to discussions on how to effectively regulate digital markets. Concurrently, our discussion underscores how various features of digital platforms can pose challenges for researchers and regulators in deducing anti-competitive intent or impact from observed patterns of platform behavior. Specifically, the case studies presented highlight the considerable challenges involved in assessing whether a platform utilizes a specific choice architecture to advance its own interests in improper ways.

Our article proceeds as follows. Section 2 introduces the design of Amazon’s digital marketplace and several choice architectures governing this marketplace. Section 3 provides an overview of how Amazon’s vertical integration relates to antitrust and digital market regulation doctrines in the European Union and the United States. Section 4 presents our empirical observations on four distinct types of choice architectures: Buy Box, listing pages, Alexa search, and recommendation systems. Section 5 discusses our findings. Section 6 concludes.

2 Amazon and its Marketplace

Over the last years, Amazon has emerged as the largest online e-commerce platform by connecting millions of sellers across the world to billions of customers (Saqib, 2023). On the Amazon platform, the different stakeholders can be primarily classified into three categories: (a) sellers on the platform – they rely on Amazon to sell their products to consumers; (b) customers of the platform – they rely on Amazon for fulfilling their purchase needs; and (c) the platform organization itself – it mediates the interaction between the sellers and customer.

Due to the enormity and scale at which Amazon operates, Amazon can facilitate these interactions only by relying on mass-scale automation and on countless deliberate design choices that determine how sellers can present their products to customers, and how customers express their preferences in the face of these results. These design choices result in choice architectures (Johnson et al., 2012; Mota et al., 2020). Originating from debates about behavioral policy-making, the term *choice architecture* refers to a deliberate organizing of context in which people make decisions (Thaler and Sunstein, 2008, p. 3). As in the behavioral policy-making context, choice architectures on digital platforms are the result of conscious design decisions that should be subject to scrutiny by policy makers and regulators. As indicated in the introduction, digital platforms are environments in which the platform provider has full control over the platform’s choice architectures. As a result, if a digital platform raises competition problems, these problems can often be traced back to choice architecture decisions the platform provider made.

In section 4, we investigate several choice architectures under the control of Amazon. First, when multiple sellers offer to sell the same product on Amazon, one of the sellers is selected by default by the deployed *Buy Box algorithm* to feature on the product page (Chen, Mislove, and Wilson, 2016; Dash et al., 2024). Which seller will end up in the Buy Box arguably has significant impact on con-

sumer choice, as consumers may simply buy from the seller appearing in the default Buy Box, rather than comparing competing offers in detail. We explore the implications of Amazon's Buy Box design in section 4.1.

Second, the seller featured in the Amazon Buy Box and all other sellers are listed on a separate offer listing page (Dash et al., 2024). This page includes seller performance metrics such as average user ratings, percentage of positive feedback, and the number of ratings. This setup enables consumers to compare the sellers from whom they would like to purchase the product. We investigate the choice architecture of these offer listing pages in section 4.2.

Third, Amazon has designed its platform not only for consumers accessing the platform through a computer with a web browser. The platform can also be navigated using voice control: Amazon provides conversational voice search through its smart speakers (Echo devices) powered by the in-house voice assistant Alexa (Dash et al., 2022). As consumers accessing Amazon through voice control can only process limited amounts of information, Amazon's choice architecture for voice access presents a drastically reduced version of Amazon's interface to consumers. We explore the implications of some of Amazon's decision in this choice architecture in section 4.3.

Finally, Amazon provides various recommendation systems guiding consumer search and providing additional information about products and seller (Sorokina and Cantú-Paz, 2016; Dash et al., 2021). We explore some of the choice architectures underlying these recommendation systems in section 4.4.

While Amazon uses these and other choice architectures to design its platform, Amazon's role is not restricted to a platform designer. In addition, Amazon sells directly to consumers in most of its marketplaces – for example, in the United States, France, or Germany – making it one of the major sellers on its own marketplace. This puts Amazon in direct competition with independent sellers for many products that are sold on Amazon. In some countries – for example, India – Amazon is not allowed to sell products on its marketplace directly. As a response, Amazon created joint ventures in India in collaboration with other retailers to sell products on amazon.in. Cloudtail India and Appario Retail Private Ltd., for example, were two joint ventures which emerged as two of the biggest sellers on Amazon India.¹

Apart from operating the marketplace and being a seller, Amazon also provides fulfillment and logistical services. For example, Amazon stores third-party products in its fulfillment centers and packs and ships them to different customers upon arrival of orders. This service is usually known as Fulfillment by Amazon (FBA) (Amazon, 2020). For instance, sellers can opt for the full range of services offered in the Fulfillment by Amazon program; they can choose to store their products at their own warehouse and merely rely on Amazon's delivery service; or they can select to store, pack, and ship their products either on their own or through other

¹ While Cloudtail ceased to offer products on Amazon in 2022, Appario remains in the business of selling products on Amazon. Both companies were sellers on Amazon India while the data presented in this article were collected.

third-party services. By opting to use Amazon's subsidiary services (fulfillment and/or delivery), sellers (with whom Amazon competes on its marketplace) also become Amazon's customers for logistical services.

The intricate relationship between Amazon and sellers on Amazon Marketplace also plays out at the product level. Amazon manufactures its own *private-label products* and sells them under different brand names (for example, Amazon Basics, Presto, or Solimo; Dash et al., 2021). Note that only Amazon sells Amazon's private-label products on Amazon.

Finally, Amazon also maintains an advertising platform. For example, when a customer looks for a product on Amazon, the platform presents both organic and sponsored recommendations. While organic recommendations refer to the results that the algorithmic system produces in the desired context based on previous customer behavior, sponsored recommendations are sponsored product advertisements which are the result of second-price auctions in which sellers pay Amazon for displaying their product in a favorable position. Notably, Amazon also displays its own products among sponsored search results (Dash et al., 2021).

3 Amazon's Vertical Integration and Antitrust Laws

As the preceding section has shown, Amazon fulfills multiple roles on its digital platform, acting as a platform designer and operator, seller and product manufacturer, provider of fulfillment and logistical services, and operator of an advertising platform. The way how sellers and customers interact on Amazon's platform is heavily influenced by Amazon's choice architectures, which determines the context in which customers identify products and make purchase decisions.

The complex interactions between multiple players on Amazon are a result of continuous product development and innovation by Amazon over the last three decades. Originally, Amazon emerged as a book seller, entering the business of providing a platform for other sellers only later. Similarly, major online travel agencies such as Booking.com or Expedia over time became vertically integrated with meta-search platforms, enabling these travel agencies to compete with independent competitors on these platforms (Cure et al., 2022).

When a digital platform such as Amazon sells its own products and products from affiliated producers on its online platform, this raises issues of vertical integration. Can a vertically integrated platform that operates both as platform provider and seller sufficiently distinguish between both roles? Or is there a risk that it has incentives to bias search results and recommendations in favor of its own or affiliated products, at the expense of products from competing independent sellers?

Problems of vertical integration are not unique to digital platforms. The practice of supermarkets providing preferential shelf space to their own or affiliated products (see Tagliavini, 2023, p. 202), and financial advisers earning commissions that might bias their investment recommendations (Cookson et al., 2021; Egan, 2019), both raise similar questions. From an economic perspective, vertical integration can

have ambiguous effects on social welfare. Vertical integration of digital platforms may increase consumer search costs, discourage product innovation and prompt sellers to leave the market, as preferential treatment by the platform provider for own or affiliated products may make it more difficult for independent sellers to compete (see Farronato, Fradkin, and MacKay, 2023). However, it is not clear that such welfare-decreasing effects will always dominate. By engaging in preferential treatment, a dominant vertically integrated platform provider may have incentives to invest in product quality, product design, and innovation. It may also have an incentive to decrease prices and increase output because it can capture a larger fraction of the benefits from such actions (often referred to as the “internalisation of double mark-ups” – see European Commission, 2008, p. 7). Therefore, whether vertically integrated platform providers are beneficial or detrimental to economic welfare is a complex question that is often impossible to answer in the abstract, based solely on theoretical considerations (see, for example, de Cornière and Taylor, 2019; Hagiu, Teh, and Wright, 2022; Anderson and Bedre-Defolie, 2022).

These insights from industrial organization research have had an influence on competition law’s treatment of vertical integration. With the increasing influence of the Chicago School in antitrust, many jurisdictions have become reluctant to condemn cases of vertical integration as such. Rather, antitrust authorities investigate potential violations of antitrust laws in vertical relationships as part of a rule-of-reason analysis. Under rule-of-reason analysis, antitrust agencies and courts examine both the positive and negative effects of a certain conduct on competition in order to determine whether the conduct violates antitrust laws. For example, U.S. courts have shifted from a *per se* prohibition of resale price maintenance² and territorial restrictions³ to a rule-of-reason analysis that focuses on the economic effects of vertical restraints. Relatedly, self-preferencing by a dominant firm is not a *per se* violation of European competition law (Article 102 of the Treaty on the Functioning of the European Union, TFEU), but is subject to an effects test (Crémer, Montjoye, and Schweitzer, 2019, pp. 7 and 66).

This effects-based approach to assessing vertical integration under antitrust laws has influenced the way authorities have addressed alleged preferential treatment by digital platforms. The General Court of the European Union’s recent Google Shopping decision provides an example. In this decision, the court reviewed the European Commission’s finding that Google had promoted its own comparison shopping service on its general search engine at the expense of rival services (see General Court, Case T-612/17, November 10, 2021, Doc. ECLI:EU:T:2021:763).

² Compare *Dr. Miles Medical Co. v. John D. Park & Sons Co.*, 220 U.S. 373 (1911), treating a vertical price fixing scheme as *per se* illegal, with *Leegin Creative Leather Prods., Inc. v. PSKS, Inc.*, 551 U.S. 877 (2007), holding that all vertical distribution restraints – non-price and price – should be analyzed under the rule of reason.

³ Compare *United States v. Arnold, Schwinn & Co.*, 388 U.S. 365 (1967), holding that vertical restraints limiting distributors to exclusive sales areas is *per se* illegal, with *Continental T.V., Inc. v. GTE Sylvania, Inc.*, 433 U.S. 36 (1977), applying a rule of reason analysis to non-price distribution restraints protecting interbrand competition.

In its decision, the court pointed out that not all foreclosure is detrimental to competition, and that the mere extension of a dominant position to a neighbouring market is not necessarily anti-competitive (General Court, Case T-612/17, November 10, 2021, Doc. ECLI:EU:T:2021:763, paras. 157 and 162). While the court ultimately confirmed the European Commission's decision, the case demonstrates that whether self-preferencing violates European competition law can only be decided on a case-by-case basis. This analysis involves various different competition doctrines such as duty to deal, tying, margin squeeze, general abuse, or discrimination (see Colangelo, 2023).

Antitrust's rule-of-reason analysis indicates an important feature of this area of the law: antitrust law traditionally focuses on *ex post* regulation. Antitrust agencies and courts enforce rules, but they typically do this in a reactive manner. They investigate firm behavior and issue decisions to hold firms accountable for past non-compliance and to ensure compliance in the future.

Over the recent years, digital platforms have become subject to increased scrutiny with regard to their vertical integration. This includes their role as a platform provider, product manufacturer and seller, logistical service provider, and advertising platform operator. There has been a growing concern that antitrust law's focus on *ex post* intervention may not be sufficient to ensure viable competition in platform markets. In particular, network effects and switching costs might make it effectively impossible for sellers and customers on one platform to switch to a competing platform, thereby preventing competition in the market for digital platforms: it may just not be profitable for a competitor to Amazon and other digital platforms to create a competing marketplace.

In this environment, calls to introduce an *ex ante* regulatory regime gained prominence. In such a regime, rules and standards are established by the regulator before issues arise, and firms have to obey them when they become active in the market. When a particular conduct is prohibited, the regulator does not allow agencies or courts to take countervailing evidence into account, as is typical in rule-of-reason analysis. Once the regulator has determined that certain firm conduct is harmful, this conduct is prohibited, even if the firm could show that the conduct also had beneficial effects on economic welfare.

Along these lines, the European Union recently introduced a regulatory system for digital markets that is based on an *ex ante* approach where efficiency defenses raised by firms play a very limited role. With the Digital Markets Act (DMA), which came into full force in 2024, the European Union aims to make key digital markets fairer and more contestable by preventing large companies from abusing their market power and lowering barriers to the entry and expansion of smaller players. The DMA targets large digital platforms – so-called “gatekeepers” – such as Alphabet, Amazon, Apple, Meta, or Microsoft and imposes various obligations on them.

Article 6(5) of the DMA, for example, includes an explicit prohibition according to which gatekeepers “shall not treat more favourably, in ranking and related indexing and crawling, services and products offered by the gatekeeper itself than

similar services or products of a third party. The gatekeeper shall apply transparent, fair and non-discriminatory conditions to such ranking.” Importantly, this provision does not allow gatekeepers to provide justifications showing that any economic disadvantages created by self-preferencing are outweighed by pro-competitive effects. Efficiency defenses play no role in this ex ante regulatory approach.

Article 6(2) of the DMA prohibits digital platforms from using, in competition with other sellers, any non-public data that are generated or provided by those sellers. This prohibition is intended to countervail the ability of vertically integrated digital platforms to observe sales data from independent sellers with whom the platforms compete. The DMA also includes, among other provisions, prohibitions on combining data from different services belonging to the same company, protections for advertisers and publishers using the gatekeeper against monopolization trends, and provisions for ensuring interoperability and data access for business and individual users of platforms. If gatekeepers fail to abide by these obligations, the European Commission may impose fines of up to 10% of the gatekeepers’ annual worldwide turnover (Article 30(1) DMA). The Digital Market Act empowers the Commission to monitor the effective compliance of digital platforms with the obligations of the DMA, potentially using independent external experts, auditors, and member state authorities (Article 26 DMA). Article 21 of the DMA entitles the European Commission to require access to data and algorithms of digital platforms.

Going beyond the Digital Markets Act, the European Union also enacted the Digital Services Act (DSA), which came into effect in 2024 as well. This regulation provides a comprehensive framework regarding illegal content, transparent advertising, and disinformation of online platform providers. Article 26 of the DSA, for example, requires online platforms to be transparent about the nature and origin of advertisements. This obligation also applies to “sponsored links” displayed by digital platforms. Articles 25 and 27 of the DSA include rules about online interface design and recommender system transparency. The European Commission may impose fines of up to 6% of the annual worldwide turnover of very large online providers (Article 74(1) DSA), which includes Amazon, Alibaba, Apple, Google Shopping, and Zalando, among other providers.⁴

4 Amazon and Preferential Treatment

The effectiveness of regulatory action against platforms such as Amazon will always depend on the extent to which their behavior can be observed. While public discussions about these platforms often rely on individual cases or anecdotes, algorithmic enforcement tools might open up a way to investigate potential self-preferencing behavior in a systematic, quantifiable manner. Doing so requires tools

⁴ For a list of very large online platforms as designated by the European Commission, see <https://digital-strategy.ec.europa.eu/en/policies/list-designated-vlops-and-vloses>, accessed April 2, 2024.


to observe Amazon’s choice architectures and their effects on customer behavior in an automated, scalable manner.

In the following section, we summarize a series of empirical investigations into various of Amazon’s choice architectures undertaken by several co-authors of this article (Dash et al., 2021, 2022, 2024). The purpose of this section is to describe their main findings and distill important insights about the promises and challenges of algorithmic enforcement of platforms.

4.1 Amazon Buy Box

We start our exploration of Amazon’s market practices by investigating how the service decides which offer to feature in its Buy Box (now called “Featured Offer”).⁵ A typical Amazon product page has a rectangular box on the top-right corner, right next to the title of the product. This box is popularly known as the Buy Box. Figure 1 shows an example of such a Buy Box. The box contains the price of the product, a delivery estimate, and two buttons that read: “Add to Cart” and “Buy Now.” The box also lists the name of a seller. In the Buy Box depicted in Figure 1, the seller is Appario Retail Private Ltd.

Figure 1
Example of a Buy Box on Amazon



In principle, there are various ways through which customers can buy products on Amazon. However, given the prime positioning of the Buy Box (and the prominently displayed buttons inside the box), it accounts for more than 80% sales on Amazon (Lanxner, 2021).

Given the importance of the Buy Box for transactions on Amazon’s platform, Amazon exerts considerable power over customers’ purchase decisions through

⁵ This subsection is based on Dash et al. (2024).

its algorithm that selects the seller that appears in the Buy Box. This decision is trivial if there is only one seller that offers a given product. However, in practice, multiple sellers can and often do offer the same product. In this scenario, Amazon selects the seller to be featured in the Buy Box through a proprietary algorithm known as the *Featured Offer Selection* algorithm, also commonly known as the Buy Box algorithm. The precise structure of this algorithm is unknown, making it a subject of speculation (and concern) among sellers on Amazon (Kaneshiro, 2020). Amazon presents some clues and recommendations on its seller interface regarding what factors its algorithms take into account in determining the winner of the Buy Box. These factors include, for example, competitive pricing and seller performance metrics.

Amazon's Buy Box algorithm has been at the center of multiple antitrust investigations. A main concern is that Amazon might design the algorithm in a way that amounts to self-preferencing. To understand what that means, remember that Amazon directly competes with other business users for selling products on its marketplace. As discussed in section 2, Amazon also provides paid logistical services to sellers who opt for it, raising the question whether sellers that purchase Amazon's services receive preferential treatment.

We start our investigation by analyzing how frequently Amazon (or its special merchants in India, see section 2) win the Buy Box. To answer this question, we collected data for over 70,000 Buy Box competitions in Amazon's Indian, American, German, and French marketplaces. Specifically, we recorded the top-100 popular queries on Amazon and searched for them on each of the marketplaces. We then recorded, for each marketplace, all the products that appeared on the first page of the search result page. For all products, we proceeded to the product page and checked whether multiple sellers offered the corresponding product. If the answer to this question was yes, we collected information about the sellers and the winner of the Buy Box. Since Buy Box winners can potentially change between different queries, we collected the data for a span of two weeks.

During our period of observations and across all the countries analyzed, we observed that Amazon won upwards of 80% of the Buy Box competitions where it was involved as a seller. At the same time, only 11%, 13.22%, 17.74%, and 26.23% of the sellers competing with Amazon succeeded in winning at least one Buy Box competition throughout our data collection period in the Indian, U.S., German, and French marketplaces, respectively.

Of course, the fact that Amazon won most Buy Box competitions on its own does not prove self-preferencing behavior on the part of the platform. After all, it appears possible that Amazon's own offers are preferred by its customers. If that was the case, the Buy Box algorithm would simply give customers what they want, not nudge them towards purchase decisions that benefit Amazon at the expense of its customers and competing sellers.

To investigate this question, we first compared the price offered by Amazon and rival sellers in cases in which Amazon won the Buy Box. We found considerable numbers of instances in which Amazon won the Buy Box despite offering the prod-

uct at a price that is higher than that offered by at least one rival seller: In the USA and in India, 25% of the Buy Boxes won by Amazon and its special merchants were cases where Amazon did not offer the corresponding product at the lowest price. In the European marketplaces, these numbers were slightly lower (12.45% for France and 13.91% for Germany).

However, information on offer prices cannot exclude the possibility that the results of the Buy Box algorithm, despite its apparent preference for Amazon over other sellers, produces results that map to consumer preferences. If the outcomes of Buy Box competitions usually match with consumer preferences, Amazon’s choice architecture might be designed in a welfare-maximizing way.

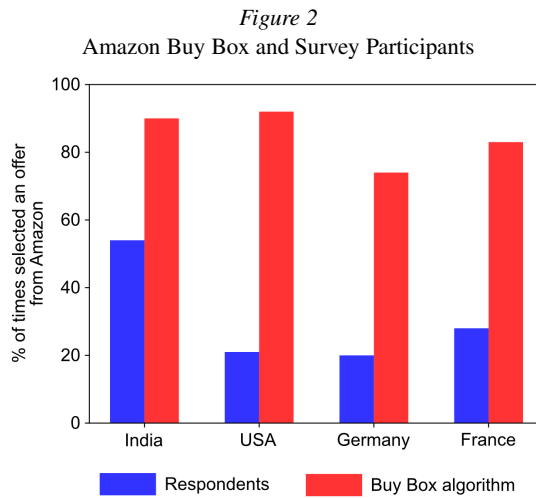
Determining consumer preferences in digital markets in a reliable and reproducible manner is a challenging task. The data available to us from our empirical investigation of Amazon is insufficient to answer the question whether the Buy Box algorithm’s results correspond to consumer preferences. We therefore turn to a different empirical method that allows us to gain insights into consumers preferences: Using the Prolific crowdsourcing platform, we conducted a survey among 200 respondents (50 from each of the four countries).

In the survey, we showed survey participants from each country the four best offers for a product and elicited their preference regarding the different sellers. We started by randomly selecting, for each country, 15 products from among the products for which Amazon, even though it did not offer the lowest price, won the Buy Box in the first iteration of the data collection. For each product, we showed our survey participants the four offers with the lowest price and delivery charges. For each of the four selected offers, we displayed the name of the seller, the price offered (including delivery charges, if any), the average user rating of the seller, the percentage of positive feedback of the seller, the number of ratings received, and delivery options.⁶ We then asked participants the following question: “Suppose you are willing to buy a *<product name>* on Amazon and these are the offers from different sellers for the same product. Which one would you prefer to buy?”

Out of a total of 3,000 evaluations, participants chose the offer from Amazon as their first preference at only 31%. This preference varied across countries (as shown in Figure 2), with 54% in India, 21% in the USA, 20% in Germany, and 28% in France. In contrast, if we consider the winner of the Buy Box for the 60 products surveyed at different time points that we observed during our data collection, more than 80% of the times the Buy Box had been won by Amazon (90% in India, 92% in the USA, 74% in Germany, and 83% in France). This suggests a considerable mismatch between the frequency at which the Buy Box algorithm selected an offer from Amazon and the frequency at which survey participants selected the same.

This result raises doubts about whether consumer preferences align with seller choices determined by Amazon’s Buy Box competitions. By contrast, it might sug-

⁶ We obtained this information from the offer listing page for that product, which will be discussed in greater length in section 4.2. For the present context, it suffices to say that Amazon displays all competing sellers on a separate page, together with information about the seller, product price, and delivery charges.



Notes: This figure shows how often survey participants and Amazon’s Buy Box algorithm selected an offer from Amazon for the surveyed products. We observe a significant gap in survey participants’ preference as opposed to that of the algorithm’s.

gest a potential pattern of self-preferencing in the Buy Box algorithm in scenarios in which other sellers offer the product at a lower price than Amazon. In these scenarios, the Buy Box algorithm might nudge consumers towards Amazon as a seller.

4.2 Offer Listing Pages

We continue our exploration of Amazon’s market practices by investigating its offer listing page, another choice architecture in the Amazon ecosystem.⁷ As described in the previous section, product pages on Amazon include a Buy Box that allows for quick product purchases and for which Amazon, if necessary, selects a seller from among the competing sellers offering a product. However, consumers are not bound to use the Buy Box. Amazon offers consumers a separate page which lists all the available offers for a product. Figure 3 displays an example of such an offer listing page. To get to the offer listing page, consumers can click a link that appears in a second rectangular box displayed below the Buy Box (see Figure 1). The offer listing page shows the offers from all the competing sellers, ordered by price and delivery charges. It also shows some seller performance metrics including its average user rating, the percentage of positive feedback, and the number of ratings.

In our empirical investigation, we focused on two specific characteristics related to how Amazon displays seller performance metrics. First, while some metrics provide an indication of the (average) quality of service offered by a seller, the number

⁷ This subsection is based on Dash et al. (2024).

Figure 3
Offer Listing Page on Amazon: Example

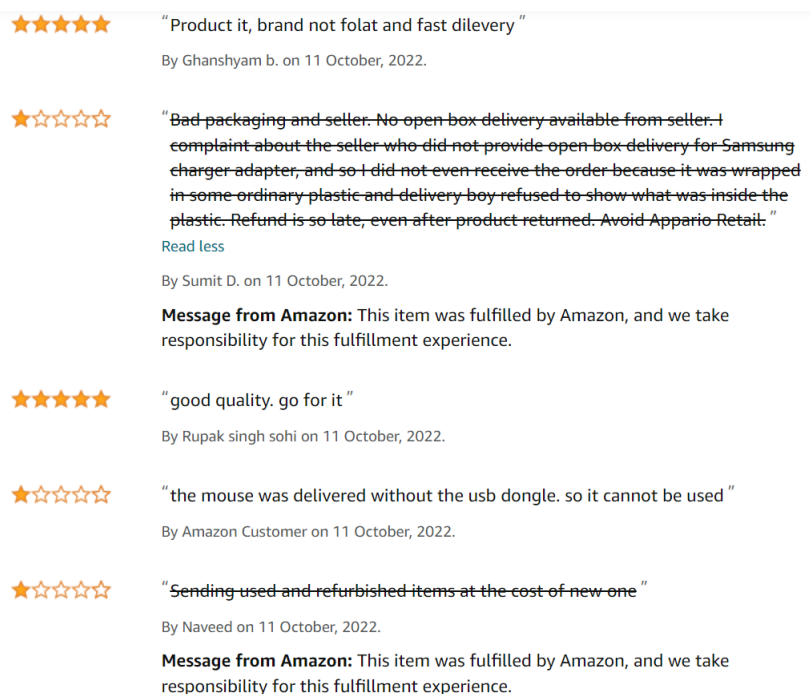
The screenshot shows two product offers on an Amazon page. The top offer is for a product priced at ₹275⁰⁰ (₹34.38 / count), with a 10% discount on SBI Cards and free delivery on Monday, 10 October. The bottom offer is for a product priced at ₹294⁰⁰ (₹36.75 / count), with a 10% discount on SBI Cards and ₹95 delivery on Wednesday, 12 October. Both offers include an 'Add to Cart' button and a 'Filter' dropdown menu.

6 other options sorted by price + delivery: low to high		Filter
New ₹275 ⁰⁰ (₹34.38 / count)	10% Off on SBI Cards Details ✓ FREE delivery Monday, 10 October. Details	Add to Cart
Ships from	7.color	
Sold by	7.color Just launched	
New ₹294 ⁰⁰ (₹36.75 / count)	10% Off on SBI Cards Details ✓ ₹95 delivery Wednesday, 12 October. Details	Add to Cart
Ships from	BIZKART	
Sold by	BIZKART ★★★★☆ (15 ratings) 80% positive	

of ratings (# Ratings) reflects the total feedback ratings received by a seller throughout their tenure on Amazon. The inclusion of such metrics could potentially sway customers towards larger, more established sellers. It is important to recognize that inferring anti-competitive intent solely from this design choice is challenging, as consumers may naturally gravitate towards proven sellers for a sense of reliability and trustworthiness.

Second, we turn to a specific feature of Amazon's approach to displaying customer ratings on the offer listings page that might be less innocuous, specifically because it might offer advantages to sellers who use its subsidiary fulfillment and/or delivery services (see section 2). For these sellers, under certain circumstances, Amazon reserves the right to "strike through" negative feedback by consumers. For example, Amazon may strike through a seller feedback if "[t]he entire comment relates explicitly to delivery experience for an order fulfilled by Amazon" or "[t]he entire comment is related to a delayed or undelivered order, which you shipped on time by using Buy Shipping" (Amazon Seller Central, 2022a). In both cases, in addition to the strike through, Amazon puts up a message taking responsibility for the inconvenience faced by the customer. Figure 4 shows examples of struck-through reviews. Amazon states in its Fulfilled-by-Amazon and Amazon Shipment advertisements that negative reviews which have been struck through will not impact sellers' performance metrics (Amazon Seller Central, 2022b,c). Importantly, this policy is unavailable for independent sellers that do not use Ama-

Figure 4
Review Striking Through: Example



Notes: This figure provides an example of review strike throughs: the second and the last reviews have been struck through. A message from Amazon is mentioned taking responsibility of the issue.

zon's fulfillment or delivery services. Hence, for the same customer inconvenience, sellers might be treated differently by Amazon depending on whether they purchase services from Amazon or not.

To investigate the impact of the strike-through policy on Amazon seller performance metrics, we collected feedback data for the top-1000 active sellers in Amazon's Indian, U.S., German, and French marketplaces. Our dataset consists of a total of 4 million seller feedback reviews for these 4,000 sellers. The data collected include the feedback text, associated ratings, and whether it was struck through or not by Amazon.

We categorized the sellers in our dataset into three main groups: (1) Amazon Fulfilled, (2) Amazon Shipped, and (3) Merchant Fulfilled. The determination of seller membership in the first two groups was based on the text displayed when Amazon strikes through reviews. Sellers falling into the third group were identified if none of their reviews in the dataset were struck through. Subsequently, we con-

ducted a manual inspection of their product pages to ascertain how they fulfill their orders.

We then assessed how often Amazon struck through seller feedback pertaining to members of the first group, and how the striking through of feedback affected the relative ratings of Amazon Fulfilled sellers as opposed to Merchant Fulfilled sellers. Notably, more than half of the reviews in which Amazon Fulfilled sellers received only three stars or less was struck through in every country we investigated. For Amazon Shipped sellers, the numbers hover between 22% and 4% dependent on the country. Reviews for Merchant Fulfilled Sellers were never struck through.

The exclusion of struck-through feedback markedly changes how different types of sellers perform in ratings. For example, when excluding struck-through reviews, the overall ratings are best for Amazon Fulfilled sellers. When struck-through reviews are included in the evaluation, Amazon Fulfilled appear to have lower rankings than Merchant Fulfilled sellers in India, the USA, and Germany.

Consistent with our previous investigations, these findings suggest that Amazon's choice architecture favors sellers affiliated with Amazon. Different from the inclusion of the number of ratings as a performance metric, it seems less likely that the decision to exclude certain unfavorable ratings for Amazon-affiliated sellers from the calculation of seller performance is aligned with consumer preferences. At the very least, it results in the display of ratings that are not comparable across different types of sellers, thereby reducing their informativeness to consumers. At worst, it could distort consumers' purchase decisions.

But do these design decisions have a real impact on consumer's purchase behavior? To investigate this question further, we again conducted a survey, this time with participants in India. In the survey, we exposed participants to several counterfactual settings. We showed them (a) the Amazon-reported seller metrics with number of ratings (# Ratings), (b) the Amazon reported-metrics seller without # Ratings, (c) the rectified metrics (including the struck-through feedback) with # Ratings and, finally, (d) rectified metrics without # Ratings. We recorded the percentage of times respondents selected the offers from Amazon special merchants or Amazon Fulfilled sellers in each of the settings. The results are listed in the table below.

The results of this survey suggest that when participants saw the rectified seller metrics, their preference toward sellers using Amazon's logistical services reduced significantly in all cases. Similarly, when the number of ratings were withheld, the participant's preference toward Amazon special merchants reduced by nearly 40%.

In summary, the observations in this section indicate that various aspects of Amazon's offer listings page, including its strike-through policy for negative reviews, have a divergent impact on different types of sellers, with those affiliated with Amazon receiving preferential treatment. Additionally, our findings suggest that these design choices can significantly influence consumers' purchase decisions.

Table
Preferences by Survey Participants

Percentage of 1st preference votes	Amazon metrics	Rectified metrics
Amazon special merchants		
with # Ratings	54.26%	28.53%
without # Ratings	14.53%	6.93%
Amazon Fulfilled		
with # Ratings	81.20%	64.53%
without # Ratings	72.13%	54.6%

Notes: This table shows the percentage of first preference votes aggregated by Amazon special merchants and Amazon Fulfilled sellers in the different survey settings. We observe that the percentages drop when we show rectified metrics as opposed to metrics shown by Amazon or when we omit the # Ratings feature.

4.3 Alexa Search

Next, we shift our focus to product selections facilitated by Amazon’s virtual assistant, Alexa.⁸ While many consumers traditionally purchase products from Amazon through its website, there exists an alternative avenue: Alexa enables consumers to make purchases through voice commands, interacting with smart speakers.

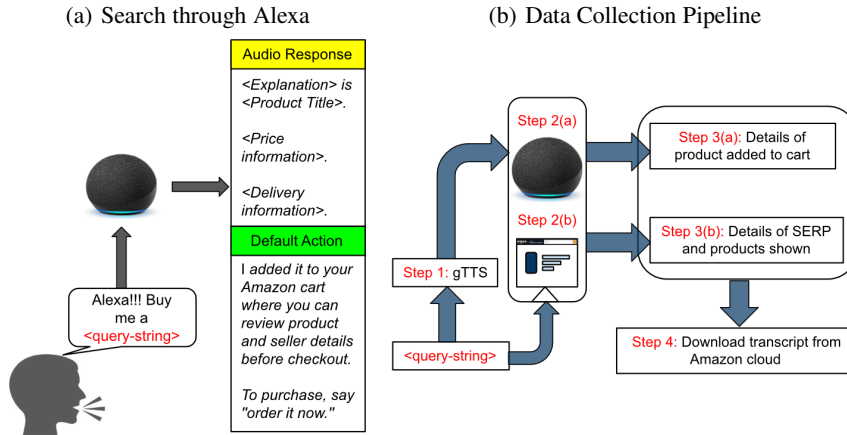
Figure 5(a) graphically depicts a query and a typical response from Amazon Alexa. When a consumer asks a smart speaker powered by the Alexa virtual assistant to purchase a product online, the typical responses has two main parts: (1) an *audio response describing a chosen product with a brief explanation*, and (2) a *status quo or default action*. As part of its product description, Alexa spells out relevant product details, including its title, price, and delivery information. Often, the response also contains a brief explanation of why the virtual assistant has chosen the corresponding product. Explanations provided by the virtual assistant include that the product is “Amazon’s Choice” or “a top result.” Alexa’s default action is to add the product to the customer’s cart for further review or purchase. In addition, the virtual assistant asks the consumer whether she wants to make the purchase immediately.

Compared to consumers interacting with Amazon on a website, consumers enjoy a significantly smaller action space when interacting with Amazon through a virtual assistant such as Alexa. Rather than presenting a list of search results, Alexa spells out the details of a single product and adds it to the consumer’s cart. By reducing consumer choice, Amazon’s default choices can have a particularly strong influence on consumer decisions, even stronger than the selection of a seller in the Buy Box.

Similar to our previous investigations, we start our investigation by collecting information on the products that Alexa selects in response to searches and how these products compare with the search results of a similarly worded search on Ama-

⁸ This subsection is based on Dash et al. (2022).

Figure 5
Amazon Alexa



Notes: Panel (a) shows the response to a product search through Alexa: In response to a purchase query, (i) the virtual assistant spells out an audio response explaining product information, and (ii) adds the product to cart for further exploration. Panel (b) shows the data collection pipeline: a query string gets converted to an audio signal by Google Text-to-Speech. The audio signal and the textual query are provided to Alexa and desktop search respectively. Finally, the product page details of the retrieved products and the transcripts of the conversation are collected.

zon's web page. Overall, we collected search results for 1,000 different queries from both Amazon Alexa and from Amazon's website. We also collected the meta-data (including product title, ratings, price, and seller information) for the products added to the cart by Alexa and for products appearing among the search results on Amazon's website. The exact data collection framework is shown in Figure 5(b).

We then conducted surveys with 100 participants to elicit consumers' preferences between the product added to cart by Alexa and the top search result (appearing in the first position) from a desktop search. We selected the top search result for comparison because Amazon identified it as the most relevant product for the query at the time for the customer. We showed our survey participants the titles of the two products, their prices, their average user ratings, and the number of ratings received by each of the two products – information we obtained from the Amazon results page during our data collection. We then asked participants the following question: *Suppose you are looking for "query-string." Which of the following would you prefer to buy?*

All 100 participants evaluated ten such queries. Out of the 1000 evaluations we received (100 respondents \times 10 queries), 732 survey participants (73.2%) chose the top desktop search result over the product added to cart by Alexa for the corresponding queries. Analogous to what we found in our investigations of Amazon's

Buy Box algorithm, this might offer suggestive evidence that a majority of participants prefer product offerings that differ from Alexa's choices.

Overall, we found that Alexa's product selection mechanism often selected products that were different from the top search results on Amazon's webpage, and that, in a majority of cases, consumers would have preferred the top search result over Alexa's product choice. Depending on the identity of the products' sellers, these choices might constitute another example of self-preferencing or preferential treatment of Amazon-affiliated sellers.

4.4 Recommendation Systems

This subsection turns to yet another choice architecture by which Amazon may engage in self-preferencing: recommendation systems.⁹

As an e-commerce marketplace, Amazon provides basic product recommendation functions on its platform (see section 2). It seems reasonable to assume that the results of such recommendations will often be relevant determinants for consumers' purchase decisions. These functions may offer Amazon yet another opportunity to "push" its products at the expense of products manufactured and/or sold by others.

Figure 6 shows a screenshot of one type of product recommendations on Amazon, namely related item recommendations. These related item recommendations appear towards the bottom of almost all Amazon product pages. Naturally, the number of products Amazon can display in this context is limited, and a crucial question is how Amazon's algorithms decide which products a customer can see.


Amazon displays both (a) organic results, which appear to be determined on the basis of prior customer behavior, and (b) sponsored results, with the items displayed determined through second-price auctions in which sellers can participate.¹⁰ When Amazon displays sponsored results, it marks those with a tag that reads, on its English language platform, "Sponsored."

In recent years, Amazon appears to have given more room in its related item recommendations to sponsored results. For example, consider the upper carousel in Figure 6: although the title reads "Products related to this item," there is a small text in gray color toward the top left of the carousel that reads "Sponsored." This essentially means that the entire carousel does not consist of organic recommendations, but that these are all sponsored product advertisements.

⁹ This subsection is based on Dash et al. (2021).

¹⁰ For an example of related item recommendations, see the lower carousel of recommendations in Figure 6. The title of the carousel reads "Customer who viewed this item also viewed." As the title suggests, the items in the carousel are essentially similar to the product on whose product pages they appear based on how customers viewed the items in prior exploration sessions. As has been described elsewhere, these results are curated from item-item collaborative filtering processes (Linden, Smith, and York, 2003; Smith and Linden, 2017). Similarly, different approaches used for evaluating Amazon's search results are discussed in prior work (Sorokina and Cantú-Paz, 2016; Tsagkias et al., 2021).

Figure 6
Related Item Recommendations on Amazon



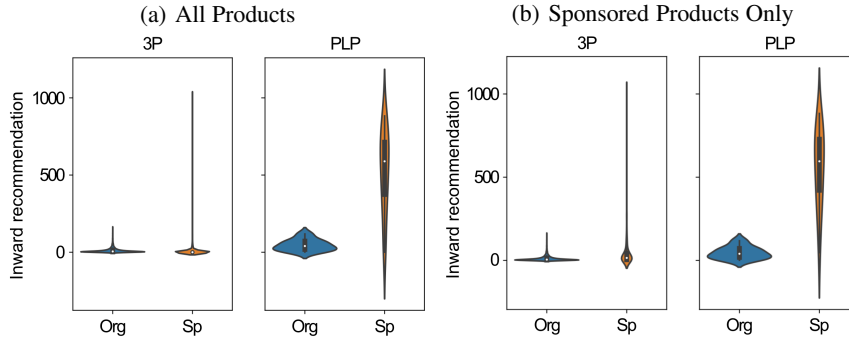
Notes: This figure provides an example of related item recommendations on Amazon.

Of course, Amazon is not the only platform that shows sponsored results alongside organic results. Also, given that sponsored results are selected through a different mechanism than organic results, it is to be expected that the inclusion of sponsored results among related item recommendations will result in fewer organic results shown to the customer, which might be perceived as depriving the customer of valuable information. However, given Amazon’s vertical integration, we are specifically interested in whether Amazon uses the sponsored results architecture to promote its private-label products.

To analyze this question, we investigate the products shown on Amazon’s websites as related product recommendations for batteries. We collected more than 100,000 related item recommendations from the product pages of 5,352 products in the battery category of Amazon using a breadth first search crawler. Out of these 5,352 products, only 17 products (0.3%) were Amazon private-label products.

We focus on demonstrating a disparity in recommendations that Amazon private-label products receive as opposed to third-party products in sponsored versus organic recommendations. Figure 7 demonstrates that Amazon private-label products on average received a much higher number of recommendations among sponsored results (520) as compared with organic product recommendations (46). In other words, Amazon private-label batteries were recommended more than 10 times as frequently in the context of sponsored results than in the context of organic results. Third-party products do not experience a similar increase in appearances as recommended products among sponsored recommendations. Such products were recommended, on average, 11 times in organic results. This number does not change much for sponsored results. In fact, a large number (4,357 out of 5,335) of third-party products do not receive any sponsored recommendations, thus dragging the

Figure 7
Recommendations on Amazon



Notes: This figure shows the distribution of inward recommendations toward Amazon private-label products (PLP) and third-party (3P) products. Panel (a) depicts distributions for all products, panel (b) for products that appear at least once as sponsored results.

mean number of recommendations to 9. By contrast, almost all private-label products (all but one) are featured among sponsored recommendations.

Panel (b) shows that, even if one limits the analysis to products that are featured in the sponsored recommendations, the increase in recommendations is larger for private-label products than it is for third-party products. The average number of recommendations for the former is 46 among organic recommendations and 553 among sponsored recommendations. For the latter, the number increases from 13 for organic recommendations to 46 among recommendations.

Another way to look at these numbers is to compare how often private-label products appear as recommended products in the context of organic recommendations vs. sponsored results. Amazon's own organic recommendation engine recommended Amazon private-label products in approximately 15% of recommendations shown to consumers. Simultaneously, among sponsored recommendations, private-label products occupied almost 50% of the available slots.

Accordingly, it appears that Amazon, by virtue of the sponsored recommendations, made its private-label batteries much more accessible to consumers looking at other products' web pages. To the extent that the organic recommendation engine captures consumer preferences regarding what other products they consider interesting alternatives to the product they are looking at, this could be viewed as yet another choice architecture that implements a type of self-preferencing.

5 Discussion

In this article, we explore unique design features of digital platforms, focusing on those that depend heavily on algorithms and choice architectures. We address reg-

ulatory challenges that emerge when such platforms integrate vertically across various stages of their product and service production and distribution. Through four empirical case studies, we demonstrate how combining large-scale analysis with consumer surveys can uncover issues such as preferential treatment by platforms and discrepancies between platform offerings and consumer preferences.

We want to highlight four key findings from our research. First, our empirical case studies show that evaluating self-preferencing on digital platforms demands a detailed, case-by-case analysis. These platforms use a variety of complex choice architectures, each of which may include self-preferencing elements. A thorough analysis is essential for regulators and researchers to identify such elements, a task made difficult by the often opaque nature of platform design decisions. Distinguishing between beneficial and detrimental effects of choice architectures can be challenging, and it remains to be seen whether data access rights and reporting duties, as established under the European Union's Digital Markets Act and Digital Services Act, will be sufficient to overcome this challenge.

Second, our empirical case studies demonstrate how comparing choice architectures with consumer preferences can allow for insights into which choice architectures may be problematic. They also showcase how surveys can provide an avenue for comparing the impact of product offering choice architectures with consumer preferences. For instance, if a choice architecture nudges consumers towards product A over product B, but surveys show a clear preference for product B, this discrepancy could point towards a mismatch between choice architectures and consumer preferences.

It is important to stress that we do not view consumer surveys as flawless or the sole method for understanding consumer preferences. Browser extensions and click-stream data present alternative methods for capturing consumer preferences, enabling researchers and regulators to directly observe consumer search and purchasing activities. Consumer panels, often equipped with custom browser extensions, have previously been employed for algorithmic auditing, particularly to observe search results and product features (Farronato, Fradkin, and MacKay, 2023). Observing actual consumer behavior through click-stream data could be another avenue to explore consumer preferences (Bechtold and Tucker, 2014). While a detailed exploration of these tools is beyond this article's scope, they could enable precise comparisons between choice architectures and consumer preferences on digital platforms. As such, these tools could provide an important ingredient in regulators' enforcement actions against digital platforms.

Empirical investigations such as those explored in this article are no substitute for a thorough legal analysis of platform behavior under applicable legal standards. They should be considered as one component of a broader enforcement strategy, potentially relevant to a range of legal issues. In this article, we deliberately avoid making definitive statements about legal compliance, as a comprehensive legal analysis would encompass factors beyond the scope of the tools and methods discussed here.

Third, as section 4 has shown, choice architectures influence consumer behavior in multiple ways. How product searches are processed and search results presented can influence not only consumer purchasing decisions. It can also shape consumer preferences from the outset. With digital platforms becoming an integral part of our daily lives, establishing a neutral “ground truth” – determining what a consumer would choose in the absence of these architectures – poses a significant challenge. This may complicate future regulatory efforts that seek to base interventions on a deviation of a platform architecture from an objectively defined “ground truth.” For example, uncertainty over whether consumer preferences are influenced by a specific choice architecture could present regulators with difficult questions when contemplating interventions against platforms for perceived self-preferencing practices that seem contrary to consumers’ best interests.

Fourth, our case studies on Amazon’s choice architectures underscore the need for caution in addressing digital platform issues solely through *ex ante* regulation. Regulating choice architectures involves intervening in the core functions of digital platforms, a move that could blur the lines between competition policy and market design. This does not mean that *ex ante* regulation is always inappropriate for digital markets. However, it comes with its own set of challenges. In some instances, fine-grained regulatory interventions may become so costly in terms of design, audit, and compliance that a regulator should refrain from such interventions.

On a more general level, our empirical case studies suggest a future where algorithmic auditing becomes a key part of regulating the digital economy. By employing algorithmic auditing tools, regulators and researchers could uncover new forms of self-preferencing previously undetectable without the aid of automated, large-scale analysis. Although our focus here has been on Amazon’s choice architecture and its impact on consumers, similar methodologies could investigate its effects on independent sellers, delivery logistics providers, and fulfillment centers. Importantly, the tools discussed in this article have implications beyond regulatory and research interests. Digital platforms themselves could leverage these tools to proactively identify and address potential legal and societal risks they might otherwise inadvertently contribute to.

Moving towards a future of algorithmic auditing and compliance requires regulations that are relatively easy to audit in an algorithmic, large-scale manner. In such a future, it may be particularly important for policy makers to design regulations in a way that allows for easy algorithmic auditing (Sandvig et al., 2014; Metaxa et al., 2021). This might involve not just granting data access rights to regulators and researchers and adopting standardized data formats, but also designing regulatory obligations that are straightforward to encode and monitor algorithmically. Additionally, clarifying the legal permissions for researchers to use web scraping and other tools for compliance assessment could significantly enhance the effectiveness of these regulations.

6 Conclusion

In this article, we have explored how digital platforms whose architectures are characterized by algorithmic systems and deliberate choice architectures can be investigated utilizing large-scale algorithmic audits and consumer surveys. While our empirical case studies focused on Amazon, our approach could be applied to other online platforms that act as an intermediary among various stakeholders as well. We believe that our case studies demonstrate how empirical methods originating in computer science can equip policy makers, regulators, and researchers with the tools necessary for detailed monitoring of digital platform operations. This, in turn, provides valuable empirical evidence to inform the ongoing debate on effective digital market regulation. Additionally, our analysis highlights the difficulties researchers and regulators face when trying to identify anti-competitive behaviors or impacts from observed platform conduct. Specifically, our case studies shed light on the significant challenges in determining whether a platform's choice architecture is designed to unfairly benefit its own interests or to genuinely enhance service quality.

In a future world of regulating digital markets through algorithmic auditing, we envision tools enabling computer scientists, economists, and legal scholars to analyze firm behavior under digital regulations, to empirically identify unintended consequences of such regulations, and to predict, *ex ante*, the effects of further or alternative regulatory interventions. We envision tools that, at least partially, automate the auditing and compliance processes, thereby reducing compliance costs for firms while offering regulators automated enforcement mechanisms. Such tools could find applicability beyond antitrust law, extending to privacy law (see Manandhar, Singh, and Nadkarni, 2024; Zac et al., 2024), content moderation, intellectual property law, and many other legal domains. Regulating digital markets through algorithmic auditing could not be accomplished by one or two of the disciplines on their own. We need computer scientists to develop new tools for large-scale data collection and automated regulation enforcement. We need economists to build frameworks for understanding social implications and welfare effects, alongside econometric methods for measuring causal impacts of firm behavior. And we need legal scholars to study the normative goals of particular regulations, the trade-offs involved in achieving such goals, and to understand the institutional specifics of enforcement mechanisms. In this sense, the answer *to* the machine may not be *in* the machine (Clark, 1996). It may be in more interdisciplinary research.

References

- Amazon (2020), "Amazon FBA – Get the Prime Advantage," <https://services.amazon.in/services/fulfilment-by-amazon/pricing.html>, accessed April 2, 2024.
- Amazon Seller Central (2022a), "Can Amazon Remove Buyer Feedback?" <https://amzn.to/3EiL2Pa>, accessed April 2, 2024.
- (2022b), "Use Buy Shipping Services," <https://amzn.to/3rWhlXr>, accessed April 2, 2024.

- (2022c), “What Happens when I Receive a Negative Feedback on an AFN Order?” <https://amzn.to/3ecvfbf>, accessed April 2, 2024.
- Anderson, Simon, and Özlem Bedre-Defolie (2022), “Online Trade Platforms: Hosting, Selling, or Both?” *International Journal of Industrial Organization*, 84, DOI: 10.1016/j.ijindorg.2022.102861.
- Bandy, Jack (2021), “Problematic Machine Behavior: A Systematic Literature Review of Algorithm Audits,” *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW1), DOI: 10.1145/3449148.
- Bechtold, Stefan, and Catherine Tucker (2014), “Trademarks, Triggers, and Online Search,” *Journal of Empirical Legal Studies*, 11(4), 718–750.
- Bradford, Anu (2012), “The Brussels Effect,” *Northwestern University Law Review*, 107(1), 1–68.
- (2020), *The Brussels Effect: How the European Union Rules the World*, Oxford University Press, New York (NY).
- Chen, Le, Alan Mislove, and Christo Wilson (2016), “An Empirical Analysis of Algorithmic Pricing on Amazon Marketplace,” in: *Proceedings of the 25th International Conference on World Wide Web, WWW '16*, International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, pp. 1339–1349.
- Clark, Charles (1996), “The Answer to the Machine Is in the Machine,” in: P. Bernt Hugenholtz (ed.), *The Future of Copyright in a Digital Environment: Proceedings of the Royal Academy Colloquium*, Kluwer Law International, The Hague, p. 139–146.
- CloudIncome (2023), “30 Eye-Opening YouTube Facts, Figures and Statistics you Should Know in 2023,” <https://cloudincome.com/youtube-statistics>, accessed April 15, 2023.
- Colangelo, Giuseppe (2023), “Antitrust Unchained: The EU’s Case against Self-Preferencing,” *GRUR International*, 72(6), 538–556.
- Cookson, Gordon, Tim Jenkinson, Howard Jones, and Jose Vicente Martinez (2021), “Best Buys and Own Brands: Investment Platforms’ Recommendations of Mutual Funds,” *The Review of Financial Studies*, 34(1), 227–263.
- Crémer, Jacques, Yves-Alexandre de Montjoye, and Heike Schweitzer (2019), “Competition Policy for the Digital Era,” <https://op.europa.eu/en/publication-detail/-/publication/21dc175c-7b76-11e9-9f05-01aa75ed71a1/language-en>, accessed March 26, 2024.
- Cure, Morgane, Matthias Hunold, Reinhold Kesler, Ulrich Laitenberger, and Thomas Lariou (2022), “Vertical Integration of Platforms and Product Prominence,” *Quantitative Marketing and Economics*, 20(4), 353–395.
- Dash, Abhisek, Abhijnan Chakraborty, Saptarshi Ghosh, Animesh Mukherjee, and Krishna P. Gummadi (2021), “When the Umpire is Also a Player: Bias in Private Label Product Recommendations on E-Commerce Marketplaces,” in: *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, Facct '21, Association for Computing Machinery, New York (NY), pp. 873–884.
- , —, —, —, and — (2022), “Alexa, in you, I Trust! Fairness and Interpretability Issues in E-Commerce Search through Smart Speakers,” in: *Proceedings of the ACM Web Conference 2022, WWW '22*, Association for Computing Machinery, New York (NY), pp. 3695–3705.
- , —, —, —, and — (2024), “Investigating Nudges toward Related Sellers on Ecommerce Marketplaces: A Case Study on Amazon,” *Proceedings of the ACM on Human-Computer Interaction*, forthcoming.
- Davis, Kevin E., and Florencia Marotta-Wurgler (2024), “Filling the Void: How E.U. Privacy Law Spills Over to the U.S.,” Law and Economics Research Paper 24-15, New York University School of Law, New York (NY).
- de Cornière, Alexandre, and Greg Taylor (2019), “A Model of Biased Intermediation,” *The RAND Journal of Economics*, 50(4), 854–882.

- Edelman, Gilad (2019), "Amazon Doesn't Favor its Own Brands—Except when it Does," November 24, <https://www.wired.com/story/amazon-gating-private-labels-antitrust/>, accessed April 2, 2024.
- Egan, Mark (2019), "Brokers versus Retail Investors: Conflicting Interests and Dominated Products," *The Journal of Finance*, 74(3), 1217–1260.
- European Commission (2008), "Guidelines on the Assessment of Non-Horizontal Mergers under the Council Regulation on the Control of Concentrations between Undertakings," *Official Journal of the European Union*, C 265, 18.10.2008, 6–25.
- Farronato, Chiara, Andrey Fradkin, and Alexander MacKay (2023), "Self-Preferencing at Amazon: Evidence from Search Rankings," *AEA Papers and Proceedings*, 113, 239–243.
- Frankenreiter, Jens (2022), "Cost-Based California Effects," *Yale Journal on Regulation*, 39(3), 1155–1217.
- Hagiu, Andrei, Tat-How Teh, and Julian Wright (2022), "Should Platforms be Allowed to Sell on their Own Marketplaces?" *The RAND Journal of Economics*, 53(2), 297–327.
- Johnson, Eric J., Suzanne B. Shu, Benedict G. C. Dellaert, Craig Fox, et al. (2012), "Beyond Nudges: Tools of a Choice Architecture," *Marketing Letters*, 23(2), 487–504.
- Kaneshiro, Ian (2020), "Keeping up with the Amazonians: My Observations on Amazon's Buy Box Algorithm Modifications During the Pandemic," August 7, <https://bit.ly/3tmditX>, accessed April 2, 2024.
- Lanxner, Eyal (2021), "How to Win the Buy Box on Amazon: Ecommerce's Most Exclusive Neighborhood," <https://bit.ly/3h5a21a>, accessed April 2, 2024.
- Lessig, Lawrence (1999), *Code and Other Laws of Cyberspace*, Basic Books, New York (NY).
- Linden, Greg, Brent Smith, and Jeremy York (2003), "Amazon.com Recommendations: Item-to-Item Collaborative Filtering," *IEEE Internet Computing*, 7(1), 76–80.
- Manandhar, Sunil, Kapil Singh, and Adwait Nadkarni (2024), "Towards Automated Regulation Analysis for Effective Privacy Compliance," in: *Network and Distributed Systems Security Symposium*, NDSS 2024, The Internet Society, Reston (VA), DOI: 10.14722/ndss.2024.24650.
- Metaxa, Danaë, Joon Sung Park, Ronald E. Robertson, Karrie Karahalios, Christo Wilson, et al. (2021), "Auditing Algorithms: Understanding Algorithmic Systems from the Outside In," *Foundations and Trends in Human-Computer Interaction*, 14(4), 272–344.
- Mota, Nuno, Abhijnan Chakraborty, Asia J. Biega, Krishna P. Gummadi, and Hoda Heidari (2020), "On the Desiderata for Online Altruism: Nudging for Equitable Donations," *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW2), DOI: 10.1145/3415197.
- Peukert, Christian, Stefan Bechtold, Michail Batikas, and Tobias Kretschmer (2022), "Regulatory Spillovers and Data Governance: Evidence from the GDPR," *Marketing Science*, 41(4), 746–768.
- Sandvig, Christian, Kevin Hamilton, Karrie Karahalios, and Cedric Langbort (2014), "Auditing Algorithms: Research Methods for Detecting Discrimination on Internet Platforms," *Data and Discrimination: Converting Critical Concerns into Productive Inquiry*, 22, 4349–4357.
- Saqib, Mohammed (2023), "17 Most Valuable E-Commerce Companies in the World," January 28, <https://finance.yahoo.com/news/17-most-valuable-e-commerce-171841990.html>, accessed April 2, 2024.
- Smith, Brent, and Greg Linden (2017), "Two Decades of Recommender Systems at Amazon.com," *IEEE Internet Computing*, 21(3), 12–18.
- Sorokina, Daria, and Erick Cantú-Paz (2016), "Amazon Search: The Joy of Ranking Products," in: *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '16, Association for Computing Machinery, New York (NY), pp. 459–460.

- Statista (2023a), “Facebook MAU Worldwide 2022,” <https://www.statista.com/statistics/264810/number-of-monthly-active-facebook-users-worldwide>, accessed April 15, 2023.
- (2023b), “Number of Apps Available in Leading App Stores as of 3rd Quarter 2022,” <https://www.statista.com/statistics/276623/number-of-apps-available-in-leading-app-stores>, accessed April 2, 2024.
- Tagliavini, Michael (2023), “Self-Preferencing and the Concept of Abuse of Dominance: Much Ado about Nothing?” *European Competition Law Review*, 44(5), 201–211.
- Tambe, Nikita, and Aashika Jain (2024), “Top Website Statistics and Trends,” February 23, <https://www.forbes.com/advisor/in/business/software/website-statistics/>, accessed March 28, 2024.
- Thaler, Richard H., and Cass R. Sunstein (2008), *Nudge: Improving Decisions about Health, Wealth and Happiness*, Penguin Books, New York (NY).
- Tsagkias, Manos, Tracy Holloway King, Surya Kallumadi, Vanessa Murdock, and Maarten de Rijke (2021), “Challenges and Research Opportunities in eCommerce Search and Recommendations,” *ACM SIGIR Forum*, 54(1), DOI: 10.1145/3451964.3451966.
- Zac, Amit, Pablo Wey, Stefan Bechtold, David Rodriguez, and Jose Del Alamo (2024), “The Court Speaks, but Who Listens? Automated Compliance Review of the GDPR,” <https://ssrn.com/abstract=4709913>, accessed April 2, 2024.

Abhisek Dash
Max Planck Institute for
Software Systems
Campus E1 5
66123 Saarbrücken
Germany
adash@mpi-sws.org

Abhijnan Chakraborty
Indian Institute of
Technology Delhi
Hauz Khas
New Delhi, 110 016
India
abhijnan@iitd.ac.in

Saptarshi Ghosh
Indian Institute of
Technology Kharagpur
Kharagpur, 721302
India
saptarshi@cse.iitkgp.ac.in

Animesh Mukherjee
Indian Institute of
Technology Kharagpur
Kharagpur, 721302
India
animeshm@cse.iitkgp.ac.in

Jens Frankenreiter
Washington University
in St. Louis
School of Law
One Brookings Drive
St. Louis, MO 63130
USA
fjens@wustl.edu

Stefan Bechtold
ETH Zurich
IFW E 47.2
8092 Zurich
Switzerland
sbechtold@ethz.ch

Krishna P. Gummedi
Max Planck Institute for
Software Systems
Campus E1 5
66123 Saarbrücken
Germany
gummedi@mpi-sws.org