

RESEARCH ARTICLE

Algorithmic attention rents: A theory of digital platform market power

Tim O'Reilly^{1,2}, Ilan Strauss² and Mariana Mazzucato²

¹O'Reilly Media, Sebastopol, CA, USA

²The Bartlett Faculty of the Built Environment, UCL Institute for Innovation and Public Purpose (UCL IIPP), London, UK

Corresponding author: Tim O'Reilly; Email: tim@oreilly.com

Received: 20 November 2023; **Accepted:** 03 January 2024

Keywords: advertising; algorithms; business models; platforms; regulation

Abstract

We outline a theory of algorithmic attention rents in digital aggregator platforms. We explore the way that as platforms grow, they become increasingly capable of extracting rents from a variety of actors in their ecosystems—users, suppliers, and advertisers—through their algorithmic control over user attention. We focus our analysis on advertising business models, in which attention harvested from users is monetized by reselling the attention to suppliers or other advertisers, though we believe the theory has relevance to other online business models as well. We argue that regulations should mandate the disclosure of the operating metrics that platforms use to allocate user attention and shape the “free” side of their marketplace, as well as details on how that attention is monetized.

Policy Significance Statement

Today's big tech platforms can use their power over the attention of billions of users to shape the markets in which they participate for their own benefit, and against the interests of rivals, their users, and firms that depend on that flow of attention. Measures of market power based only on revenues and profit are insufficient for understanding these dynamics. To better regulate these attention-gatekeepers—and their coming AI-powered successors—we must mandate regular and consistent disclosures of the internal operating metrics that measure how they allocate attention, and develop a baseline understanding of what constitutes good behavior and what constitutes abuse.

1. Introduction

The dominant policy narrative that guides the regulation of internet platforms today focuses on user data and privacy (Albrecht, 2016; Srinivasan, 2019). Platforms are said to abuse their market power by taking data from consumers (or “users”) without their permission and using it to manipulate their behavior through personalization. Zuboff (2019) calls this surveillance capitalism.

We argue instead that it may be more productive to understand platform market power and to regulate its possible abuses by measuring the ways that internet platforms control and monetize the *attention* of their users. The fairness or unfairness of the algorithmic systems by which platforms allocate user attention affects not only users but also an entire ecosystem of third-party suppliers (such as websites, content creators, or app developers), as well as advertisers.

As Simon (1971) noted, an abundance of information leads to a scarcity of attention. In the face of increasing information abundance, he predicted that we would use machines to help us better allocate our time and attention. And so it has transpired: information has become so abundant that it defies manual curation. Instead, powerful, proprietary algorithmic systems use the data they collect to match users with the answers, news, entertainment, products, applications, and services they seek.

Tirole (2017, Chapter 14) echoed Simon's idea decades later, noting: "We suffer from too much choice, not too little. Our problem now is how best to allocate time and attention to this plethora of potential activities, trades, and relationships [...] The more the other costs (transportation, customs duties, listing) fall, the more important costs associated with signaling, reading, and selecting become, and the more we need sophisticated platforms to match the buyers and sellers."

Each of the dominant internet platforms is an attention gatekeeper of one kind or another, matching requests from billions of consumers to content, services, and products from millions of suppliers. Despite their differences, these platforms have risen to prominence because each has developed an effective way to efficiently allocate user attention to the most relevant information, products, and services. Google and other search engines promise to find the most relevant web pages from millions of possibilities for each user query; e-commerce sites promise to find the best products available at the best price; social media sites promise to generate a unique, personalized feed of updates from friends; music and video recommendation services promise to deliver a feed that matches a user's taste; on-demand transportation services promise to find the closest driver. And so on. And these platforms make a corresponding promise to the suppliers of content, products, or services (web sites, app developers, merchants, creators, and even other users) on the other side of what is typically a two-sided or three-sided marketplace consisting of consumers ("users"), producers ("suppliers"), and advertisers: that if the supplier provides the most relevant information, they will be rewarded with consumer attention.

In many cases, these information matching marketplaces have proven to be remarkably efficient. Complex, data-driven algorithmic systems act as a kind of proprietary "invisible hand," making use of immense amounts of consumer, producer, and advertiser information to efficiently match supply with demand. When the algorithms are fair, they deliver services to consumers that were previously unthinkable, saving them time and effort in making better choices by providing extraordinarily relevant results despite an overwhelming number of competing options. Suppliers and advertisers find new customers and the ability to transact with existing ones. But these markets have proven to be "winner takes most" (Wu, 2013; Petit, 2020), or sometimes "winner takes all," leaving them ripe for abuse. Once a platform establishes dominance, it is in a position to extract additional time and attention from its users, and economic rents from its supplier marketplace or advertisers, by controlling that flow of attention.

Economists and policymakers have long been concerned about the power of dominant companies to extract economic rents, and there is a growing body of research arguing that an increase in rents is a major contributor to increased inequality, less vibrant entrepreneurial ecosystems, and lower levels of productivity growth and investment in modern economies (Piketty, 2014; Standing, 2016; Ryan-Collins et al., 2017; Mazzucato, 2018; Stiglitz, 2019; Christophers, 2020; Kurz, 2023).

Rents typically reflect control over a scarce factor of production. This control allows its holder to extract profits above a "normal" rate achievable in a competitive market. These profits are not the result of productive improvements that grow the economic pie; they are a reallocation of economic value from one party to another as a result of some kind of market power.

Not all rents represent abuse of power, though. As noted by Schumpeter (2013), innovation—whether protected by patents, trade secrets, or just by moving faster and more capably than the competition—provides an opportunity to receive a disproportionate share of profits until the innovation is spread more widely. A company that continues to innovate can earn disproportionate profits for a long time, especially in a growing market.

During the expansive period of a new technology cycle, market leaders do emerge through innovation, solving new problems, and creating new value not only for consumers but also for a rich ecosystem of suppliers, intermediaries, and even competitors. These market leaders can reach astonishing levels of Schumpeterian profit as they lay waste to incumbents and dominate the new market. But once the growth

of the market as a whole slows, they can no longer rely on the rising tide of new user adoption and their own innovations to maintain that level of profit. At that point, they may turn to more traditional extractive techniques, using their market power to maintain or increase their now-customary level of profits in the face of macroeconomic factors and competition that ought to be eating them away.¹

Companies like Amazon, Apple, Google, and Microsoft have been innovators, and much of the value they have received has been well earned as a return on their investments. But they are also increasingly the beneficiaries of economic rents. But what is the scarce factor of production that allows them to extract these rents? And how do you measure rents in a market where services are offered for free? This article argues that the scarce factor of production is the attention of users, and that rents can be identified by deviations from the best possible attention allocations of which a platform is capable.² These are represented by what in the search engine literature are referred to as “organic” results; that is, by the results chosen as best by the platform’s own search or recommendation algorithms before any self-serving distortions.

These “algorithmic attention rents” are rents in the classical sense. Attention is a factor of production that is limited in supply³ and can see its value appropriated by others than those who supply it.⁴ By virtue of a platform’s dominance in a given attention market, it is able to appropriate an increasing share of the return to “attention”—including by providing lower-quality results, by charging a higher price than what the attention may be worth to those buying it, by forcing ecosystem participants to pay for visibility, or by trying to monopolize vertical product or service markets.

In allocating user attention, the platform is also shaping the allocation of economic value between competing stakeholders on the platform, including itself, its users, its third-party supplier ecosystem, and its advertisers. A platform’s third-party producers compete with each other, and advertisers compete with these producers and other advertisers, for a fixed quantum of user attention. Not only is a user’s attention finite, so too is the narrow window onto abundant information provided by the screen, whose interface design is controlled by the platform. Every user attention allocation can thus lead to a pecuniary gain or harm for a firm, website owner, or content creator on another side of the platform. *Attention allocations drive value allocations.*

This understanding shifts the analysis of a platform’s abuse of market power away from prices. A platform’s dominance can be measured by its ability to shape user attention independently of user preferences,⁵ user inputs, and the relevance of its third-party ecosystem’s information.

Our approach differs from the surveillance capitalism view that “Big Tech” algorithms extract a “behavioral surplus” from users as excess data (beyond service improvement) to manipulate them (Zuboff, 2019). While it is true that platforms collect enormous amounts of data on their users, profit from it, and use it not only for their users’ benefit but also for the benefit of their advertisers, this narrative misses the mark in several important ways. Data is an essential raw material that is aggregated and made useful by internet services, and personalization is often experienced as a benefit by consumers rather than a harm. Drawing a clear line between permissible and impermissible uses of data and personalization is often difficult. Data is ultimately a means to more effective attention allocations, not an end in itself.

In our framing, it is attention that can be extracted in excess of that needed by the platform to earn a normal return on capital. And once that excess attention has been extracted from the consumer, it can be

¹ For evidence on this cycle for Google and Search, see *United States of America v. Google LLC* (2020).

² To be clear, there is no single result that is best for every user, which is why Google, for example, traditionally offered 10 results for the user to choose from. But what we see today is the shrinkage of the range of options offered to users, with the platform making more of the choices for them, often dictated by commercial considerations, not user benefit.

³ Invariant to changes in prices (vertical supply curve), for overview see: Alchian (2017) and Mazzucato (2018).

⁴ Rents may simply reflect the returns received by a factor (Alchian, 2017). But rents can also reflect a deviation from what a factor contributes to production versus what it receives. See, for example, Samuelson and Nordhaus (2010) showing the deviation between wage and marginal revenue product of each worker. For discussion of Neoclassical rent as “excess over opportunity cost” and its reliance on theories of profit, see Blaug (1997, p. 439).

⁵ See, for example, Ricks and McCrosky (2022) on the ineffectiveness of YouTube’s user preference controls.

redirected to extract pecuniary rents from suppliers or advertisers—or to allocate more value to the platform's own information. A surveillance capitalism paradigm ignores that platforms are multisided, such that every suboptimal allocation or action impacts not just users but also the other platform sides too.⁶

Our emphasis on nonpecuniary attention rents being extracted from users in order to extract pecuniary rents from suppliers is in-line with the predictions of Rochet and Tirole's benchmark economic models of platforms (Rochet and Tirole, 2003, 2006). These predict that a monopolist will charge users zero pecuniary prices to maximize profits, if cross-side network effects are large (such that each advertiser benefits considerably from each additional user).⁷ Our article elaborates on what happens next: a monopolist charging users zero pecuniary prices can increase profits further by degrading the quality of attention allocations to users below a competitive level without a loss of revenue that would make such a strategy unprofitable (Begent and Collyer, 2013). Unsurprisingly, this is what is widely observed today: such platforms have become more focused on attracting advertisers than on providing a good user experience, in order to extract excessive profits from advertisers or their producer ecosystems. Doctorow (2023) calls this “enshittification.”

Why users, suppliers, and advertisers not switch to other platforms? One answer is that it is difficult to assemble what Amazon famously called the “flywheel,” in which a critical mass of algorithmically curated content from suppliers draws users, and more users draw more suppliers, in a virtuous circle through which the marketplace provider is able to continuously improve its services.⁸ Data does play a role here. The more users that a platform has, the more data that it can collect about them and the better its algorithmic results can be. That means in practice that the market leaders are sufficiently far enough ahead of the competition that, once they have established market power, they have headroom to worsen the product in other ways without losing users to competitors.

A critical part of the monopolist's toolkit is also to raise switching costs by reducing frictions internally, and raising them externally. For example, free shipping with Amazon Prime encourages users not to shop around, and Amazon's “most favored nation” pricing contracts with its suppliers make it unlikely that lower prices will be found elsewhere (Graham, 2023). While this is not the focus of our analysis, it is a backdrop to any understanding of how a marketplace platform can reduce the quality of results without losing participants.

The identification of a dominant platform's pecuniary rents as being extracted via algorithmically manipulated attention makes it possible to better understand several different types of platform harms, including self-preferencing, excessive advertising, exploitation of third-party ecosystems, and exploitation of user click behavior. A major implication of this work is the need for greater disclosure, to allow regulators, investors, and the public to better observe, measure, and ultimately regulate potential harms stemming from how user attention is allocated.

Our primary goal here is to articulate a theory of platform market power and its abuses in the digital age that serves as a foundation for future work. In a companion paper, “Amazon's Algorithmic Rents” (Strauss et al., 2023), we take a deeper look at the legal and policy application of these ideas to Amazon's third-party marketplace. And in “Behind the Clicks: Can Amazon allocate user attention as it pleases?” (Rock et al., 2023), we demonstrate an approach to measuring algorithmic attention rents in Amazon's marketplace. This research is part of a broader effort to map modern economic rents (Mazzucato et al., 2023).

⁶ Hovenkamp (2020, p. 1961) notes that “power assessment on two-sided platforms requires considering the reactions that occur on the opposite side.”

⁷ Rochet and Tirole (2006) note that “a factor that is conducive to a high price on one side, to the extent that it raises the platform's margin on that side, tends also to call for a low price on the other side as attracting members on that other side becomes more profitable.”

⁸ In Wilkes (2012), Amazon executive Jeff Wilkes explains in an internal presentation to Amazon staff how Jeff Bezos first came up with this idea. There is now an extensive management literature explaining the concept.

2. A theory of rents in digital markets

Economists see prices as the coordinator of economic activity: they are the sinews of the market's invisible hand. Prices are thought to optimally allocate resources among competing ends when they reflect the dynamic information (preferences and scarcities) contained in the billions of daily decentralized interactions between demand and supply. As Hayek (1945, p. 1) notes, decentralized price formation solves “the problem of the utilization of knowledge which is not given to anyone in its totality.” Market-driven price formation is superior to any centralized mechanism of coordination of economy activity because it ensures that “fuller use will be made of the existing knowledge” contained in the economy (Hayek, 1945., p. 2).

In neoclassical (marginalist) theory, perfect information exists, such that prices are optimal because they reflect the subjective utility evaluations of consumers.⁹ By contrast, in Simon's view of decision-making and in more recent work in behavioral economics, perfect information does not exist. Instead, consumers and producers make decisions that are shaped not only by human limits and biases but also *by the institutions that shape the information that is available to decision-makers* (Simon, 1997). And today, what makes information imperfect is often not that there is too little of it, but too much, and the institutions that help us manage that abundance have extraordinary power to shape our decisions. Internet platforms change the institutional context and challenge the conventional view of decision-making in markets in several ways:

1. In informationally complex markets, platforms transfer much of the work of decision-making from humans to machines. Internet search, e-commerce search, social media feeds, and other algorithmically managed recommendation engines are examples of such machines.
2. These algorithmic machines are often used to match the supply and demand for non-priced goods and services, or those that are not individually priced. For example, in a free, ad-supported service such as Google Search or Facebook, or a subscription service such as Netflix or Spotify, consumers are matched with suppliers of information without considering price as a factor. Matching instead relies on other non-price factors to gauge the objectives and preferences of the consumer, the quality of the products or services on offer from suppliers, and even the reliability and reputation of the suppliers themselves. As Google founders Larry Page and Sergey Brin noted (Brin, 1998), a platform such as Google Search offers *objective* rankings based on something as seemingly *subjective* as optimizing for “relevancy.”¹⁰ Much like the decentralized markets celebrated by Hayek, platforms work their allocative magic by processing signals based on millions of decisions taken by other users on the internet and in real life, combined with data, both expressed and implied, on the user's personal preferences.¹¹ Collecting more data is an essential part of what makes these systems work. But in processing these data to produce a relevancy ranking for user search or recommendations, the algorithmic system takes on the role of the invisible hand and works *either to preserve the competitive process or to distort it*.
3. Even when price is a factor, as in e-commerce or an advertising marketplace, the platform's proprietary matching algorithms *internalize* and *centralize* the otherwise decentralized market mechanism. Furthermore, this internalized and centralized market is opaque. Rather than providing

⁹ Foley (2008, p. 160) notes that in neoclassical economic theory, “the quantity of the various commodities available to society has to be taken as given, so that their relative scarcities can determine marginal utilities and hence price.” Market prices then are “exactly analogous to the ratios of marginal utilities that an individual equalizes in making a rational allocation of resources.” Continuing (Foley, 2008, p. 171): “The idea that the goal of economic activity is the satisfaction of individual consumers is deeply rooted in the structure of marginalist thought, which sees subjective utility evaluation as the regulating factor of price and value.”

¹⁰ At the core of PageRank, the search quality algorithm developed by Google founders Larry Page and Sergey Brin when they were graduate students at Stanford, was the belief that although “*the importance of a Web page is an inherently subjective matter*,” an algorithm could “*objectively and mechanically*” rank them through “*effectively measuring the human interest and attention devoted to them*” (Brin, 1998).

¹¹ For a view of this process at Amazon, see Sorokina and Cantu-Paz (2016). While the actual algorithms described in the paper have undoubtedly been updated since the paper was written, it gives an invaluable view of the algorithmic training process.

explicit information to consumers about the basis for the ranking of products and services, an algorithmically generated ranking *implies* much of that information, with the user expected to trust the rank ordering provided by the platform. The platforms control the presentation of information, and their algorithmically populated interface designs become the context for user decision-making.

Algorithmic attention allocations thus supplement—and at times supplant—traditional markets as the key institutional mechanism coordinating economic activity and shaping the terms on which exchange takes place online. The resulting algorithmic systems decide the winner among different producers whose information is competing for the user's attention. They not only facilitate the effective delivery of the platform's information services to users but the monetization *between* platform sides (e.g., advertisers and users).

Because internet platforms have effectively internalized the market mechanism, their algorithmic allocations tend to reflect the degree of competitiveness within and between platforms. In a competitive market, the platform has a strong incentive for its algorithms to be fair; once they have market power, platforms are liable to make allocations that are self-serving.

Attention allocations involve design choices and trade-offs because attention is finite and consumable.¹² For example, a platform allocating more top screen space to advertising information can prevent the user from spending attention on more relevant organic results, leading to poorer choices. A platform providing information directly in response to a query rather than directing traffic to a third party website might yield benefit to its users even as it reduces benefit to third party suppliers.

Internet platforms are in a unique position to explore and optimize these trade-offs for their own benefit due to their access to real-time data on the participants from all platform sides. With millions of users repeating the same search, or responding to the same recommendations, the platform is able to run statistically meaningful A/B tests on thousands of different algorithmic weightings and design options. In 2022, Google claims to have run more than 800,000 search experiments, which led to more than 4,000 changes to Search (Google, 2023a).

2.1. How the limits of human cognition enable algorithmic authority

Simon's (1978a, 1995) information processing paradigm focuses on how humans make decisions in the real world: "compatible with the access to information and the computational capacities that are actually possessed by organisms, including man, in the kinds of environments in which such organisms exist" (Simon, 1955, p. 99).

Human computational capacities are limited, as is their time.¹³ These "hardware" limits help explain real-world behaviors, which tend to follow *heuristics*—informational shortcuts and strategies that allow humans to make reasonable choices in everyday complex environments (Simon, 1978b, p. 12, Simon, 2017). Heuristics reflect human "satisficing behavior" which aims for "good enough" outcomes,¹⁴ based on many unknowns. This contrasts with the economic assumption that humans "optimize" to achieve the "best" solution based on the known outcomes from every action. Simon's emphasis on the decision-making *process* contrasts strongly with the neoclassical focus on (equilibrium) *outcomes* by unconstrained actors.

Simon's insights are echoed by contemporary behavioral economics. Tversky and Kahneman (1974, 1981) and Kahneman (2011) posited that humans have two decision-making modes: what they called System 1, for decisions that need to be made quickly, and System 2, for those that depend on careful rational analysis. According to Kahneman (2011), 98% of human decision-making relies on

¹² Kahneman (2011) notes that "The often-used phrase 'pay attention' is apt: you dispose of a limited budget of attention that you can allocate to activities, and if you try to go beyond your budget, you will fail." Competition and Markets Authority CMA (2022) outlines how online choice architectures can lead to user harm.

¹³ Simon (1978b, p. 12) notes that "the scarce resource is computational capacity—the mind."

¹⁴ For Sen (2017), satisficing is "a target level of achievement" but can be interpreted as optimizing.

System 1. While these two systems are appropriate for different types of activity (immediate response to a threat, for instance, versus long-term planning), they may be applied inappropriately. System 1 is particularly subject to various cognitive biases, such as anchoring bias, by which the first piece of information presented to us frames our judgment of additional information.

Studies of actual human behavior on internet platforms bear out the predictions of both Simon’s “information abundant/attention scarce” model and behavioral economics. The majority of clicks tend to go to the first few results displayed near the top left of the screen (Keane and O’Brien, 2006; Craswell et al., 2008)—even if they will vary by some unknown amount as these results deteriorate in relevancy. This is called positional bias (Joachims et al., 2017).¹⁵ This is a form of anchoring bias that has been well documented by search engine optimization consultants. For example, a 2022 study of more than 4 million Google Search results pages found that more than 54% of clicks go to the first three organic results. Only 0.63% of users click through to the second page of the results (Dean, 2023). A 2018 report from the web tracking firm Jumpshot (2018), based on a large sample of users, noted the same behavior on Amazon: user clicks are concentrated on the first few rows of product results. In our own study, we found that 78% of the most clicked product listings are positioned within the first two rows of the search results (Rock et al., 2023). Perhaps even more strikingly, in 2022, a study by e-commerce consulting firm Feedvisor (2022) reported that 32% of shoppers simply opted to buy the first product listed on a search results page.

The value proposition and business model of the modern internet is thus entirely in line with the premise that users are *not* perfect “hedonic calculators” (Robinson, 2001). And this is why users are able to benefit from intermediators such as Google or Amazon. Time savings accrue to users when they click on algorithmically ranked results based on position, trusting that the platform has done the work for them, rather than having to evaluate intrinsic and latent product qualities themselves. This behavioral heuristic saves users enormous amounts of time and cognition in decision-making. Following a platform’s algorithm can lead the user to make better choices, since it effectively secures what Simon (1997) calls a degree of expertise in the decision-making process.

The impact of that expertise is borne out by actual observed user behavior. Clicks go to the top results not just because of human positional bias, but also because the platforms have traditionally worked very hard to make the top results the *best* results. Google went so far as to patent a ranking factor that it called “the long click” (Lopatenko et al., 2015). The patent posited that a “short click” (i.e., when a user clicks on a result and comes right back and clicks on another result) indicates dissatisfaction with the result, while a long click (when the user goes away and does not come back) represents success. Along with countless other factors measured across millions of searches, long clicks could be used to raise or lower the rank of search results, with the goal that the first result gets the long click.¹⁶ The success of these systems is precisely why they have gained the trust of users, and why that trust, once earned, can be abused.

A platform’s power to get users to click on its algorithmic outputs is imperfect, since clicks and views are also influenced by the broader relative prominence and attractiveness of the results as a whole (Yue et al., 2010), including their relative informativeness (Keane and O’Brien, 2006; Craswell et al., 2008). Yet, the power that trusted platforms have to drive user attention, and in turn clicks or views, is immense. Attention directed to new information, including advertising, can lead to immediate, almost frictionless action. A purchase, a view of an addictive TikTok, YouTube, or Facebook video, a path down a rabbit hole of attention consumption, is just a click away. In feed recommendation systems such as those offered by social media systems, even clicks may not be necessary, as platforms increasingly use autoplaying videos to capture user attention, which thus becomes “opt out” rather than “opt in.”

The attention-shaping power of online platforms is a reminder of Simon’s key assertion that organizations or institutions—he often used the terms interchangeably—shape individual behavior by

¹⁵ For the other main theoretical model of clicks, the cascade model, see Craswell et al. (2008). For critique, see Wang et al. (2018).

¹⁶ Perhaps because of the focus by regulators on harms from the collection of user clickstream data, Google denies that such data has been used as a formal ranking factor. In any case, it now notes that improvements in full text analysis with AI makes user clickstream data less important. See Goodwin (2023).

determining “the environments of information in which decisions are taken” (Simon, 1997). Institutions help consumers make decisions by establishing a division of labor in “the process of deciding,” which is just as important as a division of labor in the processes of “doing,” argues Simon. He elaborates on these ideas in *Administrative Behavior* (Simon, 1997, Chapter 7), focusing on the concept of authority in securing an effective division of labor in decision-making within the organization:

‘Authority’ may be defined as the power to make decisions which guide the actions of another. [...] That is, he holds in abeyance his own critical faculties for choosing between alternatives and uses the formal criterion of the receipt of a command or signal as his basis for choice.

This passage could be read today as applying to algorithmic recommendations made by a dominant digital platform to a user, just as much as a recommendation coming from a senior manager or expert to a worker, as outlined by Simon.

2.2. *How algorithmic authority enables attention rents*

Attention rents occur when a platform abuses its algorithmic authority and exploits its role as a trusted intermediary to direct user attention (clicks or time spent) to suboptimal—often sponsored—information. At its core, these rents exploit users’ positional bias in how they click, or what they view, by placing this suboptimal information in users’ core attentional zone. This positional bias relies heavily on the suboptimal information trying to replicate, or leverage, the authority of a platform’s organic algorithmic result (so-called “trust bias”) (Keane and O’Brien, 2006; Keane et al., 2008). The information may be embedded between the optimal organic results, or increasingly, may displace them.

One implication of this algorithmic power in allocating user attention is that it provides the platform with considerable leverage over both parties (suppliers and advertisers) looking to access user attention—the platform’s core commodity. The platform can ensure that paid (advertising) and organic results directly compete with one another for user attention (especially when placed “above the fold” as it was called in the days of newspapers, but today might be called “before the scroll.”)¹⁷ The *attention rent* is levied on the users—to get them to allocate attention to more profitable content for the platform—in order to extract a *pecuniary rent* from the supplier or advertiser side of the dominant platform’s multisided market. This rent can help create an above-normal return for the platform, especially when combined with other charges and fees already placed on its ecosystem.

Thus, greater monetization of users entails a double-sided process, impacting not just a platform’s users but also third-party firms and advertisers, as the nature of content shown to users changes. Information from some third-party firms gets demoted by additional paid or addictive content, and the demoted content must adjust to the new algorithmic optimization to survive. Potential advertisers receive more priority attention space and, in turn, greater incentive to advertise. And users must consume more paid-for content. What happens on one side of the platform impacts the other(s) (Hovenkamp, 2020)—not necessarily from network effects,¹⁸ but by virtue of the fixed screen space (i.e., attention sphere) over which the platform sides compete.

2.3. *How attention rents become pecuniary rents*

Contemporary theories of platforms focus mostly on externalities arising from network effects. This helps explain why platforms provide services to users for free, in order to generate network effects and/or economies of scale. This also explains pricing structures on platforms, that is, how the pricing burden falls

¹⁷ Petrescu (2014) notes that “We found that on average, the presence of ads on a search results page caused the organic CTR of the first position to drop by 30%—from 25.7% organic CTR in the absence of ads to 17.9% CTR when ads are displayed.” This effect is likely even more significant today, when ads have become much more prominent, occupying the positions most likely to draw attention, and less clearly marked as advertisements.

¹⁸ Same-side or cross-platform.

between the two sides of the platform. For Rochet and Tirole (2006), the optimal price structure on platforms is set indirectly to optimize total profits from both sides. This means that pricing on one side of the platform should take into account the impact it would have on participation—and in turn profits—on the other side of the platform.¹⁹ This tends to create a highly imbalanced pricing structure, whereby users pay no pecuniary price for the platform and the price burden instead falls on third-party firms or advertisers. Because advertisers value an additional user considerably, and because users prefer free services to paying directly for them, it is often more profitable for the platform to charge users indirectly, through their attention.²⁰

But this theory says less about what mature platforms, which have for some time been charging users zero (or low) prices, do next. When users face zero pecuniary prices, and their attention is monetized on the other side of the market, increasing profits initially come from gaining more users. But over time, as user growth slows, increasing profits from user participation become a function of the quantity of advertising. Weyl (2010, p. 1643) goes so far as to see “the platform’s problem as [first] choosing participation rates on the two sides rather than the prices supporting this allocation.” User responsiveness to advertising, often measured in real time by the platform, becomes key for determining the optimal business model mix such that (Behringer and Filistrucchi, 2015, p. 293): “In two-sided markets, quantities on one market side are functions of prices on that market side and quantities on the other market side.”

It follows that algorithms coordinate platform sides not just by setting the price level as Tirole notes, but also by controlling attention allocations. If the participation of sides is relatively fixed (due to user stickiness and third-party lock-in), then extracting monopoly rents from one side of the platform requires adjusting attention allocations on the other side(s), since users must consume greater quantities of the more profitable content.²¹

In a multisided digital platform, increasing profits or profitability for the platform thus tends to require changing the information content shown to users on the screen (Behringer and Filistrucchi, 2015): *What* information outputs a platform’s organic algorithms optimize for (such as recommendation algorithms optimizing for more sustained engagement); and/or the algorithmic *mix* of outputs (such as more advertising and fewer organic outputs) are essential ingredients to a platform increasing user monetization and extracting more profits from its ecosystem of firms or advertisers. Without changing the type of content shown to users, the opportunities for user monetization by the platform’s third-party firms or advertisers are limited to the traditional solution of raising prices (e.g., for advertising, subscriptions, or other fees).

These changes in information can impact value allocations if the platform chooses to algorithmically allocate more attention to paid or otherwise self-serving content. This is done constantly, by changes in the relative screen presentation of results, by changes to the underlying algorithmic ranking systems, and, as asserted in the US Department of Justice complaint against Google, by internal controls that “tweak” the results to get the desired outcome (*United States of America v. Google LLC, 2020*). Measured user behavior provides a rich stream of data by which the platform can assess the impact of changes and make further adjustments to reach its objectives (“operating metrics”).

In traditional markets, in which goods and services are directly traded for money, prices persistently higher than “normal” are seen as a sign of excess market power and rent extraction. In attention markets, the primary expression of *algorithmic market power* is the ability of the platform to profitably direct

¹⁹ Though they only take into account indirect network effects, rather than within-side network effects.

²⁰ Rochet and Tirole (2006) note that “a factor that is conducive to a high price on one side, to the extent that it raises the platform’s margin on that side, tends also to call for a low price on the other side as attracting members on that other side becomes more profitable.” The dynamics are similar to those in a competitive bottleneck model, as discussed in Jullien et al. (2021). Users on one side of the market primarily engage with just one platform (“single-home”), while users on the other side use multiple platforms (“multi-home”). Once a platform has successfully attracted the single-homing side, it has greater market power when dealing with the multi-homing side, and that is typically where it aims to monetize or profit. Advertisers (multi-home) place ads on multiple search engines, but users typically use one search engine (single-home) for most of their queries.

²¹ For Newman (2015), “customers of zero-price products paying for those products, primarily by exchanging their attention, information, or both.” For critical discussion, see Strandburg (2013).

user attention, and to produce information allocations, that are to an appreciable degree independent of consumer preferences, competitor information relevance, and its users' explicit search inputs. Algorithmic rent is part of a platform's wider exertion of market power when it can make inferior attention allocations without losing sufficient sales, or users and third-party suppliers, to make such an allocation unprofitable (Begent and Collyer, 2013; *Federal Trade Commission, et al. v. Amazon.Com, Inc.* 2023).

A core notion explored in our research is that an attention rent exists when there is *a deviation from the best attention allocations of which the platform is capable*, and on whose promise the platform drew its users in the first place. In a competitive market, platforms win by providing the best possible results to consumers. Deviations from this standard generally occur once the platform has cemented its dominance and hold over user attention. Because it is extremely difficult and costly to deliver high-quality search or recommendation results, the market leaders are frequently far more capable than their competitors, which gives them additional room to worsen their results without losing customers.

Algorithmically driven attention allocations by the platform, matching users to suboptimal information, can create rents in the sense of “returns” to a factor of production—here “attention” (proxied by screen space and user time)—which is fixed in supply, largely invariant to changes in prices, and exploited for the platform's own profit (Blaug, 1997; Alchian, 2017). A platform is increasing profits not from better matches (productivity improvements), but from information matches that are more profitable for itself, or from higher matching fees.

This pattern of rent extraction by exploiting the relative proportion and position of advertising and organic results is exactly what has been observed in practice. In its earlier years, Google Search result pages consisted of a list of ten organic search results (often referred to as the “ten blue links”) and a snippet of content from the destination site. These organic links were framed by ads²²: three above the organic results, with additional ads in a column to the right of the organic results. The ads were clearly differentiated from the organic results in a distinct color block. A remarkable series of screenshots recorded by Dutch search engine consultant Blacquière (2014)—updated by Marvin (2020)—shows how ads took over more of the Search screen and became increasingly hard to distinguish from organic search results on Google Search.

In 2016, the right-hand column of ads was removed completely (Kim, 2016), and today, ads are nearly indistinguishable from organic results, interspersed between the organic results, and far more of the most favored space at the top of the page is given over to ads. Additional space is taken up by Google's own content, often displayed in a carousel of images, or in an informational panel that Google refers to as a “OneBox” (O'Reilly, 2019), or most recently, in an AI-generated answer to the user's query.

Clearly, internet platforms are well within their rights to monetize their offerings. When services are free to users, someone has to pay the bills. The question, though, is what level of monetization is justified to recover costs and earn a fair return on investment, and when does it become excessive? When is advertising of value to the supplier ecosystem and to users, and when does it become a source of extractive rents? It is not always easy to determine the answers.

At Amazon, however, the answer to the question of whether monetization has become extractive appears unambiguous. Almost all organic recommendations,²³ which helped make Amazon's digital marketplace revolutionary, have been replaced by purely paid-for recommendations or hybrid recommendations, for example, “Trending now—Sponsored,” or “Highly rated—Sponsored” (Kaziukėnas, 2021). E-commerce research firm marketplace pulse estimates that of the first 20 products a shopper sees when searching on Amazon, only 4 are now organic results (Kaziukėnas, 2022). Because of how paid results are placed, it can sometimes take scrolling past three browser windows worth of search results to get to the fifth organic result (Kaziukėnas, 2022).

²² On pages of interest to advertisers—a large percentage of Google search result pages carry no ads at all.

²³ “Customers who bought this item also bought,” “Customers who viewed this item also viewed,” “Frequently bought together.”

Our own analysis of Amazon’s ability to allocate value (i.e., clicks) between competing products in its third-party marketplace through the algorithmic arrangement of search results shows that a product listing shifted to a sufficiently higher “attention share” position will receive more clicks, regardless of its relevance or price. A product listing in the bottom 10 for relevance, but the highest percentile for “attention share,” is as likely to be clicked as a top 5 most relevant product in an average position. This dynamic enables Amazon to trade-off the relevance of a search result for its screen position, through advertising, while maintaining a high click-through rate (Rock et al., 2023).

Users may be harmed when a platform replaces organic results with advertising because they may see fewer relevant options and may be directed away from better products or lower prices. Organic outputs and advertising both aim to be clicked on or looked at, but the ranking of the former is optimized for intrinsic relevance among all eligible information sources, while the latter optimizes directly for clicks *among those firms that bid for user attention*. While the most relevant ads can be as useful to consumers as the organic results, showing more ads means diving deeper into the inventory of ads and potentially showing worse ads that are not as useful.

The harms to Amazon’s marketplace suppliers from replacing organic content with ads are far more immediate. In a world where gatekeepers’ principal benefit to consumers is algorithmically curated access to a rich ecosystem of suppliers, and where their benefit to suppliers is algorithmically curated access to a huge market of end users, organic search results are the coin of the realm and fairness is of the essence. Paid results standing in for (rather than supplementing) organic results are the equivalent of a debased currency. When advertising replaces organic search results, the supplier ecosystem must now pay for visibility that it once earned through product quality and reputation signals.

As advertising dominates more of the screen, it has become a barrier to entry²⁴ for merchants wanting to sell on Amazon—a tax on top of referral and fulfillment fees (Morrison, 2021). “There’s fewer organic search results on the Amazon marketplace page, so that increasingly means the only way to get on the page is to buy your way on there,” said Jason Goldberg, chief commerce strategy officer at advertising megafirm Publicis (Palmer, 2021).²⁵ The result, according to Quartile (2018), a major AI-powered advertising platform, is that Amazon is now a “pay-to-play” platform for top screen positions. Three-quarters of all sellers on Amazon now “choose” to advertise (Mileva, 2022). For small- and medium-sized sellers, the figure is 79% (Jungle Scout, 2022).

This is not advertising like that offered by Google or Facebook or traditional media, which rides as a passenger on a current of attention focused on information or entertainment, but a kind of *Hunger Games*-like competition between merchants to capture the purchases of people who *are already looking to buy*. It is a zero-sum transfer of attention and value *between* sellers (Gornyi, 2023), which simultaneously increases revenue for Amazon but without “necessarily growing the sales volume” (Kaziukėnas, 2021). An attention boost to one seller necessarily comes at the expense of another.

Moreover, unlike in traditional media advertising, these ads appear in the actual decision-making interface *with an almost identical appearance to the organically recommended product choices*. The majority of Amazon marketplace ads simply duplicate the organic listings and appear on the same page, often adjacent to them, offering no additional information to consumers. And with ads and organic listings competing for user attention on the same search results screen, Amazon exploits users’ positional bias, and puts the ad rather than the organic result in the position most likely to be clicked on, thus extracting a fee from the supplier while providing no added benefit.²⁶ This behavior demonstrates exactly what we mean when we say that Amazon uses attention rents extracted from users to extract corresponding pecuniary rents from its supplier ecosystem.

²⁴ For Petit and Teece (2020), “Once they have led ecosystem partners to a strong, steady adoption outcome, platform leaders might be tempted to tweak the contractual or technological rules to their own advantage.”

²⁵ Moreover, (though Amazon denies this) since sales volume for a product is one of Amazon’s organic ranking factors, failure to advertise may drive down the organic rank as well (Mitchell, 2023).

²⁶ In Rock et al. (2023), we find that 95% of the top three most clicked ads are simply duplicates of product listings.

Through this zero-sum competition, Amazon has fostered both higher ad prices and lower return from ad spend (Jungle Scout, 2022; Soper, 2023), leading to a rent transfer from third-party firms to the platform itself. By 2022, advertising had become a highly profitable \$37.7 billion business for Amazon (Amazon.com, 2022, p. 67).²⁷ Meanwhile, the average cost per click on Amazon ads doubled from \$0.56 in 2018 to \$1.2 in 2021 (Business of Apps, 2022). Average cost of spend was 30% according to Adbadger, meaning that 30 cents now has to be spent on ads to drive \$1 of sales (The Badger, 2023).

Interestingly, as TikTok begins to mount an e-commerce challenge to Amazon, it is relying on new forms of algorithmic earned attention to attract and engage users. As with its social feed product, it uses viral cascades of user attention to surface popular content. In the case of e-commerce, that is the discovery of so-called “dupes,” lower-cost products that have the same or better features than higher cost branded products (Barinka, 2023). That is, TikTok is demonstrating once again that finding new signals that provide better organic algorithmic results provides a competitive advantage in acquiring and retaining users, and that algorithmic rents tend to be extracted by platforms only once they have cemented their dominance.

Distortion of the best algorithmically chosen results can also be seen in social media systems. Facebook, X (formerly Twitter), and Instagram all began by matching users with a unique feed of content from other users that they have chosen to follow, whether they be friends, celebrities, or news sources. During the growth period, when user acquisition is the paramount goal, these platforms align their algorithmic selections with this user promise.

In this early stage, sorting of posts in the feed is typically in reverse chronological order (newest to oldest), or through a social graph (from friends). Later, other possible arrangements are tried, such as sorting by posts with the most engagement. Eventually, the sorting is fully given over to the algorithmic recommendations. At first, users are given some control, allowing them to revert, for example, to a chronological sort, or to preference posts from those marked as friends, or from those the user has chosen to follow. Over time, these expressions of user preference are made harder to find or disappear entirely. Meanwhile, the ad load goes up, also making it harder to find and enjoy the posts that the user originally signed up for. Moreover, “dark patterns,” such as automatically switching the user to a feed of recommended videos (e.g., Instagram reels) as soon as they finish viewing a video posted by a friend, are also used to help override the user’s preferences. From there, it is a slippery slope to shaping the content shown in such a way as to extract additional attention from users by showing them whatever content will drive the most engagement.

While we have not done empirical research to measure algorithmic attention rents in social media, we believe that the notion of algorithmic attention rents provides fruitful avenues for further research and analysis of social media and other recommendation systems as well as for search-based systems.

3. Harms from algorithmic rents

While it is harm to consumers or to competitors that is most likely to draw the eye of regulators, many of the harms from algorithmic attention rents may fall more heavily on the supply side of a platform’s ecosystem.

²⁷ To put this information in context, it should be understood that Amazon reports its total e-commerce business in two different ways. In its original “first party” model, it purchases inventory from suppliers at a discount and resells it to consumers, profiting from the spread between cost and selling price, less its own operating costs, including warehousing, fulfillment, and delivery, just like other retailers. In its “third party marketplace,” suppliers pay a fee to have their product listed on Amazon, and additional fees for warehousing and fulfillment (a total of \$117.7 billion in 2022), but continue to hold the inventory risk. Starting in 2014, advertising was added as a new line of business. Advertising is not reported as revenue from third party seller service, but as a separate business line. However, most observers attribute the vast majority of Amazon’s advertising revenue to the third party marketplace as Amazon has replaced organic recommendations and search results with paid-for advertising results. In private conversations with one of the authors, the best that Amazon insiders could say of other sources of advertising, such as from Amazon Prime TV, was that they were “small, but growing.”

Thompson (2014, 2023) notes that in order to provide ranked results or recommendations to its users, a platform must aggregate information—pooling and digitizing supply, and commoditizing it in the process. This aggregation reduces the power of the platform’s suppliers, and makes them dependent on the fairness of the algorithmic rankings provided by the platform.²⁸ It is these suppliers, not just consumers, who are harmed when information allocations are distorted from their optimally relevant competitive level.

To understand why clearly defining the role of the platform as intermediary and aggregator matters, it is important to take an ecosystem view of the total investment in value creation, rather than attributing all of the value creation to the platform. Any notion of fairness depends on this ecosystem view: without websites, there would be no need for Google Search; without merchants, no Amazon; without app developers, no App Stores; without users creating content as well as consuming it, no social media. When suppliers are harmed, users too will be harmed over the long run.

These ecosystems of value creators depend on the platform’s algorithms for what in the industry is called “earned attention.”²⁹ When the platform exempts its own competitive content or services from its own algorithms, or when it displaces organic results with paid results, the ecosystem suffers a loss of incentive and reward for continuing to produce value. Eventually, this loss of value affects both users and the platform itself, as the whole virtuous circle of creation, aggregation, and curation breaks down. That which has been earned is appropriated by the platform.

Advertisers can also provide value to consumers, especially when ads are well targeted to their interests. This value too can be algorithmically misappropriated by the platform. Targeted advertising commoditizes attention, personalized by demographics, location, and interests, to be auctioned off at scale and algorithmically matched billions of times per day with content viewed by consumers. Advertisers bid for that attention, but what they get may not be what they pay for if the platform places ads near inappropriate brand-damaging content (Internet Advertising Bureau, 2020), places display ads (which advertisers pay for by a count of impressions [i.e., inclusion in pages being displayed] rather than clicks) in locations far below the scroll where they are unlikely actually to be seen by users, or doesn’t properly police problems such as click fraud, viewing by bots rather than humans, autoplaying videos that are not actually chosen by the user, and so on. Hwang (2020) calls this ad quality problem “the subprime attention crisis.”³⁰

A fair reward for value created is a useful framework within which to evaluate harms from algorithmic allocations. When considering algorithmic attention allocations, it is helpful to ask who wins, who loses, and who decides.³¹ This can be a complex calculus. Examples from Google Search, Amazon marketplace, and social media feeds, which we explore below, show the range of questions that can be explored using the analysis of algorithmic attention rents in ad-based businesses.

²⁸ Thompson’s “aggregation theory” makes a distinction between what he calls aggregators, which collect and manage information and access to a marketplace, and true platforms, which provide capabilities that provide a foundation that third parties can build on. For example, Apple’s iPhone App Store is an aggregator, while its iOS operating system is a platform. Amazon’s e-commerce marketplace is an aggregator, but Amazon Web Services is a platform. Fulfillment by Amazon, Amazon’s suite of services for third party merchants can also be considered a platform, but the Marketplace itself is an aggregator. Google Maps is a platform, while Google search is an aggregator. While this is an exceedingly useful distinction, the term “platform” is so widely used in the literature that we have adopted it here.

²⁹ Hugonenc (2022) notes that “Earned attention is far more effective at delivering results than when attention has been acquired by force. When consumers voluntarily view an ad, there is a significantly higher impact on brand metrics than ads that have been forced upon them, whatever the length of time they view the ad for.”

³⁰ There are many types of algorithmic rents that can be extracted from advertisers that are not attention rents, but purely extraction of revenue via unfair manipulation of a platform’s algorithms. For example, per the 2023 United States Department of Justice complaint against Google, the company has “designed its ad auction algorithms to include adjustable variables internally known as ‘pricing knobs’; Google then tunes the variables to increase advertiser prices.” The discovery process during the trial uncovered damaging emails showing that Google used these pricing knobs to make up a shortfall in Google’s revenue projections due to a decline in search volume. See Trial Exhibit—UPX0522: (*United States of America v. Google LLC*, 2020).

³¹ This is also the framework adopted in environmental economics (see Boyce, 2022).

3.1. Google Search: The market shaping power of attention allocations

Google does make a substantial effort to find the most relevant advertisements (Google, 2023b).³² In the pay-per-click (PPC) ad system used in Google Search (“Adwords”), the company has strong incentives to find the ads that are most likely to draw clicks. In display advertising formats (pay-per-view), in which Google is a player through its algorithmic placement of ads on third party websites and its control of the ad exchanges in which advertisements are algorithmically bought and sold, Google’s incentives are less clearly aligned with those of its advertisers.

Even in the PPC system of Adwords, though, the pool of possible results comes from those willing and able to pay. And that strongly favors some firms over others. A good example may be found in a search such as “buy new tires.” In current designs, ads populate the best screen positions, and those ads typically come from large national chains and online sellers (O’Reilly, 2022). Local merchants are shown on a map far down the page, requiring scrolling to see them. The map does provide unpaid business listings that allow users to learn more about each of the merchants, including their opening hours, address, and phone number. This is valuable to both users and to merchants even though it does not take the form of a traditional web link. But this information is far harder to find today than it was 10 years ago. This imposes a time cost on users and a visibility penalty on suppliers, such that they must now advertise to gain attention that they had previously earned by signals such as quality and a location convenient to the user.

Note the *market shaping* power on display. As consumers increasingly rely on search engines to find products, Google is essentially deciding that those most willing and able to pay are more deserving of consumer attention than those that relevancy factors such as location or reputation suggest. This example highlights the responsibility of platforms to think deeply about the consequences of their attention allocations. In short, the trade-offs are not always obvious, and weighing the harms to any party requires careful analysis.

3.2. Amazon marketplace: Advertising as extractive rent

While celebrated by shareholders and analysts as a triumph of business strategy, earning the company enormous profits, Amazon’s marketplace advertising business seems to provide little benefit to either users or suppliers (who, in the case of Amazon’s third party marketplace, are also its advertisers). Users see fewer of the results that Amazon’s own algorithms have calculated as an ideal match for their search, and instead see ads, whose relevance is often lower.

In Rock et al. (2023), we compared the organic rank of a search result with the ad rank of the same result.³³ We found that on average an ad boosted the visibility of a possibly inferior (by the judgment of Amazon’s algorithms) product by between 5 and 50 positions, with a median boost of 17 positions. We also found that occupying high value screen positions is more important for sponsored than organic listings, suggesting that users are not finding the most relevant products in the top ranking spots.

There may be some benefit from advertising to merchants trying to raise the visibility of new high quality products (which may have fewer reviews and less history of being clicked on and purchased), but for the most part Amazon’s ad business appears simply to be an additional fee levied on merchants rather than providing added value to them. It raises their cost of doing business, which may be passed along to consumers in the form of higher prices. We found that the top-3 most clicked advertised products are about 17% more expensive than organic ones (\$19.3 vs. \$16.5) and one-third less relevant (organic rank of 4 vs. 3.)

³² Its ad auction takes into account various measures of ad quality before allowing participation, though it admits that ad quality is not actually considered in the auction. Instead, a low ad quality score may mean that an advertiser is prohibited outright, or charged more to win the auction (see Google, 2023b).

³³ Note that in some systems, such as in Google search results, the ordinal position is easy to determine. On Amazon, which uses a grid format for results, ordinality is more complex. It can be read left to right and top to bottom, but it may be more productive to develop an attention-weighted model of the spots most likely to be viewed and clicked on by the user, on the assumption that Amazon has made a similar estimation of the relative value of each spot. This is the approach we take in Rock et al. (2023), for products that have an identical listing.

Why do sellers put up with this behavior? As outlined in the recent United States Federal Trade Commission complaint against Amazon (Graham, 2023),³⁴ the company has used contractual requirements to raise switching costs and to prevent sellers from offering lower prices on competing marketplaces, or even on their own sites. Additional techniques involve punishing sellers who do not purchase services such as fulfillment and advertising from Amazon by sentencing them to virtual algorithmic invisibility. We explore this subject in greater detail in the companion paper, “Amazon’s Algorithmic Rents” (Strauss et al., 2023).

Why do users put up with it? It is quite possible that they do not, but just not in sufficient numbers to affect Amazon’s calculus about the overall profitability of its strategy. According to Stone (2021), “When sponsored ads were prominently displayed, there was a small, statistically detectable short-term decline in the number of customers who ended up making a purchase... But while he [Bezos] cautioned against alienating customers by serving too many ads, he opted to vigorously move forward, saying that any deleterious long-term consequences would have to be implausibly large to outweigh the potential windfall and the investment opportunities that could result from it.”

The damage to Amazon may be a gradual downslope or a sudden cliff. When does brand and reputation damage accumulate to the point that consumers start trusting Amazon less, shopping at Amazon less, and expending the effort of trying alternatives? If Amazon is experiencing slow, incremental costs from what it is doing, there is more hope that it will change its behavior³⁵ than if it is experiencing no costs unless regulators impose them.

3.3. Social media: Engagement is a two-edged sword

Assessing optimality for social media platforms, which may be used either for utility or entertainment, is more difficult than it is for utility platforms like search or e-commerce. If the user is simply looking to be entertained, an algorithmic feed optimized for engagement may arguably be what the user wants, as the success of TikTok demonstrates. And especially when informed by personalized data, ads can be highly relevant and add information value to consumers. But while higher engagement and time spent on a social media site can, on the surface, be seen as a sign that users are finding more value on it, there are also harms, starting with addictive behavior, and when engagement is driven by posts fueling anger, self-doubt, misinformation, or controversy.

Social media feed recommendations, which contain advertising directly in the feed, give platforms an incentive to extract attention rents via increased engagement. More time spent on the platform means more surface area for advertising. Thus, consumers may see a higher proportion of stories optimized for “bad” forms of engagement, including those that are divisive, titillating, contain misinformation, or are otherwise harmful to consumers. As Stray (2021) notes, over the long term, companies have incentives to manage for “good” engagement so that users do not leave, but over the short term, they can profit even from “bad” engagement.

Advertisers are also harmed when their ads are algorithmically matched to inappropriate content (Internet Advertising Bureau, 2020). Only the platform benefits, which is why we consider these systems ripe for attention rents. More research is needed.

4. Measuring algorithmic attention rents

Today’s platforms are data-rich environments, where harms can be measured and estimated directly. Antitrust regulatory bodies now include data teams (Hunt, 2022), though they remain hampered by platforms not providing regular operating and data disclosures (Section 5.2).

³⁴ The full text of the complaint can be found in *Federal Trade Commission, et al. v. Amazon.Com, Inc.* (2023).

³⁵ Describing research on long-term versus short-term impacts of excessive advertising at Google, Hohnhold et al. (2015) note “Reducing the mobile ad load strongly improved the user experience but was a substantially short-term revenue negative change... the long-term revenue impact was shown to be neutral. Thus, with the user satisfaction improvement, this change was a net positive for both business and users.”

Measuring the extent to which sponsored listings have displaced organic search results is one way to understand whether or not these platforms are abusing their market power. We make the assumption that if, as the platforms have both promised and demonstrated, their algorithms aim to make the best possible attention allocations, the deviation between organic and paid results provides evidence of rents and harms. Using this approach, the presence of algorithmic attention rents can be inferred as follows:

1. *Comparing the organic ranking of a product with its paid ranking* to determine the extent to which the platform is preferencing results that its organic algorithm shows are inferior. For example, on Amazon, paid advertising results often rank a product higher—sometimes far higher—than the organic listing for the same product, and to other superior organically ranked products (Rock et al., 2023).
2. *Examining whether ads bring additional information to consumers.* For example, on Amazon, the organic listing for a product and an identical paid listing (ad) for the same product often appear side by side. This duplication reduces variety. No new information is provided. Instead, net new information is taken away.
3. *Comparing the quality of a dominant platform's organic algorithmic results with the organic allocations offered by other less dominant platforms that do face competitive pressures.* For example, it is possible to compare Amazon's product search results on its marketplace with those from Google Shopping; or Google Travel results with those from a site such as TripAdvisor. This entails an assessment of the relative quality of the organic results, which may be a very data intensive and complex task.
4. *Examining whether or not the information (including information quality) that a business or consumer could reasonably expect to find in a competitive market is available* (Areeda and Hovenkamp, 2023).³⁶ For example, in a competitive market a consumer would expect to find the total (and unit) price of a good or service displayed to aid in shopping comparisons. Yet, this is often hidden from view online. Airbnb, for example, only recently started showing total prices of accommodation (including cleaning fees) (Airbnb, 2022). So-called “drip pricing” (Fletcher, 2012), only showing part of the price to the consumer initially, is widespread, but may be more common on platforms with greater market power.³⁷

In another example, companies such as Yelp and TripAdvisor have complained that Google's Travel listings do not make the same efforts as these smaller platforms to filter out review spam—reflecting Google's added market power (Hawkins, 2018).³⁸ The information degraded could also be the quality of advertising results (and disclosures) that an advertiser could reasonably expect from the platform. The DoJ argued in its 2023 trial against Google that it degraded the ad quality in search results by making changes to its keyword matching. This reflects Google's monopoly over “when and where” an ad appears (*United States of America v. Google LLC*, 2020, p. 6).

5. *Examining whether ads have increased (and organic output declined) beyond the level reasonably required for the platform to earn a competitive return on capital invested.* Estimating a reasonable level of return is always difficult, especially in a new industry where norms have not been widely established, and network effects create high levels of concentration. However, a historical view of the platform operating and financial metrics can provide some perspectives. For example, if profit per user (or per search, per session, or other relevant measure of service delivered) continues to increase once user growth has leveled off, without demonstrated reductions in costs or commensurate ecosystem benefits, this would suggest increased monetization above the level that was

³⁶ Agreements pertaining to advertising and related dissemination of product information.

³⁷ To regulate this behavior, the FTC is considering requiring platforms to display “all-in pricing” (US Federal Trade Commission, 2022; FTC, 2023). The UK CMA and Department for Business and Trade has made similar recommendations (Hollinrake, 2023).

³⁸ Harsher regulation of fake reviews is under consideration in both the United Kingdom (Hollinrake, 2023; U.S. Federal Trade Commission, 2023) and the United States.

previously considered sufficient to deliver the service. This profit-driven monetization growth would need to be linked to the platform's overall level of profit margins or profit growth for it to be found to be exploitative, or above normal.³⁹

Advertising needs to be looked at in combination with the other platform fees and prices a platform charges in order to assess its reasonableness. For example, Apple justifies the 30% commission it charges for purchases on its App Store by relating it to various investments and user benefits such as privacy and quality (*Epic Games v. Apple Inc.*, 2021). Yet in addition, Apple charges app developers for visibility on its App Store by reserving prime display and result spots for advertising. These two fees combined—advertising and commission—are significantly above 30% (Kuriata, 2022). Financial disclosures by the company, which put the two revenue sources in different categories, help obscure the total cost to app developers.

Much of the data for these assessments can be found in an extensive literature from the Search Engine Optimization community, social media marketing consultants, e-commerce consultants, and the like, which provides some understanding of the ranking factors that guide algorithmic rankings and recommendations. The platforms themselves drop tantalizing clues in their announcements, blog posts, conference presentations, and annual shareholder letters (O'Reilly et al., 2023a). In addition, there are numerous one-time studies by academics, consultants, marketplace or advertising data firms, and occasionally, regulators. These studies usually consist of statistical analysis of a snapshot of web-scraped data at a specific time. For example, from a study of 1.4 billion searches by 28 million UK citizens (Goodwin, 2012), we know that in 2011, 94% of Google clicks were organic and only 6% went to ads. But we have no idea what the ratio was in different countries, what that ratio is today, or how it changed in the intervening years as Google has updated its algorithms and screen designs.

While information from studies such as this is useful, it points at a gaping hole in the regulatory apparatus: the lack of regular, mandated disclosures by platform companies of the operating metrics that actually guide the design of their algorithms, measure their results, and ultimately control the monetization of user attention.

5. Some possible regulatory interventions

The theory of algorithmic rents and data available from studies such as the one we have done suggest some relatively straightforward regulatory interventions.

5.1. Regulations of algorithmic output and preferences

For *search-based* algorithmic systems, including App Stores and e-commerce, regulators could impose the following requirements:

1. A percentage of the screen positions receiving the top share of attention could be reserved for organic results.
2. When ads duplicate organic results and provide no additional information (as in the Amazon marketplace), regulations could require that the organic result appear first, in the position most likely to be clicked on.
3. When there is an exact match for a user query (as in a search for a product by brand name in the Amazon marketplace), the exact match organic result must be the first result.

For *feed-based* algorithmic systems, such as TikTok, Twitter, and Instagram:

³⁹ For exploitation in EU law, see Whish and Bailey (2021, pp. 201–202).

1. Platforms could be required to offer “sticky” preferences to users, rather than requiring them to repeatedly assert their wishes (e.g., in Instagram, it is possible for the user to choose to see posts only from those they follow or have marked as favorites, but the feed almost immediately reverts to the platform’s own algorithmic feed preferences). In general, when users opt out of a new behavior offered by the platforms, that preference should be persistent.
2. Attention-hijacking patterns such as autoplaying videos could require an opt-in expression of preference.

However, these interventions suffer from several flaws:

1. Platform behavior is a moving target, with new features, constant algorithm updates, and design changes.
2. Regulatory interventions are likely to be countered by the platforms, much as they counter attempts by users and marketplace participants to “game” their algorithms.
3. Over-specified regulations could inadvertently strangle legitimate innovations that would benefit not only the platforms but also their users and marketplace participants.

The more fundamental problem that regulators need to address is that the mechanisms by which platforms measure and manage user attention are poorly understood. Effective regulation depends on enhanced disclosures.

5.2. The need for regular, mandated, disclosures of operating metrics

The platforms themselves collect numerous and detailed operating metrics to judge and manage the performance of their own systems for directing user attention. They know how many users they have, how much time those users spend on each of the platform’s services, how they are monetized, and the impact of new services and designs on their usage and monetization. They know the ad load. They know the ratio of organic clicks to ad clicks. They know much traffic is sent on to outside sites, by market segment. They know the gross merchandise volume of an e-commerce marketplace or app store, and they know what percentage they collect in fees. And they know how each of these measurements compares to prior periods—and how those changes result from updates of their interface designs and algorithms—just as well as they know their personnel costs, capital equipment costs, revenue, and their profit. But only the financial metrics are reported regularly and consistently, and those financial reports are almost completely disconnected from the operating metrics that are used to actually manage so much of the business (O’Reilly et al., 2023a).

The lack of disclosure of operating metrics for the free side of internet aggregators is a gaping hole in the regulatory apparatus. Costs, revenue, profit, and other financial metrics may be sufficient to understand a business based on tangible inputs and outputs, but are not fit for purpose for information businesses whose assets and activities are largely intangible and whose market power is exercised through delivery of services that are free to consumers (Mazzucato et al., 2023).

Given the size of the major internet companies, these metrics need to be highly disaggregated, both on a geographical and product basis. Google parent Alphabet alone has more than nine free products with more than a billion users, yet it reports only one major business—“advertising.” The connection between its revenues and the underlying free products and services is completely opaque. Meta too discloses little disaggregated information about products such as Facebook, Instagram, WhatsApp, and Messenger, each with billions of users. This pattern is repeated across the industry.

Even the products that are directly monetized often are not required to be broken out in detail, due to outdated segment reporting rules. US securities regulations require companies to break out financial detail for any business operation (or “segment”) that represents more than 10% of revenue or profit, yet in practice, they allow for management discretion in what segments are reported. This allowed Amazon to hide the remarkable growth and profitability of its Amazon Web Services business for years, and for

Apple to claim in response to a lawsuit from Epic Games that it did not actually know the profitability of its App Store (Mazzucato et al., 2023). But perhaps more importantly, it allows them to hide the workings of the free side of the products that underpin their enormous market power.

The most important first step is to understand how much is left out of the picture when market power is measured purely in financial terms. We live in an attention economy. A product like Google Maps, with over a billion users, has insignificant revenue relative to Alphabet's total (perhaps \$3 billion out of \$289 billion), yet it is unquestionably the most powerful player in its competitive segment, offering free products that cap the growth and opportunities for smaller mapping companies like ESRI or Garmin with more traditional business models. Segment reporting rules should be triggered by the number of active users a product has, not just by its contribution to company revenue or profit.

5.3. *Some recommended reportable operating metrics*

Understanding what operating metrics need to be reported for free products and services is in its infancy. In contrast to the financial metrics required on the money side of businesses, which are rooted in systems of accounting dating back to the 13th century, the metrics that are used to govern algorithmic attention businesses are at best a few decades old. Nonetheless, we must make a start.

All metrics should be reported quarterly, with more detail annually, as part of the existing financial disclosures required of public companies. And because a historical view would be useful, to whatever extent possible, the introduction of such reporting should require a backward look for at least several years, and ideally much longer. As noted above, metrics should also be disaggregated by product, with reporting required for any product having more than 100 million monthly users. They should also be disaggregated by country, and by device type (desktop or mobile.)

Here are some of the metrics that would be useful for examining attention rents monetized through advertising:

1. *Ad load.* Because not every page has the same number of ads—on Google, for example, many search engine results pages are noncommercial, and carry no ads at all—ad load should be reported by decile, or some other framework that highlights the ad concentration on the most highly monetized pages.
2. *Ratio of organic clicks to ad clicks.* Again, by decile or other weighted format.
3. *Average click through rate of the first organic result.* The proportion of users who visit the page who click on the first organic result.
4. *Average click through rate of the first ad.* The proportion of users who visit the page who click on the first ad.
5. *Amount of traffic sent on to third party sites.* This should be bucketed by market segment, such as news, entertainment, commerce, travel, local search, and so on.
6. *Amount of traffic sent to the company's own other products and services.* This could be further detailed by traffic source. For example, it would be useful to know how many users come to Google Search from Chrome on Apple devices vs. Chrome on Android, vs. from other browsers such as Firefox.
7. *Gross merchandise volume (for e-commerce platforms).* Without this information, it is impossible to determine the percentage of all fees levied on third party marketplace participants.
8. *Gross fee revenue, including advertising* from marketplace participants (for e-commerce platforms and app stores).
9. *A monetization narrative* explains the relationship between these various metrics describing the free side of their platform and their monetization on other platform sides.

Ideally, regulators, working with cooperative industry players, would define reportable metrics based on those that are actually used by the platforms themselves to manage search, social media, e-commerce, and other algorithmic relevancy and recommendation engines. These metrics should then be standardized

and required. There may be some metrics that can legitimately be considered trade secrets, but there are many that are common to most if not all internet businesses of the same type.

Note also that the operating metrics of big tech players are a moving target, constantly updated as the platforms continue to innovate. So this is also an opportunity to update the standards-setting process by which required reporting metrics are defined, mandating updated and timely reporting of any meaningful change in operating metrics.

Platforms will claim, with some justice, that disclosure will harm their businesses, as it will allow third parties to game their systems more easily. But this is akin to the old approach to cybersecurity, of “security through obscurity.” We have learned that it is far better to find and fix vulnerabilities than to hide them.

This is particularly important as we enter the age of large language models (LLMs) and generative AI, the next generation of attention management machines foretold by Herbert Simon. You cannot regulate what you do not understand (O'Reilly, 2023). It is not enough to rely on the assurances of powerful players that they are doing their best. Regular, reportable metrics will allow investors, the public, regulators, and the platforms themselves to better understand and operate truly free markets, which, as classical economists such as Adam Smith and David Ricardo believed, are not markets free of government intervention, but markets free of rents (Mazzucato, 2018).

6. Conclusion: AI and attention

While this article focuses on the present state of algorithmic attention rents, we are far from done with the changing institutional context of human decision-making online. Today's LLMs do not provide ranked choices but answers. They do not (at present) send traffic or other compensation to third party content, sites, or apps; nor do they depend on clicks and views. They do not (yet) have an advertising-based business model. Yet, they are quite consistent with the thrust of this article, since they depend on users accepting the increasing penetration of algorithmic authority (Choudhury and Shamszare, 2023), not just into everyday decision-making but increasingly, into everyday thinking.

Attention and cognition conserving heuristics help explain much about the present trajectory of today's LLMs and other “frontier” AI systems. Newell and Simon (1975) viewed both human and artificial intelligence inextricably bound up with the use of context and selectivity to create more efficient, guided, approximate solutions—for humans and machines.⁴⁰

As platforms evolve, AI may not only guide and replace human decision-making but human thinking (cognition.) LLMs can make use of existing information by creating new or summarized outputs from that information, not through the ranking of it. As AI improves, those outputs may become increasingly original, perhaps transcending their inputs and creating new knowledge and new possibilities. But even today:

1. The model can save users enormous cognitive costs and time. It demands little by way of users' attention and so improves upon the net time saving provided to users by existing productivity platforms and services. But at the same time, it takes even greater trust by users in the reliability and fairness of the model and of the platform providing it.
2. The model produces a probabilistic service for its users. It does so by consuming information inputs in a production process (“training”) to produce outputs that may vary each time. In the same way, an orchestra will train to gain skills, and then each performance may vary depending on the occasion. This makes it even harder to measure bad behavior from the outside, making the need for preemptive disclosures even more urgent.
3. While the first generation of LLMs were not updated in real time through interaction with users, suppliers, and advertisers, that is coming. The early instantiations of today's web applications were

⁴⁰Newell and Simon (1975, p. 13) note that “The task of intelligence, then, is to avert the ever-present threat of the exponential explosion of search.”

also relatively static captures of a moment in time, and only became rapidly updated as the technology progressed. But even today, these models depend on content created by humans—the vast corpus of human knowledge and creativity on which they have been trained.

Despite, and perhaps because of, all these differences between current search, recommendation, and feed algorithms and LLMs, the need for disclosure is paramount. Like their predecessors, these LLM systems internalize and centralize a vast marketplace of human knowledge and experience. As presently implemented, AI systems pass through neither attention nor remuneration to the providers of content used to train the model. As in present systems, human inputs are regarded as raw materials that can be appropriated by the developer of the system. If history is any guide, control over these raw materials by frontier AI platforms will eventually lead to the quest for a business model that allows for the extraction of monopoly rents.

Looking back at what we know now about present platforms, we can only wish there had been a disclosure regime that would have shown us the state of these systems when their creators were focused on serving their users and other ecosystem partners, and thus told us when and how they began to turn from that path to extract self-serving economic rents. Much like their predecessors, these frontier AI systems are managed by metrics whose details are known only to their creators and disclosed to the outside world only via generalities and sporadic, often self-serving data points. The time to establish rules for disclosure of operating metrics for frontier AI systems is now (O'Reilly, 2023).

We are still in the early stages of the attention economy, and innovation should be allowed to flourish. But this places an even greater emphasis on the need for transparency, and the establishment of baseline reporting frameworks that will allow regulators to measure whether attention allocation systems, including frontier AI systems, are getting better or worse over time. Greater public visibility into the operation of these platforms can, in conjunction with more informed policy making, lead to better behavior on the part of those who own and manage these systems, more balanced ecosystems of value creation, and the optimal use of knowledge in society.

Data availability statement. All secondary data used in this article have been cited and are available by request of those authors. Contact rufus.rock.21@ucl.ac.uk for data cited in Rock et al. (2023).

Acknowledgments. The authors acknowledge the helpful comments from Jennifer Pahlka, Bill Janeway, Greg Linden, Betsy Masiello, and Derek Slater. All errors are their own.

Author contribution. Writing – original draft: M.M., T.O., I.S.; Writing – review & editing: M.M., T.O., I.S.

Funding statement. This research was supported by the Omidyar Network.

Competing interest. No competing interests exist which have influenced in any undue matter the views advanced in this article.

Provenance statement. This article was published as a working paper (O'Reilly et al., 2023b).

References

- Airbnb** (2022) Airbnb is introducing total price display and updating guest checkout. Airbnb. Available at <https://news.airbnb.com/airbnb-is-introducing-total-price-display-andupdating-guest-checkout/>. (accessed 26 January 2024)
- Albrecht JP** (2016) How the GDPR will change the world. *European Data Protection Law Review* 2(3), 287–289.
- Alchian AA** (2017) Rent. In Durlauf S and Blume LE (eds.), *The New Palgrave Dictionary of Economics*. London: Palgrave Macmillan UK, pp. 1–6.
- Amazon.com** (2022) Amazon.com 10-K. US Securities and Exchange Commission. Available at https://www.sec.gov/Archives/edgar/data/1018724/000101872423000004/amzn-20221231.htm#icc32c5c732854b7f9975929c57cd5bd4_88. (accessed 26 January 2024)
- Areeda P and Hovenkamp H** (2023) *Antitrust Law: An Analysis of Antitrust Principles and their Application*. New York: Wolters Kluwer.
- Barinka A** (2023) TikTok starts US retail push with copycats, trinkets. *Bloomberg*. Available at <https://www.bloomberg.com/news/newsletters/2023-09-08/tiktok-starts-us-retailpush-with-copycats-trinkets>. (accessed 26 January 2024)

- Begent C and Collyer K** (2013) OECD roundtable on the role and measurement of quality in competition analysis (UK contribution). *Competition Law Journal* 12, 452.
- Behringer S and Filistrucchi L** (2015) Areeda–turner in two-sided markets. *Review of Industrial Organization* 46(3), 287–306.
- Blacquière E** (2014) Steeds meer ruimte AdWords, ten koste van SEO. EdWords.nl. Available at <https://www.edwords.nl/2014/04/27/steeds-meer-ruimte-adwords-ten-koste-van-seo/>. (accessed 26 January 2024)
- Blaug M** (1997) *Economic Theory in Retrospect*, 5th Edn. Cambridge: Cambridge University Press.
- Boyce JK** (2022) Political economy of the environment: A look Back and ahead. In Laurent É and Zwickl K (eds.), *The Routledge Handbook of the Political Economy of the Environment*. London: Routledge.
- Brin S** (1998) The PageRank citation ranking: Bringing order to the web. *Proceedings of ASIS* 98, 161–172.
- Business of Apps** (2022) CPC rates. Available at <https://www.businessofapps.com/ads/cpc/research/cpc-rates/>. (accessed 26 January 2024)
- Choudhury A and Shamszare H** (2023) Investigating the impact of user trust on the adoption and use of ChatGPT: Survey analysis. *Journal of Medical Internet Research* 25, e47184.
- Christophers B** (2020) *Rentier Capitalism: Who Owns the Economy and Who Pays for it?* London: Wiley.
- Competition & Markets Authority (CMA)** (2022) Evidence review of online choice architecture and consumer and competition harm. Available at <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>. (accessed 26 January 2024)
- Craswell N, Zoeter O, Taylor M and Ramsey B** (2008) An experimental comparison of click position-bias models. In *Proceedings of the 2008 International Conference on Web Search and Data Mining. WSDM '08*. New York: Association for Computing Machinery, pp. 87–94.
- Dean B** (2023) We analyzed 4 million google search results. Here's what we learned about organic CTR. Backlinko. Available at <https://backlinko.com/google-ctr-stats>. (accessed 26 January 2024)
- Doctorow C** (2023) The 'enshittification' of TikTok. *Wired*. Available at <https://www.wired.com/story/tiktok-platforms-cory-doctorow/>. (accessed 26 January 2024)
- Epic Games v. Apple Inc.** (2021) United States District Court Northern District of California. Case No. 20-cv-05640-YGR. Available at <https://cand.uscourts.gov/cases-e-filing/cases-of-interest/epic-games-inc-v-apple-inc/>. (accessed 26 January 2024)
- Federal Trade Commission, et al. v. Amazon.Com, Inc.** (2023) United States District Court for the Western District of Washington. Case No. 2:23-cv-01495. Available at https://www.ftc.gov/system/files/ftc_gov/pdf/1910129AmazonCommerceComplaintPublic.pdf. (accessed 26 January 2024)
- Feedvisor** (2022) The 2022 Amazon consumer behavior report. Feedvisor. Available at <https://feedvisor.com/resources/e-commerce-strategies/the-2022-amazon-consumer-behavior-report/>. (accessed 26 January 2024)
- Fletcher A** (2012) Drip pricing: UK experience. US Federal Trade Commission. Available at https://www.ftc.gov/sites/default/files/documents/public_events/economics-drip-pricing/afletcher.pdf. (accessed 26 January 2024)
- Foley DK** (2008) *Adam's Fallacy: A Guide to Economic Theology*. Cambridge: Belknap Harvard.
- FTC** (2023) FTC proposes rule to ban junk fees. FTC. Available at <https://www.ftc.gov/newsevents/news/press-releases/2023/10/ftc-proposes-rule-ban-junk-fees>. (accessed 26 January 2024)
- Goodwin D** (2012) Organic vs. paid search results: Organic wins 94% of time. *Search Engine Watch*. Available at <https://www.searchenginewatch.com/2012/08/23/organic-vs-paid-searchresults-organic-wins-94-of-time/>. (accessed 26 January 2024)
- Goodwin D** (2023) Former Googler: Google 'Using clicks in rankings'. *Search Engine Land*. Available at <https://searchengineland.com/former-googler-google-using-clicks-in-rankings-432401>. (accessed 26 January 2024)
- Google** (2023a) Rigorous testing – How Google Search Works. Google. Available at https://www.google.com/intl/en_uk/search/howsearchworks/how-search-works/rigorous-testing/. (accessed 26 January 2024)
- Google** (2023b) About Ad Quality. Google Ads Help. Available at <https://support.google.com/google-ads/answer/156066?sjid=12077254068652264551-NA#QSVAR>. (accessed 26 January 2024)
- Gornyi A** (2023) Letter: Why Amazon's advertising revenue is not what it seems. *Financial Times*. Available at <https://www.ft.com/content/e1969f6e-85ce-42aa-9149-ca21a8a895ad>. (accessed 26 January 2024)
- Graham V** (2023) FTC sues Amazon for illegally maintaining monopoly power. US Federal Trade Commission. Available at <https://www.ftc.gov/news-events/news/press-releases/2023/09/ftc-sues-amazon-illegally-maintaining-monopoly-power>. (accessed 26 January 2024)
- Hawkins J** (2018) Yelp vs Google: How they deal with fake reviews. *Search Engine Land*. Available at <https://searchengineland.com/yelp-vs-google-how-do-they-deal-with-fake-reviews-307332>. (accessed 26 January 2024)
- Hayek FA** (1945) The use of knowledge in society. *The American Economic Review* 35(4), 519–530.
- Hohnhold H, O'Brien D and Tang D** (2015) Focus on the long-term: It's better for users and business. In *Proceedings 21st Conference on Knowledge Discovery and Data Mining*. Sydney, Australia: ACM. <http://dl.acm.org/citation.cfm?doi=2783258.2788583>. (accessed 26 January 2024)
- Hollinrake K** (2023) Government clampdown on fake reviews and hidden fees to help customers cut the costs of living. *UK Department for Business and Trade*. Available at <https://www.gov.uk/government/news/government-clampdown-on-fake-reviews-and-hidden-fees-tohelp-customers-cut-the-costs-of-living>. (accessed 26 January 2024)
- Hovenkamp H** (2020) Antitrust and platform monopoly. *Yale Law Journal* 130(8), 1952–2051.
- Hugonenc C** (2022) Attention is the new metric. *Advertising Week*. Available at <https://advertisingweek.com/attention-is-the-new-metric/>. (accessed 26 January 2024)

- Hunt S** (2022) The technology-led transformation of competition and consumer agencies: The competition and markets authority's experience. UK Competition & Markets Authority. Available at <https://www.gov.uk/government/publications/the-technology-led-transformation-of-competition-and-consumer-agencies-the-cmas-experience>. (accessed 26 January 2024)
- Hwang T** (2020) *Subprime Attention Crisis: Advertising and the Time Bomb at the Heart of the Internet*. New York: FSG Originals.
- Internet Advertising Bureau** (2020) Understanding brand safety & brand suitability in a contemporary media landscape. IAB. Available at https://www.iab.com/wp-content/uploads/2020/12/IAB_Brand_Safety_and_Suitability_Guide_2020-12.pdf. (accessed 26 January 2024)
- Joachims T, Swaminathan A and Schnabel T** (2017) Unbiased learning-to-rank with biased feedback. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*. New York: ACM, pp. 781–789.
- Jullien B, Pavan A and Rysman M** (2021) Chapter 7: Two-sided markets, pricing, and network effects. In Ho K, Hortaçsu A and Lizzeri A (eds.), *Handbook of Industrial Organization*, Vol. 4. Amsterdam: Elsevier, pp. 485–592.
- Jumpshot** (2018) The competitive state of eCommerce marketplaces: Data report Q2 2018. Jumpshot.
- Jungle Scout** (2022) State of the Amazon seller. Jungle Scout. Available at <https://www.junglescout.com/amazon-seller-report/2022-results/>. (accessed 26 January 2024)
- Kahneman D** (2011) *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.
- Kaziukėnas J** (2021) Everything on Amazon is an Ad. Marketplace Pulse. Available at <https://www.marketplacepulse.com/articles/everything-on-amazon-is-an-ad>. (accessed 26 January 2024)
- Kaziukėnas J** (2022) Amazon is burying organic search results. Marketplace Pulse. Available at <https://www.marketplacepulse.com/articles/amazon-is-burying-organic-search-results>. (accessed 26 January 2024)
- Keane MT and O'Brien M** (2006) Modeling result-list searching in the world wide web: The role of relevance topologies and trust bias. *Proceedings of the Annual Meeting of the Cognitive Science Society 28*, 1–6.
- Keane MT, O'Brien M and Smyth B** (2008) Are people biased in their use of search engines? *Communications of the ACM 51*(2), 49–52.
- Kim L** (2016) Google kills off side Ads: What you need to know. WordStream. Available at <https://www.wordstream.com/blog/ws/2016/02/22/google-kills-off-right-side-ads>. (accessed 26 January 2024)
- Kuriata G** (2022) Apple search Ads cost: CPT, CPA and CR benchmarks 2022. SplitMetrics. Available at <https://splitmetrics.com/blog/apple-search-ads-cost/>. (accessed 26 January 2024)
- Kurz M** (2023) *The Market Power of Technology*. New York: Columbia University Press.
- Lopatenko A, Kim H-J, Dornbush S, Wei L, Kilbourn TP and Lopyrev M** (2015) Ranking search results based on similar queries. U.S. pat. Google LLC. Available at <https://patents.google.com/patent/US9009146B1/en>. (accessed 26 January 2024)
- Marvin G** (2020) A visual history of Google Ad labeling in search results. *Search Engine Land*. Available at <https://searchengineland.com/search-ad-labeling-history-google-bing-254332>. (accessed 26 January 2024)
- Mazzucato M** (2018) *The Value of Everything: Making and Taking in the Global Economy*. Paris: Hachette UK.
- Mazzucato M, Ryan-Collins J and Gouzoulis G** (2023) Mapping modern economic rents: The good, the bad, and the grey areas. *Cambridge Journal of Economics 47*(3), 507–534.
- Mazzucato M, Strauss I, O'Reilly T and Ryan-Collins J** (2023) Regulating big tech: The role of enhanced disclosures. *Oxford Review of Economic Policy 39*(1), 47–69.
- Mileva G** (2022) Amazon Ad revenue statistics that will blow your mind. Influencer Marketing Hub. Available at <https://influencermarketinghub.com/amazon-ad-revenue/>. (accessed 26 January 2024)
- Mitchell S** (2023) New research: Amazon's monopoly tollbooth in 2023. Institute for Local Self-Reliance. Available at <https://ilsr.org/amazonmonopolytollbooth-2023/>. (accessed 26 January 2024)
- Morrison S** (2021) Amazon's strategy to squeeze marketplace sellers and maximize its own profits is evolving. *Vox*. Available at <https://www.vox.com/recode/22810795/amazon-marketplaceprime-report>. (accessed 26 January 2024)
- Newell A and Simon HA** (1975) Computer science as empirical inquiry: Symbols and search. *ACM Turing award lectures*.
- Newman JM** (2015) Antitrust in zero-Price markets: Foundations. *University of Pennsylvania Law Review 164*(1), 149–206.
- O'Reilly T** (2019) Antitrust regulators are using the wrong tools to break up big tech. *Quartz*. Available at <https://qz.com/1666863/why-big-tech-keeps-outsmarting-antitrust-regulators>. (accessed 26 January 2024)
- O'Reilly T** (2022) Google Shopping Search. O'Reilly Media. Available at <https://www.oreilly.com/tim/antitrust/shoppingsearch.html>. (accessed 26 January 2024)
- O'Reilly T** (2023) You can't regulate what you don't understand. O'Reilly Media. Available at <https://www.oreilly.com/content/you-cant-regulate-what-you-dont-understand-2/>. (accessed 26 January 2024)
- O'Reilly T, Strauss I and Mazzucato M** (2023a) Regulating Big Tech through digital disclosures. *UCL IIPP Policy Brief No. 26 June 2023*. Available at https://www.ucl.ac.uk/bartlett/public-purpose/sites/bartlett_public Purpose/files/oreilly_strauss_mazzucato_2023_regulating_big_tech_through_digital_disclosures.pdf. (accessed 26 January 2024)
- O'Reilly T, Strauss I and Mazzucato M** (2023b) *Algorithmic Attention Rents: A Theory of Digital Platform Market Power*. UCL Institute for Innovation and Public Purpose, Working Paper Series (IIPP WP 2023-10). Available at <https://www.ucl.ac.uk/bartlett/public-purpose/wp2023-10>. (accessed 26 January 2024)
- Palmer A** (2021) Amazon is piling Ads into search results and top consumer brands are paying up for prominent placement. *CNBC*. Available at <https://www.cnbc.com/2021/09/19/amazonpiles-ads-into-search-results-as-big-brands-pay-for-placement.html>. (accessed 26 January 2024)
- Petit N** (2020) *Big Tech and the Digital Economy: The Moligopoly Scenario*. Oxford: Oxford University Press.

- Petit N and Teece DJ** (2020) Taking ecosystems competition seriously in the digital economy: A (preliminary) dynamic competition/capabilities perspective. OECD Directorate for Financial and Enterprise Affairs Competition Committee. Available at [https://one.oecd.org/document/DAF/COMP/WD\(2020\)90/en/pdf](https://one.oecd.org/document/DAF/COMP/WD(2020)90/en/pdf). (accessed 26 January 2024)
- Petrescu P** (2014) How Ads influence organic click-through rate on Google. *Search Engine Land*. Available at <https://searchengineland.com/different-types-ads-influence-organicctr-google-204676>. (accessed 26 January 2024)
- Piketty T** (2014) *Capital in the 21st Century* (Trans. by Arthur Goldhammer). Cambridge MA: Belknap Press.
- Quartile** (2018) New study reveals Amazon's quiet shift toward 'pay-to-Play' platform – and what it means for brands. *PR Newswire*. Available at <https://www.prnewswire.com/news-releases/newstudy-reveals-amazons-quiet-shift-toward-pay-to-play-platform-and-what-it-meansfor-brands-300759702.html>. (accessed 26 January 2024)
- Ricks B and McCrosky J** (2022) Does this button work? Investigating YouTube's ineffective user controls. Mozilla Foundation. Available at <https://foundation.mozilla.org/en/research/library/user-controls/report/>. (accessed 26 January 2024)
- Robinson J** (2001) *The Economics of Imperfect Competition*. London: Palgrave.
- Rochet J-C and Tirole J** (2003) Platform competition in two-sided markets. *Journal of the European Economic Association* 1(4), 990–1029.
- Rochet J-C and Tirole J** (2006) Two-sided markets: A progress report. *The Rand Journal of Economics* 37(3), 645.
- 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). UCL IIPP Working Paper. Available at <https://www.ucl.ac.uk/bartlett/publicpurpose/research/digital-economy-and-algorithmic-rents/algorithmic-attention-rents>. (accessed 26 January 2024)
- Ryan-Collins J, Lloyd T and Macfarlane L** (2017) *Rethinking the Economics of Land and Housing*. London: New Economics Foundation.
- Samuelson PA and Nordhaus WD** (2010) *Economics*, 19th Edn. McGraw-Hill Irwin.
- Schumpeter JA** (2013) *Capitalism, Socialism and Democracy*. London: Routledge.
- Sen A** (2017) Rational behaviour. In Durlauf S and Blume LE (eds.), *The New Palgrave Dictionary of Economics*. London: Palgrave Macmillan UK, pp. 1–14.
- Simon HA** (1955) A behavioral model of rational choice. *The Quarterly Journal of Economics* 69(1), 99.
- Simon HA** (1971) Designing organizations for an information-rich world. *Computers, Communications, and the Public Interest* 72, 37.
- Simon HA** (1978a) Information-processing theory of human problem solving. In Estes WK (ed.), *Handbook of Learning and Cognitive Processes*. Mahwah, NJ: Lawrence Erlbaum Associates, pp. 271–295.
- Simon HA** (1978b) Rationality as process and as product of thought. *The American Economic Review* 68(2), 1–16.
- Simon HA** (1995) The information-processing theory of mind. *The American Psychologist* 50.7.
- Simon HA** (1997) *Administrative Behavior: A Study of Decision-Making Processes in Administrative Organizations*, 4th Edn. Glencoe, IL: Free Press.
- Simon HA** (2017) Bounded rationality. In Durlauf S and Blume LE (eds.), *The New Palgrave Dictionary of Economics*. London: Palgrave Macmillan UK, pp. 1–4.
- Soper S** (2023) *Amazon is taking half of each sale from its merchants*. *Bloomberg.com*. Available at <https://www.bloomberg.com/news/articles/2023-02-13/amazon-amzn-takes-half-of-each-sale-from-2-million-small-businesses>. (accessed 26 January 2024)
- Sorokina D and Cantu-Paz E** (2016) Amazon Search: The joy of ranking products. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*. New York: ACM, pp. 459–460.
- Srinivasan D** (2019) The antitrust case against Facebook: A Monopolist's journey towards pervasive surveillance in spite of consumers' preference for privacy. *Berkeley Business Law Journal* 16(1), 39–101.
- Standing G** (2016) *The Precariat: The New Dangerous Class*. London: Bloomsbury Academic.
- Stiglitz JE** (2019) *People, Power, and Profits: Progressive Capitalism for an Age of Discontent*. London: Penguin UK.
- Stone B** (2021) *Amazon Unbound*. New York: Simon and Schuster.
- Strandburg KJ** (2013) Free fall: The online market's consumer preference disconnect the frontiers of consumer protection. *University of Chicago Legal Forum* 2013, 95–172.
- Strauss I, O'Reilly T and Mazzucato M** (2023) *Amazon's algorithmic attention rents*. UCL IIPP Working Paper. Available at <https://www.ucl.ac.uk/bartlett/publicpurpose/research/digital-economy-and-algorithmic-rents/algorithmic-attention-rents>. (accessed 26 January 2024)
- Stray J** (2021) Beyond engagement: Aligning algorithmic recommendations with prosocial goals. Partnership on AI. Available at <https://partnershiponai.org/beyond-engagement-aligning-algorithmic-recommendations-with-prosocial-goals/>. (accessed 26 January 2024)
- The Badger** (2023) Amazon advertising stats. Ad Badger. Available at <https://www.adbadger.com/blog/amazon-advertising-stats/>. (accessed 26 January 2024)
- Thompson B** (2014) Economic power in the age of abundance. *Stratechery.com*. Available at <https://stratechery.com/2014/economic-power-age-abundance/>. (accessed 26 January 2024)
- Thompson B** (2023) *Commoditizing suppliers*. *Stratechery.com*. Available at <https://stratechery.com/concept/aggregation-theory/commoditizing-suppliers/>. (accessed 26 January 2024)
- Tirole J** (2017) *Economics for the Common Good*. Princeton: Princeton University Press.
- Tversky** (1981) The framing of decisions and the psychology of choice. *Science* 211(4481), 453–458.

- Tversky A and Kahneman D** (1974) Judgment under uncertainty: Heuristics and biases. *Science* 185(4157), 1124–1131.
- U.S. Federal Trade Commission** (2023) Federal Trade Commission announces proposed rule banning fake reviews and testimonials. Available at <https://www.ftc.gov/news-events/news/pressreleases/2023/06/federal-trade-commission-announces-proposed-rule-banning-fakereviews-testimonials>. (accessed 26 January 2024)
- United States of America v. Google LLC** (2020) United States District Court for the District of Columbia. Case No. 1:20-cv-03010-APM. Available at <https://cdn.arstechnica.net/wp-content/uploads/2023/09/US-v-Google-DOJ-Pre-Trial-Brief-9-8-2023.pdf>. (accessed 26 January 2024)
- US Federal Trade Commission** (2022) Unfair or deceptive fees trade regulation rule. R207011. Available at <https://www.federalregister.gov/documents/2022/11/08/2022-24326/unfair-or-deceptive-fees-trade-regulation-rule-commission-matter-no-r207011>. (accessed 26 January 2024)
- Wang X, Golbandi N, Bendersky M, Metzler D and Najork M** (2018) Position bias estimation for unbiased learning to rank in personal search. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining. WSDM '18*. New York: Association for Computing Machinery, pp. 610–618.
- Weyl EG** (2010) A price theory of multi-sided platforms. *American Economic Review* 100(4), 1642–1672.
- Whish R and Bailey D** (2021) *Competition Law*, 10th Edn. Oxford: Oxford University Press.
- Wilkes J** (2012) *Virtuous Cycle*. YouTube. Available at <https://www.youtube.com/watch?v=5jcDIGn-tZA>. (accessed 26 January 2024)
- Wu T** (2013) *The oligopoly problem*. *The New Yorker*. Available at <https://www.newyorker.com/tech/annals-of-technology/the-oligopoly-problem>. (accessed 26 January 2024)
- Yue Y, Patel R and Roehrig H** (2010) Beyond position bias: Examining result attractiveness as a source of presentation bias in clickthrough data. In *ACM: Proceedings of the 19th International Conference on World Wide Web*. New York: ACM, pp. 1011–1018.
- Zuboff S** (2019) *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. London: Profile Books.