

Rising tide rents and robber baron rents: How innovators lose their edge and their ideals

Tim O'Reilly

Founder, CEO, and Chairman of O'Reilly Media

| Visiting Professor of Practice | UCL Institute for Innovation and Public Purpose

About the Institute for Innovation and Public Purpose

The Institute for Innovation and Public Purpose (IIPP) at University College London (UCL) aims to develop a new framework for creating, nurturing and evaluating public value in order to achieve economic growth that is more innovation-led, inclusive and sustainable.

This requires rethinking the underlying economics that has informed the education of global civil servants and the design of government policies. Our work feeds into innovation and industrial policy, financial reform, institutional change and sustainable development.

A key pillar of IIPP's research is its understanding of markets as outcomes of the interactions between different actors. In this context, public policy should not be seen as simply fixing market failures, but also as actively shaping and co-creating markets. Re-focusing and designing public organisations around mission-led, public purpose aims will help tackle the grand challenges facing the 21st century.

IIPP is housed in The Bartlett, a leading global Faculty of the Built Environment at University College London (UCL), with its radical thinking about space, design and sustainability.

ISSN 2635-0122

This working paper can be referenced as follows:

O'Reilly, T. (2024). Rising tide rents and robber baron rents: How innovators lose their edge and their ideals. UCL Institute for Innovation and Public Purpose, Working Paper Series (IIPP WP 2024-04). Available at: <https://www.ucl.ac.uk/bartlett/public-purpose/wp2024-04>

Rising tide rents and robber baron rents: How innovators lose their edge and their ideals

Tim O'Reilly*

Abstract:

During a new technology cycle, market leaders emerge, because they solve new problems and create new value, not only for consumers but also for a rich ecosystem of suppliers, intermediaries and even competitors. Even though these market leaders receive a disproportionate share of the profits, earning so-called 'Schumpeterian rents' as they dominate the emerging market, value creation is a rising tide that lifts all boats. However, this kind of virtuous rising tide rent, which benefits everyone, doesn't last. Once the growth of the new market slows, the now-powerful innovators can no longer rely on new user adoption and collective innovation from a vibrant ecosystem to maintain their extraordinary level of profit. They often turn to extractive techniques, using their market power to try to maintain their now-customary profits in the face of macroeconomic factors and competition that ought to be eating them away. They start to collect robber baron rents. This pattern has played out throughout the history of the computer and software industry. As the industry begins a new cycle fuelled by generative AI, what can we learn from this history that might guide entrepreneurs, regulators and policymakers?

*Founder, CEO and Chairman of O'Reilly Media and Visiting Professor of Practice at UCL Institute for Innovation and Public Purpose. Contact tim@oreilly.com, 11 Montague St, London WC1B 5BP, United Kingdom.

This paper aims to provide broader context for the IIPP research project on Big Tech market power and algorithmic rents. Please visit the project page for other working papers and more information: <https://www.ucl.ac.uk/bartlett/public-purpose/research/digital-economy-and-algorithmic-rents/algorithmic-attention-rents>. This project has been funded by Omidyar.

1. Introduction

Why is it that Google, a company once known for its distinctive 'Do no evil' guideline, is now facing the same charges of 'surveillance capitalism' as Facebook, a company that never made such claims? Why is it now subject to the same kind of antitrust complaints faced by Microsoft, once the 'evil empire' of the previous generation of computing? Why is it that Amazon, which has positioned itself as 'the most customer-centric company on the planet', now lards its search results with advertisements, placing them ahead of the customer-centric results chosen by the company's organic search algorithms, which prioritise a combination of low price, high customer ratings and other similar factors?

The answer can be found in the theory of economic rents and, in particular, in the kinds of rents that are collected by companies during different stages of the technology business cycle. There are many types of rents and an extensive economics literature discussing them, but for the purposes of this paper, they can be lumped into two broad categories — 'rising tide rents' that benefit society as a whole, such as those that encourage innovation and the development of new markets, and 'robber baron rents' that disproportionately benefit those with power. The history of the computer industry and the commercial internet provides a remarkable series of natural experiments that illustrate the difference.

During the expansive period of a new technology cycle, market leaders emerge because they solve new problems and create new value, not only for consumers but also for a rich ecosystem of suppliers, intermediaries and even competitors. Even though these market leaders tend to receive a disproportionate share of the profits, as they lay waste to incumbents and dominate the emerging market, value creation is a rising tide that lifts all boats.

But this kind of virtuous rising tide rent, which benefits everyone, doesn't last. Once the growth of the new market slows, the now-powerful innovators can no longer rely on new user adoption and collective innovation from a vibrant ecosystem to maintain their extraordinary level of profit. In the dying stages of the old cycle, the companies on top of the heap turn to extractive techniques, using their market power to try to maintain their now-customary level of profits in the face of macroeconomic factors and competition that ought to be eating them away. They start to collect robber baron rents. That's exactly what IBM did in the mainframe era, what Microsoft did for personal computers, and what Google, Amazon and Meta are doing today

Each time, the cycle begins again with a new class of competitors, who are forced to explore new, disruptive technologies that reset the entire market. OpenAI, Anthropic and the other entrepreneurial companies exploring the frontiers of artificial intelligence represent the latest occurrence of this pattern.

2. What is economic rent?

Not to be confused with the ordinary sense of rent as a charge for temporary use of property, economic rents are the income above a competitive market rate that is collected because of asymmetries in ownership, information or power.

Mazzucato et al. (2020) write, 'If the reward accruing to an actor is larger than their contribution to value creation, then the difference may be defined as rent. This can be due to the ownership of a scarce asset, the creation of monopolistic conditions that enable rising returns in a specific sector, or policy decisions that favour directly or indirectly a specific group of interest.'

For example, consider drug pricing. Patents – exclusive, government-granted rights intended to encourage innovation – protect pharmaceutical companies from competition and allow them to charge high prices. Once the patents expire, there is competition from so-called ‘generic drugs’ and the price comes down. That difference in price (and its impact on pharmaceutical company profits) shows the extent of the rent.

In 20th century neoliberal economics, rents are typically seen as a temporary aberration that is eventually competed away, a price that we pay for the rising tide of innovation. However, as Mazzucato (2018) points out, to the classical economists – Smith, Ricardo and Mill – who lived in a world of inherited power and privilege, rents were a pernicious and persistent consequence (and source) of inequality. At the dawn of economic theory, agriculture was still the chief source of value creation, and much of that value created by the labour of serfs and tenant farmers was appropriated by those who owned the land. When the local baron sent his troops to collect what he considered his share of the harvest, it was impossible to say no. In an unjust society, neither effort nor investment nor innovation but rents rooted in power asymmetries determine who gets what and why.

But not all rents represent an abuse of power. As noted by Schumpeter (1942), 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 new market, but eventually growth slows. The market becomes saturated and the company ceases to create new value, focusing instead on cementing its market dominance and suppressing competitors, who are forced to explore new, disruptive technologies that reset the entire market (Christensen 1997).

This pattern has played out throughout the entire history of the computer and software industry. Modern digital technology has proceeded in a series of waves, alternating periods of innovation and value creation with periods of consolidation and value extraction. In each of these cycles, technology put into the public domain or otherwise made widely available lowered the barriers to market entry, creating a rising tide of decentralised innovation. Once the new leaders consolidated market power, though, they eventually lost their innovative edge and raised barriers to protect their advantages. The industry stagnated until the entrenched leaders were surpassed by those riding the next wave of innovation. Rents were collected, but eventually competed away, just as the theory predicts.

3. A brief history of computing from 1945 to 2005

The first digital computers were developed with UK and US government funding during World War II, with key designs put into the public domain (Dyson 2012). There was a flurry of innovation as companies such as IBM, Burroughs, UNIVAC, NCR, Control Data Corporation (CDC) and Honeywell rushed to commercialise digital computing. IBM achieved a dominant position, which it locked in by providing its equipment through leasing contracts that required customers to use only IBM-supplied parts and service. At its peak, IBM controlled as much as 70% of the computer market (Conigliaro et al. 1996). During this period, software was also tightly integrated with the hardware, so third-party software suppliers were rare and satellite to the hardware platform.

The next wave of decentralised innovation came with the introduction of the microprocessor and the rise of the personal computer. The Altair 8800, released in 1974 in kit form, quickly gave rise to hobbyist groups such as the Homebrew Computer Club and, by 1977, hobbyists had become entrepreneurs, with the release of pre-assembled computers such as the Apple II, the Commodore PET and the Tandy (Radio Shack) TRS-80.

Trying to catch up in the emerging market for microcomputers, in 1981 IBM released a personal computer with an architecture designed to support peripherals such as disks, memory boards, screens and printers from third-party manufacturers (Koenig 2011). IBM even published the PC's design documents so that it would be easy for manufacturers to create compatible peripheral devices and for developers to write software, which at that time required them to understand the underlying hardware (Chposky and Leonsis 1988). These open standards led to a period of intense competition, but entrepreneurs didn't stop at the boundary IBM had drawn. Independent manufacturers began to build what came to be called 'clones' (Bajarin 2021). IBM initially expected the market to be small — 240,000 machines over five years (Schwartz 2020) — but instead it exploded, with 200,000 a month sold within a year after the first release (Cortada 2021) and with PCs eventually numbering in the billions. However, only a tiny fraction of them came from IBM and eventually none at all.

There was a frenzy of innovation as thousands of companies built computers using the IBM design, which became a de facto standard, and eventually turned computer hardware into a low-priced commodity. Computer hardware became cheap, ubiquitous and powerful. Software, which in the age of the mainframe had been tightly bound to the hardware (much as it still is today with Apple's products), became decoupled from hardware and an industry in its own right. New categories of software, including spreadsheets, word processing, databases, desktop publishing, image manipulation and illustration became the basis of entrepreneurial fortunes.¹

The cycle of monopoly then began again. IBM had licensed the PC operating system from a small company called Microsoft, allowing it to retain ownership (Maher 2017). Little did IBM realise that personal computer hardware would become a commodity and that software would be the locus of control in the new industry. Initially, Microsoft was an innovator, turning that first operating system into a platform that enabled thousands of small software entrepreneurs and taking personal computing mainstream, with the expansive vision of 'a computer on every desk and in every home' (Choudhary 2015).

Over time, though, Microsoft used its position as operating system provider to advantage its own applications and took over one desktop application category after another. Over the next two decades, the frenzy of innovation of the early PC era ended with a dominant monopoly. Microsoft replaced IBM as extractor of outsized profits (DOJ 2006). While Microsoft continued to invest heavily in R&D, innovation as a whole slowed. Microsoft began to hobble its own internal innovators, imposing what came to be called a 'strategy tax' by which every new product had to protect or reinforce Microsoft's lucrative Windows and Office monopolies (Obasanjo 2004). The threat that Microsoft would compete with any successful new software company left venture capitalists attending company briefings about what areas might be safe to invest in.

Lacking opportunity in the PC market, computer enthusiasts and entrepreneurs began to look elsewhere. In the 1990s, the rise of open source software and the open protocols of the internet began to undermine Microsoft's lock on the industry and, by the time Google went public in 2004, it was clear that the industry had entered a new cycle.

¹ According to Wu (2020), this decoupling was at least in part due to IBM's restraint amid 'antitrust phobia', due to the long-running US antitrust lawsuit, emphasising that openness, modularity and competition generally lead to more robust innovation.

The 1990s were a hotbed of internet competition. Initially that competition was focused on what we now recognise as the dying gasps of the previous generation of computer entrepreneurship: PC-based software applications. Tim Berners-Lee's original web browser, which had been put into the public domain, was supplanted first by Mosaic, then by Netscape, then by Microsoft's Internet Explorer. A decade later, Internet Explorer's dominance was challenged in turn by the open source Mozilla browser, then toppled by Google Chrome. Netscape and Microsoft and open source web servers like Apache were locked in a battle for control of the web (O'Reilly 1999), but this was just the endgame of the era in which dominance had been achieved through control over software applications and operating systems.

A new game was afoot, dominated not by software per se, but by websites delivering what we now call software as a service, supported by business models like advertising and e-commerce, and underpinned by the collection of massive amounts of data.

During the 1990s, entrepreneurs went down many blind alleys, with investors pouring billions of dollars into 'dot com' startups delivering a bewildering array of services, many of them with no clear business model. Bill Janeway refers to such periods of experimentation as a 'productive bubble' driven by 'Schumpeterian waste' (2018). This bubble ended abruptly with the 'dotcom bust' of 2001, when many of these companies went bankrupt. However, those that survived had truly made enormous advances, not just in the scale and style of computing applications, but also in business models. In 2004, my company began using the term 'Web 2.0' (O'Reilly 2005) to describe the second coming of the web after the dotcom bust, explaining why some companies survived while so many others had gone down in flames. What was it that separated the winners from the losers? And what was it that eventually became the scarce resource that was the source of monopoly control?

4. Attention is all you need

What is the source of Big Tech market power? What is the limited resource that they control and monopolise? It's not our data. It's not the price of the services we purchase from them – they give those away for free. It's our attention (O'Reilly et al. 2024).

In a talk called 'Designing organizations for an Information-rich world', political scientist Herbert Simon (1971) noted that the cost of information is not just money spent to acquire it, but the time it takes to consume it:

'In an information-rich world, the wealth of information means a dearth of something else: a scarcity of whatever it is that information consumes. What information consumes is rather obvious: it consumes the attention of its recipients. Hence a wealth of information creates a poverty of attention and a need to allocate that attention efficiently among the overabundance of information sources that might consume it.'

In the discussion following his presentation, Simon commented that in the future, information would be so abundant that we would need machines to help us manage our attention - and that has indeed been the secret to success in the information age. Google was founded with the promise of finding the right web page out of billions, giving you just what you want and then sending you on your way. Amazon aimed to help customers find the best quality and price for any one of millions of products. Even social media started with the promise of information triage:

for each person, a unique feed of updates from only the friends they had chosen to follow. These are all astonishing tools for making our limited capacity for attention more efficient.

In the early idealistic days of internet expansion, the leading companies all earned outsized profits by solving the attention allocation problem. As the internet grew, the amount of information available to consumers became so vast that it outran traditional human means of curation and selection. Attention allocation was outsourced to the machines. Algorithms for search, recommendations, social media feeds, entertainment and news became the foundation of an enormous new economy.

The internet giants succeeded by doing what they are now too often reviled for: extracting signal from massive amounts of data. Google not only crawled and indexed virtually every page on the web, it looked at how sites linked to each other, tracked which of the ten top links they offered were clicked on the most, which ones led people to come back and try another, and which sent them away satisfied. It used location data and past searches to make answers more relevant and personalised. Amazon used everything from price, user reviews, popularity and your individual purchase history to bring to the top the products they believed best matched their customers' needs.

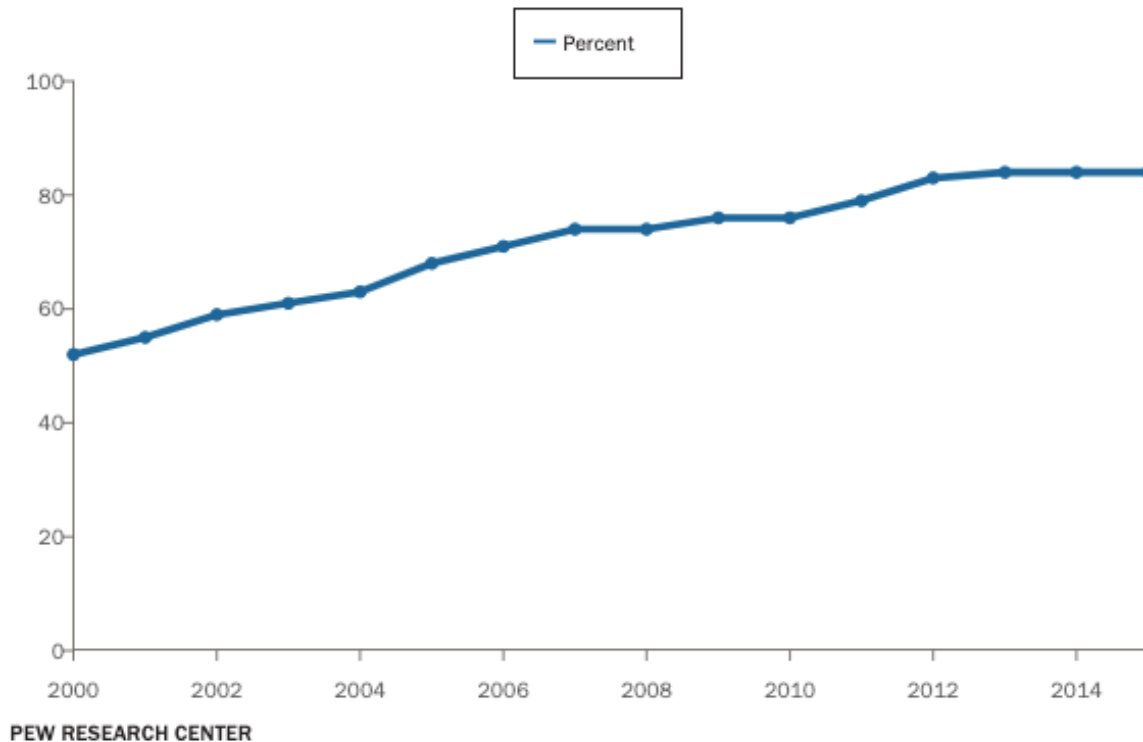
In *What is Web 2.0?* (O'Reilly 2005), I made the case that the companies that had survived the dotcom bust had all in one way or another become experts at 'harnessing collective intelligence.' Or perhaps a more direct way to say this in the context of economic value creation is that companies such as Amazon, Google and Facebook had developed a set of remarkable advances in networked and data-enabled market coordination.

Over time, though, something went very wrong. Instead of continuing to deploy their attention optimisation algorithms for their users' and suppliers' benefit, the tech giants began to use them to favour themselves. It first became obvious with social media: recommended posts and amplification of addictive, divisive content in order to keep users scrolling, creating additional surface area for advertising. Google began to place more and more advertising ahead of 'organic' search results, turning advertising from a complementary stream of useful information that ran beside search results into a substitute that pushed those results further down the page. Amazon was late to the party, but once it discovered advertising, it went all in. Now a typical page of Amazon product search results consists of 16 ads and only four organic results.

Google and Amazon were still atop their respective hills of web search and ecommerce in 2010, and Meta's growth was still accelerating, but it was hard to miss that overall internet growth had begun to slow. The market was maturing. From 2000 to 2011, the percentage of US adults using the internet had grown from about 60% to nearly 80%. By the end of 2012, it was up to 82%, But in 2013 and 2014, it remained stuck at 83% (Perrin et al. 2015; Pew 2024). While in the ten years since, it has reached 95% (Poushter et al. 2024), it had become clear that the easy money that came from acquiring more users was ending. Penetration in Europe, the other lucrative market, was on a similar track to the US and, while there was lots of user growth still to be found in the rest of the world, the revenue per user was much lower (O'Reilly 2021). What are now-gigantic companies to do when their immense market capitalisation depends on rapid growth and the expectation of growing profits to match?

84% of American Adults Use the Internet

% of all American adults who use the internet



Source: <https://www.pewresearch.org/internet/2015/06/26/americans-internet-access-2000-2015>

These companies did continue to innovate. Some of those innovations, like Amazon's cloud computing business, represented enormous new markets and a new business model, but the internet giants also came to focus on extracting more usage and time spent, and thus more revenue, from a relatively stable base of existing customers. Often this was done by making their products more addictive, getting more out of their users by nefarious means. Doctorow (2024) calls this the 'enshittification' of Big Tech platforms.

Fast forward to the present and Amazon has clearly given up on the goal of finding the best result for its users. Since launching its Marketplace advertising business in 2016, Amazon has chosen to become a 'pay to play' platform where the top results are those that are most profitable for the company. Research firm Marketplace Pulse (Kaziukenas 2022) notes:

Of the first twenty products a shopper sees when searching on Amazon, only four are organic results. There is little space left for organic results at the top of the page, the real estate that drives most sales. Few purchases happen beyond the first page of search results. And not many shoppers scroll to the bottom of even the first page...

It takes scrolling past three browser windows' worth of search results to get to the fifth organic result. It takes even more swipes to see the fifth organic result on mobile.

This is what we mean by a 'robber baron' rent: 'Pay us or you'll effectively disappear from search.' The harm to users isn't just lost time scrolling through ads to find the best results. In a recent research project here at University College London's Institute for Innovation and Public Purpose, my colleagues and I found that users still tend to click on the product results at the top of the page, even when they are no longer the best results. Amazon abuses the trust that users have come to place in its algorithms and instead allocates user attention and clicks to inferior quality, sponsored information (Strauss et al. 2024). The most-clicked sponsored products were 17% more expensive and 33% lower ranked according to Amazon's own quality, price and popularity optimising algorithms (Rock et al. 2023). What's more, because product suppliers must now pay for the product ranking that they previously earned through product quality and reputation, their profits go down as Amazon's go up and prices rise as some of the cost is passed on to customers.

It appears to have worked – for now. Amazon's most recent quarterly disclosures (Amazon 2023), for example, show year-on-year growth in online sales revenue of 9%, but growth in fees of 20% (third-party seller services) and 27% (advertising sales). However, the historical lessons from the downfall of both IBM's mainframe monopoly and Microsoft's stranglehold on the personal computer suggest that the company will be forced to renew its commitment to value creation or face decline and challenges from new, disruptive market entrants, who are focused on providing the kind of value to users and suppliers that Amazon once did. 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 to try alternatives? If history is any judge, it will happen sooner or later unless Amazon dials back the rents.

A similar dark pattern is visible in the evolution of Google search (O'Reilly 2019). Starting around 2011, advertising, which once framed the organic results and was clearly differentiated from them by colour, gradually became more dominant and the signaling that it was advertising became more subtle. Today, especially on mobile, the user may have to do a lot of scrolling to get to the first organic result. The result is less striking than on Amazon, since a very large percentage of Google searches carry no advertisements at all, but for commercial searches the best result for users (a local merchant, for example) can often be found only after scrolling through pages of ads from internet sellers and national chains.

Advertising can provide a useful complement to organic search results, surfacing additional products that might otherwise have been missed, but when it supplants organic search it becomes a kind of rent extraction, forcing vendors to advertise if they want their product or service to be visible at all.

We do know that organic search results are getting pushed further and further down the page. Search engine optimisation firm Moz.com found that in 2013, across a basket of 10,000 searches, the first organic result was found at an average of 375 pixels from the top of the window, but by 2020 that first result appeared, on average, at 616 pixels down (Meyers 2020). Not all this movement is due to increased ad load. Google has introduced many new features more beneficial to users in their quest for quick answers than traditional links to third-party websites. There are entire classes of commodity information, such as weather, stock prices, information about public figures and so on, where Google has determined that the best thing to do is simply to provide the information. In some cases, such as music lyrics, they may have licensed data from a third party so that they can present it directly and in others, such as travel search, they may run a booking service that competes with those operated by internet third parties.

Most often, though, the culprit is the increased prominence – both in size and position – of advertising. Keep in mind that this is only true on search result pages that have value to advertisers. According to Google, about 80% of search result pages have no ads at all. Of the remaining 20%, some may have only a single ad, while others, especially those where the commercial intent of the search can be inferred, carry a very high ad load.

The harms to users are thus less than they appear to be at Amazon, where advertising distorts the results of every search, but there are still serious concerns. Both Google and Amazon are gatekeepers controlling the visibility of a vast ecosystem of suppliers. It is important to take an ecosystem view of the total investment in value creation, rather than attributing all the value creation to the platform. Without websites, there would be no need for Google search or raw material for its results; without merchants, no Amazon. The same is true of other internet gatekeepers: without app developers, there would be no app stores; without users creating content as well as consuming it, no social media.

When these suppliers are harmed, in the long run users, too, will be harmed. These ecosystems of value creators depend on the platform's fairness in allocating attention to the most relevant results. When the platform displaces organic results with paid results, preferences its own applications, products or services, or provides information directly to the consumer in competition with the originators of that information, 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, and the whole virtuous circle of creation, aggregation and curation breaks down.

The company itself is also harmed, as even its own innovations may be held back to protect lucrative existing lines of business. Google, for example, invented the large language model (LLM) architecture that underlies today's disruptive AI startups. They published the original Transformer paper, *Attention is All You Need*, in 2017 (Vaswani et al. 2017) and released BERT, an open source implementation, in late 2018 (Devlin and Chang 2018), but they never went so far as to build and release anything like OpenAI's GPT line of services. It's unclear whether this was a lack of imagination or Google's own version of a 'strategy tax'. It was certainly obvious to outsiders how disruptive BERT could be to Google Search. In 2020, when my own company released O'Reilly Answers, a plain language search engine based on BERT, for the content on the O'Reilly platform, I was struck by how, for the first time, we could search our own content better than Google could.

It was left to startups to explore the broader possibilities of generative AI and chatbots.

5. Will history repeat itself?

The enshittification of Amazon and Google is old news to most users. We remember how good these services used to be and lament their decline, but we have slowly become used to the fact that results are not as good as they once were.

Antitrust authorities in Europe and the US have woken up, and are questioning abuses of market power by Big Tech companies, albeit not always successfully. Regulators may force better behaviour. In responding to this competition, companies themselves may wake up and pull back from the brink before it's too late, but it's already clear that LLMs may offer the greatest competition that Google, Amazon, and other current Internet giants have ever faced.

While the results are as yet inferior to those offered by Google and Amazon, users are already

asking questions of ChatGPT that would once have gone to a search engine. The lower quality of the results is typical for the early days of a disruptive technology. It doesn't matter, because disruptive technologies start out by solving new problems, serving new markets and creating new opportunities. However, their disruptive quality also comes because *they draw outside the lines that have been drawn to protect the business model of the existing players*. They are eager to surprise and delight their users; the focus in the early days is always on value creation. Mature and declining companies, by contrast, tend to hobble their products as they focus on value extraction, eventually alienating both their customers and their suppliers, and opening the door to competition.

We are in those early days once again. Leadership comes to those who create the most value for the most users. It is only later, after the market consolidates, that the value extraction phase begins. At that point, will the new market leaders also turn to more traditional extractive techniques? Just like today's incumbents, will they end up using their market power to protect their now-customary level of profits in the face of macroeconomic factors and competition that ought to be eating them away?

Regulators would be wise to get ahead of this development. The current generation of algorithmic overlords shape the attention of their users, helping to decide what we read and watch, and buy, whom we befriend and whom we believe. The next generation will shape human cognition, creativity and interaction even more profoundly.

There is a great deal of discussion about the risks and benefits of AI, but it is generally focused narrowly on the technical capabilities of AI tools and whether continued advances will eventually put AI beyond human control, leading to possible disaster. Closer to the present, risk analysis focuses on social problems like bias, misinformation, and hate speech, or the potential spread of biological and nuclear capabilities.

Yet many of the most pressing risks are economic, embedded in the financial aims of the companies that control and manage AI systems and services. (O'Reilly 2023). Are AI companies going to be immune to the incentives that have made today's current tech giants turn against their users and their suppliers; the incentives that led financial institutions to peddle bad assets, pharmaceutical companies to promote opioids, cigarette companies to hide the health risks of smoking and oil companies to deny climate change? I think not.

Rather than blaming the moral failings of company leadership, look instead to the economic incentives that rule public companies. Financial markets (including venture capitalists considering valuation of the next round) reward companies handsomely for outsized growth of revenue and profit, while brutally punishing any slowdown. Since stock options are a large part of executive compensation — and all compensation at Silicon Valley companies — failing to deliver the required growth comes at a very high cost to company leadership and employees.

It is too early to know best how to regulate AI, but one thing is certain: you can't regulate what you don't understand. Economic abuses by companies typically hide in plain sight for years, with whistleblowers, researchers, regulators and lawyers struggling to prove what the companies continue to deny. This is going to be even more true of an inscrutable black box like AI.

AI safety and governance will be impossible without robust and consistent institutions for disclosure and auditing. To achieve prosocial outcomes, AI model and application developers need to define the metrics that explicitly aim for those outcomes, and then measure and report the extent to which they have been achieved. These are not narrow technical disclosures of model capabilities, but the metrics the companies use to manage AI as a business, including what processes and metrics they use to reduce the risks that have been identified. *If they begin to twist its objectives for their own benefit, we should be able to see it in the numbers.*

The time to do this is now, when AI developers are still in the virtuous stage of innovation and rising tide rents, and while the companies are exploring the possibilities of AI regulation. It is important to understand what 'good' looks like while companies are still putting their best foot forward, developing services to delight and serve users and suppliers and society, so that if (or perhaps when) the incentives to take advantage of others take over, we can look back and see when and how things began to go wrong.

Let's not wait until the robber barons are back.

References

- Amazon. (2023). Amazon.com announces fourth quarter results. Business Wire. Available at: https://s2.q4cdn.com/299287126/files/doc_financials/2023/q4/AMZN-Q4-2023-Earnings-Release.pdf
- Bajarin, T. (2021). Attack of the clones: How IBM lost control of the PC market. Forbes. Available at: <https://www.forbes.com/sites/timbajarin/2021/08/25/attack-of-the-clones-how-ibm-lost-control-of-the-pc-market/?sh=740f79df5b81>
- Choudhary, R. (2015). Bill Gates' mail to Microsoft employees on its 40th anniversary. Available at: <https://twitter.com/roychoudhary/status/584044739000868865>
- Chposky, J. and Leonsis, T. (1988). Blue Magic: The People, Power and Politics Behind the IBM Personal Computer. New York: Facts on File.
- Christensen, C. (1997). The Innovator's Dilemma. Cambridge, MA: Harvard Business School Press.
- Conigliaro, A., Elman, J., Schreiber, J. and Small, T. (1996). Rise of Monopolies: The Building of Big Blue. Available at: https://cs.stanford.edu/people/eroberts/cs201/projects/corporate-monopolies/development_ibm.html
- Cortada, J. (2021). How the IBM PC won, then lost, the personal computer market. IEEE Spectrum, July 2021. Available at: <https://spectrum.ieee.org/how-the-ibm-pc-won-then-lost-the-personal-computer-market>
- Devlin, J. and Chang, M. (2018). Open sourcing BERT: State-of-the-art pre-training for natural language processing. Available at: <https://blog.research.google/2018/11/open-sourcing-bert-state-of-art-pre.html>
- Doctorow, C. (2024). 'Enshittification' is coming for absolutely everything. The Financial Times. Available at: <https://www.ft.com/content/6fb1602d-a08b-4a8c-bac0-047b7d64aba5>
- DOJ. (2006). U.S. V. Microsoft: Court's findings of fact. Available at: <https://www.justice.gov/atr/us-v-microsoft-courts-findings-fact>
- Dyson, G. (2012). Turing's Cathedral: The Origins of the Digital Universe. New York: Vintage.
- Hoogland W. and Weber, H. (2006). Statement concerning CERN W3 release into the public domain. Available at: <https://cds.cern.ch/record/1164399?ln=en>
- Janeway, W. (2018). Doing Capitalism in the Innovation Economy, 2nd edition. Cambridge, UK: Cambridge University Press.
- Kaziukenas, J. (2022). Amazon is burying organic search results. Available at: <https://www.marketplacepulse.com/articles/amazon-is-burying-organic-search-results>
- Koenig, J. (2011). The history of the IBM personal computer. Available at: <https://web.archive.org/web/20110921112147/http://computemagazine.com/the-history-of-the-ibm-personal-computer/>
- Maher, J. (2017). The complete history of the IBM PC, part one: The deal of the century. Available at: <https://arstechnica.com/gadgets/2017/06/ibm-pc-history-part-1/>
- Mazzucato, M. (2018). The Value of Everything. London, UK: Penguin.

Mazzucato, M., Ryan-Collins, J. and Gouzoulis, G. (2020). Theorising and mapping modern economic rents. UCL Institute for Innovation and Public Purpose, Working Paper Series (IIPP WP 2020-13).

Meyers, P. (2020). How low can #1 go? (2020 edition) moz.com. Available at: <https://moz.com/blog/how-low-can-number-one-go-2020>

O'Reilly, T. (1999). How the web was almost won. Salon. Available at: https://www.salon.com/1999/11/16/microsoft_servers/

O'Reilly, T. (2005). What is Web 2.0? Available at: <https://www.oreilly.com/pub/a/web2/archive/what-is-web-20.html>

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>

O'Reilly, T. (2021). Data is the new sand. The Information. Available at: <https://www.theinformation.com/articles/data-is-the-new-sand>

O'Reilly, T. (2023). The alignment problem is not new. Available at: <https://www.oreilly.com/radar/the-alignment-problem-is-not-new/>

O'Reilly, T., Strauss, I. and Mazzucato, M. (2024). Algorithmic attention rents: A theory of digital platform market power. Data & Policy, 2 February 2024, Volume 6, 2024, e6.

Obasanjo, D. (2004). The new strategy tax. Available at: <http://www.25hoursaday.com/weblog/2004/10/24/TheNewStrategyTax.aspx>

Perrin, A. and Duggan, M. (2015). Americans' Internet Access: Percent of Adults 2000-2015. Pew Research Center. Available at: <https://www.pewresearch.org/internet/2015/06/26/americans-internet-access-2000-2015/>

Pew Research Center. (2024). Internet/broadband fact sheet. Available at: <https://www.pewresearch.org/internet/fact-sheet/internet-broadband/>

Poushter, J., Gubbala, S. and Austin, S. (2024). 8 charts on technology use around the world. Pew Research Center. Available at: <https://www.pewresearch.org/short-reads/2024/02/05/8-charts-on-technology-use-around-the-world/>

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 2023-10).

Schumpeter, J.A. (1942). Capitalism, Socialism and Democracy. New York: Harper & Row.

Schwartz, P. (2020). Scenario planning for the long term. Long Now Foundation. Available at: <https://longnow.org/ideas/scenario-planning-for-the-long-term/>

Simon, H. (1971). Designing organizations for an information-rich world. In M. Greenberger (Ed.), Computers, Communications, and the Public Interest. Baltimore, MD: The Johns Hopkins Press. Available at: https://veryinteractive.net/pdfs/simon_designing-organizations-for-an-information-rich-world.pdf

Strauss, I., O'Reilly, T. and Mazzucato, M. (2023). Amazon's algorithmic rents: The economics of information on Amazon. UCL Institute for Innovation and Public Purpose, Working Paper Series (IIPP WP 2023-12).

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A., Kaiser, L. and

Polosukhin, I. (2017). Attention is all you need. Available at: <https://doi.org/10.48550/arXiv.1706.03762>

Wu, T. (2020). Tech dominance and the policeman at the elbow. Available at: https://scholarship.law.columbia.edu/cgi/viewcontent.cgi?article=3292&context=faculty_scholarship



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

@IIPP_UCL
ucl.ac.uk/IIPP