

To understand the risks posed by AI, follow the money

April 10 2024, by Tim O'Reilly, Ilan Strauss, Mariana Mazzucato and Rufus Rock



Credit: AI-generated image

Time and again, leading scientists, technologists, and philosophers have



made <u>spectacularly terrible guesses</u> about the direction of innovation. Even Einstein was not immune, <u>claiming</u>, "There is not the slightest indication that nuclear energy will ever be obtainable," just 10 years before Enrico Fermi completed construction of the first fission reactor in Chicago. Shortly thereafter, the consensus switched to fears of an imminent nuclear holocaust.

Similarly, today's experts warn that an <u>artificial general intelligence</u> (AGI) doomsday is imminent. Others retort that <u>large language models</u> (LLMs) have already reached <u>the peak</u> of their powers.

It's difficult to argue with David Collingridge's <u>influential thesis</u> that attempting to predict the risks posed by new technologies is a fool's errand. Given that our leading scientists and technologists are usually so mistaken about technological evolution, what chance do our policymakers have of effectively regulating the emerging technological risks from artificial intelligence (AI)?

We ought to heed Collingridge's warning that technology evolves in uncertain ways. However, there is one class of AI risk that is generally knowable in advance. These are risks stemming from misalignment between a company's economic incentives to profit from its proprietary AI model in a particular way and society's interests in how the AI model should be monetised and deployed.

The surest way to ignore such misalignment is by focusing exclusively on technical questions about AI model capabilities, divorced from the socioeconomic environment in which these models will operate and be designed for profit.

Focusing on the economic risks from AI is not simply about preventing



"monopoly," "self-preferencing," or "Big Tech dominance". It's about ensuring that the economic environment facilitating innovation is not incentivizing hard-to-predict technological risks as companies "move fast and break things" in a race for profit or market dominance.

It's also about ensuring that value from AI is widely shared, by preventing premature consolidation. We'll see more innovation if emerging AI tools are accessible to everyone, such that a dispersed ecosystem of new firms, start-ups, and AI tools can arise.

OpenAI is already becoming a dominant player with US\$2 billion (£1.6 billion) in annual <u>sales</u> and <u>millions</u> of users. Its GPT store and developer tools need to return value to those who create it in order to ensure ecosystems of innovation remain viable and dispersed.

By carefully interrogating the system of economic incentives underlying innovations and how technologies are monetised in practice, we can generate a better understanding of the risks, both economic and technological, nurtured by a market's structure. Market structure is not simply the number of firms, but the cost structure and economic incentives in the market that follow from the institutions, adjacent government regulations, and available financing.

Degrading quality for higher profit

It is instructive to consider how the algorithmic technologies that underpinned the aggregator platforms of old (think Amazon, Google and Facebook among others) initially deployed to benefit users, were eventually reprogrammed to increase profits for the platform.

The problems fostered by <u>social media, search, and recommendation</u> <u>algorithms</u> was never an engineering issue, but one of financial incentives (of profit growth) not aligning with algorithms' safe,



effective, and equitable deployment. As the saying goes: history doesn't necessarily repeat itself but it does rhyme.

To understand how platforms allocate value to themselves and what we can do about it, <u>we investigated</u> the role of algorithms, and the unique informational set-up of digital markets, in extracting so-called economic rents from users and producers on platforms. In economic theory, rents are "super-normal profits" (profits that are above what would be achievable in a competitive market) and reflect control over some scarce resource.

Importantly, rents are a pure return to ownership or some degree of monopoly power, rather than a return earned from producing something in a competitive market (such as many producers making and selling cars). For <u>digital platforms</u>, extracting digital rents usually entails degrading the quality of information shown to the user, on the basis of them "owning" access to a mass of customers.

For example, Amazon's millions of users rely on its product search algorithms to show them the best products available for sale, since they are unable to inspect each product individually. These algorithms save everyone time and money: by helping users navigate through thousands of products to find the ones with the highest quality and the lowest price, and by expanding the market reach of suppliers through Amazon's delivery infrastructure and immense customer network.

These platforms made markets more efficient and delivered enormous value both to users and to product suppliers. But over time, a misalignment between the initial promise of them providing user value and the need to expand profit margins as growth slows has driven bad platform behavior. Amazon's advertising business is a case in point.

Amazon's advertising



In <u>our research</u> on Amazon, we 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 but instead paid advertising placements. Amazon abuses the habituated trust that users have come to place in its algorithms, and instead allocates user attention and clicks to inferior quality, sponsored, information from which it <u>profits immensely</u>.

We found that, on average, the most-clicked sponsored products (advertisements) were 17% more expensive and 33% lower ranked according to Amazon's own quality, price, and popularity optimizing algorithms. And 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.

Amazon is one the most striking examples of a company pivoting away from its <u>original "virtuous" mission</u> ("to be the most customer-centric company on Earth") towards an extractive business model. But it is far from alone.

Google, Meta, and virtually all other major online aggregators have, over time, come to preference their economic interests over their <u>original</u> <u>promise to their users</u> and to their ecosystems of content and product suppliers or application developers. Science fiction writer and activist Cory Doctorow calls this the <u>"enshittification"</u> of Big Tech platforms.

But not all rents are bad. According to the economist Joseph Schumpeter, rents received by a firm from innovating can be <u>beneficial</u> <u>for society</u>. Big Tech's platforms got ahead through highly innovative, superior, algorithmic breakthroughs. The current market leaders in AI are doing the same.

So while Schumpeterian rents are real and justified, over time, and under



external financial pressure, market leaders began to use their algorithmic market power to capture a greater share of the value created by the ecosystem of advertisers, suppliers and users in order to keep profit growing.

User preferences were downgraded in algorithmic importance in favor of more profitable content. For <u>social media</u> platforms, this was addictive content to increase time spent on platform at any cost to <u>user</u> <u>health</u>. Meanwhile, the ultimate suppliers of value to their platform—the content creators, website owners and merchants—have had to hand over more of their returns to the platform owner. In the process, profits and profit margins have become concentrated in a few platforms' hands, making innovation by outside companies harder.

A platform compelling its ecosystem of firms to pay ever higher fees (in return for nothing of commensurate value on either side of the platform) cannot be justified. It is a red light that the platform has a degree of market power that it is exploiting to extract unearned rents. <u>Amazon's most recent quarterly disclosures</u> (Q4, 2023), shows year-on-year growth in online sales of 9%, but growth in fees of 20% (third-party seller services) and 27% (advertising sales).

What is important to remember in the context of risk and innovation is that this rent-extracting deployment of algorithmic technologies by Big Tech is not an unknowable risk, as identified by Collingridge. It is a predictable economic risk. The pursuit of profit via the exploitation of scarce resources under one's control is a story as old as commerce itself.

Technological safeguards on algorithms, as well as more detailed disclosure about how platforms were monetising their algorithms, may have prevented such behavior from taking place. Algorithms have become market gatekeepers and value allocators, and are now becoming producers and arbiters of knowledge.

Risks posed by the next generation of AI

The limits we place on algorithms and AI models will be instrumental to directing economic activity and human attention towards productive ends. But how much greater are the risks for the next generation of AI systems? They will shape not just what information is shown to us, but how we think and express ourselves. Centralisation of the power of AI in the hands of a few profit-driven entities that are likely to face future economic incentives for bad behavior is surely a bad idea.

Thankfully, society is not helpless in shaping the economic risks that invariably arise after each new innovation. Risks brought about from the economic environment in which innovation occurs are not immutable. Market structure is shaped by regulators and a platform's algorithmic institutions (especially its algorithms which make market-like allocations). Together, these factors influence how strong the network effects and economies of scale and scope are in a market, including the rewards to <u>market dominance</u>.

Technological mandates such as <u>interoperability</u>, which refers to the ability of different digital systems to work together seamlessly; or "side-loading", the practice of installing apps from sources other than a platform's official store, have shaped the fluidity of user mobility within and between markets, and in turn the ability for any dominant entity to durably exploit its users and ecosystem. The internet protocols helped keep the internet open instead of closed. Open source software enabled it to escape from under the thumb of the PC era's dominant monopoly. What role might interoperability and open source play in keeping the AI industry a more competitive and inclusive market?

<u>Disclosure</u> is another powerful market-shaping tool. Disclosures can require technology companies to provide transparent information and explanations about their products and monetisation strategies. Mandatory



disclosure of ad load and other operating metrics might have helped to prevent Facebook, for example, from exploiting its users' privacy in order to maximize ad dollars from harvesting each user's data.

But a lack of data portability, and an inability to independently audit Facebook's algorithms, meant that Facebook continued to benefit from its surveillance system for longer than it should have. Today, OpenAI and other leading AI model providers refuse to disclose their training data sets, while <u>questions arise about copyright infringement</u> and who should have the right to profit from AI-aided creative works. Disclosures and open technological standards are key steps to try and ensure the benefits from these emerging AI platforms are shared as widely as possible.

Market structure, and its impact on "who gets what and why", evolves as the technological basis for how firms are allowed to compete *in* a market evolves. So perhaps it is time to turn our regulatory gaze away from attempting to predict the specific risks that might arise as specific technologies develop. After all, even Einstein couldn't do that.

Instead, we should try to recalibrate the economic incentives underpinning today's innovations, away from risky uses of AI technology and towards open, accountable, AI algorithms that support and disperse value equitably. The sooner we acknowledge that technological risks are frequently an outgrowth of misaligned economic incentives, the more quickly we can work to avoid repeating the mistakes of the past.

We are not opposed to Amazon offering advertising services to firms on its third-party marketplace. An appropriate amount of advertising space can indeed help lesser-known businesses or products, with competitive offerings, to gain traction in a fair manner. But when advertising almost entirely displaces top-ranked organic product results, advertising becomes a rent extraction device for the platform.



An Amazon spokesperson said, "We disagree with a number of conclusions made in this research, which misrepresents and overstates the limited data it uses. It ignores that sales from independent sellers, which are growing faster than Amazon's own, contribute to revenue from services, and that many of our advertising services do not appear on the store.

"Amazon obsesses over making customers' lives easier and a big part of that is making sure customers can quickly and conveniently find and discover the products they want in our store. Advertisements have been an integral part of retail for many decades and anytime we include them they are clearly marked as 'Sponsored'. We provide a mix of organic and sponsored search results based on factors including relevance, popularity with customers, availability, price, and speed of delivery, along with helpful search filters to refine their results. We have also invested billions in the tools and services for sellers to help them grow and additional services such as advertising and logistics are entirely optional."

This article is republished from <u>The Conversation</u> under a Creative Commons license. Read the <u>original article</u>.

Provided by The Conversation

Citation: To understand the risks posed by AI, follow the money (2024, April 10) retrieved 20 May 2024 from <u>https://techxplore.com/news/2024-04-posed-ai-money.html</u>

This document is subject to copyright. Apart from any fair dealing for the purpose of private



study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.