

The Coming Enshittification of AI

Will AI follow internet search and e-commerce down the path of enshittification, or can we finally have nice things?

James Ryan

DOI: <https://doi.org/10.66241/i5yoq>



Abstract

Will AI serve the vast majority of internet users by making it easier to find the information, products, or services they are looking for? Or will it require us to wade through a cesspool of sponsored results, fake reviews, advertiser-influenced answers, and inefficient interfaces? Will it usher in an even worse internet, where AI is used to generate individualized content on the fly that is so subtly fine-tuned to

steer users towards sponsored content that they don't even notice they are being manipulated?

I wrote this piece to expand on Cory Doctorow's concept of "enshittification," a term he uses to describe how internet platforms inevitably go from being user-focused to monetizing every aspect of the experience, ultimately cannibalizing their own users and business customers before collapsing. I argue that generative AI tools like ChatGPT have succeeded so spectacularly because they're still in the early phase—where they provide straightforward, ad-free answers that genuinely help people, rather than steering them into sponsored links or frustratingly inaccurate search results. They are a breath of fresh air compared to the user experience of today's mainstream search engines.

I also share my vision of what could replace today's stagnating platforms. I imagine a future where AI "buyer's agents" and "seller's agents" allow individuals

and businesses to interact seamlessly, without user-hostile tactics driven by ad revenue. That may sound optimistic—enshittification has a lot of momentum behind it—but if AI can help us sidestep the entrenched models that drive pointless friction, maybe we can nudge the internet toward a genuinely user-first paradigm.

Pardon my French

In an article for Wired magazine in 2023, Cory Doctorow defined the concept of “enshittification,”^[1] which he had introduced in a blog post the previous year. If you haven’t come across the term previously, “enshittification” may strike you as overly vulgar, not befitting serious professional discourse, or just a joke in bad taste. However, I think it is a perfect term to capture both the horrid state of affairs at many of the most popular internet services, and the frustration users feel when using them. If you haven’t read Doctorow’s engaging and insightful article, it’s worth reading thoroughly from beginning to end. For purposes of this article, though, you only need to know Doctorow’s key concepts of how internet platforms tend to evolve:

1. First, they are good to their users;
2. Then they abuse their users to make things better for their business customers;
3. Finally, they abuse those business customers to claw back all the value for themselves.
4. Then, they die.

I argued in my last article^[2] that the rise of generative AI tools will soon cause most of today’s popular internet search and e-commerce sites to rapidly move into phase 4. Search engines and behemoth e-commerce marketplaces have stagnated in Phase 3 for years for the very reasons Doctorow describes in his Wired article. First, they were great at doing what users wanted. Search engines gave you the most accurate results in an easy-to-read format (remember the first few years of Google?), and e-commerce sites helped you find exactly the product you were looking for. Then, having captured huge audience groups, they began monetizing

those users through advertising and “sponsored” search results. It gradually became harder and harder to find what you were looking for. Today, the search experience on most major search engines and e-commerce sites is frustrating and needlessly inefficient. It’s no accident that the filtering options on e-commerce sites are so limited, making it virtually impossible to search for items by exact size, capabilities, formats, materials, etc. The more specifically they allow you to search, the harder it is to show “sponsored results” that match your criteria. It’s also no accident that e-commerce sites, by default, revert your sorting preference to “Featured” or “Recommended,” requiring you to manually change it back to “Price: Low to High” every single time. It’s a terrible user experience, but it makes these platforms tons of money.

And that’s the way it will remain until a new player disrupts these platforms, hastening their transition from Phase 3 to Phase 4. The problem is that the barrier to entry is so high, given the considerable investment needed to challenge these established players, that they (and their shareholders) have been enjoying the financial rewards of being firmly in Phase 3 of enshittification for at least a decade.

Then Came ChatGPT

Why was ChatGPT such an instant hit? Sure, the natural, conversational style of interacting with ChatGPT was novel and compelling, but I don’t think that’s the primary reason for its popularity. Despite its initial 2022 launch being a “research preview,” ChatGPT grew to 1 million users in only five days, and surpassed 100 million monthly active users in only two months. Today, two years after its preview launch, ChatGPT has nearly 500 million monthly active users.^[3]

Clearly, the novelty of a natural conversational interface is not a sufficient reason to explain such growth. I believe the primary reason ChatGPT (and other AI chatbots) took off so quickly and continue to attract both new and returning visitors is simple: they give you the information you are looking for, simply and accurately, without muddying up the experience with “sponsored” or manipulated results. In other words, they are firmly in Phase 1; **they are being good to their users.**

The End of the Ad-Based Business Model?

In my previous article, I predicted that generative AI tools would rapidly bring about the demise of the advertising-based business model supporting most of the internet economy. While I am not the first to predict this sort of disruption, most others have focused on a different aspect of the potential disruption that AI may bring about. Specifically, if users can just get answers to their questions directly in the chatbot interface, rather than searching for a subject on a search engine and then clicking on the resulting links to read the individual pieces of content on external websites, then traffic to these websites might plummet, threatening their economic viability^{[4][5][6]}. This phenomenon is already well advanced in particular niches where websites are designed to answer specific questions, such as the coding advice site Stack Overflow, whose traffic has declined from 213 million visits in March 2021 to just 90 million in November 2024.^[7]

While many authors have focused on the economic impact on content creators (not surprising, given that they are themselves content creators), I predict something bigger: the demise of traditional internet search altogether.

To understand this, one needs to ask a simple question. Why didn't Google own the AI chatbot market from the beginning? After all, it was Google who, in 2017, invented the transformer model that makes all modern chatbots possible, including ChatGPT.^[8] The "T" in ChatGPT stands for "transformer." So why, five years after Google invented the transformer, did OpenAI pull off a coup of historic proportions, beating Google to become the undisputed champion of the transformer-based chatbot market? ChatGPT had nearly 500 million active users and 1.4 billion visits in November 2024, while Google Gemini had only 80 million active users and 135 million visits in the same month (according to SEMRUSH).^[9] And those ChatGPT users aren't using ChatGPT *in addition to* traditional search, they are largely using it *instead of* traditional search. In fact, 56% of teens and 55% of parents say they prefer getting answers from an AI chatbot over a traditional search engine.^[10]

Many articles attribute ChatGPT's early success to first-mover advantage, combined with simply being a superior product. When ChatGPT was released, Google was struggling to get Bard (the precursor to Gemini) to work reliably.

Although Google responded quickly to the release of ChatGPT, pushing Bard out to the market about two months later, its launch was plagued by quality and accuracy issues, leading to Google's share price dropping by 7.7% (worth \$100 billion).^{[11][12]}

So why was Google so far behind an (admittedly well-funded) upstart? They had the financial resources and the talent, so what happened?

It's possible that Google's team just had bad luck, while OpenAI's team had good luck. Google's team may have just run into more technical hurdles or pursued more dead ends than OpenAI's team. However, given Google's overwhelming technical head start as the inventor of the transformer, I think we need to look elsewhere.

To be clear, I have not spoken with any current or former Google employees, so what follows is pure speculation on my part (though Jeremy Kahn at Fortune published an early analysis along these same lines over a year ago).^[13] Google fell behind, and has remained behind, upstart OpenAI because of Clayton Christensen's classic Innovator's Dilemma.^[14] Search is Google's cash cow—fully three-quarters of Google's revenue comes from search and advertising. This is a powerful incentive to deprioritize and underfund anything that might threaten that business.

I can't imagine that no one at Google recognized the enormous disruption potential of transformers six years ago when Google invented them. What would have happened if Google had taken a page out of IBM's playbook and set up a secret, off-site project to explore the full potential of AI, without regard for its effects on current revenue models, as IBM did when inventing the Personal Computer back in 1981? Despite this example, and the popularity of Christensen's book, it is still exceedingly rare for established companies to recognize and rise to the challenge of their own Innovator's Dilemmas.

The Second Law of Shittydynamics

I therefore humbly propose an addition to Doctorow's concept, along the lines of the Second Law of Thermodynamics. That is, there is an entropy to enshittification. It runs very easily in the forward direction, but requires

enormous amounts of energy to run in the reverse direction. That energy can be supplied in the form of competitive pressure, regulation, or (rarely) visionary leadership. However, the forces of enshittification entropy driving the process forward, including user “lock-in,” shareholder revenue expectations, and technological and organizational path dependence, are almost impossible to overcome. Viewed in this light, it’s not at all surprising that the AI chatbot market was pioneered by a new player, a majority of whose funding came from an existing player (Microsoft) whose search business could be generously described as “anemic,” but who had a strong desire to topple Google.

Past and Future Predictions

In my article on how AI will disrupt internet business models from May of this year, I made two predictions—that AI would enter both the search and the e-commerce markets with disruptive new products and business models. In the six months since that time, both have happened. Although Bing had already launched real-time AI search as of the time of my article, this is simply an AI interface tacked on to an existing search engine.

Then, just a few weeks ago, AI search startup Exa announced the launch of its first consumer-facing AI-based web search engine that does not rely on existing search engines to do the actual searching.^[15] Exa’s “Websets” search engine turns the entire internet into an indexed database that you can search using natural language prompts via its AI interface. Exa’s back-end infrastructure actually crawls the web and indexes content similarly to how traditional search engines work, but using its own algorithm that is specifically tuned to giving the most precise results, and that indexes the information in a format that is custom-designed to make it easily processed by large language models. This would make it one of the first “native” AI search engines, and represents a true paradigm shift.

Similarly, Perplexity has launched its own shopping service called “Shop with Perplexity.” When I first read about this, I had dreams of finally being able to search for specific products with specific features and specifications, with the best reviews and ratings, and finding the best deals on the products without having to wade through sponsored results and results that don’t match my search

parameters. The reality of Shop with Perplexity is... not quite there yet, but a good start. You can ask it, in natural language, to look for specific products. It does make mistakes. It will give you, for example, 4-foot and 6-foot ladders when you've searched for a 12-foot ladder. This may be due to the poor quality of the product descriptions it is searching through.

Perplexity does not (to my knowledge) serve any sponsored results within the response itself (however, it does serve side-bar ads and sponsored "follow-up queries"). I asked Perplexity itself (the chatbot) how they make money from Shop with Perplexity, and its response was eerily honest about the potential for future enshittification:

"Although Perplexity isn't currently taking commissions, future plans may include affiliate marketing or other forms of revenue generation related to e-commerce. For now, Perplexity seems to be focusing on growing its user base and increasing search queries rather than immediate monetization of the shopping feature. Perplexity has launched a free merchant program for larger retailers, which may lead to future monetization opportunities."

Future monetization opportunities, indeed. Time will tell, but given that Perplexity is already embracing legacy monetization strategies like sponsored side-bar content and sponsored follow-up queries tacked to the end of searches, the outlook for them avoiding enshittification seems somewhat pessimistic.

As for my next prediction, I would prefer to classify it as an aspiration, since I think the probability of it coming to fruition is low. But if it does...

Imagine all the Agents Living Life in Peace

John Lennon asked us to "Imagine" a utopian world where people just get along. While my aspiration is not quite so grandiose, it is analogous since it sees a future internet where there don't have to be winners and losers.

Our current system for obtaining useful information is flawed. While users simply want to find the most pertinent, accurate, and useful information as quickly and efficiently as possible, the forces of enshittification work against this. As discussed, forces that propel enshittification are almost impossible to resist, even

when founders and executives genuinely want to do so. If you believe the purpose of the internet is to optimize the extraction of value from users on behalf of shareholders, then it is extremely efficient. However, if you view the internet as a tool to help businesses and individuals produce and accomplish more with less time and resources, then this is at odds with the first purpose.

Currently, the first purpose is dominant. But AI has the potential to change that.

Today, most manufacturers don't sell their products directly to consumers. There are various reasons for this, but primarily, it is due to the cost of running a business-to-consumer operation. You have to set up an e-commerce site where customers can easily purchase your products using a variety of payment methods, you have to manage packaging and shipping individual units to individual residential addresses around the world, you have to provide customer support to potentially millions of individual customers, and you have to get the word out to your potential customers about all of your amazing products. For most manufacturers, it's easier to just rely on an e-commerce marketplace to do all of that for you. You make less money, but it's quicker and easier, and all the customers are already there. Until recently, it's been a fairly good bargain.

However, according to Cory Doctorow, sellers on Amazon now give up more than 45% of the sale price of their items to Amazon in various "junk fees." He goes on to say, "The company's \$31 billion 'advertising' program is really a payola scheme that pits sellers against each other, forcing them to bid on the chance to be at the top of your search."^[16] (I have not verified this independently with Amazon.) For manufacturers and other third-party sellers, in the words of Lando Calrissian, "this deal is getting worse all the time." At some point, they may decide, like Calrissian, that the downsides of maintaining the deal with "the Empire" outweigh the benefits.

Imagine, instead, a world of AI agents—buyers' agents and sellers' agents. This applies equally to buying and selling information as it does to products and services. Imagine every producer of information, services and products has AI agents that act as an autonomous sales and marketing team, getting the word out to buyers' agents about what they have on offer. Imagine every consumer (individual and business) of information, products, and services having their own

team of automated AI buyer's agents who know exactly what the buyers want, their budget, etc. The buyers' and sellers' agents could negotiate billions of times a second, scouring the planet for the best deals for the parties they represent, cross-referencing with other buyers' agents to confirm product quality and seller reputation. Dishonest agents would be punished with bad reputations and eventually find it hard to do business. The buyers' agents could then work with fulfillment agents to find the optimum way to purchase and ship the products based on the buyer's priorities. All of this would happen autonomously and, from the perspective of the human user, instantaneously, enabled by open APIs that agents can access to order, pay for, and arrange shipping for products and services.

It would be the ultimate disintermediation. E-commerce sites promised to disintermediate us from the layers of distributors, wholesalers, and retailers who drove up the cost of goods and services, only to become extractive intermediaries themselves. Meet the new boss, same as the old boss...

To be sure, there would be a constant state of war in the background, with agents on both sides trying to gain an advantage. The difference is that it would be a fair battlefield. Unlike today, where platform capture enables enshittification, leading to inefficiency and increased cost to both buyers and sellers, under an agent model, there is no platform to be locked into.

Amazon is in the best position to own this future of e-commerce, if they can overcome the Innovator's Dilemma. Jeff Bezos' famous "Bezos Mandate," which he issued in 2002, ordered that all teams must expose their data and functionality via APIs, and that all APIs must be designed to be easily exposed to external third parties. Furthermore, the corporate culture at Amazon encourages risk-taking and a "Day 1 Mentality," meaning it should always act like a startup. Although Amazon's current e-commerce business model delivers enviable shareholder returns, there is reason to believe that Amazon may be at the forefront of deploying new business models that overturn its existing ones. At any rate, if they don't, someone else likely will.

Perhaps this is what the rise of AI agents will lead to. This could even "solve" the age-old problem of micropayments for content monetization, which I described in

my previous article. The biggest reason we ended up with the advertising-based internet is that, in the 1990s, it was too complicated and expensive to implement the micropayments that would enable users to pay tiny amounts of money for each website they visited.^[17] Today, even without requiring the use of cryptocurrencies, it is viable to offer a clearinghouse payments service that aggregates micropayments and distributes them periodically when the balance exceeds a threshold that justifies the transfer cost. So, even if a content creator charged \$0.25 per view to read their content, the payment clearinghouse could “hold on” to those payments until they exceeded \$25, for example, and then transfer the total using the existing financial system.

As I said, these are aspirations, more than predictions. It is also possible that AI agents will be deployed in a way that reinforces enshittification. The big players could build audiences by offering valuable services (Phase 1) and then start “monetizing” with sponsored links, sidebar ads, and then, eventually, answers influenced by sponsors in subtle ways that the user isn’t even aware of. If these tactics are successful, advertisers would have no choice but to jump on board. Welcome to the AI enshittification dystopia where all results are sponsored!

But I would rather live in a world where, instead of having to close four popups, shut off a video ad, and scroll through five other ads just to read a single article, I could just read it, knowing the creator was paid fairly. I would rather live in a world where search engines and e-commerce sites give me exactly the results I am looking for, the first time, rather than the results the highest bidding advertiser paid them to show me.

Two roads diverged in a yellow wood...

And so, like Robert Frost over 100 years ago, our AI agents stand at a fork in the path. In one direction lies the familiar road to enshittification. In the other, something new. A new business model for the internet that doesn’t sacrifice user experience for advertiser dollars; that doesn’t sell user frustration and wasted time for shareholder returns, but, instead, a model that works for everyone—users, advertisers, and shareholders. Let’s hope they take the path less traveled by; it will make all the difference.

References

1. [The 'Enshittification' of TikTok | WIRED](#) ↑
2. [Will AI Finally Destroy the Advertising-Based Business Model of the Internet? When AI agents replace search engines, will they give us "sponsored content" or the best answer to our queries? | The Journal of Business and Artificial Intelligence](#) ↑
3. [openai.com: Overview, Traffic Analytics](#) ↑
4. [Will AI Kill Web Traffic? | Shelly Palmer](#) ↑
5. [Gartner Predicts Search Engine Volume Will Drop 25% by 2026, Due to AI Chatbots and Other Virtual Agents](#) ↑
6. [AI search could break the web | MIT Technology Review](#) ↑
7. [Stack Overflow is ChatGPT Casualty: Traffic Down 14% in March | Similarweb](#) ↑
8. [Attention is All You Need](#) ↑
9. [openai.com: Compare Mode, Traffic Analytics](#) ↑
10. [What's an AI Chatbot? Everything to Know - CNET](#) ↑
11. [Google Bard: A Rival to ChatGPT or Just Another Shallow Imitation?](#) ↑
12. [ChatGPT vs Google's Bard: A Side-by-Side Comparison and Why It Matters | Join Me for an AI Happy Hour in SF!](#) ↑
13. [Will Google A.I. and Bard threaten Google Search? | Fortune](#) ↑
14. Christensen, Clayton M. *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*. Boston: Harvard Business School Press, 1997. ↑
15. [The startup trying to turn the web into a database | MIT Technology Review](#) ↑
16. [The 'Enshittification' of TikTok | WIRED](#) ↑

About the Author



James Ryan

Editor in Chief at Journal of Business and Artificial Intelligence

James Ryan is Editor-in-Chief of the Journal of Business and Artificial Intelligence. He is an AI Advisor at the Harvard Innovation Labs, where he advises Harvard startups on AI strategy, product development, and market entry. He is Senior AI Advisor at AI Transformations, and on the board of directors of the Keep AI Safe Foundation. James has over 30 years of experience building and running global P&Ls for high-tech startups focused on leveraging technology to improve operational efficiency and revenue growth. He holds a master's degree in behavioral sciences from Osaka University and a bachelor's degree in East Asian languages from Harvard University, and has worked in 39 countries as a native speaker of Japanese and English.

[ORCID: 0009-0005-6720-0642](#) · [LinkedIn](#) · [Website](#)

Copyright Notice

All articles are published in the Journal of Business and Artificial Intelligence under the Creative Commons 'CC BY' ("Gold Open Access") license, where authors retain the copyright of their articles. The author grants JBAI a license to publish the article under a Creative Commons 'CC BY' license, which allows the work to be freely accessed, shared, and used under certain conditions.

About the Journal

The Journal of Business and Artificial Intelligence (ISSN: 2995-5971) is the leading publication at the nexus of artificial intelligence (AI) and business practices. Our primary goal is to serve as a premier forum for the dissemination of practical, case-study-based insights into how AI can be effectively applied to various business problems. The journal focuses on a wide array of topics, including product development, market research, discovery, sales & marketing, compliance, and manufacturing & supply chain. By providing in-depth analyses and showcasing innovative applications of AI, we seek to guide businesses in harnessing AI's potential to optimize their operations and strategies.