

PUBLIC NOTE

Where are the Customers' Bots?

Peter Lewis

peter@plraresearch.com



For investors in consumer-facing sectors, the most important question about agentic AI is where it will drive increased price discovery. But even as this discourse heats up, it still has some major blind spots. In this note I'll try to frame the topic from the point of view of a real consumer, and ask some of the uncomfortable questions that aren't being asked.

We'll start with an example. If you want to get the best possible rate at a particular hotel, here are some things you might want your AI agent to do for you:

- Log onto the websites for all your credit cards, checking their coupon-book special offers to see if there's one that applies. (" \$100 back on \$500 of spend at Hyatt resorts in Latin America during March and April.")
- Log into the brand's website, and every online travel platform, in order to see the best rates or specials available to logged-in members
- Create these accounts wherever they're free, and you don't already have one
- Sign you up for email and text marketing with all of the above, but use filtering (and masked emails, etc.) to monitor those offers while hiding them from you
- Compare notes with other customers' agents on how deals and special offers are being timed and targeted. Somewhat like the subreddits and forums for frequent flyers and travel hackers, which of course your agent would also be following so you don't have to...

In other words, you want an agent that will take YOUR side in this algorithmic swamp of surveillance pricing, and work to save you money in the same ways that the swamp is working to extract it.

Needless to say, that is not what any of these companies want you to do, and for the most part they're still pretending that it's not even what you want. For example, this is Amazon CEO Andy Jassy, on their most recent earnings call:

We've invested [in] our own shopping assistant—and if you haven't checked out Rufus recently, I really encourage you to do so... at the same time, we will have relationships with third-party horizontal agents that can go shopping as well. [But] we have to try to find a customer experience together that's better, and **a value exchange to make sense for both parties...**

[And] you have to ask **how many consumers are going to prefer using a horizontal agent, where it's kind of a middle person between the retailer and the consumer, vs. wanting to use a great agent from that retailer** that has all its shopping history and then has all the data right there...

Now, I'm not saying there's anything sinister about this. No incumbent ever wants to be commoditized. Nor am I suggesting that non-price factors like convenience don't matter; of course they do. But let's spell out the subtext here:

1. The main reason that older comparison-shopping tools failed is not that consumers didn't want them. It's that they were obfuscated, blocked and sued by the sites they were trying to scrape, or taken out in anti-competitive acquisitions... or most often, simply coopted with advertising and partnerships.
2. The main plan to stop new AI comparison-shopping tools is to run the same playbook again, not to outcompete them in a transparent market. Here is Chris Nassetta of Hilton, putting it a little more bluntly:

I mean, look at the US alone, we're 13-plus percent of the market. If you look at the quality market, [then] we're well over 20% of the market. **We have complete control over rate, inventory, pricing, availability. And if we don't want to share it, nobody can get it.**

3. This playbook is already working in the first phase of agentic AI, because the AI providers have raised so much capital that they can't really afford to be disruptors, right? Even when they're not actively selling out their users already, they have every incentive to respect bot restrictions and terms of service, if only to avoid a legal battle.

Again, I'm not picking on Amazon or Hilton. If anything, I've chosen two current winners where much of their success really has been aligned with consumer preferences, and more of the returns have come out of competitors or other counterparties. In any case, I'm not writing a pro-consumer polemic,

like Cory Doctorow's (excellent) "enshittification" article. We're trying to think about what will happen, not what should happen.

So let's just skip all the marketing baby talk about this imaginary customer that's actually happy about every purchase turning into an algorithmic slot machine. Let's stop pretending that "agentic AI" is mainly about proprietary customer-facing chatbots, or that these tools are the kind of "game changer" that would justify all the AI hype.

Let's even skip the more serious intermediate questions, like the ones about human ad impressions. Let's get straight to the second phase, when the real disruption arrives. Where will *those* horizontal agents come from?

1. THEY'RE ALREADY HERE

This is the scenario I just dismissed above, but maybe I'm being too cynical. For example, I just saw that Claude can now take over a logged-in browser state, via Cowork and a Chrome extension. I haven't tried it for the kind of comparison shopping routine that we laid out at the top, but it's a partial step in that direction.

Now, if it works smoothly on a site like Amazon or Hilton that doesn't want it to work, then I would not expect it to work for very long, because that would be an odd hill for Anthropic to die on. As one of many recent Claude Code junkies, I wouldn't even want them to.

That said, maybe all this stuff is moving so fast that we'll have a meaningful window where enough power users can set up that kind of true agentic workflow for price comparison, without moving the needle enough for large sellers to react. Maybe by that point, they won't know how to react. You never know.

2. NEW STARTUPS

But more likely, the threat would come from newer AI providers that are built to push the envelope on licensing and terms of use, with less to lose on the legal front. This would be more like Uber, Airbnb and other "break the rules to change them" strategies from prior cycles. Rather than targeting municipal governments, they'd be deploying the same three-pronged legal/lobbying/PR attack against the online platforms trying to shut them down.

And the PR element would be even more important, because these platforms are sensitive about customer perception, and you can always find weak ones to start with. Imagine an ad campaign that says "Why is Target.com blocking our agent from logging in for you and finding the best deals?" After everything else Target has been through recently, is that a battle they want to fight?

You also don't have to get 100% buy-in, just a tipping point on customer volume where the *next* holdout can't ignore you. If you don't remember these dynamics

from the peak of the online travel wars, just think of all the major brands like Nike that are now caving in to Amazon distribution, after holding out as long as they could.

In the end, the “regulation” might come from the same few dominant platforms like Apple or Google, who have more plausible deniability in selling out their users. For example, why isn’t there a popular meta-app that opens Uber and Lyft on your iPhone, and checks the best price for every ride? If it’s really just a sandboxing/security thing, then shouldn’t Apple be building that kind of feature directly into Apple Intelligence, rather than all this sad little slop like summarizing your text messages? Don’t hold your breath.

3. ARBITRAGE

From my 2024 note here on e-commerce, and the “algorithmic swamp” of third-party marketplaces:

One of my recent Walmart orders arrived in a Home Depot box, from a Home Depot warehouse. I’ve had a few similar funny experiences with drop shippers on Amazon or eBay, and maybe you have too. In this case, it looks like they scanned for products offered by Home Depot with free shipping, and marked them up as a third-party Walmart seller.

With a small item like this one, where the markup was only a dollar, you might wonder if this arbitrage is worth anyone’s time. But of course, there’s automation on this side of things too, and it’s a volume game. So this is one of the simplest possible examples, but there are many other mechanisms keeping the ecosystem of this online swamp in balance... [and] **even for customers who have given up on doing any comparison shopping, there’s still a rough limit to the amount of price discrimination that any brand can realize just by tweaking their offering across different sites.**

This is the kind of meta-competition that’s always happening in the background, and it’s certainly already being impacted by AI. It’s the hardest of these scenarios to forecast or even measure, because it’s not even a discrete “scenario” as much as a property of the system; the question is how AI adoption could be turning the dials and speeding it up. Candidly, I’m still not sure how to think about that.

4. THE WILD WEST

This would be something like the Napster era: a rare moment of distributed mass adoption for buggy gray market tech (at least among younger cohorts) that ultimately forced the record studios into a paradigm shift they weren’t ready for. For online retail or travel booking in 2026, it might look like an open-source or foreign LLM that’s optimized for a continuous arms race against bot detection. If enough consumers are fed up by AI tools that fail to live up to their promises, it’s not crazy to imagine something like this catching on.

It's also the scariest scenario in some ways, right? Comparison shopping is a more sympathetic use case than college students downloading free music, but most current bot use cases are less sympathetic—like cheating at video games, scalping tickets, fake reviews, political propaganda, and so on.

Some potential use cases are even scarier—and if we're trying to think about the equilibrium, with AI bot detection fighting AI bots, then we're really asking whether this nets out to a more offensive or defensive technology, like the evolving military theory about drones. From what I've read so far, AI certainly seems to favor the attacker.

But at the same time, this is the grim new Turing test that emerges naturally from our current walled garden ecosystem. If every recognized AI agent is either shut out of that system or compromised by it, then the usefulness of an agent will be directly proportional to its ability to convince another computer that it's a human.

CONCLUSIONS

Hopefully it's clear that this is more of a thought exercise than the kind of piece with a firm conclusion. I don't have a strong view on the odds of each scenario above—and of course, there are many non-consumer facing AI applications that could ultimately be more important.

But on the consumer side, it's remarkable how much corporate and investment “thought leadership” on AI is still just a long-winded way of saying that it doesn't matter, and it won't lead to any real competitive disruption—just a new wrapper for the same market structures. The better these tools get, the less serious that sounds.

UPDATE: For more on scenario #4, see [The New Consumer Turing Test](#)

PLRA is a paid research publication for institutional investors, but you can share these public notes freely, and sign up at plraresearch.com to receive them directly.