

## PUBLIC NOTE

## Three Consumer AI Anecdotes

**Peter Lewis**  
*peter@plraresearch.com*



There's a shortage of real anecdotal evidence in AI discourse, and it's a funny contrast with other business topics. I guess we all have deep thoughts about AI that we'd rather debate instead, or maybe a little FOMO about how others are using it.

So I want to give you three examples of recent experiences that happened sooner than I would have thought, with a little context but no grand conclusions. They are small and relatable, and I'm not sure what to make of them myself.

### 1. SHOPPING

A few weeks ago I asked ChatGPT for headphone recommendations, and looked at its shopping links for the model it recommended. They're still a work in progress:

|                                                                                                                                                          |                                 |                       |
|----------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|-----------------------|
|  Walmart - Venus                                                      | \$99.99                         | <a href="#">Visit</a> |
|  Target<br>In stock                                                   | <del>\$119.99</del><br>\$109.99 | <a href="#">Visit</a> |
|  eBay<br>In stock online, Delivery \$44<br><a href="#">Best price</a> | \$65.00                         | <a href="#">Visit</a> |
|  TikTok Shop<br>In stock online, Free delivery between<br>Fri - Sat   | \$84.99                         | <a href="#">Visit</a> |

For example, they pull in this outsized shipping charge from eBay and free shipping from the TikTok shop, but they don't reflect them in the "best price." And you might think "Venus" is the town with the nearest Walmart, though it's actually the name of the seller on their third-party marketplace.

Well, by the standards of our current [algorithmic swamp](#), it could be even worse. And I almost clicked the Walmart link... but then I remembered it was Amazon Prime Week, and thought "this is the kind of thing they'd have on a flash sale." Which they did, and I paid \$63.

I'm not a representative shopper, if only because I follow these companies for work. I know a little about their OpenAI integrations and why Amazon doesn't have one; I've [written about](#) how third-party agents could drive increased competition. But I don't think any of that changed my behavior here, and I could spin this anecdote to support any number of other takes on agentic commerce.

Amazon is the most apparent winner: not only capturing this sale to a Prime member, but also reinforcing (by a mile) my prior behavior of going straight to Amazon for something like this. On the other hand, they almost lost the sale to my overlapping Walmart+ membership, and I know Walmart may have had a similar flash discount that they didn't expose here.

Whoever wins the transaction, they're also both losing the search: it's been a long time since I would *begin* a generic search on Amazon, given the ad-driven degradation in their search results. But even if those long-degraded search results at Google, Amazon et al. are helping LLMs take over quickly as a search/discovery surface... is it a positive sign for ChatGPT in particular that I was even looking at these price cards, or a negative sign that I'm already learning not to trust them?

We're not even talking about a static price table, right? This Amazon seller ("BSD wholesale") is likely tweaking the same listing across multiple platforms, and learning that they might have gotten another \$20 out of me on a different one, perhaps by out-tweaking another seller like Venus.

So the truth is that we're layering AI into an increasingly complex market with accelerating feedback, and more ways to be wrong than right. For example, one easy way to be wrong a year ago was to say that consumers won't trust an agent to shop independently or complete a purchase, as a way to dismiss the whole stack; we know that even if that's true, our purchase decisions can still be influenced in all these other ways. But if you asked me even three months ago, I would not have expected to feel the ground shift quite this soon, with LLMs already capturing this much of my own purchase funnel.

---


## 2. CREDIT CARDS

---

This one is from a friend in tech who had read my prior notes on consumer agents. She sent me a spreadsheet that her mother (74) was building with Claude to optimize spending across two premium cards, in order to maximize rewards/

miles. It wasn't perfect, and at first glance it looks like the kind of comparison that you'd already find at sites like ThePointsGuy. But it's very different to help an individual user with their current cards, without those affiliate-driven incentives to keep pushing new ones.

Now, if you think Amex (or Delta, Marriott, whoever) has been squeezing customers too hard with pointflation and time-wasting gimmicks, then it's no surprise that they're vulnerable to this kind of AI-driven counter-optimization...



...but again, that pro-consumer story is not the only one you can tell. The incumbents will try to co-opt the commercial agents the same way they did with Google, and it's even possible that these individual optimizers will become *easier* to exploit with targeted offers and surveillance pricing. All I can tell you is that even as someone who was predicting this kind of counter-optimization, it seems to be picking up faster than I expected.

### 3. ONLINE WRITING

I do a lot of [paid research](#) and writing for institutional investors, but I've never been organized about public writing like this, and I haven't stuck to a single platform or mailing list. I'm grateful when it finds the right readers anyway, as in point #2 above. But I've recently been setting up some basic cross-posting scripts with Claude Code, so at least the business-oriented notes go up on a few platforms at once.

The part that's surprised me is how rapidly all these platforms have realigned for me—almost overnight—between the ones like Ghost that expose a clean API, and the ones like Substack that don't.

It doesn't map perfectly to open vs. closed source, and a few (like LinkedIn) are sort of in between. But Medium is an especially funny case; it's been declining

for years and filling up with pre-AI slop, but that wasn't enough to make me stop using it for the occasional post that didn't fit elsewhere.

And yet, as soon as I learned that Claude couldn't tee up a whole Medium post for me... well, I could still use [browser automation](#) or a third-party cross-posting tool, but you might say my laziness has been harnessed in the other direction. And my perception of the Medium platform has gone from well-intentioned and bumbling to adversarial and extractive, even though nothing has changed on their end.

Is this a third sign that walled gardens will struggle to keep their walls up in an agentic world? I don't know, and they will certainly try. But this perceptual shift has been the biggest surprise of the three from my end.

*PLRA is a paid research publication for institutional investors, but you can share these public notes freely, and sign up at [plraresearch.com](https://plraresearch.com) to receive them directly.*

Cover photo by [Aleksandar Spasojevic](#)