



Published July 2026 | Authorship Aidan Smith

MEMO

AGENTIC AI ADOPTION, PLATFORM POWER, AND CONSUMER RISK

ASSESSING COMPETITION
CONCERNS AND THE NEED
FOR PRIVACY GUADRAILS



LABYRINTH
INSIGHTS

TABLE OF CONTENTS

I. Introduction	2
II. Competition Implications	3
CONFLICT OF INTEREST ISSUES	3
SURVEILLANCE PRICING CONCERNS	3
III. Privacy and Safety Risks	5
OPERATIONAL RISKS	5
THREATS TO USERS' FINANCIAL DATA	5
FRAMEWORK FOR FEDERAL GUARDRAILS	6
IV. Conclusion	7
AUTHORSHIP AND PROJECT INFORMATION	7

I. Introduction

The release of ChatGPT in 2022 and the subsequent proliferation of AI chatbots prompted immediate public scrutiny, with concerns ranging from misinformation and mental health-related harms to more dire scenarios, such as their potential use for acts of [terrorism](#). In the year that followed, public and lawmaker scrutiny was strong enough to warrant White House [executive orders](#), congressional [hearings](#), and attention from agencies such as the [FTC](#). Though this did not culminate in meaningful legislative reform, a result unfortunately consistent with Washington's long history of [inaction](#) on tech policy, it is clear that the dawn of the generative AI era made tech policy issues far more salient for the general public.

By 2024, the tech industry had pivoted from chatbots to agentic AI services, systems capable of executing complex tasks and transactions with minimal user intervention. Since 2025, Big Tech companies including Amazon (via its agentic [revamp](#) of Alexa), Google ([Universal Cart](#)), and Meta Platforms (via its [enterprise](#) agentic venture) are among those who have debuted major agentic tools. But despite mounting [consumer](#) and enterprise adoption of AI agents alike, public concern over agentic AI's potentially harmful impact is seemingly minimal compared to those that have accompanied generative AI services like chatbots and video generators.

In 2022 and 2023, early generative AI development led commentators to raise concerns about more extreme harmful use cases, such as chatbots being used to assist [bioterrorism](#) efforts. It is somewhat peculiar, then, that similar alarm has not accompanied the proliferation of AI agents, which are capable of taking consequential actions *without* user participation. (Neither catastrophic scenario should be the focal point of AI policy-making.)

As far as immediate, real-world harm is concerned, consumer advocates are increasingly drawing attention to agentic AI's potential to work against the financial interests of the very users it purports to serve, all while proving lucrative to the platforms that deploy them. Beyond financial harms, the proliferation of agentic AI without strong regulatory safeguards poses a direct threat to users' personal privacy and security. In this context, a new discussion draft proposal for a federal agentic AI [framework](#) being circulated by Sen. Mark Warner (D-VA) is both welcome and overdue.

II. Competition Implications

CONFLICT OF INTEREST ISSUES

In 2022, most commentators predicted that the generative AI boom would disrupt the dominance of existing tech giants and help build a more dynamic tech sector in the process. On the contrary, it's clear that AI development has served to [entrench the dominance](#) of the very same companies who dominated the industry at the beginning of the generative AI era.

As it stands, agentic AI deployment by Big Tech companies stands to have a similar impact. Agentic AI services are marketed as tools to empower consumers by helping users find ideal products at the best available prices. However, this premise is undermined by the fact that the corporations operating the agents are incentivized to steer users to shopping outcomes that benefit the company, not the user.

As Lindsay Owens and Nia Law [argued](#) in a piece in the *Capitol Forum*, the established doctrine of agency holds that a representative cannot serve two principals with opposing interests; AI agents deployed by Big Tech companies appear to operate in direct violation of these principles. A report by the Open Markets Institute in June 2026 warned that, while 'most people assume that AI agents also work for them', this arrangement is inherently undermined by AI agents' control by companies who [profit](#) from "surveillance, addiction, manipulating users with hyper-targeted content, or extracting tolls".

Alongside these concerns, the proliferation of AI agents is also likely to lead to separate questions relevant to competition policy. Just as Meta has received antitrust scrutiny for blocking rivals' [AI chatbots](#) from WhatsApp, Amazon's decision to block Perplexity's Comet shopping agent from access to its marketplace (which Amazon's own AI agents can [access](#)) suggests that dominant platforms will use agentic AI as a new tool of [self-preferencing](#).

SURVEILLANCE PRICING CONCERNS

Algorithmic collusion, identified as a [competitive threat](#) long before the public debut of DALL-E or GPT-3.5, has only become a more acute threat in the generative AI era. Antitrust attorneys including [Ashley Walters](#), assistant attorney general for the District of Columbia, have raised concerns that agentic AI services will further enable algorithmic collusion, and in the process undermine independent price assessment at consumers' expense.

Survey data suggests that consumers see price comparison as agentic AI's most valuable use case. A [Radial survey](#), for instance, found that 47% of respondents would use AI agents to find the best price for a particular product, outpacing any other use case.

At a time when [rampant surveillance](#) pricing has undermined consumers' trust in [fairness in online commerce](#), it is understandable that an automatic tool for price-comparisons would provide objectivity in pricing. However, advocates have warned that, on the contrary, agentic AI is more likely to accelerate surveillance pricing by giving platforms an unprecedented window into users' preferences and purchasing intent.

Experts including economist Ted Tatos have raised [concerns](#) that services like Google's Universal Commerce Protocol (UCP) will serve to entrench "exploitative practices that mirror those already challenged under antitrust statutes". In other words, while AI agents are positioned as tools to find objective and fair prices in a time when consumers worry their choice to use or not use Incognito mode will [affect](#) their plane ticket price, they are likely to have the opposite impact by providing companies with more intimate insights into consumers' preferences.

III. Privacy and Safety Risks

OPERATIONAL RISKS

Agentic AI systems have already demonstrated a tendency to act outside the boundaries users set for them, and the legal framework governing them is equally unsettled. Cybersecurity experts have noted that agentic AI services are inherently [vulnerable](#) to attack by harmful actors due to the so-called “lethal trifecta” of properties: first, agents’ access to users’ private data; second, agents’ ability to externally communicate; and third, agents’ exposure to uncontrolled content.

AI agents acting outside of specified boundaries have already produced serious real-world consequences, including for tech industry professionals best equipped to manage it. In April 2026, software company PocketOS [experienced a 30-hour outage](#) after an AI coding agent deleted its entire database in just nine seconds. These risks have also been felt by employees at Big Tech companies who operate major agentic AI services. Earlier in 2026, a director at Meta's superintelligence lab publicly [documented](#) an instance in which an agent she had instructed to “confirm before acting” began deleting emails en masse. After Amazon instituted an employee mandate for utilizing its [Kiro AI tool](#) for coding projects, the company would experience a 13-hour Amazon Web Services (AWS) outage, along with 6.3 million [lost orders](#).

THREATS TO USERS’ FINANCIAL DATA

Given that agentic AI systems have caused serious disruption for tech professionals, including those at companies that build them, the consequences for ordinary users with less technical recourse are likely to be more severe. As AI-driven agents play an increasingly major role in large transactions, the possibility of them making costly mistakes at a users' expense (such as the “[Wrong Paris](#)” scenario posited by data rights expert Dazza Greenwood) makes the need for regulatory certainty more important.

Generative AI development has already enabled financial scams, such as [deepfake voice call scams](#) that have a particularly devastating impact on vulnerable groups such as [elderly people](#). The proliferation of AI agents without meaningful regulatory safeguards puts users’ intimate personal data, including financial data such as credit card information, [at risk](#). Research has found that AI agents are [capable](#) of conducting multi-turn scam conversations with prospective victims, putting more people at threat of extortion.

FRAMEWORK FOR FEDERAL GUARDRAILS

Sen. Mark Warner (D-VA) [circulated a discussion draft](#) on June 30, 2026 of the **AI Agent Act** (Artificial Intelligence Access, Gatekeeper Exchange and Nondiscriminatory Transfer Act), which would represent the first proposed federal framework for agentic AI. The bill would empower the FTC to establish baseline privacy and security standards for vendors, an overdue measure given the current rapid rate of agentic adoption.

Notably, the bill would [mandate](#) that agentic AI providers link individual AI agents to the identity of their human operator, along with requiring that users have accessible methods of authorizing or deauthorizing agents. The National Institute of Standards and Technology (NIST) will be empowered to develop interoperability standards for the technology.

The bill's primary value is in its proposal for strong standards governing agentic AI vendors' handling of sensitive user data, namely email access, payment details, or accounts on websites they shop on. The proposed measure's impact is limited by the fact that, if it were to pass, the FTC would not be able to prevent platforms that do not meet standards from certification from operating agents (it can, however, deregister them from the commission's list).

It stands to reason that it is unlikely that legislation with these provisions will make it into law under the Trump administration, particularly given the administration's active efforts to prevent states from implementing AI [regulations](#). Nevertheless, Warner's legislation is a strong step towards providing an alternative to the present hands-off approach to agentic AI.

IV. Conclusion

The early record of agentic AI deployment indicates that the technology's benefits are accruing primarily to the platforms operating these services and not the consumers using them. And absent real safeguards, this gap is likely to only widen as deployment accelerates further. This can be seen in the apparent financial boon Amazon experienced as a result of its now-defunct Rufus agent during the 2025 holiday [season](#), as well as Walmart executives' own admission that agentic AI is leading to an increase of 35% in consumer [spending](#). That these services have been held up as internal successes by the corporations deploying them is consistent with the concern that agentic AI is designed to serve platform interests rather than those of the consumers who use them.

In the realm of privacy, it is understandable that many consumers are worried about how agentic AI may mishandle their data, particularly given that AI chatbots have built a [poor record](#) on the matter. A [survey by Quad and The Harris Poll](#) found that 73 percent of Gen Z and millennials are worried about how AI tools will use their personal shopping data, and nearly three-quarters said algorithm-driven pricing makes it harder to confirm they are receiving the best price.

The regulatory window to establish guardrails before agentic AI becomes further embedded in everyday commerce is narrowing. As our analysis has [previously noted](#), the case for self-preferencing and tech antitrust legislation has grown stronger in the AI era, a cause made more urgent by dominant platforms' deployment of agentic AI services. Consumer skepticism is likely to mount as the gap between the pace of deployment and the development of governance frameworks becomes more apparent.

AUTHORSHIP AND PROJECT INFORMATION

Aidan Smith is the founder of Labyrinth Insights LLC, a boutique strategic research platform focused on policy and public accountability. Labyrinth Insights regularly publishes [independent research](#) concerning a broad range of policy areas, including competition policy, technology regulation, and consumer protection issues.