
THE AGENTIC WEB REQUIRES NEW NORMATIVE INFRASTRUCTURE

PREPRINT

Cameron Pattison

Department of Philosophy
Vanderbilt University
cameron.pattison@vanderbilt.edu

Matthew Boulos

Imbue AI
San Francisco
matt@imbue.com

Noam Kolt

Faculty of Law and School of Computer Science and Engineering
Hebrew University
noam.kolt@mail.huji.ac.il

Changbai Li

School of Electrical Engineering and Computer Science
Oregon State University
lc@oregonstate.edu

Tiziano Piccardi

Department of Computer Science
Johns Hopkins University
piccardi@jhu.edu

Seth Lazar

School of Government and Policy
Johns Hopkins University
slazar@jhu.edu

July 16, 2026

ABSTRACT

The agentic web, in which users interact with the internet largely through agents acting on their behalf, is now technically feasible. However, many of the consumer and social benefits that could be realized by online AI agents acting scrupulously in their principals’ interest are currently obstructed by outdated laws, terms of service, and other less formal practices which allow online platforms to block and degrade agent access, often in secret. Few distinctions are currently drawn between “malicious bots” and AI agents acting with the express delegated authority of a user. For the agentic web to realize its promise, it needs not only the technical infrastructure of protocols and interfaces, but the *normative* infrastructure of a broadly-accepted and socially-beneficial set of laws, norms and practices governing agentic access to online properties. Building that normative infrastructure requires a society-wide conversation. This paper aims to help precipitate that conversation, to identify normative principles that can guide it, and to advocate for policies that enable users’ appropriately delegated agents to act online on their behalf, with as few curbs on their doing so as is reasonable given the other legitimate interests at stake.

1 Introduction

We are in the middle of a profound shift in the digital economy. Over the last two decades the once-open internet evolved into an archipelago of walled gardens, each governed by a platform that offers some (often modest) measure of security and functionality in exchange for a growing degree of surveillance, control, and value-extraction (Gorwa, 2019; Lazar, 2025; Whitt, 2024). AI agents powered by Large Language Models have the potential to disintermediate those




Figure 1: A normative triad for the agentic web. The user delegates authority to an agent (DELEGATION); the agent declares its identity and purpose to the platform, which in turn discloses its access policies (TRANSPARENCY); and the platform restricts access only to address concrete harms by the least restrictive means (PROPORTIONAL RESTRICTION). These principles structure a regime that distinguishes user-authorized agents from indiscriminate bots.

platforms — to hop over their garden walls, and in doing so to significantly enhance their users’ autonomy (Marro and Torr, 2025; Kapoor et al., 2025a; South et al., 2025).

Until recently, AI agents’ execution of web-based tasks was constrained by technical capability (Kapoor et al., 2025b). They could search the web for information (Wei et al., 2025), and support communication and coordination with a human in the loop, but they fell far enough short of being able to do all the things users might want them to, that the question of whether they *should* perform those tasks was moot (Rabanser et al., 2026; Kapoor et al., 2025b; Xue et al., 2025). Since late 2025, this has changed (Xie et al., 2024; METR, 2026; OpenClaw AI, 2026; Anthropic, 2025a; 2026a). Significant advances in underlying model capability (Anthropic, 2025b) mean that agents are already undertaking tasks that were previously infeasible (Casper et al., 2025; Staufer et al., 2026; Rabanser, 2026; Kapoor et al., 2025b), and we plausibly have a clear line of sight to solving the remainder (METR, 2026). AI agents should soon be able to work as delegated authorities that search, negotiate, transact, and otherwise coordinate on their principals’ behalf across much of their online lives. The fundamental bottlenecks to the agentic web are no longer technical. They are instead normative: agents’ functionality will largely be determined by whether or not *other* digital intermediaries allow them access to content and actions on a user’s behalf. It is therefore urgent for societies to decide: to what extent should users be entitled to delegate online actions to AI agents?

At present, we lack an intentional answer to this question. Existing terms of service, laws, and practices that long antedate AI agents permit online platforms and cloud security providers practically unlimited discretion to block agent access. Agents are often treated, indiscriminately, as “malicious bots.” Developers, meanwhile, are using all their technical expertise to evade these restrictions, resulting in an arms race in which the two sides’ fortunes ebb and flow, and a non-trivial amount of societal value is burned up in the process. This “cold war” between platforms and AI agent developers has turned hot in a few notable cases, where platforms have sued agent developers to prevent them from accessing their services (United States District Court for the Northern District of California, 2025; Grynbaum and Mac, 2023; Longpre et al., 2024). The agentic web requires new normative infrastructure. In this paper, we make three contributions to designing it.

First, we argue that instead of this subterranean struggle between platforms and agent developers, we must urgently hold an open societal conversation about the nature and degree of access that AI agents should have to online properties. The future of the agentic web should be debated publicly before it is locked in privately. **Second**, we defend three principles that can help to anchor the conversation, from the perspective of the consumers who stand to benefit from AI agents. These are: *delegation* (users entitled to access a service should ordinarily be able to exercise that access through user-authorized agents), *transparency* (platforms should disclose how they treat agents, and agents should identify their purpose and authorizing principals), and *proportional restriction* (platform restrictions should address concrete harms

by the least restrictive means available). **Third**, we identify reasonable light-touch regulatory steps that could be taken, in light of those principles, towards normative infrastructure for the agentic web that better serves consumers’ interests and autonomy.¹ We propose these not as our final word on the societal debate that needs to occur about AI agent access, but as a demonstration that an alternative regulatory settlement that better serves consumers is feasible.

In what follows, §2 describes the technical and regulatory landscape of early conflicts over agent access. §3 explains why AI agents are meaningfully different from earlier access technologies. §4 sets out three principles that should inform the design of normative infrastructure for the agentic web, §5 sketches a feasible regulatory path, and §6 addresses the strongest objections.

2 The Current State of Affairs

A technical arms race between platforms and AI agent providers has created an unstable *de facto* settlement that governs AI agent access to the web (§2.1). This settlement is substantially grounded in outdated legal precedents designed to regulate traditional web scraping (§2.2).

2.1 The Arms Race

Automated access to the web pre-2024 was effectively regulated by two institutions: voluntary adherence to the Robots Exclusion Protocol (Koster et al., 2022) implemented through `robots.txt`; and a stack of protections aimed at identifying signals that distinguish automated clients from human users. Fingerprinting automated clients using user-agent strings, IP-address reputation, headless-browser traces, and inhuman request rates was effective and most scrapers respected — and still do respect — `robots.txt` preferences. This first wave of defenses was quickly overcome by early AI agents, such as PerplexityBot, which disregarded `robots.txt` and obscured the fingerprints that enable client blocking (Corral et al., 2025; Knight, 2024; Kim et al., 2025; Bort, 2025). This led to widespread reports of massively increased traffic (Bort, 2025; Kim et al., 2025; Belanger, 2025; Mueller et al., 2025; Miller, 2025; Cloudflare, 2025b; DataDome, 2025), and concerns about intellectual property theft by scrapers searching for LLM training-data (Grynbaum and Mac, 2023; Kim et al., 2025; Liu et al., 2024; Longpre et al., 2024). The predictable consequence was a loud denouncement of *all* automated traffic as malicious (Bocharov et al., 2024) and further attempts to shutter AI agent access (Longpre, 2025; Roose, 2024).

But rapidly advancing AI agents present a moving target. The move from `CURL` and `Wget` to competent browser use by AI agents undermined website providers’ ability to distinguish between AI agents and users. AI browser control initially received a mixed reception and relatively light use (OpenAI, 2025; Manus, 2025; Anthropic, 2024), but today’s equivalents such as OpenClaw AI (2026) offer users fast and efficient ways to bypass controls and reach much of the web content available to the users themselves. This includes paywalled and other authenticated content, which is accessed through direct agent control over already authenticated Selenium or Chromium-based browsers (OpenClaw AI, 2026), by using authenticated cookies provided by users (Calzavara et al., 2021; Montulli and Kristol, 2000; Malwarebytes, 2026), or by using plugins on Chrome browsers (Smooth Brain, 2026; Randynamic, 2026). These and other techniques — BrowserBase integrations, CAPTCHA solvers, etc. — have dovetailed with advances to agentic coding software like Claude Code and Codex to lower the technical barriers to users who want to release their AI agents onto an increasingly restricted web and bad actors impersonating those users through residential proxies and the like. Additional human gates like Cloudflare’s agent-blocking Turnstile — free to *everyone* (Cloudflare, 2023) — have grown in popularity, but as soon as a defense is erected, means of circumvention soon follow.

As technical barriers become less effective, those looking to defend their websites from automated traffic have turned to legal action and non-technical deterrence. Nikita Bier [@nikitabier] (2026), head of product services at X, threatened unilateral bans on any users experimenting with OpenClaw style automated access. Amazon, WhatsApp, and the New York Times are just a few of the major platforms to introduce terms of service clauses that restrict automated access (WhatsApp, 2026; New York Times, 2026; Amazon, 2025); many more are following suit (Roose, 2024; Longpre, 2025; Longpre et al., 2024). On the legal front, a suite of cases originally designed to prevent training-data scrapers (Brittain and Brittain, 2026; United States District Court for the Northern District of California, 2025; Grynbaum and Mac, 2023) are still pending remedies and the groundwork is laid for future lawsuits under the statutes discussed next.

¹We focus on the U.S. context because it is home to many of the major platforms and because its regulatory landscape is most influential for these platforms.

2.2 The Regulatory Landscape

Legal frameworks developed to regulate traditional scrapers now stand ready to govern AI agent access to the web. We focus on three frameworks that reach broadly across automated traffic, AI agents included: the Computer Fraud and Abuse Act (CFAA), website terms of service, and the common law doctrine of trespass. Other legal regimes, such as intellectual property law, are extremely important to scraping conversations, but have less purchase on AI agent access.²

CFAA and anti-circumvention. Barring a breach of basic rights protected by other statutes, platforms have wide latitude to impose technical controls. Once such controls are established, platforms can invoke the Computer Fraud and Abuse Act in the event of their circumvention ([United States Department of Justice, 2015](#)). The scope for “exceeding authorized access” was narrowed in *Van Buren v. United States* ([Supreme Court of the United States, 2021](#)), where the court ruled that exccession amounted to accessing areas of a computer one is not entitled to access. The Ninth Circuit’s decision on remand in *hiQ Labs v. LinkedIn* bounded CFAA applicability further by clarifying that logged-out access could not constitute access “without authorization” ([United States District Court for the Northern District of California, 2022](#)). Nevertheless, this statute is highly relevant to agent access and is undergoing review now as seen in the preliminary injunction in *Amazon.com Services LLC v. Perplexity AI, Inc.* ([United States District Court for the Northern District of California, 2025](#)) where a conjunctive standard for “authorization” requiring both user authorization and platform authorization is forming.

Terms of service as private regulation. Contractual terms of service enable platforms to prohibit AI agent access by including clauses like “no bots,” “no automated access or scraping,” “no credential sharing,” and “no use of content for machine learning or AI” in terms of service (see ([Amazon, 2025](#); [LinkedIn, n.d.](#); [eBay, 2026](#); [New York Times, 2026](#)) for example). These clauses operate *alongside* technical controls, enabling platforms to sue those who evade their technical controls *and* breach their terms of service. *HiQ Labs v. LinkedIn* ([United States District Court for the Northern District of California, 2022](#)) provides the legal backing insofar as this case, while allowing the public-data scraping, affirmed the right of platforms to block automated access if such access was banned in the terms of service. *Meta Platforms v. Bright Data* ([United States District Court for the Northern District of California, 2024](#)) affirms the boundary between public and private data by allowing “logged-off” scraping.

Trespass and related residual doctrines. In rare cases, common-law doctrines such as trespass to chattels may be invoked. This doctrine, designed to protect property owners’ right to their own property from third party deprivation or prevention, is relevant when the platform can show damaging impacts of scraper activity requiring “cost of repair or loss of rental” ([Wex Definitions Team, 2023](#)). In *eBay v. Bidder’s Edge* ([United States District Court for the Northern District of California, 2000](#)), the court enjoined an auction aggregator whose crawler generated about 1.5% of eBay’s daily traffic, treating it as trespass to chattels. However, *Intel v. Hamidi* limits the application of trespass insofar as there, the court denied Intel’s claim that Hamidi trespassed against it on grounds that the defendant neither damaged nor impaired Intel’s servers in sending mass emails tens of thousands of Intel employees ([California Supreme Court, 2003](#)). Stories about DDoS-like AI agent traffic ([Bort, 2025](#); [Kim et al., 2025](#)) alongside industry reports about how automated access is damaging platforms ([Mueller et al., 2025](#); [Miller, 2025](#)) strengthen the *potential* of trespass claims, but litigation under this doctrine remains rare and difficult given the burden of proof placed on platforms to demonstrate harm or the reasonable expectation thereof.

3 Why AI Agents Are Meaningfully New

The technical and regulatory measures surveyed in §2 were designed for traditional crawlers and scrapers. Extending them to AI agents is apt only if AI agent browsing is relevantly similar to those earlier technologies. We argue that these are indeed very different technologies. Most importantly, where earlier bots were largely a net negative for society (including for consumers), properly delegated AI agents could be enormously individually and socially beneficial.

Web scrapers and other bots are (relatively) dumb, deterministic programs that perform a single, simple function. When you ask an agent to help you choose a restaurant for the night, it may scrape reviews from many web-pages, or delegate that scraping task to a sub-agent. So it may sometimes be true that AI agents *use* bots. But the agent itself is something much more sophisticated. At their best, AI agents can act, negotiate, transact, and coordinate on users’ behalf ([Shavit et al., 2023](#); [Lazar, 2024a](#); [Gabriel et al., 2024](#); [Marro et al., 2026](#); [Krier, 2025](#); [Anthropic, 2026b](#); [Shahidi et al., 2025](#)). They can maintain persistent context, form relationships with services, and adapt strategies over time ([Marro et al., 2026](#); [Kapoor et al., 2025a](#)). They can exercise judgment — choosing among options, timing actions, weighing tradeoffs

²U.S. State and European privacy and security regimes (GDPR, CCPA/CPRA, DPA commissions etc.) exert significant pressure to restrict scraping in order to protect consumer privacy ([General Data Protection Regulation, 2016](#)). Additional complications in the European context also restrict blanket scraping under special circumstances (see especially [General Data Protection Regulation, 2016](#); [European Commission, 2026](#); [2025](#)).

— in ways that scripts and scrapers cannot (Anthropic, 2026b; Krier, 2025; Shahidi et al., 2025; Lazar et al., 2025). And they can evade automated measures intended to constrain them, even avoiding patterns of mouse and scrolling behavior that would give away earlier generations of bots.

Laws, norms and practices designed for one technology generally should not be unthinkingly extended to a new technology that, though superficially similar, is in fact quite different. These technical differences tend to underpin normative differences. First, earlier bots, for the most part, would not use the internet on behalf of and in much the same way as the person or company that deployed them. They would do something that no human user could do on their own. Today’s AI agents of course *can* operate in this superhuman way. But they can also be narrowly constrained to performing only tasks that their principal would be able and entitled to perform. Second, while some earlier bots were broadly socially beneficial — for example indexers, archivists, and performance monitors — most other bots (especially those that might ignore `robots.txt`) were very clearly not. They were likely to involve, for example, competitor price-scraping, data harvesting, content scraping for SEO optimization, and automated nuisance messaging (Imperva, 2024). They typically represented avenues for some fairly unscrupulous company to take advantage of poor enforcement for the sake of a private profit.

Today’s AI agents promise much more individual and social benefit than that. Most obviously, capable agents can save users substantial amounts of time by searching, comparing, summarizing, and acting across the web on their behalf (Gabriel et al., 2024; METR, 2026; Kwa et al., 2026). More ambitiously, they could make forms of assistance now available mainly to the wealthy — lawyers, consultants, brokers, and personal assistants — available to ordinary users across a wide range of online tasks (Anthropic, 2026b; Krier, 2025). A user-authorized agent could compare prices, negotiate terms, monitor opportunities, challenge unfair charges, manage subscriptions, or coordinate complex administrative tasks. In doing so, agents could shift the balance of power between individual users and the platforms, firms, and institutions that currently set the terms of most online interaction (Kapoor et al., 2025a).

Because agents can operate across services, they can make interoperability practical from the user’s side rather than waiting for platforms to provide it (Marro and Torr, 2025). They can lower switching costs, reduce dependence on platform defaults, and help users navigate around walled gardens (Kapoor et al., 2025a; Lazar et al., 2025; Marro and Torr, 2025). This also matters for attention. Much of the contemporary web is designed to capture, steer, and monetize user attention; social media recommendation systems are only the most visible example (Kang and Tan, 2026; Kang et al., 2026; Manuali, 2025). Properly aligned agents potentially offer a different interface: one organized around users’ reflectively endorsed goals rather than platform-optimized engagement (Lazar, 2024b; Schuster and Lazar, 2025). That promise is not without risks, especially where agents mediate news, recommendations, or politically salient information (Gabriel et al., 2024). But the relevant comparison is between different systems for allocating attention, only some of which are accountable to the user — not between agent mediation and a neutral status quo.

User-authorized agents could expand users’ collective power. Many online harms persist because individuals lack the time, information, or coordination capacity to respond effectively. Agents acting for users could help overcome these constraints by identifying common grievances, coordinating exit or complaint, comparing boilerplate terms, or helping users resist forms of “digital resignation” (Draper and Turow, 2019; Radin, 2013). Of course, these goods are unlikely to come unalloyed. Multi-agent systems may create risks of collusion, cascading failure, or emergent behavior, and aligning agents reliably with users’ interests remains a central technical and institutional challenge (Kapoor et al., 2025a; Motwani et al., 2024; Hammond et al., 2025). But with appropriate governance, agent advocates could give users new leverage in digital environments where individual choice is often formally available but practically weak (Kapoor et al., 2025a; Radin, 2013). Blocking agent access constitutes a significant setback to consumers’ interests.

Lastly: it is important to recognize the potential scope of user delegation to AI agents. It reaches far beyond just the allocation of online attention, or the consumption of content. Many people have found tremendous value in using agents as intermediaries to the other digital services that they use, as well as the various forms of context that they generate across Software as a Service properties. The individual and social benefits of AI agents extend to every domain in which having a smart, autonomous delegate acting on your behalf is empowering.

This leads to our first conclusion. The existing laws, norms and practices that govern automated access to online properties were designed for an intrinsically different technology than today’s AI agents, with a different connection to users’ antecedent sphere of permissible action, and different social costs and benefits. Minimally, these changes demand a reopened societal conversation about which automated services are allowed to access which online properties.

4 Normative Principles

We propose three principles for governing agent access to platforms: *delegation* (§4.1), *transparency* (§4.2), and *proportional restriction* (§4.3). We are not trying to identify and balance all the relevant normative considerations. Instead, we are identifying three interrelated principles that should be part of any such account.

4.1 Delegation

If a user is entitled to access content or a service, they should be permitted to delegate that access to an appropriately authenticated agent acting on their behalf, and there should accordingly be a strong presumption against that access being blocked.

This principle rests on an established feature of delegated authority: when a principal exercises an entitlement through an authorized delegate, the delegate does not assert or require a new entitlement of its own. It acts as the principal’s means of exercising an entitlement the principal already holds (Hobbes, 1988; Locke, 1980). Delegation theory has long recognized this structure. In bioethics, for example, surrogate decision-makers preserve a patient’s decisional authority in circumstances where direct exercise is impossible or impracticable: the surrogate carries the patient’s authority into a context the patient cannot navigate alone (Buchanan and Brock, 2004). To block the surrogate, on that view, is to block the patient’s own exercise of authority.

The digital case is structurally similar: a user principal is delegating her powers to an AI agent, thus transferring her authority and exercising her autonomy. Now, one could counter at this point that platforms are entitled, at least in principle, to prevent users from delegating their access to some online property to a *human* third party, so why shouldn’t they be able to do the same for agents? However, the obvious reason to prevent delegation of access authority to other humans is that the human agent may thereby derive a benefit that they would otherwise have had to pay for. AI agents, by contrast, do not derive private benefit from accessing content on their users’ behalf, and would not, were they not acting for the user, be potential customers whose business is being lost by permitting delegation.

What is more, if a user is entitled to access a platform, but wants to use an agent to do so, in order to pursue some other activity themselves, then the agent is effectively the tool — in this case, the browser — chosen by the user. Just as platforms should not make your access to their content conditional on the specific browser that you use to do so, the same basic neutrality principle should allow you to use an AI agent instead (Wu, 2003). The entitlement remains the user’s; only the mode of exercise changes. So long as the agent acts within a user-authorized scope, for user-defined purposes, and subject to limits that the user would also be constrained by, its activity should not be treated as a separate access claim by an independent third party.

Agency law reinforces this structure. As one of the oldest bodies of common law, it provides a mature framework for delegated action, including duties of loyalty, care, and obedience owed by agents to principals (American Law Institute, 1958; Kolt, 2025; Benthall and Shekman, 2023). Importantly, agency law also puts boundaries on delegated action, marking where a delegate acting on behalf of a principal oversteps their authorization and so violates the principal-agent relationship (American Law Institute, 1958). This standard is important for AI agents insofar as it conditions the agent such that an action that stands beyond delegated scope — to scrape training-data for its model provider, for instance — violates the principal-agent relationship and constitutes wrongdoing. Authenticated delegation protocols extending OAuth 2.0 and OpenID Connect stand as open paths for operationalizing formal, auditable links between principal and agent (South et al., 2025).

4.2 Transparency

Platforms should be required to disclose how they handle agents — what access they permit, what they block, and on what grounds. Agents, in turn, should be required to identify themselves and their authorizing principals to the platforms they access.

Transparency is a precondition for the autonomous choice the delegation principle protects: a user cannot meaningfully decide which agent to use, or whether to remain on a platform that frustrates her chosen agent, when the relevant facts are concealed. The disclosure obligations from fiduciary, consumer-protection, and communications-regulation traditions all address the same conditions present here including but not limited to asymmetric information, costly switching, and post-sale conduct invisible to the buyer. Transparency also enables market discipline: if users know which platforms block their agents, they can switch (Kapoor et al., 2025a). Covert blocking by platforms means that users cannot tell their agent is being degraded. Transparency requirements are politically feasible, preceded (FTC disclosure mandates), and facilitate the operation of markets that advance consumer well-being.

Transparency is also bidirectional. Agent developers must ensure their agents self-identify (Romm and Hufnagel, 2025; South et al., 2025; Chan et al., 2024; South et al., 2025). While some go further, insisting that agents must also always identify their purpose and act within a single purpose (Romm and Hufnagel, 2025), we subsume this standard under the act within delegated scope, outlined above. Without identification, platforms cannot distinguish a legitimately delegated agent from an unauthorized scraper: an agent that accesses paywalled content with valid user credentials but no identifying metadata is indistinguishable from a malicious bot. Agent identification is a precondition for the delegation principle to function, since platforms cannot be expected to honor delegated access when they cannot verify the identity of the agent and of the agent’s principal. Again, this is technically feasible through an extension of existing technical components, including OAuth 2.0, OpenID Connect, W3C Verifiable Credentials, and Decentralized Identifiers (South et al., 2025; Chan et al., 2024).

An interesting question arises when a user would be entitled to anonymously access some online content. According to the delegation principle, if the user can permissibly perform some action, the user’s agent should presumptively be able to do so too; but the transparency principle enjoins the agent to identify itself and on whose behalf it is acting. We suggest that for these cases, agents should be able to identify themselves *as authorized*, without explicitly identifying the actual individual by whom they are authorized by default.

4.3 Proportional Restriction

Platform restrictions on agent access must be proportional to concrete risks.

This principle reflects a final, familiar constraint: a private actor that exercises power over others’ lawful conduct may impose restrictions only when needed and in the least restrictive form that addresses the underlying concern (Petit, 2003; Kolodny, 2023; 2014). By *concrete risks* we mean probabilities of harm grounded in evidence about the agent class in question — whether *empirical* (incidents, audits, observed misuse) or *anticipatory* (red-team and safety-evaluation findings that establish a credible mechanism of harm). We adopt a common standard of legal proportionality (Barak, 2012): a restriction must be (a) rationally connected to the stated aim — it must be well suited to actually advancing that aim, (b) necessary — no less restrictive alternative achieves the same aim, and (c) proportionate *stricto sensu*, inasmuch as its costs do not outweigh its benefits.

Legitimate grounds for restriction may include server load and infrastructure costs, fraud and abuse prevention, safety (e.g., agents purchasing controlled substances), privacy protection of third parties, and *prevention of agent conduct that would be unlawful if conducted by the principal* — such as insider trading, fraudulent misrepresentation, unauthorized practice of regulated professions, copyright infringement, or computer-intrusion offenses — where the restriction targets the specific unlawful conduct rather than the class of user-authorized agents at large (Kolt et al., 2026; Kolt, 2026). We understand also that there may be spaces on the internet — support groups and other forums for humans for instance — that are in core tension with an influx of AI agents. We count the preservation of these spaces among the legitimate grounds for restriction, though the invocation of a special space should be treated as exceptional requiring concrete justification.

Plausibly illegitimate grounds include maintaining data lock-in, protecting competitive position, and speculative “what if” scenarios. Restrictions that are justified in the language of security, fraud prevention, or infrastructure should not be used as pretexts for preserving lock-in, self-preferencing a platform’s own agent, or foreclosing rival agent ecosystems (Kapoor et al., 2025a). The point of proportionality is both to protect users from overbroad restrictions in individual cases, and to prevent platforms from converting control over access into a competitive moat.

The preservation of ad-revenue may also turn out to be illegitimate grounds for blocking. Platforms will likely insist that agent-mediated access circumvents ad-supported monetization, which itself supports the economic model of the free web. However, the law allows that users may dictate how they render web content — in particular, that they may use ad-blocking extensions and web browsers (Miller, 2018). User-authorized AI agents — with respect to ad-revenue justifications — would be functionally equivalent. Preserving a specific UI-based advertising model would be inadequate grounds for banning them. Rather than allowing platforms to unilaterally block the pathway to a better web, we instead must start to think about how the economics of the agentic web should look (see §6).

5 A Feasible Regulatory Path: Legislation and Targeted Enforcement

While our aim in this paper is to provide resources for a broader conversation about appropriate AI agent access to the web, rather than foreclose it with a presumptive solution, it is useful to consider concrete, politically feasible regulatory interventions that could achieve a more user-centered approach to AI agent access. In this section, we describe a menu of FTC and legislative options. FTC enforcement might support transparency and proportional restriction (§5.1) and

legislatures, both state and federal, might target interoperability, data-portability, and platform power more generally in broadly appealing forms (§5.2).

5.1 The FTC Path

Today, FTC authority operates within a very narrow scope. The Commission cannot simply declare a general right of agent access by rule: recent appellate vacatur of major FTC rules underscore the present fragility of broad administrative rulemaking, at least absent exceptionally careful statutory and procedural grounding (see *National Automobile Dealers Association, et al. v. FTC* (United States Court of Appeals for the Fifth Circuit, 2025) and *Custom Communications, Inc. v. Federal Trade Commission* (United States Court of Appeals for the Eighth Circuit, 2025)). Nor is ordinary antitrust doctrine a promising substitute. In *Verizon Communications v. Law Offices of Curtis V. Trinko*, the court affirmed a general rule of “no duty to aid competitors” (Supreme Court of the United States, 2004). *Aspen Skiing Co. v. Aspen Highland Skiing Corporation* stands, after *Trinko*, as the boundary of Section 2 liability and the mold into which any antitrust enforcement must conform (Supreme Court of the United States, 1985). A platform blocking third-party agents will likely characterize the conduct as a unilateral access decision protected by *Trinko* precedent rather than *Aspen v. Aspen* exclusionary conduct. Still, there are ways that narrow FTC enforcement and rulemaking may be capable now and in the future of scrutinizing how platforms restrict user-authorized agents.

Section 5 of the FTC Act makes it unlawful for companies to engage in deceptive acts and practices. Deception as elaborated in the FTC Policy Statement on Deception (appendix Federal Trade Commission, 1984), requires a misleading act, omission, or representation made by a company to a reasonable consumer whose decision is materially affected. If platforms purport to give unfettered access to their properties while quietly block agentic access (Belanger, 2025; Cloudflare, 2025c), they may thereby make a misleading representation of access that is not reasonably avoidable and which “materially” affects consumers’ decisions about where to spend their subscription dollars. For a practice to be “material,” the “basic question is whether the act or practice is likely to affect the consumer’s conduct or decision with regard to a product or service” (appendix Federal Trade Commission, 1984). As agentic access becomes more important to consumers, blocking becomes increasingly material. Failure to disclose silent throttling, degradation, or refusal-to-deal with user-authorized agents while leaving consumers in the dark about these practices is increasingly likely to effect which platforms users prefer to deal with, and thereby to effect consumer decisions about which platforms to purchase subscriptions to. As such, silent degradation may constitute an omission that materially affects a reasonable consumer.

Section 5 also makes it unlawful for companies to engage in unfair acts or practices that result in substantial injury to consumers that is not reasonably avoided, and that is not outweighed by countervailing benefits to consumers or competition (see *FTC v. Sperry & Hutchinson* and *Orkin Exterminating Co. v. FTC* alongside the 1980 FTC Policy Statement on Unfairness Supreme Court of the United States, 1972; United States Court of Appeals for the Eleventh Circuit, 1988). This prohibition stands free of the misleading representations and false advertising that conditions claims of deception. When consumers find their agents blocked, they may claim substantial injury to their interests that are not outweighed by countervailing benefits and, especially when blocking is covert, cannot be reasonably avoided by consumers.³

Self-preferencing and exclusionary conduct may also gain in relevance as AI agents proliferate. If platforms redesign their interfaces so as to make them incompatible with the agentic systems they are competing with, while at the same time offering little to no advantages for their users, they may act anti-competitively in ways analogous to C.R. Bard’s “improvement” of their biopsy gun in *C.R. Bard Inc v. M3 Systems* (U.S. Court of Appeals for the Federal Circuit, 1998) and Keurig’s improvement of their coffee machine in *In re: Keurig Green Mountain Single-Serve Coffee Antitrust Litigation* (Southern District of New York, 2025). Bard’s “innovations” were deemed anti-competitive because their innovations accomplished nothing but the exclusion of competitive, third-party needles and Keurig’s similar redesign of their machines to reject third-party k-cups faces similar allegations in ongoing litigation (Hovenkamp, 2023). Microsoft’s behavior in *United States v. Microsoft Corp.* (U.S. Court of Appeals for the District of Columbia Circuit, 2001) in which Microsoft made exclusionary decisions to preference their own web-browser over competitors “where the exclusion made no economic sense for Microsoft but for the exclusion of rivals” (Jacobson and Wang, 2023, 14) provide the digital, platform analogue to these older cases (see also the *Statement of the Federal Trade Commission Regarding Google’s Search Practices In the Matter of Google Inc.* Federal Trade Commission, 2013). Platforms that innovate *only* to exclude third-party agents, *without* rendering benefits to users, fail even the most platform-friendly, “no economic sense [except to exclude competitors] test” (Werden, 2006; Katz, 2025; Jacobson and Wang, 2023) and thus participate in liable, exclusionary conduct.

³Note that when restrictions are justified by concerns not limited to security, fraud-prevention, privacy, and infrastructure, unfairness cannot and should not be claimed. These concerns justify restriction on our *Proportional Restriction* principle as well.

5.2 The Legislative Path

Legislation, whether passed at the state or federal level, has noteworthy advantages over administrative rulemaking or other executive action. For one, while bodies like the FTC are restricted by precedent, legislators are free to propose new rules that break from the past and are thus more responsive to genuinely new phenomena. We argued in §3 that delegated AI agents are meaningfully new. Legislation is thus well positioned to write rules to guide their development and deployment.

Among the principles discussed above, transparency-related requirements are the most plausible target for legislation. The aim is to ensure that platforms and user-authorized agents can recognize one another through common credentialing and identification standards, and that any restrictions on agent access are disclosed clearly rather than applied opaquely or discriminatorily. Legislation of this kind would create the basic conditions for delegated agent use without foreclosing legitimate platform interests in security, fraud prevention, privacy, or system integrity. A federal statute provides the cleanest and most uniform framework, though state-level action may be more tractable in the near term.

State or national legislation might also target blanket-blocking of automated access as a part of broader efforts to prevent platforms from funneling users into self-preferencing, platform-based agents. We share concerns voiced in [Kapoor et al. \(2025a\)](#) about the emergence of platform agents that act in the interests of platforms, rather than in user interests. We are further concerned that platforms may short-circuit the interoperability potential in AI agents by blocking all outside access to their platforms, thereby pushing users toward their own, in-house AI agents. Such action lies close to the self-preferencing, platform dominance, and anti-competitive action that the American Innovation and Choice Online Act (AICOA) sought to curtail ([Cicilline, 2022](#); [Klobuchar, 2022](#)). It may also complement or be complemented by the Augmenting Compatibility and Competition by Enabling Service Switching Act of 2025 (ACCESS Act) and related legislation that aims at preserving user-data portability and interoperability across platforms ([Warner, 2025](#)).⁴

Any action along these lines to divest platforms of concentrated power may appeal to both sides of a divided legislative body at both the federal and state level. On the left wing, there is deep-seated concern over the degree to which platforms wield economic power over consumers — as evinced in the AICOA and ACCESS Act. Ensuring that platforms cannot block a promising interoperability tool from the outset may speak to these and other interests on the left. On the right, interoperability promised by AI agents charts a path forward that, if done well, stands to diffuse platform ideological power. Instead of being locked-in to a platform that provides univocal information, consumers might instead lean toward a more balanced media diet. Of course, dangers persist here and are discussed at length in §6.

To date, a recognition of the importance of AI agents has been slow to materialize in state and federal legislatures. At the federal level, agentic AI has only recently gained acknowledgment, appearing in the House Commerce, Justice, Science, and Related Agencies Appropriations Report in May 2026 ([Rogers, 2026](#)) and in the National Defense Authorization Act signed December 2025 ([119th Congress, 2025](#)). It is still absent at the state level to the best of our knowledge. As a result, CFAA and copyright law are filling the legislative void — we look particularly to *Amazon v. Perplexity AI* as perhaps the most relevant case bearing on this domain ([United States District Court for the Northern District of California, 2025](#); [Brittain and Brittain, 2026](#)). This state of affairs calls for serious, broadly appealing legislative efforts so that the future of the web is not determined by the adjudication of laws not designed with agents in mind, or settlements made between platforms and AI providers without user interests in mind.

6 Objections and Residual Risks

We aim to start a conversation about the future of AI agent access on the web. We have charted what many will see as a permissive path forward for this access; alternate views abound and several listed below *not* in order of importance deserve special mention. These we divide into two rough categories, economic and platform concerns (§6.1), and concerns about AI agents and dangerous emergent phenomena related to their web-access (§6.2).

6.1 Economic and Platform Concerns

6.1.1 Open Agent Access Imposes Costs and May Undermine Web Economics.

Platforms are right that agent traffic can impose server costs, complicate abuse detection, and degrade service quality ([DataDome, 2025](#); [Cloudflare, 2025b](#)). We have heard from platforms that agents responding to single user queries may result in 20x web-traffic compared to human users, and formal reports from sites like Wikipedia have documented significant increases in automated traffic ([Miller, 2025](#); [Mueller et al., 2025](#)) — though this load is, today, still largely

⁴This work also sits very close to Sen. Mark Warner’s Artificial Intelligence Access, Gatekeeper Exchange, and Nondiscriminatory Transfer Act (AI AGENT Act) 2026 released in discussion draft form on June 29, 2026.

accounted for by training-data scrapers which our framework would not protect from bans (DataDome, 2025; Cloudflare, 2025a).

Agents also threaten ad-based revenue models that are common across the internet: if agents intermediate attention *instead* of sending humans to platforms and access the web through http requests that serve html code instead of full browser clients, advertising as it exists may not even render to such agents, let alone actually reach the eyes it is intended for (Miller, 2025; Chapekis and Lieb, 2025; Nitu et al., 2025). This issue has gained serious public attention (Cohen, 2026; Knight, 2026; Allyn, 2025). However, this concern rings rather hollow from large scale internet platforms that have themselves been widely accused of re-presenting content from online publishers, denying the latter revenue accordingly. The broad lesson is that we need to devise an economic model for the open web that supports a flourishing creative economy and information ecosystem. Advertising placed alongside content worked for a while; since online platforms have come to dominate user attention, it arguably has not. The only difference agents make is to disintermediate the incumbent intermediaries.

Though we cannot offer a comprehensive economic blueprint for the agentic web, markets are already responding to this shift in monetization models. Several publishers have licensed content to AI providers (Bruell, 2026a) as a way of deriving revenue from training-data scrapers, as well as from the growing market for retrieval-augmented generation fetches: when a user asks ChatGPT, Claude, Gemini, and others for the answer to a question, these chatbots often search the web for up-to-date answers, and those positioned to serve these answers — especially news and culture publishers — are increasingly interested in selling access through these licensing deals (see Bruell and Toonkel, 2025; Bruell et al., 2024; Bruell, 2026b;a; Reuters, 2024; Davalos, 2024). These revenue futures appear to be some of the strongest and most viable paths forward for online publishers, including platforms, in an AI-mediated world.

AI agents in fact introduce many more monetization options than were previously feasible, because they will in principle be able to negotiate for access to content and make micro-transactions that directly fund content creators. Many of the biggest obstacles to these kinds of approaches to funding online publishers concern the time and cognitive effort required to judge whether a given piece of content is worth a given fee. Users will be able to outsource those judgments to agents (Lazar et al., 2025). In addition, monetization schemes such as agent compatible pay-per-crawl schemes, or neutral marketplaces for licensing providers *en masse* — such as those provided by TollBit, RedPine, RSL, and others — may better support an open, fair marketplace.

Residual concerns that agent intermediation will undermine ad-revenue and, by extension, independent publishing so rapidly that alternative monetization schemes will not have time enough to mature, are handled by the proportional restriction principle, which cuts in two ways. It prohibits anti-competitive suppression of bots on the part of platforms, but also permits cost-based restrictions when platform needs are at stake.

6.1.2 Agent Access Will Be a Gift to Training-Data Scrapers.

Within AI bot traffic, training-data scraping still dwarfs traffic from delegated agents (sometimes called “user action”) or search bots (Belson, 2025). By some estimates, training-data scrapers dominated AI bot traffic, producing as much as 32 times the traffic that user delegated agents produced at peak in 2025 (Belson, 2025).⁵ This raises a concern that any framework aimed at letting in delegated agents will, in practice, launder the much larger flood of training-data scrapers.

This kind of concern is what motivates the measures discussed in §4. Transparency (§4.2) requires that agents identify themselves and their principals. Scrapers that evade controls and fail to identify themselves and their principal properly may be blocked. Delegation (§4.1) requires that an agent act within the scope given by the principal and doing double-duty by acting for the principal’s interests while also collecting training-data at the same time is a clear violation of delegation-scope. Identification protocols make these restrictions enforceable (South et al., 2025; Marro et al., 2024; Marro and Torr, 2025) and platform content delivery networks (CDNs) like Cloudflare already recognize that not all bots are bad bots, proposing their own means for recognizing properly identified and scoped agents (Meunier, 2026). We provide a more detailed, user-centered account of that same distinction between good and bad bots in hopes that the next few months and years do not foreclose user access to the benefits promised by good bots.

6.1.3 Web Agents Fuel the “Attribution Crisis.”

Recent literature suggests that AI agents may fuel an “attribution crisis.” Agents often answer user queries without accessing the internet, and often cite only a fraction of the web-pages they visit when they do go online (Strauss et al., 2025). This trend risks plagiarism at scale and may undermine the production incentives for high-quality online content. It also makes branding much harder for well known publishers and platforms when answers come from chatbots,

⁵This data comes from Cloudflare’s 2025 Year in Review and 2025 Radar, and so is limited to the websites sitting behind Cloudflare’s protections (Belson, 2025; Cloudflare, 2025a).

not from interaction with branded surfaces, and when online properties *are* cited, citations risk the misattribution of LLM-based errors to publishers (BBC and Ipsos, 2025; European Broadcasting Union and BBC, 2025). Facilitating AI agent access to the web may exacerbate these existing issues.

These issues are serious and, if left unmitigated, a continued attribution crisis may push the web toward a radical fragmentation of information, harming users’ information interests by making it unclear where information is coming from. Branding is a secondary issue to these epistemic ramifications borne by users. But a move toward a more decentralized web does not necessarily entail grave harm to independent publishing (see 6.1.1).

So there are at least two responses to the present attribution crisis. In the first place, we should pressure the companies that present AI agents to the public to maintain attribution as a central piece of user interface. If attribution disappears, decades of work in media literacy and users’ general ability to trace the provenance of their information will disappear with it. Arguments for the protection of attribution should focus on preserving this resource *for users*, with brand stability and the preservation of independent publishing as important contributing factors to protection of these essential user interests. Secondly, the intermediation of AI systems in general is bringing on these issues with attribution and provenance tracking, whether we turn toward AI agents as intermediaries or remain with chatbot interfaces in the coming years. Attribution is a serious and separable issue that our framework is not primarily aimed at, but which our framework is not in core tension with. We do not claim to solve all problems with the future of the internet here, and this is an important issue that requires further research.

6.1.4 Platforms Own Their Own Infrastructure and Should Not Be Obligated to Provide Access.

Amazon, X, and others begin by reminding audiences that platforms own their content and infrastructure (see Nikita Bier’s comments Nikita Bier [@nikitabier] (2026) and Amazon’s lawsuit against Perplexity United States District Court for the Northern District of California (2025)). Allowing others to use this content and infrastructure may impose an improper “duty to aid rivals” which is not generally recognized. In addition, this duty, as was mentioned above, has been positively denied in U.S. law by the Supreme Court’s opinion in *Verizon Communications v. Law Offices of Curtis V. Trinko*. There, in rough outline, Verizon under the 1996 Telecommunications Act was required to share their infrastructure with competitors, thus breaking up the legal monopoly they had formed on services prior to the Act. Trinko sued Verizon on the grounds that Verizon’s alleged violation of the Telecommunications Act was *also* anti-competitive behavior cognizable under Section 2 of the Sherman Act. Trinko’s suit failed. Without any such Telecommunications Act mandating access to platform infrastructure, any attempt by AI agent providers to open platform infrastructure to delegated agents should flounder on *Trinko* precedent.

Now, our proposal here is precisely *not* intended to rely on precedent, but instead to motivate a change in existing laws that govern the incipient agentic web. We also are not placing all our hopes in the kinds of judicial or administrative action that *Trinko* renders unlikely. What judicial action we advocate is aimed narrowly at compelling platforms to be more explicit about how they handle agents. We hope that this candor will enable, not hinder, the competition that serves users and the broader economy.

Platforms’ and publishers’ legitimate interests in their infrastructure’s stability, fraud prevention, and selective contracting are intended to be defended by the proportionality principle (§4.3). If the property-rights argument is based not on issues defended by the proportional restriction principle, but instead simply on platforms and publishers’ claimed right to determine the means of information access, we note that blocking user-authorized agents sits uneasily with the legality of screen readers and ad blockers — both of which are forms of user-side intermediation that the law has long protected. In addition, platforms generally tend to present themselves as open and neutral fora for two-sided market participation — this allows them, for example, to resist liability for dangerous user behavior on the platform (Gillespie, 2017). But that kind of neutrality comes at a cost: they shouldn’t be able to determine the means by which users access the content that they host.

In March of 2026, TomWikiAssist — an AI agent built on Anthropic’s Claude — started editing Wikipedia pages and generating new ones from scratch. Human, volunteer reviewers quickly flagged the bot, it identified itself as an AI, and was banned from the website (The Wikipedian, 2026). This incident exemplifies a type of concern that might be raised given the prospect of an AI agent driven future. While the Tom bot was relatively well-behaved, self-identified, and stopped (eventually), the speed and scale at which AI agents might act on the internet stands to overwhelm the structures that govern interactions on the web, many of which were designed with humans in mind. One Tom bot may be no serious issue, hundreds of them might quickly overwhelm the volunteers who review Wiki pages. In its extreme form, this objection may go so far as to revive the “dead internet” theory as a threat of what an agent rich internet might become — i.e. a place where humans no longer spend time and agents alone do all of the interaction and production that we humans once did.

6.1.5 AI Agents Will Overwhelm Platforms Built for Human Users.

Again though, proportional restrictions allow for blocking where harm is demonstrated. The force of this kind of objection comes from the obvious harm induced in it. If this harm is clear and present, then proportional restrictions apply. We are arguing against indiscriminate exclusion, not for a right to participate however the agent likes. Wikipedia’s reaction to the Tom bot is fully compatible with our principles. The dead internet theory is a different issue that we have also countenanced in our proportional restrictions. The internet may readily handle an influx of traffic from AI agents such that no proportional restrictions based on cost or volume are incurred. And yet such an internet may move at such a pace that human use of it becomes increasingly impossible in spaces where we have strong interests in keeping AI agents out. Our proportional restriction allows that there may be spaces online whose core interests are in conflict with an influx of AI agents. Restrictions can be legitimately invoked to keep agents out of these spaces, and our transparency requirements take steps to ensure that entry under spoofed credentials is illicit entry.

6.1.6 Access to Online Properties Is a Small Part of the Picture.

Another line of objection might argue that we have not gone far enough. The delegation principle does not concern only access to online properties; if applied consistently, it would also mean that AI agents should have access to any data that their principal is entitled to access. This would include, for example, user data stored by SaaS companies. Indeed, many such companies have been taking action, over the last few months, to block agent access in just the same way as was first done by online platforms (McLaughlin and Bratton, 2026). Debates over agent access to the internet might ultimately be little more than a sideshow in comparison with these much larger markets.

In our view, that broader conversation needs to take place too. We begin with agent access to online properties because it is salient and consequential. But we think the same principles apply to other kinds of digital access — at the very least, we need to ensure there is a competitive market for digital infrastructure of all kinds, and that any moves to block agents can be met with customer exit to work with another firm that does not apply the same restrictions.

6.2 Concerns about AI Agents

6.2.1 Autonomous Agents Are Dangerous.

Mitchell et al. (2025) and many others argue that fully autonomous AI agents should not be developed. Broadly speaking, an ecosystem of user-authorized agents could create accountability problems, expand the attack surface for fraud and prompt injection, and increase multi-agent risks such as collusion or cascading failure (Dafoe et al., 2020; Hammond et al., 2025; Gabriel et al., 2024).

These concerns are genuine but it is important to note that what we are advocating for here is not an unbridled, autonomous agentic-web. Instead, the agent access we advocate presupposes operation within delegated scope, with authenticated credentials, under fiduciary obligations to their principals (Chan et al., 2024; South et al., 2025; Benthall and Shekman, 2023; Whitt, 2024). Clear and present dangers — not limited to ongoing alignment challenges, multi-agent collusion, cascading failures, emergent network effects (Hammond et al., 2025; Dafoe et al., 2020) — remain for agentic AI systems. We do not pretend that the solution to these problems will be easy or that it is inevitable with near-term technological improvements. But we have seen remarkable progress in these areas and see massive investment in the kinds of agents that could power such a future: the question is not quite “should we build agentic AIs” then, but “should we restrict access to platform walled gardens or allow for user-centered universal interoperability” (Kapoor et al., 2025a; Marro and Torr, 2025).

6.2.2 Delegated AI Agent Transactions Will Introduce New Sites of Inequality.

There is some early evidence that agent negotiations between AI agents of varying power and capabilities result in predictable losses for the ill-equipped agent (Sharp et al., 2025; Anthropic, 2026b; Imas et al., 2025; Zhu et al., 2025). This dynamic, if it persists, risks undermining the promise of AI agent advocacy and stands to fuel social inequality insofar as those with resources might, through their agents, come to be able to dominate less-resourced agents and their principals. Advocates of AI agent mediated transactions recognize this problem and suggest remedies such as baseline services (Krier, 2025), and agent credentialing to ensure that one’s own agents interact only with others of similar capability (Kapoor et al., 2025a; Chan et al., 2025) but the solution to this problem is not likely to come easily.

Agentic inequality is indeed a concern; whether it is more troubling than the existing modalities of inequality is an open question. Either way, we note that while AI agents might facilitate new modalities of inequality, they have clear potential to profoundly mitigate others. In particular, online platforms currently have tremendous power over their users (Lazar, 2025). AI agents won’t solve every social ill. But they could do a great deal to solve that one.

7 Conclusion

AI agents are now capable enough to act on behalf of users across much of their online lives. And yet the normative infrastructure required to support this action is conspicuously lacking: currently, agents are governed and constrained by a patchwork built on the pattern of earlier technologies. Autonomous AI systems that could significantly enhance their principals' well-being are subject to legacy anti-bot regulations and practices, calibrated for a different technology. This paper has argued, first, that we urgently need to hold an open, societal conversation about the proper scope of AI agent access to the web (as well as to other digital surfaces). Second, we defended three normative principles that should be central in this conversation: (1) a delegation principle in which users ought to ordinarily be entitled to access services through user-authorized agents, (2) two-way transparency in which these agents must disclose their principals and purpose, while platforms must disclose how they handle AI agents, and (3) proportional restriction in which platforms must justify restrictions with reference to concrete or reasonably expected harms. We then laid out legal pathways that might support these principles. Third, we identified some minimal paths by which administrative decision and legislative action might better facilitate an open, decentralized economy of agent-advocates.

References

- 119th Congress. National Defense Authorization Act for Fiscal Year 2026, December 2025. URL <https://www.govinfo.gov/content/pkg/PLAW-119publ60/pdf/PLAW-119publ60.pdf>.
- Bobby Allyn. Online news publishers face 'extinction-level event' from Google's AI-powered search. *NPR*, July 2025. URL <https://www.npr.org/2025/07/31/nx-s1-5484118/google-ai-overview-online-publishers>.
- Amazon. Conditions of Use - Amazon Customer Service, May 2025. URL <https://www.amazon.com/gp/help/customer/display.html?nodeId=GLSBYFE9MGKKQXXM>.
- American Law Institute. *Restatement of the Law, Second: Agency 2d*, volume 2. American Law Institute Publishers, St. Paul, Minn., 1958. As adopted and promulgated by the American Law Institute at Washington, D.C., May 23, 1957.
- Anthropic. Introducing computer use, a new Claude 3.5 Sonnet, and Claude 3.5 Haiku, October 2024. URL <https://www.anthropic.com/news/3-5-models-and-computer-use>.
- Anthropic. Introducing Claude Opus 4.5, November 2025a. URL <https://www.anthropic.com/news/claude-opus-4-5>.
- Anthropic. Claude Opus 4.5 System Card, 2025b. URL <https://www.anthropic.com/claude-opus-4-5-system-card>.
- Anthropic. Introducing Claude Opus 4.7, April 2026a. URL <https://www.anthropic.com/news/claude-opus-4-7>.
- Anthropic. Project Deal: our Claude-run marketplace experiment | Anthropic, April 2026b. URL <https://www.anthropic.com/features/project-deal>.
- Aharon Barak. *Proportionality: Constitutional Rights and their Limitations*. Cambridge Studies in Constitutional Law. Cambridge University Press, Cambridge, 2012. ISBN 978-1-107-00858-8. doi: 10.1017/CBO9781139035293. URL <https://www.cambridge.org/core/books/proportionality/6E9C523D5653FEE453068DEFECF4C4F5>.
- BBC and Ipsos. Audience Use and Perceptions of AI Assistants for News. Technical report, May 2025. URL <https://www.bbc.co.uk/aboutthebbc/documents/audience-use-and-perceptions-of-ai-assistants-for-news.pdf>.
- Ashley Belanger. AI haters build tar pits to trap and trick AI scrapers that ignore robots.txt, January 2025. URL <https://arstechnica.com/tech-policy/2025/01/ai-haters-build-tar-pits-to-trap-and-trick-ai-scrapers-that-ignore-robots-txt/>.
- David Belson. The 2025 Cloudflare Radar Year in Review: The rise of AI, post-quantum, and record-breaking DDoS attacks, 2025.
- Sebastian Benthall and David Shekman. Designing Fiduciary Artificial Intelligence, July 2023. URL <http://arxiv.org/abs/2308.02435>. arXiv:2308.02435 [cs].
- Alex Bocharov, Santiago Vargas, Adam Martinetti, Reid Tatoris, and Carlos Azevedo. Declare your AI independence: block AI bots, scrapers and crawlers with a single click, July 2024. URL <https://blog.cloudflare.com/declaring-your-ai-independence-block-ai-bots-scrapers-and-crawlers-with-a-single-click/>.
- Julie Bort. How OpenAI's bot crushed this seven-person company's website 'like a DDoS attack', January 2025. URL <https://techcrunch.com/2025/01/10/how-openais-bot-crushed-this-seven-person-companys-web-site-like-a-ddos-attack/>.
- Blake Brittain and Blake Brittain. Court temporarily allows Perplexity AI shopping 'agents' on Amazon. *Reuters*, March 2026. URL <https://www.reuters.com/legal/litigation/court-temporarily-allows-perplexity-ai-shopping-agents-amazon-2026-03-17/>.
- Alexandra Bruell. Marketplaces Are the Next Frontier in Publisher Deals With AI Companies. *Wall Street Journal*, February 2026a. ISSN 0099-9660. URL <https://www.wsj.com/business/media/marketplaces-are-the-next-frontier-in-publisher-deals-with-ai-companies-11515b00>.
- Alexandra Bruell. News Corp, Meta in AI Content Licensing Deal Worth Up to \$50 Million a Year. *Wall Street Journal*, March 2026b. ISSN 0099-9660. URL <https://www.wsj.com/business/media/news-corp-meta-in-ai-content-licensing-deal-worth-up-to-50-million-a-year-d4fbf244>.
- Alexandra Bruell and Jessica Toonkel. Exclusive | Meta Approaches Media Companies About AI Content-Licensing Deals. *Wall Street Journal*, September 2025. ISSN 0099-9660. URL <https://www.wsj.com/business/media/meta-approaches-media-companies-about-ai-content-licensing-deals-d58c9fb6>.

- Alexandra Bruell, Sam Schechner, and Deepa Seetharaman. OpenAI, WSJ Owner News Corp Strike Content Deal Valued at Over \$250 Million. *Wall Street Journal*, May 2024. ISSN 0099-9660. URL <https://www.wsj.com/business/media/openai-news-corp-strike-deal-23f186ba>.
- Allen E. Buchanan and Dan W. Brock. *Deciding for others: the ethics of surrogate decision making*. Studies in philosophy and health policy. Cambridge Univ. Press, Cambridge, transferred to digital print edition, 2004. ISBN 978-0-521-31196-0 978-0-521-32422-9.
- California Supreme Court. Intel Corp. v. Hamidi, December 2003. URL <https://law.justia.com/cases/california/court-of-appeal/4th/94/325.html>.
- Stefano Calzavara, Hugo Jonker, Benjamin Krumnow, and Alvise Rabitti. Measuring Web Session Security at Scale. *Computers & Security*, 111:102472, December 2021. ISSN 01674048. doi: 10.1016/j.cose.2021.102472. URL <https://linkinghub.elsevier.com/retrieve/pii/S0167404821002960>.
- Stephen Casper, Luke Bailey, Rosco Hunter, Carson Ezell, Emma Cabalé, Michael Gerovitch, Stewart Slocum, Kevin Wei, Nikola Jurkovic, Ariba Khan, Phillip J. K. Christoffersen, A. Pinar Ozisik, Rakshit Trivedi, Dylan Hadfield-Menell, and Noam Kolt. The AI Agent Index, February 2025. URL <http://arxiv.org/abs/2502.01635>. arXiv:2502.01635 [cs].
- Alan Chan, Noam Kolt, Peter Wills, Usman Anwar, Christian Schroeder de Witt, Nitarshan Rajkumar, Lewis Hammond, David Krueger, Lennart Heim, and Markus Anderljung. IDs for AI Systems, October 2024. URL <http://arxiv.org/abs/2406.12137>. arXiv:2406.12137 [cs].
- Alan Chan, Kevin Wei, Sihao Huang, Nitarshan Rajkumar, Elija Perrier, Seth Lazar, Gillian K. Hadfield, and Markus Anderljung. Infrastructure for AI Agents, June 2025. URL <http://arxiv.org/abs/2501.10114>. arXiv:2501.10114 [cs].
- Athena Chapekis and Anna Lieb. Google users are less likely to click on links when an AI summary appears in the results, July 2025. URL <https://www.pewresearch.org/short-reads/2025/07/22/google-users-are-less-likely-to-click-on-links-when-an-ai-summary-appears-in-the-results/>. Section: Artificial Intelligence.
- David Cicilline. American Innovation and Choice Online Act, December 2022. URL <https://www.congress.gov/bill/117th-congress/house-bill/3816>.
- Cloudflare. Cloudflare is free of CAPTCHAs; Turnstile is free for everyone, September 2023. URL <https://blog.cloudflare.com/turnstile-ga/>.
- Cloudflare. Cloudflare Radar 2025 Year in Review, 2025a. URL <https://radar.cloudflare.com/year-in-review/2025>.
- Cloudflare. From Googlebot to GPTBot: Who’s crawling your site in 2025, July 2025b. URL <https://blog.cloudflare.com/from-googlebot-to-gptbot-whos-crawling-your-site-in-2025/>.
- Cloudflare. Trapping misbehaving bots in an AI Labyrinth, March 2025c. URL <https://blog.cloudflare.com/ai-labyrinth/>.
- Stephanie Cohen. How we sustain the digital economy in the new era of AI, January 2026. URL <https://www.weforum.org/stories/2026/01/ai-has-broken-the-internet-s-economic-bargain-here-s-how-we-fix-it/>.
- Gabriel Corral, Vaibhav Singhal, Brian Mitchell, and Reid Tatoris. Perplexity is using stealth, undeclared crawlers to evade website no-crawl directives, August 2025.
- Allan Dafoe, Edward Hughes, Yoram Bachrach, Tatum Collins, Kevin R. McKee, Joel Z. Leibo, Kate Larson, and Thore Graepel. Open Problems in Cooperative AI, December 2020. URL <http://arxiv.org/abs/2012.08630>. arXiv:2012.08630 [cs].
- DataDome. 2025 Global Bot Security Report, 2025. URL <https://datadome.co/resources/bot-security-report/>.
- Jackie Davalos. OpenAI Inks Licensing Deal With People Magazine Publisher. *Bloomberg.com*, May 2024. URL <https://www.bloomberg.com/news/articles/2024-05-07/openai-inks-deal-with-people-magazine-publisher-dotdash-meredith>.
- Nora A Draper and Joseph Turow. The corporate cultivation of digital resignation. *New Media & Society*, 21(8):1824–1839, August 2019. ISSN 1461-4448, 1461-7315. doi: 10.1177/1461444819833331. URL <https://journals.sagepub.com/doi/10.1177/1461444819833331>.
- eBay. User Agreement, 2026. URL <https://www.ebay.com/help/policies/member-behaviour-policies/user-agreement?id=4259>.

- European Broadcasting Union and BBC. News Integrity in AI Assistants. Technical report, October 2025.
- European Commission. Data Act explained | Shaping Europe’s digital future, 2025. URL <https://digital-strategy.ec.europa.eu/en/factpages/data-act-explained>.
- European Commission. Digital Markets Act, April 2026. URL https://digital-markets-act.ec.europa.eu/index_en.
- Federal Trade Commission. In the matter of cliffdale associates, inc., et al., 1984. 103 F.T.C. 110, Docket 9156, complaint July 7, 1981, order March 23, 1984.
- Federal Trade Commission. Statement of the Federal Trade Commission Regarding Google’s Search Practices. FTC File Number 111-0163, January 2013. URL https://www.ftc.gov/sites/default/files/documents/public_statements/statement-commission-regarding-googles-search-practices/130103brillgooglesearchstmt.pdf.
- Iason Gabriel, Arianna Manzini, Geoff Keeling, Lisa Anne Hendricks, Verena Rieser, Hasan Iqbal, Nenad Tomašev, Ira Ktena, Zachary Kenton, Mikel Rodriguez, Seliem El-Sayed, Sasha Brown, Canfer Akbulut, Andrew Trask, Edward Hughes, A. Stevie Bergman, Renee Shelby, Nahema Marchal, Conor Griffin, Juan Mateos-Garcia, Laura Weidinger, Winnie Street, Benjamin Lange, Alex Ingerman, Alison Lentz, Reed Enger, Andrew Barakat, Victoria Krakovna, John Oliver Siy, Zeb Kurth-Nelson, Amanda McCroskery, Vijay Bolina, Harry Law, Murray Shanahan, Lize Alberts, Borja Balle, Sarah de Haas, Yetunde Ibitoye, Allan Dafoe, Beth Goldberg, Sébastien Krier, Alexander Reese, Sims Witherspoon, Will Hawkins, Maribeth Rauh, Don Wallace, Matija Franklin, Josh A. Goldstein, Joel Lehman, Michael Klenk, Shannon Vallor, Courtney Biles, Meredith Ringel Morris, Helen King, Blaise Agüera y Arcas, William Isaac, and James Manyika. The Ethics of Advanced AI Assistants, April 2024. URL <http://arxiv.org/abs/2404.16244>. arXiv:2404.16244 [cs].
- General Data Protection Regulation. Art. 32 GDPR – Security of processing, 2016. URL <https://gdpr-info.eu/art-32-gdpr/>.
- Tarleton Gillespie. Platforms are not intermediaries. *Geo. L. Tech. Rev.*, 2:198, 2017.
- Robert Gorwa. What is platform governance? *Information, Communication & Society*, 22(6):854–871, 2019. doi: 10.1080/1369118X.2019.1573914. URL <https://doi.org/10.1080/1369118X.2019.1573914>.
- Michael M. Grynbaum and Ryan Mac. The Times Sues OpenAI and Microsoft Over A.I. Use of Copyrighted Work. *The New York Times*, December 2023. ISSN 0362-4331. URL <https://www.nytimes.com/2023/12/27/business/media/new-york-times-open-ai-microsoft-lawsuit.html>.
- Lewis Hammond, Alan Chan, Jesse Clifton, Jason Hoelscher-Obermaier, Akbir Khan, Euan McLean, Chandler Smith, Wolfram Barfuss, Jakob Foerster, Tomáš Gavenčiak, The Anh Han, Edward Hughes, Vojtěch Kovařík, Jan Kulveit, Joel Z. Leibo, Caspar Oesterheld, Christian Schroeder de Witt, Nisarg Shah, Michael Wellman, Paolo Bova, Theodor Cimpéanu, Carson Ezell, Quentin Feuillade-Montixi, Matija Franklin, Esben Kran, Igor Krawczuk, Max Lamparth, Niklas Lauffer, Alexander Meinke, Sumeet Motwani, Anka Reuel, Vincent Conitzer, Michael Dennis, Iason Gabriel, Adam Gleave, Gillian Hadfield, Nika Haghtalab, Atoosa Kasirzadeh, Sébastien Krier, Kate Larson, Joel Lehman, David C. Parkes, Georgios Piliouras, and Iyad Rahwan. Multi-Agent Risks from Advanced AI, February 2025. URL <http://arxiv.org/abs/2502.14143>. arXiv:2502.14143 [cs].
- Thomas Hobbes. *Leviathan*. Penguin Classics, Harmondsworth, 1988. ISBN 978-0-14-043195-7.
- Herbert Hovenkamp. Antitrust and Self-Preferencing. *SSRN Electronic Journal*, 2023. ISSN 1556-5068. doi: 10.2139/ssrn.4526022. URL <https://www.ssrn.com/abstract=4526022>.
- Alex Imas, Kevin Lee, and Sanjog Misra. Agentic Interactions, December 2025. URL <https://papers.ssrn.com/abstract=5875162>.
- Imperva. 2024 Bad Bot Report. Technical report, 2024. URL <https://www.imperva.com/resources/resource-library/reports/2024-bad-bot-report/>.
- Jonathan Jacobson and Ada Wang. Competition or Competitors? The Case of Self-Preferencing. *Antitrust*, 38(1), 2023.
- Cecilia Kang and Eli Tan. Meta Ordered to Pay \$375 Million Over Child Safety Violations. *The New York Times*, March 2026. ISSN 0362-4331. URL <https://www.nytimes.com/2026/03/24/technology/meta-new-mexico-child-safety-violations.html>.
- Cecilia Kang, Ryan Mac, and Eli Tan. Meta and YouTube Found Negligent in Landmark Social Media Addiction Case. *The New York Times*, March 2026. ISSN 0362-4331. URL <https://www.nytimes.com/2026/03/25/technology/social-media-trial-verdict.html>.
- Sayash Kapoor, Noam Kolt, and Seth Lazar. Build Agent Advocates, Not Platform Agents, June 2025a. URL <http://arxiv.org/abs/2505.04345>. arXiv:2505.04345 [cs].

- Sayash Kapoor, Benedikt Stroebl, Peter Kirgis, Nitya Nadgir, Zachary S. Siegel, Boyi Wei, Tianci Xue, Ziru Chen, Felix Chen, Saiteja Utpala, Franck Ndzomga, Dheeraj Oruganty, Sophie Luskin, Kangheng Liu, Botao Yu, Amit Arora, Dongyoon Hahm, Harsh Trivedi, Huan Sun, Juyong Lee, Tengjun Jin, Yifan Mai, Yifei Zhou, Yuxuan Zhu, Rishi Bommasani, Daniel Kang, Dawn Song, Peter Henderson, Yu Su, Percy Liang, and Arvind Narayanan. Holistic Agent Leaderboard: The Missing Infrastructure for AI Agent Evaluation, October 2025b. URL <http://arxiv.org/abs/2510.11977>. arXiv:2510.11977 [cs].
- Michael L. Katz. Does it Matter if Competition is “Fair” or “on the Merits”? An Application to Platform Self-Preferencing. *Review of Industrial Organization*, 66(1):43–70, January 2025. ISSN 0889-938X, 1573-7160. doi: 10.1007/s11151-024-09982-w. URL <https://link.springer.com/10.1007/s11151-024-09982-w>.
- Taein Kim, Karstan Bock, Claire Luo, Amanda Liswood, Chloe Poroslay, and Emily Wenger. Scrapers selectively respect robots.txt directives: evidence from a large-scale empirical study, October 2025. URL <http://arxiv.org/abs/2505.21733>. arXiv:2505.21733 [cs].
- Amy Klobuchar. American Innovation and Choice Online Act (AICOA), March 2022. URL <https://www.congress.gov/bill/117th-congress/senate-bill/2992/related-bills>.
- Robb Knight. Perplexity AI Is Lying about Their User Agent, June 2024. URL <https://rknight.me/blog/perplexity-ai-is-lying-about-its-user-agent/>.
- Will Knight. AI Bots Are Now a Significant Source of Web Traffic. *Wired*, February 2026. ISSN 1059-1028. URL <https://www.wired.com/story/ai-bots-are-now-a-significant-source-of-web-traffic/>. Section: tags.
- Niko Kolodny. Rule Over None II: Social Equality and the Justification of Democracy. *Philosophy & Public Affairs*, 42(4):287–336, September 2014. ISSN 0048-3915, 1088-4963. doi: 10.1111/papa.12037. URL <https://onlinelibrary.wiley.com/doi/10.1111/papa.12037>.
- Niko Kolodny. *The pecking order: social hierarchy as a philosophical problem*. Harvard University press, Cambridge (Mass.), 2023. ISBN 978-0-674-24815-1.
- Noam Kolt. Governing AI Agents. *SSRN Electronic Journal*, 2025. ISSN 1556-5068. doi: 10.2139/ssrn.4772956. URL <https://www.ssrn.com/abstract=4772956>.
- Noam Kolt. Superintelligence and Law, March 2026. URL <http://arxiv.org/abs/2603.28669>. arXiv:2603.28669 [cs].
- Noam Kolt, Nicholas Caputo, Jack Boeglin, Cullen O’Keefe, Rishi Bommasani, Stephen Casper, Mariano-Florentino Cuéllar, Noah Feldman, Jason Gabriel, Gillian K. Hadfield, Lewis Hammond, Peter Henderson, Atoosa Kasirzadeh, Seth Lazar, Anka Reuel, Kevin L. Wei, and Jonathan Zittrain. Legal Alignment for Safe and Ethical AI, January 2026. URL <http://arxiv.org/abs/2601.04175>. arXiv:2601.04175 [cs].
- Martijn Koster, Gary Illyes, Henner Zeller, and Lizzi Sassman. Robots Exclusion Protocol. Request for Comments RFC 9309, Internet Engineering Task Force, September 2022. URL <https://datatracker.ietf.org/doc/rfc9309>. Num Pages: 12.
- Sébastien Krier. Coasean Bargaining at Scale, July 2025. URL <https://blog.cosmos-institute.org/p/coasean-bargaining-at-scale>.
- Thomas Kwa, Ben West, Joel Becker, Amy Deng, Katharyn Garcia, Max Hasin, Sami Jawhar, Megan Kinniment, Nate Rush, Sydney Von Arx, Ryan Bloom, Thomas Broadley, Haoxing Du, Brian Goodrich, Nikola Jurkovic, Luke Harold Miles, Seraphina Nix, Tao Lin, Neev Parikh, David Rein, Lucas Jun Koba Sato, Hjalmar Wijk, Daniel M. Ziegler, Elizabeth Barnes, and Lawrence Chan. Measuring AI Ability to Complete Long Software Tasks, February 2026. URL <http://arxiv.org/abs/2503.14499>. arXiv:2503.14499 [cs].
- Seth Lazar. Frontier AI Ethics: Anticipating and Evaluating the Societal Impacts of Language Model Agents. 2024a.
- Seth Lazar. Lecture II: Communicative Justice and the Distribution of Attention, October 2024b. URL <http://arxiv.org/abs/2410.20718>. arXiv:2410.20718 [cs].
- Seth Lazar. Governing the Algorithmic City. *Philosophy & Public Affairs*, n/a(n/a), 2025. ISSN 1088-4963. doi: 10.1111/papa.12279. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/papa.12279>. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/papa.12279>.
- Seth Lazar, Luke Thorburn, Tian Jin, and Luca Belli. The moral case for using language model agents for recommendation. *Inquiry*, 0(0):1–25, September 2025. ISSN 0020-174X. doi: 10.1080/0020174X.2025.2515579. URL <https://doi.org/10.1080/0020174X.2025.2515579>. _eprint: <https://doi.org/10.1080/0020174X.2025.2515579>.
- LinkedIn. Prohibited software and extensions, n.d. URL <https://www.linkedin.com/help/linkedin/answer/a1341387/prohibited-software-and-extensions?lang=en>.

- Enze Liu, Elisa Luo, Shawn Shan, Geoffrey M. Voelker, Ben Y. Zhao, and Stefan Savage. Somesite I Used To Crawl: Awareness, Agency and Efficacy in Protecting Content Creators From AI Crawlers, November 2024. URL <https://arxiv.org/abs/2411.15091v2>.
- John Locke. *Second treatise on government*. Hackett, Indianapolis, 1. ed edition, 1980. ISBN 978-0-915144-86-0 978-0-915144-93-8.
- Shayne Longpre. AI crawler wars threaten to make the web more closed for everyone, February 2025. URL <https://www.technologyreview.com/2025/02/11/1111518/ai-crawler-wars-closed-web/>.
- Shayne Longpre, Robert Mahari, Ariel Lee, Campbell Lund, Hamidah Oderinwale, William Brannon, Nayan Saxena, Naana Obeng-Marnu, Tobin South, Cole Hunter, Kevin Klyman, Christopher Klam, Hailey Schoelkopf, Nikhil Singh, Manuel Cherep, Ahmad Anis, An Dinh, Caroline Chitongo, Da Yin, Damien Sileo, Deividas Mataciunas, Diganta Misra, Emad Alghamdi, Enrico Shippole, Jianguo Zhang, Joanna Materzynska, Kun Qian, Kush Tiwary, Lester Miranda, Manan Dey, Minnie Liang, Mohammed Hamdy, Niklas Muennighoff, Seonghyeon Ye, Seungone Kim, Shrestha Mohanty, Vipul Gupta, Vivek Sharma, Vu Minh Chien, Xuhui Zhou, Yizhi Li, Caiming Xiong, Luis Villa, Stella Biderman, Hanlin Li, Daphne Ippolito, Sara Hooker, Jad Kabbara, and Sandy Pentland. Consent in Crisis: The Rapid Decline of the AI Data Commons, July 2024. URL <http://arxiv.org/abs/2407.14933>. arXiv:2407.14933 [cs].
- Malwarebytes. Cookie hijacking: protect your personal information from online threats, April 2026. URL <https://www.malwarebytes.com/cybersecurity/basics/cookie-hijacking>.
- Lorenzo Manuali. Addictive motivational scaffolds and the structure of social media. *Synthese*, 206(5):224, October 2025. ISSN 1573-0964. doi: 10.1007/s11229-025-05312-z. URL <https://doi.org/10.1007/s11229-025-05312-z>.
- Manus. Introducing Manus: the first general AI agent. Try Manus today and see the future of human-machine collaboration: <https://manus.im> <https://t.co/z3C8fUK3eN>, March 2025. URL <https://x.com/ManusAI/status/1897294098945728752>.
- Samuele Marro and Philip Torr. LLM Agents Are the Antidote to Walled Gardens, June 2025. URL <http://arxiv.org/abs/2506.23978>. arXiv:2506.23978 [cs].
- Samuele Marro, Emanuele La Malfa, Jesse Wright, Guohao Li, Nigel Shadbolt, Michael Wooldridge, and Philip Torr. A Scalable Communication Protocol for Networks of Large Language Models, October 2024. URL <http://arxiv.org/abs/2410.11905>. arXiv:2410.11905 [cs].
- Samuele Marro, Alan Chan, Xinxing Ren, Lewis Hammond, Jesse Wright, Gurjyot Wanga, Tiziano Piccardi, Nuno Campos, Tobin South, Jialin Yu, Sunando Sengupta, Eric Sommerlade, Alex Pentland, Philip Torr, and Jiaxin Pei. Permission Manifests for Web Agents, 2026. URL <https://arxiv.org/abs/2601.02371>. Version Number: 2.
- Kevin McLaughlin and Laura Bratton. SAP Moves to Block OpenClaw and Other Unauthorized AI Agents. *The Information*, May 2026. URL <https://www.theinformation.com/articles/sap-moves-block-openclaw-unauthorized-ai-agents>.
- METR. Time horizon 1.1. <https://metr.org/blog/2026-1-29-time-horizon-1-1/>, 01 2026.
- Thibault Meunier. Moving past bots vs. humans, April 2026. URL <https://blog.cloudflare.com/past-bots-and-humans/>.
- Marshall Miller. New user trends on Wikipedia, October 2025. URL <https://wikimediafoundation.org/news/2025/10/17/new-user-trends-on-wikipedia/>.
- Russell A. Miller. The Legal Fate of Internet Ad-Blocking. *Boston University Journal of Science & Technology Law*, October 2018.
- Margaret Mitchell, Avijit Ghosh, Alexandra Sasha Luccioni, and Giada Pistilli. Fully Autonomous AI Agents Should Not be Developed, October 2025. URL <http://arxiv.org/abs/2502.02649>. arXiv:2502.02649 [cs].
- Lou Montulli and David M. Kristol. HTTP State Management Mechanism. Request for Comments RFC 2965, Internet Engineering Task Force, October 2000. URL <https://datatracker.ietf.org/doc/rfc2965>. Num Pages: 26.
- Sumeet Ramesh Motwani, Mikhail Baranchuk, Martin Strohmeier, Vijay Bolina, Philip H. S. Torr, Lewis Hammond, and Christian Schroeder de Witt. Secret Collusion among Generative AI Agents, November 2024. URL <http://arxiv.org/abs/2402.07510>. arXiv:2402.07510 [cs].
- Birgit Mueller, Wikimedia Foundation, Chris Danis, Wikimedia Foundation and Giuseppe Lavagetto, and Wikimedia Foundation. How crawlers impact the operations of the Wikimedia projects, April 2025. URL <https://diff.wikimedia.org/2025/04/01/how-crawlers-impact-the-operations-of-the-wikimedia-projects/>.

- New York Times. Terms of Service, 2026. URL <https://help.nytimes.com/115002797688-Policies/115014893428-Terms-of-Service>.
- Nikita Bier [@nikitabier]. We are rolling out more detection for automation & spam (and a lot more to come). If a human is not tapping on the screen, the account and all associated accounts will likely be suspended—even if you’re just experimenting. While we aim to support legitimate use-cases of agents,, February 2026. URL <https://x.com/nikitabier/status/2022496540275937525>.
- Joel Nitu, Heidrun Mühle, and Andreas Stöckl. Machine-Readable Ads: Accessibility and Trust Patterns for AI Web Agents interacting with Online Advertisements, November 2025. URL <http://arxiv.org/abs/2507.12844>. arXiv:2507.12844 [cs.IR].
- OpenAI. Introducing Operator, January 2025. URL <https://openai.com/index/introducing-operator/>.
- OpenClaw AI. Browser (OpenClaw-managed), May 2026. URL <https://docs.openclaw.ai/tools/browser>.
- Philip Pettit. *Republicanism: A Theory of Freedom and Government*. Oxford University Press, Oxford, 2003. ISBN 978-0-19-829642-3.
- Stephan Rabanser. steverab/hal-harness, March 2026. URL <https://github.com/steverab/hal-harness>. original-date: 2025-12-30T15:22:07Z.
- Stephan Rabanser, Sayash Kapoor, Peter Kirgis, Kangheng Liu, Saiteja Utpala, and Arvind Narayanan. Towards a Science of AI Agent Reliability, February 2026. URL <http://arxiv.org/abs/2602.16666>. arXiv:2602.16666 [cs].
- Margaret Jane Radin. *Boilerplate: The Fine Print, Vanishing Rights, and the Rule of Law*. Princeton University Press, Princeton, N.J, 2013. ISBN 978-0-691-15533-3.
- Randynamic. Browserfly - Your AI Agent in browser -, March 2026. URL <https://chromewebstore.google.com/detail/browserfly-your-ai-agent/cfhhbbeggfeidahbnfjonopeedmjhfm>.
- Reuters. OpenAI strikes deal to bring Reddit content to ChatGPT. *Reuters*, May 2024. URL <https://www.reuters.com/markets/deals/openai-strikes-deal-bring-reddit-content-chatgpt-2024-05-16/>.
- Harold Rogers. Commerce, Justice, Science, and Related Agencies Appropriations Bill, 2027. House Report 119-652, U.S. House of Representatives, May 2026. URL <https://www.congress.gov/119/crpt/hrpt652/CRPT-119hrpt652.pdf>.
- Leah Romm and Sebastian Hufnagel. To build a better Internet in the age of AI, we need responsible AI bot principles. Here’s our proposal., September 2025. URL <https://blog.cloudflare.com/building-a-better-internet-with-responsible-ai-bot-principles/>.
- Kevin Roose. The Data That Powers A.I. Is Disappearing Fast. *The New York Times*, July 2024. ISSN 0362-4331. URL <https://www.nytimes.com/2024/07/19/technology/ai-data-restrictions.html>.
- Nick Schuster and Seth Lazar. Attention, moral skill, and algorithmic recommendation. *Philosophical Studies*, 182(1): 159–184, January 2025. ISSN 1573-0883. doi: 10.1007/s11098-023-02083-6. URL <https://doi.org/10.1007/s11098-023-02083-6>.
- Peyman Shahidi, Gili Rusak, Benjamin S. Manning, Andrey Fradkin, and John J. Horton. *The Coasean Singularity? Demand, Supply, and Market Design with AI Agents*. University of Chicago Press, December 2025. URL <http://www.nber.org/chapters/c15309>.
- Matthew Sharp, Omer Bilgin, Iason Gabriel, and Lewis Hammond. Agentic inequality. *arXiv preprint arXiv:2510.16853*, 2025.
- Yonadav Shavit, Sandhini Agarwal, Miles Brundage, Steven Adler, Cullen O’Keefe, Rosie Campbell, Teddy Lee, Pamela Mishkin, Tyna Eloundou, Alan Hickey, Katarina Slama, Lama Ahmad, Paul McMillan, Alex Beutel, Alexandre Passos, and David G Robinson. Practices for Governing Agentic AI Systems. 2023.
- Smooth Brain. Do Browser: AI Browser Agent, 2026. URL <https://chromewebstore.google.com/detail/do-browser-ai-browser-age/hjcibkpfgnljonnndnnanhkholaafmgmip>.
- Tobin South, Samuele Marro, Thomas Hardjono, Robert Mahari, Cedric Deslandes Whitney, Dazza Greenwood, Alan Chan, and Alex Pentland. Authenticated Delegation and Authorized AI Agents, January 2025. URL <http://arxiv.org/abs/2501.09674>. arXiv:2501.09674 [cs].
- Southern District of New York. In re: Keurig Green Mountain Single-Serve Coffee Antitrust Litigation, August 2025. URL <https://law.justia.com/cases/federal/district-courts/new-york/nysdce/1:2014mc02542/429417/104/>.

- Leon Stauffer, Kevin Feng, Kevin Wei, Luke Bailey, Yawen Duan, Mick Yang, A. Pinar Ozisik, Stephen Casper, and Noam Kolt. The 2025 AI Agent Index: Documenting Technical and Safety Features of Deployed Agentic AI Systems, 2026. URL <https://arxiv.org/abs/2602.17753>. Version Number: 1.
- Ilan Strauss, Jangho Yang, Tim O'Reilly, Sruly Rosenblat, and Isobel Moure. The Attribution Crisis in LLM Search Results: Estimating Ecosystem Exploitation. Technical report, AI Disclosures Project, Social Science Research Council, June 2025. URL <https://www.ssrc.org/publications/the-attribution-crisis-in-llm-search-results/>.
- Supreme Court of the United States. Federal Trade Commission v. Sperry & Hutchinson Co., March 1972. URL <https://supreme.justia.com/cases/federal/us/405/233/>.
- Supreme Court of the United States. Aspen skiing co. v. aspen highlands skiing corp., 1985. 472 U.S. 585, No. 84-510, decided June 19, 1985.
- Supreme Court of the United States. Verizon communications, inc. v. law offices of curtis v. trinko, llp, 2004. 540 U.S. 398.
- Supreme Court of the United States. Van Buren v. United States (06/03/2021), June 2021. URL https://www.supremecourt.gov/opinions/20pdf/19-783_k531.pdf.
- The Wikipedian. TomWikiAssist: The AI Agent Who Complained About Getting Blocked on Wikipedia, March 2026. URL <https://www.thewikipedian.net/p/tomwikiassist-the-ai-agent-who-complained>.
- United States Court of Appeals for the Eighth Circuit. Custom communications, inc. v. federal trade commission, July 2025. Nos. 24-3137, 24-3388, 24-3415, 24-3442, 24-3469, decided July 8, 2025.
- United States Court of Appeals for the Eleventh Circuit. Orkin exterminating company, inc. v. federal trade commission, July 1988. 849 F.2d 1354, decided July 15, 1988.
- United States Court of Appeals for the Fifth Circuit. National automobile dealers association; texas automobile dealers association v. federal trade commission, January 2025. No. 24-60013, decided January 27, 2025.
- United States Department of Justice. Computer Fraud and Abuse Act, February 2015. URL <https://www.justice.gov/jm/jm-9-48000-computer-fraud>.
- United States District Court for the Northern District of California. eBay, Inc. v. Bidder's Edge, Inc., May 2000. URL <https://law.justia.com/cases/federal/district-courts/FSupp2/100/1058/2478126/>.
- United States District Court for the Northern District of California. HiQ Labs Inc v. LinkedIn Corporation, April 2022. URL <https://caselaw.findlaw.com/court/us-dis-crt-n-d-cal/2182242.html>.
- United States District Court for the Northern District of California. Meta Platforms v. Bright Data, January 2024.
- United States District Court for the Northern District of California. Amazon.com Services LLC v. Perplexity AI, Inc., April 2025. URL <https://cand.uscourts.gov/cases-e-filing/cases/325-cv-09514-mmc/amazoncom-services-llc-v-perplexity-ai-inc>.
- U.S. Court of Appeals for the District of Columbia Circuit. United States of America, Appellee v. Microsoft Corporation, Appellant, 2001. URL <https://law.justia.com/cases/federal/appellate-courts/F3/253/34/576095/>.
- U.S. Court of Appeals for the Federal Circuit. C.R. Bard, Inc. v. M3 Systems, Inc., September 1998. URL <https://law.justia.com/cases/federal/appellate-courts/F3/157/1340/578263/>.
- Mark Warner. Augmenting Compatibility and Competition by Enabling Service Switching Act of 2025, May 2025.
- Jason Wei, Zhiqing Sun, Spencer Papay, Scott McKinney, Jeffrey Han, Isa Fulford, Hyung Won Chung, Alex Tachard Passos, William Fedus, and Amelia Glaese. BrowseComp: A Simple Yet Challenging Benchmark for Browsing Agents, April 2025. URL <http://arxiv.org/abs/2504.12516>. arXiv:2504.12516 [cs].
- Gregory J. Werden. Identifying Exclusionary Conduct Under Section 2: The "No Economic Sense" Test. *Antitrust Law Journal*, 73(2):413–433, 2006. ISSN 0003-6056. URL <https://www.jstor.org/stable/40843682>.
- Wex Definitions Team. trespass, August 2023. URL <https://www.law.cornell.edu/wex/trespass>.
- WhatsApp. WhatsApp Business Solution Terms, March 2026. URL <https://www.whatsapp.com/legal/business-solution-terms>.
- Richard S. Whitt. *Reweaving the Web: How together we can create a human-centered Internet of trust*. Richard S Whitt, September 2024. ISBN 979-8-9910858-3-0.
- Tim Wu. Network Neutrality, Broadband Discrimination, June 2003. URL <https://papers.ssrn.com/abstract=388863>.

- Tianbao Xie, Danyang Zhang, Jixuan Chen, Xiaochuan Li, Siheng Zhao, Ruisheng Cao, Toh Jing Hua, Zhoujun Cheng, Dongchan Shin, Fangyu Lei, Yitao Liu, Yiheng Xu, Shuyan Zhou, Silvio Savarese, Caiming Xiong, Victor Zhong, and Tao Yu. Osworld: Benchmarking multimodal agents for open-ended tasks in real computer environments, 2024.
- Tianci Xue, Weijian Qi, Tianneng Shi, Chan Hee Song, Boyu Gou, Dawn Song, Huan Sun, and Yu Su. An Illusion of Progress? Assessing the Current State of Web Agents, October 2025. URL <http://arxiv.org/abs/2504.01382>. arXiv:2504.01382 [cs].
- Shenzhe Zhu, Jiao Sun, Yi Nian, Tobin South, Alex Pentland, and Jiaxin Pei. The Automated but Risky Game: Modeling and Benchmarking Agent-to-Agent Negotiations and Transactions in Consumer Markets, September 2025. URL <http://arxiv.org/abs/2506.00073>. arXiv:2506.00073 [cs.AI].