Renegotiating the Social Contract of Search

Ryen W. White

Microsoft, Redmond, WA 98052, USA

Synopsis: *Artificial Intelligence (AI)-mediated search is quietly rewriting the implicit social contract between searchers, search providers, and content creators that has governed information seeking for over two decades. This column argues for deliberate renegotiation along six critical dimensions before the terms are set by default.*

Motivating example: A software engineer queries an AI-mediated search system on whether a popular open-source library has a known authentication vulnerability. The system synthesizes information from public security advisories, online developer forums, and software library documentation. It generates a confident, fluent answer and displays it immediately. In doing so, it omits a contradictory finding from a recently published Common Vulnerabilities and Exposures (CVE) advisory. It offers to update the dependency automatically. The engineer never sees the original advisories, never encounters the disagreement in the evidence, and never evaluates the sources. In the course of generating this answer and performing the update, every element of what once made search a tool for informed human decision-making, i.e., source visibility, searcher evaluation, and human agency, has been quietly overridden.

This scenario captures dynamics already in play. It reflects the emerging reality of AI-mediated search, in which search systems powered by Large Language Models (LLMs) do not merely retrieve documents from an online collection such as the World Wide Web, they also *mediate*—interpreting searcher intent, synthesizing information across sources, engaging in multi-turn conversations, and increasingly acting (with human confirmation or fully autonomously) on behalf of searchers [1–3]. While these new capabilities represent genuine technological advances, they also mark a fundamental shift in the implicit social contract that has governed search for over two decades.

That contract was simple: searchers posed queries, search systems returned ranked lists of results from identifiable sources, and humans retained the agency to evaluate, interpret, and decide what action should follow. The terms of a new contract are now being defined by the rapidly accelerating pace of AI innovation, often faster than scientists, policymakers, and the public can adapt to the changes and fully assess their implications. The window to participate in shaping these terms is narrowing. Younger cohorts are already adopting general-purpose AI systems such as ChatGPT as their primary means of online information seeking [4], and there have been early experiments with integrating advertising directly into AI-generated responses, signaling the possible emergence of a business model that could entrench this new paradigm for decades [5].

The Old Contract

The social contract of search was never codified in law or regulation. It emerged as a set of implicit norms, reinforced by searcher expectations, market competition, and occasional regulatory action, e.g., the U.S. Federal Trade Commission's guidelines on search advertising disclosure.[1] This informal contract rested on five identifiable terms.

First, *human agency*: searchers formulated queries, scanned retrieved results, and decided which sources to consult. The system retrieved; the human interpreted. Second, *source transparency*: results comprised clickable hyperlinks to identifiable, verifiable sources, enabling searchers to directly assess provenance and source credibility. Third, *distributed trust*: because searchers evaluated multiple sources, no single entity served as the source of ground truth. Trust was a judgment exercised by the searcher, not a property conferred by the system. Fourth, *content reciprocity*: content creators produced knowledge and received traffic and attribution in return, sustaining a virtuous cycle that funded journalism, scholarship, and public-interest information [6]. Fifth, *separation of commerce and information*: advertising was visually distinguished from organic results, allowing searchers to differentiate paid placement from editorially neutral retrieval.

The informality of this contract is precisely what makes it vulnerable. Each of its terms is now being strained, not through deliberate breach, but through the emergent

[1] https://www.ftc.gov/news-events/news/press-releases/2013/06/ftc-consumer-protection-staff-updates-agencys-guidance-search-engine-industry-need-distinguish

properties of new AI-powered systems whose capabilities outpace the norms that searchers have long relied upon. This is creating three significant fractures in the social contract of search that must be acknowledged and addressed.

Three Fractures

Agency inversion. The most consequential fracture concerns human agency. AI-mediated search systems now interpret latent intent, inferring what the searcher left unsaid, synthesize across sources without requiring the searcher to directly evaluate any of them, and increasingly act directly, e.g., booking reservations, completing purchases, and filing online forms [2][3]. The searcher's role in these processes is shifting from principal to observer. Where the previous contract positioned the human centrally as the decision-maker who used search as a tool, the emerging paradigm positions the AI as an intermediary that operates on the human's behalf, as best the system can infer from observed interactions. This shift is driven partly by searcher demand for reduced friction and partly by incremental system design choices, often made without explicit searcher consent. That many searchers welcome this shift, preferring a fluent answer to the effort of evaluating online sources, does not repair the fracture; it widens it, because a social contract shaped by convenience alone will not protect all the interests that it should—not least the livelihoods of content creators whose work sustains the very ecosystem that it depends on.

Knowledge commons collapse. AI-mediated search depends on human-created content: the articles, databases, analyses, insights, and reporting that provide the knowledge base to first train the LLMs and to power the retrieval-augmented generation that grounds LLMs in current, domain-relevant knowledge beyond their training data. Yet by synthesizing this content into direct answers, AI-mediated search systems reduce the incentive for searchers to visit original sources, eroding the traffic and revenue models that sustain content creation [6–8]. This is a classic tragedy of the commons: individually rational behavior (accepting an AI-synthesized answer rather than accessing an original source) produces collectively irrational outcomes. If the economic foundations of content creation collapse, the knowledge base that all AI systems (search-related or otherwise) depend on degrades, creating a self-undermining spiral.

Trust centralization. Under the old social contract, trust was distributed. Searchers assessed the credibility of individual sources (e.g., a medical journal, a government agency, a news outlet) and formed their judgments accordingly. Under AI-mediated search, this evaluative step is eliminated. Searchers receive a single synthesized response from a system whose reasoning can be opaque, whose source selection can be invisible, and whose confidence can be uncalibrated [9]. Experimental evidence demonstrates a measurable downstream consequence of these changes: searchers who receive LLM-generated answers rather than traditional search results develop systematically shallower knowledge, showing that the loss of evaluative engagement has significant cognitive consequences [10].

The common thread across these three fractures is a shift from distributed to centralized control with the advent of AI-mediated search systems: control over interpretation (of queries and sources), control over economic value, and control over epistemic authority. This concentration is the core policy concern raised in this column.

Not Just Another Interface Change

A natural objection to the challenges discussed thus far is that search has always evolved, from fixed web directories to dynamic ranked result lists, from desktop to mobile, and the social contract has continued to adapt accordingly. This objection underestimates the current shift in three respects.

First, previous transitions changed how searchers *accessed* information; this transition changes who interprets and acts on it. The move from Yahoo's content directory to Google's PageRank altered the retrieval mechanism but preserved the searcher's evaluative role. AI-mediated synthesis transfers that role to the system.

Second, every prior paradigm preserved source visibility. Searchers reviewed and selected among sources. In AI-mediated search, sources are obscured behind a system-generated response. The evaluative step that anchored the old contract (i.e., assessing provenance, comparing perspectives, exercising judgment) is largely removed by design.

Third, the pace of adoption outstrips societal adaptation. The transition from directories to PageRank unfolded over several years. AI-mediated search adoption is

occurring at an unprecedented pace, leaving no time for norms, regulations, or economic models to adjust organically.

The question, then, is not whether the social contract of search will change. It already has changed. The question is whether the new contract will emerge through deliberate, inclusive negotiation or whether it will take shape by default—incrementally, and without broad input.

Six Dimensions of a New Contract

There are six dimensions along which a new social contract for search could be negotiated. Each addresses one or more of the three fractures identified above, and each has precedent in existing regulatory or governance models. Table 1 summarizes these dimensions alongside the original and evolving contract terms.

1. Provenance disclosure. AI-mediated search systems could be designed to disclose how a response was constructed: which sources contributed, how they were considered and weighted, what was excluded, and what overall confidence the system assigns to its synthesis. Analogies exist in nutrition labeling and financial disclosure. Provenance is a system-facing requirement: it holds the AI accountable for the reasoning behind its output.

2. Unmediated source access. Separately, searchers should retain the ability to bypass the synthesis entirely and consult the underlying sources directly. This is a searcher-facing requirement: it preserves the evaluative capacity that was central to the old contract, ensuring that AI-mediated search remains a tool that searchers can scrutinize, not a tool they must blindly trust. The principle is analogous to rights of access in data protection law.

3. Graduated consent for autonomous action. Frameworks could define the degree of autonomy AI search systems exercise along a spectrum: *inform, suggest, act with confirmation* (from the searcher), or *act autonomously* (of its own accord) [11]. Giving searchers meaningful control over where on this spectrum their AI-mediated search systems operate, with defaults that favor human agency, especially in high-stakes

domains such as healthcare, law, and personal finance, would preserve a core element of the old contract while accommodating emergent AI capabilities.

4. Content ecosystem sustainability. The long-term viability of AI-mediated search depends on the vitality of the knowledge ecosystem it draws from. Sustaining that ecosystem—through licensing frameworks, revenue-sharing arrangements, or contributions to public knowledge commons (e.g., Wikipedia, open data portals, public domain datasets)—merits serious exploration. Partial precedents exist in creative industries, where streaming platforms contribute to content production funds, and in news media bargaining codes adopted in Australia and Canada, which require digital platforms to compensate news publishers for content surfaced on their platforms [7][8].

5. Cognitive autonomy protections in high-stakes domains. For consequential decisions in, e.g., health, legal, financial, and civic contexts, AI-mediated search systems could be designed to surface source material, alternative perspectives, and expressions of uncertainty before facilitating action, analogous to informed consent in medicine or cooling-off periods in consumer protection. The empirical evidence that AI-generated answers produce measurably shallower knowledge (noted earlier in this column) [10] suggests that this is a practical concern, not merely a theoretical one.

6. Contestability and redress. When an AI-mediated search system provides inaccurate guidance, misrepresents a source, or omits material evidence, affected parties currently have no formal mechanism to challenge or correct the output. The five dimensions listed above are preventive, while this one is corrective. Searchers, content creators, and other affected parties should have meaningful avenues to flag errors, request corrections, and seek accountability for consequential harms. Precedents exist in credit reporting law (e.g., the U.S. Fair Credit Reporting Act, which grants individuals the right to dispute inaccurate records) and in content moderation, where online platforms have developed structured appeals processes. Without such a redress mechanism, the other dimensions of a new social contract lack enforcement.

Who Negotiates?

The old social contract of search was set informally, through market dynamics and emergent norms. The new social contract will require more deliberate negotiation.

This will be challenging, because the stakeholders' interests may conflict and the ability to shape the terms is not equally distributed among them. Search providers benefit from consolidated attention, advertising revenue, and market power; searchers prefer reduced friction over informed control; and content creators bear the cost with little collective bargaining power. Yet the scale of impact, affecting billions of searchers, millions of content creators, and the epistemic infrastructure of democratic societies, makes a multi-stakeholder approach essential.

Renegotiation requires the scientific community (to define evaluation standards and identify harms), policymakers (to establish legal frameworks and enforce compliance), content creators and knowledge institutions (to represent the supply side of the knowledge commons), civil society (to advocate for equity and public interest), and technology companies (to implement technical requirements and participate in standard-setting) [12][13]. Precedents for governing shared infrastructure through negotiated frameworks exist in internet governance (e.g., ICANN), financial regulation (e.g., Basel Accords), and international environmental agreements (e.g., Paris Agreement).

Conclusion

The implicit social contract of search served humanity well for two decades, not because it was perfect, but because it preserved human agency at the center of the information-seeking process. AI-mediated search offers genuine advances: broader access to knowledge, reduced interaction costs, and powerful new capabilities. However, these advances entail a fundamental redistribution of agency, trust, and economic value. The six dimensions outlined here (provenance disclosure, source access, graduated consent, content sustainability, cognitive autonomy protections, contestability and redress) offer a framework to initiate that negotiation. Change will not wait for consensus. The question now is not whether the social contract of search will be rewritten, but by whom and in whose interest.

Table 1. The social contract of search: Old, evolving, and possible new terms.

Dimension	Old Contract: Established – as seen in traditional search engines	Current State: Evolving – as seen in new AI-mediated search systems	Possible New Terms
Human agency	Searcher evaluates, interprets, decides	AI interprets, synthesizes, acts	Graduated consent; searcher controls level of automation
Source transparency	Ranked list of identifiable sources	Sources often obscured by AI synthesis	Provenance disclosure
Trust model	Distributed across many sources	Centralized in AI system	Right to unmediated source access
Content economics	Traffic to creators sustains ecosystem	AI synthesis may reduce traffic to sources	Sustainability mechanisms (e.g., licensing, revenue-sharing)
Commerce and information boundary	Advertising separated from organic results	Advertisements embedded in AI-generated answers	Disclosure and separation requirements
High-stakes safeguards	Searcher caution (self-regulated)	AI-mediated action across all domains	Cognitive autonomy protections for consequential decisions
Accountability and redress	Searcher verifies claims against visible sources	Errors in AI-generated responses difficult to identify or challenge	Contestability mechanisms and formal redress processes

Acknowledgments

AI tools were used to assist with editing and refinement of this manuscript. The views expressed are the author's own and do not necessarily reflect those of Microsoft.

References

[1] R. W. White. "Advancing the search frontier with AI agents". Commun. ACM, vol. 67, pp. 54–65, 2024.

[2] C. Shah and R. W. White. "From to-do to ta-da: Transforming task-focused IR with generative AI". Proc. ACM SIGIR, 2025, pp. 3911–3921.

[3] Google DeepMind. "Project Mariner," 2024. [Online]. Available: https://deepmind.google/models/project-mariner/

[4] P. Dhanalakota. "The rise of the personal AI advisers". Fast Company, 2025.

[5] A. Capoot. "OpenAI to begin testing ads on ChatGPT in the U.S.". CNBC, 2026.

[6] N. Vincent. "The paradox of reuse, language models edition". Nmvg, 2022.

[7] E. Katsamakas and J. M. Sanchez-Cartas. "Generative artificial intelligence, content creation, and platforms". J. Ind. Compet. Trade, vol. 24, p. 19, 2024.

[8] T. Whithorn. "Google AI Overviews are here to stay". Coll. Res. Libr. News, vol. 86, p. 214, 2025.

[9] P.N. Venkit et al. "Search engines in the AI era: The false promise of factual and verifiable source-cited sources". Proc. ACM FAccT, 2025, pp. 1325–1340.

[10] S. Melumad and J. H. Yun. "Experimental evidence of the effects of large language models versus web search on depth of learning". PNAS Nexus, vol. 4, pgaf316, 2025.

[11] J. Heer. "Agency plus automation: Designing artificial intelligence into interactive systems". Proc. Natl. Acad. Sci. U.S.A., vol. 116, pp. 1844–1850, 2019.

[12] N. Helberger, J. Pierson, and T. Poell. "Governing online platforms: From contested to cooperative responsibility". Inf. Soc., vol. 34, pp. 1–14, 2018.

[13] T. Gillespie. Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions That Shape Social Media. New Haven, CT, USA: Yale Univ. Press, 2018.

Biography

Ryen W. White is a Vice President of Applied Science at Microsoft. Contact him at ryenw@microsoft.com.