

Impossible Consent: Generative Search, Browsewrap, and the Rights-Information Harm to Cognitively Disabled Users

Travis Gilly

ORCID: 0009-0007-2954-6313

Real Safety AI Foundation

Illinois, United States

Working paper, June 2026 (v1)

ABSTRACT

A person searches for the law that protects them and receives, at the top of the page, an authoritative answer that is wrong, accompanied by a cited source that confirms the error rather than correcting it. This paper examines that event as a distinct kind of harm, separate from the commercial deception that dominates discussion of generative search. When an automated system delivers false information about a person's own legal rights to a member of the protected class those rights exist to serve, the injury is not a lost sale but a right unasserted, a wrongful denial accepted, a deadline misjudged. I argue that this harm falls hardest on people with neurodevelopmental conditions, intellectual and developmental disabilities, and cognitive decline, who are least able to independently evaluate the answer and most exposed to acting on it, and I name the mechanism that defeats their last line of defense: counterfeit verifiability, in which the citation that is supposed to let a careful reader check the claim instead confirms the false claim. I then turn to the platform's anticipated defenses, which rest on terms a logged-out search user is said to accept by mere use. Drawing on the current acceptance language of a major provider, I show that this is browsewrap, the weakest form of online assent, and that for cognitively disabled users it is not merely weak but structurally unattainable. I call this condition impossible consent: agreement imputed to a person through a channel that the nature of their disability makes it impossible to satisfy. Because the provider's liability disclaimers and any dispute-resolution terms rest entirely on that imputed assent, they fail precisely against the population the harm reaches. The inaccessibility of the consent channel is itself an access problem, which moves the analysis from contract into disability law. I close by distinguishing what this clears, the defenses, from what remains open, the affirmative cause of action.

Keywords: browsewrap, contractual capacity, impossible consent, disability rights, generative search, AI Overviews, access to justice, ADA, counterfeit verifiability

I. INTRODUCTION

Consider an ordinary act. A person wants to know whether the law protects them in a situation they are living through, a workplace accommodation, a housing question, a benefits denial, the rights of a disabled child at school. They type the question into a search engine. At the top of the page, above the links, a generated summary answers in a confident,

neutral voice. The answer is wrong. It misstates the protection, or the deadline, or whether the person is covered at all. Beneath it sits a citation, a link to a source, presented as the basis for the answer. The person, doing the responsible thing, clicks the source to check. The source says the same wrong thing, or does not say what the summary claimed, and the error is confirmed rather than caught.

Most analysis of generative search treats its failures as a consumer problem: the model recommends an inferior product, or repeats a marketing claim as fact, and the harm is measured in misdirected purchases. That analysis is real, but it is not the gravest case, and it is not the case this paper is about. The event described above is a different kind of harm, because the subject is not a product but the reader's own legal rights, and the reader is frequently a member of the very class those rights were enacted to protect.

This paper makes an argument in two parts. The first concerns the harm itself. Automated misinformation about a person's legal rights, delivered with the authority of a neutral system and confirmed by its own cited source, is an injury distinct from commercial deception, it is regressive in its incidence, falling hardest on people with cognitive disabilities, and it is worsened by a mechanism I call counterfeit verifiability, in which the citation defeats the check it appears to invite. The second concerns the defense the platform will raise. A provider whose generated answer misinforms a logged-out user will point to terms of service the user is said to have accepted by using the service. I argue that for a logged-out user this is browwrap, the weakest form of online assent, and that for a cognitively disabled user it is not weak but impossible: a consent the offered channel makes structurally unattainable. Because the provider's liability limits and any dispute terms rest on that assent, they fail against the exact population the harm reaches, and the inaccessibility of the consent channel becomes an access problem in its own right.

II. THE HARM: DISTORTION OF RIGHTS INFORMATION

The commercial doctrines do not reach this case, and seeing why locates the harm precisely. A claim under the Lanham Act requires a false statement in commercial advertising of the defendant's own goods, and a deceptive-practices action under Section 5 of the Federal Trade Commission Act requires a deceptive act in commerce. A wrong statement of disability rights law is not a claim about a product and is not made to sell one. The commercial apparatus has nothing to grasp. This is not a weaker instance of false advertising. It is a different category, and the absence of a ready doctrine is part of what makes it worth naming.

What it is, stated plainly, is the distortion of a person's information about their own rights, delivered by an automated system, with authority, to someone who did not bargain for legal advice and frequently cannot tell that the advice is wrong. The consequences do not look like a refund problem. They look like a person not asserting a protection they hold, accepting a denial they were entitled to challenge, believing a filing deadline that is incorrect, or concluding that they fall outside a statute that in fact covers them. Each of these is a concrete loss of a legal entitlement, and each follows from acting on an answer that arrived wearing the costume of a settled fact.

The incidence of this harm is not uniform. It concentrates on the people least equipped to detect the error and most exposed to its effects: those navigating the legal system without counsel, and within that group, disproportionately, people with neurodevelopmental conditions, intellectual and developmental disabilities, and cognitive decline. These are the readers for whom independent verification is hardest, for whom the authority of a confident answer is most difficult to discount, and for whom the stakes of a missed right are often highest. The difficulty is documented. Self-represented litigants routinely struggle with legal procedures and concepts they do not understand (Jean-Louis, 2021), and people with cognitive impairments encounter barriers to justice when legal information reaches them in inaccessible forms, which breeds uncertainty and distrust of formal processes even when they want to use them (Harding and Keeling, 2024). A harm whose severity rises as the reader's capacity to check falls is regressive in the most literal sense, and it

lands on a protected class in the domain of the protections meant for them.

III. COUNTERFEIT VERIFIABILITY

The defense the industry offers against a model stating falsehoods is grounding: attach a citation, let the output rest on a retrievable source, and let the careful reader verify. The citation is sold as the antidote to fabrication and as the handle by which a diligent person checks a claim before relying on it. The premise of that defense is that the source, once consulted, will either confirm a true claim or expose a false one. That premise is fragile even on its own terms. Purpose-built legal research tools that use retrieval to ground their answers have been found to hallucinate in a substantial share of cases, between roughly 17 and 33 percent in one preregistered evaluation, despite being marketed as hallucination free (Magesh et al., 2025).

That premise fails in a specific and underexamined way. When the cited source is itself wrong, or does not in fact say what the summary attributes to it, the act of checking does not catch the error. It confirms it. The reader who does exactly the responsible thing, who does not take the answer on faith but follows the citation to its source, is rewarded with the same falsehood a second time, now apparently corroborated. I call this counterfeit verifiability: the presence of the apparatus of verification, a source, a link, a citation, around a claim that the apparatus does not actually support, such that the verification step becomes a laundering step.

Counterfeit verifiability is worse than a bare fabrication, and the reason matters. A naked false statement can be checked, and a careful reader can catch it. A false statement wrapped in a confirming citation defeats the check itself. It closes the one escape route that the diligent reader was supposed to have. This pattern has a close analog in human factors research, where an automated aid that erroneously validates a user's own initial mistake, a condition termed incorrect reassurance, produces worse final accuracy than other error patterns (Kim et al., 2024). And it compounds a long-documented tendency, automation bias, in which people presented with a computer generated answer accepted as correct disregard or fail to search for information that would contradict it (Cummings, 2004). And it does its worst to the reader who behaves best, because the person who shrugs and trusts the summary and the person who carefully follows the source arrive at the same wrong conclusion, the second more confident than the first. The harm is therefore sharpest for a reader who has the diligence to verify but not the domain expertise to evaluate the quality of the source, which describes a great many people researching their own legal rights, and describes with particular force the cognitively disabled reader who has been taught, correctly, to check sources but cannot independently assess whether a source is authoritative or current.

IV. THE CONSENT POSTURE OF LOGGED-OUT GENERATIVE SEARCH

When a generated answer harms a logged-out user, the provider's first move is not to defend the answer but to invoke its terms. It is therefore necessary to establish what, if anything, a logged-out search user actually agreed to. The current general terms of service of a major provider supply a clean specimen, and their structure is instructive precisely because the provider is sophisticated and the terms are recent.

The acceptance theory is browwrap. The terms state that a person is bound by using the services, and they extend that binding, in terms, to use whether or not the person is signed in to an account (Google, 2026). There is no checkbox, no acknowledgment, no moment at which the user clicks to accept. The agreement is imputed from the act of searching. This is the weakest recognized form of online assent. Courts subject it to heightened scrutiny and frequently refuse to enforce it, on the ground that a user may have no awareness of the terms at all; one provider had its entire terms of service voided for relying on browwrap, and the Ninth Circuit declined to enforce arbitration terms that a site failed to make conspicuous and failed to require users to acknowledge (*In re Zappos.com*, 2012; Berman, 2022). The leading

statement of the rule holds that even a conspicuous hyperlink to the terms, without some affirmative act of assent, does not create constructive notice (*Nguyen v. Barnes and Noble*, 2014), and courts have refused to enforce such terms where it was not clear the user knew of them at all (*Specht v. Netscape*, 2002).

Two further features of the specimen terms matter. First, the general terms do not send disputes to arbitration. They route disputes to the state and federal courts of a named California county and contain no class waiver (Google, 2026). Binding arbitration appears only in separate, product-specific agreements attached to things a person buys and signs into, and even there a provider seeking to enforce arbitration has had to prove, with records, that the user accepted the terms (*In re Google Digital Advertising*, 2025). A logged-out search user generates no such record, because there is no act of acceptance to record. Second, the general terms cap the provider's total liability at a small fixed sum or the fees paid, which for a free service is the small fixed sum, while expressly preserving liability for gross negligence and willful misconduct and disclaiming any intent to limit rights that cannot be limited by contract (Google, 2026).

The upshot is that the entire defensive structure available against a logged-out user, the liability cap and the disclaimers, rests on a single foundation: browsewrap assent imputed from the act of searching. There is no arbitration shield for this conduct and no class waiver. There is only the claim that by searching, the user agreed. Everything turns on whether that imputed agreement holds.

V. IMPOSSIBLE CONSENT

Browsewrap, where it is enforced at all, is enforced on a theory of constructive notice: a court asks whether a reasonable, prudent user would have been on notice of the terms, given their placement and conspicuousness. The standard is objective, and it is built around a particular imagined person. That person can perceive a hyperlink, infer that it carries legal significance, locate and open the terms, read dense contractual language, and understand that the act of using a free service has surrendered a legal right. The standard never asks whether the actual user could do any of this. It asks what a reasonable user would have noticed.

Examine who that reasonable user is. He is cognitively typical by construction. The entire chain of inferences the doctrine requires, from noticing the link to understanding the waiver, presumes a baseline of cognitive function that the standard treats as universal. A person with a neurodevelopmental condition, an intellectual or developmental disability, or cognitive decline cannot complete that chain through the channel the provider offered. The barrier is not inattention, which the reasonable-user standard already tolerates and binds. It is incapacity to perform the very operations the standard assumes anyone can perform. The doctrine sweeps such a person in anyway, because it never inquired into the actual person at all.

This is the condition I name impossible consent, and it should be distinguished sharply from a familiar cousin. Phantom consent is consent that is fictional in the ordinary way: the terms existed, the reasonable user could in principle have read them, and the law imputes agreement to a person who simply did not. The ordinary critique of boilerplate already presses on this fiction, observing that those who click are consenting to terms that are unread or, even if read, not understood, and may thereby surrender legal rights they do not know they hold (*Radin*, 2013; *Becher*, 2008). Phantom consent is a fiction about what a capable person did. Impossible consent is a fiction about what an incapable person could do. The reasonable user could have noticed and understood the terms and chose not to. The cognitively disabled user the doctrine binds could not have noticed and understood them through the offered channel, by reason of the disability itself. One cannot waive a right through a door that, by the nature of one's condition, one cannot find, open, or read. To call that constructive notice is to impute an act that was never available to be performed.

A second doctrine reinforces the point and does not depend on it. The law of contractual capacity holds that an agreement entered by a person who lacks the mental capacity to understand it is voidable (*Restatement (Second) of*

Contracts). So even setting the notice theory aside, an agreement imputed to a cognitively disabled person by mere use is weak twice over: there was no affirmative assent under the browsewrap theory, and there was no capacity to assent had assent been sought.

Two seams must be marked honestly, because the argument is strong enough not to need overstatement. First, the capacity defense as it currently stands is individualized. A litigant proves that a particular person lacked capacity; there is no rule that browsewrap categorically fails against every disabled person, and the flat claim that consent simply never applies to people with disabilities is the thesis being advanced here, not a holding that can be cited. The defensible and sharper claim is that the constructive-notice standard is doctrinally blind to disability, presumes a capacity it never tests, and, applied to a cognitively disabled user, manufactures a consent that was impossible to give. Second, a provider will argue that continued use supplies assent regardless. After the recent skepticism toward weak online assent, that argument is fragile without conspicuous notice and a genuine opportunity to understand, and its strength turns on the actual presentation of the terms on the logged-out page, which is an evidentiary fact a litigant should preserve as it appears.

VI. THE ACCESS DIMENSION

At this point the analysis leaves contract and enters disability law, because the defect is not only that a particular agreement is weak. The defect is that the mechanism of agreement is inaccessible to a class of people by design, and is then used against them.

A service is made available to everyone, and its generated answers are presented to everyone, signed in or not. The only purported agreement governing that universal service is one whose validity depends on a person performing cognitive operations that a recognized class of people cannot perform. The agreement is, in effect, gated behind an accessibility barrier, except that the barrier does not keep the disabled user out of the service. It keeps them out of the terms while binding them to the consequences. They receive the harm of the service, the rights-distorting answer, and are denied the comprehension that the terms presuppose, while still being held to the terms.

That structure is an access problem of the kind disability law exists to address. The right of people with cognitive disabilities to equal access to web content is itself a recognized concern under the Americans with Disabilities Act and allied instruments (Blanck, 2014), and the barrier here belongs to that family: terms and information presented in a form a cognitively disabled person cannot use, with the consequences imposed regardless (Harding and Keeling, 2024). The relevant question is not only whether a specific disabled person lacked capacity, but whether a provider may impute a binding waiver to a population through a channel that the provider knows, or should know, that population cannot use. Framed this way, the imputation of consent through an inaccessible channel is not a neutral contract default. It is the allocation of legal risk onto exactly the people least able to bear or even perceive it, in a domain, the assertion of one's own rights, where the protections of disability law are at their most central. The same standard that produces impossible consent in contract reappears here as a failure of access, and the two descriptions are of one underlying fact.

VII. WHAT THIS CLEARS, AND WHAT REMAINS OPEN

It is important to be exact about what the impossible-consent argument accomplishes, because it is easy to claim too much.

What it clears is the defense. A provider facing a claim that its generated answer misinformed a logged-out, cognitively disabled user about that user's legal rights will reach first for the terms of service: the liability cap, the warranty disclaimers, and, where it exists, the arbitration and class waiver. Every one of those rests on the user's assent.

For this user, on this channel, assent is impossible, the capacity to give it is absent, and, for general search, no arbitration clause governs the conduct in the first place. The defensive scaffolding therefore does not stand against the population the harm reaches. The claim proceeds in court, and because the affected population is definable, it is the kind of claim that can proceed collectively rather than being fractured into individual proceedings.

What remains open is the affirmative cause of action, and it is the harder problem. Clearing a defense is not the same as stating a claim. The substantive theory under which a provider owes a duty regarding the accuracy of generated statements about a person's legal rights is unsettled. Candidate vehicles exist, including state unfair and deceptive practices statutes, negligence theories built on a duty of reasonable care in the design and deployment of a system that delivers consequential information to vulnerable users, and disability-access theories aimed at the inaccessible consent channel itself. Each faces real obstacles, including the unresolved question of whether and when generated output is protected expression, which courts have so far declined to decide (*Garcia v. Character Technologies*, 2025). The contribution of this paper is not to resolve that question. It is to establish that when the claim is brought, the consent-based defenses that would otherwise dispose of it quickly do not reach the people most likely to bring it, and to name why.

VIII. CONCLUSION

Strip the case to its frame and it is this. A system delivers authoritative, wrong information about a person's legal rights to someone least able to detect the error, and confirms the error through its own citation when that person tries to check. The same system then disclaims responsibility by pointing to an agreement the person is said to have made by searching, an agreement that the person, by reason of a cognitive disability, could not have read or understood through the only channel offered. The system manufactures agreement from a person who could not give it, and then uses the manufactured agreement to disclaim the harm.

Several interventions follow without requiring the full argument to be accepted. Any terms meant to bind a user, and certainly any that limit a provider's liability for harm, should require affirmative and accessible assent rather than imputed assent, so that a person who cannot navigate the channel is not bound by it. A provider that delivers consequential information about legal rights to the public should bear a duty of reasonable care in the accuracy of that information and in the integrity of the citations that purport to support it, so that the verification handle is not counterfeit. And the constructive-notice standard's blindness to disability should be treated as a defect to be corrected rather than a convenience to be preserved, because a standard that imputes consent to people who cannot give it is not measuring agreement. It is dispensing with it. The people for whom this matters most are precisely the ones the standard renders invisible, and they are the ones the law of rights, of all laws, should be least willing to lose.

ACKNOWLEDGMENTS

The author discloses the use of speech-to-text software (Wispr Flow) and Anthropic’s Claude as assistive technology for a diagnosed neurodevelopmental condition. All legal analysis, case synthesis, theoretical framing, doctrinal arguments, and conclusions are solely the author’s own.

Correspondence concerning this article should be addressed to Travis Gilly, Real Safety AI Foundation, Illinois, United States. Email: t.gilly@ai-literacy-labs.org.

REFERENCES

- Becher, S. I. (2008). Asymmetric information in consumer contracts: The challenge that is yet to be met. *American Business Law Journal*, 45(4), 723–774.
- Berman v. Freedom Financial Network, LLC, 30 F.4th 849 (9th Cir. 2022). [Reporter cite, court, and date verified via CourtListener; verify any pinpoint page before print.]
- Blanck, P. (2014). The struggle for web eQuality by persons with cognitive disabilities. *Behavioral Sciences & the Law*, 32(1), 4–32.
- Cummings, M. L. (2004). Automation bias in intelligent time critical decision support systems. *AIAA 1st Intelligent Systems Technical Conference*, paper 2004-6313. [Verified via Scite; conference paper, locator is the AIAA paper number.]
- Garcia v. Character Technologies, Inc., 785 F. Supp. 3d 1157 (M.D. Fla. 2025). [Reporter cite, court, and date verified; holding that the court declined to rule AI output protected speech confirmed against the opinion.]
- Google. (2026). *Google terms of service*. Retrieved June 2026 from <https://policies.google.com/terms> [Quoted acceptance language and dispute, liability, and arbitration provisions current as of retrieval; reverify text and effective date before reliance.]
- Harding, R., & Keeling, A. (2024). Raising relational legal consciousness through co-production research? Making law more accessible. *Journal of Law and Society*, 51(S1), S102–S117.
- In re Google Digital Advertising Antitrust Litigation, 763 F. Supp. 3d 563 (S.D.N.Y. 2025) (also reported at 2025 WL 289726). [Published reporter cite confirmed via a primary federal source (J.P.M.L. transfer order, MDL No. 3010, Aug. 7, 2025); full opinion and docket are free on CourtListener (master docket 1:21-md-03010). No paid database used or needed.]
- In re Zappos.com, Inc., 893 F. Supp. 2d 1058 (D. Nev. 2012). [Reporter cite, court, and date verified via CourtListener; canonical short name used.]
- Jean-Louis, S. (2021). I don’t know what I’m doing: Using limited license legal technicians in family court to improve access to justice. *Family Court Review*, 59(3), 599–611.
- Kim, J. Y., Richie, O., & Yang, X. J. (2024). Evaluating the influence of incorrect reassurances on trust in imperfect automated decision aid. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 68(1), 81–87.
- Magesh, V., Surani, F., Dahl, M., Suzgun, M., Manning, C. D., & Ho, D. E. (2025). Hallucination-free? Assessing the reliability of leading AI legal research tools. *Journal of Empirical Legal Studies*, 22(2), 216–242.
- Nguyen v. Barnes & Noble Inc., 763 F.3d 1171 (9th Cir. 2014). [Reporter cite, court, and date verified via CourtListener;

verify any pinpoint page before print.]

Radin, M. J. (2013). *Boilerplate: The fine print, vanishing rights, and the rule of law*. Princeton University Press.

Restatement (Second) of Contracts §§ 12–16 (Am. L. Inst. 1981) (capacity to contract; § 15 governs mental illness or defect, rendering the contract voidable). [Black-letter authority; not verifiable through case-law databases.]

Specht v. Netscape Communications Corp., 306 F.3d 17 (2d Cir. 2002). [Reporter cite, court, and date verified via CourtListener.]