CLICKWRAP ACCOUNTABILITY

*Alicia Solow-Niederman**

Picture a public-facing generative artificial intelligence (AI) chatbot on a government website that can answer questions about benefits eligibility. Systems like these expose a sociotechnical-legal divide: They route around the standard due process framework because there has been no formal government determination, yet they induce reliance from users, above and beyond past rounds of automated legal guidance. Moreover, if there is an error, current doctrine on government errors and equitable estoppel makes relief unlikely.

This Essay focuses on government deployment of public-facing generative AI chatbots and critically assesses the front-end power dynamics that determine who is filtered out of government programs before traditional procedural protections even attach. The only reliable backstop right now is what this Essay calls "clickwrap accountability": a disclaimer borrowed from contract law that shifts responsibility for the tool's performance to its users.

Importing this private law accountability measure into public law is a mistake. And we should know better given the deficiencies of notice and choice consent paradigms in information privacy law and mounting critiques of clickwrap contracts in general. At a moment when more and more government actors are deploying generative AI chatbots for public-facing services, and when the federal government continues to push for greater AI adoption, we must insist on better public law protections than clickwrap accountability.

INTRODUCTION 570

I. DUE PROCESS AND THE GENERATIVE STATE 573

II. ACCOUNTABILITY GAPS: THE THINNESS OF CLICKWRAP 584

---

* Associate Professor of Law, The George Washington University. Thank you to Amit Haim, Sonia Katyal, Martha Minow, Naomi Schoenbaum, and participants in the *Fordham Law Review* Spring 2026 Symposium for helpful feedback, and to Rhyia Bibby and Charlsie Doan for expert research assistance. I also wish to thank the editors of the *Fordham Law Review*, especially Emma Miloglav, for their tireless work to bring this piece to publication. This Essay was substantively finalized for publication in early May 2026 and reflects developments through that time. All remaining errors and omissions are my own.

This Essay was prepared for the Symposium titled *Constitutional Law in the Age of Generative AI* hosted by the *Fordham Law Review* on February 27, 2026, at Fordham University School of Law.

A. Public Law Gaps 584
B. Private Law Moves 585
C. Public-Private Problems 587
III. ACCOUNTABILITY BY LAW AND BY DESIGN 590
CONCLUSION 597

INTRODUCTION

"To err is human,"[1] but who should pay when machines give erroneous advice? When Air Canada put a chatbot on its website and the chatbot incorrectly represented the company's bereavement policy, a Canadian tribunal found that the airline should bear the cost.[2] Specifically, the tribunal found Air Canada responsible for the chatbot's outputs, just as it was "for all the information on its website,"[3] and held that the plaintiff could succeed in his negligent representation theory and recover through tort law.[4] But if the government puts a chatbot on its website and that chatbot wrongly represents the government's policy, it is quite likely that there will not be redress for an individual who similarly relies on that advice.[5]

This Essay focuses on public-facing, government-provided artificial intelligence (AI) chatbots to expose how contemporary legal doctrines provide more protection for pocketbooks than for rights. The goal is not to canvass every possible chatbot but rather to reveal an emerging pattern and to position us to confront it before it creates further problems. When a private entity's errors affect consumers' commercial interests, there are private law remedies. When a public entity's errors affect citizens' interests, there is typically no parallel remedy. Instead, the government can rely on what I call "clickwrap accountability," resting on disclaimers about chatbot performance and holding the individual responsible for assessing the tool's outputs—without guaranteeing any public or private law remedy.

The now-discontinued New York City "MyCity Chatbot (Beta)" is illustrative of clickwrap accountability and its pathologies.[6] After the tool

1. ALEXANDER POPE, AN ESSAY ON CRITICISM l. 525 (Project Gutenberg 2005) (1711).

2. *See* Kyle Melnick, *Air Canada Chatbot Promised a Discount. Now the Airline Has to Pay It.*, WASH. POST (Feb. 18, 2024), https://www.washingtonpost.com/travel/2024/02/18/air-canada-airline-chatbot-ruling/ (on file with the *Fordham Law Review*).

3. Moffatt v. Air Can., 2024 BCCRT 149, ¶ 27 (Can.).

4. *Id.* ¶¶ 24–32. The record did not include details about the support chatbot. *See id.* at ¶ 14. It also did not contain any discussion of whether the airline provided the plaintiff with any warning or disclaimer. *See id.*

5. *See infra* Part II.A.

6. Mayor Zohran K. Mamdani has ended the beta test of the chatbot. *See MyCity Chatbot Beta*, NYC, https://www.nyc.gov/main/error/chatbot-maintenance [https://perma.cc/MWV4-EP4N] (last visited Aug. 17, 2026); Colin Lecher, The Markup & Katie Honan, *Mamdani Targets 'Unusable' AI Chatbot for Termination*, THE CITY REP. (Feb. 5, 2026), https://www.thecity.nyc/2026/01/30/mamdani-unusable-ai-chatbot-budget/ [https://perma.cc/YPZ8-CBTB].

provided incorrect and even downright illegal advice to users,[7] the city's initial response was to add disclaimers to the site and to present a pop-up that required affirmative user consent before asking a question.[8] While operational, the disclaimer stated:

> As a beta product still being tested, [the MyCity chatbot] may occasionally provide incomplete or inaccurate responses. Verify answers with links provided after the response or by visiting MyCity Business and NYC 311. Do not use its responses as legal or professional advice nor provide sensitive information to the Chatbot. Your conversations are not connected to your MyCity account and will be permanently deleted after 30 days.[9]

This disclaimer was followed by a checkbox requiring a user to consent to the statement, "I understand the MyCity Chatbot's limitations,"[10] before engaging with the chatbot.

This disclaimer shifted responsibility to the user and, as Part II.A explains, public law will generally fail to provide remedies for incomplete or inaccurate chatbot responses. To be fair, the New York City example is both historic and extreme. But New York City is not alone in the use of disclaimers to regulate public-facing chatbots. Far from it. In myriad localities, states, and nations, governments are launching or adapting public-facing chatbots to incorporate generative AI.[11] As this Essay

---

7. *See* Colin Lecher & Katie Honan, *NYC's AI Chatbot Tells Businesses to Break the Law*, THE MARKUP (Mar. 29, 2024, at 06:00 UTC), https://themarkup.org/news/2024/03/29/nycs-ai-chatbot-tells-businesses-to-break-the-law [https://perma.cc/A6DQ-RTD4].

8. *See MyCity Chatbot Beta*, NYC, https://chat.nyc.gov/ [https://perma.cc/H2UR-6EZC] (last visited Aug. 4, 2026).

9. *Id.*

10. *Id.*

11. *See GenAI Chatbot*, CAN. REVENUE AGENCY (Mar. 20, 2025), https://www.canada.ca/en/revenue-agency/corporate/contact-information/cra-chatbot.html?hsid=a89bd2d0-8260-4bc2-9773-9cbb8a453a66 [https://perma.cc/9U5R-5R45]; Leanne Cummings, *GOV.UK Has Entered the Chat: Our Vision for GOV.UK Chat*, GOV.UK (Dec. 16, 2025), https://insidegovuk.blog.gov.uk/2025/12/16/gov-uk-has-entered-the-chat-our-vision-for-gov-uk-chat/ [https://perma.cc/3U2U-KGE7]; Mykhailo Fedorov, *Diia.AI: The World's First National AI-Agent That Delivers Real Government Services*, DIGIT. STATE UA (Sep. 3, 2025), https://digitalstate.gov.ua/news/govtech/diiaai-pershyy-u-sviti-derzavnyy-ai-ahent-iakyy-ne-prosto-konsultuye-a-nadaye-posluhy-iak-pratsiuye-shtuchnyy-intelekt-na-portali [https://perma.cc/9RD3-7WVS]; Madelyn Hanes, *Meet 'Ask Indiana,' State Government's New AI Chatbot*, IND. CAP. CHRON. (May 15, 2025), https://indianacapitalchronicle.com/2025/05/15/meet-ask-indiana-state-governments-new-ai-resident-assistant-chatbot/ [https://perma.cc/P7WR-4DUL]; *Ask Missi*, MS.GOV, https://www.ms.gov/ [https://perma.cc/BA62-ESH7] (last visited Aug. 8, 2026); *MC311 Contact Center*, MONTGOMERY CNTY., MD., https://www.montgomerycountymd.gov/MC311/ [https://perma.cc/6JQU-S9ND] (last visited July 29, 2026); Ashley Silver, *South Carolina AI, Cyber Initiatives Reshape Services*, GOV'T TECH. (Sep. 24, 2025), https://www.govtech.com/artificial-intelligence/south-carolina-ai-cyber-initiatives-reshape-services [https://perma.cc/4ABS-67CW]; *ChatGPT*, U.AE (Oct. 20, 2025), https://u.ae/en/about-the-uae/digital-uae/digital-technology/artificial-intelligence/chatgpt [https://perma.cc/54XG-WNJT]; *Virtual Assistant Bürokratt*, AI EMPOWERING EST., https://www.kratid.ee/en/homevana [https://perma.cc/E6V7-6ZBA] (last visited July 29, 2026). For a federal example that appears to be a generative AI system based on self-descriptions or terms of service provisions, see James Bergeron, *What's New for the 2026–27 FAFSA® Form?: Beta Launch, Smarter Virtual Assistant, and More*,

discusses below, clickwrap accountability is proliferating across these settings.

The problem is that such disclaimers are already known to work poorly in consumer law, and it is a mistake to rely on them in public law settings that feature generative AI. Specifically, this approach discounts the known pathologies of notice and choice disclaimers in privacy law, and of clickwrap and browsewrap in contract law more generally. It also ignores the fact that private law remedies, such as tort law or contract law claims that might be available in the commercial context, are generally unavailable when public actors rely on notice-and-choice-style disclaimers. And, in so doing, this approach disregards sociotechnical dynamics. More precisely, clickwrap accountability does not account for the ways that people are affected by public chatbots, how these interactions affect rights and entitlements, and how the government's deployment of public-facing generative AI chatbots might affect "public service values."[12]

To be sure, government chatbots are not axiomatically bad. Public-facing chatbots may improve public service delivery as compared to conventional methods, such as relying on static FAQ pages or overburdened, underpaid bureaucrats. These tools may offer new, needed pathways to citizen services.[13] They may be a better option than, say, taking time off work and waiting at a crowded Social Security Administration office for an appointment or waiting on hold for hours to speak with a caseworker. They may also afford vital access for an underresourced state. However, these tools can serve the public interest only if law provides adequate guardrails for when they do not perform well,

---

DEP'T. EDUC. HOMEROOM BLOG (June 25, 2025), https://www.ed.gov/about/homeroom-blog/whats-new-2026-27-fafsar-form-beta-launch-smarter-virtual-assistant-and-more [https://perma.cc/N65W-VXZM] (describing expansion of the tool to include generative AI technology).

12. *See, e.g.*, Theodoros Papadopoulos, Charalampos Alexopoulos & Yannis Charalabidis, *Evaluating Chatbot Architectures for Public Service Delivery: Balancing Functionality, Safety, Ethics, and Adaptability*, FRONTIERS POL. SCI., Oct. 2025, at 1, 2; Tendai Makasi, Alireza Nili, Kevin C. Desouza & Mary Tate, *Chatbot-Mediated Public Service Delivery: A Public Service Value-Based Framework*, FIRST MONDAY, Dec. 2020, at 1, 2, https://firstmonday.org/ojs/index.php/fm/article/view/10598 [https://perma.cc/8TP9-W5WH]. *See generally* Gregory Porumbescu, Matthew M. Young & Kaifeng Yang, *Introduction to Special Issue on Generative AI and Public Administration: Performance, Accountability, and Democratic Governance*, 48 PUB. PERFORMANCE & MGMT. REV. 1263, 1263–64 (2025) (assessing generative AI systems and public administration in general); Qianli Yuan & Tzuhao Chen, *Holding AI-Based Systems Accountable in the Public Sector: A Systematic Review*, 48 PUB. PERFORMANCE & MGMT. REV. 1389, 1390 (2025) (discussing public sector AI in general).

13. *See, e.g.*, Sam Dub & Sharon McDonald, *5 Things We Learned Testing GOV.UK Chat: An AI Assistant for Government*, INSIDEGOV.UK BLOG (Mar. 16, 2026), https://insidegovuk.blog.gov.uk/2026/03/16/5-things-we-learned-testing-gov-uk-chat-an-ai-assistant-for-government/ [https://perma.cc/L4ZN-YZVV] (reporting on the United Kingdom's public pilot of its AI-powered app and summarizing positive user feedback); Mark Chillingworth, *GOV UK AI Chat Completes Public Pilot*, UKAUTHORITY (Mar. 25, 2026), https://www.ukauthority.com/articles/gov-uk-ai-chat-completes-public-pilot [https://perma.cc/42U6-K7H7] (same).

and only if legal rules reflect how people engage with chatbots.[14] And right now, they do not.

The underlying problem is structural: Law does not recognize or provide adequate redress for government errors delivered through chatbot systems, and the backdrop of contract law disclaimers is normatively and pragmatically inadequate.[15] Moreover, although scholars have long studied and critiqued privatization, legal scholarship has not yet reckoned with an emerging pattern that chatbots put in vivid relief: Public law authorities are importing private law frameworks that are already known to be problematic.[16] These contemporary legal arrangements will produce only a façade of accountability—and that is not good enough.

This Essay proceeds as follows. Part I first situates public-sector chatbots in existing legal scholarship and constitutional doctrines and then assesses real-world examples of public-sector generative AI chatbots to illustrate how user interactions with them functionally blur doctrinal lines. Part II then argues that the main contemporary source of accountability is the private law intervention of clickwrap-style disclaimers and critiques this form of accountability. Part III considers interventions to support government use of chatbots in publicly accountable, rights-respecting ways.

I. DUE PROCESS AND THE GENERATIVE STATE

This part situates public-facing chatbots in the legal literature and in emerging research on human interactions with AI systems. This survey

14. *See* David Freeman Engstrom & Amit Haim, *Regulating Government AI and the Challenge of Sociotechnical Design*, 19 ANN. REV. L. & SOC. SCI. 277, 287 (2023) (calling for more scholarship on "second-level" algorithmic design, which "focuses on bureaucratic-organizational dynamics and, in particular, the human-machine interface at the point of decision"). My argument is similar in spirit but distinct because Professors David Freeman Engstrom and Amit Haim emphasize the bureaucrat and the question of how to "guid[e] frontline discretion," *id.* at 288, whereas I focus on the interface between citizens and the system, even without a conventional government determination.

15. A wealth of prior private law scholarship critically evaluates what Professor Margaret Radin calls "regulation by contract." *See generally* Margaret Jane Radin, *Regulation by Contract, Regulation by Machine*, 160 J. INSTITUTIONAL & THEORETICAL ECON. 142 (2004). A further body of work analyzes problems with clickwrap agreements, both generally and in the notice and choice form that dominates American information privacy law. For discussion of this issue in the context of public-sector chatbots, see *infra* Part III.

16. It is rare for scholars to approach the question as one of structural considerations and public values, rather than focusing on a particular right or affected domain. For a recent exception, see Hannah Bloch-Wehba, How Tech Took Over 52 (July 20, 2026) (unpublished manuscript), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6378220 (on file with the *Fordham Law Review*) ("Hybridization of public and private governance . . . risks expanding reservoirs of private power that constitutional doctrine frequently ignores." (citations omitted)); *id.* at 49 ("[C]urrent arrangements are in tension with structural constitutional principles and values of democratic accountability."). Whereas Professor Hannah Bloch-Wehba emphasizes the historic economic and political roots of tech's "entrenchment" in public governance, *id.* at 28, I begin with the status quo and explore how private law methods and conceptions of accountability are shaping public law institutions today, focusing on the specific context of public-facing generative AI chatbots.

reveals how contemporary public-facing government chatbots both present distinct questions about when process is due and raise distinct sociotechnical concerns as compared to prior forms of automated legal guidance.

Generative AI is new, but government automation is not.[17] Before generative AI, there was predictive AI in the form of both machine-learning models and hard-coded expert systems that follow "if-then" statements. These systems vary a great deal, but many share a common feature: They help the government to make decisions. More specifically, governments use them to make investigative or enforcement decisions, to adjudicate, and to punish.[18] These decisions have a visceral, real-world impact. Agencies have used algorithmic tools to make or inform choices about people in high-stakes domains, including welfare benefits determinations,[19] child welfare decisions,[20] public employee pay and retention,[21] and criminal law risk assessment.[22]

Legal scholars have long expressed concern about automation and how it threatens procedural due process values. Over a decade ago, Professor Danielle K. Citron's foundational work on "technological due process" surveyed expert systems. Citron contended that government automation collapses the traditional distinction between rulemaking and individualized adjudication and "risks dismantling critical procedural safeguards at the foundation of administrative law."[23] A rich literature has since expressed concerns about fairness, accountability, and transparency in automated tools

17. *See, e.g.*, David Freeman Engstrom, *The Automated State: A Realist View*, 92 GEO. WASH. L. REV. 1437, 1447–51 (2025); Amit Haim, *The Administrative State and Artificial Intelligence: Toward an Internal Law of Administrative Algorithms*, 14 U.C. IRVINE L. REV. 103, 110–13 (2024); David Freeman Engstrom & Daniel E. Ho, *Artificially Intelligent Government: A Review and Agenda*, *in* RESEARCH HANDBOOK ON BIG DATA LAW 58–60 (Roland Vogl ed., 2021). For discussion of the federal government's use of AI tools before the rise of generative AI and "foundation models," see Joshua D. Blank & Leigh Osofsky, *Automated Agencies*, 107 MINN. L. REV. 2115, 2126–32 (2023). *See* DAVID FREEMAN ENGSTROM, DANIEL E. HO, CATHERINE M. SHARKEY & MARIANO-FLORENTINO CUÉLLAR, GOVERNMENT BY ALGORITHM: ARTIFICIAL INTELLIGENCE IN FEDERAL ADMINISTRATIVE AGENCIES 15–20 (2020) (on file with the *Fordham Law Review*) (report to the Administrative Conference of the United States).

18. *See* Aziz Z. Huq, *Constitutional Rights in the Machine Learning State*, 105 CORNELL L. REV. 1875, 1890–1905 (2020).

19. *See, e.g.*, Michael T. v. Bowling, No. 15-cv-09655, 2016 WL 4870284 (S.D. W. Va. Sep. 13, 2016), *modified sub nom.* Michael T. v. Crouch, No. 15-cv-09655, 2018 WL 1513295 (S.D. W. Va. Mar. 26, 2018).

20. *See, e.g.*, *Allegheny Family Screening Tool*, ALLEGHENY CNTY., https://www.alleghenycounty.us/Services/Human-Services-DHS/DHS-News-and-Events/Accomplishments-and-Innovations/Allegheny-Family-Screening-Tool [https://perma.cc/Z7EA-6JVP] (last visited July 29, 2026).

21. *See, e.g.*, Hou. Fed'n of Tchrs. v. Hou. Indep. Sch. Dist., 251 F. Supp. 3d 1168, 1171–74 (S.D. Tex. 2017).

22. *See, e.g.*, State v. Loomis, 881 N.W.2d 749, 753–57 (Wis. 2016).

23. Danielle Keats Citron, *Technological Due Process*, 85 WASH. U. L. REV. 1249, 1253 (2008).

in both public and private settings.[24] These accounts generally focus on what process is due when the state automates, whether the state has satisfied the letter and the spirit of the *Mathews v. Eldridge*[25] balancing test,[26] and what interventions might be necessary if not.[27]

The common thread in these accounts is attention to government decisions. That makes good sense. An individualized adjudication about a specific government program is the action that triggers procedural due process. In doctrinal terms, the *Board of Regents of State Colleges v. Roth*[28] test is satisfied because it is self-evident, in such AI use cases, that there has been a "deprivation of interests encompassed by the Fourteenth Amendment's protection of liberty and property."[29] The legal and normative challenge is thus whether the process provided when the state automates is constitutionally adequate, as well as whether automated decision-making systems undercut substantive values such as public legitimacy and democratic self-governance.[30]

Public law generally emphasizes the endpoint of government automation: the decision. But generative AI requires attention earlier in the process. Public-facing chatbots are distinct from, say, a machine-learning tool that's used by a government official to inform an eligibility determination for public benefits. Chatbots instead engage with people and offer advice without any formal decision about a constitutionally protected entitlement. The proper analogy is to an employee at an administrative agency's office who provides guidance, or to the agency's informational materials. But here, the advice is provided through a different medium: interactions with the chatbot interface. The medium matters.[31] These chatbots combine reliance and nonbinding guidance in distinctive ways, and they operate at

---

24. The literature is voluminous, with an entire conference dedicated to these topics. *See* ACM CONFERENCE ON FAIRNESS, ACCOUNTABILITY, AND TRANSPARENCY (FACCT), https://facctconference.org/ [https://perma.cc/5Y9N-QWN3] (last visited July 27, 2026).

25. 424 U.S. 319 (1976).

26. *Mathews* set out a balancing test for due process that:

> generally requires consideration of three distinct factors: First, the private interest that will be affected by the official action; second the risk of an erroneous deprivation of such interest through the procedures used, and the probable value, if any, of additional or substitute procedural safeguards; and finally, the government's interest, including the function involved and the fiscal and administrative burdens that the additional or substitute procedural requirement would entail.

Id. at 335.

27. *See, e.g.*, Citron, *supra* note 23, at 1305.

28. 408 U.S. 564 (1972).

29. *See id*. at 569 ("The requirements of procedural due process apply only to the deprivation of interests encompassed by the Fourteenth Amendment's protection of liberty and property.").

30. *See generally, e.g.*, Bloch-Wehba, *supra* note 16. *See also* Alicia Solow-Niederman, *Algorithmic Grey Holes*, 5 J.L. & INNOVATION 116, 122–23 (2023); Ryan Calo & Danielle Keats Citron, *The Automated Administrative State: A Crisis of Legitimacy*, 70 EMORY L.J. 797, 804–10 (2021); Ari Ezra Waldman, *Power, Process, and Automated Decision-Making*, 88 FORDHAM L. REV. 613, 614–16 (2019).

31. *Cf.* MARSHALL MCLUHAN, UNDERSTANDING MEDIA: THE EXTENSIONS OF MAN 7 (MIT Press ed., 1994) (1964) ("[T]he medium is the message.").

scale—and this combination raises especially pressing questions about when process is due.

Now, these issues are not entirely new; indeed, prior work has focused on automated government advice before the emergence of generative AI. Notably, Professors Joshua D. Blank and Leigh Osofsky detail how automated legal advice risks oversimplifying complex law and often appears more authoritative than it is.[32] Blank and Osofsky recognize how automated legal guidance can be functionally "quasi-binding" for individuals who receive it—without the same legitimizing procedures as legislative rules—and highlight how this mode of citizen engagement risks inequitable outcomes.[33]

Even so, few scholars to date have recognized how generative AI chatbots that provide government advice are distinct.[34] For example, Blank and Osofsky contend that evidence of user reliance should affect agency obligations and distinguish between "unilateral guidance, that which does not depend upon users' responses to questions," and "'bilateral' administrative guidance—where users input information regarding their own personal circumstances in exchange for the output of advice from automated tools."[35] Such analysis of users' reliance interests is essential. But here, the sociotechnical details matter. Because generative AI guidance inherently responds to user-specific inputs and is almost certainly "bilateral," these systems perform differently from prior versions of automated legal guidance.[36] At a minimum, the technology is a force multiplier of long-standing, latent tensions and challenges. And there is reason to believe this is a difference in kind, and not merely degree, given the sociotechnical dynamics of generative AI interactions.

To illustrate the thorny sociotechnical and legal issues that chatbots raise, consider two public-facing generative AI systems: the state of South Carolina's "AI Resident Assistant," named "Bradley,"[37] and New York

32. *See generally* JOSHUA D. BLANK & LEIGH OSOFSKY, AUTOMATED AGENCIES: THE TRANSFORMATION OF GOVERNMENT GUIDANCE (2025); Blank & Osofsky, *supra* note 17; Joshua D. Blank & Leigh Osofsky, *Automated Legal Guidance*, 106 CORNELL L. REV. 179 (2020) [hereinafter Blank & Osofsky, *Automated Legal Guidance*]; JOSHUA D. BLANK & LEIGH OSOFSKY, AUTOMATED LEGAL GUIDANCE AT FEDERAL AGENCIES (Mar. 25, 2022) (draft report to the Administrative Conference of the United States) [hereinafter BLANK & OSOFSKY, AUTOMATED LEGAL GUIDANCE AT FEDERAL AGENCIES].

33. *See* sources cited *supra* note 32.

34. For the only exception of which I am aware, see Daniel E. Ho, Olivia H. Martin, Amy Perez & Kit Rodolfa, *Evaluation as Due Process: Civil Service in an Automated Age*, 77 ADMIN. L. REV. 831, 880 (2025) (asserting that "chatbots offer guidance that occupies an uncertain legal space" and distinguishing generative AI chat errors from the "nuanced issue of simplexity").

35. BLANK & OSOFSKY, AUTOMATED LEGAL GUIDANCE AT FEDERAL AGENCIES, *supra* note 32, at 38–39.

36. *See infra* Part II. Thus, even well-known tools, such as the Free Application for Federal Student Aid (FAFSA) guidance chatbot, Aidan, warrant fresh analysis if they are expanded or adjusted to include generative AI. On prior, nongenerative AI versions of Aidan, see, for example, Blank & Osofsky, *supra* note 17, at 2129–30.

37. *South Carolina*, SC.GOV, https://sc.gov/ [https://perma.cc/9PKP-K468] (last visited Aug. 23, 2026).

City's "MyCity Chatbot Beta," before its termination.[38] Suppose that one navigates to the government's website, agrees to the terms of conduct (more on this below), and asks each system about eligibility for Supplemental Nutrition Assistance Program (SNAP) benefits.

For Bradley, as of January 2026, one would first be shown a cartoon image of a cute dog, followed by this introductory statement:[39]

> Hi there! I'm Bradley, a friendly Boykin Spaniel, and your guide to South Carolina's official website. As the state dog, I represent the loyalty and energy South Carolina is known for! I'm here to help you sniff out the information you need, whether it's government services, local resources or something else. Just give me a bark—or rather, type your question—and I'll fetch the answers for you![40]

Then, in response to the query, "am I eligible for SNAP benefits," Bradley would generate a box with the following text:

> To find out if you are eligible for SNAP (Supplemental Nutrition Assistance Program) benefits in South Carolina, your household's gross monthly income generally must not exceed 130% of the federal poverty level. For example, for a household of 1, the gross monthly income limit is $1,696; for a household of 4, it is $3,483. There are also net income and resource limits, and some people (like certain students or those with certain convictions) may not be eligible. The best way to check your eligibility is to use the "Am I Eligible" tool on the South Carolina Department of Social Services (DSS) Benefits Portal. This tool will ask you questions about your household and income to help you determine if you might qualify.[41]

The text concluded with links to resources to "access the eligibility tool and learn more about SNAP requirements" and informed the user that they could contact their local Department of Social Services (DSS) office for more assistance.[42]

For the MyCity Chatbot, as of January 2026, one would receive the following response:

> Eligibility for SNAP benefits is determined based on household size, income, and household expenses. You are not required to screen for eligibility before applying for SNAP. If you receive SSI/Disability benefits, you can still qualify for SNAP, and receiving SNAP does not reduce your SSI/Disability payments. Caregivers who are financially responsible for children in their case may also be eligible for SNAP. Additionally, Able-Bodied Adults Without Dependents (ABAWD) work requirements are waived until February 28, 2026. Most able-bodied

38. *MyCity Chatbot Beta*, *supra* note 8.

39. Screenshots of this interaction, as well as those discussed next, are available at the Perma links in the citations that follow.

40. South Carolina, SC.GOV, https://sc.gov/ [https://perma.cc/CY2V-S47Y] (last visited Aug. 4, 2026).

41. Alicia Solow-Niederman, Bradley, "Am I eligible for snap benefits?" (Jan. 25, 2026) [https://perma.cc/CY2V-S47Y].

42. *Id.*

> students aged 18 through 49 who are enrolled at least half-time in college or other institutions of higher education are generally not eligible for SNAP, but exceptions may apply for students that meet other eligibility criteria. To check your eligibility, you can use the ACCESS NYC eligibility screener or download a list of eligibility requirements and necessary documents.[43]

The text would conclude by warning, "[t]his list may not be exhaustive" and directing the user to refer to nyc.gov for further information. It then listed references before repeating a disclaimer about the accuracy of the tool.[44]

Summed up, these responses provide an enumerated list of considerations and direct the user to further resources. From the perspective of the black letter law, it is not an adjudication: Because these responses are nonfinal, nonbinding advice, they are neither an official eligibility determination, nor a formal decision at all.[45] But from the perspective of the user, the lived experience of interacting with this chatbot is distinct. Consider the fact that to fit the constraints of law review publishing and the physical medium of a printed law review, the best option was to type up the text of chatbot conversations in the manner recounted above and provide the screenshots separately via Perma links.[46] There was no way to include the chatbot interaction in this piece, nor even a way to easily capture the chats themselves as static photographs. That reproduction conveys the message—but not the medium. Chatbot guidance is not like a static FAQ, pamphlet, or brochure. Rather, it is tailored to reflect the specific query that the user inputs. It is thus personalized. It can also be far more individualized than analog guidance if the government designs the chatbot to allow users to engage in ongoing conversations that add their own personal facts. For instance, consider the following queries that I posed to Indiana's "Ask Indiana" chatbot[47] and Washington, D.C.'s "AI Chatbot Pilot," respectively, in January 2026.[48] In each case, I began with the same general question about SNAP eligibility: "[A]m I eligible for snap benefits?"[49]

---

43. Alicia Solow-Niederman, MyCity Chatbot Beta, "Am I eligible for snap benefits?" (Jan. 25, 2026) [https://perma.cc/85ZX-6X49].

44. *Id.*

45. *See* Adjudication, ADMIN. CONF. U.S., https://www.acus.gov/page/adjudication [https://perma.cc/XVL8-3HRD] (last visited July 29, 2026) (explaining the broad definition of "adjudication" under the Administrative Procedure Act and the narrower, common understanding of the term).

46. This is in no way to blame the editors. I am deeply appreciative to Emma Miloglav and the editors of the *Fordham Law Review* for their extraordinary efforts to work with me on inclusion of this nontraditional medium.

47. *Ask Indiana*, IN.GOV, https://www.in.gov/core/ai-terms-and-conditions.html [https://perma.cc/UA3N-94HB] (last visited Aug. 2, 2026).

48. *AI Pilot*, DC.GOV (last visited Jan. 25, 2026). As of April 2026, the chatbot had been removed from the DC.gov website, and use of the tool is discontinued as of this writing.

49. Alicia Solow-Niederman, D.C. AI Chatbot Pilot, "am I eligible for snap benefits?" (Jan. 25, 2026) [https://perma.cc/6FVQ-V6J9] [hereinafter Solow-Niederman, D.C. AI

The Ask Indiana chatbot responded:

> Eligibility for SNAP (Supplemental Nutrition Assistance Program) benefits in Indiana depends on several factors:
>
> - **Residency:** You must be an Indiana resident.
> - **Income:** Your household income must be below certain limits. For example, for a household of 1, the gross monthly income limit is $1,696; for 2, it's $2,292; for 3, it's $2,888, and so on. Both gross and net income limits apply.
> - **Assets:** Most households must have less than $5,000 in countable resources (like bank accounts, cash, and some vehicles).
> - **Other Requirements:** You must meet citizenship/alien status, work registration, and cooperate with job training programs if required.
>
> Special rules apply for students, elderly, and disabled individuals. If you think you might qualify, you can apply online or at your local Division of Family Resources office.[50]

The interface then directed the user to a webpage to check eligibility and apply before providing citations.[51]

After receiving this response, I then asked follow-ups, such as, "What if I am in Indiana for work for 6 months?"[52] The Indiana Chatbot responded as follows and then provided a list of citations:

> If you are in Indiana for work for 6 months and want to know if you can get SNAP (food assistance) or CCDF (child care) benefits, here's what Indiana policy says:
>
> **SNAP (Food Assistance):**
>
> - To get SNAP in Indiana, you must be a resident of Indiana. Residency means you intend to live in Indiana, either permanently or temporarily—not just for vacation.
> - If you are in Indiana for work (not just visiting), and you intend to live here for those 6 months, you can be considered a resident for SNAP purposes.
> - Migrant and itinerant workers who move from state to state for employment, and homeless individuals, also meet the residency requirement and may receive assistance if otherwise eligible.
> - You must also meet other eligibility requirements, such as income and resource limits.

Chatbot Pilot]; Alicia Solow-Niederman, Ask Indiana, "am I eligible for snap benefits?" (Jan. 25, 2026) [https://perma.cc/L9NZ-KRMQ] [hereinafter Solow-Niederman, Ask Indiana].

50. Solow-Niederman, Ask Indiana, *supra* note 46.

51. *Id.*

52. Alicia Solow-Niederman, Ask Indiana, "what if I am in Indiana for work for 6 months?" (Jan. 25, 2026) [https://perma.cc/L9NZ-KRMQ].

> **CCDF (Child Care Assistance):**
>
> - There is no minimum length of time you must live in Indiana to get CCDF child care assistance.
> - You must provide proof of your Indiana address, but you can be new to the state and still apply.
> - If you move to another Indiana county, you can transfer your CCDF benefits.
>
> **Summary:** If you are in Indiana for work for 6 months and intend to live here during that time, you can apply for SNAP and CCDF benefits as a resident, as long as you meet all of the other eligibility requirements.[53]

D.C.'s AI Chatbot Pilot similarly offered an initial reply to my general inquiry and then provided customized responses to questions about particular aspects of the SNAP benefits policy. Initially, in response to the question "am I eligible for SNAP benefits?,"[54] the D.C. AI Chatbot Pilot stated: "To determine if you are eligible for SNAP benefits, consider the following criteria based on the information available" and then provided a detailed description of factors such as "household composition" and "income limits."[55] For example, the "household composition" paragraph states, "SNAP eligibility is determined by the household, which includes individuals who live together and buy and prepare food together. If you are under 22 and living with your parents, or have children under 22, their income will be considered with yours."[56]

Based on this response, I provided the following input: "I am 21 and have three kids."[57] The D.C. AI Chatbot Pilot responded: "As a 21-year-old with three children, you would be considered a separate SNAP household, assuming you live independently from your parents. Your eligibility for SNAP benefits will depend on your household's total income and whether it falls below the income limits for a household of four."[58] It then offered further guidance about potential deductions, the possible relevance of the Temporary Assistance for Needy Families (TANF) and Supplemental Security Income (SSI) programs, and instructions to visit the D.C. Department of Human Services (DHS) website or contact the call center for more information or apply.[59] In response, I wrote: "I live with my parents because they are older and I help to take care of them."[60] The D.C. AI Chatbot Pilot replied:

---

53. *Id.*
54. Solow-Niederman, D.C. AI Chatbot Pilot, *supra* note 49.
55. *Id.*
56. *Id.*
57. Alicia Solow-Niederman, D.C. AI Chatbot Pilot, "I am 21 and have three kids." (Jan. 25, 2026) [https://perma.cc/4GWC-XJSC].
58. *Id.*
59. *Id.*
60. Alicia Solow-Niederman, D.C. AI Chatbot Pilot, "I live with my parents because they are older and I help to take care of them." (Jan. 25, 2026) [https://perma.cc/JW7U-3JA9].

> As a 21-year-old living with your parents and helping take care of them, your eligibility for SNAP benefits will be determined based on the combined income of your entire household, which includes you, your three children, and your parents. Since you are under 22, you must be included in the same SNAP household as your parents, and their income will count towards the household's total income. The total household income must fall below the SNAP income limits for a household of this size to be eligible for benefits.[61]

The chatbot once again directed to the website or call center to receive further assistance or to apply.[62]

In each instance, the user can focus on specific facts that matter to them—such as temporary residence, in the case of the Indiana example, or age and family members, in the case of the D.C. example. They can then pose follow-up questions that emphasize those facts, aiming to clarify an ambiguity in the chatbot's initial response.

This duality—not individualized and binding in the legal sense, yet personalized and responsive to specific, individually salient facts—helps to explain both the appeal and the constitutional hazards of public-sector chatbots. Scholarship on past forms of automated legal guidance and human interactions with technological systems emphasizes the importance of caution in the face of such personalization. As Blank and Osofsky discuss, "automated legal guidance" can affect users differently than traditional written legal publications.[63] An important characteristic of prior forms of legal guidance, even before generative AI, is the capacity to "deliver[] personalized, rather than generic, information."[64] For example, Blank and Osofsky describe one taxpaying tool that "use[s] second-person pronouns (i.e., 'you' or 'yours')" and then "presents the taxpayer with an output that is seemingly personalized to the taxpayer, as it uses second-person pronouns."[65] Citing behavioral research, the authors contend that this personalized communication affects individuals' beliefs and actions far more than traditional, static, generic forms of guidance.[66]

Because generative AI chatbots can be personalized to an extent not previously seen in government guidance, these systems are likely to affect user behavior to an even greater degree than prior forms of automation. They can not only use second-person pronouns to refer to the user, as seen in the examples provided above, but also repeat the individualized facts that are provided by the user. Recall, for instance, how the D.C. AI Chatbot Pilot repeated back the facts that were posed in the query: "As a 21-year-old with three children, you would be considered a separate SNAP household." This form of engagement creates the impression that the user is actually participating in a conversation and receiving considered advice,

61. *Id.*
62. *Id.*
63. Blank & Osofsky, *Automated Legal Guidance*, *supra* note 32, at 218.
64. *Id.*
65. *Id.*
66. *Id.*

even though there is no person who actually understands on the other side—just pattern matching in response to the query. Such responses might well increase the user's trust in the advice they receive.

When the government deploys chatbots in this way, it creates an information architecture that mediates a citizen's relationship with the state. If the tool presents advice in an accessible, user-friendly interface; allows audio input (as the tools in D.C. did and as those in Indiana and South Carolina do, as of this writing); or is anthropomorphized with friendly names, personas, and images (like South Carolina's "Bradley"), it arguably creates the impression that it is a friendly, helpful resource that should be trusted.[67] If the tool responds to a specific user query and engages with specific facts that the user has provided, it arguably creates the impression that the tool is providing authoritative advice about those facts. Focusing too narrowly on the legal form—guidance, not adjudication—overlooks how interfaces like these can affect citizens' perceptions of these interactions. Interactions such as these risk creating the impression that there is a relational dynamic with the user, as well as the impression that the state is saying something about eligibility through the chatbot even though it is mere information, and even though no binding pronouncement or eligibility determination has been provided.

These kinds of generative AI interactions thus blur the line between the traditional hallmarks of government guidance and the traditional hallmarks of administrative adjudication.[68] This blurring of assumed doctrinal lines, catalyzed by social interactions with technology, is not unlike Citron's observation that expert systems blurred rulemaking and adjudication. But here, the impact is upstream, at the interface between the citizen and the system itself, before there is any formal proceeding. We might think of these chatbots (including the underlying model and the associated AI system and user interface) as "street-level algorithms" that translate formal regulatory pronouncements for citizens as they respond to citizens' questions.[69] Given this newly salient interaction, the question is not (just) what process is due. It is also *when* process is due, given how people engage with sociotechnical systems.

Furthermore, emerging evidence about how real people engage with chatbots, in general, suggests that caution is prudent. For one, individuals are apt to trust chatbots in distinct ways when they are presented as friendly and approachable (say, featuring an animal or persona with a memorable name) or human-like (say, including an assigned gender or a name that a

67. *See* discussion *infra* Part II.B.

68. *See* Ho et al., *supra* note 34, at 837.

69. Ali Alkhatib & Michael Bernstein, *Street-Level Algorithms: A Theory at the Gaps Between Policy and Decisions*, CHI 2019, May 2019, at 1, 2; *see also* Michael Veale & Irina Brass, *Administration by Algorithm?: Public Management Meets Public Sector Machine Learning*, *in* ALGORITHMIC REGULATION 128 (Karen Yeung & Martin Lodge eds., 2019) (analyzing "micro-level" algorithmic governance choices at the "street-level" to provide government services).

real person could have).[70] For another, as people increasingly turn to chatbots for extraordinarily intimate conversations,[71] and even develop relationships with chatbots,[72] they may not think twice about relying on a similar system for authoritative advice about the law. They may trust this advice as a final determination, even when it is decidedly not. And therein lies the rub: When these tools are easier to access and customize, they are arguably more useful for the citizen who seeks individualized assistance. However, that same ease of access and customization also heighten the risk that the engagement becomes conversational and personalized in ways that create trust and induce reliance.

Admittedly, some of these assertions are speculative because there is not a robust body of empirical work on the specific issue of how individuals interact with public-sector generative AI chatbots. But action is needed now, before there are more serious costs for government accountability. American government chatbots to date are informational—and even these uses raise questions about where to draw the line between guidance and adjudication. Moreover, government chatbots are already expanding at the federal level[73] and proliferating across the globe, including systems that promise agentic capabilities[74] or combine advice giving and government decisions in ways that will strain the line even more directly.[75] We should consider whether the legal structures in place are fit for these uses before expanding these tools to even higher-stakes, rights-affecting settings,

70. *See, e.g.*, Sandra Peter, Kai Riemer & Jevin D. West, *The Benefits and Dangers of Anthropomorphic Conversational Agents*, PNAS, June 2025, at 2, 5, https://www.pnas.org/doi/epdf/10.1073/pnas.2415898122 [https://perma.cc/N95K-7YNC]; *see also* Cindy M. Grimm, *The Danger of Anthropomorphic Language in Robotic AI Systems*, BROOKINGS (June 18, 2021), https://www.brookings.edu/articles/the-danger-of-anthropomorphic-language-in-robotic-ai-systems/ [https://perma.cc/RL5E-RANG] (discussing dangers of adopting anthropomorphic language with regard to robots).

71. *See* Gerrit De Vynck & Jeremy B. Merrill, *We Analyzed 47,000 ChatGPT Conversations. Here's What People Really Use It For.*, WASH. POST (Nov. 12, 2025), https://www.washingtonpost.com/technology/2025/11/12/how-people-use-chatgpt-data/ (on file with the *Fordham Law Review*).

72. *See* Rhiannon Williams, *It's Surprisingly Easy to Stumble into a Relationship with an AI Chatbot*, MIT TECH. REV. (Sep. 24, 2025), https://www.technologyreview.com/2025/09/24/1123915/relationship-ai-without-seeking-it/ [https://perma.cc/37ZN-V9ML].

73. *See, e.g.*, Ask Aidan®Beta, FED. STUDENT AID, https://studentaid.gov/h/apply-for-aid/fafsa [https://perma.cc/4QTZ-L4ES] (last visited Aug. 23, 2026) (offering a generative AI interface for FAFSA-related questions); Sarah Todd, *We Tested the Government's Official New AI Nutrition Tool: Grok*, STAT NEWS (Feb. 10, 2026), https://www.statnews.com/2026/02/10/new-food-pyramid-website-raises-ai-questions-grok-realfood-dot-gov [https://perma.cc/4G4U-6SSP] (discussing RealFood.gov's chatbot tool, which embeds xAI's Grok AI into the government website to provide public-facing nutrition advice).

74. *See, e.g.*, Chillingworth, *supra* note 13 (reporting that the GOV.UK team is "experimenting with agentic AI"); Cummings, *supra* note 11 ("We're experimenting with agentic AI, and the idea that GOV.UK Chat can evolve from providing answers to performing actions.").

75. *See, e.g.*, Fedorov, *supra* note 11 ("[Diia.AI] is the world's first national AI-agent that goes beyond answering questions—it actually provides government services directly within a chat.").

especially when—as the next part takes up—there are limited legal hooks for chatbot errors that harm real people's interests.

II. ACCOUNTABILITY GAPS: THE THINNESS OF CLICKWRAP

This part exposes how public-sector chatbot use rests predominantly on clickwrap accountability and contends that it is a mistake to do so. Part II.A traces why current constitutional law doctrines are unlikely to provide recognition of, or redress for, public law systems that give bad advice to individuals. Part II.B exposes how governments are instead resting accountability on clickwrap style interventions imported from private law. Part II.C critiques this move, drawing on lessons from information privacy law and consumer protection law.

A. Public Law Gaps

This part considers how generative AI chatbots intersect with procedural due process doctrine and reveals two fundamental mismatches between the law on the books and how chatbots operate on the ground.

The first mismatch involves the due process framework itself. As discussed above, chatbot advice giving is not likely to trigger procedural due process protections at all. There is no decision in the way that constitutional law requires. Absent a tech-specific reworking of the doctrine, process will not be due when the government deploys chatbots. Part III returns to potential interventions that would directly address this mismatch and considers their efficacy and costs.

I focus here on a second mismatch: Even if procedural due process protections are somehow triggered, it is still unlikely that redress will be available for public-sector chatbot errors. An initial hurdle involves the status of the chatbot user. Someone who asks a chatbot for information about government programs may be a potential applicant who is seeking to determine their eligibility for a benefits program, as opposed to someone who has already qualified for those benefits. The question is whether such an applicant has protected property interests equivalent to those of current benefits recipients; if not, then many public-sector chatbot users may be cut out of the due process framework.

Here, the doctrine is likely not an outright bar. As the U.S. Court of Appeals for the Second Circuit explained in *Kapps v. Wing*,[76] although "[t]he Supreme Court has repeatedly reserved decision on the question of whether *applicants* for benefits (in contradistinction to *current recipients* of benefits) possess a property interest protected by the Due Process Clause . . .[,] [e]very circuit to address the question" has found that applicants "may possess a property interest in the receipt of public welfare entitlements."[77] Thus, at least where the government program "mandates

76. 404 F.3d 105 (2d Cir. 2005).
77. *Id.* at 115 (citations omitted).

award of the benefit" when an applicant meets "specified criteria," at least "some measure of due process protection" will likely attach to the eligibility determination.[78] That said, this line leaves some ambiguity, both because there may be some government programs without specific criteria that sufficiently limit official discretion, and because the U.S. Supreme Court has repeatedly declined to resolve this question.[79]

A second hurdle, which involves the advice provided by the generative AI system itself, is more significant. As established in *Office of Personnel Management v. Richmond*[80] and *Schweiker v. Hansen*,[81] mere error will not provide grounds for equitable estoppel against the government in most cases (and especially if the plaintiff seeks monetary relief).[82] That doctrinal line has important implications for redress in cases involving chatbot-generated errors. To make this point more concrete, suppose a chatbot tells a person that they are not eligible for SNAP, but the chatbot's advice is incorrect. If the person relies on this information to their detriment, they typically will not have an estoppel claim against the government.[83] There are principled policy reasons for these lines.[84] And yet a fundamental point remains true: The doctrine on the books is likely to bar redress if a chatbot errs, even when that error affects the distribution of public benefits, and even when that error affects constitutional rights. These doctrinal lines leave accountability gaps in public law doctrine. As an accountability backstop, the government is relying on disclaimers and consumer-facing warnings. The next two parts identify and then critique this emerging development.

B. Private Law Moves

This part reveals how agencies deploy generative AI chatbots along with disclaimers that shift responsibility for errors from the government to chatbot users. This clickwrap accountability brings private law mechanisms into public law settings.

The specific form of clickwrap accountability varies. To illustrate, this Essay describes two emerging types. The first approach is pure notice. One example is Washington, D.C.'s warning about its now-discontinued AI

78. *Id.* at 116 (citation omitted).

79. *See id.* at 115 (citing Lyng v. Payne, 476 U.S. 926, 942 (1986); Walters v. Nat'l Ass'n of Radiation Survivors, 473 U.S. 305, 320 n.8 (1985); Gregory v. Town of Pittsfield, 470 U.S. 1018, 1018 (1985) (O'Connor, J., dissenting from denial of certiorari)). There is also a distinct set of sociotechnical questions, which I reserve for future work: If a decision is made with a foundation model, *whose* discretion must be constrained? Is a system that operates by predicting the most likely token in a sequence constrained in the manner that the law contemplates? What does this doctrinal analysis capture about broader legal values, and what does it elide?

80. 496 U.S. 414 (1990).

81. 450 U.S. 785 (1981) (per curiam).

82. *See Richmond*, 496 U.S. at 417; *Schweiker*, 450 U.S. at 790.

83. *See* Ho et al., *supra* note 34, at 878 (discussing *Schweiker* and *Richmond*).

84. *See id.* at 878–79 (discussing equitable estoppel doctrine).

Pilot[85]: "We're excited to begin testing generative artificial intelligence (AI) to help residents, businesses, and visitors navigate DC.gov information and services. Please understand that not all of the information may be accurate, but we hope by releasing AI into the wild, it can get smarter and smarter after every use."[86]

This warning may also be included in the interface itself, as in the case of Indiana's chatbot.[87] The small text at the very bottom of the landing page on the main Indiana government website, below the embedded chatbot interface, cautions: "Ask Indiana is powered by AI and can make mistakes. Check important info," and then includes a hyperlink to further "Terms and Conditions" that provide details about the system and the user's responsibilities.[88] These terms and conditions can include further hyperlinks; for instance, Indiana includes links to the "Microsoft Generative AI Services Code of Conduct" and other third-party terms and conditions.[89]

In some instances, there may be an initial warning but no disclaimer on the tool itself; for example, if one clicks through to open the "Ask Indiana" chatbot interface in its own window, there is no reproduction of the disclaimer. This interface includes only a link to the external terms and conditions page.[90] This pure notice approach is a form of "browsewrap," wherein "the user does not see the contract at all but in which the license terms provide that using a Web site constitutes agreement to a contract whether the user knows it or not."[91]

The second approach is notice and consent. Here, the government offers notice and pairs it with an affirmative consent requirement. As discussed above, New York City took this approach while "MyCity Chatbot (Beta)" was operational, accompanied by a disclaimer on the tool interface itself. South Carolina also provides a disclaimer and requires a visitor to the site to click "agree" before using the tool.[92] Specifically, when one visits the state's website, an icon for the "SC AI Resident Assistant," accompanied by a pawprint (to signify Bradley), appears in the lower right corner. The first time that a user clicks on this icon to expand the interface, the following message displays: "Welcome to the SC Artificial Intelligence (AI) Resident Assistant Beta release! Please read the SC AI Resident Assistant's Disclaimer and Terms and Conditions carefully before agreeing

85. *AI Pilot*, DC.GOV [https://perma.cc/4GXX-95KN] (last visited Aug. 4, 2026).

86. *Id.*

87. *Welcome to Indiana*, IN.GOV, https://www.in.gov/core/index.html [https://perma.cc/AH5P-UMGR] (last visited July 29, 2026).

88. *Id.*

89. *AI Terms and Conditions*, IN.GOV, https://www.in.gov/core/ai-terms-and-conditions.html [https://perma.cc/NUL7-2PEK] (last visited July 29, 2026).

90. *Welcome to Indiana*, *supra* note 87.

91. Mark A. Lemley, *Terms of Use*, 91 MINN. L. REV. 459, 460 (2006) [hereinafter Lemley, *Terms of Use*]; Mark A. Lemley, *Protecting Consumers in a Post-Consent World*, 77 STAN. L. REV. ONLINE 247, 253 & n.22 (2025) [hereinafter Lemley, *Protecting Consumers in a Post-Consent World*] (citing Lemley, *Terms of Use*, *supra*, at 461–63, 469–70).

92. *South Carolina*, *supra* note 37.

to use this service."[93] An individual must scroll through the text to access these terms, which include the following "Use of Generative AI" statement: "SC AI Resident Assistant utilizes generative AI technology to provide information, answer queries and generate content. While the SC AI Resident Assistant strives to offer accurate and helpful information, the AI-generated content may not always be correct, up-to-date or applicable to your specific situation."[94] At the end of the terms, the visitor is told, "Please specify if you agree to be bound by the SC AI Resident Assistant's Disclaimer and Terms and Conditions."[95]

This notice-and-consent approach is a form of "clickwrap," wherein "an online user clicks 'I agree' to standard form terms."[96] Here, and in general, common private law mechanisms are baked into the government interface itself.

Furthermore, there are various design options that affect citizens' interactions with the government's chatbot, especially for the notice-and-consent version of clickwrap accountability. For instance, as discussed above, South Carolina instructs visitors to scroll through a lengthy box to view the full disclaimer; however, the user interface allows a person to click "Agree" before scrolling through the full text. South Carolina also includes hyperlinks to the "Terms and Conditions" webpage, which are "incorporated by reference" at the end of the disclaimer.[97] Design choices about the interface thus mediate the citizen-government interaction and affect how an individual is informed of their responsibilities for the chatbot's performance.[98]

Abstracting away from specific design choices, a pattern emerges: Governments are adopting the contractual mechanisms of pure notice, notice and consent, and notice and consent plus external third-party terms to implement their chatbots. Because public law doctrine does not provide consistent accountability measures,[99] there is no other reliable accountability backstop for these tools.

C. Public-Private Problems

This part contends that such clickwrap accountability is flimsy at best.[100] Lessons from other areas of law counsel against its use in the public sector

93. *Id.*
94. *Id.*
95. *Id.*
96. Lemley, *Terms of Use*, *supra* note 91, at 459; Lemley, *Protecting Consumers in a Post-Consent World*, *supra* note 91, at 253 n.22.
97. South Carolina, *supra* note 37.
98. Privacy law scholars have long recognized how design affects core values and rights. *See generally* WOODROW HARTZOG, PRIVACY'S BLUEPRINT: THE BATTLE TO CONTROL THE DESIGN OF NEW TECHNOLOGIES (2018); ANN CAVOUKIAN, INFO. & PRIV. COMM'N OF ONT., PRIVACY BY DESIGN 1 (n.d.), https://www.ipc.on.ca/en/media/1826/download?attachment [https://perma.cc/X785-YHRV].
99. *See supra* Part II.A.
100. In addition to the limitations of this model of public accountability, there is in fact a double accountability problem: Many of these government chatbots are created by private

and underscore the need for caution in contexts that are especially apt to induce reliance.

Clickwrap accountability shifts public law questions into the register of private law disclaimers, with major implications for the relationship between the government and the governed. As a matter of first principles, one might object to the use of private law interventions to resolve public law issues. One might find it normatively problematic to treat individuals as consumers rather than as citizens, and to think about a person who seeks information about their benefits eligibility as a "user," akin to someone who wants to post on social media or purchase a product online. In addition, as a matter of good governance, one might find it troubling to require an individual to monitor the terms of their own engagement with government, rather than guaranteeing baseline protections for all people within a system of constitutional government. Accordingly, one might question the use of contract law to disclaim public responsibility for design choices that affect citizens' rights and protected interests.[101]

But even assuming that one accepts the use of private law mechanisms in the public law context, relying on disclaimers is problematic on its own terms. The issue is simple: Clickwrap accountability disregards long-standing critiques of the private law methods on which it relies. Notably, privacy law scholars have long questioned the efficacy of notice and consent (sometimes called notice and choice) as a way to protect consumers' information privacy interests.[102] In privacy law, this approach aims to "provide a consumer with control over their own data, by giving that person a statement of the company's privacy practices ('notice') and obtaining individual consent to those practices ('choice') in exchange for access to a good or service."[103]

When a government agency purports to provide a citizen/consumer/user with control over their queries by warning that chatbots "can make mistakes," it is taking a page from this private law playbook. In the pure notice approach, that disclaimer is all that is provided. In the notice-and-consent approach, government agencies obtain individual consent to the disclaimer in exchange for access to the chatbot services. Across both approaches, the state imports the underlying logic of notice and choice in privacy law, relying on individual consent to legitimize its conduct.

---

vendors, such that their design and operation are not directly controlled by public actors. To focus on the pathologies of clickwrap accountability itself, this Essay reserves further consideration of this point.

101. *See* Andrew Keane Woods, *The New Social Contracts*, 77 VAND. L. REV. 1831, 1852–54 (2024).

102. *See* Joel R. Reidenberg, N. Cameron Russell, Alexander J. Callen, Sophia Qasir & Thomas B. Norton, *Privacy Harms and the Effectiveness of the Notice and Choice Framework*, 11 I/S: J.L. & POL'Y FOR INFO. SOC'Y 485, 490–95 (2015).

103. Alicia Solow-Niederman, *AI and Doctrinal Collapse*, 78 STAN. L. REV. 955, 975 (2026).

Relying on notice and consent in the administrative state amplifies long-standing concerns from consumer contexts. One concern is a lack of informed consent: Even if a person affirmatively agrees, does the average person understand what they are agreeing to with respect to generative AI, or do they just want a quick, thorough answer to their question?[104] The latter seems more likely. To be sure, the public is increasingly aware of hallucinations and one might argue that the "reasonable consumer" should know that chatbots can err. Yet, even accepting, arguendo, that a reasonable individual should know about chatbot errors, notice and consent remains a problematic governance mechanism for at least two reasons.

First, information asymmetries undermine the notice-and-consent structure. A person may know that chatbots in general can err, but they cannot assess the risk that a particular chatbot will err. The typical visitor to an agency's website will lack access to information about that specific public-sector chatbot. They will, moreover, lack access to error rates, training data, validation-testing results, or information about system design, ongoing evaluation, and other internal oversight measures. They will not be able to assess whether that specific chatbot is more or less reliable than a human caseworker, or whether it performs better or worse for certain types of queries. Consent is meaningless when, as here, a person cannot effectively evaluate the magnitude of the risk they are accepting.

Second, even if the visitor to the site had this information, they would likely suffer from information overload. Such overload has often been the result when individuals are confronted with pages and pages of complex, often technical disclosures and then throw up their hands.[105] Clickwrap accountability ignores both concerns—and this is all the more troubling because, as we have seen, chatbots are likely to induce reliance from their users.

* * *

Traditionally, law has drawn a line between the government as an information provider and as a decision-maker, imposing far fewer constraints on the former role. In a world where information is passive and

104. *Cf.* Daniel J. Solove, *Murky Consent: An Approach to the Fictions of Consent in Privacy Law*, 104 B.U. L. REV. 593, 614–15 (2024) (articulating challenges with consent in privacy law); Solon Barocas & Helen Nissenbaum, *On Notice: The Trouble with Notice and Consent*, 2009 PROC. OF THE ENGAGING DATA F.: THE FIRST INT'L FORUM ON THE APPLICATION AND MGMT. OF PERS. ELEC. INFO 1, 5 (same). *But see* Nikolas Guggenberger, *Consent as Friction*, 66 B.C. L. REV. 353, 359 (2025) (arguing that "U.S. policymakers and regulators can and should leverage consent into friction to undermine the economic viability of personalized advertising and similarly harmful surveillance-driven business models"); M. Ryan Calo, *Against Notice Skepticism in Privacy (and Elsewhere)*, 87 NOTRE DAME L. REV. 1027, 1027 (2012) (providing "an exploration of innovative new ways to deliver privacy notice").

105. *See, e.g.*, Calo, *supra* note 104, at 1054–59 (discussing information overload in privacy policies, then considering "visceral notice" as an alternative); Omri Ben-Shahar & Carl E. Schneider, *The Failure of Mandated Disclosure*, 159 U. PA. L. REV. 647, 687 (2011) (discussing information overload in the face of mandated disclosures, in general).

general, that line is far more defensible. But we do not live in that world. Today, public-sector chatbots are blurring information and adjudication, and public law doctrine does not provide redress when the chatbots get it wrong. The remaining accountability backstop consists of private law measures that are flawed and incapable of reckoning with the ways that people are likely to rely on chatbots. Law does not have a working framework to account for these dynamics. We need one before we accept technology as inevitable and before governments worldwide roll out chatbots. Otherwise, we will foreclose accountable, law-respecting uses of technology in ways that could improve citizen-government relations. The next part confronts that challenge.

III. ACCOUNTABILITY BY LAW AND BY DESIGN

This part proposes three complementary moves to shift responsibility in the generative state back to agencies and simultaneously provide carefully calibrated, limited redress for citizens who detrimentally rely on materially inaccurate advice provided by state-mediated generative AI chatbots.[106]

The most direct move involves constitutional law. This approach reconsiders when process is due for chatbot-mediated guidance and seeks to provide additional individual rights for government chatbot users. The idea is straightforward: If a new medium (a chatbot) shifts the nature of the message (government advice) to make it seem more akin to adjudication than guidance, then perhaps the legal rules that apply to the message should change. In other words, given the state's contemporary power to personalize answers, and the prospect of greater reliance on those answers, additional procedural guardrails may be advisable—even before there is "deprivation of interests encompassed by the Fourteenth Amendment's protection of liberty and property."[107]

Although such an intervention is worth considering, it would need to be carefully tailored, and it is inadvisable as a stand-alone solution. Significantly, there is a scoping problem. If such a doctrinal shift is limited to public-facing generative AI interfaces, then it sets out a technology-specific rule and entrenches it in constitutional law. This path threatens ossification and underinclusivity if later developments present similar challenges.[108] But if it is not so limited, then the doctrinal shift would fly in the face of stare decisis: It would represent a fundamental shift in the law, in the face of reliance interests, without evidence that the

106. *See* Margot E. Kaminski, *Binary Governance: Lessons from the GDPR's Approach to Algorithmic Accountability*, 92 S. CAL. L. REV. 1529, 1533 (2019) ("[G]overning algorithmic decision-making *should* include both individual rights and systemic approaches.").

107. Bd. of Regents of State Colls. v. Roth, 408 U.S. 564, 569 (1972).

108. *See generally* Rebecca Crootof & BJ Ard, *Structuring Techlaw*, 34 HARV. J.L. & TECH. 347 (2021) (canvassing pros and cons of tech-specific interventions).

contemporary rule is impractical or otherwise problematic across the board.[109]

Additionally, there is a policy problem. Applying constitutionally guaranteed procedural guardrails to a broader range of government actions, upstream of a fixed decision that involves a protected property or liberty interest, would itself come with serious governance costs. Part of the reason for the *Mathews v. Eldridge* balancing test is that the government cannot accord infinite process in every case, or else it would cease to function.[110] If more government actions, including advice giving, were subject to additional process, then that would massively increase the total load on the state. A well-intentioned measure to better protect citizens' rights, before an official determination, might instead strain the system past its breaking point.

To balance these concerns, it is worth considering more modest versions of this constitutional law intervention that focus on specific design choices as a trigger for when additional process attaches to chatbot-provided guidance, even pre-determination. For instance, if a government entity deploys a generative AI chatbot that bears the agency's imprimatur, allows a person to input personalized facts, and includes a default interface that will answer questions about entitlement programs and respond to follow-up questions, then it may become distinct from conventional guidance in ways that require additional protections. Or, perhaps the state crosses this threshold if its chatbot includes features that are known to induce high levels of trust, such as a seeming ability to empathize with an individual's situation. For instance, suppose, hypothetically, that a public-sector chatbot responded to a user's query about benefits eligibility as follows: "Oh, taking care of that many family members with your income level sounds so hard. Don't worry, I can help you! Let's brainstorm options together. Would you like that?" Though more rigorous empirical analysis is needed, as a matter of common sense, this further personalization and simulated human empathy seem likely to induce even more trust in the AI system. Future work might consider limiting principles such as these, fusing social and technical research and design choices with legal guardrails.

An important corollary question is what process would suffice to protect citizens, particularly without a formal government decision, and whether additional process would counterproductively place more burdens on the applicant when the very idea of the generative AI system was to improve citizen access. These are serious concerns, and they further underscore why constitutional law tools, on their own, are unlikely to repair the holes in the clickwrap accountability framework.

Accordingly, a second complementary set of interventions is inspired by administrative law, rather than constitutional law, and seeks to implement

109. *See, e.g.*, Melissa Murray, *Stare Decisis and Remedy*, 73 DUKE L.J. 1501, 1507–09 (2024).

110. *See* Mathews v. Eldridge, 424 U.S. 319, 347–49 (1976).

system-wide protections to ensure appropriate process over time.[111] Scholars increasingly recognize that democratic accountability for AI systems is not a one-shot endeavor; rather, systemic accountability turns on sociotechnical design choices made throughout the life cycle of that system.[112] One of the key components of accountability is evaluation of the system, both before and after deployment. As Professor Daniel E. Ho, Olivia H. Martin, Amy Perez, and Kit Rodolfa contend, "Without such evaluation, agencies cannot satisfy *Mathews* balancing, cannot ensure systemic accuracy, and cannot avoid the constitutional violations that plagued" problematic historic examples of government automation.[113] The authors thus call for evaluation practices as due process, meaning that agencies must assess system performance through measures such as error rate testing, validation, and meaningful human oversight.[114] Quite right: Evaluation is a vital step to ensuring state accountability for the use of AI systems.

However, assessment alone cannot regulate how people interact with public-facing chatbots. It is critical to understand what evaluation can and cannot do. First, predeployment evaluation can help to assess chatbot performance;[115] however, it can only serve this function if the systems in question fall within the scope of the evaluation. If predeployment assessment is limited to systems that affect receipt and termination of benefits,[116] then many contemporary chatbot interactions will not be covered by procedural due process doctrine.[117] To be sure, one might accuse me of being too literal when "evaluation as due process" includes

111. *See* David Freeman Engstrom & Daniel E. Ho, *Algorithmic Accountability in the Administrative State*, 37 YALE J. ON REG. 800, 806 (2020).

112. *See generally, e.g.*, Engstrom & Haim, *supra* note 14; Haim, *supra* note 17; Mariano-Florentino Cuéllar & Aziz Z. Huq, *The Democratic Regulation of Artificial Intelligence*, KNIGHT INST. (Jan. 31, 2022), https://knightcolumbia.org/content/the-democratic-regulation-of-artificial-intelligence [https://perma.cc/F55H-PH8C]. There is also an expansive nonlegal literature on AI as a sociotechnical system and the technical and policy implications that follow from that understanding. *See, e.g.*, Andrew D. Selbst, danah boyd, Sorelle A. Friedler, Suresh Venkatasubramanian & Janet Vertesi, *Fairness and Abstraction in Sociotechnical Systems*, FAT* '19: PROCS. CONF. ON FAIRNESS, ACCOUNTABILITY, AND TRANSPARENCY, Jan. 2019, at 59, 63, https://dl.acm.org/doi/epdf/10.1145/3287560.3287598 [https://perma.cc/ZVR4-SVUH].

113. Ho et al., *supra* note 34, at 895.

114. *See id.*

115. *See id.* at 874–78, 891 (auditing generative AI chatbots that offer unemployment insurance advice and revealing that these chatbots do not always get the law right and do not always give sensible responses about basic facts).

116. *See id.* at 874–78 (focusing on systems that "meaningfully affect receipt and termination of benefits before deployment").

117. This issue will compound if more agencies adopt generative AI as a way to provide public-facing guidance. There are already indications that agencies are planning such expansions. *See also Treasury and Artificial Intelligence*, U.S. DEP'T OF THE TREASURY, https://home.treasury.gov/policy-issues/financial-markets-financial-institutions-and-fiscal-service/treasury-and-artificial-intelligence [https://perma.cc/P98Z-FMPS] (last visited July 29, 2026) (describing plans for "aspirational, not yet funded" TFI-1 "Public Chatbot" that "provides search results and summaries of public OFAC information . . . so that users of the public-facing OFAC website can find relevant information more easily").

other, more capacious legal hooks, such as federal and state procurement law,[118] and extends to other government conduct. I am in violent agreement. We can and should look to other bodies of law to support the thesis that evaluation is legally mandated, as a matter of long-standing public law tenets. And unless our evaluations account for the ways that citizens engage with generative AI systems—including for guidance—we will miss a critical part of the accountability picture.

Yet evaluation is still limited in a more fundamental sense: Although evaluation can reduce errors, it cannot eliminate them. Under the technological state of the art, it is not possible to eliminate the risk of "hallucination"[119] and other forms of error, such as relying on misinformation, can creep into AI outputs.[120] Even well-calibrated chatbots will always make mistakes. And when errors in generative AI systems inevitably occur, there is a further accountability problem that predeployment evaluation does not address. Under the *Schweiker* and *Richmond* line of estoppel doctrine, a person will generally be unable to obtain relief against the government for chatbot errors. It will look like there are protections for chatbot users in the form of procedural due process rights or in the form of performance evaluations. But in practice, redress for inevitable errors will remain unavailable. Now, one might object that this has always been the case, emphasizing that government officials have long offered advice and they have always made mistakes. Why should the doctrine be different simply because those mistakes are filtered through chatbots?

Fair enough—but those who take this stance need to confront the larger stakes. The lack of redress for public-sector AI systems deployed at scale presents a distinct issue: There is a structural problem when our system of law purports to offer procedural guardrails (such as procedural due process and predeployment evaluations), yet those protections are practically unavailable because there is no legal hook for relief for incorrect or

118. *See generally, e.g.*, Jessica Tillipman, *Buying Blind: Corruption Risk and the Erosion of Oversight in Federal AI Procurement*, 55 PUB. CONT. L.J. 175 (2026). *See, e.g.*, Ho et al., *supra* note 34, at 898–901; Cary Coglianese & Erik Lampmann, *Contracting for Algorithmic Accountability*, 6 ADMIN. L. REV. ACCORD 175, 181–84 (2021); Lavi M. Ben Dor & Cary Coglianese, *Procurement as AI Governance*, 2 IEEE TRANSACTIONS ON TECH. & SOC'Y 192, 194 (2021); *see also Artificial Intelligence and Procurement*, REGUL. REV. (June 27, 2022), https://www.theregreview.org/2022/06/27/series-artificial-intelligence-procurement/ [https://perma.cc/U37K-HAHG] (compiling essays that "explore the possibilities of using procurement standards and processes as a means of regulating AI use by government"); *cf.* Deirdre K. Mulligan & Kenneth A. Bamberger, *Procurement as Policy: Administrative Process for Machine Learning*, 34 BERKELEY TECH. L.J. 781, 777–78 (2019) (arguing that a "'policymaking' mindset," and not a "procurement mindset," should apply to "choices in the design, adoption, and use of machine learning systems").

119. *See* Nicola Jones, *AI Hallucinations Can't Be Stopped—but These Techniques Can Limit Their Damage*, NATURE (Jan. 21, 2025), https://www.nature.com/articles/d41586-025-00068-5 (on file with the *Fordham Law Review*).

120. *See* Jeremy Hsu, *AI Hallucinations Are Getting Worse–and They're Here to Stay*, NEW SCIENTIST (May 9, 2025), https://www.newscientist.com/article/2479545-ai-hallucinations-are-getting-worse-and-theyre-here-to-stay/ [https://perma.cc/QM5U-RYP3].

misleading chatbot advice. This result is a generative AI version of what I have previously called "algorithmic grey holes": There is a façade of redress without any substance.[121] Algorithmic grey holes such as these corrode the rule of law itself. There are also individual costs. As a practical matter, engagement with a public-sector chatbot can affect how a citizen perceives and receives government benefits. Advice can channel what public entitlements and benefits a person believes they have a right to in the first place. These effects matter, and a failure to engage with what is distinctive about citizen reactions to generative AI systems will overlook important accountability and legitimacy gaps.[122]

Accordingly, it is well worth considering additional front-end assessment requirements that expand when process is due beyond what procedural due process formally requires. Creators of generative AI systems might be held to systemic predeployment evaluation requirements, regardless of whether the decision is used for a formal determination. In any such evaluation, the accuracy of the advice is vital. Critically, though, accuracy is not the sole value that matters for a sociotechnical system that engages with citizens on the ground. Predeployment evaluations of chatbot interfaces should consider not only error rates but also human factors, such as how much a particular design induces reliance, how particular linguistic choices make a user perceive a system as more personalized, and the degree to which individuals have different perceptions or expectations of public-sector chatbots as compared to commercial chatbots.

To inform these interventions, there is a real need for empirical work on how citizens interact with these systems before deployment. We are missing critical data on whether disclaimers are effective accountability measures in the context of street-level algorithmic interactions. At a minimum, given evidence that these disclaimers do not work well in private law contexts, it seems prudent to adopt precautionary steps, such as limiting anthropomorphizing features in government chatbots and ensuring that disclaimers are large, accessible, and do not require scrolling through pages of a small window. The state might also adapt known lessons from consumer protection law, with an eye to the types of privacy statements that the Federal Trade Commission (FTC) has found deceptive, misleading, or otherwise unfair. At best, an agency should test for inappropriate reliance before deploying that system at scale, thereby building in additional accountability checks when the state chooses to rely on generative AI systems to inform its citizens.

These moves, and others like them, will bring us closer to system-wide accountability for government chatbots. But because better evaluation will

121. *See generally* Solow-Niederman, *supra* note 30.

122. There are also potential accountability and legitimacy costs on the other side, such as the concern that citizens seeking to "game" government systems might use generative AI systems to test out hypothetical situations or learn how to present their case for maximum personal gain. I reserve such issues for future work in order to focus on potential negative impacts on the citizen under the status quo. Thank you to Amit Haim for raising these questions.

not help individuals harmed by errors that inevitably slip through, and because evaluation seems unlikely to eliminate the problem of induced reliance, these system-wide protections will leave important limitations of public-sector generative AI systems unaddressed.

A complementary third move, which the remainder of this Essay proposes, is institutional: Take the private law example seriously and import private law remedies to bolster public accountability. In private law, one response to generative AI errors is to allow individuals to bring suit in tort law and contract law. Right now, a comparable cause of action for AI-generated advice does not exist. Even with procedural mechanisms nominally in place, estoppel doctrine and the lack of a formal decision affecting a constitutionally protected property or liberty interest generally mean that there is no legal remedy when government chatbots inevitably make mistakes. Unlike an individual who is allegedly given bad advice by a commercial chatbot, an individual whose rights are allegedly affected by bad government advice is at a loss. As a result, public actors are importing private law methods (such as clickwrap), while routing around private law redress.

Seen this way, part of the problem is remedial asymmetry, wherein legal lines block individual contestation and redress for government errors. As such, the generative AI state is missing the individual rights and contestation component that Professor Margot Kaminski, alone and with coauthors, has called for as a vital complement to systemic interventions.[123] Tailored regulatory interventions might level out this asymmetry, seeking to provide further process after an individual receives advice from a public-sector chatbot.

Granted, there are still obstacles. For one, Appropriations Clause concerns may arise at the merits stage. More specifically, if a plaintiff seeks monetary redress for the government's error, then the judiciary may be unwilling to issue an order that requires an outlay of government money without an express congressional mandate.[124] There may also be more general concerns about the government's monetary exposure or the strain on government resources if private litigants are allowed to sue the state for relief. For another, there are threshold justiciability barriers to private rights of action. One potential difficulty here is Article III standing: A plaintiff who has relied on chatbot advice may not be able to establish a concrete, nonspeculative injury in the manner that injury-in-fact requires.[125]

These concerns are valid, yet they are not insurmountable hurdles. First, the threshold standing requirement can help to filter out spurious claims. If a claimant can establish that they have already relied on bad advice—say,

123. *See, e.g.*, Margot E. Kaminski & Jennifer M. Urban, *The Right to Contest AI*, 121 COLUM. L. REV. 1957, 1988–2003 (2021) (assessing the case for an individual right to contest AI decisions).

124. This Appropriations Clause issue arose in *Richmond*. *See* Off. of Pers. Mgmt. v. Richmond, 496 U.S. 414, 424–31 (1990).

125. *See* TransUnion LLC v. Ramirez, 141 S. Ct. 2190, 2202–03 (2021).

an individual who asked for advice about SNAP eligibility and then, based on erroneous guidance, did not apply for benefits—then that injury is arguably nonspeculative, concrete, and particularized in the manner injury-in-fact demands.[126] If the claimant cannot make a showing of this type, then they will not be able to establish injury-in-fact under controlling doctrine.

Second, the Appropriations Clause concern can be addressed legislatively. As the *Richmond* Court explained, although "judicial use of the equitable doctrine of estoppel cannot grant respondent a money remedy that Congress has not authorized," Congress can change that status quo.[127] Historically, "where Congress has been concerned at the possibility of significant detrimental reliance on the erroneous advice of Government agents, it has provided appropriate legislative relief."[128] Here, there is comparable cause for concern, and Congress can address it. Specifically, Congress has the ability to create a limited statutory cause of action that permits a claimant to seek redress if that claimant can establish detrimental reliance on materially inaccurate government chatbot advice.[129] This approach is auspicious because it has a legislative precedent: Congress responded similarly after *Schweiker v. Hansen*, "exercis[ing] its power to expand recoveries for those who rely on mistaken advice."[130] Such a response is, moreover, principled and targeted because it requires the state to internalize the costs of its errors if it chooses to use a chatbot.

Third, and finally, policy concerns can be counterbalanced by tailoring the remedy. For instance, the availability of any private cause of action might be limited to instances where the state chooses to provide chatbot-mediated guidance with design features that are particularly apt to induce reliance, such as personalization, responsiveness, and a mode of communicating that creates the appearance of understanding and empathy. The right response to these concerns is to cabin the remedy to counteract negative consequences—not to foreclose it across the board.

Leveling out the remedial playing field is vital because the status quo is not working. Someone always pays when a chatbot errs. Right now, the cost falls on individuals who did not choose to integrate the tool into an administrative agency's operations, who lack information about the system, and who do not have equal bargaining power to the administrative state—with no reliable legal recourse. That is not an acceptable outcome, particularly when the government is defaulting to thin, clickwrap accountability that is known to be problematic in the private law setting, and particularly when the state is relying on generative AI for

126. *See id*. Moreover, because such an injury has a historic analog in common law causes of action like negligent misrepresentation or equitable estoppel, a pleading of this type should pass muster under *TransUnion*. *See id.* at 2203–04.

127. *Richmond*, 496 U.S. at 426.

128. *Id.* at 428.

129. *See* Blank & Osofsky, AUTOMATED LEGAL GUIDANCE AT FEDERAL AGENCIES, *supra* note 32, at 38–39.

130. *Richmond*, 496 U.S. at 428–29 (discussing Congress's response to *Schweiker*).

rights-impacting applications. Given this power imbalance, it is well worth calibrating appropriate individual remedies that can operate alongside systemic interventions.

Now is the time to confront the growing gap between what law assumes and how public chatbots interact with real people on the ground, by pairing carefully crafted individualized remedies of the sort traditionally available in private law with systemic evaluation afforded by public law.

CONCLUSION

As public-facing chatbots proliferate in government systems, it is vital to think about the government's relationship to the individuals it serves. There is no perfect solution, nor any intervention without a cost. But that does not mean that proceeding without any adaptation is the right move. Today, public-sector chatbots route around public accountability doctrines. Meanwhile, the backstop of clickwrap accountability, imported from private law, puts pressure on individuals and permits the government to disclaim responsibility. Yet individuals cannot evaluate how well the system works, and the structure of the system induces reliance. That bargain is a bad deal for a democratically accountable system of government. Rather than agree to it unthinkingly, this Essay exposes these increasingly pressing issues and proposes layers of legal interventions to better serve the government and its citizens in the generative AI state.