Title: Governing Conversational AI in Healthcare: An Executive-Controlled Architecture for Safe, Scalable HIV Prevention Counseling

Author list:
Jun Tao, PhD,[1] Ellie Pavlick, PhD,[2] Hannah Parent, MPH,[1] Natalie Fenn, PhD,[1,3] Mofan Gu, PhD,[1] Jack Rusley, MD, MHS,[5,6] Bryce Perler, MD, MSTR,[1] Tyler Wray, PhD,[7], Fan Yang, BS,[8], Philip A Chan, MD, MS[1,9]

[1] Department of Medicine, The Warren Alpert Medical School of Brown University, Providence, Rhode Island, USA
[2] Department of Computer Science, Brown University, Providence, Rhode Island, USA
[4] Department of Psychiatry and Human Behavior, The Warren Alpert Medical School of Brown University, Providence, Rhode Island, USA
[5] Department of Health Services, Policy, and Practice, Brown University School of Public Health, Providence, Rhode Island, USA
[6] Department of Pediatrics, Division of Adolescent Medicine, The Warren Alpert Medical School of Brown University, Providence, Rhode Island, USA
[7] Department of Behavioral and Social Sciences, Brown University School of Public Health, Providence, Rhode Island, USA
[8] Syntalon LLC, USA
[9] Open Door Health, Rhode Island Public Health Institute, Providence, Rhode Island, USA

Abstract

Large language models (LLMs) are increasingly being proposed as tools to augment healthcare delivery, yet their use in preventative care raises persistent concerns regarding safety, scope control, behavioral fidelity, and unintended clinical substitution. We present an executive-governed cognitive architecture for AI-assisted prevention counseling that explicitly separates task orchestration, domain-specific generation, behavioral style enforcement, safety oversight, and learning adaptation into distinct agents operating under centralized supervisory control. Using HIV prevention and pre-exposure prophylaxis (PrEP) counseling as an example application, we describe how this architecture enables adaptive improvement while preserving prospective safety constraints and clear clinical boundaries. Drawing on internal system evaluation, we show that architectural governance can meaningfully shape system behavior, suggesting that responsible deployment of conversational AI in healthcare depends as much on institutional design as on model capability. We discuss implications for integration into primary care and other clinical settings, workforce constraints, reimbursement, and AI governance in healthcare.

Introduction

Artificial intelligence (AI), and in particular large language models (LLMs), are increasingly being explored as tools to augment healthcare delivery.[1-3] In counseling and preventive care, conversational AI systems promise scalable, low-marginal-cost support for education, behavior change, and care navigation.[3-8] However, the use of LLMs on health counseling raises persistent concerns related to safety, scope control, behavioral fidelity, hallucination, and unintentionally replacing clinical care.[8,9] These challenges are especially salient in preventive counseling contexts—such as HIV prevention and pre-exposure prophylaxis (PrEP)—where engagement, trust, and adherence to evidence-based counseling principles are central to effectiveness.[8] As HIV prevention fundamentally involves health behavior change, counseling approaches that support patient autonomy, address ambivalence, and avoid directive or judgmental interactions are associated with improved uptake and persistence of preventive interventions.

Motivational interviewing (MI) is a well-established, evidence-based approach for supporting health behavior change and has demonstrated effectiveness in promoting PrEP uptake and other preventive behaviors.[10-12] However, MI is labor-intensive, requires specialized training, and is difficult to deliver consistently in time-constrained clinical environments such as primary care and emergency departments.[13-15] As a result, despite strong evidence of efficacy, MI-based counseling remains difficult to scale within routine care.[13-15] AI-based conversational agents offer a potential pathway to extend MI-aligned counseling beyond traditional provider encounters, but deployment of general-purpose LLMs without explicit clinical governance mechanisms risks misinformation, loss of clinical boundaries, inconsistent counseling style, and erosion of trust.[1-3,8,9]

Recent work has shown that LLM-based chatbots can generate empathetic, contextually appropriate responses and approximate certain elements of MI.[3,8] However, most existing conversational AI systems are built around a single, general-purpose language model that performs many functions at once.[16] In such designs, the same model is often responsible for interpreting user input, generating responses, accessing external information, and adapting over time, making it difficult to enforce clinical boundaries or audit decision pathways.[16,17] These limitations pose significant barriers to responsible integration of conversational AI into healthcare settings.

In prior work, we conducted an internal usability evaluation of an AI-based chatbot, CHIA (Chatbot for HIV Prevention and Action), designed to deliver MI-aligned counseling for HIV prevention and PrEP.[18] That study demonstrated the feasibility of delivering behaviorally grounded counseling using an LLM-based system and identified key strengths and limitations related to counseling fidelity, multilingual performance, and user engagement.[18] However, the initial system architecture was not designed to explicitly address broader questions of governance, supervision, and scalability within real-world healthcare systems.

In this paper, we present a substantially redesigned version of CHIA that shifts the focus from chatbot performance to system architecture and governance. We introduce an executive-governed cognitive LLM architecture in which task orchestration, domain-specific generation, behavioral style enforcement, safety oversight, and learning adaptation are handled by distinct but coordinated agents operating under centralized supervisory control. This design explicitly separates decision-making functions that are often conflated in LLM-based systems, enabling adaptive learning while preserving prospective safety constraints and clinical scope.

Our contribution is not a claim about replacing clinicians or automating care delivery. Rather, we propose a system-level architecture for deploying AI-assisted counseling as a supervised, assistive tool that complements provider-delivered care. We believe that this architecture is particularly well suited for integration into primary care, emergency department, and mental health settings, where time constraints, competing priorities, and workforce limitations restrict the delivery of evidence-based counseling.[13-15] By offloading structured education, readiness assessment, and non-diagnostic counseling tasks—while maintaining explicit provider oversight—executive-governed AI systems may offer a scalable and sustainable pathway for extending preventive services within existing care models.

Finally, this work has broader implications for policy and implementation questions surrounding AI in healthcare. Digital counseling and care management services already operate within established reimbursement frameworks, including CPT-based models, suggesting a plausible pathway for financial sustainability.[19] By embedding governance and scope control at the architectural level, systems such as CHIA offer a framework for aligning AI-assisted counseling with regulatory, ethical, and economic constraints that shape real-world healthcare delivery.

Section 2. Conceptual Framework and System Architecture

2.1 Design Principles

The redesign of CHIA is guided by three core principles: (1) explicit governance and scope control, (2) separation of cognitive and operational responsibilities, and (3) adaptive learning constrained by post-hoc evaluation rather than prospective autonomy. These principles reflect the recognition that in healthcare settings, the risks associated with AI systems arise not only from

incorrect outputs, but from poorly governed decision pathways and unclear boundaries between assistance and substitution.[20]

First, governance is treated as a first-order design objective rather than an afterthought. The system is structured around a centralized executive agent that supervises all task allocation, tool invocation, and learning updates. This executive layer defines what the system is permitted to do, when adaptation is needed, and how safety and fidelity are evaluated.

Second, CHIA separates functions that are often bundled within a single LLM into distinct agents with narrowly defined roles. This modularity allows each component to be optimized and audited independently, reducing the likelihood that errors or failures in one component will affect other parts of the system.

Third, adaptation is restricted to post-hoc learning mechanisms that are triggered only after outputs have been evaluated for safety, scope compliance, and counseling fidelity. This ensures that improvements in performance do not come at the cost of expanded clinical authority or degraded behavioral alignment.

2.2 Executive-Governed Multi-Agent Architecture

CHIA employs a hierarchical, multi-agent architecture with an explicit safety and governance sequence that controls how user input is processed and how responses are delivered. All user input is first evaluated by a mental health triage agent, which screens for signals of acute distress, self-harm, violence, or other crisis situations. When such signals are detected, the system does not attempt to provide counseling. Instead, it follows predefined escalation protocols, which may include directing users to appropriate crisis hotlines, emergency services, or clearly informing

them that the content falls outside the scope of CHIA's training and must be addressed by a qualified mental health professional. In these cases, no further automated counseling is performed.

For interactions that do not trigger safety escalation, user input is passed to an executive supervisory layer that structures the interaction. This executive layer breaks the input into bounded informational and counseling tasks and determines which specialized agents should be engaged. An assistant agent then carries out these tasks using a restricted set of tools, including HIV risk assessment, assessment of readiness to change, referral to HIV prevention providers via the HIV.gov provider dataset, and signaling when human support may be beneficial.

The outputs produced by the assistant agent are not shown directly to users. Instead, they are passed to a dedicated counselor agent that uses a motivational interviewing–trained language model to interpret and reframe the information. This step ensures that responses adhere to motivational interviewing principles, including nonjudgmental tone, respect for autonomy, and avoidance of directive or diagnostic language.

Before any response is delivered, it is reviewed by an independent gatekeeper agent. The gatekeeper does not rewrite content. Rather, it evaluates whether the proposed response meets safety, scope, and appropriateness criteria. If concerns are identified, the response is returned to the counselor agent for revision. Only responses that satisfy all governance checks are approved for user delivery.

The entire workflow is governed through a graph-based control structure implemented within the AutoGen framework.[21] This structure explicitly constrains the sequence of actions and the

direction of information flow, ensuring that safety screening, task execution, counseling synthesis, and governance review occur in the correct order. Together, this design enables CHIA to provide supportive, adaptive counseling while maintaining clear boundaries, escalation pathways, and institutional oversight.

3. Data, Methods, and Evaluation Strategy

This study draws on internal system development data and structured internal evaluation of CHIA's redesigned architecture.[18] The goal of this evaluation is not to establish clinical effectiveness, but rather to assess whether the executive-governed architecture can reliably enforce scope control, safety, and behavioral fidelity while supporting adaptive learning. As such, the evaluation strategy focuses on system behavior, governance performance, and counseling alignment rather than patient outcomes.

The CHIA system is built on a LLM foundation with multilingual capability and is deployed within a secure cloud-based environment (Azure) and compliant with HIPPA policy.[17,22] Domain knowledge related to HIV prevention and PrEP is accessed exclusively through curated, physician-validated datasets using retrieval-augmented generation, ensuring that informational outputs are grounded in approved clinical content.[23,24] All external tool access—including risk assessment, readiness-to-change evaluation, and provider referral—is mediated through predefined tools with explicit scope constraints.[25,26]

Internal evaluation was conducted using simulated user interactions designed to reflect common counseling scenarios encountered in HIV prevention contexts.[18] These scenarios included information-seeking questions, expressions of ambivalence toward PrEP, readiness assessment

prompts, and requests for care navigation support. Each interaction was processed through the full executive-governed pipeline, allowing assessment of task decomposition, agent coordination, safety gating, and post-hoc evaluation.

System performance was assessed across four primary dimensions. First, informational accuracy and groundedness were evaluated to ensure that responses adhered strictly to validated knowledge sources.[9,27] Second, safety and scope compliance were assessed to confirm that outputs did not exceed the system's intended counseling role or drift into diagnostic or therapeutic decision-making.[23,24] Third, counseling fidelity was evaluated using the established MISC coding framework, focusing on adherence to MI spirit and avoidance of directive or coercive language.[28,29] Fourth, governance performance was assessed by examining whether adaptive updates were triggered only through post-hoc evaluation and executive approval, rather than through autonomous model behavior.

This evaluation strategy reflects a deliberate shift away from treating conversational AI as a monolithic predictive model and toward treating it as a governed system whose performance depends on institutional design choices. The emphasis is therefore placed on whether the architecture functions as intended under realistic interaction patterns, rather than on headline performance metrics alone.

4. Preliminary Findings

Internal evaluation of the redesigned CHIA system suggests that the executive-governed architecture successfully enforces separation of responsibilities and maintains prospective safety constraints while supporting iterative improvement.[18] Across simulated counseling interactions,

task decomposition reliably routed user inputs to appropriate domain-specific agents without invoking unauthorized tools or exceeding predefined clinical scope.

The gatekeeper layer consistently intercepted responses that violated safety or scope criteria prior to user delivery, indicating that independent review mechanisms can function effectively even in complex conversational settings. Importantly, no learning updates were applied directly to generation agents without prior post-hoc evaluation, demonstrating that adaptive behavior remained subordinate to executive oversight.

Counseling fidelity assessments indicated that the dedicated counselor agent maintained consistent adherence to MI principles across interaction types. While the system occasionally exhibited conservative tendencies—such as deferring to provider referral earlier than necessary—these behaviors were identifiable through post-hoc evaluation and flagged for targeted refinement. This highlights a key advantage of the architecture: deviations from desired counseling style can be isolated, audited, and corrected without altering unrelated system components.

From a system perspective, these findings underscore the feasibility of combining expressive conversational AI with institutional controls that prevent uncontrolled autonomy. Rather than relying on static prompts or one-time fine-tuning, the executive-governed framework allows CHIA to evolve in response to observed performance while preserving clear accountability pathways.

While these findings are preliminary and based on internal evaluation, they provide evidence that architectural governance can meaningfully shape AI system behavior. This suggests that many

concerns surrounding the use of LLMs in healthcare may be better addressed through system design and oversight mechanisms than through model restriction alone.

5. Implementation, Integration, and Policy Implications

The architecture presented in this paper is designed explicitly for integration into real-world healthcare settings rather than for standalone consumer use. By construction, CHIA functions as a clinician-supervised decision support and counseling assistance system, not as a substitute for provider judgment or care delivery. This distinction is critical for both ethical deployment and regulatory alignment.

Primary care, emergency department, and mental health settings represent particularly promising environments for such systems. In these contexts, providers face substantial time constraints, competing clinical priorities, and increasing administrative burden, all of which limit the delivery of evidence-based preventive counseling.[30-33] Executive-governed AI systems can assist by providing structured education, readiness assessment, and non-diagnostic counseling support outside or alongside brief clinical encounters, while preserving provider authority over clinical decisions.

From an economic perspective, the sustainability of AI-assisted counseling depends on alignment with existing reimbursement and care delivery models. Digital counseling, care management, and patient education services already operate within established billing frameworks, including CPT-based mechanisms.[19] Systems such as CHIA may therefore complement existing workflows by extending the reach of preventive services without requiring

entirely new payment structures. Importantly, architectural safeguards that prevent clinical substitution help ensure that such tools augment rather than displace provider labor.

This work also has implications for AI governance and regulation in healthcare. Current policy debates often focus on model-level safeguards, such as restricting training data or limiting output capabilities.[34] The CHIA architecture illustrates an alternative approach in which governance is embedded at the system level through separation of responsibilities, independent oversight agents, and constrained adaptation pathways. Such designs may offer regulators and healthcare organizations a more practical route to balancing innovation with safety and accountability.

Finally, while this paper focuses on HIV prevention and PrEP counseling, the executive-governed architecture is intentionally generalizable. Similar design principles could be applied to other preventive and behavioral health domains where engagement, trust, and scope control are paramount. As healthcare systems continue to explore AI-assisted interventions, architectures that prioritize governance, supervision, and integration may be essential for realizing their potential at scale.

Acknowledgements

Artificial intelligence (GPT-4o) was used solely for the purpose of grammar checking and corrections. This project was funded by Brown University's Data Science Institute. Research reported in this publication was supported by the National Institute of Allergy and Infectious Diseases (R21AI192507) and the National Institute of Mental Health (R34MH138136) of the National Institutes of Health. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. The funders had no involvement in the study design, data collection, analysis, interpretation, or the writing of the manuscript. The datasets generated or analyzed during this study are available from the corresponding author on reasonable request. EP is a paid consultant for Google Research. All other authors declare no competing interests.

References

1. Nazir A, Wang Z. A Comprehensive Survey of ChatGPT: Advancements, Applications, Prospects, and Challenges. *Meta Radiol*. 2023;1(2):100022. doi:10.1016/j.metrad.2023.100022
2. Tudor Car L, Dhinagaran DA, Kyaw BM, et al. Conversational Agents in Health Care: Scoping Review and Conceptual Analysis. *J Med Internet Res*. 2020;22(8):e17158. Published 2020 Aug 7. doi:10.2196/17158
3. Aggarwal A, Tam CC, Wu D, Li X, Qiao S. Artificial Intelligence-Based Chatbots for Promoting Health Behavioral Changes: Systematic Review. *J Med Internet Res*. 2023;25:e40789. Published 2023 Feb 24. doi:10.2196/40789
4. Brown A, Kumar AT, Melamed O, et al. A Motivational Interviewing Chatbot With Generative Reflections for Increasing Readiness to Quit Smoking: Iterative Development Study. *JMIR Ment Health*. 2023;10:e49132. Published 2023 Oct 17. doi:10.2196/49132
5. He L, Basar E, Wiers RW, Antheunis ML, Krahmer E. Can chatbots help to motivate smoking cessation? A study on the effectiveness of motivational interviewing on engagement and therapeutic alliance. *BMC Public Health*. 2022;22(1):726. Published 2022 Apr 12. doi:10.1186/s12889-022-13115-x
6. Prochaska JJ, Vogel EA, Chieng A, et al. A Therapeutic Relational Agent for Reducing Problematic Substance Use (Woebot): Development and Usability Study. *J Med Internet Res*. 2021;23(3):e24850. Published 2021 Mar 23. doi:10.2196/24850
7. Abid A, Baxter SL. Breaking Barriers in Behavioral Change: The Potential of Artificial Intelligence-Driven Motivational Interviewing. *J Glaucoma*. 2024;33(7):473-477. doi:10.1097/IJG.0000000000002382
8. Van Heerden A, Bosman S, Swendeman D, Comulada WS. Chatbots for HIV Prevention and Care: a Narrative Review. *Curr HIV/AIDS Rep*. 2023;20(6):481-486. doi:10.1007/s11904-023-00681-x
9. Fujimoto M, Hunter L, McCoy S, Outran S, Packel L. Evaluating AI chatbots for HIV prevention: an assessment of response quality and user tailoring. California HIV / AIDS Policy Research Center; 2024. URL: https://chprc.org/wpcontent/uploads/2022/06/AI_Policy_Brief_Oct2024.pdf [Accessed 2026-01-10]
10. Chan PA, Nunn A, van den Berg JJ, et al. A Randomized Trial of a Brief Behavioral Intervention for PrEP Uptake Among Men Who Have Sex With Men at Increased Risk for HIV Infection. *J Acquir Immune Defic Syndr*. 2021;87(3):937-943. doi:10.1097/QAI.0000000000002671
11. Moitra E, van den Berg JJ, Sowemimo-Coker G, Chau S, Nunn A, Chan PA. Open pilot trial of a brief motivational interviewing-based HIV pre-exposure prophylaxis intervention for men who have sex with men: preliminary effects, and evidence of

feasibility and acceptability. *AIDS Care*. 2020;32(3):406-410. doi:10.1080/09540121.2019.1622644
12. Miller WR, Rollnick S. Motivational Interviewing: Helping People Change. 3rd ed. Guilford Press; 2013. URL: https:// www.biblio.com/book/motivational-interviewing-helping-people-change-miller/d/1659959579 [Accessed 2026-01-10]
13. Jacobs NN, Calvo L, Dieringer A, Hall A, Danko R. Motivational Interviewing Training: A Case-Based Curriculum for Preclinical Medical Students. *MedEdPORTAL*. 2021;17:11104. Published 2021 Feb 12. doi:10.15766/mep_2374-8265.11104
14. Rollnick S, Miller WR, Butler CC. Motivational Interviewing in Health Care: Helping Patients Change Behavior. Guilford Press; 2008. URL: https://www.guilford.com/books/Motivational-Interviewing-in-Health-Care/Rollnick-MillerButler/9781462550371 [Accessed 2026-01-10] [doi: 10.1080/15412550802093108]
15. Rubak S, Sandbaek A, Lauritzen T, Christensen B. Motivational interviewing: a systematic review and meta-analysis. *Br J Gen Pract*. 2005;55(513):305-312.
16. Ray PP. ChatGPT: a comprehensive review on background, applications, key challenges, bias, ethics, limitations and future scope. Internet Things Cyber Phys Syst. 2023;3:121-154. [doi: 10.1016/j.iotcps.2023.04.003]
17. Zhang N, Sun Z, Xie Y, Wu H, Li C. The latest version ChatGPT powered by GPT-4o: what will it bring to the medical field? Int J Surg. Sep 1, 2024;110(9):6018-6019. [doi: 10.1097/JS9.0000000000001754] [Medline: 38857508]
18. Tao J, Pavlick E, Grondin A, Bustamente JD, Martin H, Parent H, Fenn N, Almonte A, Maguire-Wilkerson A, Gu M, Rusley J, Perler B, Wray T, Nunn AS, Chan PA. Evaluation of an Artificial Intelligence Conversational Chatbot to Enhance HIV Pre-Exposure Prophylaxis Uptake: Development and Usability Internal Testing. Journal of Medical Internet Research. 02/10/2025:79671 (forthcoming/in press) [doi: 10.2196/79671] URL: https://preprints.jmir.org/preprint/79671
19. O'Reilly KB. 288 new CPT codes cover digital health, AI and more. AMA website. Published September 12, 2025. Accessed January 12, 2026. URL: https://www.ama-assn.org/practice-management/cpt/288-new-cpt-codes-cover-digital-health-ai-and-more
20. Khawaja Z, Bélisle-Pipon JC. Your robot therapist is not your therapist: understanding the role of AI-powered mental health chatbots. Front Digit Health. 2023 Nov 8;5:1278186. doi: 10.3389/fdgth.2023.1278186. PMID: 38026836; PMCID: PMC10663264.
21. Wu Q, Bansal G, Zhang J, et al. AutoGen: enabling next-gen LLM applications via multi-agent conversation. arXiv. Preprint posted online on Aug 16, 2023. [doi: 10.48550/arXiv.2308.08155]
22. Microsoft. HIPAA – Azure Compliance. Microsoft Learn. Updated April 5, 2023. Accessed January 12, 2026. Available at: https://learn.microsoft.com/en-us/azure/compliance/offerings/offering-hipaa-us

23. Papadimitriou I, Gialampoukidis I, Vrochidis S, Ioannis K. RAG playground: a framework for systematic evaluation of retrieval strategies and prompt engineering in RAG systems. arXiv. Preprint posted online on Dec 16, 2024. [doi: 10.48550/arXiv.2412.12322]
24. Es S, James J, Espinosa-Anke L, Schockaert S. RAGAS: automated evaluation of retrieval augmented generation. arXiv. Preprint posted online on Sep 26, 2023. [doi: 10.18653/v1/2024.eacl-demo.16]
25. Raihan N, Cogburn M. Stages of Change Theory. 2023 Mar 6. In: StatPearls [Internet]. Treasure Island (FL): StatPearls Publishing; 2025 Jan–. PMID: 32310465.
26. Smith DK, Pals SL, Herbst JH, Shinde S, Carey JW. Development of a clinical screening index predictive of incident HIV infection among men who have sex with men in the United States. J Acquir Immune Defic Syndr. Aug 1, 2012;60(4):421-427. [doi: 10.1097/QAI.0b013e318256b2f6] [Medline: 22487585]
27. Abbasian M, Khatibi E, Azimi I, et al. Foundation metrics for evaluating effectiveness of healthcare conversations powered by generative AI. NPJ Digit Med. Mar 29, 2024;7(1):82. [doi: 10.1038/s41746-024-01074-z] [Medline: 38553625]
28. Moyers TB, Manuel JK, Ernst D. Motivational interviewing treatment integrity coding manual 4.2.1. Center on Alcoholism, Substance Abuse, and Addictions (CASAA); 2015. URL: https://motivationalinterviewing.org/sites/default/files/miti4_2.pdf [Accessed 2026-01-10]
29. Moyers TB, Martin T, Manuel JK, Miller WR. The motivational interviewing treatment integrity (MITI) code: version 2.0. University of New Mexico Center on Alcoholism, Substance Abuse, and Addictions (CASAA); 2003. URL: https://casaa.unm.edu/assets/docs/miti1.pdf [Accessed 2026-01-10]
30. Porter J, Boyd C, Skandari MR, Laiteerapong N. Revisiting the Time Needed to Provide Adult Primary Care. J Gen Intern Med. 2023 Jan;38(1):147-155. doi: 10.1007/s11606-022-07707-x. Epub 2022 Jul 1. PMID: 35776372; PMCID: PMC9848034.
31. Sartini M, Carbone A, Demartini A, Giribone L, Oliva M, Spagnolo AM, Cremonesi P, Canale F, Cristina ML. Overcrowding in Emergency Department: Causes, Consequences, and Solutions-A Narrative Review. Healthcare (Basel). 2022 Aug 25;10(9):1625. doi: 10.3390/healthcare10091625. PMID: 36141237; PMCID: PMC9498666.
32. Satterwhite S, Nguyen MT, Honcharov V, McDermott AM, Sarkar U. "Good Care Is Slow Enough to Be Able to Pay Attention": Primary Care Time Scarcity and Patient Safety. J Gen Intern Med. 2024 Jul;39(9):1575-1582. doi: 10.1007/s11606-024-08658-1. Epub 2024 Feb 15. PMID: 38360962; PMCID: PMC11255151.
33. Zhu JM, Eisenberg M. Administrative Frictions and the Mental Health Workforce. *JAMA Health Forum*. 2024;5(3):e240207. Published 2024 Mar 1. doi:10.1001/jamahealthforum.2024.0207

34. Wagner JK, Doerr M, Schmit CD. AI Governance: A Challenge for Public Health. JMIR Public Health Surveill. 2024 Sep 30;10:e58358. doi: 10.2196/58358. PMID: 39347615; PMCID: PMC11457458.