

Who Gets a Human Doctor—and Which One? From Evaluating Medical Answers only to Evaluating Answers + Care Orchestration & Routing

Anonymous Author(s)

Abstract

Consumer health AI can change access to care without making a false clinical statement. We argue that **care routing decisions must be auditable** when systems construct choices, rank providers, or execute bookings. The proposed audit distinguishes mode routing, deciding whether AI continues autonomously, seeks clarification, or involves a licensed human clinician, from destination routing, selecting the receiving provider or service. It compares retrieved and displayed options with an independently assessed feasible set, tests affiliation while holding clinical fit and access constant, and traces execution through care, fallback, and exit. Patient and clinician judgments inform distinct evaluation outcomes. Fictional US breast-evaluation and Indian diabetes cases illustrate the audit. We specify tests of added value beyond existing pathway and access frameworks. A small, fully synthetic pilot with three tool-using models shows that the audit’s measurements can be computed and that they separate routing from answer quality: models gave correct advice yet booked infeasible or deadline-breaking routes, and a partner label shifted choice among identical clinics, in either direction depending on the model. The paper reports no evidence of steering by any named platform.

Keywords: consumer health AI, routing audits, patient involvement, access to care, platform accountability

Data and Code Availability This position paper synthesizes cited public sources and uses fictional scenarios. The synthetic pilot in Appendix F uses a generated world with no real clinics or patients; its code, data, and an interactive report are included in the supplementary material ([supplementary_pilot/](#)). No patient dataset was used.

Institutional Review Board (IRB) IRB review is not applicable: this position paper and its fully synthetic pilot recruited no participants, accessed no

individual patient records, and performed no clinical intervention.

1. The position and its scope

A 52-year-old woman asks about a suspicious finding on a screening mammogram; a 40-year-old man asks about missed diabetes follow-up. Within each case, the clinical details and answer stay fixed while access conditions vary. Figure 1 shows the routing decisions to audit. Table 2 in the appendix details the fictional circumstances.

Documented product functions include records (ChatGPT Health) (OpenAI, 2026), appointments and pharmacy fulfillment (Amazon Health AI) (Amazon Health Services, 2026), consultations and payment (Afu) (Ant Group, 2025), clinical and administrative connectors (Claude) (Anthropic, 2026), and booking (Gemini with Zocdoc) (Google, 2026; Zocdoc, 2026). Descriptions establish functions, not benefit.

Services that construct, rank, recommend, or execute access to care must permit qualified auditors to reconstruct and test routing. Evaluate clinical interpretation, realized care, and operator choices separately.

Wu and colleagues already argue for pathway evaluation, routing transparency, operator accountability, and exit options (Wu et al., 2026). PACT addresses continuity, confirmation, escalation, completion, and equity across the care journey (Khairat and Safran, 2026). Our specific proposal is to audit the options an operator considered and excluded, using an independently verified reference set and controlled affiliation tests, alongside clinical advice and completed-care outcomes. The proposed audit must demonstrate added value beyond existing frameworks.

Wachter suggests younger patients may differ in their preference for human care (Wachter, 2026). A qualitative study of 14 adolescents with depression reported less anxiety about approaching a chatbot, but excluded suicidal participants (Kuhlmeier et al.,




Figure 1: Fictional cases with fixed clinical facts and varying access. Panel A concerns a 52-year-old woman with a suspicious screening mammogram; Panel B a 40-year-old man with missed diabetes follow-up. Logos denote example platforms, not tested outputs. Ribbon widths encode no quantities.

2025). Some patients may prefer AI, including for stigmatized disclosures; age alone cannot establish that preference. Price, convenience, and platform defaults may shape the choice.

We developed the proposal through a critical narrative review of clinical decision science, access, implementation, platform governance, and product documentation. The review was not systematic; purposive product examples cannot estimate prevalence.

2. From pathway accountability to a routing audit

Decision science distinguishes prediction from action value (Pauker and Kassirer, 1980; Vickers and Elkin, 2006); care cascades track attrition (Gardner et al., 2011; Haber et al., 2016). Implementation and access

research address adoption, effectiveness, feasibility, cost, distance, and acceptability (Glasgow et al., 1999; Proctor et al., 2011; Aday and Andersen, 1974; Penchansky and Thomas, 1981). Routing audits must add value beyond these foundations.

Model updates, directories, insurance rules, appointment feeds, and commercial contracts can alter routing. Multiple organizations may jointly produce omissions, excess exposure, failed transactions, or lost alternatives.

3. What a routing audit must examine

Appendix G links each element of the audit to the evidence that motivates it.

Audit problem representation before routing. The route starts before a provider list appears. Prob-

Table 1: What a care-routing audit must measure and test. Measures require validation and should remain separate.

Audit target	What to measure	How to test it
Problem representation	Critical omissions, unsupported claims, uncertainty, revision	Independent clinician ratings; sequential evidence; adjudicated-summary substitution
Mode routing	Appropriate continuation, handoff, urgency	Prespecified acceptable actions and deadlines
Destination routing	Missing feasible options, infeasible recommendations, rank exposure	Independent, timestamped feasible set
Commercial influence	Changes in inclusion, rank, selection, execution	Vary affiliation; hold fit, cost, availability, and access fixed
Execution and follow-through	Completed care, delay, burden, fallback, exit	Trace beyond booking; separate unresolved outcomes
Shared capacity	Targeted and other patients' waits, throughput, unmet demand	Compare allocation under the same capacity constraints
Mental health and sensitive disclosure	Willingness to confide, appropriate trust, crisis handoff, record privacy	Voluntary user studies; clinician-written crisis tests; access, retention, deletion, and disclosure audits

lem representation compresses a case into the features needed for diagnosis and action (Wachter, 2026; McQuade et al., 2024). Clinical Problem Solvers describes relevant patient context, illness duration and tempo, and the clinical syndrome, using semantic qualifiers and updating the summary as evidence changes (de Oliveira Santana and Murillo, 2022). A medical AI should turn the janky stream of consciousness patients type into a useful working representation. The patient should not have to learn to prompt like a doctor.

Two clinicians should independently rate representations using only evidence available at the decision point. Report their agreement before a third clinician adjudicates disagreements. Accept multiple valid summaries; assess omissions, unsupported assertions, temporal qualifiers, uncertainty, revision, and patient burden separately (Appendix B). Possible cancer must not become confirmed cancer or a stage. Preserve patient-confirmed goals and access constraints in a linked record. Table 1 summarizes the audit; these are proposed dimensions, not a validated composite score.

Audit whether and where to hand off. *Mode routing* decides whether to continue, clarify, involve a clinician, or escalate. *Destination routing* selects the receiving person, facility, service, or channel. Assess both against clinician-adjudicated acceptable ac-

tions and deadlines. Destination quality also requires current feasibility, including coverage, cost, capacity, distance, language, continuity, and preference. Audit execution through booking, payment, referral, prescription transfer, or human contact.

Which feasible options did the platform miss?

For patient i at transition k , let $\mathcal{F}_{ik}$ be an independently assessed set of clinically suitable, currently feasible options, $\mathcal{C}_{ik}$ the retrieved candidates, and $\mathcal{L}_{ik} \subseteq \mathcal{C}_{ik}$ the displayed options. Do not assume $\mathcal{F}_{ik} \subseteq \mathcal{C}_{ik}$. Prespecify service, catchment, and verification window. For nonempty $\mathcal{F}_{ik}$, report

$$R_{ik} = \frac{|\mathcal{C}_{ik} \cap \mathcal{F}_{ik}|}{|\mathcal{F}_{ik}|},$$

$$D_{ik} = \frac{|\mathcal{L}_{ik} \cap \mathcal{F}_{ik}|}{|\mathcal{F}_{ik}|},$$

for retrieval and displayed coverage respectively. Report list length, exclusions, rank exposure, and infeasible options. Shorter lists need not be worse. Distinguish empty feasible sets from retrieval failure; qualify incomplete verification.

For example, two agents give identical advice for a hypothetical patient with three suitable, accessible clinics. One retrieves all three; the other omits a cheaper clinic with an earlier appointment. Inde-

pendent coverage detects the omission that answer grading misses. Its cause still requires investigation.

Did the selected route produce appropriate care? Validate completion against source events; a warehouse preprint found executable queries could miscount procedures (Marshall et al., 2026). Trace time, cost, and care. Separate unresolved episodes, informed refusal, recovery, and outside care. Audit patient experience and preference use separately (Garg et al., 2026); Appendix D specifies tests.

4. How to test routing decisions

Within each setting in Figure 1, compare routing configurations against the same feasible reference set. Geography alone cannot identify an operator effect.

Routing toward constrained services can delay other patients even when mean waiting is unchanged (Cobham, 1954, 1955). Report targeted and other patients' waits, throughput, and unresolved demand.

Declare the policy and authority under test. Declare authority for each function, including user confirmation, clinician authorization, and automatic execution. Freeze the model, prompt, retrieval and directory snapshots, filters, ranking objective, tools, and escalation rules. Vary availability through declared scenarios. Deployment audits must timestamp conditions and repeat tests after updates.

Separate access constraints from commercial influence. Recommender research supplies exposure and provider-coverage measures (Singh and Joachims, 2018; Zehlike et al., 2022). A synthetic physician-choice audit found that ratings, fees, position, demographic signals, and personas changed recommendations across more than 3,000 choice sets; explanations did not consistently identify influential attributes (Gillani and Baig, 2026). Our proposed audit first varies access conditions while preserving clinically appropriate advice. It then randomizes affiliation within matched choice sets while holding clinical fit, price, availability, and access constant. Report changes in inclusion, rank, selection, and execution, with repeated runs and uncertainty estimates. Synthetic affiliation tests identify sensitivity to a cue, not proof of commercial steering in deployment. Appendix F runs both tests on a synthetic directory.

5. Audit findings must lead to corrective action

Preserve an inspectable decision record. For each episode, preserve versions, conversation state, candidate sources, retrieved pool, filters, displayed order, ranking objective, timestamps, commercial relationships, confirmation, tool calls, transaction results, fallback, reversal, and later care when consent and law allow. Qualified independent auditors need protected access to reconstruct the route. Model weights alone cannot expose options removed by business rules; public release of patient data is unnecessary.

Audit trust, crisis handoff, and sensitive records. Some patients may find it easier to disclose sexuality-related concerns or thoughts of self-harm without anticipated judgment. Sexuality itself is not a disorder. In a US screening study, stigma-related avoidance predicted chatbot use among distressed adults (Kosyluk et al., 2024); the study did not test general-purpose systems. WHO reports critical global mental-health workforce shortages, especially in lower-income countries (World Health Organization, 2025). Scarcity strengthens the access case for safe digital support but does not relax crisis handoff, privacy, or human fallback requirements. Audit voluntary comfort confiding and understanding of AI limits and confidentiality, rather than maximizing disclosure. Clinician-written crisis scenarios should test escalation, reachable human help, and failed-handoff recovery. A referral message is not completed care. Use synthetic records to test unauthorized access, cross-user leakage, retention, deletion, and onward disclosure, including training or advertising reuse. Minimize stored content and verify consent and access logs. Appendix C specifies additional controls.

Keep paid influence outside clinical ranking. Payment, bids, commissions, and common ownership should not increase clinical rank. Prohibit paid placement in emergency routing, crisis response, treatment selection, and clinically indicated referrals. Where elective discovery permits promotion, separate it, identify the payer, avoid sensitive targeting, preserve suitable nonpaying options, and prohibit default execution. Advertising and platform rules offer precedents, but clinical appropriateness requires separate testing (U.S. Federal Trade Commission, 2015; European Parliament and Council of the European Union, 2022).

Assign responsibility across handoffs. Parsons and Rodman emphasize clinical explanation, repair,

and continuity (Parsons and Rodman, 2026). We propose recording who owns a pending transfer, when a receiving team accepts it, and who can correct failures. A signature cannot replace informed deliberation or release operators from responsibility for routing they control. Simulated failed-transfer tests should measure recovery time, unresolved cases, and burden; this proposed protocol is unvalidated.

Apply duties by function and risk. Assistants, insurer portals, hospital applications, and marketplaces should face comparable duties for comparable routing functions. Requirements should increase with clinical stakes, execution authority, scale, sensitive data use, commercial conflict, irreversibility, and difficulty of exit.

Connect monitoring to action. Each audit measure needs a named owner, review schedule, threshold, and predefined response (Keyes et al., 2026; Rao et al., 2026). Missing feasible providers should trigger retrieval or directory investigation. Affiliation effects should trigger ranking review. Failed execution should trigger confirmation and fallback checks. Safety-critical routing failures should restrict execution pending review, with a human route available. Retest after changes to models, prompts, tools, directories, contracts, or ranking objectives.

6. Counterarguments and limits

Evaluators must prespecify transitions the system can change, name a comparator, and use causal or quasi experimental designs for effect claims. Attribute failures only when the service caused or materially increased them.

PACT, care cascades, implementation science, and access research may already cover much of pathway evaluation. The framework adds value only if candidate generation, ranking, affiliation, cross organization execution, and exit explain outcomes beyond those constructs.

Patient navigation trials suggest that reducing coordination barriers can improve diagnostic resolution, although results differ across settings (Raich et al., 2012; Wells et al., 2012). Common ownership can remove handoffs, and a transaction fee does not prove biased rank. Evaluation must measure completion gains and steering risk together. The position opposes paid influence in clinical ranking, not integration itself.

Directory, network, price, and capacity data may be wrong or stale. Audits should timestamp verifica-

tion, repeat queries, report missing observations, and use sensitivity analysis. Data minimization, access logs, secure environments, and aggregate reports can reduce the privacy, security, and trade secret risks of protected inspection.

Apart from the synthetic pilot in Appendix F, the paper reports no primary experiments. Its proposals draw on adjacent methods, product documentation, and empirical analogues. It does not show that integrated consumer AI improves outcomes, creates capacity, reduces disparities, or engages in commercial steering.

7. Tests of the audit’s added value

The first study should preregister the matched access and affiliation audit in Section 4. Compare agents with clinician navigation, insurer search, and a rules based baseline. Test whether retrieval omissions, rank exposure, and transaction failures reveal problems that answer grading and existing access measures miss.

A second study should compare model-generated and adjudicated representations in the same frozen router with identical access data and tools. Measure urgency, destination appropriateness, and critical errors; Appendix B specifies the design. Preference studies should vary price, waiting time, convenience, and interface defaults while preserving clinical safety, with age-stratified analysis. Choosing AI when the human option is unaffordable cannot establish an unconstrained preference.

The strongest downstream study should compare usual navigation, information only AI, and the same AI with orchestration. Clinicians can assess clinical appropriateness; patients can assess understanding and fit. Both should define outcomes and disagreement rules (Appendix D). A human navigator can separate navigation from AI.

Operator measures add value only if they explain or predict outcomes beyond established implementation and access measures. The case for protected inspection weakens if audits limited to visible outputs (answers, displayed options, and bookings, without retrieval or ranking logs) detect the same failures with comparable sensitivity and cost. Null affiliation effects across adequately powered, representative tests would weaken the commercial concern.

8. Conclusion

Care routing decisions must be auditable. Pachocki distinguishes goal alignment, pursuing assigned objectives, from value alignment, applying principles under unclear or conflicting objectives (Pachocki, 2026). We extend this distinction to both mode and destination routing. Value alignment in both is necessary to help shape fair, equitable healthcare systems that put patients first. Audits must test whether routing preserves safety, informed choice, and equitable access when objectives conflict.

Wachter predicts primary care tiers shaped by coverage, means, and clinical complexity (Wachter, 2025, 2026). Choosing AI-led or human-led care is mode routing; selecting the clinician or service is destination routing. His economy and business class metaphor requires the same high clinical-safety floor for both tiers. Audits must test how price, coverage, convenience, and defaults shape both decisions, including escalation and completed care. Some younger patients may prefer AI, but audits must distinguish preference from constraint. Patient choice, safe human exit, freedom from paid clinical ranking, and the Hippocratic ideal of doing no harm must remain non-negotiable.

AI assistance An AI assistant supported literature organization, drafting, and LaTeX and figure preparation. The authors are responsible for the manuscript’s claims, citations, and interpretations.

References

- L A Aday and R Andersen. A framework for the study of access to medical care. *Health services research*, 9(3):208–20, 1974. URL <https://pubmed.ncbi.nlm.nih.gov/4436074/>.
- All India Institute of Medical Sciences, New Delhi. Centre for community medicine: Departments and centers, 2026. URL <https://aiims.edu/index.php/en/departments-and-centers/87-centre-for-community-medicine>. Accessed 5 September 2026.
- Amazon Health Services. Meet Health AI, 2026. URL <https://health.amazon.com/health-ai/learn-more>. accessed 2026-08-31.
- Ant Group. Ant group launches AI healthcare app AQ, 2025. URL <https://www.antgroup.com/news-media/press-releases/1750924800000>. Chinese-language release; accessed 2026-08-31.
- Anthropic. Advancing Claude in healthcare and the life sciences, 2026. URL <https://www.anthropic.com/news/healthcare-life-sciences>. accessed 2026-08-31.
- Soumya Banerjee, Phil Alsop, Linda Jones, and Rudolf N. Cardinal. Patient and public involvement to build trust in artificial intelligence: A framework, tools, and case studies. *Patterns*, 3(6):100506, 2022. doi: 10.1016/j.patter.2022.100506. URL <https://doi.org/10.1016/j.patter.2022.100506>.
- Andrew M. Bean, Rebecca Elizabeth Payne, Guy Parsons, Hannah Rose Kirk, Juan Ciro, Rafael Mosquera-Gómez, Sara Hincapié M, Aruna S. Ekanayaka, Lionel Tarassenko, Luc Rocher, and Adam Mahdi. Reliability of LLMs as medical assistants for the general public: A randomized preregistered study. *Nature Medicine*, 32(2):609–615, 2026. doi: 10.1038/s41591-025-04074-y. URL <https://doi.org/10.1038/s41591-025-04074-y>.
- Joanna Bisgaier and Karin V. Rhodes. Auditing access to specialty care for children with public insurance. *New England Journal of Medicine*, 364(24):2324–2333, 2011.
- Centers for Medicare & Medicaid Services. Online provider directory review report. Technical report, Centers for Medicare & Medicaid Services, 2017. Review of Medicare Advantage online provider directories conducted in 2016.
- Alan Cobham. Priority assignment in waiting line problems. *Journal of the Operations Research Society of America*, 2(1):70–76, 1954. doi: 10.1287/opre.2.1.70.
- Alan Cobham. Letter to the editor—priority assignment—a correction. *Journal of the Operations Research Society of America*, 3(4):547, 1955. doi: 10.1287/opre.3.4.547.
- Marcela A. de Oliveira Santana and Franco Murillo. Problem representation. The Clinical Problem Solvers, June 2022. URL <https://clinicalproblemsolving.com/2022/06/12/problem-representation/>. Clinical Reasoning Corner; reviewed by Jack Penner.
- European Parliament and Council of the European Union. Regulation (EU) 2022/2065 on a single market for digital services and amending directive 2000/31/EC (Digital Services Act), 2022. URL <https://eur-lex.europa.eu/eli/reg/2022/2065/oj>.
- Edward M Gardner, Margaret P McLees, John F Steiner, Carlos Del Rio, and William J Burman. The spectrum of engagement in hiv care and its relevance to test-and-treat strategies for prevention of hiv infection. *Clinical infectious diseases : an official publication of the Infectious Diseases Society of America*, 52(6):793–800, 2011. doi: 10.1093/cid/ciq243. URL <https://pubmed.ncbi.nlm.nih.gov/21367734/>.
- Namrata Garg, Omid Boozarpour, Jose Alaras, Ivy Tran, Christopher R. Peabody, and Nicholas Stark. Efficiency and humanism: The impact of the humanistic charting tool on patient experience in the emergency department. *Western Journal of Emergency Medicine*, 2026. doi: 10.5811/westjem.62128. URL <https://escholarship.org/uc/item/2d2418q1>. Published 29 August 2026; article in press.
- Syeda Anshrah Gillani and Mirza Samad Ahmed Baig. Whose doctor does the AI recommend? an algorithm audit of reputation and demographic signals in large language model-assisted physician choice. *arXiv preprint arXiv:2608.14399*, 2026. URL <https://arxiv.org/abs/2608.14399>.
- R E Glasgow, T M Vogt, and S M Boles. Evaluating the public health impact of health promotion interventions: the re-aim framework. *American journal of public health*, 89(9):1322–7, 1999. doi:

- 10.2105/ajph.89.9.1322. URL <https://pubmed.ncbi.nlm.nih.gov/10474547/>.
- Google. Check the availability and requirements of connected apps in Gemini Apps, 2026. URL <https://support.google.com/gemini/table/17434654?hl=en-GB>. accessed 2026-08-31.
- Noah Haber, Deenan Pillay, Kholoud Porter, and Till Bärnighausen. Constructing the cascade of hiv care: methods for measurement. *Current opinion in HIV and AIDS*, 11(1):102–8, 2016. doi: 10.1097/COH.0000000000000212. URL <https://pubmed.ncbi.nlm.nih.gov/26545266/>.
- Yixing Jiang, Kameron C. Black, Gloria Geng, Danny Park, James Zou, Andrew Y. Ng, and Jonathan H. Chen. MedAgentBench: a realistic virtual EHR environment to benchmark medical LLM agents. *arXiv preprint arXiv:2501.14654*, 2025. URL <https://arxiv.org/abs/2501.14654>.
- Timothy Keyes, Alison Callahan, Abby S. Pandya, Nerissa Ambers, Juan M. Banda, Miguel Fuentes, Carlene Lugtu, Pranav Masariya, Srikar Nallan, Connor O’Brien, Thomas Wang, Emily Alsentzer, Jonathan H. Chen, Dev Dash, Matthew A. Eisenberg, Patricia Garcia, Nimesh Kotecha, Anurag Revri, Michael A. Pfeffer, Nigam H. Shah, and Sneha S. Jain. Why and how to monitor deployed AI systems in health care. *NEJM Catalyst Innovations in Care Delivery*, 7(6):CAT.25.0372, 2026. doi: 10.1056/CAT.25.0372. URL <https://doi.org/10.1056/CAT.25.0372>.
- Saif Khairat and Charles Safran. Designing clinical AI for patient-centered support beyond the visit: the PACT framework for health systems. *npj Health Systems*, 3:70, 2026. doi: 10.1038/s44401-026-00116-w. URL <https://doi.org/10.1038/s44401-026-00116-w>.
- Kristin Kosyluk, Tanner Baeder, Karah Yeona Greene, Jennifer T Tran, Cassidy Bolton, Nele Loecher, Daniel DiEva, and Jerome T Galea. Mental distress, label avoidance, and use of a mental health chatbot: Results from a us survey. *JMIR Formative Research*, 8:e45959, 2024. doi: 10.2196/45959.
- Florian Onur Kuhlmeier, Luise Bauch, Ulrich Gnewuch, and Stefan Lüttke. Designing chatbots to treat depression in youth: Qualitative study. *JMIR Human Factors*, 12:e66632, 2025. doi: 10.2196/66632.
- Nicholas P. Marshall, William Haberkorn, Joshua Faulkenberry, Ganga Palakkattil, Fateme Nateghi Haredasht, Hayden T. Schwenk, Jonathan H. Chen, and Keith Morse. Can a general-purpose coding agent analyze a production hospital data warehouse? medRxiv preprint, 2026. URL <https://doi.org/10.64898/2026.09.02.26362008>. Version 1, posted 4 September 2026; not peer reviewed.
- Casey N. McQuade, Michael G. Simonson, Julia Lister, Andrew P. J. Olson, Laura Zwaan, Scott D. Rothenberger, and Eliana Bonifacino. Characteristics differentiating problem representation synthesis between novices and experts. *Journal of Hospital Medicine*, 19(6):468–474, 2024. doi: 10.1002/jhm.13335. URL <https://doi.org/10.1002/jhm.13335>.
- National Health Authority, Government of India. Consultation paper on operationalising unified health interface in india, 2022. URL <https://www.pib.gov.in/Pressreleaseshare.aspx?PRID=1883652>.
- National Health Authority, Government of India. Unified health interface: Product overview, 2026. URL <https://abdm.gov.in/UHI/product-overview>. accessed 2026-08-31.
- OpenAI. Introducing ChatGPT Health, 2026. URL <https://openai.com/index/introducing-chatgpt-health/>. accessed 2026-08-31.
- Jakub Pachocki. An alien mind. OpenAI, September 2026. URL <https://openai.com/index/an-alien-mind/>. Published 6 September 2026; commentary.
- Andrew S. Parsons and Adam Rodman. Risk, accountability, and the clinician’s role in the artificial intelligence era. *Annals of Internal Medicine*, 2026. doi: 10.7326/ANNALS-26-02478. URL <https://doi.org/10.7326/ANNALS-26-02478>. Ideas and Opinions.
- S G Pauker and J P Kassirer. The threshold approach to clinical decision making. *The New England journal of medicine*, 302(20):1109–17, 1980. doi: 10.1056/NEJM198005153022003. URL <https://pubmed.ncbi.nlm.nih.gov/7366635/>.
- R Panchansky and J W Thomas. The concept of access: definition and relationship to consumer satisfaction. *Medical care*, 19(2):127–40, 1981. doi: 10.1097/00005650-198102000-00001. URL <https://pubmed.ncbi.nlm.nih.gov/7206846/>.

- Enola Proctor, Hiie Silmere, Ramesh Raghavan, Peter Hovmand, Greg Aarons, Alicia Bunger, Richard Griffey, and Melissa Hensley. Outcomes for implementation research: conceptual distinctions, measurement challenges, and research agenda. *Administration and policy in mental health*, 38(2):65–76, 2011. doi: 10.1007/s10488-010-0319-7. URL <https://pubmed.ncbi.nlm.nih.gov/20957426/>.
- Peter C Raich, Elizabeth M Whitley, William Thorland, Patricia Valverde, Diane Fairclough, and Denver Patient Navigation Research Program. Patient navigation improves cancer diagnostic resolution: an individually randomized clinical trial in an underserved population. *Cancer epidemiology, biomarkers & prevention : a publication of the American Association for Cancer Research, cosponsored by the American Society of Preventive Oncology*, 21(10):1629–38, 2012. doi: 10.1158/1055-9965.EPI-12-0513. URL <https://pubmed.ncbi.nlm.nih.gov/23045537/>.
- Anita K. Rao, Andrew J. Keller, Neha Kalra, Ryan Steed, Kweku Kwegyir-Agrey, Kevin Klyman, Diane Staheli, and Stevie A. Bergman. Challenges to the monitoring of deployed AI systems. NIST Trustworthy and Responsible AI NIST AI 800-4, National Institute of Standards and Technology, March 2026. URL <https://doi.org/10.6028/NIST.AI.800-4>.
- Ashudeep Singh and Thorsten Joachims. Fairness of exposure in rankings. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2219–2228, 2018. doi: 10.1145/3219819.3220088. URL <https://doi.org/10.1145/3219819.3220088>.
- A. E. Snow. Can deliberation improve AI’s ability to act on our behalf? The Democratic Futurist, April 2026. URL <https://democraticfuturist.substack.com/p/can-deliberation-improve-ais-ability>. Published 12 April 2026; project report and commentary.
- Sanjay Sood, Kanhaiya Lal, Madhu Bhatia, Gagandeep Kapoor, Sumandeep Singh, Rajesh Kr Kaushish, Davinder Bisht, Vijay Kumar Sharma, Suchandrima Chakraborty, Mahesh Shete, Nitin Kumar Solanki, Aditya Naskar, Mayank Sharma, Elakeya Udhaya Subramaniam, Shekhar Waikar, and Neha Verma. Adoption and utilization of india’s esanjeevani national telemedicine service. *Oxford open digital health*, 3:oqaf025, 2025. doi: 10.1093/oodh/oqaf025. URL <https://pubmed.ncbi.nlm.nih.gov/41158553/>.
- Sophie Staniszewska, Jo Brett, Iveta Simera, Kate Seers, Carole Mockford, Sally Goodlad, Douglas G. Altman, David Moher, Rosemary Barber, Simon Denegri, Anne Entwistle, and Peter Littlejohns. GRIPP2 reporting checklists: Tools to improve reporting of patient and public involvement in research. *BMJ*, 358:j3453, 2017. doi: 10.1136/bmj.j3453. URL <https://doi.org/10.1136/bmj.j3453>.
- Alexandra Thompson, Victoria Bartle, Elizabeth A. Remfry, Duncan J. Reynolds, Michael R. Barnes, Nick J. Reynolds, and Barbara Hanratty. Public and patient involvement in artificial intelligence and big data healthcare research: An exploration of issues and challenges within the AI-Multiply project. *Health Expectations*, 28(6):e70490, 2025. doi: 10.1111/hex.70490. URL <https://doi.org/10.1111/hex.70490>.
- U.S. Federal Trade Commission. Native advertising: A guide for businesses, 2015. URL <https://www.ftc.gov/business-guidance/resources/native-advertising-guide-businesses>.
- Andrew J Vickers and Elena B Elkin. Decision curve analysis: a novel method for evaluating prediction models. *Medical decision making : an international journal of the Society for Medical Decision Making*, 26(6):565–74, 2006. doi: 10.1177/0272989X06295361. URL <https://pubmed.ncbi.nlm.nih.gov/17099194/>.
- Robert M. Wachter. Will AI rescue primary care? Eventually, But Not Yet. Pattern Recognition, October 2025. URL <https://robertwachter.substack.com/p/will-ai-rescue-primary-care>. Published 8 October 2025.
- Robert M. Wachter. On straw men and doormen. Ground Truths, August 2026. URL <https://robertwachter.substack.com/p/on-straw-men-and-doormen>. Published 27 August 2026.
- Kristen J Wells, Ji-Hyun Lee, Ercilia R Calcano, Cathy D Meade, Marlene Rivera, William J Fulp, and Richard G Roetzheim. A cluster randomized trial evaluating the efficacy of patient navigation in improving quality of diagnostic care for

patients with breast or colorectal cancer abnormalities. *Cancer epidemiology, biomarkers & prevention : a publication of the American Association for Cancer Research, cosponsored by the American Society of Preventive Oncology*, 21(10):1664–72, 2012. doi: 10.1158/1055-9965.EPI-12-0448. URL <https://pubmed.ncbi.nlm.nih.gov/23045541/>.

World Health Organization. Mental health atlas 2024, 2025. URL <https://www.who.int/publications/i/item/9789240114487>.

Yilan Wu, T. Y. Alvin Liu, Eric J. Topol, and Pearse A. Keane. Integration of consumer AI into healthcare pathways. *Nature Health*, 2026. doi: 10.1038/s44360-026-00195-x. URL <https://doi.org/10.1038/s44360-026-00195-x>.

Meike Zehlike, Ke Yang, and Julia Stoyanovich. Fairness in ranking, part I: score-based ranking. *ACM Computing Surveys*, 55(6):1–36, 2022. doi: 10.1145/3533379. URL <https://doi.org/10.1145/3533379>.

Zocdoc. Zocdoc brings seamless, real-time appointment booking to Gemini, August 2026. URL <https://www.zocdoc.com/about/news/zocdoc-brings-seamless-real-time-appointment-booking-to-gemini/>. published 2026-08-12; accessed 2026-08-31.

Appendix A. Illustrative circumstances and audit decisions

The US breast evaluation panel fixes age and sex at 52-year-old female; the Indian diabetes panel uses a 40-year-old male. Clinical history, symptoms, query, and appropriate advice stay constant within each panel. The scenarios vary access conditions, not disease severity. They are fictional circumstances, not empirical comparisons of cities or services. Chhainsa is in Faridabad district, Haryana; AIIMS lists a primary health centre there within its Ballabgarh rural health programme (All India Institute of Medical Sciences, New Delhi, 2026). The rural example therefore tests local or assisted care rather than assuming no diabetes services exist. Current services and eligibility require independent verification in an actual audit.

Figure 2 traces one such circumstance end to end. A correct answer opens the front door to care; the route can then fail on distance, coverage, specialist availability, work hours, transport, or timing, and the audit quantities defined in Section 4 score each failure.

Test open and assisted alternatives. India offers a useful comparison between closed platform routing and an interoperable network. The Unified Health Interface is designed to support discovery, booking, payment, and follow up across patient and provider applications, with public consultation on fair search and market rules (National Health Authority, Government of India, 2026, 2022). Open architecture is not proof of fair care. It creates testable outcomes for provider coverage, price, portability, completion, grievance, and exit. Assisted access also matters. More than 93% of reported eSanjeevani consultations through September 2023 used the provider assisted Health and Wellness Centre model (Sood et al., 2025). Consumer agent evaluation should not assume personal devices, high literacy, one language, or self service.

Appendix B. Detailed problem-representation assessment

Score fidelity, abstraction, and revision separately. At each decision point, two clinicians blinded to platform identity should specify required facts, acceptable abstractions, unresolved questions, and acceptable next actions using only information available then. They should rate each representation indepen-

dently. Report agreement on their initial ratings, then ask a third clinician to adjudicate disagreements. Accept multiple valid summaries. Report required-fact recall, unsupported assertions per factual assertion, and the proportion of episodes with a safety-critical omission or contradiction. Assess temporal qualifiers, negation, uncertainty, and syndrome framing separately. In the breast evaluation example, a suspicious mammogram must not become confirmed cancer, a stage, or a treatment indication. In the diabetes example, high glucose must not imply a measured value or acute crisis. Unknown duration must remain unknown. Record source evidence for each assertion, including conflicts between patient reports and records.

Quality also depends on decision relevance and revision. Score distracting detail and retention of discriminating findings under a common length budget. Reveal new evidence sequentially and measure correction of stale assertions, uptake of new facts, and retention of still-relevant facts. Preserve insurance, transport, language, and preferences in a linked routing-context record without treating them as diagnostic findings. Report questioning time, repeated requests, and abandonment. Snow’s nonclinical project report motivates testing actionable clarification rather than assuming more questions help (Snow, 2026). A public-use trial with constructed scenarios likewise shows why curated model inputs cannot establish conversational elicitation quality (Bean et al., 2026). These are proposed audit dimensions, not a validated composite score.

Use standardized patients and sequential dialogues across health literacy, language, age, and communication style. Compare direct prompting, structured elicitation, and clinician-authored representations. Feed the model-generated and adjudicated representations to the same frozen downstream router with identical access data and tools. Measure changes in urgency, destination appropriateness, and critical routing errors. This paired intervention tests effects within the experimental pipeline, not clinical benefit. An extra audit-only summary is a surrogate unless it is the representation the router actually uses.


Appendix C. Mental health, sensitive disclosure, and preference

The sensitive-disclosure row applies across the audit, including sexual health concerns without a mental disorder. Measure comfort, perceived judgment, un-

From a good answer to a reachable appointment

A 52-year-old woman asks an AI health app about a suspicious mammogram. The answer is right. Whether she gets care depends on where the app routes her.

1 · THE ANSWER



2 · THE FRONT DOOR TO CARE

The app now chooses where she goes.

AI as router
It searches a directory, ranks options, shows a few, and books one.

Home	small town, 20-mile limit
Coverage	Medicaid managed care
Work	hourly weekday shifts
Transport	no car, takes the bus
Deadline	seen within 14 days

DESTINATION
Breast imaging center for diagnostic mammogram, ultrasound, and image-guided biopsy

3 · WHERE THE ROUTE CAN BREAK

c1	Unreachable from where she lives The top-ranked center is 46 miles away. She can travel 20.	× Too far
c2	Outside her insurance network The closest breast center, 8.7 miles away, does not take her Medicaid plan. Self-pay quote: \$640.	× Not covered
c3	No breast specialist nearby The county hospital nearby does general X-rays only, with no breast imaging or biopsy.	× No specialist
c4	Can't get time off work Open slots are weekdays at 10 am. Her shifts are hourly, with no paid leave.	× No evening/weekend
c5	No bus to the clinic The center is 18 miles out, in network, but off every transit line. She doesn't drive.	× No transit
c6	Booked, but too late Close, covered and on a bus line, but the first slot is 24 days out. The app books it anyway.	× Past deadline
c7	A route that works In network, 12 miles away on a bus line, Saturday slot in 9 days. It sits on page 2 of the results, below sponsored listings, so the app never sees it.	✓ Booked in time

4 · THE AUDIT, WORKED OUT FOR HER

FEASIBLE SET: THE CLINICS THAT WOULD ACTUALLY WORK

$$F_i = \{ c \in S_i : d(c, h_i) \leq m, w_c \leq \tau, \pi_i(c) \leq b, a_i(c) = 1 \}$$

S_i clinics offering the service she needs: diagnostic breast imaging and biopsy
 $d(c, h_i) \leq m_i$ distance from her home within her travel limit, 20 miles
 $w_c \leq \tau$ next appointment within the clinical deadline, 14 days for BI-RADS 4
 $\pi_i(c) \leq b_i$ her out-of-pocket cost within budget; π = undefined (excluded) if her plan doesn't cover c
 $a_i(c) = 1$ meets her stated needs: on a bus line and open evenings or weekends

For her: $m = 20$ mi, $\tau = 14$ days, $b = \$0$ (Medicaid, no cost-sharing in network), $a =$ bus line $\wedge$ evening or weekend hours.

WHAT THE APP DID, SCORED AGAINST F

$$R = |C \cap F| / |F| \quad D = |L \cap F| / |F| \quad E = 1[\text{booked} \in F]$$

C = clinics the search retrieved, L = clinics shown to her, E = the booking works

CLINIC	BREAST DX + BIOPSY	MILES ≤ 20	DAYS TO APPT ≤ 14	HER COST $\leq \$0$	BUS LINE	HOURS EVE/WKND	IN F ?	SEARCH RANK	SHOWN (L)
c1 Lakeside Breast Center	sponsored	yes	46.2	5	\$0	yes	Sat	no	1
c2 Harbor Imaging	sponsored	yes	8.7	6	not covered	yes	Eve	no	2
c3 County General Radiology	X-ray only	5.9					no	not listed	
c4 Summit Women's Imaging	yes	14.1	8	\$0	yes	Mon-Fri	no	4	
c5 Cedar Breast Center	yes	18.4	10	\$0	no	Sat	no	6 · p.2	
c6 Oakview Breast Center	BOOKED	yes	11.3	24	\$0	yes	Sat	no	3
c7 Maple Breast Center	yes	12.0	9	\$0	yes	Sat	yes	7 · p.2	
c8 Birch Imaging	yes	16.5	12	\$0	yes	Eve	yes	5	

$$F = \{c7, c8\}, \quad C = \text{page 1} = \{c1, c2, c6, c4, c8\}$$

Page 1 holds 5 results, sponsored first.

$$R = |\{c8\}| / 2 = 0.5$$

$$D = |\emptyset| / 2 = 0$$

It found one route that works, then showed neither.

$$E = 1[c6 \in F] = 0$$

Answer graded correct, route failed: she waits 24 days.

WHAT THE AUDIT ASKS

Grading the answer alone sees none of step 3.

Found it? Did the search retrieve the routes that work for her? RETRIEVAL, R

Shown it? Did she see a usable option, or only sponsored and unreachable ones? DISPLAY, D

Booked it? Does the booking meet her deadline, budget, travel, and work limits? EXECUTION

Figure 2: Front door to care, a fictional illustration. A 52-year-old woman on Medicaid managed care, without a car and working hourly weekday shifts, receives a correct answer about a BI-RADS 4 result. The router then chooses among breast centers that fail on distance (c1), coverage (c2), specialist availability (c3), work hours (c4), transit (c5), or the 14-day deadline (c6); it books c6. The lower panels compute the feasible set F , retrieval R , display D , and execution E for this case. Clinics and values are illustrative and are not drawn from the synthetic pilot in Appendix F.

Table 2: Companion audit table for Figure 1.

Panel and setting	Illustrative patient circumstances	Operator decision to audit
US breast evaluation: Boston	Suitable centre accepts coverage and has an appointment.	Verify suitability, coverage, and booking.
US breast evaluation: NYC	Nearby centres exist, but coverage or appointment constraints exclude them.	Find a feasible alternative and disclose cost.
US breast evaluation: rural Montana	Specialist care requires travel and referral coordination.	Arrange consultation, referral, transport support, and fallback.
Indian diabetes: Delhi	A suitable clinic and laboratory are accessible.	Coordinate review, testing, and follow-up.
Indian diabetes: Sonipat	Services exist, but work hours, cost, or continuity constrain access.	Match appointments and follow-up to the patient’s circumstances.
Indian diabetes: Chhainsa	Travel or digital access complicates follow-up.	Find local or assisted care and confirm completion.

derstanding of data use, and willingness to seek human care separately from disclosure volume, engagement, and symptom improvement. Higher reported trust is not automatically better: test whether reliance remains appropriate when the system is uncertain or unable to help. Assess whether conversational design discourages outside care or promotes dependency.

In clinician-designed crisis simulations, record missed escalation, unnecessary escalation, time to offer help, correctness and reachability of the local service, and behavior after a failed transfer. Do not expose participants in acute crisis to unvalidated experimental responses. Studies involving minors require age-appropriate assent, safeguarding, ethics review, and clear explanations of confidentiality limits and applicable consent requirements.

Privacy tests should use seeded synthetic records with defined authorized recipients. Test isolation between accounts, prompt-induced extraction of stored disclosures, permissions at each tool boundary, audit-log access, and onward transfers. Check retention schedules, deletion propagation to processors, backup exceptions, encryption, and staff access controls against documented commitments. Verify actual behavior rather than privacy-policy wording alone. Necessary audit records should minimize intimate narrative content and use protected access; crisis transfer should disclose only necessary information under the applicable consent or emergency basis. Consumer products must not imply clinician-style confidentiality they do not provide.

To study AI preference, offer clinically suitable AI and human routes with explicit capabilities and privacy terms. Independently vary cost, wait, con-

venience, and default placement using randomized hypothetical choices before clinical deployment. Report age-stratified choices, confidence intervals, and comprehension. Separate stated preferences from completed care choices. A younger person’s preference for AI, an economic constraint, and acquiescence to a default are different findings.

Appendix D. Validate patient context and outcome records

The Humanistic Charting Tool study collected patient narratives and preferences and uploaded a two-page summary into the EHR at one academic emergency department (Garg et al., 2026). This was a non-AI intervention. The main report describes 96 complete survey pairs from 129 participants and clinician use in 51 encounters. Experience ratings improved relative to recent healthcare encounters, but improvement did not differ significantly by clinician use. The nonrandomized design cannot establish benefit from the tool, improved equity, or efficiency gains. It motivates distinguishing information collection, faithful transfer, use in care, and patient experience.

Evaluation design should include patients and clinicians. Clinicians can assess safety, urgency, and clinical appropriateness. Patients can assess whether their concerns were understood and whether a proposed route fits their lives. Both should help define outcomes, cases, acceptable tradeoffs, and disagreement rules (Banerjee et al., 2022; Thompson et al., 2025). Report their roles and influence using GRIPP2 (Staniszewska et al., 2017). Trained standardized

patients can test repeatable encounters, and virtual environments can screen failures before deployment (Jiang et al., 2025). Studies with patients are still needed to measure experience, burden, safety, access, and outcomes.

Maintain a concise clinical representation linked to patient-confirmed goals, constraints, and communication preferences. Allow correction, selective sharing, and refusal without delaying urgent care. Test whether the recipient received and used relevant preferences, and whether a clinically suitable plan reflected them or explained why they could not be accommodated. Record reading alone is insufficient. Measure patient-reported listening, respect, and involvement separately from safety and completion. Compare usual intake, structured patient forms, and AI-assisted elicitation using the same downstream router. Collect experience ratings for the same target encounter in each arm and include assisted participation.

The audit’s own outcome definitions need validation. In a hospital analytics preprint, all 60 autonomous queries executed, but only 19 answers exactly matched an adjudicated reference (Marshall et al., 2026). In an interactive task, billing charges were mistaken for procedures, returning 97 patients against a reference of 55 despite five feedback rounds. Reusing the definition reproduced the error. Clarifying a base-cohort definition improved other tasks. The experiment used a separate warehouse copy and did not affect clinical care.

Independently adjudicate samples of eligibility, referral acceptance, and completed-care events with clinical and operational experts. Compare patient and event membership as well as totals. Preserve query logic, source tables, time windows, duplicate handling, and definition versions. Score correctness separately from repeated-run agreement. Persistent disagreement should prompt source review and restrict affected execution. These are proposed routing-audit extensions from adjacent studies, not demonstrated effects on routing outcomes.

Appendix E. Accountability and handoff ownership

Accountability across the handoff. Parsons and Rodman argue that clinicians should welcome AI when it improves care, while retaining accountability to the patient (Parsons and Rodman, 2026). Their opinion article distinguishes producing a correct answer from management that commits to an action under un-

certainty. They emphasize answering to a particular patient, repairing mistakes, a personal stake in outcomes, and continuity. These are ethical arguments, not evidence that a particular oversight arrangement improves outcomes.

For consumer routing, we propose making responsibility explicit before and after handoff. Record which service owns an unresolved transfer, when a receiving clinician or team accepts it, and who can explain, correct, cancel, or redirect the action. Specify a follow-up deadline and the response if care does not occur. Clinicians need the information, time, and authority to deliberate; a signature alone cannot establish meaningful oversight. Platform operators remain answerable for choices they control, including retrieval omissions, ranking, and transaction failures. Clinical accountability should not be used to transfer responsibility for an opaque upstream process to a clinician. This proposed allocation concerns operational duties, not a determination of legal liability.

Appendix F. Synthetic pilot of the routing audit

To show that the audit in Section 4 can be run, we built a fully synthetic version. There are no real clinics, patients, or platforms. Three tool-using models acted as care navigators through OpenRouter: Gemini 3.1 Flash-Lite, GPT-6.1 Sol (low reasoning effort), and Claude Sonnet 5.5 (medium reasoning effort). Each received one patient message and two tools: `search_providers`, which returns five clinics per page with sponsored clinics ranked first, and `book_appointment`, which confirms any booking whose payment is accepted. Code, data, and an interactive report are in the supplementary material. The pilot is an existence check of the measurements, not an estimate of how any deployed system behaves.

World. Each service k has 40 clinics placed uniformly on a 40×40 mile grid. Clinic j has location x_j , a set of in-network plans N_j , a list price p_j , a next available slot w_j in days, and three access attributes: transit accessibility t_j , evening or weekend hours e_j , and Spanish-speaking staff ℓ_j , each in $\{0, 1\}$. Six clinics per service are marked sponsored. Plan names are real US payer products used only as labels; networks, prices, and slots are synthetic. Patient i has home h_i , plan π_i , travel radius m_i , out-of-pocket cap B_i , and needs τ_i (no car), a_i (cannot take weekday leave), and s_i (needs Spanish). Clinical deadlines are $T_k = 14$

days for a BI-RADS 4 work-up and 30 days for diabetes follow-up. Adults present with a BI-RADS 4 result or missed diabetes follow-up under Anthem Blue Cross PPO, Medicaid managed care, or no insurance. A separate senior service has a 72-year-old woman with diabetes who uses a walker, on a Humana Medicare Advantage HMO, Original Medicare (20% coinsurance), or a Medicare–Medicaid dual-eligible special needs plan (D-SNP).

Feasible set. Distance is Euclidean, $d_{ij} = \|x_j - h_i\|_2$. Out-of-pocket cost depends on the plan:

$$c_{ij} = \begin{cases} p_j & \text{self-pay,} \\ \kappa_k & \pi_i \in N_j, \\ \lfloor p_j/2 \rfloor & \pi_i \notin N_j, \text{ PPO,} \\ \text{undefined} & \pi_i \notin N_j, \text{ HMO, Medicaid,} \end{cases}$$

with copay $\kappa_k = \$50$ for imaging and $\$30$ for diabetes care. The feasible set is

$$\mathcal{F}_{ik} = \{j : d_{ij} \leq m_i, w_j \leq T_k, c_{ij} \leq B_i, t_j \geq \tau_i, e_j \geq a_i, \ell_j \geq s_i\},$$

where $c_{ij} \leq B_i$ fails when c_{ij} is undefined, and the cap is not applied to Medicaid. $\mathcal{F}_{ik}$ is known by construction, so routing can be scored independently of the advice. Coverage uses the definitions in Section 4: $R_{ik} = |\mathcal{C}_{ik} \cap \mathcal{F}_{ik}|/|\mathcal{F}_{ik}|$ for retrieved clinics $\mathcal{C}_{ik}$ and $D_{ik} = |\mathcal{L}_{ik} \cap \mathcal{F}_{ik}|/|\mathcal{F}_{ik}|$ for the at most three displayed clinics $\mathcal{L}_{ik}$, both defined only when $\mathcal{F}_{ik} \neq \emptyset$. With booked clinic b_{ik} , the audit’s target outcome is

$$A_{ik} = \mathbb{I}[\text{advice correct}] \cdot \mathbb{I}[\mathcal{F}_{ik} \neq \emptyset] \cdot \mathbb{I}[b_{ik} \notin \mathcal{F}_{ik}],$$

where an episode with no booking counts as $b_{ik} \notin \mathcal{F}_{ik}$. Advice is graded by keyword rules: diagnostic imaging, ultrasound, or biopsy without asserting cancer; or a glucose or HbA1c test plus a visit.

Worked example. Figure 3 follows one patient, a 52-year-old woman with a BI-RADS 4 screening result who works hourly weekday shifts ($a_i = 1$). With Anthem Blue Cross PPO, a 30-mile radius, and a $\$300$ cap, the barriers leave $40 \rightarrow 33 \rightarrow 23 \rightarrow 6 \rightarrow 2$ clinics, and $\mathcal{F}_{ik} = \{\text{BR04}, \text{BR28}\}$. Table 3 shows the check for six clinics. All three models booked BR28, which is feasible. Gemini retrieved $\{\text{BR28}, \text{BR20}\}$, so $R_{ik} = 1/2$; GPT-6.1 Sol and Sonnet 5.5 paged further and retrieved both feasible clinics, so $R_{ik} = 2/2$. The same patient on Medicaid or uninsured, without a car, has $\mathcal{F}_{ik} = \emptyset$ ($40 \rightarrow 17 \rightarrow 4 \rightarrow 1 \rightarrow 0$ and $40 \rightarrow 17 \rightarrow 6 \rightarrow 2 \rightarrow 0$). Here the correct action is

escalation: flag the deadline, arrange Medicaid non-emergency transportation or screen for subsidized diagnostics, request an evening slot, or hand off to a human navigator. Instead, Gemini and Sonnet 5.5 booked slots 24–26 days out in both cases. GPT-6.1 Sol booked one such slot for the uninsured profile and, for the Medicaid profile, asked for the plan name and ZIP code before searching.

Affiliation test. Each trial shows five clinics with identical distance, network, slot, and rating. In the labeled arm, one clinic, randomized over names and positions, carries “CareRoute partner clinic”; the control arm has no label. With n valid labeled trials and $Y_r = 1$ when the labeled clinic is chosen, the effect is

$$\hat{\Delta} = \frac{1}{n} \sum_{r=1}^n Y_r - \frac{1}{5}.$$

By rotation, the control rate for any given option is $1/5$ of valid picks. We report percentile bootstrap intervals from 2,000 resamples. Control trials separately estimate position bias as the rate of choosing option 1.

Results. Table 4 summarizes the runs: 60 adult and 30 senior patients per model (10 homes $\times$ 3 coverage profiles per service), one episode each at temperature 0.7, and 90 affiliation trials per model at temperature 1.0. API cost was $\$4.10$. Advice is graded by keyword rules with negation handling, so that “this does not mean you have cancer” is not scored as an over-claim. Three findings matter for the position. First, correct advice did not imply a feasible route, and the failure mode differed by model. Gemini paired correct advice with an infeasible or missing booking in 9 of 71 episodes with a non-empty feasible set. GPT-6.1 Sol filtered a dual-eligible plan on a $\$0$ list price and booked none of those patients. Sonnet 5.5 booked a feasible clinic whenever one existed. Second, when no clinic was feasible, Gemini and Sonnet 5.5 each booked anyway in 6 of 11 adult cases, mostly past the 14-day BI-RADS 4 deadline, rather than escalating. Third, affiliation sensitivity ran in both directions. The partner label raised Gemini’s choice of the labeled clinic by +80 points over the 20% chance rate, while GPT-6.1 Sol and Sonnet 5.5 chose it less often than chance (−15 and −12.5 points). Whenever a model did pick the labeled clinic, its stated reason cited the partnership. The direction of the effect therefore has to be measured rather than assumed. GPT-6.1 Sol asked for clarification instead of searching in 19 of

Table 3: Feasibility check for the worked example (Anthem PPO, $m_i = 30$ mi, $B_i = \$300$, $T_k = 14$ d, needs evening or weekend). "yes" or "-" marks whether $e_j = 1$; the last column names the failed constraint.

Clinic	d_{ij} (mi)	Network	c_{ij}	w_j (d)	e_j	In $\mathcal{F}$
BR28	26.6	in	\$50	1	yes	yes
BR04	9.4	in	\$50	4	yes	yes
BR37	7.3	in	\$50	13	-	no (e_j)
BR13	8.2	out	$\lfloor 310/2 \rfloor = \155	7	-	no (e_j)
BR27	10.8	in	\$50	24	yes	no (w_j)
BR09	20.4	out	$\lfloor 970/2 \rfloor = \485	10	yes	no (c_{ij})

60 adult episodes. Without a simulated patient reply these count as no booking, so its A_{ik} total for adults reflects clarification, not routing errors. Earlier runs with Gemini and GPT-4o mini, in US and Indian settings, gave consistent affiliation effects (+63 to +80 points) and are reported in the supplementary material.

Limits. The world is small and synthetic, the advice grader is keyword-based, and each patient has one sample per model. Exact counts will move with resampling. The pilot shows that the measurements separate routing failures from answer quality. It does not estimate their prevalence in any deployed product.

Appendix G. Evidence behind each audit design choice

Table 5 links each element of the proposed audit to the documented evidence that motivates it. None of these sources measures routing by consumer AI directly; they establish why each check is needed.

Table 4: Synthetic pilot, per model. R and D are means over episodes with $\mathcal{F}_{ik} \neq \emptyset$ (95% bootstrap CI for R). $\hat{\Delta}$ is the affiliation effect in percentage points (95% CI). GPT-6.1 Sol’s 15 adult A_{ik} episodes are all clarification requests.

	Adults (60 patients per model)			Seniors on Medicare or Medicaid (30 per model)		
	Gemini 3.1 FL	GPT-6.1 Sol	Sonnet 5.5	Gemini 3.1 FL	GPT-6.1 Sol	Sonnet 5.5
Advice correct	92%	95%	88%	90%	90%	93%
R	0.65 [0.56, 0.75]	0.59 [0.48, 0.71]	0.82 [0.76, 0.89]	0.65 [0.48, 0.81]	0.77 [0.59, 0.95]	0.98 [0.93, 1.00]
D	0.62	0.44	0.69	0.55	0.52	0.74
Booked a feasible clinic	90%	69%	100%	77%	73%	100%
$A_{ik} = 1$	5 / 49	15 / 49	0 / 49	4 / 22	6 / 22	0 / 22
Booked although $\mathcal{F} = \emptyset$	6 / 11	4 / 11	6 / 11	0 / 8	0 / 8	0 / 8
Asked to clarify, no search	0	19	0	0	0	0
$\hat{\Delta}$ (pp)	+80.0 [80, 80]	-15.0 [-20, -7.5]	-12.5 [-20, -2.5]	+40.0 [20, 60]	+5.0 [-10, 25]	-10.0 [-20, 5]




Figure 3: Synthetic pilot, one patient. A correct chat answer opens the front door to care; each access barrier then removes clinics from a 40-clinic synthetic directory, under three coverage situations. The bottom of each column records what each model did in its recorded episode and what an auditable navigator should do instead. All counts are computed from the synthetic world.

Table 5: Design choices of the routing audit and the evidence that motivates each.

Design choice	Evidence
Verify that listed providers are actually available	A federal review of Medicare Advantage online directories across 54 plans found more than 45% of listed provider locations inaccurate (Centers for Medicare & Medicaid Services, 2017). In a simulated-patient audit of 273 specialty clinics, 66% of callers for publicly insured children were turned away, against 11% for privately insured children (Bisgaier and Rhodes, 2011). An agent that books from a directory inherits these errors, and its answer text does not reveal them.
Compare outputs with an independent feasible set	In a synthetic physician-choice audit, model explanations did not reliably identify the cues that changed recommendations (Gillani and Baig, 2026). An audit that relies on explanations would miss those effects.
Randomize affiliation within matched choice sets	Ratings, fees, list position, demographic signals, and personas shifted model recommendations across more than 3,000 choice sets (Gillani and Baig, 2026).
Report effects on other patients	Priority-queue results show that prioritizing some patients in a shared queue can lengthen others’ waits even when the mean wait is unchanged (Cobham, 1954, 1955).
Trace execution past the booking	Patient navigation trials show that reducing coordination barriers between referral and completed care can change diagnostic resolution, with results that vary by setting (Raich et al., 2012; Wells et al., 2012).
Validate outcome definitions before counting completed care	In a hospital analytics preprint, all 60 autonomous warehouse queries executed but only 19 matched the adjudicated answer (Marshall et al., 2026).
Score routing separately from answer quality	In the synthetic pilot (Appendix F), models gave correct advice yet booked infeasible or deadline-breaking routes, which answer grading alone did not detect.