AI-Regulation Papers 2026-03-3

The Health AI Agent Rush

Five Companies, Your Health Data,

and the Governance Questions Nobody Is Asking

Theodore Christakis

Chair Responsible AI - Legal

March 2026

Theodore Christakis is Professor of International, European and Digital Law at University Grenoble Alpes (France), Director of Research for Europe with the Cross-Border Data Forum, Member of the Board of Directors of the Future of Privacy Forum and a former Distinguished Visiting Fellow at the New York University Cybersecurity Centre. He is co-Chair with Héber Arcolezi on "Responsible AI: Design, Regulation and Conformity" at the Multidisciplinary Institute in Artificial Intelligence.

Acknowledgments:
The author would like to thank Shadée Pinto, Research Fellow with AI-Regulation.com, for her careful reading and very useful comments. Any remaining errors are the author's own.

Cover Illustration:
Original illustration conceived and directed by the author and produced with generative AI assistance.
Artistically inspired by the surrealist aesthetic associated with René Magritte.

To cite this study: T. Christakis, "The Health AI Agent Rush: Five Companies, Your Health Data, and the Governance Questions Nobody is Asking", AI Regulation Papers, 26-03-3, AI-Regulation.com, March 2026.

These statements are attributable only to the author, and their publication here does not necessarily reflect the view of the other members of the AI-Regulation Chair or any partner organisations.

This work has been partially supported by MIAI @ Grenoble Alpes, (ANR-23-IACL-0006) and by the Interdisciplinary Project on Privacy (IPoP) of the Cybersecurity PEPR (ANR 22-PECY-0002 IPOP).

EXECUTIVE SUMMARY

THE HEALTH AI AGENT RUSH

Five Companies, Your Health Data, and the Governance Questions Nobody Is Asking

Theodore Christakis

Between January and March 2026, five major technology companies (OpenAI, Anthropic, Microsoft, Amazon, and Perplexity) launched health-specific AI products connecting medical records, wearables, and wellness apps to their chatbot platforms. Every one followed the same playbook. Every one is restricted to the United States. This article is the first comparative academic analysis of what is arguably the most consequential AI development of 2026 that has not yet received serious academic or regulatory scrutiny.

Context. This follow-up builds on my March 2026 study *You Trust Your Chatbot With Everything. Should You? Part 1: How The Controller Uses Your Chat Data*, which proposed Sealed Mode: a confidentiality framework for sensitive chatbot conversations built into product architecture (no training, no advertising, siloed personalisation, strict retention, minimised human review, cryptographic hardening). The article analyses all five health AI products against Sealed Mode's six components through a comparative table covering sixteen dimensions, and maps the broader data protection/governance questions this convergence creates.

Key Findings and Questions the Article Raises

1. The market has validated Sealed Mode's core intuition. Five companies independently concluded that health conversations cannot be governed by the same defaults as ordinary chat. But every product has converged on health data integration hubs (connecting records, apps, wearables) rather than on protecting the conversational disclosure that hundreds of millions of users already make without connecting anything.

2. The European exclusion reveals a missed opportunity. Every product bundles privacy protections with data integration features. It is the integration dimension that triggers regulatory hurdles in Europe (GDPR Article 9, the Medical Device Regulation, the AI Act). Providers could have offered the privacy-protective dimension as a standalone feature globally. By bundling both, they excluded Europeans from both. The paradox: **the users most protected by data protection law are the ones denied access to the most privacy-protective chatbot feature currently available.**

3. None of the five products meets the full Sealed Mode standard. All five offer no-training commitments and data isolation. But other protective elements are missing. And the governance frameworks remain under each provider's unilateral control, and have not been tested by time, commercial pressure, or independent scrutiny.

4. There is no free lunch. Several of these products are free or bundled with subscriptions. But the strategic logic is clear: Amazon channels users to One Medical and Amazon Pharmacy; OpenAI deepens engagement on a platform it is simultaneously monetising through advertising; Perplexity and Anthropic restrict health features to paying subscribers. Once a user connects years of medical records and wearable data to one platform, switching costs become

prohibitive. The question is whether privacy commitments made during the trust-building phase of launch will survive the monetisation phase that follows.

5. The gold mine question. “Not used to train foundation models” does not mean “not used.” It does not preclude product analytics, quality improvement, or what Amazon calls training on “abstracted patterns.” These platforms could constitute the largest aggregation of health data in history. Under US law, since these consumer products are not HIPAA-covered entities, there is no federal constraint on secondary use beyond each company’s voluntary commitments. Privacy policies can be modified at any time.

6. Private health hubs vs. the European Health Data Space. The EU has spent years building the EHDS, a public infrastructure for secondary use of health data with strict institutional gatekeeping. A private AI platform aggregating consumer health data under consent could derive comparable epidemiological and pharmaceutical insights without passing through that governance. Even under GDPR, this structural asymmetry would persist. It risks creating a two-track system where the most demanding governance requirements apply to public institutions and academic researchers, while the largest health data aggregations sit in private hands under less structured, consent-based frameworks..

7. Cybersecurity concentration risk. Five companies are centralising health data linked to personal identities, medical records, and years of behavioural patterns. A single breach could expose health data at a scale never seen before. Security architectures remain undisclosed; third-party intermediaries (b.well, HealthEx, Terra API) each add a point of vulnerability.

8. Structural tensions with core GDPR principles. The architectural choices behind these products raise questions under several foundational principles of EU data protection law, including data minimisation, storage limitation, purpose limitation, controller and processor complexity, data portability, and the protection of children's data.

* * *

These questions are not reasons to block health AI products whose potential to help people better understand and manage their health is real. They are reasons to get the governance right before the gap between the pace of product launches and the pace of oversight becomes impossible to close.

The article also acknowledges that data integration in the health context is not driven solely by business logic; it can serve a genuine clinical function, since generic health guidance that ignores a user’s medical history, medications, or pre-existing conditions may be not merely unhelpful but actively harmful. This reinforces rather than undermines the case for strong governance: the more data these systems ingest, the more consequential their outputs become.

Policymakers and regulators in Europe should work proactively, and in dialogue with providers, to find constructive and protective solutions. One avenue worth exploring is the use of regulatory sandboxes, already provided for under the AI Act, to allow health AI products to operate in Europe under supervised conditions with strong, verifiable privacy safeguards, including the kind of architectural protections that Sealed Mode envisions. The current outcome, where Europeans are excluded entirely, serves neither innovation nor protection.

The real question is no longer whether differentiated privacy for sensitive chatbot conversations is conceivable. Five companies answered that in three months. The question is whether the privacy-protective core can be extracted from the integration products it is bundled with and offered as a standalone standard, available to every user, everywhere.

THE HEALTH AI AGENT RUSH

Five Companies, Your Health Data, and the Governance Questions Nobody Is Asking

Theodore Christakis

Introduction

My recent study, *You Trust Your Chatbot With Everything. Should You?*, proposed Sealed Mode as a confidentiality framework for especially sensitive chatbot conversations: a clearly identified consumer pathway with stronger defaults - no training, no advertising, minimised human review, siloed personalisation, strict retention, and cryptographic hardening - protections embedded in product architecture rather than left to policy language.[1]

The study was drafted between December 2025 and late February 2026 and published in early March 2026. Shortly after publication, I discovered a pilot initiative that had been announced during the very period in which I was drafting the study and developing the Sealed Mode proposal: OpenAI's ChatGPT Health, launched on 7 January 2026.[2] ChatGPT Health creates a dedicated, compartmentalised health space inside ChatGPT with additional privacy protections, including a no-training commitment and data isolation[3]. As this follow-up article will show, however, OpenAI's initiative was driven by a different objective than Sealed Mode: its primary aim was to build a health data integration hub - a "trusted AI health agent" connecting medical records, wellness apps, and wearables - rather than to create a privacy-protective pathway for ordinary health conversations.[4] The privacy protections are instrumental to the connectivity objective, not the objective itself.

This difference in starting point also explains why ChatGPT Health was launched as an experimental programme available only through a staged waitlist, and explicitly excluded from the European Economic Area, Switzerland, and the United Kingdom[5] - the jurisdictions where I work and where the study was written.[6] The European exclusion is not driven by the privacy protections (adding stronger privacy defaults to conversations that already take place would create no new regulatory hurdle). It is driven by the data integration dimension: the ingestion of structured health data from medical records and wellness apps triggers Article 9 GDPR special-category requirements, raises questions under the EU Medical Device Regulation if the

[1] T. Christakis, You Trust Your Chatbot With Everything. Should You? Part 1: How The Controller Uses Your Chat Data, AI Regulation Papers, 26-03-2, AI-Regulation.com, March 3, 2026. See also the companion article: T. Christakis, "New Study Maps the Privacy Gap in Consumer AI — and Proposes a Fix", IAPP, 4 March 2026.

[2] OpenAI, "Introducing ChatGPT Health", 7 January 2026.

[3] *Ibid*. See also OpenAI Help Center, "What is ChatGPT Health?".

[4] See CNBC, "OpenAI launches ChatGPT Health", 7 January 2026; Fortune, "OpenAI launches ChatGPT Health in a push to become a hub for personal health data", 7 January 2026; TechCrunch, "OpenAI unveils ChatGPT Health, says 230 million users ask about health each week", 7 January 2026.

[5] OpenAI Help Center (n 3). ChatGPT Health is available for users on Free, Go, Plus and Pro plans "in all our supported countries (except EEA, CH, and the UK)". Medical record integration (via b.well) is US-only and requires users to be over 18. As of late March 2026, no public timeline for European availability has been announced.

[6] See Euronews, "OpenAI launches dedicated ChatGPT Health feature with medical record integrations", 8 January 2026. ("Notably excluded from the early rollout are users in the European Economic Area, Switzerland and the United Kingdom, where local health and data regulations are more robust"); Heise Online, "ChatGPT Health: OpenAI launches AI health assistant", 8 January 2026. ("The exclusion of the EU, Switzerland, and the UK points to regulatory hurdles").

software is deemed to have a medical purpose, and may engage the AI Act's high-risk classification for healthcare AI systems (see analysis in Section 1.4 below).

The discovery of this pilot program prompted me to do two things. First, to examine ChatGPT Health closely and assess how it compares with the Sealed Mode framework I proposed. Second, to research systematically whether any of the other four chatbot providers analysed in my study - Anthropic (Claude), Google (Gemini), xAI (Grok), and DeepSeek -, or, indeed, *any* other chatbot provider, had launched comparable pilot initiatives for sensitive conversations.

My research revealed that ChatGPT Health was only the first in what has rapidly become an industry-wide movement. During these last weeks, Anthropic (Claude for Healthcare), Microsoft (Copilot Health), Amazon (Health AI), and Perplexity (Perplexity Health) all launched comparable health-specific AI products, each following a strikingly similar pattern of health data integration, privacy commitments, and US-only availability. The scale and speed of this convergence transforms the context for the Sealed Mode proposal: it is no longer a question of whether one company will differentiate health conversations, but of whether an entire industry's approach to that differentiation addresses the right problem.

The result is this follow-up article, which sits between Part 1 and the forthcoming Part 2 of the study. It does not replace or delay Part 2 (which will examine the external boundary: civil discovery, government-compelled access, and breach exposure); it responds to market developments that emerged too quickly to wait. It also identifies a set of broader data protection questions that the health agent rush leaves open, from the secondary use of what may become the world's largest health dataset to the relationship between private health AI hubs and the European Health Data Space.

1. OpenAI's ChatGPT Health: a close analogue, but with important differences

Although the ChatGPT Health pilot was launched in the U.S. only a few weeks ago, OpenAI has stated that it had been laying the groundwork for ChatGPT Health for approximately two years, drawing on the work of over 260 physicians across 60 countries.[7] Let's examine the parallels and differences with the "Sealed Mode" proposal and also enquire why ChatGPT Health is not available (not even for testing) in Europe.

1.1 The parallels

ChatGPT Health is the closest existing market analogue to Sealed Mode, and the surface-level parallels are striking. Both start from the recognition that health conversations deserve a separate, more protective space inside the AI chatbot. Both commit to not training on health conversations. Both isolate health-specific data from the ordinary chat experience. Both create a distinct memory system that does not does not merge into the main AI chatbot experience . If a user starts a health-related conversation in the ordinary ChatGPT interface, the system nudges them to switch to the Health space for additional protections.[8]

[7] Fortune (n 4). OpenAI hired Nate Gross (co-founder and former chief strategy officer of Doximity) to lead healthcare strategy and Ashley Alexander (former co-head of product at Instagram) to lead healthcare product. Karan Singhal, who leads health AI at OpenAI, stated that the company had "been laying the groundwork for ChatGPT Health for about two years".

[8] OpenAI, "Introducing ChatGPT Health", 7 January 2026 ("If you start a health-related conversation in ChatGPT, we'll suggest moving into Health for these additional protections"). See also OpenAI Help Center, "What is ChatGPT Health?".

1.2 A fundamentally different starting point

Beneath these parallels, however[9], the two approaches start from fundamentally different places and the difference matters.

ChatGPT Health's primary framing is connectivity and utility. The pilot is designed as a health data integration hub or, in OpenAI's own framing, "another step toward turning ChatGPT into a personal super-assistant"[10] and creating "AI as a Healthcare Ally".[11] Users can securely connect medical records (via a partnership with b.well for US electronic health records), wellness apps (Apple Health, MyFitnessPal, Weight Watchers, Peloton, AllTrails, Function, Instacart), and uploaded files, so that ChatGPT's responses are grounded in the user's own health data.[12] The product logic follows what might be called a *health agent model*: the leading proposition is what the system can *do* with the user's health data: interpret test results, prepare for appointments, suggest dietary adjustments, compare insurance options. The privacy protections (no training, data isolation, purpose-built encryption) *are the necessary foundation that makes this connectivity safe*, but they are instrumental to the connectivity objective, not the objective itself. As one commentator noted, AI firms are "leaning hard" into finding ways to bring more personalisation to their services to boost value.[13]

As this article will show, this health agent logic is not unique to OpenAI. Every subsequent entrant (Anthropic, Microsoft, Amazon, Perplexity…) has adopted the same framing: the product's value proposition is what it can do with the user's health data, and the privacy protections are the necessary foundation that makes the connectivity safe.

Sealed Mode's starting point is the opposite. It begins not from connectivity but from the user's conversational disclosure, for instance the moment someone types "I'm terrified it might be something serious" or "I haven't told my wife because she's already stressed". The proposal is rooted in the observation, documented extensively in my study, that hundreds of millions of people already confide health conditions, fears, and emotional distress to chatbots through ordinary conversation, without uploading any medical record or connecting any app. The privacy risk does not arise from data integration; it arises from the conversation itself. Sealed Mode is designed to protect that conversational intimacy by constraining what the system and its operators can do with it downstream.

This difference in starting point has practical consequences. ChatGPT Health is designed for users who affirmatively want to engage with their health data through an AI assistant. Sealed Mode is designed to protect users who may not even realise how much they are disclosing in

[9] Several other protections not directly related to the six components of my my "Sealed Mode" proposal appear in Open AI's notice. For instance, toggling to Health for the first time triggers a consent flow for users. The notice also links to a dedicated ChatGPT Health Help Centre article, and provides users with equally prominent options to either "Accept" or "Decline" use of Health mode. The notice also contains a disclosure informing users that Health does not provide medical advice or replace healthcare professionals.

[10] TechCrunch (n 4). OpenAI's CEO of Applications, Fidji Simo, described the feature as "another step toward turning ChatGPT into a personal super-assistant". See also MobiHealthNews, "OpenAI launches ChatGPT Health, partners with b.well", 7 January 2026.

[11] OpenAI explains that more than 5% of all ChatGPT messages globally are about healthcare, averaging billions of messages each week. More than 200 million Open AI distinct global regular users submit a prompt about healthcare every week. More than 40 million turn to ChatGPT every day with healthcare questions. See Open AI, "AI as a Healthcare Ally", January 2026.

[12] OpenAI (n 2). See also BBC, "ChatGPT launches health feature, raising new questions about data safety", January 2026.

[13] See A. Crawford (Center for Democracy and Technology), quoted in BBC (n 10): "Since it's up to each company to set the rules for how health data is collected, used, shared, and stored, inadequate data protections and policies can put sensitive data at risk". See also TIME, "Is Giving ChatGPT Health Your Medical Records a Good Idea?", 9 January 2026; The Register, "ChatGPT Health wants access to sensitive medical records", 8 January 2026.

what feels like a private conversation. The two approaches are *complementary*, not contradictory – but they address different populations and different risk profiles.

1.3 Other differences

Beyond this foundational difference in starting point, two further gaps between ChatGPT Health and the full Sealed Mode framework stand out.

First, scope. ChatGPT Health is limited to health. Sealed Mode was proposed as a broader framework for high-stakes contexts, including crisis-adjacent conversations, legal consultations, and other disclosures where confidentiality expectations are acute. Health is the right starting point (my own study proposes it as the first domain) but the question is whether the principle can/will extend beyond it.

Second, transparency about the internal boundary. OpenAI's public materials describe meaningful protections, but they do not yet articulate the full level of specificity that a robust Sealed Mode framework would imply: a publicly documented retention duration for sealed conversations, explicit constraints on human access with audit trails, and longer-term guardrails against monetisation drift. This last point is particularly relevant given that OpenAI has simultaneously begun testing advertising in the U.S. in the standard ChatGPT experience, with personalisation enabled by default based on past chats and memories.[14] OpenAI states that ads do not appear in chats about sensitive topics including health, but the structural proximity between an ad-supported product and a health-data-rich compartment makes the question of durable boundaries more pressing, not less.

1.4 Why ChatGPT Health is not available in Europe – and what this reveals

ChatGPT Health remains an experimental programme distributed through a staged waitlist. It is explicitly not available in the European Economic Area, Switzerland, or the United Kingdom.[15] OpenAI has not provided a specific public explanation, but the exclusion may be widely attributed to regulatory hurdles in these jurisdictions.[16] Understanding which regulatory hurdles are at play is important, because it illuminates the nature of the product and, by contrast, the nature of Sealed Mode.

One obstacle is Article 9 GDPR, which classifies "data concerning health" as a special category of personal data whose processing is in principle prohibited unless one of the conditions listed in Article 9(2) is met. Where a chatbot actively ingests structured health data from external sources (medical records, wellness apps, wearable devices…) the processing almost certainly falls within this regime, requiring either explicit consent under strict conditions or another applicable derogation. This creates compliance obligations, including the drafting of a DPIA, significantly more demanding than those applicable to ordinary conversational use (see also below Section 6.5).

[14] ChatGPT Release Notes. OpenAI states that during the advertising test, "ads do not appear in chats about sensitive topics or regulated topics, including health, mental health, or politics". Check also this freshly published OpenAI, "Ad Policies".
[15] See n5.
[16] See Euronews, 'OpenAI launches dedicated ChatGPT Health feature', 8 January 2026 ('Notably excluded from the early rollout are users in the European Economic Area, Switzerland and the United Kingdom, where local health and data regulations are more robust'); Heise Online, 'ChatGPT Health: OpenAI launches AI health assistant', 8 January 2026 ('The exclusion of the EU, Switzerland, and the UK points to regulatory hurdles').

The EU Medical Device Regulation (MDR) could apply if ChatGPT Health were classified as software with a "medical purpose".[17] Unlike the United States, where certain clinical decision support software functions may fall outside the FDA's device regime if they meet the narrow exclusion in section 520(o)(1)(E) of the Federal Food, Drug, and Cosmetic Act (FD&C Act),[18] the EU MDR contains no equivalent general carve-out for decision-support software. In Europe, if the system is used to detect medically relevant patterns in health data, generate patient-specific suggestions that may inform diagnosis or treatment, or otherwise perform a diagnostic, monitoring, or prognostic function, regulators may argue that it falls within the MDR's scope.[19] The decisive issue under the MDR is the manufacturer's intended purpose. Accordingly, disclaimers that the tool is "not intended for diagnosis or treatment" would be relevant, but not necessarily conclusive, if the system's actual functions, outputs, and presentation indicate a medical purpose.

The AI Act[20] adds a further layer: certain AI systems used in health-related contexts may qualify as high-risk under the EU AI Act, triggering conformity assessment obligations including risk management, data governance, human oversight, and registration requirements. The combined weight of GDPR Article 9, the MDR, and the AI Act creates a regulatory environment that is significantly more complex for a health data integration product than for an ordinary chatbot.

This analysis reveals something important about the relationship between ChatGPT Health and Sealed Mode. The regulatory obstacles that block ChatGPT Health in Europe are driven almost entirely by the data integration dimension: the ingestion of medical records, the connectivity with wellness apps, the cross-border transfer of structured health data. They are *not* driven by the privacy protections that ChatGPT Health shares with Sealed Mode (no training, data isolation, compartmentalised memory). On the contrary: adding stronger privacy defaults to health conversations that already take place every day on ChatGPT would create no new regulatory hurdle. It would, if anything, bring the service closer to GDPR compliance, not further from it.

This points to a **missed opportunity**. OpenAI could have proceeded in **two steps**. *First*, offer a Sealed Mode for health conversations globally, the same conversations that 230 million people already have on ChatGPT every week,[21] but with stronger default protections (no training, isolated memory, no advertising, minimised human access). This step would have been available everywhere, including Europe, since it only adds protections to existing functionality. *Second*, layer the medical record and app integrations on top, in jurisdictions where regulatory compliance permits. Instead, by bundling both into a single product, OpenAI ended up excluding Europeans from both the privacy protections and the data integrations. **The result**

[17] Regulation (EU) 2017/745 of the European Parliament and of the Council of 5 April 2017 on medical devices (MDR).
[18] See FDA, Clinical Decision Support Software, January 2026.
[19] Regulation (EU) 2017/745 expressly includes "software" within the definition of a medical device where it is intended by the manufacturer to be used for a medical purpose; the listed purposes include "diagnosis, prevention, monitoring, prediction, prognosis, treatment or alleviation of disease" and the "investigation … of a physiological or pathological process or state" (Art. 2(1)). The MDR further defines "intended purpose" by reference to the manufacturer's statements on the label, in the instructions for use, in promotional or sales materials, and in the clinical evaluation (Art. 2(12)). Annex VIII, Rule 11 then specifically addresses software intended to provide information used to take decisions for diagnostic or therapeutic purposes, and software intended to monitor physiological processes. See also Medical Device Coordination Group (MDCG), Guidance on Qualification and Classification of Software in Regulation (EU) 2017/745 (MDR) and Regulation (EU) 2017/746 (IVDR), MDCG 2019-11 Rev. 1, at 4, 8, 17-18 (June 2025) (explaining that software must have a medical purpose on its own to qualify as MDSW; that qualification depends on the manufacturer's intended purpose; and that software intended to provide information used for diagnostic or therapeutic decisions, or to monitor physiological processes, falls under Rule 11).
[20] Regulation (EU) 2024/1689 (AI Act).
[21] TechCrunch, "OpenAI unveils ChatGPT Health, says 230 million users ask about health each week", 7 January 2026.

is that the users most protected by data protection law are the ones denied access to the most privacy-protective chatbot feature currently available.

1.5 The enterprise dimension

Separately, OpenAI also announced OpenAI for Healthcare on 8 January 2026, a set of enterprise-grade, HIPAA[22]-ready products for healthcare organisations, already rolling out to institutions including Boston Children's Hospital, Cedars-Sinai, and Stanford Medicine Children's Health.[23] This is a distinct enterprise offering, falling outside the consumer AI chatbot scope of my study, but it confirms OpenAI's broader strategic investment in health-specific infrastructure.

2. Anthropic: a related direction, through a different product mix

Anthropic has also moved toward differentiated handling of health-related interactions, though through a more mixed enterprise-consumer form. On 11 January 2026, days after OpenAI's announcement, and timed to coincide with the JPMorgan Healthcare Conference, Anthropic launched Claude for Healthcare, combining HIPAA-ready offerings for enterprise customers (providers, payers, and health-tech companies) with consumer-facing tools intended to help individuals understand and navigate their own health data.[24]

The consumer-facing component is also in beta, and also limited to the United States. Claude Pro and Max subscribers in the US can connect their personal medical records to Claude through HealthEx, a startup that consolidates electronic health records from over 50,000 health systems via the federal TEFCA interoperability framework. Additional connectors to Function Health, Apple Health, and Android Health Connect are also rolling out in beta.[25] Anthropic states that Claude requests only the categories of information relevant to a specific question rather than pulling an entire medical record, that users must explicitly opt in, can disconnect permissions at any time, and that health data is not used to train models.[26] This confirms that the idea of giving health-related AI interactions a distinct, more protective treatment is not unique to one company.

But Anthropic's current consumer implementation differs from OpenAI's in an important structural respect: there is no dedicated "Health" tab in the Claude interface, no separate memory space for health conversations, and no system that nudges users to switch lanes when they begin discussing health topics in ordinary chat. Independent analysis has noted that "it doesn't seem that Anthropic has invested much product development capacity to go deep here (at least yet)" and that "the experience is slightly clunkier in Claude" compared to ChatGPT Health's more integrated design.[27] The direction of travel is similar; the product expression and the degree of compartmentalisation remain different.

[22] The Health Insurance Portability and Accountability Act (HIPAA) is a US federal law enacted in 1996 that establishes national standards for the protection of individually identifiable health information ('protected health information' or 'PHI').

[23] OpenAI, "Introducing OpenAI for Healthcare", 8 January 2026. This is a separate enterprise offering with HIPAA-ready infrastructure, a BAA, and role-based access controls.

[24] Anthropic, "Advancing Claude in healthcare and the life sciences", January 2026. The page was edited on 7 February 2026 to clarify scope. See also Fortune, "Anthropic rolls out new healthcare and life science features for Claude", 11 January 2026.

[25] See Fierce Healthcare, "JPM26: Anthropic launches Claude for Healthcare to turbocharge AI efficiency at health systems, payers", 11 January 2026; HealthEx/HLTH, "HealthEx and Anthropic Partner to Bring Personal Health Records Directly to Claude", 12 January 2026.

[26] TechCrunch, "Anthropic announces Claude for Healthcare following OpenAI's ChatGPT Health reveal", 12 January 2026. Both companies have stated that health data connected through these integrations is not used to train their models.

[27] Health API Guy, "Another One: Anthropic's Healthcare Debut", 16 January 2026.

3. The acceleration: Microsoft, Amazon, Perplexity, and the emerging industry pattern

What appeared in January 2026 to be an initiative by two companies has, within less than three months, become an industry-wide movement. Consumer health AI has been described, with justification, as "the year's fastest-moving product category".[28] Since the publication of my study in early March, three additional major technology companies have launched health-specific AI products that follow a strikingly similar pattern.

Microsoft Copilot Health was announced on 12 March 2026 as a "separate, secure space within Copilot where medical intelligence makes sense of your information and delivers personalized health insights."[29] The architecture mirrors ChatGPT Health in several key respects: health data and conversations are stored separately from general Copilot interactions, encrypted, subject to stricter access controls, and not used for model training.[30] Like ChatGPT Health, Copilot Health connects to electronic health records (via HealthEx, the same partner Anthropic uses), wearables (Apple Health, Oura, Fitbit), and lab results. Microsoft states that its consumer AI products already handle over 50 million health questions per day.[31] The product is launching through a waitlist, in English, for US adults only.[32]

Perplexity Health launched on 19 March 2026, building on Perplexity's AI search platform.[33] It connects to Apple Health, electronic health records from over 1.7 million care providers via b.well (the same partner OpenAI uses), Fitbit, Ultrahuman, Withings, and other wearables. Health data is encrypted, not used for training, and not sold to third parties.[34] The product is available for Pro and Max subscribers in the United States only.

Amazon Health AI was expanded in March 2026 from its initial launch for One Medical members to the broader US public through Amazon.com and the Amazon app.[35] Amazon's approach goes furthest toward a true health agent model: Health AI can not only interpret medical records and answer health questions, but can also book appointments with One Medical clinicians, manage prescription renewals with Amazon Pharmacy, and connect users directly to licensed healthcare professionals. It accesses patient data through state health information exchanges. Amazon describes this as "personalized medicine at consumer scale." Of all the products reviewed in this article, Amazon's is the only one integrated with a company that owns its own healthcare provider network, creating a closed loop from AI conversation to clinical care.

The pattern across all five launches is remarkably consistent. Every product follows the health agent model: connecting medical records, wellness apps, and wearables to deliver personalised health intelligence. Every product promises not to train on health data. Every

[28] The Next Web, "Perplexity has launched Perplexity Health", 19 March 2026.

[29]Microsoft, "Introducing Copilot Health", 12 March 2026. See also Fortune, "Microsoft launches Copilot Health, a dedicated space for personal health data and AI-driven insights", 12 March 2026.

[30] Ibid. Microsoft ads that "the feature was developed with input from over 230 physicians across 24 countries, in collaboration with AARP and the National Health Council".

[31] See The Next Web, "Microsoft launches Copilot Health", 12 March 2026 citing Dominic King, VP of Health at Microsoft AI.

[32] Copilot Health connects to HealthEx for EHRs from over 50,000 US hospitals and provider organisations, wearables (Apple Health, Oura, Fitbit), and lab results from Function. It also serves expert-written answer cards from Harvard Health and connects to real-time US provider directories. See Healthcare Brew, "Microsoft launches AI platform, Copilot Health", 12 March 2026.

[33] Perplexity, "Introducing Perplexity Health", 19 March 2026. See also The Next Web, "Perplexity has launched Perplexity Health", 19 March 2026.

[34] See 9to5Mac, "Apple Health integrates with newly announced Perplexity Health AI feature", 19 March 2026.

[35] Amazon, "Amazon launches Health AI agent on Amazon website", March 2026. See also TechCrunch, "Amazon launches its healthcare AI assistant on its website and app", 10 March 2026. Fierce Healthcare, "Amazon launches health AI agent on its website, expands free virtual care to 200M Prime members", 13 March 2026.

product stores health conversations separately from ordinary chat. Every product is US-only or US-first. And every product is positioned as an intermediary between the user and the healthcare system, not as a privacy-protective lane for conversations that are already happening.

This convergence powerfully reinforces the central observation of this article. **The market has recognised, unanimously, that health conversations deserve differentiated treatment**.

But it has converged on one specific form of differentiation: the health data integration hub. None of these products addresses the question that Sealed Mode was designed to answer: what happens to the hundreds of millions of health-related conversations that users already have with chatbots every day, without connecting any medical record or downloading any app, in jurisdictions around the world including those where none of these pilot programmes is available?

The speed of the convergence is itself significant. Five major health AI products in less than three months suggests that the industry views consumer health as a competitive battleground where first-mover advantage matters. But the uniformity of the approach (connectivity and utility first, privacy protections as instrumental foundation) also suggests that a different kind of initiative, one that starts from the conversation rather than from the medical record, has not yet found its champion.

Table 1: The Health Agent Rush Comparative Overview

Assessed against the six components of Sealed Mode (T. Christakis, You Trust Your Chatbot With Everything. Should You?, March 2026)

PANEL A: Architecture, connectivity, and functionality

	ChatGPT Health	Claude Healthcare	Copilot Health	Health AI (Amazon)	Perplexity Health
Provider	OpenAI	Anthropic	Microsoft	Amazon	Perplexity
Launch date	7 Jan. 2026	11 Jan. 2026	12 Mar. 2026	22 Jan. 2026 (One Medical); 10–13 Mar. (Amazon.com)	19 Mar. 2026
Availability	Staged waitlist (Free, Go, Plus, Pro)	Beta (Pro & Max)	Staged waitlist	Live (One Medical); waitlist (Amazon.com/app)	Phased launch (Pro & Max, web + iOS)
Available in EEA / CH / UK?	✗ Explicitly excluded	✗ US only	✗ US only	✗ US only	✗ US only
Framing and connectivity					
Primary framing	Health data integration hub	Enterprise-consumer health tools	Health data integration hub ("medical superintelligence")	Health agent with provider integration	Health data search engine + dashboard
Medical record / EHR connectivity	✓ Via b.well	✓ Via HealthEx	✓ Via HealthEx (50K+ providers)	✓ Via state HIEs	✓ Via b.well (1.7M providers)
Wellness apps / wearables	Apple Health, MyFitnessPal, Peloton, AllTrails, Function, Instacart, Weight Watchers	Apple Health, Android Health Connect, Function (beta)	Apple Health, Oura, Fitbit, Function + 50 sources	✗ None; integrates One Medical, Amazon Pharmacy, health purchases	Apple Health, Fitbit, Ultrahuman, Withings, Clue; Oura, Function coming
Dedicated health space	✓ Yes	✗ No dedicated tab	✓ Yes	~ Separate page on Amazon	✓ Yes (hub + files repository)
Can take action (book, prescribe)	✗ No	✗ No	~ Provider directory	✓ Yes (book, prescriptions, connect to providers)	✗ No
Nudges user to health lane	✓ Yes	✗ No	~ Unclear	✗ No	✓ Auto-triggers on health queries

Panel B: Privacy and Data Protection Assessment

	ChatGPT Health	Claude Healthcare	Copilot Health	Health AI (Amazon)	Perplexity Health
Training & advertising					
No-training commitment	✓ Yes	✓ Yes	✓ Yes	Trains on "abstracted patterns" without identifying info	✓ Yes
Advertising-free in health space	✓ No ads in Health; health info not used for ads	No ads in Claude generally (not health-specific)	Not addressed publicly	PHI not used to market; may recommend Amazon health products	Not addressed publicly
Data isolation & access controls					
Isolated health memory	✓ Health memories never flow back to main chats	✗ Not documented	Data isolated from general Copilot; separate memory not documented	Unclear	Health memories stored separately; referenced only when relevant
Retention limits documented	✗ Not specified	✗ Not specified	✗ Not specified	✗ Not specified	Raw data queried on demand, not stored permanently
Human access constraints	"Limited authorised personnel" for safety (Privacy Notice); "more restricted" access (spokesperson)	✗ Not specified	✗ Not specified	HIPAA environment; strict access controls mentioned	✗ Not specified
Cryptographic hardening					
Zero-access architecture (provider cannot read plaintext)	"Purpose-built encryption and isolation" but provider retains access to content	Standard encryption; no health-specific cryptographic architecture	Encryption at rest/in transit; ISO 42001; provider retains access	Encryption within HIPAA environment; provider retains access	Encryption at rest/in transit; standard access controls

Note: Assessments based on publicly available materials as of 22 March 2026. All five products are available in the United States only. Product features may have changed since review. Colour coding: ✓ green = documented protection present; ~ orange = partial or unclear; ✗ red = absent or not publicly documented. The six components of Sealed Mode are: (1) no training; (2) no advertising; (3) siloed personalisation; (4) strict retention; (5) minimised human access with audit trails; (6) cryptographic hardening.[36] *All five products describe encryption at rest and in transit. None offers a zero-access cryptographic architecture, where even the provider cannot read user data. This distinguishes conventional security (protection against external attackers) from cryptographic hardening that would offer protection against the provider itself.*

Sources: OpenAI Help Center and Health Privacy Notice; Anthropic announcement; Microsoft AI launch page; Amazon News and TechCrunch; Perplexity Help Center and blog.[37]

4. Google, Grok and DeepSeek: no comparable consumer health lane

Against this backdrop of rapid market convergence, the remaining three consumer chatbots examined in my study (Google's Gemini, DeepSeek, and xAI's Grok) have not yet, in the materials reviewed for this article, launched a consumer-facing health compartment comparable to the products described above.

[36] See T. Christakis, "*You Trust Your Chatbot With Everything. Should You? Part 1: How The Controller Uses Your Chat Data*", AI-Regulation.com, *March 2026*

[37] *Sources:* ***OpenAI:*** *OpenAI, 'Introducing ChatGPT Health', 7 January 2026; OpenAI Help Center, 'What is ChatGPT Health?', accessed 22 March 2026; OpenAI, 'Health Privacy Notice', accessed 22 March 2026; OpenAI Help Center, 'Ads in ChatGPT', accessed 22 March 2026.* ***Anthropic:*** *Anthropic, 'Introducing Claude for Healthcare', 11 January 2026.* ***Microsoft:*** *Microsoft AI, 'Introducing Copilot Health', 12 March 2026.* ***Amazon:*** *Amazon News, 'Amazon One Medical introduces agentic Health AI assistant', 22 January 2026; Amazon News, 'Amazon launches Health AI agent on Amazon website and app', 10 March 2026; A. Malik, 'Amazon launches its healthcare AI assistant on its website and app', TechCrunch, 10 March 2026.* ***Perplexity:*** *Perplexity, 'Introducing Perplexity Health', 19 March 2026; Perplexity Help Center, 'What is Perplexity Health?', accessed 22 March 2026; Perplexity, 'Privacy Policy'.* All assessments based on publicly available materials as of 22 March 2026. Product features and availability may have changed since this review.

This does not mean that none of these providers has invested in health-related capabilities. Google has developed significant privacy-preserving infrastructure (Private AI Compute[38], a confidential-computing architecture comparable to Apple's Private Cloud Compute), introduced consumer personalisation guardrails that limit proactive use of health data in Gemini, launched, as other companies, Temporary Chats that are not used for training, and released MedGemma[39], a collection of open-source medical AI models for developers. These are meaningful investments, but they operate at the infrastructure and developer level rather than creating a consumer-facing health lane with the kind of compartmentalisation, separate memory, and health-specific no-training commitment that characterise ChatGPT Health or that Sealed Mode proposes.

For DeepSeek, the approach is different still. DeepSeek's privacy policy expressly states that its services are "not designed or intended to process" sensitive personal data, including health data, and tells users, in bold characters, **"not to provide such information"[40]** - even though, as my study documents, users routinely do. For Grok, xAI's public materials describe a general consumer service without a distinct high-assurance lane for sensitive conversations.

The three providers thus arrive at the same outcome - no dedicated health compartment for the time being - through different routes: Google through infrastructure investment that has not yet translated into a consumer product, DeepSeek through a policy of disclaiming responsibility for sensitive use, and Grok through the absence of any visible initiative in this area.

What unites all three cases is the gap they leave. When five major technology companies launch health-specific AI products within the space of three months, **the assumption that all chatbot conversations can be governed by roughly the same defaults becomes not merely difficult to sustain but increasingly anachronistic**.

5. What this means for the Sealed Mode debate

The most important lesson is not that one provider has "implemented" Sealed Mode and others have not. None of these developments corresponds precisely to the full framework proposed in my study.[41]

The **lesson is that the market has now decisively recognised** the basic intuition behind Sealed Mode: consumer chatbot conversations are not all the same. Some concern ordinary everyday tasks. Others concern symptoms, emotional fragility, crises, legal fears, work conflicts, or intimate personal circumstances. It is increasingly artificial to pretend that these exchanges can all be governed by the same privacy defaults, the same training settings, and the same downstream-use logic.

But the analysis in this article also reveals that the market's approach and Sealed Mode's approach, while convergent in their recognition that health conversations deserve differentiation, remain fundamentally oriented differently. The market has converged, with remarkable speed and unanimity, on **health data integration hubs**: trusted AI health agents that connect records, apps, and wearables to deliver personalised intelligence. Five major

[38] Google, "Private AI Compute", November 2025 (updated January 2026).
[39] See here. For Google's approaches on the issue see also : See also Google, "Helping healthcare move from data to agentic action", March 5, 2026.
[40] DeepSeek Privacy Policy, February 10, 2026.
[41] Christakis (n 1).

products launched in under three months, all following the same model. The privacy protections bundled with these hubs are welcome and significant, but they are designed to support the integration objective and to create user trust for this. Sealed Mode, by contrast, starts from the millions of intimate conversations that already happen every day, without any data integration, and asks a simpler question: should those conversations be protected by architecture rather than by fine print?

The **European exclusion/postponement** brings this distinction into even sharper focus. Every single one of the health AI products reviewed in this note is US-only or US-first. A Sealed Mode feature, based on stronger privacy defaults for health conversations that already take place, could be offered globally today, including in the most regulated jurisdictions. A health data integration hub cannot, because it triggers a cascade of additional regulatory requirements that have nothing to do with the privacy protections and everything to do with the data connectivity. The two should not be conflated. Providers **could, and in my view should, offer the first while working toward the second**.

But the health agent rush raises questions that go beyond the Sealed Mode framework itself.

6. Beyond Sealed Mode: the data protection questions that the Health Agent Rush leaves open

The speed with which five major technology companies have launched health-specific AI products raises a set of data protection questions that go well beyond the scope of Sealed Mode. This article does not attempt to answer them all. But it would be incomplete without naming them, because they define the regulatory and academic agenda that this convergence has created.

6.1.The business model question: *"there is no free lunch"*

Several of these health products are offered for free or bundled with existing subscriptions. But the strategic logic is transparent. Amazon's Health AI channels users toward One Medical consultations and Amazon Pharmacy prescriptions. Microsoft frames Copilot Health as a pathway to "medical superintelligence". Perplexity and Anthropic restrict health features to paying subscribers, creating a direct conversion incentive. OpenAI offers ChatGPT Health across all tiers including its free plan, but the health feature deepens user engagement and enriches the profile of a user base that OpenAI has already started to simultaneously monetise through advertising in the standard ChatGPT experience.[42] Once a user connects years of medical records, wearable history, and wellness data to a single AI platform, switching costs become prohibitive. As one cybersecurity expert observed: *"Nothing is free. Companies monetize your data"*.[43] The question for regulators is not whether these products will be monetised, but how - and whether the privacy commitments made during the "trust-building" phase of launch will survive the monetisation phase that follows.

[42] See Christakis (n1).

[43] See e.g. Van Steel, quoted in BankInfoSecurity, "ChatGPT Health: Top Privacy, Security, Governance Concerns", 8 January 2026.

6.2 The gold mine question: secondary use of the world's largest health dataset

Every provider states that health data in the new offerings is "not used to train our foundation models". But this commitment is narrower than it appears. It does not necessarily preclude using health data for product analytics, quality improvement, aggregate statistical research, or what Amazon calls training on "abstracted patterns without directly identifying information".[44] The gap between "not training foundation models" and "not using this data for any form of learning or improvement" is significant and largely undisclosed.

More fundamentally, a platform that aggregates hundreds of millions of health conversations linked to medical records and wearable data would constitute the largest health dataset in human history. The value of such a dataset for epidemiological research, drug discovery, clinical pattern recognition, and pharmaceutical development would be immense. The temptation for providers - or for third parties willing to pay for access - will only grow. Current privacy policies can be modified at any time. Under US law, since these consumer products are not HIPAA-covered entities[45], there is no federal constraint on secondary use beyond each company's voluntary commitments.[46] Under GDPR, the legal basis for processing health data for research purposes (Article 9(2)(j)) does not necessarily require individual consent if national law provides for it, though it does require appropriate safeguards.[47] The question is whether the privacy commitments of today will withstand the commercial pressures of tomorrow.

6.3 The institutional competition question: private health hubs and the European Health Data Space

The European Union has spent years building the European Health Data Space (EHDS), an ambitious public infrastructure project for both primary use of health data (patient care) and secondary use (research, innovation, policy-making), governed by strict rules on access,

[44] See Amazon, "Amazon launches Health AI agent on Amazon website", March 2026. Amazon explains that "For example, if multiple patients ask about medication interactions, we may use those patterns—without patient names—to improve how Health AI responds to similar questions. We only use protected health information for purposes permitted under HIPAA". See also TechCrunch, "Amazon launches its healthcare AI assistant on its website and app", 10 March 2026. The distinction between training on "abstracted patterns" and training on user data raises questions about the practical scope of no-training commitments.

[45] The Health Insurance Portability and Accountability Act (HIPAA) is a US federal law enacted in 1996 that establishes national standards for the protection of individually identifiable health information ('protected health information' or 'PHI'). HIPAA applies to "covered entities" (healthcare providers, health plans, and healthcare clearinghouses) and their "business associates" (entities that perform functions involving the use or disclosure of PHI on their behalf). Crucially, HIPAA does not apply to technology companies that receive health data directly from consumers, because such companies are neither covered entities nor business associates unless they have signed a Business Associate Agreement (BAA) with a covered entity. This means that when a user voluntarily uploads medical records to a consumer AI product like ChatGPT Health, that data falls outside HIPAA's protective framework and is governed only by the company's own privacy policy and applicable state consumer protection laws. This is a **fundamental structural difference** from the EU framework, **where the GDPR applies to any controller processing personal data**, regardless of whether the controller operates within the healthcare system.

[46] See also S. Geoghegan (EPIC), quoted in The Record, "ChatGPT Health feature draws concern from privacy critics over sensitive medical data", 8 January 2026: "ChatGPT is only bound by its own disclosures and promises, so without any meaningful limitation on that, like regulation or a law, ChatGPT can change the terms of its service at any time".

[47] Under Article 9(2)(j) GDPR, processing of special categories of personal data for scientific research purposes is permitted subject to appropriate safeguards, including, where feasible, pseudonymisation, and provided that the processing is proportionate and respects the essence of the right to data protection. This derogation does not necessarily require individual consent if national law provides for it.

purpose limitation, and oversight by national health data access bodies.[48] France's Health Data Hub (Plateforme des données de santé) pursues a similar objective at national level.[49]

The emergence of private health AI hubs creates a parallel ecosystem that, in the United States where all five products currently operate, functions largely outside any health-specific regulatory framework. If and when these products expand to Europe, they would of course be subject to the GDPR, including the strict requirements for processing health data under Article 9. But even under GDPR, a structural tension with the European Health Data Space would remain. The EHDS imposes a dedicated institutional governance layer for secondary use of health data: access through designated health data access bodies, purpose-specific permits, and structured oversight that goes beyond general GDPR compliance. A private AI platform that aggregates consumer health data under consent-based processing could potentially derive similar insights - epidemiological patterns, pharmaceutical intelligence, population-level health trends - without passing through the EHDS's institutional gatekeeping. The GDPR would apply; the EHDS's structured governance would not. This **asymmetry deserves urgent attention** from European policymakers, because it risks creating a two-track system in which the most demanding governance requirements apply to public institutions and academic researchers, while the largest aggregations of health data sit in private hands under less structured, consent-based frameworks.

This is not to suggest that private health AI products should be blocked from the European market. For millions of patients who currently piece together their health picture from scattered portals, contradictory apps, and fifteen-minute consultations where the doctor has not read the file, an AI that connects the dots could be genuinely transformative. It is to suggest that policymakers should work proactively, and in dialogue with providers, to ensure that the governance frameworks for private and public health data ecosystems are coherent, interoperable, and mutually reinforcing rather than competing. The goal should not be to choose between innovation and governance, but to design a framework in which both can thrive.

6.4 The cybersecurity question: concentration of risk at unprecedented scale

Five companies are now building centralised repositories of health data linked to personal identities, medical records, and longitudinal behavioural patterns. A single breach at any of these providers could expose health data of a magnitude never seen before. Unlike hospital breaches, which are geographically contained, a breach of a global consumer health AI platform could simultaneously expose the health data of users across dozens of countries.

The protections described in the public materials are worryingly vague. All five companies mention "encryption at rest and in transit" and "strict access controls," but none has published a detailed security architecture or announced it would allow independent verification.[50] Each product also relies on third-party data intermediaries (b.well, HealthEx, Terra API), each of

[48] Regulation (EU) 2025/327 on the European Health Data Space (EHDS). The EHDS creates a harmonised framework for both primary use of health data (patient care) and secondary use (research, innovation, policy-making), with strict governance requirements including purpose limitation, data minimisation, and oversight by national health data access bodies. For more info and the timeline see European Commission, EHDS.

[49] See Health Data Hub.

[50] See R. D'Ovidio (Drexel University), quoted in HuffPost, "Experts Warn Sharing Medical Information With ChatGPT Health Is Dangerous", 15 January 2026 ("people also have to be concerned about what's going on at b.well because b.well is now part of the equation"). See also Dark Reading, "ChatGPT Health Raises Big Security, Safety Concerns", 21 January 2026.

which adds a point of vulnerability in the data supply chain.[51] Furthermore, as recent litigation has demonstrated, data stored in these systems may be legally compellable through subpoenas, warrants, or government access requests, a question that will be addressed in detail in Part 2 of my study.[52]

6.5 The structural data protection question: design choices that challenge core GDPR principles

Even leaving aside the specific health data issues under Article 9, the architectural choices behind these products sit in tension with several foundational principles of data protection law.

Data minimisation. These products are designed to aggregate as much health data as possible from as many sources as possible: medical records, wellness apps, wearables, lab results, pharmacy data, insurance information.[53] This is structurally the opposite of the principle that personal data should be limited to what is necessary for the stated purpose.

Storage limitation. No product reviewed in this note (with the partial exception of Perplexity, which states that raw health data is queried on demand rather than stored permanently) publicly specifies retention duration for health data. This absence is striking given the sensitivity of the information.

Purpose limitation. Health data is collected for "personalised health insights". But can it be used for product improvement? For de-identified analytics? For developing new health features? The boundaries between primary and secondary purposes are poorly defined and could expand over time without meaningful user awareness.

Controller and processor complexity. The intermediaries involved (b.well, HealthEx, Terra API, and others) create layered data controllership arrangements that users cannot meaningfully assess or audit. When a user connects their medical records to ChatGPT Health through b.well, who is the controller of what, and under which terms?

Data portability. If a user decides to move their health AI relationship from one provider to another (from ChatGPT Health to Copilot Health, for instance) can they take their accumulated health data, memories, and conversation history with them? Nothing in the current product architectures supports this, creating a lock-in effect that further concentrates market power.

Children's data. Public materials do not evidence a uniform hard age-verification model across these products, which is significant in a health context because such systems may collect or infer sensitive data about minors.

[51] See E.F. McSherry (EFF), quoted in Dark Reading (*idid*): "If you give your data to any third party, you are inevitably giving up some control over it and people should be extremely cautious about doing that when it's their personal health information".

[52] As we will discuss in Part 2 of my study, a federal judge in the U.S. recently upheld an order requiring OpenAI to produce 20 million anonymised ChatGPT conversation logs to news organisations as part of a consolidated copyright case. It is plausible that health data could be sought in future legal proceedings or demanded by government authorities.

[53] GDPR, Article 5(1)(c) (data minimisation).

6.6 Getting it right: why the answer is governance, not exclusion

All these questions are not reasons to reject the health agent trend. Nobody disputes that better-informed patients make better health decisions. And the potential benefits for users who struggle to navigate fragmented, and often overloaded healthcare systems are real as physicians themselves acknowledge.[54]

A fair assessment must also acknowledge that data integration in the health AI context is not driven solely by business logic or platform strategy. It can serve a genuine clinical function. Generic health guidance from a chatbot that knows nothing about the user's medical history, medications, or pre-existing conditions may be not merely unhelpful but actively harmful: encouraging vigorous exercise for someone with an undiagnosed cardiovascular condition, for instance, or suggesting a dietary supplement that interacts with a current prescription. As Google Research has argued in a recent study on the architecture of personal health agents, integrating personal health data (wearables, medical records, biomarkers) is a prerequisite for safe, context-aware responses in a domain where getting it wrong can carry serious consequences.[55] In this sense, the data integration dimension of the health AI products reviewed in this article should probably not seen as merely a business proposition. It is also, when and if properly governed, a safety feature. This reinforces rather than undermines the case for strong governance: the more data these systems ingest, the more consequential their outputs become, and the more essential it is that the privacy and security frameworks surrounding them are built to match.

The convergence documented in this article has created, in less than three months, a new category of consumer service that processes the most intimate data people possess - their health, their fears, their bodies - under governance frameworks that remain under the unilateral control of each provider, and whose durability has not been tested by time, by commercial pressure, or by independent scrutiny. The regulatory gap is already wide. If the pace of product launches continues to outstrip the pace of governance, it will only grow.

This article has focused on the Sealed Mode dimension of the health agent rush: the gap between conversational privacy and data integration, and the missed opportunity to offer the first without being blocked by the second. But the questions raised in this section suggest that the agenda is far larger. They **call for serious academic and regulatory attention**: from privacy

[54] See American Medical Association, Physician Survey on Augmented Intelligence (2026), March 2026. The AMA's annual survey found that 81 per cent of US physicians now use AI in their practices. Physicians report growing confidence in AI's clinical and administrative benefits. Physicians are also increasingly comfortable with patients using AI for certain purposes: 68 per cent support patients using AI for questions about medications and side effects, and 64 per cent for general health questions. However, nearly half would never or rarely want patients using AI to interpret pathology (49 per cent) or radiology results (46 per cent). This nuanced picture maps directly onto the health AI products analysed in this article, which are designed to help consumers do all of these things: answer general health questions, explain medication interactions, but also interpret test results and medical records. The physician community, in other words, welcomes the category but draws a line within it. That line deserves attention from both providers and regulators.

[55] See A. Ali Heydari et al., "The Anatomy of a Personal Health Agent", Google Research, 2025. See also here. The paper presents a multi-agent framework built on Gemini 2.0 that reasons about multimodal data from wearables and personal health records to provide personalised health recommendations, and argues that effective health AI requires integration of diverse personal health data sources to deliver safe, contextualised guidance. The paper also argues that effective health AI requires a multi-agent architecture that decouples medical reasoning from user-facing communication, with specialised sub-agents (a Data Science Agent, a Domain Expert Agent, a Health Coach Agent) operating under distinct constraints. The authors contend that this modular design enhances explainability because outputs can be traced to the specific function of each agent, which is described as 'an essential requirement in sensitive health contexts.' While the implications of such architectures for AI Act transparency and explainability obligations deserve separate analysis, they reinforce the broader point that health AI systems require purpose-built governance frameworks rather than the general-purpose defaults that currently apply to consumer chatbots.

scholars, from health law specialists, from cybersecurity researchers, and from the institutions responsible for the European Health Data Space.

7. What comes next

Going back for now to my Sealed Mode proposal, it seems clear that the real question is no longer whether differentiated privacy for sensitive chatbot conversations is conceivable. Five major technology companies have answered that question in less than three months. The question is whether the privacy-protective core (no training, isolated memory, no advertising, minimised human access, strict retention, cryptographic hardening) can be extracted from the health data integration products it is currently bundled with and offered as a **standalone standard, available to every user, everywhere**.

This is not to deny that data integration serves important purposes in the health context, including clinical safety. The objective is to argue instead that the privacy-protective dimension should not be held hostage to the integration dimension: the first can and should be offered everywhere, while the second is deployed where governance permits.

If these products remain US-only pilots inseparable from their medical record integration features, and if no provider offers the privacy-protective dimension to the hundreds of millions of users in Europe (and other continents) who confide health concerns to chatbots every day without connecting any app or uploading any record, then the lesson will be limited. If, by contrast, providers recognise that the privacy-protective dimension **can and should be separated from the integration dimension** - offered first, offered globally, offered as a baseline rather than a premium - then the debate will have entered a new phase.

That would be the real moment when Sealed Mode meets the market. Not when a provider builds a health hub with privacy features attached. But when a provider decides that some conversations deserve stronger protections simply because of what the user is saying, regardless of whether any medical record has been uploaded, any app has been connected, or any subscription has been purchased. And that, ultimately, is the point. The issue is not whether the label travels. It is whether the principle does.

As for the "health agents" dimension, policymakers and regulators in Europe should work proactively, and in dialogue with providers, to find constructive and protective solutions. One avenue worth exploring is the use of regulatory sandboxes, already provided for under the AI Act, to allow health AI products to operate in Europe under supervised conditions with strong, verifiable privacy safeguards. The current outcome, where the regulatory complexity of deploying a health data integration product in Europe results in Europeans being excluded entirely, serves neither innovation nor protection. A supervised pathway that conditions market access on demonstrated compliance with robust privacy standards, including the kind of architectural safeguards that Sealed Mode envisions, would be preferable to the status quo, in which European users receive no additional protections at all.