



AN OATH FOR GenAI IN HEALTHCARE

I swear (or affirm) by all I hold dear and sacred, and in the presence of those assembled here:

To commit to life-long learning, sharing knowledge with others, and improving the standards by which my profession operates, as it incorporates GenAI, knowing its risks as well as its benefits.

To do my personal best and encourage others on my teams and throughout my organization, to offer our skills to serve patients and their wellbeing.

To ensure my personal and organizational interests and financial gains do not eclipse the care of my patients.

Intending only to do good, ensure the expertise of those involved, commit to improving life, contribute what I know, and collaborate with others to ensure the best outcomes for those I serve.

To encourage those in my care, their clinicians, and others engaged in healthcare to appropriately use the tools I build and use, welcoming their feedback to support improvements to benefit them and others.

To be always mindful that healing is a human enterprise, and humans must be informed, aware, and in control of each step, lest we lose the trust upon which we depend for the therapeutic alliance.

To maintain the highest moral, ethical, technical, financial, legal, and research conduct, including:

Not making false claims about my product.

Keeping information entrusted to me private, confidential, and secure lest my patients suffer further from any kind of bias and the whims of online multitudes.

Recognizing that the tools of GenAI do greatly benefit patients, yet may also cause harm.

If I keep this oath faithfully, may I enjoy my life and work and always be respected. If I violate this oath, may the reverse be my lot.

Carrie Brubaker

Carrie Brubaker, PhD

Glenna Crooks

Glenna Crooks, PhD

Executive Summary

In *Profiles of the Future: An Inquiry into the Limits of the Possible*, Sir Arthur C. Clarke wrote, “Any sufficiently advanced technology is indistinguishable from magic.”

Generative artificial intelligence (GenAI) is not the only advanced, seemingly magical technology to appear in medicine, but it is among the most recent and may be the most impactful.

It is also the first in many decades to bring a large number of new players into the healthcare arena who do not view themselves as healers. As happened when post-World War II European nations created healthcare systems and when the U.S. created Medicare 50 years ago, these people do not ascribe to the oaths typical of the medical, nursing, and pharmacy professions.

GenAI is not the only seemingly magical technology to appear in medicine, but it is the most recent and may be the most impactful.

Lacking formal training and enculturation in the healing arts, these new players may be unaware of the centrality of the therapeutic alliance between patients and clinicians. It is likely they are unfamiliar with deeply-held sociocultural expectations imposed on any healing enterprise. They are inexperienced in dealing with societal backlash should they violate those norms, particularly when doing so harms patients. They may not realize that since before the dawn of recorded time, societies held healers in high regard, not solely due to their wizardry, but because theirs was not a trade. It was a sacred profession that pledged its duty to the divine before it did so to patients, as described in the Hippocratic Oath and similar oaths that followed. In those oaths, healers reassured society they would practice their craft only for the good of patients, never for evil, expedient, or political purposes.

When any new actor enters the healthcare arena – as those embracing GenAI do today – and their art is a mystery and seems magical, they should ascribe to similar oaths. Patients, communities, and others in today’s healing enterprises deserve to know that this newest tool is in the hands of skilled, trustworthy developers and users who will uphold the highest standards and never violate the long-held expectations of patients.

This is why we propose an *Oath for GenAI in Healthcare*. It aligns with similar oaths taken by physicians, nurses, and pharmacists and adheres to essential criteria and familiar structure of those oaths.

An *Oath for GenAI in Healthcare* supplements GenAI policy and regulatory initiatives already underway to help close the chasm between GenAI risks and opportunities. It helps address any ambiguity in GenAI applications as we refine and more widely use them. It helps remind those who develop and/or use GenAI that they are participating in a healing enterprise, even if their work never brings them into direct contact with patients.

An *Oath for GenAI in Healthcare* supplements policy and regulatory initiatives to help close the chasm between GenAI risks and opportunities.

We propose that any person, team, or organization developing or using GenAI in healthcare should ascribe to this oath or to one tailored to their specific situation.

Opportunities and Risks

Artificial Intelligence (AI) is a multidisciplinary field of computer science that aims to create machines capable of performing tasks that typically require human intelligence. It encompasses a wide range of technologies and techniques that enable computers and software to mimic human-like cognitive functions, such as learning, reasoning, pattern recognition, problem-solving, and decision-making. Various branches of AI help do that, including natural language processing (NLP), computer vision, and reinforcement learning, which enable machines to understand text, interpret visual data, and learn from their environment.

Large Language Models (LLM), one type of AI, use deep learning techniques and large data sets to understand, summarize, generate, and predict new textual content. Machine Learning (ML) uses LLMs to improve outputs over time without explicit programming by building algorithms that learn from and make predictions or decisions based on data.

Generative models, of which GenAI is a type, are ML models that generate new data – such as text or images – similar to data used to train it. Though



impressive and despite its potential – and perhaps because of it – it will create significant challenges. That is because its ability to replicate complex human emotions, intuition, and nuanced decision-making is limited. Hence, it might one day harm individuals and populations.

Within healthcare, GenAI offers unprecedented opportunities to augment the intellectual capacity of individuals, teams, and entire organizations. It, therefore, holds immense potential to transform healthcare by accelerating drug development, enhancing decision-making, improving communication, increasing accessibility, and revolutionizing all data-related practices.

However, GenAI also creates risks for any sector that deploys it; healthcare will not be immune. Today, as trust in the internet, digital tools, public health, and experts is at an all-time low, managing potential risks to patient and population health is important.

Lagging Oversight

People who lack sufficient knowledge of both AI *and* the content areas using AI tools are engaged in deploying it. In healthcare, for example, very few AI platform developers are clinicians, and very few clinician users are trained platform developers. That, combined with uncertainty about the source, correctness, and validity of its training data, limits the ability of many users to make fully informed decisions about whether, where, and how to integrate it correctly and safely.

At this stage of its lifecycle, GenAI's regulatory landscape remains underdeveloped, with many unanswered questions about its training data, ethical

guidelines, liability, and intellectual property rights. Governments and regulatory bodies currently struggle to keep pace with rapid advancements.

Likely, they always will. As its developments outpace the ability of users, regulators, and the public to keep up, there is a danger that "hype by example" will oversell its value and overlook its risks in the meantime.

As GenAI outpaces our ability to keep up, there is a danger that "hype by example" will oversell its value and overlook its risks.

Unpredictable Outcomes

GenAI must be *creative* to be *generative*. It works by recognizing patterns in the data used to train it and extrapolating those to entirely new outputs. That makes it unlike more traditional logic-based software, where all inputs can be controlled, and outputs can be known in advance. GenAI outputs will be more variable and less predictable than other data-based technologies.

In addition, GenAI – and LLMs in particular – exhibit *emergent* properties, behaving in ways and with capabilities even its creators did not expect. This means that alongside unknown potential are unpredictable risks. This occurs partly because GenAI top-level programming, including user input (prompts), is written in natural language rather than computer code. End-users have as much influence on the outputs as the initial developers. This blurs the distinction between developers and users in a way that traditional logic-based software cannot.

Uncertain Accountability

Unlike Predictive AI, which can generally be proven right or wrong, it is harder to rate the correctness of GenAI results, and its "black-box" nature means it is difficult to explain or interpret how it arrived at the result. This raises concerns about who we will hold accountable for errors or misuse. Unregulated predictive algorithms lacking transparency about training data and scientific rigor may go unchallenged, partly because they emerge faster than anyone can keep up.

Physicians are aware of this risk and are understandably concerned that GenAI-driven medical decisions will increase legal liability. As knowledge grows, diagnosis codes change, or cyber criminals cause malicious actions, new risks will appear entirely unforeseen by regulators, developers, and users today. As vulnerabilities emerge and we resolve them, new ones will emerge.

Privacy Violations

GenAI tools process vast amounts of patient data from multiple sources, presenting new privacy and security challenges with serious ethical, legal, and practical ramifications when breaches occur, as they most probably will. Despite anonymizing, deidentifying, and masking training data, it can unintentionally expose



sensitive information, such as proprietary or personal data, which may be embedded within patterns recognized by the AI.

Correcting the issues presented when personal data leaks into GenAI training datasets can prove extremely difficult, if not impossible, as the behavior of the GenAI model will reflect its training data even after that personal data has been deleted. GenAI models do not retain “memory” as specific data points an operator can scrub, as would be the case in other databases.

GenAI offers great promise in rare diseases, for example. Extrapolating from small datasets has already been shown to support the statistical significance needed for clinical research and the development of new treatments. Managing such small datasets requires even greater and more careful handling to ensure that GenAI does not immediately enable specific patients to be identified, whether directly or indirectly.

Perpetuated Biases

ML processes used in GenAI reinforce whatever patterns are present in the training data, including those already known to be biased based on historical and current practices in healthcare systems. Limitations of this type and lack of knowledge about the source of training data are why 86% of respondents in a late 2023 survey agreed that using GenAI in healthcare is problematic.¹

Because GenAI produces results so quickly, efficiently, and confidently, there is a risk of making wrong decisions based on flawed logic, unverified sources, and inherent biases. GenAI’s outputs, therefore, need rigorous scrutiny to avoid perpetuating unfair practices and those that will create a risk to patient’s health and lives. Without active efforts to correct it, GenAI not only reproduces biases but does so at speed and scale. Applications built on inherent biases will exacerbate health disparities, not alleviate them.

The recent speculation that AI may be able to speak for incapacitated patients in end-of-life decision-making² might further undermine the already fragile relationship between underrepresented and marginalized persons and their healthcare providers.

Unearned Trust

GenAI is built to be *fast and convincing*, not *cautious and accurate*, even though it is capable of outright

errors.³ Its creativity and generativity, coupled with the warmth, patience, and kindness of its user interface, can lead people to grant it unearned trust and authority.

It may appear to be human and sentient, but it is neither. Unlike humans, an LLM does not understand the meaning of the words it strings together. It does not know if those words are right or wrong. It can appear to be genius, but the term for its mistakes – “hallucinations” – belies the truth; its mistakes can be outright wrong and dangerously so. It may appear to be nice, but it has no heart. It is not conscious; it is a phenotype of consciousness.

It can replicate a patient’s voice, but it will never replace the human ability to characterize the burden of disease in their daily life.

Authoritative – and currently, premature – “anointing” of GenAI as the next panacea⁴ exacerbates this risk.

Job Displacement

GenAI can potentially improve healthcare system *efficiency*, especially by streamlining workflows. However, it is the knowledge, wisdom, and well-being of its workforce that makes health systems *effective*. Excluding front-line workers from GenAI implementation planning will limit its effectiveness and potential to be beneficial.

Failing to account for workforce impact may accelerate a brain drain underway even before the COVID-19 pandemic, which has worsened since. It will impact patient care quality because healing is, after all, a human-to-human experience.

Malicious Actors

As is the case for all healing practices, we will experience accidents caused by GenAI. We will most certainly also experience malicious actions by bad actors because it creates new offensive capabilities.

Criminals are typically the first to adopt new technologies. Medical equipment, life-support systems,

GenAI is built to be *fast and convincing*, not *cautious and accurate*, even though it is capable of outright errors. Like a scalpel, it can be a lifesaving tool in the hands of a skilled surgeon; in the hands of a toddler, it can be deadly.



and patient record errors and privacy violations will be targets for those intending to harm or create hostages for financial gain, as real-world experience has already demonstrated.⁵

Policy and Regulatory Initiatives to Mitigate Risks

A growing number of organizations – governments, professional associations, think tanks, and civil society groups – are working to achieve AI’s promise and mitigate the risks. It would be impossible to catalog them all, but several are notable.

American Association for the Advancement of Science

Its [Decision Tree for the Responsible Application of Artificial Intelligence](#) is a guide to help operationalize ethical principles in developing and implementing AI. It frames appropriate questions and identifies steps to avoid integrating negative patterns and outcomes.

U.S. National Academy of Medicine

Its [Leadership Consortium](#) of health, tech, research, and bioethics leaders is a 3-year project launched in June 2023. It will produce a code of conduct for developing and using AI in health, medical care, and health research. It will also describe the national architecture required to support AI’s equitable and responsible use.

European Union

Its [AI Act](#) is a comprehensive legal framework governing the development, deployment, and use of AI systems within the European Union. It establishes consistent standards for AI member states, focusing on safeguarding health, safety, and fundamental rights. It is expected to impact global AI governance significantly.

U.S. Office of the President

Its [Executive Order On Artificial Intelligence](#) emphasizes eight key principles, including safety, security, and mitigating risks, to harness AI’s potential while minimizing societal harms. It aligns with digital regulations in the EU, such as the *General Data Protection Regulation* and the *Digital Services Act*.

The Light Collective

They aim to advance patient communities’ rights, interests, and voices in health technology. Its call to action and code of conduct in [AI Rights for Patients](#) outlines seven rights healthcare must protect as it adopts AI.

Coalition for Health AI

This [coalition](#) of technology companies, healthcare systems, and others plans to develop guidelines and guardrails to promote credible, fair, and transparent health AI systems that drive quality care.

World Health Organization

Its [Ethics and Governance of Artificial Intelligence for Health](#) is the WHO Guidance on all aspects of the use and governance of AI within the health sector.

Others focus on GenAI:

U.S. National Academy of Medicine

Its [Issue Framing Conversation on LLMs / Generative AI in Health and Medicine](#) identifies those applications and implications for health and health care and explores considerations for AI policy and oversight.

American Medical Association

Its [Principles for Augmented Intelligence Development, Deployment, and Use](#) focuses on appropriate deployment and disclosure of GenAI in medical settings, with views on physician liability, data privacy, cybersecurity, and payer use.

Medical Affairs Professional Society

Its [MAPS Vision for Generative Intelligence in Medical Affairs](#) describes how GenAI will reshape Medical Affairs’ possibilities and practice, focusing on strategy and leadership, evidence and insights generation, evidence and insights communication, and engagement and partnerships.

World Economic Forum

Its [Patient-First Health with Generative AI: Reshaping the Care Experience](#) outlines use cases for GenAI in healthcare. It also identified three barriers to safe, effective, patient-facing generative AI: data issues, mistrust of outputs, and barriers to scaling outside wealthy countries.



Organization for Economic Co-operation and Development

Its **Initial policy considerations for generative artificial intelligence** inform policy-making, including in healthcare and scientific research, and support decision-makers in addressing them.

UK National Health Service

The NHS provides top-level guidance on the use of GenAI in healthcare. Its guidance was compiled by the **Health and Care Information Governance Panel**, including the Information Commissioner's Office (ICO) and National Data Guardian (NDG). It includes advice for patients, service users, healthcare workers, and Information Governance professionals. For the latter, the guidance considers Data Protection and Impact Assessment matters, purpose and legal basis, controllers and processors, statistical accuracy, fairness, transparency, security, decision automation, and minimum data use principle.

These policy and regulatory initiatives share important common themes.

Safety, given the life and death implications in healthcare. The potential for harm must be carefully weighed against confirmed or expected benefits. This includes creating clear guidelines for development, deployment, and use.

Validity, with data sources, algorithm training, and maintenance rooted in reality and clinical expertise to deliver accurate results. This ensures sound decisions can be made confidently and enables quality use at scale. This includes methods to evaluate validity as GenAI and other healthcare tools (e.g., ICD codes) evolve.

Fairness and equity, to avoid proliferating existing biases in care delivery. This includes being aware of and sensitive to both conscious and unconscious biases known to exist in health care today, particularly related to underrepresented populations.

Transparency, eliminating "black box" approaches as much as possible to assure users know the source of training data so they are confident in that knowledge. This includes making the investments necessary to ensure safe adoption on a level playing field.

User training, to understand GenAI's possibilities and limitations. This includes creating clear guidelines concerning appropriate and inappropriate uses and clarity about human- *versus* GenAI-mediated responsibilities.

Privacy, to ensure individual and patient data protection through mechanisms that build trust among developers, users, and beneficiaries of GenAI. This includes robust privacy measures to safeguard against breaches that could lead to discrimination, identity theft, or other harm and swift remediation of harm when it occurs.

Human oversight and accountability, by individuals and organizations at every step in the development and use of GenAI. This includes recognizing that the technology is only a tool and cannot take accountability: humans must be accountable. Humans should demand GenAI to provide specific evidence when making a factual assertion.

Rationale for an *Oath for GenAI in Healthcare*

Policy and regulatory initiatives are worthy and *necessary* efforts. However, we believe those efforts – on their own – are *not sufficient* for four reasons, each one rooted in the longstanding – and currently relevant – views of human societies about health, disease, and healing enterprises.

Healing is a Sacred Endeavor

Sir Arthur C. Clarke's statement⁶ that "...advanced technology is indistinguishable from magic" was consistent with mankind's historical view of the healing arts. Dating back to our earliest civilizations, healing was not just magic. It was more than that. In both monotheistic and polytheistic cultures, it was one of the two gifts – the other being the law – granted to humankind by the divine.

law – granted to humankind by the divine.

Since those ancient times, both divine gifts have played crucial roles in providing a sense of security in the face of harsh unknowns, contributing to the orderly course of daily life.

Dating back to our earliest civilizations, healing was not just magic. In both monotheistic and polytheistic cultures, it was one of the two gifts – the other being the law – granted to humankind by the divine.



Laws governed how one person would engage with another, structuring contracts for property and commercial exchange. Laws guided judges as they resolved conflicts, set matters right, and provided restitution when one person suffered personal or financial harm at the hands of another.

The healing arts were important for other reasons. That's because health and disease were far more mysterious, and both were tools of divine forces. Good health was the reward for aligning with a god's will; disease resulted from personal or collective sin. Given that, to whom could one turn when illness struck? Who could intervene when lesser deities sent an illness? Could a mere human intervene to relieve the fever caused by the goddess *Febris*? Yes. However, their power to do so did not come solely from their knowledge, experience, or special potions. It came from their ability to call upon an even more powerful, benevolent deity.

By the time second-century Romans attributed fevers to *Febris*, the healer-divine partnership was already well established. Written evidence first appeared in 4,000 BCE Sumerian texts and in Buddhist texts of the same era. The idea stretches back before that, though; it was based on more ancient oral traditions. Centuries later, YHWH provided detailed health practices to Moses; by 900 BCE, Greeks had already built temple hospitals for various deities. Apollo, their most powerful god-physician, healed wounds as a reward, especially when patients were soldiers. In 300 BCE, the goddesses Scarlet and White gave acupuncture to China's Yellow Emperor. In 30 CE, Jesus of Nazareth directed others, as recorded by the physician-apostle Luke, to "cure the sick...". In 600 CE, the Prophet Muhammad called for compassionate care of the sick and instituted measures similar to those recorded by Moses. In 1500 CE, missionaries colonizing North America found Native and Shaman healing practices similar to European Christianity.

Thus, all known earlier civilizations had legends and lore describing divine actions in human health. Our modern world is no exception. More than 850 studies document the relationship between spirituality, health, disease, and recovery, and ironically, though today's clinicians would never ascribe a fever's origin to the goddess *Febris*, they continue to credit her each time they describe a patient's high fever as a *febrile* condition.

Healers Must Demonstrate Worthiness to Practice Their Arts

Given the uncertainties of health and illness and the power of invoking the mysteries of life and death, healers needed to demonstrate they were worthy of their community's trust. Tribal societies and small villages could attest to a healer's skill and ethical character by virtue of personal experience. Healers held a privileged place within the community; they earned it.

As societies grew large – as in ancient Greece's city-states – that was no longer possible. How could someone trust a healer without direct, personal experience? Even Hippocrates was an itinerant physician, traveling from one town to another, tending to the sick. How could he engender trust in each new location? To gain trust and consent to practice their arts, healers in that era needed to prove they were worthy. The method Hippocrates chose was an oath, a type of covenant statement familiar to the people of that era.

Covenants structured important relationships at the time. They existed between kings and their subjects, husbands and wives, and fathers and sons. Covenants had also long since structured the relationship between the Israelites and YHWH. The *Oath of Hippocrates* was the first of the covenants for those involved in the healing arts and remains the most famous. It reassured patients, communities, and other healers that, though a particular healer was unknown, they were worthy of trust because they ascribed to it. Oath-takers were special because they "rose above" ordinary human commerce to provide care for others regardless of personal, financial, or political concerns.

First framed in BCE 500, the *Oath of Hippocrates* evolved to suit the times. One ancient version calls for physicians to live with one another as in a fraternity, which would be unheard of and seems strange today. Other oaths for physicians followed, including the *Prayer of Maimonides* in 1793 and the *Oath of a Muslim Physician* in 1977. Similar oaths emerged for nurses in the 1893 *Nightingale Pledge* and pharmacists in the 1983 *Oath of a Pharmacist*. Though the specifics of each Oath might evolve with the era, as oaths, each one follows a typical framework faithful and unique to covenant statements.



Covenant Oaths Underpin the Healing Enterprise

Most people today have experience with both covenants (oaths) and contracts. However, few appreciate the differences between them and why covenants – not contracts – form the basis for a healthcare system that patients will trust and within which healers can thrive and do their best work. These differences are key to understanding why we

Ascribing to a covenant – that is, taking an oath – permanently and fundamentally alters the identity of the oath-taker.

propose an *Oath for GenAI in Healthcare* rather than a set of ethical guidelines or a regulatory framework.

A contract is a time-limited agreement between relatively equal parties that specifies the duties of each one. Only a mutual agreement can change them, and they determine how the parties, or judges whose help they seek, will resolve disputes if either party fails to perform. An example would be an agreement with a tailor to alter a garment. You expect the tailor to measure correctly and complete the work, paying only if the work is satisfactory. You do not expect the tailor to offer fashion advice about whether the color and style suit you, and if they do, you do not feel obligated to pay additionally for it. You also do not expect your tailor to interrupt a family dinner to help if you rip a seam as you dress for an evening event or to ship another set of clothes if an airline loses your luggage.

Covenants differ from contracts in meaningful ways. Covenants structure relationships when parties are unequal, such as between kings and subjects, parents and children, or physicians and patients. Ascribing to a covenant – that is, taking an oath – permanently and fundamentally alters the identity of the oath-taker. For example, a person becomes a parent when a child is born, remaining a parent whether they are awake, asleep, at work, or at rest. Even if their child should die, they will forever be a parent. The life-long nature of a covenant presumes the oath-taker's life will change in unanticipated ways, which is why duties can never be as detailed as those in a contract. It would be impossible, for example, for a parent to know what a child might need as they grow up; nevertheless, the covenant requires they meet those needs. Covenant oaths are 24/7/365 performance agreements that cover all eventualities.

It is why physicians – unlike tailors – will interrupt a

family dinner, leave a warm bed at midnight, and argue with insurance companies to secure coverage for a medicine. That is why so many physicians, nurses, and pharmacists who took similar oaths risked their own lives to care for COVID-19 patients. The obligations of a healer are too numerous and unpredictable to list in a covenant oath statement; the nature of the oath implies them.

It is easy to lose sight of this in today's commercialized (contractual) healthcare systems. A physician may have a *contractual relationship* with a healthcare system employer or payer, but they have a *covenantal relationship with their patients*.

Healers are Accountable to the Divine

Beyond the already-mentioned distinctions between contracts and covenants are two other important differentiators: witnesses and accountability.

First, contracts do not require witnesses. Covenant agreements, on the other hand, follow a format that invites witnesses, first by invoking the presence of the divine and then by naming the gathering of people at the oath-taking event. This will sound familiar to those who attend traditional wedding ceremonies, which create covenant agreements between spouses today.

Ancient versions of Hippocrates' Oath called upon "...Apollo the physician, by Aesculapius, Hygeia, and Panacea, and ... all the gods..." Today's medical oaths no longer invoke the names of Greek gods. Rather, they name a divinity or, in modern parlance, "all I hold sacred" or "all I hold dear." In nearly all U.S. medical schools today, this is done by first-year students who swear to an oath at "white coat" ceremonies. Next, the oath-taker invokes all those present. In medicine, when students at "white coat" ceremonies do likewise, those witnesses are family and friends; the event is second only to graduation in importance.

To be clear, the healers do not make promises to patients. They make promises first to the divine or what they hold sacred or dear and only then to a community of witnesses. Those promises have evolved to adapt to clinical knowledge and culturally-valued principles, but

Healers do not make promises to *patients*. They make promises first to the divine or what they hold sacred or dear and *only then* to witnesses.



the core principles are timeless: first, continuous improvement via education, training, leadership, and standard-setting; second, patient-centeredness, respecting privacy, confidentiality, and agency; third, collaboration with other healers and with patients; and fourth, ethical practices of beneficence, justice, and equitable access.

The oath concludes by calling upon the divine or what they hold sacred or dear, and the community witnesses to hold the healer accountable, saying, “If I keep this oath faithfully, may I be rewarded, and if I violate this oath, may the reverse be my lot.” In other words, the oath-taker invites rewards or punishments based on the degree to which they kept the promises they made publicly.

Unlike contracts, covenants do not require a judge to determine and enforce accountability. The oath-taker invites the divine, what they hold sacred or dear, and the witnessing parties’ community to pass judgment. Each component of an oath is meaningful, but the first and last elevate the relationship beyond those of transactional contracts. Ironically, those two sections – which make oaths self-enforcing – are the most underappreciated features of an oath and the source of its deepest meaning.

Covenants do not require a judge to determine and enforce accountability. The oath-taker invites the divine, what they hold sacred or dear to pass judgment.

This is why we believe an oath can be consequential for GenAI. As GenAI navigates toward broad acceptance, we believe its challenges are similar to those of the fledgling medical

profession in BCE 500 when it was newly emergent, based on knowledge known only to an elite few, and at a time before laws, regulations, or litigation could control it.

Our Proposal: An *Oath for GenAI in Healthcare*

Patients today are very much like their ancient counterparts. Our modern, sophisticated biomedical understanding of how a virus is replicating or where cancer is metastasizing provides no comfort to them and their loved ones. Rather, diseases remain as mysterious and frightening today as they were in

times past; many are as uncontrollable for us as they were for our ancestors. People need not be religious, honor the sacred, or pray; when they are sick, they want help. Better yet, they want it from someone they can trust to align with more powerful, benevolent forces than those of illness they face. Oaths provide that reassurance.

We direct our proposal equally to individuals and organizations, whether they are developing and/or deploying GenAI in healthcare, regardless of whether or not they engage directly with patients.

The Hippocratic Oath and similar oaths taken by nurses and pharmacists are important pledges modern clinicians continue to honor. Those oaths outline principles that remain relevant in our 21st century and for GenAI applications. Therefore, we have built upon those versions of oaths to craft an *Oath for GenAI in Healthcare*.

Though we believe most GenAI actors may not realize it, they stepped into a healing stream that has been flowing since before recorded history. Whether they use GenAI for basic research, health system management, technology development, market research, patient advocacy, journalism, thought leadership or policy-making does not matter. Nor does it matter if they are AI company executives or investors. They will significantly impact patients and populations, so we believe they should embrace an identity as healers and ascribe to a healer’s oath.

We direct our proposal equally to individuals and organizations, whether they are developing and/or deploying GenAI in healthcare and regardless of whether or not they engage directly with patients. We stop short of recommending how and when they should take this *Oath*. We believe it is more important to explore the nature of an oath first. We know there will be differences of opinion about whether one is necessary and, if so, what it should entail. For that reason, despite our strong views on the matter, we propose this as a starting point for what we feel is a worthy discussion and hope it is productively catalytic.

Merely ascribing to an *Oath for GenAI in Healthcare* will not resolve all the issues that need to be addressed, but it will help tremendously. Failure to do so will surely contribute to the further decline of patients’ trust in the healthcare system to care for



them. GenAI may be one of the most far-reaching opportunities to do good, but perhaps not. That being the case, let's make it work for the good of those who depend upon the healers they need. They have entrusted us to care for them and those they love as if it were our sacred obligation. They may not be able to articulate perspective as we have here, but they feel it instinctively and viscerally.

Our proposal for an *Oath for GenAI in Healthcare* might seem radical to those unfamiliar with the history of healing. In reality, it is quite conservative, with a millennia-long tradition. It might also seem daunting

and burdensome, and though it may be, it is the only way to own a "place at the table" for those using GenAI who desire to serve patients.

Reimagining and enlivening covenants and applying them to GenAI in healthcare is essential. Without such an attempt, we will continue to drift in a direction that strays from the sound covenantal origins of healing that underpins all of healthcare globally. Such a direction risks GenAI not realizing its promises to heal but seeing it collapse under the weight of conflict, confusion, litigation, and regulatory barriers that will ensue.

Our proposal for an *Oath for GenAI in Healthcare* might seem radical to those unfamiliar with the history of healing. In reality, it is quite conservative, with a millennia-long tradition. It might also seem daunting and burdensome, and though it may be, it is the only way to own a "place at the table" for those using GenAI who desire to serve patients.



Oath Principles and Text Comparisons

Source Oath	Text	Principle	GenAI
Hippocrates	I swear by Apollo, the physician, by Aesculapius, Hygeia, and Panacea, and I take to witness all the gods, all the goddesses, to keep according to my ability and my judgment the following Oath.	Sacred Witness	I swear (or affirm) by all I hold dear and sacred, and in the presence of those assembled here:
Maimonides	All-bountiful! Thou hast formed the human body in Thy complete wisdom.		
Islamic	I swear by God the Great; To regard God in carrying out my profession.		
Nightingale	I solemnly pledge myself before God and in the presence of this assembly:		
Hippocrates	To consider dear to me as my parents, he who taught me this art; to look upon his children as my own brothers, to teach them this art if they so desire without fee or written promise; to impart to my sons and the sons of the master who taught me who have agreed to the rules of the profession and its precepts.	Continuous Improvement Through Education, Training, Leadership, and Standard Setting	To commit to life-long learning, sharing knowledge with others, and improving the standards by which my profession operates, as it incorporates GenAI cognizant of its risks as well as benefits.
Maimonides	To revere my teacher, teach my junior, and be brother to members of the medical profession joined in piety and charity; To strive in the pursuit of knowledge and harnessing it for the benefit, but not the harm, of mankind. Permit not the thought to awaken in me: You know enough; but grant me strength, leisure and the urge to enlarge my accomplishments and to add to others. If wiser artists seek to improve and instruct me, let my spirit be thankful; for great is the field of the art. When, however, conceited fools berate me, then let the love of the art steel my spirit and insist on truth, regardless of age, fame or standing, for to retract would mean death and disease.		
Pharmacy	I will keep abreast of developments and maintain professional competence in my profession of pharmacy. I will embrace and advocate change in the profession that improves patient care.		
Nightingale	I will do all in my power to maintain and elevate the standard of my profession		
American College of Healthcare Executives	Contribute to the advancement of our profession by exemplifying competence and leadership; Commit to lifelong learning by maintaining a personal program of continuing education; Enhance our profession through leadership in a wide range of community and professional activities.		



Source Oath	Text	Principle	GenAI
Hippocrates	I will prescribe regimen for the good of my patients according to my ability and my judgment and never do harm to anyone. To please no one will I prescribe a deadly drug, nor give advice which may cause his death.	Patient-Centeredness	To do my personal best and encourage others on my teams and throughout my organization, to offer our skills to serve patients and their wellbeing.
Nightingale	I will devote myself to the welfare of those committed to my care.		To ensure my personal and organizational interests and financial gains do not eclipse the care of my patients.
Islamic	To protect human life in all stages and under all circumstances, doing my utmost to rescue it from death, malady, pain and anxiety;		
Hippocrates	I will preserve the purity of my art and my life. I will not cut for stone, leaving this operation to be performed by specialists in this art. In every house where I come, I will enter only for the good of my patients, keeping myself far from all intentional ill-doing.	Collaboration with Other Healers Working Within the Sphere of Knowledge	Intending only to do good, ensure the expertise of those involved, commit to improving life, contribute what I know, and collaborate with others to ensure the best outcomes for those I serve.
Nightingale	With loyalty, will I endeavor to aid the physician in his work,		
Maimonides	Grant my patients confidence in me and my art, and imbue them with obedience to follow my precepts and directions. Ban from their bedside all quacks and the army of advice-giving relatives and too-wise nurses, for they are a terrible band, who, through their vanity, harm the best intentions of the healing art and frequently cause the death of Thy creatures.	Collaboration with Patients	To encourage those in my care, their clinicians, and others engaged in healthcare to appropriately use the tools I build and use, welcoming their feedback to support improvements to benefit them and others. To be always mindful that healing is a human enterprise, and humans must be informed, aware, and in control of each step, lest we lose the trust upon which we depend for the therapeutic alliance.



Source Oath	Text	Principle	GenAI
Hippocrates	All that may come to my knowledge in the exercise of my profession or outside of my profession or in daily commerce with men, which ought not be spread abroad, I will keep secret and will never reveal.	Practice Ethically	To maintain the highest moral, ethical, technical, financial, legal, and research conduct, including: Not making false claims about my product. Keeping information entrusted to me private, confidential, and secure lest my patients suffer further from any kind of bias and the whims of online multitudes. Recognizing that the tools of GenAI do greatly benefit patients, yet may also cause harm. Therefore, I will guard that my art is not reduced to an algorithm, and I will ensure the compassion for my patients is not eclipsed by a machine or technology
Islamic	To keep people’s dignity, cover their privacies, lock up their secrets		
Pharmacy	I will maintain the highest principles of moral, ethical and legal conduct. I will consider the welfare of humanity and relief of suffering my primary concerns. I will apply my knowledge, experience, and skills to the best of my ability to assure optimal drug therapy outcomes of the patients I serve.		
Nightingale	To pass my life in purity and to practice my profession faithfully; I will abstain from whatever is deleterious and mischievous and will not take or knowingly administer any harmful drug; I will hold in confidence all personal matters committed to my keeping and all family affairs coming to my knowledge in my calling		
American College of Healthcare Executives	Abide by its Code of Ethics		
Hippocrates	If I keep this oath faithfully, may I enjoy my life and practice my art, respected by all men and at all times, but if I swerve from it or violate it, may the reverse be my lot.	Recognize and Accept Consequences	If I keep this oath faithfully, may I enjoy my life and work and always be respected. If I violate this oath, may the reverse be my lot.
Pharmacy	I take these vows voluntarily with the full realization of the responsibility with which I am entrusted by the public.		

With gratitude, we acknowledge the many patients, caregivers, clinicians, healthcare company executives, academics, and visionaries who reviewed this work and offered informed and insightful feedback during its development, especially Tatiana Barakshina, PhD, Sandra Beatty, Trishna Bharadia, Bill Bloom, Dan Bonner, MBA, Denise Brown, Janet Carlson, former Ambassador Lisa Gable, Tim Glennon, Dorothy Hoffman, MPP, Inder Jaggi, EMBA, Richard 'RJ' Kedziora, MBA, Aman Khera, EMBA, Matt Kicinski, Jonathan Manierre, Ian McNeil, Jill Michal, Thomas Smith, Robin Strongin, MPA, Robert Sundelius, FACHE, Mike Taylor, Frank Tennigkeit, PhD, Martin Wasserman, MD, JD and Shadi Yazdan, MBA.

We extend special thanks to Sidi Lemine, MSc in AI, for expert advice and counsel.

Carrie Brubaker, PhD and Glenna Crooks, PhD

November 2024

To learn more, we invite you to visit oathgenai.com and to connect with us directly at carrie.brubaker@medaffairsconsultancy.com | glenna@glennacrooks.com



ENDNOTES

- 1) Wolters Kluwer Health survey, Generative AI in Healthcare: Gaining Consumer Trust, conducted online between October 27 and November 2, 2023, with 1,000 U.S. adults ages 18 and older. <https://www.wolterskluwer.com/en/news/wolters-kluwer-survey-finds-americans-believe-genai-is-coming-to-healthcare-but-worry-about-content>. Last accessed 2 September 2024.
- 2) Brender, T. D., Smith, A. K., & Block, B. L. (2024). Can Artificial Intelligence Speak for Incapacitated Patients at the End of Life? *JAMA Internal Medicine*, 184(9), 1005–1006. <https://doi.org/10.1001/jamainternmed.2024.2676>
- 3) One study posed 39 medication questions to ChatGPT, resulting in no or incomplete responses in 74% of its answers. It missed drug-drug interactions and cited fabricated references and nonexistent studies. Grossman, S., Zerilli, T., & Nathan, J. P. (2024). Appropriateness of ChatGPT as a Resource for Medication-Related Questions. *British Journal of Clinical Pharmacology*, 90(10), 2691–2695. <https://doi.org/10.1111/bcp.16212>
- 4) Panacea (PANAKEIA) was the Greek goddess of cures and panaceas in the form of medicines, salves, and other curatives. Her father was the medicine-god Asclepius, and her sisters were Laso (curing) and Hygeia (good health), from which we draw our notion of hygiene.
- 5) Marchal, N., Xu, R., Elasmir, R., Gabriel, I., Goldberg, B., & Isaac, W. (2024). Generative AI Misuse: A Taxonomy of Tactics and Insights from Real-World Data. arXiv preprint arXiv:2406.13843.
- 6) Clark, Sir Arthur, *Profiles of the Future: An Inquiry into the Limits of the Possible*, Henry Holt & Co, 1984.

