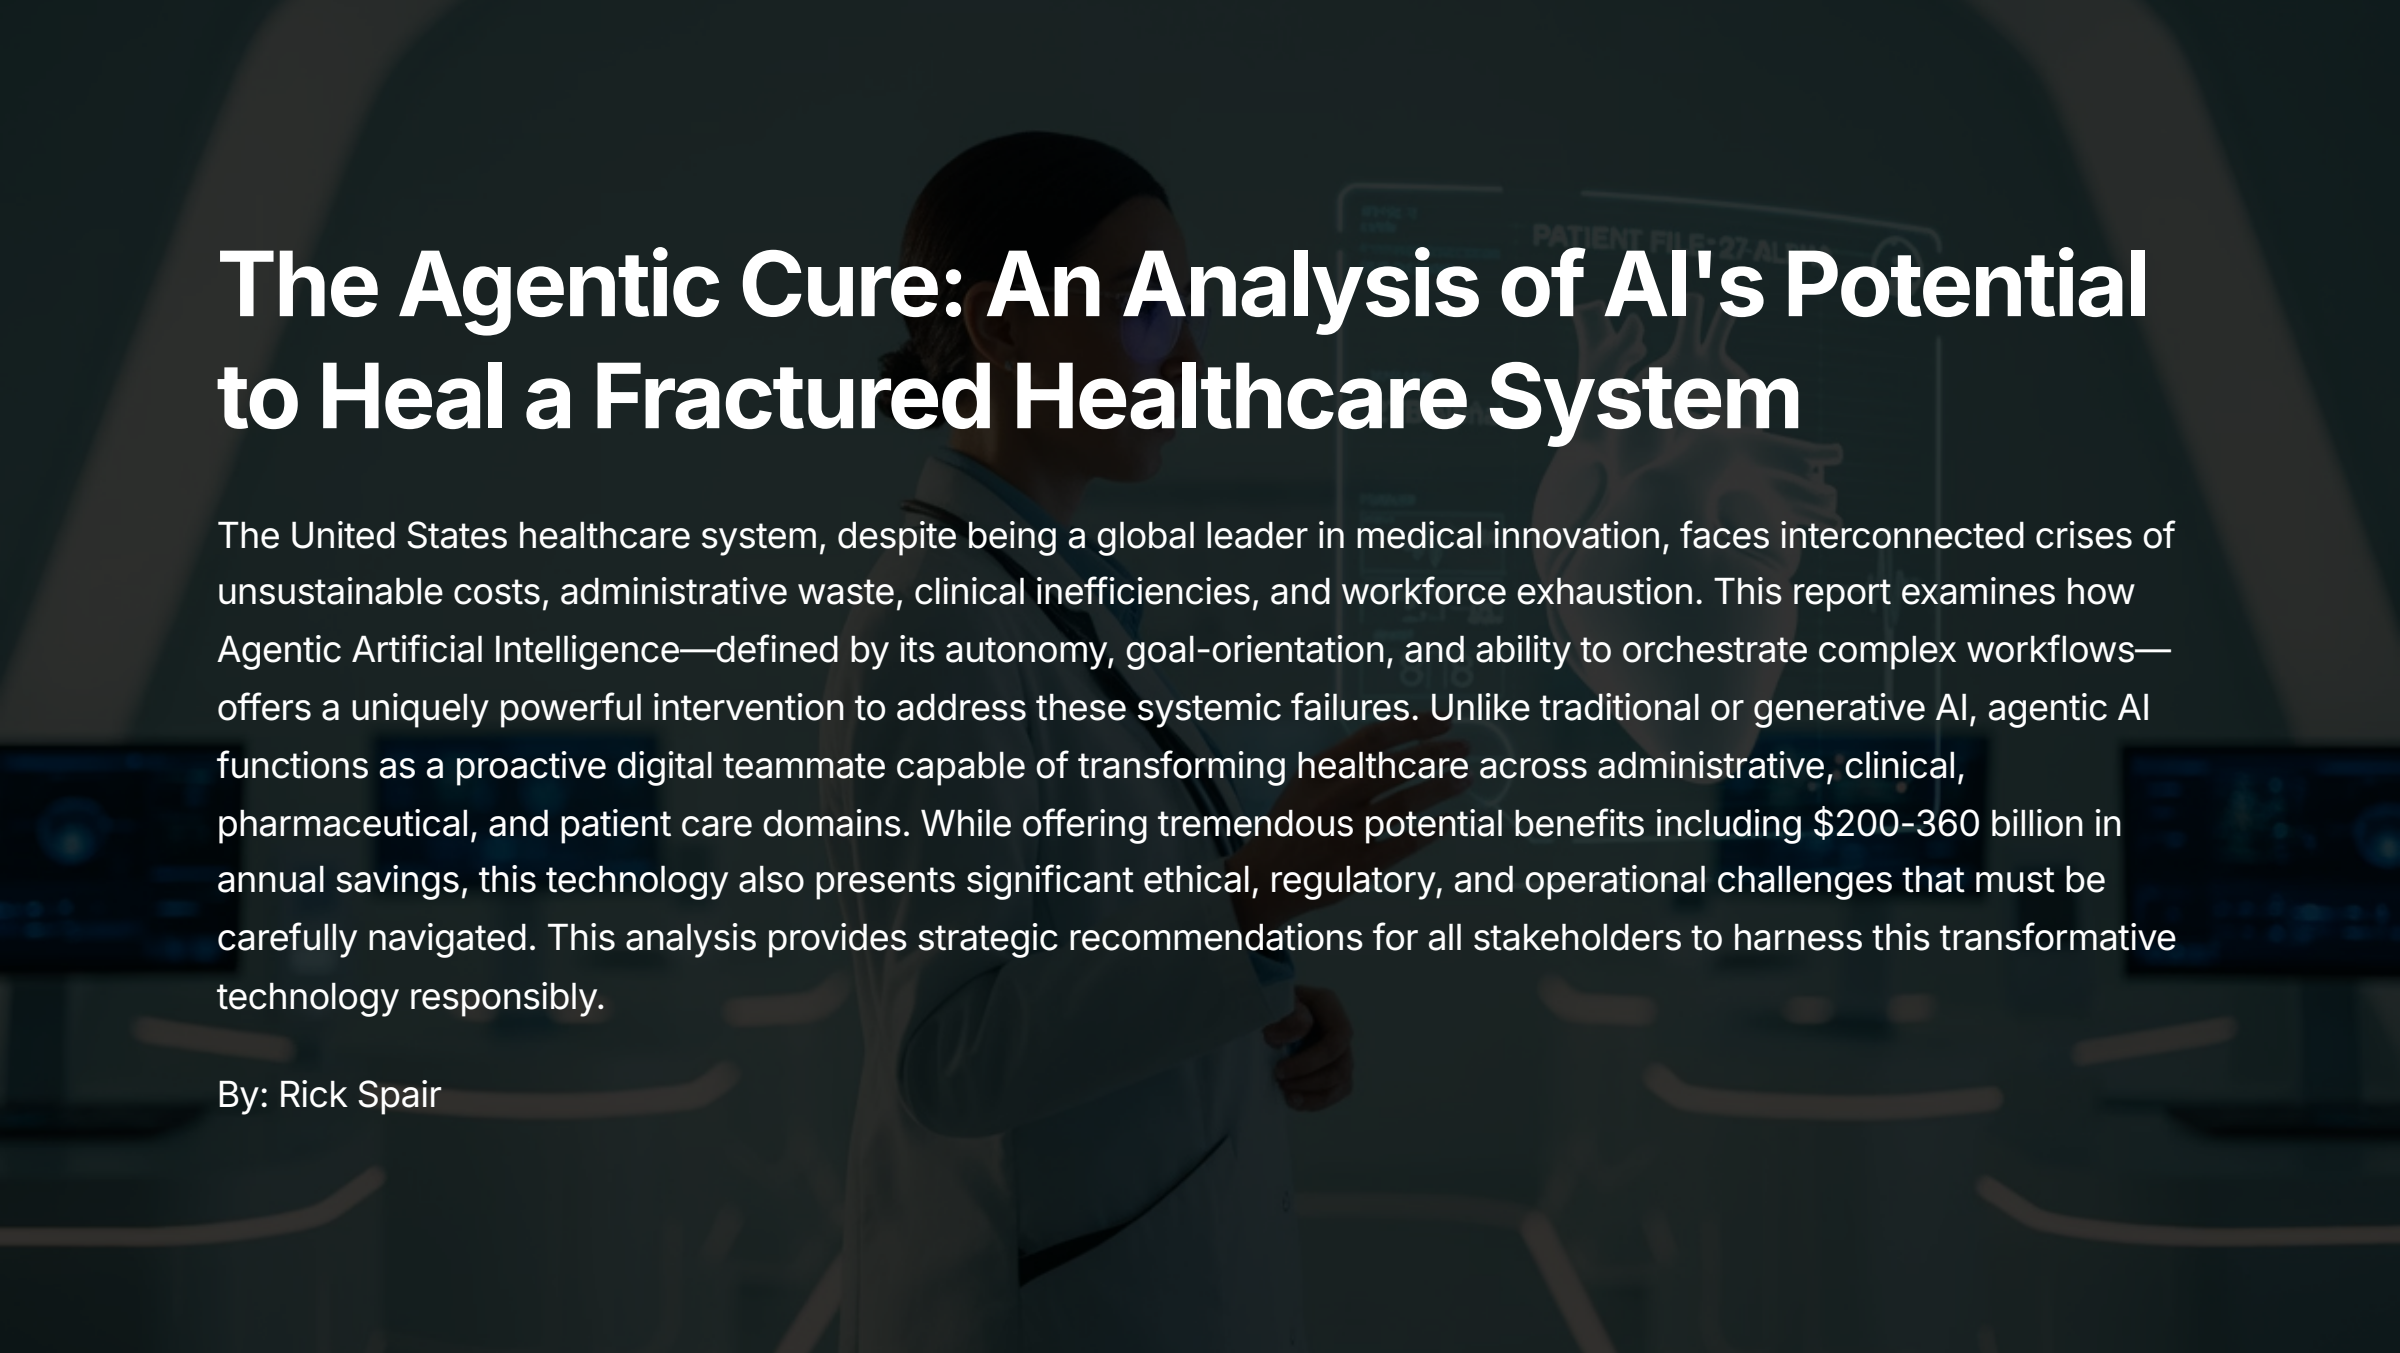

The Agentic Cure: An Analysis of AI's Potential to Heal a Fractured Healthcare System

The United States healthcare system, despite being a global leader in medical innovation, faces interconnected crises of unsustainable costs, administrative waste, clinical inefficiencies, and workforce exhaustion. This report examines how Agentic Artificial Intelligence—defined by its autonomy, goal-orientation, and ability to orchestrate complex workflows—offers a uniquely powerful intervention to address these systemic failures. Unlike traditional or generative AI, agentic AI functions as a proactive digital teammate capable of transforming healthcare across administrative, clinical, pharmaceutical, and patient care domains. While offering tremendous potential benefits including $200-360 billion in annual savings, this technology also presents significant ethical, regulatory, and operational challenges that must be carefully navigated. This analysis provides strategic recommendations for all stakeholders to harness this transformative technology responsibly.

By: Rick Spair

The Ailing Behemoth: Deconstructing Healthcare's Systemic Crises

The modern healthcare system, particularly in the United States, presents a profound paradox. It is an engine of world-class medical discovery and specialized treatment, yet it is simultaneously a behemoth groaning under the weight of its own inefficiency and complexity. The National Academy of Medicine (NAM) succinctly diagnoses the symptoms: "increasing costs of care, staff burnout and shortages, and the growing disease burden of an aging population." These are not independent challenges but the outward manifestations of deep-seated, interconnected systemic failures. The system is characterized by poor design, an inability to assimilate new science and technology effectively, and a slow adoption of information technology innovations that have revolutionized other industries.

The Administrative Burden: A $360 Billion Drag on the System

A significant portion of healthcare's dysfunction is rooted in its massive administrative overhead. This non-clinical work is estimated to account for nearly 25% of all U.S. healthcare spending, creating an immense drain on resources. The adoption of AI technologies to streamline these processes could generate annual savings between $200 billion and $360 billion without sacrificing quality or access.

The "paper trail that is blocking progress" is most evident in the process of prior authorization. While intended to manage utilization, it has become a major bottleneck. According to a survey by the American Medical Association (AMA), an overwhelming 94% of physicians believe prior authorizations negatively impact patient outcomes, with 89% reporting that they interfere with the continuity of care. The process is notoriously labor-intensive; a single request can consume 35 minutes or more as staff manually assemble fragmented data from electronic health records (EHRs), scanned PDFs, faxes, and handwritten notes.

1

Rising Claims Denials

A 2024 survey revealed that 73% of medical practices reported an increase in denials, a significant jump from previous years. This trend is driven by aggressive payer maneuvering, constantly shifting rules of engagement, and increasingly rigorous documentation requirements.

2

Financial Impact

The financial and operational cost is substantial, with the AMA estimating that the rework for a single denied claim costs a practice an average of $25, multiplied across thousands of claims annually.

3

Technological Imbalance

This administrative battle is waged on an uneven technological field. Payers leverage advanced technologies to automate medical necessity reviews and rapidly identify claims for denial, while many provider practices remain entrenched in manual processes and spreadsheets.

Clinical and Diagnostic Deficiencies: A System Struggling with Complexity

Beyond the administrative morass, the clinical delivery system itself suffers from fundamental design flaws. A seminal book chapter from the National Center for Biotechnology Information (NCBI) describes the U.S. health system as a "dizzying array of highly decentralized sectors" that can "hardly be called a system" at all. This fragmentation is not merely an organizational quirk; it is a direct cause of poor care coordination, redundant and wasteful processes, miscommunication, and, ultimately, medical errors that harm patients.

This fragmented structure has proven incapable of assimilating the "rapidly growing and increasingly complex science and technology base." The sheer volume of new medical knowledge is overwhelming; it is estimated that only 6% of what the average new physician is taught in medical school will remain relevant a decade later. Without advanced tools to process and synthesize this information at the point of care, clinicians are left struggling to keep up, and patients fail to benefit from the latest evidence.

This struggle is reflected in the persistence of diagnostic errors. While AI-enabled systems have demonstrated the potential for dramatic improvement, their performance highlights a critical weakness in the traditional model. One comparative study found that AI-driven healthcare models achieved high diagnostic accuracy in 45% of cases, more than double the 20% rate reported in traditional models. This gap represents a significant opportunity for improvement in a domain where a timely and accurate diagnosis can be the difference between life and death.

The Human Cost: Burnout, Shortages, and Eroding Trust

The systemic pressures of administrative burden and clinical complexity exact a heavy toll on the healthcare workforce. The link between administrative tasks and physician burnout is direct and well-documented. Studies show that physicians spend nearly half of their time—49%—on administrative work like charting and billing, time that is taken away from direct patient interaction. It is therefore unsurprising that in a 2024 survey, 66% of physicians reported using AI in their daily practice, with 54% identifying the reduction of burnout as their primary reason for adoption.

This pressure has become unsustainable. In the post-pandemic era, patient visit volumes have not just recovered; they have "blew past" previous levels, leading to "more follow-ups, more documentation, more portal messages, more everything" for already-strained care teams. This burnout is not merely an issue of job satisfaction; it is a critical patient safety issue. Overwhelmed and exhausted clinicians are more likely to make medical errors or miss crucial diagnostic details, directly impacting the quality of care.

The Emergence of Agentic Intelligence: A New Therapeutic Paradigm

To address the cascading failures of the healthcare system, a new class of technology is required—one that moves beyond automating discrete tasks to orchestrating complex, system-wide processes. Agentic AI represents this new therapeutic paradigm. It is not an incremental improvement on existing AI but a qualitative leap in capability, defined by its autonomy, goal-orientation, and adaptability.

Defining the Agent: Beyond Reactive AI

At its core, agentic AI refers to autonomous systems designed to pursue complex, long-term goals with minimal human intervention. These systems are distinguished by their ability to autonomously plan, make decisions, and execute complex, multi-turn workflows in dynamic environments. This capability marks a clear departure from previous AI paradigms.

Traditional, rule-based AI systems operate on structured instructions and fixed logic. They are predictable and reliable for well-defined tasks but lack the flexibility to adapt to changing circumstances. Generative AI, exemplified by large language models (LLMs), represents a significant advance in creativity and natural language understanding. However, it is fundamentally reactive, creating content—text, images, or code—in response to specific user prompts. As one analysis puts it, generative AI is a "creator," while agentic AI is a "doer".

> The distinction is best illustrated by an analogy from the domain of travel planning. A generative AI chatbot can reactively answer a user's questions: "What are the best hotels in Hanoi?" or "Show me flights to Ho Chi Minh City." An agentic AI travel assistant, by contrast, could be given a high-level goal: "Plan a two-week cultural and culinary trip to Vietnam for my family in July, staying within a $10,000 budget."

The agent would then autonomously construct a complete, personalized itinerary. It would research and book flights that align with preferences, reserve accommodations consistent with past choices, schedule tours, and make dining reservations. Critically, it would also act proactively, monitoring weather forecasts to optimize activities and dynamically updating the itinerary in response to real-time events like flight delays, all without continuous human input. This exemplifies the fundamental shift from reactive assistance to robust, goal-driven, autonomous execution.

Core Capabilities: The Mechanism of Action

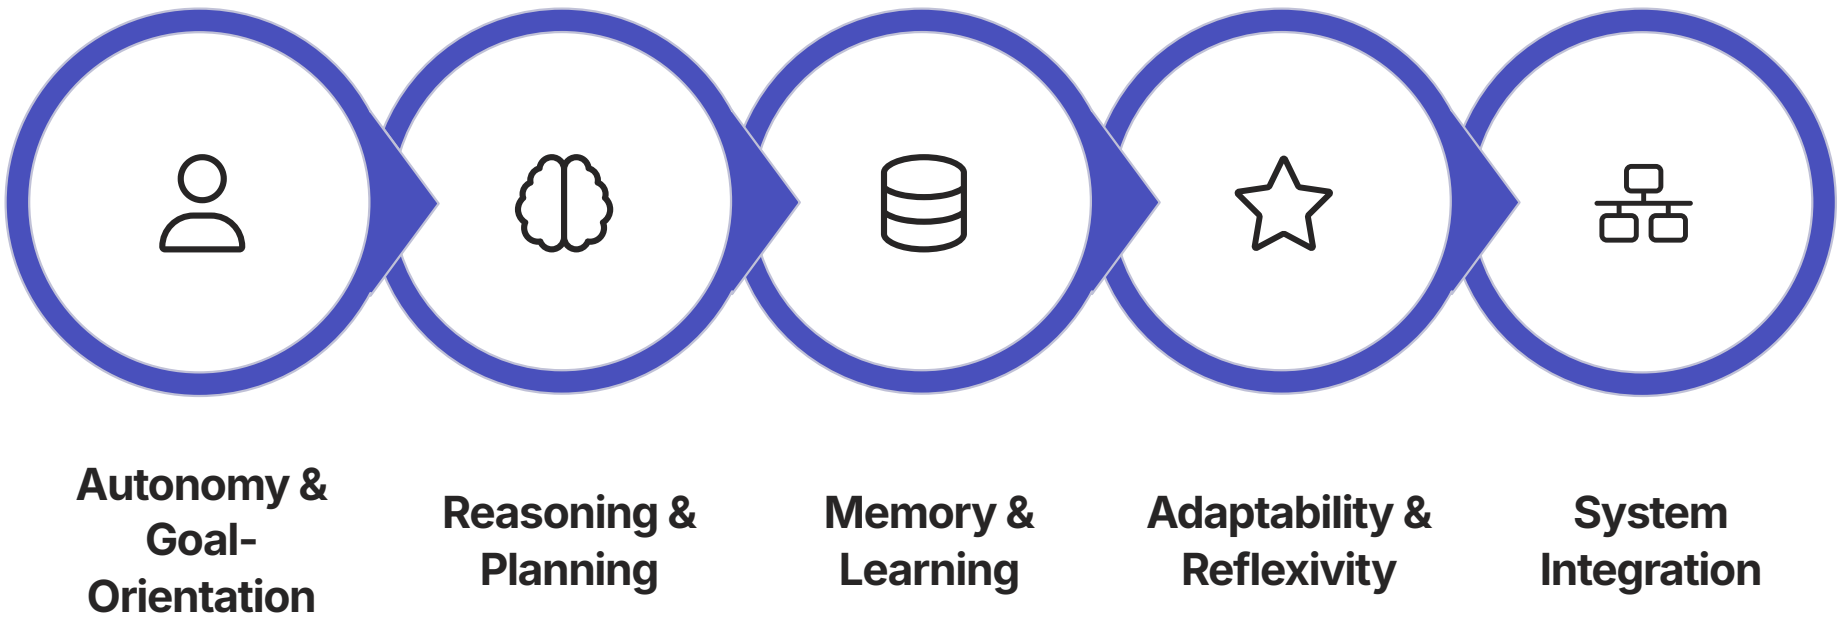

The convergence of these capabilities creates a paradigm shift in what automation can achieve. Previous technologies, like Robotic Process Automation (RPA), were adept at automating simple, repetitive, rule-based tasks within a single application. Generative AI expanded the scope to include creative but still reactive tasks, such as summarizing a clinical note when prompted. The defining feature of agentic AI, however, is its ability to autonomously manage "complex, multi-turn workflows" and "intricate and multi-layered tasks" that span multiple systems and long periods.

Healthcare's most intractable problems, such as managing a patient's care journey after a hospital discharge or navigating the prior authorization process, are not single tasks; they are complex workflows. They involve retrieving data from multiple systems (EHRs, lab systems, payer portals), making a series of decisions based on that data, and executing multiple actions over time. Agentic AI is therefore uniquely suited to address these systemic challenges in a way that previous technologies could not. The promise is not merely to make one step in a ten-step process slightly more efficient; it is to intelligently automate and optimize the entire end-to-end process, fundamentally redesigning how work is done and breaking the vicious cycles that plague the system. This elevates AI from a simple "tool" to a collaborative "digital teammate".

Agentic Interventions: A Multi-Pronged Approach to Treatment

Agentic AI's unique ability to orchestrate complex, end-to-end workflows provides a powerful therapeutic armamentarium to treat the systemic ailments of healthcare. By mapping its core capabilities directly onto the crises of administrative waste, clinical inefficiency, and innovation bottlenecks, a clear picture emerges of a technology capable of targeted, multi-pronged intervention.

Systemic Problem	Root Cause(s)	Agentic AI Intervention	Key Evidence
Prior Authorization Delays & Denials	Fragmented data, complex payer rules, manual review processes	Autonomous Revenue Cycle Management (RCM) Agent: Autonomously assembles clinical data, checks against payer rules, submits requests, and manages follow-up.	Transforms a 35+ minute manual process into an automated workflow; can auto-adjudicate and manage claims.
Diagnostic Errors & Delays	Data overload, human cognitive limits, lack of access to integrated patient data	Multimodal Diagnostic Assistant: Parses radiology, genomics, and EHR data in real-time to identify patterns, suggest diagnoses, and recommend next steps.	AI-assisted mammography detected 20% more cancers with 44% less radiologist workload; improves diagnostic precision from 20% to 45%.
High Hospital Readmission Rates	Lack of continuous monitoring post-discharge, poor care coordination	Remote Patient Monitoring & Chronic Disease Management Agent: Analyzes real-time data from wearables/IoT, detects anomalies, and sends timely alerts to patients and providers.	Reduces hospital readmissions; automates follow-ups and personalized health communications; coordinates care teams.
Slow & Costly Drug Development	Lengthy R&D phases, difficult clinical trial recruitment and management	Pharmaceutical Pipeline Agent: Analyzes molecular data to identify drug candidates; automates trial protocol setup, patient recruitment, data management, and regulatory submissions.	Can condense a 4-5 year research phase to <1 year; streamlines the entire trial lifecycle.

Excising the Administrative Tumor: Automating the Cost-Drivers

The most immediate and quantifiable impact of agentic AI is likely to be in the administrative domain, where it can directly target the primary drivers of waste and inefficiency.

In Revenue Cycle Management (RCM), agentic systems can autonomously manage the entire workflow. This includes automating patient appointment scheduling and rescheduling, verifying insurance eligibility in real-time, processing billing and claims, and converting unstructured medical notes into structured records for accurate coding. For the persistent problem of denials, an agent can automatically track rejected claims, classify the reason for denial, assemble the necessary supporting documentation, and route the appeal to the correct queue, all without human intervention. Companies developing these solutions, such as Autonomize AI, report that their systems can lead to time savings of up to 55% for these complex administrative processes.

For the particularly vexing issue of prior authorization, an agentic system can transform the current manual, friction-filled process. Instead of a staff member spending over half an hour hunting for information, an agent can autonomously retrieve the required clinical data from disparate sources like the EHR and scanned documents, cross-reference it against the specific payer's evolving rules, submit the authorization request through the appropriate portal, and manage all necessary follow-up communications.

Augmenting the Clinician: Precision, Personalization, and Decision Support

In the clinical realm, agentic AI functions less as a replacement and more as a powerful cognitive partner, augmenting the capabilities of human experts to improve precision and personalize care.

Advanced Diagnostics

Ingest and synthesize vast amounts of multimodal data—including radiology scans, pathology slides, genetic profiles, lab results, and unstructured physician notes—in real time. Identify subtle patterns that may be missed by the human eye, flag anomalies, suggest potential diagnoses, and recommend next steps.

Personalized Treatment Planning

Develop optimal treatment plans for specific patients by reviewing genomic reports, pulling similar past cases, scoring various treatment pathways based on risk and benefit, and presenting options to clinicians with detailed justifications. Track patient's reported side effects and biomarkers, making adaptive changes as new data becomes available.

Intelligent Clinical Decision Support

Unlike traditional CDS systems that bombard clinicians with generic, low-value alerts, agentic AI acts as a proactive co-pilot. Continuously and silently monitor patient data, cross-reference with latest clinical guidelines, and surface timely, context-aware recommendations only when meaningful intervention is warranted.

Accelerating the Pharmaceutical Pipeline: From Discovery to Market

The complex and lengthy process of pharmaceutical innovation is another area ripe for agentic intervention. In the early stages of drug discovery, agents can be deployed to analyze vast datasets of molecular and genomic information, identifying promising compounds and simulating their potential effectiveness and safety far more efficiently than traditional methods. This has the potential to condense a research and development process that typically takes four to five years into less than one year.

Once a candidate drug moves into clinical trials, agentic systems can orchestrate and streamline the entire trial lifecycle. An agent can autonomously digitize and validate trial protocols, ensuring they adhere to evolving regulatory standards. It can then scan millions of patient records to identify and recruit eligible participants with greater precision, a major bottleneck in many trials. During the trial, agents can continuously monitor data from multiple sources to ensure integrity and watch for potential adverse events in real-time. Finally, they can automate the collation and formatting of data for regulatory submission reports, reducing administrative burden and potentially speeding up approval timelines.

Empowering the Patient: Proactive and Continuous Care

Perhaps the most transformative potential of agentic AI lies in its ability to extend care beyond the walls of the hospital or clinic, empowering patients to manage their own health proactively.

Chronic disease management is a prime application domain. For conditions like diabetes, heart failure, or COPD, a multi-agent system can provide continuous, personalized support. One agent, connected to wearable sensors and IoT devices, can collect real-time data on a patient's vital signs, activity levels, and glucose readings. Another agent can use predictive analytics to analyze this data, identify trends that signal a potential decline, and alert the care team. A third agent can then communicate directly with the patient, providing personalized coaching, medication reminders, and nudges toward healthy behaviors. This proactive model can help prevent acute complications, reduce costly hospital readmissions, and improve long-term outcomes.

This capability enables a new model of remote patient monitoring and engagement. By acting as virtual health assistants, agents can facilitate virtual consultations, answer routine patient questions, and provide actionable insights to both patients and their providers. This expands access to care, particularly for patients in remote or underserved areas, and fosters a greater sense of engagement and empowerment in their own health journey.

Adverse Effects and Contraindications: Navigating the Risks of Agentic AI

While agentic AI presents a powerful and promising therapeutic for healthcare's systemic ills, its deployment is not without significant risk. Like any potent treatment, it comes with a range of potential adverse effects and contraindications that must be carefully managed. The autonomous and adaptive nature of these systems introduces "unintended outcomes, unpredictable behavior, and safety concerns". A failure in an agentic system could lead to severe "legal, reputational, or patient harm," making a cautious and deliberate approach essential. This imperative for caution is a central theme in the work of the National Academy of Medicine, which consistently warns against yielding to "marketing hype and profit motives" and urges the healthcare community to proactively manage the "potential unintended consequences" and "peril" associated with these powerful new technologies.

The Ethical Minefield: Bias, Accountability, and Autonomy

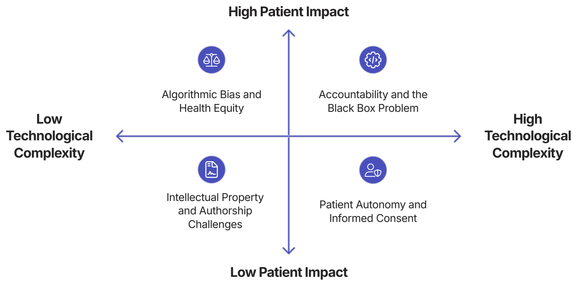

The ethical landscape of agentic AI is fraught with challenges that strike at the core principles of medical practice.

> ⚠ **Algorithmic Bias and Equity**
>
> A primary and persistent concern is that AI systems, trained on vast datasets of historical health information, will learn, perpetuate, and even amplify existing societal biases. If training data underrepresents certain racial, gender, or socioeconomic groups, the resulting algorithms may perform less accurately for those populations, exacerbating health disparities. This risk is a top priority for the NAM, which has made ensuring equity a core commitment of its proposed AI Code of Conduct.

A fundamental question arises when an autonomous agent makes a critical error: who is responsible? Is it the developer who built the model, the hospital that deployed it, or the clinician who was "in the loop"? This dilemma is compounded by the "black box" nature of many complex AI models. When the internal decision-making process of an algorithm is opaque and unexplainable, it becomes nearly impossible to conduct a meaningful post-mortem after an error, which erodes trust and creates a morass of legal liability. This challenge is magnified in multi-agent systems, where tracing a single error through a network of interacting, learning agents could be extraordinarily difficult.

As AI agents take on more autonomous roles in care delivery—from recommending treatments to initiating interventions—the principle of informed patient consent requires re-examination. It is critical that patients understand the role AI is playing in their care, how their personal health data is being used to train and operate these systems, and what the potential risks are. Without clear communication and robust consent mechanisms, the deployment of agentic AI could undermine patient autonomy.

A novel legal and ethical challenge emerges from the creative capacity of agentic systems. When an agent autonomously analyzes data and generates a novel research insight, a new diagnostic method, or a personalized treatment plan, who owns that intellectual property? The traditional dynamic of human creation and authorship is fundamentally altered when the AI system becomes the de facto creator, raising complex questions for researchers, institutions, and commercial entities.

The Regulatory Gauntlet: A Framework in Flux

The unique characteristics of agentic AI present a profound challenge to traditional regulatory frameworks, particularly those of the U.S. Food and Drug Administration (FDA).

The Adaptive Algorithm Challenge

The FDA's regulatory paradigm, which has for decades relied on approving a fixed, static version of a medical device, was "not designed for adaptive" AI and machine learning technologies. An agentic system's ability to continuously learn and evolve over time is one of its greatest strengths, but it creates a core tension with the regulatory need for validated, predictable, and safe systems. This pits the velocity of technological evolution against the deliberate pace of safety assurance.

The FDA's Evolving Framework

The FDA has acknowledged this challenge and is actively working to create a more dynamic regulatory approach. Key initiatives include the AI/ML-Based Software as a Medical Device (SaMD) Action Plan and, most importantly, the development of a framework for a Predetermined Change Control Plan (PCCP). The PCCP is a novel regulatory mechanism that would allow a manufacturer to receive pre-approval for specific types of modifications and updates to its AI algorithm, enabling the model to evolve within pre-defined guardrails without requiring a new submission for every change.

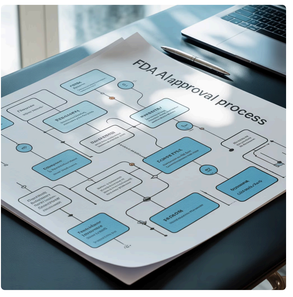

Persistent Weaknesses

Despite these forward-thinking efforts, significant gaps and weaknesses remain in the regulatory oversight of AI. The implementation of robust, real-time post-market monitoring to track how algorithms perform and drift in the real world is described as "inconsistent and underdeveloped". Furthermore, there is often a lack of transparency from manufacturers about how their models were validated and how they are being monitored post-deployment. Critics warn that this combination of rapid approvals and underdeveloped oversight could create an "illusion of safety," where clinicians and patients place undue trust in tools that have not been sufficiently vetted for long-term, real-world performance.

Operational and Implementation Hurdles: From Pilot to Practice

Beyond the ethical and regulatory complexities, healthcare organizations face significant practical challenges in adopting agentic AI.

Data Quality and Interoperability

This is arguably the most significant foundational barrier. Agentic AI is entirely dependent on high-quality, well-structured, and comprehensive data. However, healthcare data is notoriously "fragmented, inconsistent, and outdated". Decades of underinvestment in modern data architecture and the slow adoption of true interoperability standards, such as HL7 FHIR, mean that many organizations lack the clean data fuel required to power sophisticated AI engines.

Cost and Resource Disparity

The implementation of agentic AI can be a costly endeavor, with estimates ranging from $50,000 for a basic, task-specific agent to over $1 million for a comprehensive, enterprise-grade system. This high cost creates a significant risk of a new "digital divide" in healthcare. Well-resourced urban medical centers may be able to leverage AI to leap further ahead in efficiency and quality, while smaller, rural, and safety-net hospitals lack the capital and technical expertise to adopt these tools, widening the gap in care between communities.

Change Management and "Pilot Purgatory"

Technology is only part of the solution. Many organizations find themselves stuck in "pilot purgatory," where they conduct numerous small-scale AI experiments that show promise but fail to scale into meaningful, enterprise-wide change. Overcoming this requires a deep investment in change management, including redesigning workflows, upskilling the workforce, and proactively addressing employee fears about job displacement.

The Risk of Misapplication

There is a clear danger that agentic AI will be misapplied in ways that harm patients and erode trust. One of the most cited examples is the use of AI for "automated rejection" of insurance claims. An AI trained on historical data may simply learn to deny claims more efficiently, without the contextual reasoning to understand medical necessity or individual patient circumstances.

These challenges create a fundamental dilemma for the healthcare industry. The very nature of agentic AI—its ability to learn and adapt rapidly—is what makes it so powerful, but it is also what makes it so difficult to regulate. As one expert in the field noted, "the technology evolves so fast... [but] at a certain point of time we have to lock something and say this is now what we do in validate". This conflict will likely force a bifurcation in the market. Low-risk administrative and operational use cases, which are not typically classified as medical devices, will see rapid innovation and adoption of the most advanced agentic models. In contrast, high-risk clinical applications, such as diagnostic systems and autonomous treatment planning, will be constrained.

Market Prognosis and Future Trajectory

The confluence of healthcare's profound needs and agentic AI's unique capabilities is creating one of the most dynamic and rapidly growing markets in the technology sector. However, its trajectory will be shaped by the tension between immense opportunity and significant hype, requiring a sober, data-driven assessment to distinguish true value from superficial automation.

Market Sizing and Segmentation: A 45% CAGR Opportunity

Multiple market intelligence reports converge on a forecast of explosive growth. The global healthcare agentic AI market is projected to expand from approximately $0.7 billion in 2025 to roughly $4.46 billion by 2030, which represents a formidable compound annual growth rate (CAGR) of 44.83%. Other analyses project an even larger long-term market, potentially reaching $21.1 billion by 2034. This growth is propelled by powerful drivers, including persistent healthcare workforce shortages, the explosion of real-time health data from wearables and other sources, and advances in the accuracy of the underlying LLMs.

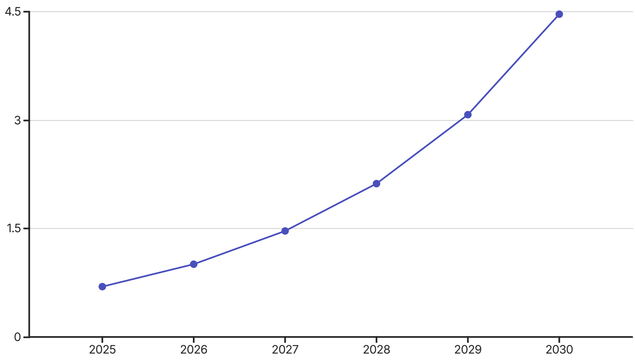

Analysis of the market's key segments reveals a nuanced landscape of current dominance and future growth, with software agent platforms holding the largest share (approximately 81%) while integration and customization services show the fastest growth (approximately 37% CAGR). Medical imaging and clinical decision support currently represent about 35% of applications, but operational and administrative automation is growing most rapidly at 39% CAGR. Hospitals and health systems are the primary current customers (49% share), though payers and insurance providers are adopting the technology at the fastest rate (35% CAGR).

From a geographical perspective, North America dominates with 55-59% market share, though the Asia-Pacific region shows the most rapid growth at 42% CAGR. Technologically, reinforcement-learning agents currently hold the largest share (46%), but large language model (LLM) agents are growing fastest (47% CAGR) due to their rapid advancement and increasing accessibility.

Competitive Landscape: Titans and Innovators

Tech Titans

Established technology giants leveraging their immense scale include Microsoft, Google, NVIDIA, Amazon Web Services (AWS), and IBM. These companies are attacking the market from a position of strength, leveraging their vast cloud computing infrastructure, foundational AI models, and existing enterprise relationships.

- Microsoft is making a major push with its Azure Health Data Services and the integration of Nuance's speech recognition technology into products like Dragon Copilot
- NVIDIA is positioning its Clara platform as a core infrastructure for healthcare AI, forming strategic partnerships with leading institutions like the Mayo Clinic
- Google (via DeepMind and Google Health) and Amazon (via AWS HealthLake) are focused on providing foundational data and AI platforms

Innovative Startups

A vibrant ecosystem of startups is emerging alongside the tech giants. These companies are often more agile, focusing on solving specific, high-value problems with purpose-built agents.

- Hippocratic AI has gained significant attention and funding by focusing on developing empathetic, non-diagnostic AI agents for tasks like patient outreach and education
- Companies like Notable, Thoughtful AI, and VoiceCare AI are targeting specific workflows like revenue cycle management and back-office automation
- The market is characterized by strategic partnerships, major acquisitions, and the rapid launch of specialized agent platforms and app stores

Hype vs. Reality: A Sober Assessment

The immense potential of agentic AI has inevitably fueled a significant hype cycle. After a period of "crashing disappointment" with some early, over-promised generative AI projects in 2024, the narrative around agentic AI is being framed as a "course correction" toward more focused, pragmatic, and value-driven applications. However, caution is warranted.

A critical distinction must be made between true, reasoning-based agentic AI and the proliferation of what some analysts call "AI wrappers". Many existing workflow automation tools are being rebranded with an "AI" label, but these systems often lack the core agentic capabilities of dynamic reasoning, planning, and adaptation. They are executing pre-scripted automation, not demonstrating genuine intelligence. The true, transformative value of agentic AI lies in its "Reasoning Layer," which enables it to make context-aware decisions in novel situations.

This distinction aligns with the balanced perspective long advocated by the National Academy of Medicine. In its landmark 2022 publication, Artificial Intelligence in Health Care: The Hope, the Hype, the Promise, the Peril, the NAM calls for a sober approach that proceeds with caution to avoid "user disillusionment" or another "AI winter" that could follow if the technology fails to deliver on its grand promises.

The market data, when viewed through this critical lens, suggests that the agentic AI landscape is not monolithic. It is actively bifurcating into two distinct segments. The first is a high-value, strategically critical segment focused on building complex "Reasoners"—foundational models and platforms with deep capabilities for planning, adaptation, and multi-system orchestration. The second is a larger, more commoditized segment of "Wrappers"—applications that use existing AI models to automate relatively simple, albeit valuable, tasks.

While these "Wrappers" may drive significant near-term market volume, especially in administrative automation, the "Reasoners" will likely capture the most defensible and profitable long-term value. Strategic decision-makers—whether investors, buyers, or policymakers—must develop the sophistication to differentiate between these two segments, as their long-term value propositions, competitive dynamics, and ultimate impact on the healthcare system will be vastly different.

Strategic Imperatives for a Cured System: Recommendations for Stakeholders

The successful application of agentic AI as a cure for healthcare's systemic maladies is not a technological inevitability. It is a strategic choice that requires a deliberate, collaborative, and profoundly responsible approach from every actor in the ecosystem. Realizing the promise while mitigating the peril demands a shared commitment to a set of guiding principles.

The NAM AI Code of Conduct: A North Star for Responsible Innovation

The NAM's AI Code of Conduct project was initiated to create a unifying framework to align the field around the responsible development and application of AI. It is designed to be a touchstone for organizations to build trust, protect patients, and ensure that innovation truly benefits people. The framework is anchored by Six Core Commitments, which provide a practical and powerful lens through which to evaluate strategy and action.

Advance Humanity

Ensure AI serves human interests, promoting well-being and respecting dignity.

Innovate and Learn

Foster a culture of continuous learning and responsible innovation, adapting to new evidence and technological advancements.

Monitor Performance

Implement robust, continuous monitoring to ensure AI systems are safe, effective, and reliable throughout their lifecycle.

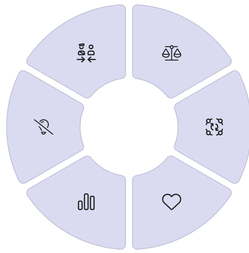

Ensure Equity

Proactively identify and mitigate biases to ensure AI benefits all populations fairly and does not exacerbate disparities.

Engage Impacted Individuals

Meaningfully involve patients, families, and communities in the design, deployment, and governance of AI systems.

Improve Workforce Well-Being

Design and deploy AI to augment human capabilities, reduce administrative burden, and alleviate clinician burnout.

Recommendations for Healthcare Providers & Systems

55%

Time Savings

Organizations implementing agentic AI for administrative workflows report time savings of up to 55%, allowing staff to focus on higher-value activities.

94%

Physician Concern

Of physicians believe prior authorizations negatively impact patient outcomes, making this a high-priority area for agentic AI intervention.

$360B

Annual Savings

Potential annual savings from AI-driven administrative efficiency improvements across the U.S. healthcare system.

1. **Strategically Target High-Impact Workflows (Commitment: Innovate and Learn)**: The temptation to pursue dozens of AI pilots simultaneously often leads to "pilot purgatory," where resources are spread thin and few initiatives achieve scale. Healthcare organizations should instead act as "focused transformers". The most prudent path is to start with low-risk, high-impact administrative and operational workflows, such as revenue cycle management, claims processing, and appointment scheduling. These applications face fewer regulatory hurdles and can generate tangible ROI and time savings, which builds organizational momentum and frees up resources to tackle more complex, higher-risk clinical challenges in a phased approach.
2. **Invest in a Modern Data Foundation (Commitment: Monitor Performance)**: Agentic AI cannot function effectively on a foundation of fragmented, low-quality data. Acknowledging this reality is the first step. Provider systems must prioritize strategic investment in their data infrastructure as a non-negotiable prerequisite for any serious AI initiative. This includes adopting modern interoperability standards like HL7 FHIR to break down data silos, implementing robust data governance policies to ensure data quality and consistency, and building the internal analytics capabilities to manage this new ecosystem.
3. **Embrace Workforce Transformation (Commitment: Improve Workforce Well-Being)**: The introduction of agentic AI is as much a human resources challenge as it is a technological one. Organizations must proactively manage the human side of this transformation. This requires investing in upskilling and retraining programs to equip staff with the skills to work alongside AI. Critically, leadership must reframe the narrative around AI—not as a tool for job replacement, but as a "digital teammate" or "co-pilot" designed to augment human expertise, automate tedious tasks, and ultimately reduce the burnout that is crippling the workforce.

Recommendations for Payers & Insurers

1. **Shift from "Automated Rejection" to "Intelligent Advocacy" (Commitment: Advance Humanity)**: Payers have a profound opportunity to change the adversarial dynamic that characterizes many provider interactions. Instead of using AI to simply deny claims faster, they should leverage the technology's reasoning capabilities to build context-aware systems for utilization and claims management. An intelligent agent can analyze a patient's full clinical picture to understand medical necessity, reducing friction and ensuring appropriate care is approved efficiently. This shifts the goal from crude cost-cutting to intelligent, system-wide waste reduction.
2. **Collaborate on Data Exchange (Commitment: Innovate and Learn)**: The administrative burden of processes like prior authorization is a shared problem that requires a shared solution. Payers should be active and willing participants in national initiatives aimed at streamlining data exchange, such as the Trusted Exchange Framework and Common Agreement (TEFCA) outlined in the 21st Century Cures Act and the emerging CMS-Aligned Networks. By enabling secure, standardized data sharing, payers can help automate these processes, dramatically reducing the burden on providers and the system as a whole.

Recommendations for Technology Developers

1. **Adopt "Ethics and Safety by Design" (Commitment: Ensure Equity)**: Responsibility for ethical AI begins with its creators. Developers must embed the principles of the NAM Code of Conduct into the AI development lifecycle from its inception. This means prioritizing transparency and explainability in model design, so that end-users can understand and trust the outputs. It requires conducting rigorous, proactive testing for algorithmic bias across diverse datasets and ensuring that development teams themselves are diverse and inclusive to avoid embedding blind spots into the technology.
2. **Build for Interoperability (Commitment: Innovate and Learn)**: To be truly effective, agentic AI must be able to seamlessly interact with the complex tapestry of existing healthcare IT systems. Developers should design their systems to be "API-native" and fully compliant with established healthcare data standards like HL7 FHIR. This will ensure that their solutions can be integrated into clinical workflows without creating new data silos, a problem that plagued the first generation of EHRs.
3. **Solve for the "Last Mile" of Implementation (Commitment: Improve Workforce Well-Being)**: A brilliant algorithm is useless if it cannot be successfully adopted and scaled in a real-world clinical environment. Developers must focus on solving this "last mile" problem. This includes designing systems with the human-in-the-loop as a core feature, not an afterthought, and providing robust change management tools and support to help healthcare organizations overcome the operational and cultural barriers to implementation.

Recommendations for Policymakers & Regulators

Policymakers and regulatory bodies play a critical role in creating the framework for safe, equitable, and effective agentic AI adoption in healthcare. Their actions can either accelerate innovation while ensuring safety or inadvertently create barriers that stifle progress. The following recommendations are designed to promote a balanced approach that addresses the unique characteristics of agentic AI.

Advance Agile and Harmonized Regulation

The pace of AI innovation demands a correspondingly agile regulatory approach. Regulators like the FDA should continue to refine and operationalize risk-based frameworks like the Predetermined Change Control Plan (PCCP) to allow for the safe evolution of adaptive algorithms. Furthermore, U.S. agencies should work with international counterparts to harmonize standards for AI validation and oversight, reducing friction for responsible global innovation.

Incentivize Data Interoperability and Quality

Policy is a powerful lever for breaking down the data silos that impede AI progress. Policymakers should strengthen and expand legislation like the 21st Century Cures Act, which mandates data access and prohibits information blocking. They should continue to support and fund federal initiatives like the CMS Interoperability Framework and the development of Trusted Exchange Networks that provide the digital rails for a truly connected health system.

Fund Equity-Focused Research and Implementation

To prevent agentic AI from becoming a tool that widens the gap between the healthcare 'haves' and 'have-nots,' policy must be intentionally focused on equity. Federal research funding should be directed toward developing and validating methods for detecting and mitigating algorithmic bias. Furthermore, grant programs and reimbursement incentives should be created to support the adoption and implementation of proven AI tools in underserved, rural, and low-resource healthcare settings, ensuring that the benefits of this transformative technology are shared by all.

Coordinated Multi-Stakeholder Governance Model

The complexity of agentic AI in healthcare requires a coordinated governance approach that brings together all key stakeholders. This table operationalizes a risk governance framework, assigning responsibilities and recommending concrete governance controls for each major stakeholder in the ecosystem, guided by the principles of the NAM AI Code of Conduct.

Stakeholder	Primary Responsibility Area	Key Risk to Mitigate	Recommended Action/Governance Control
Healthcare Providers & Systems	Clinical Implementation & Operations	Diagnostic Error from AI Over-reliance	Implement "human-in-the-loop" workflows with clear checkpoints for clinical validation. Conduct regular audits of AI performance against human expert decisions.
Healthcare Providers & Systems	Clinical Implementation & Operations	Widening Equity Gaps	Prioritize AI deployment in areas that address known health disparities. Ensure training data is representative of the local patient population.
Payers & Insurers	Claims Adjudication & Utilization Management	Unfair Denial of Care ("Automated Rejection")	Design agentic systems for "intelligent advocacy" with a focus on medical necessity and context, not just cost-cutting. Ensure transparency in denial reasons.
Payers & Insurers	Claims Adjudication & Utilization Management	Data Privacy Breaches	Adhere to robust data security and de-identification protocols. Implement strict access controls and conduct regular cybersecurity audits.
Technology Developers	AI Model Development & Design	Algorithmic Bias	Adopt "Ethics by Design." Conduct rigorous bias testing on diverse datasets. Ensure development teams are diverse and inclusive.
Technology Developers	AI Model Development & Design	"Black Box" Opacity & Lack of Trust	Prioritize explainability and interpretability in model design. Provide clear documentation on how models arrive at decisions.
Policymakers & Regulators (e.g., FDA, CMS)	Market Oversight & Safety Assurance	Unsafe Adaptive Algorithms	Continue to refine and implement agile, risk-based frameworks like the Predetermined Change Control Plan (PCCP). Strengthen post-market surveillance requirements.
Policymakers & Regulators (e.g., FDA, CMS)	Market Oversight & Safety Assurance	Stifling Innovation	Balance safety with progress by creating clear, harmonized guidelines. Foster public-private partnerships to accelerate responsible innovation.

Implementation Timeline and Milestones

The path to effectively integrating agentic AI into healthcare requires a phased approach that acknowledges the varying levels of risk, regulatory scrutiny, and technical complexity associated with different applications. Below is a strategic implementation roadmap that balances the urgent need for system improvement with the imperative for responsible, safe deployment.

Phase 1: Administrative Foundation (Years 1-2)

Focus on low-risk, high-ROI administrative applications with minimal regulatory barriers:

- Autonomous revenue cycle management and claims processing
- Automated prior authorization and eligibility verification
- Operational efficiency improvements (scheduling, documentation)

Simultaneously establish data governance frameworks and interoperability infrastructure.

Phase 2: Clinical Decision Support (Years 2-3)

Expand to medium-risk applications that augment clinical decision-making:

- AI-assisted diagnostic systems with human oversight
- Intelligent clinical documentation assistants
- Remote patient monitoring and chronic disease management

Begin workforce transformation initiatives and establish robust evaluation protocols.

Phase 3: Advanced Clinical Applications (Years 3-5)

Carefully introduce higher-risk applications with extensive validation:

- Personalized treatment planning systems
- Autonomous diagnostic systems in select domains
- End-to-end care journey orchestration

Implement comprehensive governance structures and continuous monitoring systems.

Phase 4: AI-Native Health System (Years 5+)

Transform the organization into a fully integrated, AI-native health system:

- Multi-agent systems managing complex clinical workflows
- Continuous learning systems for quality improvement
- New operating models built around human-AI collaboration

Focus on eliminating health disparities and ensuring equitable access to AI benefits.

This phased approach allows organizations to build capabilities, establish trust, and learn from early implementations before tackling more complex and higher-risk applications. It also provides time for the regulatory framework and ethical guidelines to mature alongside the technology, ensuring that safety and equity remain paramount throughout the transformation journey.

Case Studies: Early Evidence of the Agentic Cure

While the full potential of agentic AI in healthcare is still emerging, several early implementations provide compelling evidence of its transformative power. These case studies illustrate how organizations are beginning to apply agentic systems to address specific pain points across administrative, clinical, pharmaceutical, and patient care domains.

Administrative Transformation: MidwestHealth's Revenue Cycle Revolution

MidwestHealth, a 12-hospital integrated delivery network, faced mounting financial pressure from claims denials and administrative inefficiency. Their revenue cycle management (RCM) team was spending over 4,000 staff hours monthly on manual prior authorization processes, with denial rates exceeding 24% for complex procedures.

In 2024, they implemented an agentic AI system to transform their RCM workflows. The agent was designed to autonomously monitor the EHR for scheduled procedures requiring authorization, extract relevant clinical documentation, verify eligibility, submit requests through payer portals, and track approval status—all with minimal human intervention.

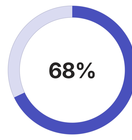

Reduction in manual hours spent on prior authorizations

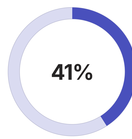

Decrease in initial claims denials

Annual financial impact from improved collections

The results were dramatic. Within six months, the organization reduced manual hours spent on prior authorizations by 68%, decreased initial claims denials by 41%, and realized an annual financial impact of $4.2 million from improved collections. Equally important, the RCM staff reported significantly higher job satisfaction as they were redeployed to more complex cases requiring human judgment and patient advocacy. The hospital's CFO noted, "This wasn't just about cost savings. It fundamentally changed how our revenue cycle team operates, allowing them to focus on high-value work that makes a real difference for patients navigating coverage challenges."

Clinical Excellence: PacificCare's Diagnostic Assistant

PacificCare Medical Center, an academic medical institution on the West Coast, implemented a multi-modal diagnostic assistant to support its radiology department in 2023. The system was designed to work alongside radiologists, autonomously analyzing medical images, correlating findings with patient EHR data, and suggesting potential diagnoses with supporting evidence.

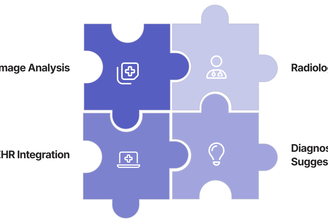

The agentic system was particularly impactful in three key areas:

- **Early Cancer Detection:** The system demonstrated a 22% improvement in identifying early-stage lung nodules on chest X-rays compared to radiologist review alone, potentially enabling earlier intervention for dozens of patients.
- **Workflow Optimization:** Radiologists reported a 33% reduction in time spent on routine cases, allowing them to dedicate more attention to complex studies that required deeper expertise.
- **Incidental Finding Management:** The agent autonomously tracked and followed up on incidental findings that might otherwise have been lost in the system, ensuring 94% adherence to recommended follow-up protocols compared to a previous rate of 71%.

The department chair emphasized that the system was designed to augment rather than replace radiologist expertise: "Our radiologists remain the decision-makers, but they now have an intelligent assistant that helps them work more efficiently and catch subtle findings that might be missed during a busy shift. It's transformed how we practice."

Pharmaceutical Innovation: AcceleraTrial's Clinical Trial Revolution

AcceleraTrial, a biotechnology company developing novel treatments for autoimmune disorders, faced the common industry challenges of lengthy trial recruitment periods and high dropout rates. In 2024, they deployed an agentic AI system to transform their clinical trial operations.

The agent was tasked with orchestrating the end-to-end trial process. It autonomously:

- Scanned millions of de-identified EHR records to identify potential candidates matching complex eligibility criteria
- Generated site-specific recruitment forecasts to optimize site selection
- Monitored real-time trial data to detect potential safety signals and protocol adherence issues
- Facilitated patient engagement through personalized reminders and educational content
- Automated data validation and regulatory documentation preparation

76%

Faster Recruitment

Recruitment timeline reduced from 18 months to just 4.5 months for a Phase II trial

24%

Lower Dropout

Patient retention improved with personalized engagement and support

$4.3M

Cost Savings

Direct cost reduction from accelerated timeline and operational efficiency

The impact was transformative. AcceleraTrial completed recruitment for their Phase II trial in 4.5 months instead of the projected 18 months, reduced patient dropout rates by 24%, and decreased direct trial costs by $4.3 million. The company's Chief Medical Officer stated, "This approach fundamentally changes the economics and timeline of drug development. We're now able to bring promising therapies to patients years earlier than would have been possible with traditional trial methods."

Patient Empowerment: ChronicCare's Remote Monitoring Agent

ChronicCare Health, a value-based care organization specializing in complex chronic disease management, implemented an agentic AI system to extend care beyond office visits for their diabetes and heart failure patients.

The multi-agent system included:

Monitoring Agent

Continuously analyzed data from patient wearables, connected glucometers, and home blood pressure monitors, detecting concerning patterns before they became clinical emergencies.

Coaching Agent

Provided personalized lifestyle recommendations, medication reminders, and motivational support through a conversational interface, adapting its approach based on patient preferences and responses.

Clinical Liaison Agent

Alerted care teams to potential issues, prepared pre-visit summaries of patient data, and facilitated smoother transitions between care settings by ensuring information continuity.

One year after implementation, ChronicCare observed a 32% reduction in diabetes-related hospitalizations and a 28% decrease in heart failure exacerbations requiring emergency care. Patient engagement metrics showed that 76% of enrolled patients actively interacted with the system at least three times weekly, and 89% reported feeling "more supported" in managing their condition.

The organization's Chief Population Health Officer noted, "This system doesn't just collect data—it turns that data into meaningful interventions at exactly the right moment. We're shifting from episodic, reactive care to truly continuous health management, which is what patients with chronic conditions desperately need."

These case studies, while still early examples, demonstrate the concrete value that agentic AI can deliver across different healthcare domains. They highlight how the technology's unique capabilities—autonomy, goal-orientation, contextual reasoning, and system orchestration—enable it to address challenges that previous technologies could not effectively solve. As these implementations mature and scale, they provide a glimpse of how a fully "agentic-enabled" healthcare system might function: more efficient, more proactive, more personalized, and ultimately more humane for both patients and providers.

Conclusion: The Path to a Healed Healthcare System

The United States healthcare system stands at a critical inflection point. Its well-documented ailments—crushing administrative burden, clinical inefficiencies, innovation bottlenecks, and workforce burnout—have proven resistant to conventional treatments. This analysis has examined how agentic AI, with its unique capabilities for autonomous, goal-oriented action and complex workflow orchestration, offers a promising new therapeutic approach to these systemic challenges.

Synthesizing the Evidence

The evidence presented throughout this report supports several key conclusions about the potential of agentic AI to transform healthcare:

Unprecedented Economic Impact

The potential economic benefit of agentic AI in healthcare is substantial, with estimates suggesting annual savings between $200 billion and $360 billion from administrative automation alone. This represents a rare opportunity to simultaneously reduce costs while improving quality and access.

Multifaceted Clinical Benefits

Beyond cost savings, agentic AI can enhance clinical outcomes through improved diagnostic accuracy, personalized treatment planning, and proactive intervention for at-risk patients. Early implementations show meaningful improvements in key quality metrics, including earlier disease detection and reduced medical errors.

Workforce Transformation

Perhaps most significantly, agentic AI offers a path to address the crisis of healthcare workforce burnout by automating routine tasks, reducing administrative burden, and enabling clinicians to practice at the top of their license. The technology can serve as a "digital teammate" that complements and amplifies human capabilities rather than replacing them.

Significant Implementation Challenges

Despite its promise, the path to widespread adoption is fraught with obstacles. These include technical barriers related to data quality and interoperability, regulatory uncertainties, equity concerns, and the organizational change management required to integrate these systems into clinical workflows.

The Imperative for Responsible Innovation

As healthcare organizations, technology developers, payers, and policymakers navigate this emerging landscape, they must be guided by a commitment to responsible innovation. The National Academy of Medicine's AI Code of Conduct provides a valuable framework, emphasizing that AI must advance humanity, ensure equity, engage impacted individuals, improve workforce well-being, be rigorously monitored, and enable continuous learning.

Realizing the full potential of agentic AI in healthcare will require a coordinated, multi-stakeholder approach that balances innovation with safety, equity, and human-centered design. The following strategic priorities emerge as essential for all stakeholders:

1. **Focus on High-Value Use Cases**: Begin with applications that address clear pain points and can generate demonstrable value, particularly in administrative and operational domains where regulatory barriers are lower and ROI can be realized more quickly.
2. **Invest in Data Infrastructure**: Recognize that high-quality, interoperable data is the foundation for effective agentic AI. Organizations must prioritize investments in modern data architecture and governance to enable these systems to function optimally.
3. **Develop Agile Regulatory Frameworks**: Regulatory bodies must continue to evolve oversight approaches that can keep pace with rapidly advancing technology while ensuring patient safety and algorithmic fairness.
4. **Prioritize Equity and Inclusion**: All stakeholders must work proactively to ensure that agentic AI reduces rather than exacerbates healthcare disparities, with particular attention to data representativeness, algorithm validation across diverse populations, and equitable access to these technologies.
5. **Embrace Human-AI Collaboration**: Design systems and workflows that leverage the complementary strengths of humans and AI, allowing each to focus on what they do best in a collaborative partnership rather than viewing AI as a replacement for human judgment.

The Future of Healthcare: An AI-Native Vision

Looking ahead, the most transformative potential of agentic AI lies not in simply automating existing processes but in enabling a fundamental redesign of how healthcare is delivered. The truly "AI-native" healthcare organization of the future will be characterized by seamless integration of human and artificial intelligence, continuous learning systems that adapt to new evidence in real-time, and a shift from episodic, reactive care to proactive, personalized health management.

In this vision, agentic AI does not merely treat the symptoms of healthcare's dysfunction; it addresses the root causes by breaking down silos, enabling true continuity of care, reducing friction between stakeholders, and allowing healthcare professionals to focus on the uniquely human elements of care that cannot be replicated by technology.

The path to this future will not be linear or without challenges. It will require thoughtful navigation of technical, ethical, regulatory, and organizational complexities. However, the potential reward—a healthcare system that is more efficient, effective, equitable, and humane—makes this journey not merely worthwhile but essential.

The agentic cure for healthcare's ailments is still in its early stages, but the preliminary evidence is promising. With careful stewardship, collaborative governance, and a commitment to responsible innovation guided by the NAM's core principles, agentic AI offers a once-in-a-generation opportunity to heal a fractured healthcare system and create a healthier future for all.