Review

Medicare and Medicaid Healthcare Access and Affordability Using Agentic Generative AI and AGI: Policy Implications and Guidelines

Satyadhar Joshi *

Posted Date: 16 September 2025

doi: 10.20944/preprints202509.1269.v1

Keywords: Agentic AI; Generative AI; Artificial General Intelligence (AGI); medicare; medicaid; healthcare access; healthcare affordability; health policy; algorithmic bias; administrative automation

 doi:10.20944/preprints202509.1269.v1

Review

Medicare and Medicaid Healthcare Access and Affordability Using Agentic Generative AI and AGI: Policy Implications and Guidelines

Satyadhar Joshi [1,2,3]

[1] Independent Researcher; satyadhar.joshi@gmail.com
[2] Alumnus, International MBA, Bar-Ilan University, Israel
[3] Alumnus, Touro College MSIT, NY, USA

Abstract

The United States' public health insurance programs, Medicare and Medicaid, face persistent challenges in ensuring equitable access to care and containing escalating costs, with an estimated quarter of all healthcare spending lost to administrative inefficiency. The emergence of advanced artificial intelligence, particularly Agentic AI and Generative AI (GenAI), presents a paradigm-shifting opportunity to address these systemic issues. Agentic AI, characterized by its ability to perform complex, multi-step tasks autonomously with human oversight, moves beyond mere automation to enable intelligent reasoning and workflow management. This paper explores the transformative potential of these technologies in streamlining administrative burdens, enhancing clinical decision support, personalizing member engagement, and optimizing resource allocation. We analyze specific use cases—including automated eligibility verification, prior authorization, and claims processing—and quantify the substantial economic impact and return on investment achievable through their implementation. Furthermore, we critically examine the significant policy implications, including algorithmic bias, regulatory compliance, data privacy, and the necessity for human-in-the-loop oversight. The paper concludes with a set of proposed guidelines and concrete policy recommendations for lawmakers and regulators to ensure the ethical, equitable, and effective deployment of Agentic and Generative AI, ultimately aiming to expand access and improve affordability for millions of beneficiaries.

Keywords: Agentic AI; Generative AI; Artificial General Intelligence (AGI); medicare; medicaid; healthcare access; healthcare affordability; health policy; algorithmic bias; administrative automation

1. Introduction

The Medicare and Medicaid programs form the cornerstone of public healthcare in the United States, providing coverage for over 150 million elderly, low-income, and disabled Americans [1]. Despite their critical role, these programs are besieged by twin challenges: ensuring timely and equitable access to quality care and controlling the unsustainable growth of healthcare expenditures. Administrative complexity alone accounts for a significant portion of wasteful spending, estimated at nearly one-quarter of total healthcare costs [2]. This bureaucratic burden impedes patient access, contributes to provider burnout, and diverts resources from direct patient care.

The advent of sophisticated Artificial Intelligence (AI) offers a beacon of hope for systemic reform. While earlier AI and robotic process automation (RPA) tools focused on discrete, repetitive tasks, the latest evolution—Agentic AI—represents a fundamental leap. Agentic AI systems are defined by their capacity to understand high-level goals, break them down into sub-tasks, utilize tools (including GenAI models), execute a sequence of actions, and adapt based on outcomes, all with a degree of autonomous reasoning [3,4]. When combined with the content-generation capabilities of Generative AI, these systems can tackle complex, knowledge-intensive workflows that were previously the exclusive domain of human experts.

 doi:10.20944/preprints202509.1269.v1

This paper investigates the application of Agentic and Generative AI to ameliorate access and affordability issues within Medicare and Medicaid. We posit that these technologies can serve as a powerful "escape hatch" from the current administrative quagmire [2]. The structure of this paper is as follows: Section 2 provides necessary background on Agentic AI and the specific challenges of Medicare/Medicaid. Section 3 details use cases for improving healthcare access. Section 4 explores applications for enhancing affordability. Section 5 delves into the critical policy and regulatory implications. Section 8 proposes concrete guidelines for implementation, and Section 11 offers concluding remarks.

2. Background and Definitions

2.1. The Medicare and Medicaid Challenge

Medicare and Medicaid are large, complex systems with distinct populations and operational challenges. Medicare, serving seniors and the disabled, grapples with chronic disease management and rising costs for advanced treatments. Medicaid, a state-federal partnership for low-income individuals, faces challenges related to member churn, eligibility complexity, and disparities in care access [5]. Both programs suffer from cumbersome processes for eligibility determination, prior authorization, claims adjudication, and member communication, creating friction for patients and providers alike [6].

2.2. Evolution from Generative AI to Agentic AI

Generative AI (GenAI) refers to models capable of creating new content—text, images, code—based on learned patterns from training data. In healthcare, GenAI has shown promise in drafting clinical notes, summarizing patient records, and powering educational chatbots [7,8].

Agentic AI builds upon this foundation. It is not a single model but an architectural framework where an AI ägentörchestrates multiple steps. As depicted in Figure 1, an agent can plan a workflow (e.g., p̈rocess a prior authorization request), use tools like a GenAI model to draft a justification, query a database for clinical guidelines, execute an API call to check eligibility, and then present a final recommendation to a human for review [9,10]. This capacity for autonomous, multi-step reasoning is what distinguishes agentic systems from traditional automation and simple chatbots [11].

2.3. The Aspiration of AGI

While current Agentic AI is powerful, it operates within predefined domains and goals. Artificial General Intelligence (AGI) remains a theoretical frontier, representing AI with human-like cognitive abilities to understand, learn, and apply knowledge across a broad range of tasks. Although true AGI is not yet realized, the trajectory of Agentic AI points towards systems with increasingly generalized capabilities, making the consideration of its potential impact prudent for long-term policy planning [12].

2.4. Structural Complexities as a Barrier to Efficiency

The operational inefficiencies of Medicare and Medicaid are not accidental but are deeply embedded in their program structures. Medicare is a primarily federal program with a well-defined population (individuals 65 and older, and those with certain disabilities), yet it is fragmented into Parts A, B, C (Medicare Advantage), and D, each with its own rules, premiums, and cost-sharing structures [1]. This complexity creates confusion for beneficiaries and administrative overhead for providers navigating different plans and coverage rules.

Medicaid, in contrast, is a state-federal partnership where each of the 50 states and territories operates its own program within federal guidelines. This results in 56 different sets of eligibility rules, benefits packages, and reimbursement rates [5]. This heterogeneity makes scaling solutions incredibly difficult and creates significant administrative burden for multi-state providers and insurers. Furthermore, Medicaid populations experience high "churn"—frequent transitions on and off coverage due to fluctuating incomes and complex renewal processes. This churn disrupts care continuity, leads

 doi:10.20944/preprints202509.1269.v1

to worse health outcomes, and generates immense administrative costs for states and managed care plans [13].

2.5. The Magnitude of the Administrative Burden

The financial scale of the administrative problem is staggering. Of the nearly $4.5 trillion spent on U.S. healthcare in 2022, an estimated $1 trillion was consumed by administrative costs, representing the highest administrative overhead of any major healthcare system globally [2]. A significant portion of this waste occurs within public programs. For providers, this manifests as staff dedicated solely to navigating prior authorization requirements, claims submission, and denial management. For payers, it appears in the form of large teams processing manual claims and eligibility verifications. This represents a massive opportunity cost, diverting resources away from direct patient care and toward bureaucratic functions.

2.6. Data Silos and Interoperability Challenges

The effectiveness of any AI system is predicated on access to comprehensive, high-quality data. The U.S. healthcare system, however, is characterized by profound data fragmentation. Patient information is locked in silos across hundreds of different Electronic Health Record (EHR) systems, payer databases, pharmacy systems, and state Medicaid IT systems. This lack of interoperability prevents a holistic view of a patient's health history and creates inefficiencies, as the same data must be manually re-entered or faxed between systems. Agentic AI's potential to automate complex workflows is therefore contingent on its ability to navigate these silos through APIs and emerging interoperability standards like FHIR, making the push for open data architectures a prerequisite for its success [14].

3. Enhancing Healthcare Access

A primary barrier to care is administrative complexity. Agentic AI can streamline these processes, reducing delays and simplifying navigation for vulnerable populations.

3.1. Streamlining Eligibility and Enrollment

Medicaid eligibility is notoriously complex, varying by state and involving frequent redeterminations that lead to coverage gaps [13]. Agentic AI can automate this process. AI agents can guide applicants through forms, parse documents to verify income and residency, interface with state eligibility systems, and proactively alert members to upcoming redeterminations, drastically reducing administrative burden and preventing lapses in coverage [5,15]. Similar agents can help seniors navigate the complex options during Medicare Annual Enrollment Periods, providing personalized plan comparisons [16,17].

3.2. Intelligent Prior Authorization

The prior authorization process is a major source of clinician burnout and care delays. Agentic AI systems can transform this workflow. An agent can receive a request, automatically retrieve the patient's clinical records from an EHR, cross-reference the indicated treatment against the latest payer-specific clinical guidelines using a GenAI module, and either auto-approve straightforward cases or compile all necessary evidence for a human reviewer to make a rapid decision [18,19]. This reduces denial rates and speeds up access to necessary care [20].

3.3. 24/7 Multilingual Member Services

Agentic AI-powered chatbots and virtual assistants can provide scalable, empathetic support. Unlike simple rule-based chatbots, these agents can handle complex, multi-turn conversations, understand context, and perform actions like scheduling appointments, explaining benefits, or submitting service requests [15,21]. This provides equitable access to information for members with limited English proficiency or those in different time zones, reducing barriers to care [22].

 doi:10.20944/preprints202509.1269.v1

4. Improving Healthcare Affordability

Beyond improving access, Agentic AI can directly and indirectly reduce costs for programs and beneficiaries.

4.1. Automating Administrative Workflows

A significant portion of healthcare costs is administrative. Agentic AI can automate high-volume, repetitive tasks across the revenue cycle, including claims processing, coding, and denials management. An AI agent can review a claim for errors, suggest correct codes, submit it to the payer, track its status, and even manage the appeals process if denied [23,24]. This reduces operational costs for payers and providers, savings that can be passed on to the system [25].

4.2. Proactive Care Management and Prevention

For Medicare Advantage and Medicaid Managed Care plans, Agentic AI can drive value-based care. Agents can continuously analyze population health data to identify high-risk patients, predict adverse health events, and proactively orchestrate interventions—such as automatically messaging a patient with CHF about weight management or scheduling a diabetic member for a retinal exam [26,27]. This shift from reactive to proactive care improves health outcomes and reduces expensive emergency department visits and hospitalizations [28].

4.3. Optimizing Resource Allocation

At a systemic level, Agentic AI can help public health officials and plan administrators optimize resource allocation. By analyzing trends in claims, social determinants of health (SDOH) data, and public health records, AI systems can predict disease outbreaks, identify regions with unmet needs, and suggest targeted interventions [29]. This ensures that resources are directed to where they can have the greatest impact on population health and cost containment.

5. Policy and Regulatory Implications

The deployment of AI in public healthcare is not merely a technical challenge but a profound policy imperative. Without careful governance, AI risks exacerbating the very disparities it seeks to solve.

5.1. Algorithmic Bias and Equity

A paramount concern is the potential for AI systems to perpetuate or amplify existing biases present in historical training data. An algorithm trained on data where certain demographic groups received less care might wrongly conclude that less care is appropriate, leading to discriminatory outcomes [30]. For Medicare and Medicaid, which serve vulnerable populations, this is unacceptable. CMS has explicitly stated that AI must not impede equitable access [30,31]. Policies must mandate rigorous bias auditing, transparency in algorithm design, and ongoing monitoring for discriminatory patterns [32].

5.2. Regulatory Compliance and Oversight

Medicare Advantage plans are already under scrutiny for using algorithmic tools in coverage decisions [6]. Regulatory frameworks must evolve to ensure that Agentic AI systems comply with all existing rules (e.g., Medicare Conditions of Participation, Medicaid Managed Care Regulations). This includes clear accountability: a human must ultimately be responsible for clinical and coverage decisions. Policies must enforce a ẗhuman-in-the-loopm̈odel for high-stakes decisions, ensuring that AI is an assistive tool, not an autonomous decider [31,33].

5.3. Data Privacy and Security

Agentic AI systems require access to vast amounts of Protected Health Information (PHI). This raises significant concerns under HIPAA and other privacy laws. Policies must mandate stringent data

 doi:10.20944/preprints202509.1269.v1

governance protocols, including data encryption, access controls, and audit trails. Furthermore, the use of PHI for training AI models must be governed by clear consent mechanisms and de-identification standards to protect beneficiary privacy [34,35].

5.4. Liability and Accountability

When an AI agent makes an error that leads to patient harm—such as a wrongfully denied prior authorization—who is liable? The provider using the tool? The health plan that deployed it? The developer that created it? Existing legal frameworks are ill-equipped to handle these questions. New policy guidelines must establish clear chains of accountability and liability for adverse events involving AI to protect patients and ensure there is recourse for harm [36].

6. Economic Impact: Quantifying Cost Savings and ROI

The adoption of Agentic AI in Medicare and Medicaid is not merely a technological upgrade but a strategic financial imperative. The potential for significant cost savings and a strong return on investment (ROI) is a primary driver behind its implementation. This section analyzes the economic rationale, drawing on estimates of current waste and projections of future savings.

6.1. Addressing the Trillion-Dollar Administrative Burden

The U.S. healthcare system is plagued by enormous administrative costs, a significant portion of which is classified as waste. This waste represents a tangible "past loss" that Agentic AI is uniquely positioned to recapture.

- **Magnitude of the Problem:** As cited in the literature, nearly one-quarter of all healthcare spending—amounting to hundreds of billions of dollars annually—is wasted on inefficient administrative processes, redundant paperwork, and manual data entry [2]. For public programs like Medicare and Medicaid, which operate on taxpayer funds, this inefficiency directly translates into reduced services for beneficiaries and higher costs for the government.
- **Target for Automation:** Agentic AI directly targets this waste by automating complex, knowledge-intensive workflows. Use cases such as automated prior authorization, claims processing, and eligibility verification are not incremental improvements but fundamental re-engineerings of processes that have historically consumed vast resources [18,23]. For instance, automating just the prior authorization process could save plans and providers billions of dollars annually in administrative overhead [18].

6.2. Projected Future Savings and Efficiency Gains

Beyond recapturing losses from waste, Agentic AI generates new value and future savings through enhanced efficiency, improved outcomes, and optimized resource allocation.

6.2.1. Operational Efficiency for Payers and Providers

The implementation of AI agents leads to direct operational cost reductions:

- **Reduced Labor Costs:** Automating tasks such as generating prior authorization letters, populating EHR fields, and answering routine member inquiries allows existing staff to focus on higher-value, complex tasks. This increases throughput without a proportional increase in labor costs [20,24].
- **Faster Reimbursement:** AI-driven claims processing that is accurate and compliant reduces denial rates and accelerates reimbursement cycles, improving cash flow for providers and reducing the administrative cost of managing denials and appeals for payers [26].

6.2.2. Clinical and Care Management Savings

The most significant long-term savings are generated by improving patient health, thereby reducing the need for expensive acute care.

 doi:10.20944/preprints202509.1269.v1

- **Proactive Care Management:** Agentic AI systems can analyze population data to identify high-risk patients and proactively orchestrate interventions. Preventing a single hospitalization for a patient with congestive heart failure or poorly managed diabetes can save tens of thousands of dollars. Scaling this across a plan's membership represents a massive financial saving [25,27].
- **Improved Medication Adherence:** AI agents can personalize patient communication, provide reminders, and identify barriers to adherence. Improved adherence to medication regimens leads to better-controlled chronic conditions and fewer complications, directly reducing overall healthcare costs [28].

6.3. Return on Investment (ROI) Considerations

Investing in Agentic AI technology requires upfront capital, but the literature suggests a compelling ROI.

- **For Health Plans:** The ROI is realized through reduced administrative overhead, lower rates of expensive hospitalizations among members, improved Star Ratings (which are tied to financial bonuses in Medicare Advantage), and increased member satisfaction and retention [21,26].
- **For Providers:** ROI is realized through reduced administrative staffing needs, faster reimbursement, and fewer denied claims. Furthermore, reducing physician burnout by automating bureaucratic tasks can save costs associated with clinician turnover [20].
- **For Government Programs:** On a macro scale, the ROI for Medicare and Medicaid is measured in the slowed rate of cost growth, better health outcomes for the population, and a more sustainable financial model for these critical public programs [1].

Conclusion of Economic Analysis

The economic case for Agentic AI is robust. It offers a direct path to recapturing hundreds of billions of dollars lost to administrative waste annually while simultaneously generating future savings through improved health outcomes and operational efficiencies. For CFOs and policymakers alike, the question is not whether the healthcare system can afford to invest in this technology, but whether it can afford not to, given the immense financial pressure on public and private payers. The implementation of Agentic AI represents a strategic investment in a more financially sustainable future for Medicare and Medicaid.

7. Workforce Training and Adaptation for the AI Era

The successful integration of Agentic and Generative AI into healthcare workflows is not solely a technological challenge; it is a human one. The effective deployment of these systems is contingent upon a workforce that is adequately trained to collaborate with, manage, and oversee AI tools. The literature consistently highlights that without strategic investment in education and change management, the full potential of AI will remain unrealized, and new risks may emerge.

7.1. The Imperative for Upskilling and Reskilling

The advent of Agentic AI will inevitably transform existing roles rather than simply replace them. The focus shifts from performing manual, repetitive tasks to overseeing automated processes and interpreting AI-generated outputs.

- **Shifting Roles:** As Agentic AI automates administrative burdens such as prior authorization, claims processing, and data entry, administrative staff must be upskilled to manage these AI systems, handle complex exceptions that the AI flags, and intervene in edge cases [20,24]. Their role evolves from data *processors* to AI *supervisors*.
- **Clinical Oversight:** For clinicians, the core of their work remains patient care. However, they require training to critically evaluate AI-generated recommendations (e.g., for treatment plans or prior authorization justifications), understand the limitations of the models, and maintain

 doi:10.20944/preprints202509.1269.v1

ultimate accountability for clinical decisions [33,37]. This "human-in-the-loop" model is only effective if the human is empowered and knowledgeable.

7.2. Addressing the Technical Knowledge Gap

A significant barrier to adoption is the knowledge gap between AI developers and healthcare professionals. Bridging this gap requires targeted educational initiatives.

- **AI Literacy:** There is a growing need for basic AI literacy across the healthcare workforce. This includes understanding fundamental concepts like how models are trained, the importance of data quality, and the potential for algorithmic bias [29,32]. This knowledge is essential for building trust and ensuring ethical use.
- **Specialized Training Programs:** The literature calls for the development of specialized training programs to equip healthcare leaders, IT staff, and clinicians with the skills to evaluate, implement, and govern AI solutions effectively [14]. This includes training on new regulatory requirements and compliance standards specific to AI in healthcare.

7.3. Implementing Effective Change Management

Introducing AI into complex healthcare environments is a significant organizational change that must be managed carefully to avoid resistance and ensure smooth adoption.

- **Workflow Integration:** Training must extend beyond abstract concepts to focus on the practical integration of AI tools into daily clinical and administrative workflows. Users need to understand how the AI agent will function as a part of their team and how it will make their jobs easier rather than more complicated [4,10].
- **Promoting a Collaborative Culture:** Success depends on fostering a culture of collaboration between humans and AI. This involves moving away from a fear of replacement and towards a mindset of augmentation, where AI handles computational tasks, freeing humans to focus on empathy, complex problem-solving, and patient interaction [9].

Conclusion and Recommendation

The consensus in the literature is clear: technological investment must be matched by investment in human capital. For policymakers and healthcare organizations, this means:

1. **Allocating Funding:** Congress and state governments should appropriate funding for grants aimed at developing AI training curricula for healthcare professionals and for upskilling existing staff within public programs like Medicare and Medicaid.
2. **Developing Standards:** Professional medical and administrative associations should collaborate to develop standardized competencies and certification programs for working with AI in healthcare settings.
3. **Prioritizing Change Management:** Healthcare organizations must dedicate resources to structured change management programs that include comprehensive training, continuous support, and clear communication about the role of AI.

Ultimately, the most sophisticated Agentic AI system will fail without a trained workforce that can use it effectively, ethically, and safely. Training is not a peripheral concern but a central pillar of successful AI implementation in healthcare.

8. Proposed Guidelines for Implementation

To harness the benefits of Agentic AI while mitigating its risks, we propose the following guidelines for policymakers, payers, and providers:

1. **Principle of Equity First:** Any deployment of AI must be preceded by a health equity impact assessment. Systems must be continuously monitored for disparate impacts on protected classes, and developers must employ techniques like debiasing and fairness constraints [29,30].

 doi:10.20944/preprints202509.1269.v1

2. **Transparency and Explainability:** AI systems must not be "black boxes." For any decision, there must be an auditable trail and an explanation comprehensible to a human reviewer. This is essential for trust, debugging, and regulatory compliance [31].
3. **Human-in-the-Loop Mandate:** Establish clear guardrails that require meaningful human oversight for all clinical decisions and coverage determinations. AI should augment human judgment, not replace it [33,37].
4. **Interoperability and Open Standards:** Promote the development of AI agents based on open standards (e.g., FHIR) to ensure they can work across different EHR and payer systems, preventing vendor lock-in and fostering innovation [14].
5. **Robust Evaluation and Validation:** Before deployment, AI systems must undergo rigorous third-party testing for safety, efficacy, and bias. There should be a defined process for post-market surveillance to catch issues that emerge in real-world use [38,39].
6. **Workforce Training and Transition:** Invest in training programs for the healthcare workforce to work effectively alongside AI agents. This includes upskilling administrative staff to manage and oversee AI systems and training clinicians to interpret AI-generated recommendations critically.

9. The Agentic AI Ecosystem: Key Technologies and Vendors

The theoretical potential of Agentic AI is being rapidly operationalized by a growing ecosystem of technology firms, startups, and specialized software platforms. An analysis of the current landscape reveals several key players and technologies specifically targeting the healthcare sector, including Medicare and Medicaid administration.

9.1. Major Technology Platforms and Frameworks

The development of Agentic AI applications is often built upon foundational frameworks and platforms that provide the necessary infrastructure for building, deploying, and managing AI agents.

- **Microsoft:** Through its **Microsoft Cloud for Healthcare**, Microsoft is actively promoting Agentic AI solutions. Their technical blog posts detail architectures where AI agents orchestrate workflows across healthcare data silos, emphasizing security and compliance capabilities crucial for handling Protected Health Information (PHI) [**?**].
- **Salesforce:** A major player in customer relationship management (CRM), **Salesforce** is applying its **Einstein AI** platform to healthcare. They specifically market Ḧealthcare Agentic AIf̈or use cases in patient engagement, clinical decision support, and streamlining workflows for both providers and payers [40,41].
- **Writer:** This company offers an AI platform that has published extensive guidance specifically for healthcare payers. Their B̈ig Book of AI Agents for Healthcare Payorsöutlines use cases for member engagement, process streamlining, and customer support, positioning themselves as a specialist in the payer space [21].

9.2. Specialized Healthcare AI Vendors

Beyond large tech platforms, a niche of vendors focuses exclusively on creating AI solutions for healthcare's unique challenges, including those of public payers.

- **HealthcareGPS:** This vendor has developed a suite of AI-powered tools specifically for the Medicare market. Their products are designed to assist agents with plan selection and navigation, directly targeting the challenge of improving access and understanding for seniors [16].
- **Inovaare:** Focused on the payer side, **Inovaare** markets AI agents designed to automate complex compliance tasks, audits, and operational workflows for health plans. This directly addresses administrative affordability and accuracy in programs like Medicare Advantage and Medicaid Managed Care [23].
- **Authenticx:** This company specializes in analyzing customer interaction data (e.g., call center transcripts). They have published analyses on the use of AI specifically with Medicare and

 doi:10.20944/preprints202509.1269.v1

Medicaid data, highlighting both positive and negative developments, thus operating in the data analytics and insight generation layer [34].

- **Productive Edge** and **Accelirate:** These are examples of digital consulting and implementation firms that have established dedicated ÄI Agents for Healthcareïresource centers and practices. They help payer and provider organizations design and implement custom Agentic AI solutions, indicating a mature consulting ecosystem around this technology [42,43].

9.3. Enabling Technologies and Architectural Concepts

The functionality of Agentic AI is enabled by specific technical paradigms and architectural patterns frequently mentioned in the literature.

- **Model Context Protocol (MCP):** This emerging technical standard, referenced in discussions of agentic architecture, provides a framework for AI agents to seamlessly connect to and use various software tools, data sources, and APIs. This interoperability is critical for building agents that can work across disparate EHR systems and payer databases [44].
- **Big Data Processing:** The effectiveness of any Agentic AI system is predicated on its ability to process and reason over massive datasets. This includes structured claims data, unstructured clinical notes, and real-time interaction streams. The agents developed by the aforementioned vendors all rely on underlying big data pipelines to ingest, clean, and process this information before the AI can act upon it [5,34].
- **Robotic Process Automation (RPA) Integration:** Many described Agentic AI workflows represent an evolution beyond traditional RPA. While RPA automates repetitive keyboard clicks, Agentic AI adds a layer of cognitive reasoning to manage and trigger those automations based on complex input, creating a more intelligent and flexible automation system [5].

Implications for Policy and Implementation

The presence of this vibrant and specialized ecosystem demonstrates that Agentic AI is not a distant future concept but a present-day reality with significant commercial investment. For policymakers, this means that:

1. **Regulatory targets are clear:** Guidance from CMS must be applicable to both large tech platforms (e.g., Microsoft, Salesforce) and specialized healthcare AI vendors (e.g., HealthcareGPS, Inovaare).
2. **Interoperability is non-negotiable:** Policies should encourage or mandate the use of open standards (like MCP or FHIR) to prevent vendor lock-in and ensure that AI tools can work across the heterogeneous IT landscape of U.S. healthcare.
3. **The market is evolving rapidly:** Any policy framework must be adaptable to accommodate new entrants and technological advancements without requiring a complete regulatory overhaul.

10. Policy Recommendations for Government Action

The transformative potential of Agentic AI in Medicare and Medicaid is contingent upon a robust policy framework that fosters innovation while rigorously protecting beneficiary rights and promoting equity. Based on the analysis of opportunities and risks, we propose the following concrete recommendations for legislative and regulatory action.

10.1. Establish a Pre-Deployment AI Equity Audit Mandate

To proactively prevent algorithmic discrimination, policymakers must mandate a rigorous pre-deployment testing and auditing regime. The Centers for Medicare & Medicaid Services (CMS) should issue a rule requiring that any Agentic AI tool used for coverage determinations, risk stratification, or prior authorization by Medicare Advantage (MA) plans or Medicaid Managed Care Organizations (MCOs) must undergo and pass a standardized Health Equity Impact Assessment (HEIA) [29,30].

This HEIA, conducted by an independent, certified third-party, must:

 doi:10.20944/preprints202509.1269.v1

- Evaluate the model's performance accuracy and prediction rates across subgroups defined by race, ethnicity, language, sex, disability status, and socioeconomic factors.
- Use synthetic or carefully anonymized real-world data to simulate edge cases and potential discriminatory outcomes.
- Be a prerequisite for CMS approval of the tool's use within public programs.

This mandate would operationalize the principle that "AI must be used in a manner that preserves equitable access" [30].

10.2. Formalize the "Human-in-the-Loop" Regulatory Standard

The autonomous nature of Agentic AI necessitates a clear, non-negotiable requirement for human oversight. We recommend that CMS explicitly amend its guidance to MA plans and MCOs to define and enforce a "meaningful human review" standard for all adverse coverage decisions [31,33].

This regulation must specify that:

- **Final Determination:** Any denial of coverage, prior authorization, or service reduction recommended by an AI system must be reviewed and made final by a qualified clinical professional employed by the plan.
- **Overruling AI:** The human reviewer must have the authority, information, and incentive to overrule the AI's recommendation. The rate of AI recommendation overrules should be a key metric in CMS audits.
- **Audit Trail:** A complete audit trail, documenting the AI's output and the human reviewer's final decision and rationale, must be maintained and made available for regulatory review and appeals processes.

This creates a clear chain of accountability and prevents the ceding of ultimate clinical decision-making to algorithms [6].

10.3. Create Transparency and Explainability Requirements

To build trust and ensure fairness, beneficiaries and providers have a right to understand AI-driven decisions. Congress should consider legislation akin to a "Right to Explanation" for AI-assisted determinations in public health insurance programs.

Regulatory guidance from CMS should require that:

- Any communication to a beneficiary regarding a denial of service based on an AI tool must include a clear, plain-language, and specific reason for the denial, citing the actual clinical criteria used—not a generic reference to an "internal protocol" or "algorithm" [31].
- Upon request, beneficiaries and their providers must be granted access to a more detailed explanation of the factors that contributed to the AI's output, ensuring the explanation is clinically relevant and useful for formulating an appeal.

This transparency is critical for upholding due process and allows for the identification and correction of flawed logic or biased data patterns [32].

10.4. Develop a Certified AI Vendor Program

To manage the complex landscape of AI vendors and ensure a baseline of quality and safety, CMS should establish a voluntary certification program for AI tools used in Medicare and Medicaid. This program, potentially modeled on the FDA's pre-certification concept for digital health, would:

- Set baseline technical standards for fairness, security, privacy, and interoperability (e.g., FHIR-based data access) [14].
- Create a public list of "CMS-Certified" AI tools, giving payers and providers confidence in their procurement choices and simplifying their own compliance efforts.
- Streamline the review process for certified tools, creating a incentive for developers to adhere to the highest standards of safety and equity from the outset.

 doi:10.20944/preprints202509.1269.v1

10.5. Allocate Funding for AI Literacy and Workforce Training

The successful integration of Agentic AI requires a workforce capable of working alongside it. We recommend that Congress appropriate funding through CMS and the Office of the National Coordinator for Health IT (ONC) for grant programs focused on:

- **Workforce Upskilling:** Training programs for administrative and clinical staff in healthcare settings, MA plans, and Medicaid agencies to effectively manage, oversee, and interpret the outputs of Agentic AI systems.
- **Beneficiary Education:** Developing and disseminating educational materials to help Medicare and Medicaid beneficiaries understand how AI might be used in their care, their rights regarding automated decisions, and how to appeal them.

This proactive investment ensures that the human elements of the healthcare system remain empowered and central [37].

Conclusion of Recommendations

The implementation of these policy recommendations will not stifle innovation but rather channel it in a direction that is equitable, accountable, and trustworthy. By establishing clear guardrails, the government can unlock the immense benefits of Agentic AI—reduced administrative waste, improved access, and better health outcomes—while safeguarding the foundational principles of the Medicare and Medicaid programs. The time for proactive governance is now, before these powerful technologies become further entrenched without necessary protections.

11. Conclusions and Future Work

This paper has articulated a comprehensive framework for leveraging Agentic and Generative AI to tackle the most persistent challenges facing Medicare and Medicaid: equitable access and systemic affordability. Our analysis demonstrates that these are not merely incremental technologies but represent a fundamental shift from automation to intelligent, autonomous reasoning. By deploying AI agents to manage complex workflows—from streamlining eligibility verification and prior authorization to enabling proactive care management—we can directly address the staggering administrative waste that currently consumes a quarter of healthcare spending [2].

The economic case is compelling. The implementation of Agentic AI offers a clear pathway to recapture hundreds of billions of dollars lost to inefficient processes, while simultaneously generating value through improved health outcomes and optimized resource allocation. This creates a powerful return on investment for payers, providers, and the public programs themselves [25]. However, as our examination of the ecosystem reveals, this potential is being rapidly operationalized by a diverse set of vendors, from tech giants to specialized startups, making the need for coherent policy more urgent than ever [23**?**].

Realizing this future is contingent upon overcoming significant hurdles. The risks of algorithmic bias, data privacy breaches, and unaccountable decision-making are profound, especially for the vulnerable populations served by Medicare and Medicaid. Therefore, our proposed guidelines and policy recommendations—centered on mandatory equity audits, a formalized human-in-the-loop mandate, transparency requirements, and workforce investment—are not ancillary concerns but foundational prerequisites for ethical and effective implementation [30,31]. The successful integration of this technology is as much a human capital challenge as a technical one, necessitating a national commitment to upskilling the workforce to collaborate with and oversee AI systems [20].

In conclusion, Agentic AI presents an opportunity to build a more efficient, accessible, and equitable healthcare system. By embracing a strategy of proactive governance, rigorous oversight, and strategic investment, policymakers, payers, and providers can ensure that these powerful technologies fulfill their promise as a force for good, safeguarding the integrity of Medicare and Medicaid for future generations. The time to establish the guardrails for this transformation is now.

 doi:10.20944/preprints202509.1269.v1

Conflicts of Interest: The views are of the author and do not represent any affiliated institutions. Work is done as a part of independent research. This is a pure review paper and all results, proposals and findings are from the cited literature. Author does not claim any novel findings.

References

1. Can Medicare and Medicaid Promote More Efficient Health Care?
2. The Administrative Reckoning in Health Care: Why Agentic AI May Be the System's Only Escape Hatch. https://www.medicaleconomics.com/view/the-administrative-reckoning-in-health-care-why-agentic-ai-may-be-the-system-s-only-escape-hatch, 2025.
3. Bresnick, J. What Is Agentic AI and What Does It Mean for Healthcare?, 2025.
4. is a freelance writer who specializes in the healthcare industry, E.L.T.E.L. What Is Agentic AI and How Can It Be Used in Healthcare? https://healthtechmagazine.net/article/2025/05/what-is-agentic-ai-in-healthcare-perfcon.
5. Bag, U. Bridging the Healthcare Gap: AI's Role in Expanding Access to Medicaid & CHIP.
6. mshepard. When Artificial Intelligence in Medicare Advantage Impedes Access to Care: A Case Study, 2022.
7. Generative AI in Healthcare: Emerging Use for Care | McKinsey. https://www.mckinsey.com/industries/healthcare/our-insights/tackling-healthcares-biggest-burdens-with-generative-ai.
8. Landi, H. Hype or Hope? Patients Optimistic about Generative AI in Healthcare. https://www.fiercehealthcare.com/ai-and-machine-learning/us-patients-believe-generative-ai-can-improve-access-affordability, 2023.
9. Padmanabhan, G. Agentic AI and the Architecture of Healthcare Transformation, 2025.
10. From Generative AI to Agentic AI | HIMSS. https://www.himss.org/resources/from-generative-ai-to-agentic-ai/.
11. Uppili, S. Agentic AI Vs Generative AI: The Ultimate Comparison Guide, 2025.
12. How Generative AI and Agentic AI Will Revolutionize Healthcare in 2030. https://www.kellton.com/kellton-tech-blog/generative-agentic-ai-healthcare-2030.
13. Beavins, E. CHAI Convenes Team to Explore AI for Medicaid Work Requirements. https://www.fiercehealthcare.com/ai-and-machine-learning/ai-could-be-used-new-medicaid-work-requirements-chai-convening-tiger-team, 2025.
14. Seizing Healthcare's Agentic AI Opportunity. https://www.cognizant.com/us/en/insights/insights-blog/automation-and-ai-agents-for-healthcare.
15. AI Agents in Healthcare: Scaling Member Service for Faster, More Equitable Access.
16. HealthcareGPS Products: AI-Powered Medicare Solutions. https://healthcaregps.ai/products.
17. Utilizing AI in Medicare Enrollment Strategies for Agents. https://www.ema.co/additional-blogs/addition-blogs/ai-medicare-enrollment-strategies-agents.
18. 4 Top Use Cases for Agentic AI in Healthcare | TechTarget. https://www.techtarget.com/healthtechanalytics/feature/4-top-use-cases-for-agentic-AI-in-healthcare.
19. Agentic AI in Healthcare: 4 Game-Changing Use Cases. https://www.edstellar.com/blog/agentic-ai-healthcare.
20. Team, M. AI for Healthcare Administration: How Agentic AI Drives Efficiency and Patient Engagement | Motics Blog. https://www.motics.ai/blog/ai-for-healthcare-administration-how-agentic-ai-drives-efficiency-and, 2025.
21. The Big Book of AI Agents for Healthcare Payors.
22. AI Facilitates Equitable Access to Public Health Services, 2023.
23. manager.inovaare. How AI Agents for Payers Are Driving Smarter Compliance and Operations | Inovaare Blog, 2025.
24. Dubey, N. Agentic AI in Healthcare: Transforming Care Delivery, 2025.
25. religroup2025. Lowering Medicare Costs and Improving Beneficiary Health with GenAI Solutions, 2025.
26. How Agentic AI Helps Improve Health Outcomes, Reimbursement, Star Ratings.
27. Plessas, D. The Agentic AI Revolution: Transforming MLR/HBR Management for Payers in Value-Based Care, 2025.
28. Sankrityayan, V.D. Can AI Make Healthcare Affordable?, 2025.
29. Driving Equitable Healthcare Innovations through AI for Medicaid.
30. CMS Proposes Rules to Ensure AI Does Not Impede Equitable Access to Services. https://www.akingump.com/en/insights/ai-law-and-regulation-tracker/cms-proposes-rules-to-ensure-ai-does-not-impede-equitable-access-to-services.

 doi:10.20944/preprints202509.1269.v1

31. Manning, Charles, K.P. CMS Clarifies Boundaries on AI in Coverage Determination for Medicare Advantage Organizations (via Passle). https://quickreads.ext.katten.com//post/102j00y/cms-clarifies-boundaries-on-ai-in-coverage-determination-for-medicare-advantage-o, 2024.
32. Lei, Carolyn Metnick, J.K.T. Recent Healthcare-Related Artificial Intelligence Developments, 2024.
33. Office, U.S.G.A. Science & Tech Spotlight: Generative AI in Health Care | U.S. GAO. https://www.gao.gov/products/gao-24-107634.
34. Authenticx. Using AI in Healthcare: Medicare and Medicaid Data, 2024.
35. Lei, Carolyn Metnick, J.K.T. Recent Healthcare-Related Artificial Intelligence Developments, 2024.
36. Group, P.L.I. Agentic AI and the Architecture of Healthcare Transformation, 2025.
37. Agentic AI: How Can It Help Physicians? https://www.medicaleconomics.com/view/agentic-ai-how-can-it-help-physicians-, 2025.
38. Signify Research | Agentic AI in Healthcare: A Game-Changer or Just Hype? https://www.signifyresearch.net/insights/generat ai-news-round-up-january-2025/.
39. Mika, S. Is Agentic AI in Healthcare Overhyped?, 2025.
40. Healthcare Agentic AI: Benefits & Use Cases. https://www.salesforce.com/healthcare-life-sciences/healthcare-artificial-intelligence/healthcare-agentic-ai/.
41. How Agentic AI in Pharma Is Revolutionizing Healthcare. https://www.salesforce.com/healthcare-life-sciences/life-sciences-artificial-intelligence/agentic-ai-in-pharma/.
42. AI Agents for Healthcare Resource Center | Productive Edge. https://www.productiveedge.com/ai-agents-for-healthcare-resource-center.
43. How Agentic AI Is Transforming Healthcare: Benefits & Use Cases - Accelirate. https://www.accelirate.com/agentic-ai-in-healthcare/.
44. Revolutionizing Healthcare: MCP and Agentic AI in Action | by Amit Parihar | Medium. https://medium.com/@amitperf/revolutionizing-healthcare-mcp-and-agentic-ai-in-action-0c1d6bc9f3b4.