THE HEDIS INSURGENCY: BUILDING AI AI AGENT ARMY TO HACK HEALTHCAR QUALITY MEASUREMENT

OCT 13, 2025 • PAID

TABLE OF CONTENTS

ABSTRACT

Disclaimer: The views and opinions expressed in this essay are my own and do n reflect the positions, strategies, or opinions of my employer or any affiliated organizations.

This dossier presents a technical and business framework for a health technolog company built around deploying specialized AI agents to systematically improve Healthcare Effectiveness Data and Information Set (HEDIS) measures across hea plan populations. The core innovation involves creating an "agent army" where e agent specializes in a specific HEDIS measure or clinical pathway, operating sem autonomously while coordinating through a central orchestration layer. The com would generate revenue through performance-based contracts tied to quality me improvement, targeting the approximately $60 billion annual flow between CMS Medicare Advantage plans based on Star Ratings. Key technical components incl multi-modal data ingestion from claims, EHR, and pharmacy systems; specialize large language models fine-tuned for clinical workflows; real-time decisioning engines; and integration with existing care management platforms. The business model leverages risk-sharing arrangements where the company captures 30-50% incremental quality bonus payments generated through measurable HEDIS improvements. Primary risks include regulatory scrutiny of AI in clinical workfl integration complexity with legacy health IT systems, and the challenge of maintaining AI model performance across diverse health plan populations.

Introduction

The American healthcare system runs on measures that nobody particularly like everyone must obsess over. Chief among these are the Healthcare Effectiveness I and Information Set measures, those seventy-some metrics that determine wheth health plans receive applause or penalties from the National Committee for Qua Assurance. If you have spent any time inside a Medicare Advantage plan during September and October, you have witnessed the annual HEDIS panic, where ent organizations mobilize to close care gaps before the measurement period ends.

Clinical teams make frantic outreach calls to patients who need colorectal cance screenings. Pharmacy teams scramble to improve medication adherence rates. C coordinators chase down hemoglobin A1c results for diabetic patients. The who enterprise feels simultaneously critical and absurd, like watching a hospital try t simultaneously treat patients while also making sure the right boxes get checkec regulators.

This frenzy exists because HEDIS measures directly translate into revenue throu Medicare Advantage Star Ratings and increasingly through Medicaid managed c quality bonuses. A half-star improvement in Star Ratings can mean hundreds of millions of dollars in additional revenue for a large Medicare Advantage plan. T stakes are extraordinary. Yet the fundamental approach to HEDIS improvement remained remarkably unchanged for two decades. Health plans employ armies o nurses to review charts and make outreach calls. They build care gap reports tha physicians largely ignore. They implement clunky clinical decision support syste that interrupt workflows without providing real value. Billions of dollars get spe this infrastructure every year, yet many health plans still struggle to move their metrics meaningfully.

What if instead of this traditional approach, we built something completely diffe What if we created an actual army of artificial intelligence agents, each one specialized and relentless, each one focused on a single HEDIS measure, each on capable of operating semi-autonomously while coordinating with hundreds of si in a shared mission to systematically dismantle care gaps across an entire health population? This is not a metaphor. This is a technical architecture and busines model for a company that could fundamentally change how health plans approac quality measurement.

The HEDIS Measurement Problem

Before diving into the solution, we need to understand the problem with sufficie granularity. HEDIS measures come in several flavors, each with distinct technica challenges. Some measures are purely administrative, calculable entirely from cl

data. Others are hybrid measures requiring both claims data and medical record review. Still others are survey-based measures dependent on patient-reported outcomes. For our purposes, we focus on the seventy-two measures that matter most for Medicare Advantage Star Ratings and Medicaid managed care quality programs.

Consider the complexity of a single measure like Diabetes Care - Blood Sugar Controlled. This measure requires identifying all plan members with diabetes, determining whether they received a hemoglobin A1c test in the measurement year, extracting the test result, and categorizing control status. Simple in theory. In practice, this requires parsing International Classification of Diseases codes from medical and pharmacy claims to identify the denominator population, integrating with laboratory information systems or medical records to retrieve test results, handling missing data when results are not available electronically, coordinating medical record retrieval from hundreds of provider practices, abstracting values from unstructured clinical notes when structured data is unavailable, and managing the logic for exclusions when patients have specific conditions that make the measure inapplicable.

Multiply this complexity by seventy-two measures. Add the reality that each measure has different numerator and exclusion criteria that change annually. Layer on the fact that health plans typically have provider networks spanning hundreds or thousands of independent practices with wildly varying levels of health IT sophistication. Factor in that many critical data elements live in unstructured clinical documentation rather than discrete fields. The result is a measurement and improvement problem of staggering technical complexity.

The current market solution involves mostly human labor. Health plans employ armies of nurses who manually review medical records, either on-site at provider practices or through electronic record retrieval. They use care management platforms from companies like Arcadia, Innovaccer, or Lightbeam to generate care gap reports. They contract with specialized HEDIS vendors like Inovalon or HealthEdge for chart retrieval and abstraction services. Some forward-thinking plans have begun experimenting with natural language processing to extract data from clinical notes, but these efforts remain nascent and narrow in scope. The entire ecosystem is built

around humans doing most of the cognitive work, with technology serving prima as data aggregation and workflow management infrastructure.

The cost structure is punishing. Large Medicare Advantage plans spend tens of millions of dollars annually on HEDIS-related activities. This includes salaries fo internal quality teams, contracts with chart retrieval vendors at thirty to sixty do per record, fees for care management platform licenses, and payments to provide practices for quality improvement initiatives. Meanwhile, physician practices fac their own HEDIS burden in the form of additional documentation requirements quality reporting workflows, and participation in endless health plan quality programs. The friction and expense in the system are enormous, yet the measure themselves are not going away. If anything, quality measurement is expanding in new domains as value-based care becomes more prevalent.

Market Opportunity and Economics

The total addressable market for HEDIS improvement is actually quite tractable estimate with reasonable precision. Start with Medicare Advantage. Approximat twenty-nine million beneficiaries are enrolled in Medicare Advantage plans as of 2024. CMS pays Medicare Advantage plans roughly five percent bonus payments achieving four or more stars, with additional quality bonus payments for five-sta plans. The average Medicare Advantage per-member-per-month payment is arou one thousand dollars, meaning approximately three hundred fifty billion dollars annually from CMS to Medicare Advantage plans. A conservative estimate sugge that approximately sixty billion dollars of this is directly or indirectly tied to qua performance through Star Ratings bonuses and quality-adjusted benchmarks.

The incremental revenue opportunity from Star Ratings improvement is substan Moving from 3.5 stars to 4.0 stars can generate an additional five percent revenu increase for a plan. For a plan with one million Medicare Advantage members, th translates to approximately six hundred million dollars in annual run-rate reven increase. The willingness to pay for technology and services that can reliably del this improvement is correspondingly high. Health plans will readily pay twenty t

thirty percent of incremental quality bonus revenue for solutions that demonstra move Star Ratings.

Medicaid managed care represents a parallel opportunity. Approximately forty-t million beneficiaries receive Medicaid benefits through managed care plans. Wh Medicaid quality bonus programs vary by state, an increasing number of states h implemented meaningful quality-based payment programs. States like Pennsylva Ohio, and Michigan have quality withhold programs where five to ten percent of capitation payments are at risk based on HEDIS performance. The total Medicai managed care market exceeds three hundred billion dollars annually, with approximately fifteen to twenty billion dollars likely tied to quality performance

Beyond the direct financial incentive, there is a secondary market in provider-sponsored health plans and risk-bearing provider organizations. As more physic groups take on full or partial risk through Accountable Care Organizations and contracting arrangements, they face the same HEDIS improvement challenges as traditional payers. The number of beneficiaries attributed to risk-bearing provid entities exceeds fifteen million in Medicare alone, creating an adjacent market fo quality improvement solutions.

The competitive landscape for capturing this opportunity is fragmented. Existin point solutions address narrow slices of the problem. Care management platforn provide data aggregation and workflow tools but rely on humans for interventio Natural language processing vendors extract some data from clinical notes but d drive interventions. Traditional HEDIS vendors focus on chart retrieval and abstraction but do not close care gaps. Disease management companies provide outreach but lack the technical sophistication for measure-specific optimization company has built a comprehensive, measure-specific, AI-agent-driven platform can systematically drive improvement across the full HEDIS measure set.

This represents a classic market inefficiency. The total pain is enormous, the willingness to pay is high, the existing solutions are inadequate, and the technic capabilities to build something better now exist in ways they did not five years a

The window is open for a well-capitalized entrant with strong technical capabili and deep healthcare domain expertise.

Technical Architecture and AI Agent Design

The core technical innovation is building an army of specialized AI agents where agent is purpose-built for a specific HEDIS measure or clinical workflow. This is single general-purpose AI trying to do everything. Instead, imagine seventy disti agents, each one an expert in its domain, each one continuously operating across health plan population to identify opportunities, prioritize interventions, execut outreach, coordinate with providers, and measure impact.

The architecture has several layers. At the foundation is a data ingestion and harmonization layer that connects to all relevant data sources within a health pla This includes claims data warehouses, electronic health record systems via FHIR HL7 interfaces, pharmacy systems, laboratory information systems, health information exchanges, and care management platforms. The data layer must ha both structured and unstructured data, must resolve patient identity across disp systems, must maintain data freshness with near-real-time updates for critical workflows, and must ensure appropriate privacy and security controls throughou This is table-stakes health IT infrastructure, difficult but well-understood.

Above the data layer sits the agent orchestration platform. This is where things g interesting. Each agent is implemented as a specialized large language model fin tuned on clinical documentation, medical coding, and workflows specific to its t measure. For example, the Colorectal Cancer Screening agent would be trained o gastroenterology notes, pathology reports, and colonoscopy documentation. It w understand the nuances of when a colonoscopy is truly diagnostic versus screeni how to interpret CT colonography results, and when fecal immunochemical test appropriate alternatives. The agent would maintain a dynamic model of every eli member in its population, understanding current gap status, prior outreach atte barriers to completion, and optimal intervention strategies.

The agents operate through a continuous sense-and-respond loop. Every day, eac agent reviews its population, identifies new opportunities based on fresh data, prioritizes interventions based on likelihood of success and impact on overall m performance, selects appropriate intervention modalities, executes outreach thr various channels, monitors responses, adjusts strategies based on outcomes, and coordinates with other agents when interventions overlap. This is not batch processing. This is continuous, real-time operation with agents making thousan micro-decisions daily.

The model architecture for each agent combines several components. A base lar language model, likely something in the GPT-4 or Claude-3 class, provides gene language understanding and reasoning capabilities. A fine-tuning layer specializ model for clinical documentation and workflows specific to the target measure. retrieval-augmented generation system allows the agent to query relevant clinica guidelines, measure specifications, and patient-specific data to ground its decisi policy layer encodes business rules, compliance constraints, and escalation crite reinforcement learning component allows the agent to improve over time based outcomes.

The intervention toolkit available to agents is multimodal. Agents can generate natural language outreach via SMS, email, or interactive voice response systems. can create faxes to provider offices requesting specific documentation or actions can populate task lists within provider EHR systems where integration exists. Tl can schedule care gap closure visits by interfacing with practice management sy They can identify transportation or language barriers and coordinate appropriat support services. They can escalate complex cases to human care managers with detailed briefings on the clinical situation and recommended approaches.

A critical technical challenge is coordination between agents. A single patient m have care gaps across fifteen different HEDIS measures. We cannot have fifteen different agents simultaneously bombarding that patient with outreach. The orchestration layer manages these conflicts through a priority ranking system ba on clinical urgency, likelihood of success, and impact on overall Star Ratings. Ag negotiate resources through an internal market mechanism where higher-priorit

interventions get preferential access to outreach channels and care management support. This prevents both patient overwhelm and resource waste while ensuri that the most valuable opportunities receive attention.

The system requires robust monitoring and evaluation infrastructure. Each agen tracks detailed performance metrics including population size, current perform rate, care gap closure rate, cost per intervention, and projected end-of-year performance. The platform provides real-time dashboards showing aggregate performance across all measures, identifying underperforming agents or measur highlighting emerging opportunities, and projecting final Star Ratings based on current trajectories. This creates a flywheel where underperforming measures automatically receive additional resources and attention.

Integration points with existing health plan systems are essential for practical deployment. The platform must integrate with existing care management system avoid creating a parallel workflow universe. It must connect to provider portals EHR systems to enable seamless information exchange. It must interface with m engagement platforms to coordinate outreach. It must feed into quality reportin systems to ensure measure performance is properly captured. The business reali that health plans will not rip out their existing infrastructure, so the new system wrap around and enhance what already exists.

Agent Army Operations and Coordinati

The operational model for the agent army is where the metaphor becomes concr Each agent operates with a high degree of autonomy within defined parameters, all agents coordinate through a central command structure that manages resour resolves conflicts, and ensures overall mission success. Think of it as a military command structure applied to healthcare quality measurement.

Individual agents maintain situational awareness of their entire eligible populat For a health plan with one million members, the Diabetes Care agent might be responsible for eighty thousand members with diabetes. The agent maintains a detailed profile for each member including current gap status, clinical complexi

social determinants of health, prior engagement history, and predicted likelihoo gap closure. Every morning, the agent reviews overnight data feeds to identify ne closed gaps, new eligible members, and changes in clinical status. Based on this updated intelligence, the agent reprioritizes its intervention queue.

The prioritization algorithm balances multiple objectives. Clinical urgency com first—a patient with uncontrolled diabetes who has not seen a provider in eighte months receives higher priority than a patient with well-controlled diabetes wh needs an annual eye exam. Likelihood of success factors in—patients who have responded to prior outreach and have scheduled appointments are higher priori than patients who have never engaged. Impact on measure performance matters- the measurement year progresses, interventions that can still impact the current rating take precedence over interventions that will only help future years. Cost-effectiveness plays a role—high-touch interventions requiring care manager involvement are reserved for situations where they are most likely to succeed.

Agents execute interventions across multiple channels simultaneously. For a pat who needs colorectal cancer screening, the agent might send an educational SMS message about screening options, generate a reminder notice to be included in t patient's next mailing, populate a best practice alert in the patient's primary car physician's EHR, send a fax to the physician's office with a pre-filled screening o and flag the patient for a care manager call if no action is taken within thirty day This multi-touch, multi-channel approach dramatically increases conversion rat compared to single-mode outreach.

The agents learn continuously from outcomes. When an intervention succeeds o the agent updates its model of what works for different patient segments. It mig learn that Spanish-speaking patients over seventy respond better to phone calls t text messages, or that patients in rural areas need transportation assistance to complete screening appointments, or that certain physician practices are highly responsive to faxed reminders while others ignore them entirely. This learning propagates across the agent network, so insights gained by one agent can inform strategies of other agents working with similar populations.

Coordination between agents prevents conflicts and optimizes resource utilizati The central orchestration layer maintains a master calendar of all planned interventions for every patient. When multiple agents want to contact the same patient, the orchestration layer sequences these contacts to avoid overwhelming patient. It might batch several reminders into a single communication, or it mig space outreach across several weeks. The orchestration layer also manages acces limited resources like care manager time, transportation assistance, and high-to interventions.

The agent army scales elegantly. Adding a new HEDIS measure means developin deploying a new specialized agent, but the orchestration infrastructure and data foundations remain unchanged. Onboarding a new health plan client means configuring the agents for that plan's specific population characteristics and pro network, but the core agent logic is reusable. This creates powerful economies o as the platform grows across multiple health plans.

Clinical Workflows and Integration Poi

The ultimate test of any health IT system is whether it integrates into actual clin workflows without creating additional burden. Physicians are already overwheln with electronic alerts, quality measures, and administrative tasks. Adding anoth system that interrupts their workflow with more requests is a recipe for failure. agent platform must be designed from the ground up to augment rather than interrupt clinical workflows.

The primary integration point is the electronic health record. For health systems practices using major EHR platforms like Epic or Cerner, the agents can write d into the EHR using FHIR APIs or proprietary integration mechanisms. This allo agents to populate care gap reminders, screening orders, and best practice alerts appear within the physician's normal workflow. The key is making these alerts contextually relevant and actionable. An alert that appears when a diabetic patie in the office for an annual physical and reminds the physician to order an A1c te helpful. An alert that fires randomly when the patient is not present is noise.

The agents use natural language generation to create clinical documentation tha physicians can easily incorporate into their notes. For example, if a patient need medication reconciliation for a diabetes medication adherence measure, the age generate a pre-filled medication list with suggested changes based on claims dat showing which medications the patient is actually filling. The physician can revi modify, and sign this documentation rather than creating it from scratch. This t the quality measure from a burden into a tool that actually helps the physician provide better care.

Provider practices without sophisticated EHR integration capabilities receive su through traditional channels that are enhanced by AI. The agents generate intell faxes that include patient-specific information, clear action requests, and pre-fil order forms. Unlike generic care gap reports that list dozens of patients and measures, these faxes are targeted, specific, and actionable. A practice might rec fax saying that Mrs. Johnson is scheduled for an appointment next Tuesday and a colorectal cancer screening, along with a pre-filled order for a colonoscopy inc the appropriate CPT code and suggested gastroenterology referrals.

The agents also integrate with scheduling systems to proactively create opportu for care gap closure. If a patient needs multiple preventive services, the agent ca coordinate with the practice's scheduling system to book an extended visit that a time to address all outstanding items. The agent can identify schedule openings suggest appointment times to patients through automated outreach. For patient have not been seen recently, the agent can trigger proactive outreach offering to schedule an appointment and providing multiple convenient time slots.

Pharmacy integration creates additional intervention opportunities. Many HEDI measures relate to medication adherence, which requires understanding not just whether prescriptions are written but whether they are being filled and taken. T agents integrate with pharmacy claims data and can trigger interventions when adherence drops. They can coordinate with pharmacies to set up automatic refill remind patients when medications are due, and alert prescribers when patients repeatedly fail to fill prescriptions. For diabetes medication adherence measures

agents can identify patients who are sub-optimally controlled and work with prescribers to intensify therapy.

Laboratory integration closes a major data gap. Many hybrid HEDIS measures re lab results that are not always available through claims data. The agents integrat with major laboratory information systems to retrieve test results directly. Wher results are not available electronically, the agents orchestrate medical record ret abstracting results from clinical documents using natural language processing. T eliminates the need for expensive manual chart review for many measures.

The platform includes a provider portal where practices can view all outstandin gaps for their patients, submit documentation to close gaps, and track their performance on quality measures. The portal provides value to practices beyond HEDIS reporting by helping them succeed in value-based payment arrangement by identifying patients who need outreach for chronic disease management. This creates a positive incentive for providers to engage with the platform rather thar viewing it as just another health plan reporting burden.

Business Model and Go-to-Market Strategy

The business model is structured around performance-based risk-sharing contra where the company captures a percentage of incremental quality bonus revenue generated through measurable HEDIS improvement. This aligns incentives perfe with health plan customers and ensures that the company only makes money wh delivers real value.

A typical contract structure works as follows. The company establishes a baselin performance level for each HEDIS measure based on the health plan's most rece reported performance. The company then commits to delivering measurable improvement above baseline, typically targeting improvements of five to fifteen percentage points on key Star Ratings measures over an eighteen-month engage The company receives payment based on the incremental quality bonus revenue

generated by this improvement, typically capturing thirty to fifty percent of the created.

For a Medicare Advantage plan with one million members, moving from 3.5 star 4.0 stars generates approximately six hundred million dollars in incremental rev If the company can demonstrate that its interventions were responsible for the improvement, a thirty percent share would yield one hundred eighty million dol revenue. Even after accounting for technology costs, data integration expenses, a operational overhead, the unit economics are extremely attractive.

The payment structure includes both upfront platform fees and success-based performance payments. The upfront platform fee covers data integration, agent deployment, and ongoing operational costs. This might be five to ten million dol for a large health plan, paid as a monthly subscription over the contract term. Tl performance payments are paid quarterly in arrears based on actual measured improvement in HEDIS rates. This structure ensures the company has sufficient capital to operate while aligning long-term incentives with customer success.

The go-to-market strategy focuses initially on Medicare Advantage plans where financial incentives for quality improvement are most direct and substantial. Th target customer profile is a plan with five hundred thousand to two million Med Advantage members, current Star Ratings between 3.0 and 4.0 stars, and demonstrated strategic commitment to quality improvement. These plans have t scale to justify the implementation investment while having sufficient room for improvement to generate meaningful financial returns.

The sales process is technical and consultative, typically requiring six to twelve months from initial engagement to signed contract. The process begins with a detailed analysis of the health plan's current HEDIS performance, identifying sp measures where the agent platform can drive improvement. The company provid detailed financial model projecting the expected revenue impact of improvemen across each measure and demonstrating the ROI of the engagement. This analysi phase builds credibility and allows both parties to align expectations before committing to a contract.

Proof of concept engagements provide a lower-risk entry point for health plans hesitant to commit to a full deployment. A typical proof of concept focuses on th five measures where the plan is underperforming and where the technical integr is relatively straightforward. The proof of concept runs for six months and inclu partial deployment of the agent platform, limited integration with health plan systems, and focused interventions on a subset of the member population. Succe the proof of concept creates a pathway to full deployment.

The customer success organization is critical to long-term retention and expansi Each health plan customer is assigned a dedicated customer success team includ strategic account director, clinical operations lead, and technical integration ma This team works closely with the health plan's quality and clinical operations te ensure smooth deployment, address technical issues, optimize agent performanc demonstrate value through regular business reviews. The goal is to become deep embedded in the customer's operations and to expand from initial deployments additional measures, member populations, and product lines.

Expansion opportunities within existing customers are substantial. Initial deployments might focus on Medicare Advantage, but successful engagements c opportunities to expand into Medicaid, commercial populations, and provider-sponsored health plans. The platform can also expand beyond HEDIS into other quality measurement frameworks like Merit-based Incentive Payment System measures for providers or condition-specific quality measures for disease manag programs.

The competitive moat deepens over time through several mechanisms. The agen continuously learn from interventions across all customer deployments, improvi their effectiveness through accumulated experience. The platform's integration health plan and provider systems creates switching costs. The performance-base business model aligns incentives in ways that traditional software-as-a-service m do not. The clinical credibility built through successful deployments creates str word-of-mouth marketing effects in the relatively small community of Medicare Advantage quality leaders.

Competitive Landscape and Differentiation

The competitive landscape includes several categories of companies addressing p of the HEDIS improvement problem. Understanding how this platform differen from existing solutions is essential for positioning and for anticipating competit responses.

Care management platforms like Arcadia, Innovaccer, and Lightbeam provide da aggregation, analytics, and workflow management tools for health plans and risk bearing providers. These platforms excel at identifying care gaps and providing dashboards and reports to care teams. However, they are fundamentally passive t that require human users to drive interventions. They do not autonomously exec outreach, do not continuously optimize intervention strategies, and do not provi measure-specific intelligence that specialized agents can deliver. The agent platf can integrate with these existing care management platforms, enhancing their va rather than replacing them entirely.

Natural language processing vendors like Apixio, Talix, and Health Fidelity focu extracting data from unstructured clinical documentation to improve risk adjust and HEDIS reporting. These solutions are valuable for identifying previously uncaptured diagnoses and for retrieving data needed for hybrid measures. Howe they do not close care gaps or drive interventions. They are data extraction tools rather than intervention platforms. The agent platform incorporates similar natu language processing capabilities as one component of a much broader interventi system.

Traditional HEDIS vendors like Inovalon, HealthEdge, and Advantmed provide c retrieval, abstraction, and coding services. These companies employ large teams nurses who manually review medical records to capture data for HEDIS reportin This is labor-intensive and expensive, typically costing thirty to sixty dollars per record reviewed. The agent platform can automate much of this work through n language processing while also driving interventions to close gaps before they re

expensive manual chart review. This creates a cost advantage while also improvi
outcomes.

Disease management and care coordination companies like Livongo, Omada, anc
Health provide coaching and engagement programs for specific conditions. Thes
programs can help improve outcomes for enrolled members but typically reach c
small percentage of the eligible population and focus on patient engagement rat
than comprehensive HEDIS measure improvement. The agent platform operates
population scale and addresses the full HEDIS measure set rather than focusing
narrow condition-specific interventions.

Point solution vendors address specific HEDIS measures or workflows. For exam
companies like Memora Health or Luma Health focus on appointment reminder
care gap outreach through SMS and voice channels. These tools are useful but na
in scope and lack the measure-specific intelligence and multi-modal interventio
capabilities of specialized agents. The agent platform incorporates similar
communication capabilities while adding layers of clinical intelligence, provider
integration, and continuous optimization.

The fundamental differentiation is the comprehensive, measure-specific, AI-driv
approach. No existing competitor has built specialized agents for each HEDIS
measure that operate autonomously, learn continuously, coordinate with each ot
and integrate across the full spectrum of clinical workflows. The vision is not to
another dashboard or reporting tool but to build an actual autonomous system t
drives measurable improvement without requiring proportional increases in hur
labor.

Secondary differentiation comes from the performance-based business model. M
competitors charge subscription fees or per-member-per-month rates regardless
whether quality measures actually improve. This creates misaligned incentives a
makes it difficult for health plans to justify the investment when outcomes are
uncertain. The risk-sharing model ensures that the company succeeds only when
customers succeed, creating a true partnership rather than a vendor relationship

The clinical credibility of the team and advisory board provides another competi advantage. The company must be founded and led by individuals with deep expe in both healthcare operations and AI technology. Having former health plan chie medical officers, quality directors, and payer executives on the team or advisory provides credibility with customers and ensures that the product is designed arc real clinical workflows rather than technology for its own sake. The founding tea should include engineers who have built production large language model syster healthcare data architects with experience integrating disparate health IT systen and clinical informaticists who understand the nuances of quality measurement.

Implementation Roadmap and Scaling Strategy

The implementation roadmap for a new customer follows a structured process designed to minimize disruption while maximizing speed to value. The process typically takes three to six months from contract signature to full production deployment, with incremental value being delivered throughout.

Phase one is data integration and platform configuration. This involves establish secure connections to all relevant health plan data sources, mapping data elemer the platform's standardized schema, configuring the agents for the specific popu characteristics and provider network, and establishing user accounts and access controls. This technical work happens largely in the background and requires cl coordination between the company's technical team and the health plan's IT and security organizations. The goal is to have complete, accurate, real-time data flo into the platform before any interventions begin.

Phase two is pilot deployment on selected measures. Rather than attempting to c all seventy agents simultaneously, the initial deployment focuses on five to ten measures where the health plan has the greatest opportunity for improvement ar where the technical integration is most straightforward. This might include mea like Colorectal Cancer Screening, Diabetes Care - Blood Sugar Controlled, and I Cancer Screening. The agents begin operating on these measures, executing outr

coordinating with providers, and tracking outcomes. This pilot phase typically r for sixty to ninety days and provides an opportunity to refine workflows, adjust intervention strategies, and demonstrate early wins.

Phase three is full deployment across the remaining measures. Based on learning from the pilot, the remaining agents are deployed in waves, typically adding ten fifteen measures every month until full deployment is complete. Each wave inclu training for health plan staff, integration with additional provider systems, and refinement of intervention strategies. By the end of phase three, all agents are operational and the full platform is running at scale.

Phase four is continuous optimization and expansion. Once the platform is fully deployed, the focus shifts to optimization. Agents continuously learn from outcc and refine their intervention strategies. The orchestration layer improves its coordination algorithms based on observed resource utilization and outcomes. N integration points with provider systems are added to increase automation and r friction. Additional services like transportation assistance or interpreter service incorporated to address barriers to care gap closure.

The scaling strategy balances breadth and depth. The initial focus is on achievin deep success with a small number of flagship customers, proving that the platfo can deliver measurable improvement in HEDIS measures and generate real finan returns. Success with these initial customers creates case studies and references accelerate sales cycles with new customers. Once the platform is proven with fiv ten customers representing several million members, the focus shifts to rapid customer acquisition.

Scaling requires building operational capabilities in parallel with technology development. The customer success organization must grow to support an expar customer base. The clinical operations team must be able to handle increasing volumes of escalations and edge cases that require human intervention. The integration team must be able to onboard new customers quickly without compromising on integration quality. The data engineering team must ensure th

platform can handle growing data volumes as the member population under management expands.

The company should target managing ten million covered lives by year three and thirty million covered lives by year five. At these volumes, the platform would be managing a material percentage of the total Medicare Advantage and Medicaid managed care populations. The revenue model scales elegantly with covered lives the performance-based payments are directly tied to the size of the population a magnitude of improvement achieved.

Risks, Challenges, and Mitigation

No business plan is complete without honest assessment of risks and challenges. Several categories of risk must be carefully managed for the company to succeed

Regulatory risk is substantial. The use of AI in clinical decision-making is an ev area with uncertain regulatory requirements. While the agents are primarily driv administrative and outreach functions rather than making clinical decisions, regulators may still scrutinize the platform for potential patient safety issues or algorithmic bias. Mitigation strategies include maintaining rigorous clinical ove of agent recommendations, implementing human-in-the-loop review for high-st interventions, ensuring transparency in how agents make decisions, and engagin proactively with regulators to shape appropriate oversight frameworks.

Integration risk is perhaps the most immediate challenge. Health IT systems are notoriously difficult to integrate, with inconsistent data standards, limited API availability, and resistance from entrenched vendors. Every new customer requir custom integration work, and some critical data sources may prove impossible t access electronically. Mitigation requires building a flexible integration architec that can accommodate diverse systems, investing in strong relationships with ma EHR vendors to facilitate integration, and having fallback processes for manual collection when automated integration is not feasible.

Clinical efficacy risk is the possibility that the agents' interventions do not actua improve HEDIS measures at the scale and speed needed to justify the business m Early pilots might show promising results that do not replicate at full scale, or interventions might be less effective in certain populations or market conditions Mitigation involves rigorous measurement and evaluation of intervention effectiveness, rapid iteration based on outcomes, building diverse clinical expert into the product development process, and maintaining realistic projections abo achievable improvement rates.

Algorithmic bias risk is critical in healthcare applications. If the agents learn pa that disadvantage certain demographic groups or inadvertently deprioritize vulnerable populations, the platform could perpetuate or worsen existing health disparities. Mitigation requires continuous monitoring for disparate impact acr demographic groups, incorporating health equity metrics into agent optimizatio objectives, maintaining diverse clinical and technical teams that can identify pot bias issues, and being transparent with customers about how the platform addre equity considerations.

Business model risk relates to the performance-based payment structure. If improvements in HEDIS measures are difficult to attribute to the platform's interventions versus other concurrent quality initiatives, customers may be reluc to pay performance bonuses. Similarly, if external factors like changes in measur specifications or member population mix confound measurement of improvemer disputes over payment may arise. Mitigation requires establishing clear measure methodologies upfront, using control groups or statistical modeling to isolate th platform's impact, maintaining strong relationships with customers built on tru transparency, and being willing to adjust payment terms when external factors materially impact outcomes.

Competitive response risk is the possibility that incumbents or well-capitalized entrants rapidly replicate the approach. Large care management platform vendo could build similar agent capabilities, or major payers could decide to build proprietary solutions in-house. Mitigation comes from moving quickly to establi strong market position, building integration depth and clinical credibility that a

difficult to replicate, continuously innovating to stay ahead of competitors, and potentially pursuing strategic partnerships or acquisitions to consolidate the ma

Talent risk is significant given the specialized expertise required. The company people who understand healthcare operations, quality measurement, AI technol and clinical workflows. This talent is scarce and highly sought after. Mitigation requires competitive compensation, equity participation for key team members, mission-driven culture that attracts people passionate about improving healthca and strong partnerships with academic institutions to access emerging talent.

Conclusion

The vision of an AI agent army systematically dismantling care gaps across the healthcare system is ambitious but achievable with current technology and favor market conditions. The fundamental insight is that HEDIS improvement is not a single monolithic problem but rather a collection of seventy-two distinct challer each with specific data requirements, clinical workflows, and intervention strate By building specialized agents for each measure rather than trying to create a ge purpose solution, the platform can achieve superhuman performance on the nar tasks that actually matter for quality measurement.

The economics are compelling for all stakeholders. Health plans gain access to technology that can deliver measurable Star Ratings improvements worth hundr millions of dollars while paying only for demonstrated results. Physicians receiv intelligent support that helps them provide better care without adding administ burden. Patients benefit from more proactive outreach and fewer missed preven services. The company captures a meaningful share of the value it creates while solving a real problem that the market has struggled to address effectively.

The technical capabilities required to build this platform exist today in ways the not even three years ago. Large language models have reached a level of sophisti where they can understand complex clinical documentation, reason about appro interventions, and generate natural language communications that feel personal rather than robotic. Integration technologies have matured to the point where

connecting to disparate health IT systems, while still challenging, is a solved prc for well-capitalized companies with strong technical teams. The infrastructure f deploying and managing fleets of AI agents at scale has been built by cloud prov and is available as a service.

The timing is favorable because the pressure on health plans to improve quality performance continues to intensify. As Medicare Advantage enrollment grows a CMS tightens Star Ratings methodology, the stakes for quality improvement onl increase. Meanwhile, the traditional approaches to HEDIS improvement are reac diminishing returns. Health plans have already hired as many nurses as their bu allow, implemented care management platforms across their organizations, and launched countless provider quality programs. Incremental improvement from t traditional approaches is becoming harder and more expensive. The market is re for a fundamentally different approach.

The company that successfully executes this vision will not just build a valuable business but will meaningfully improve healthcare delivery for millions of peopl Better HEDIS performance is not just about regulatory compliance and financia bonuses. These measures exist because they represent evidence-based preventive and chronic disease management that actually improves health outcomes. Helpir more diabetic patients achieve good blood sugar control prevents complications. Increasing cancer screening rates saves lives through early detection. Improving medication adherence reduces hospitalizations and emergency department visits agent army is ultimately in service of helping people live healthier lives, which is mission worth pursuing.

The path from concept to successful operating company is long and filled with challenges. Building the technology will require years of sustained engineering e and iteration based on real-world feedback. Achieving product-market fit will demand deep customer engagement and willingness to adapt the solution to mar needs. Scaling the business will test operational capabilities and require buildin organization that can deliver consistent quality at increasing volumes. Navigatin regulatory scrutiny and maintaining clinical credibility will require constant vig and investment in governance and oversight capabilities. None of this is easy, bu

of it is possible for a team with the right combination of technical capability, healthcare expertise, and entrepreneurial determination.

The future of healthcare quality measurement will likely involve increasing automation and AI-driven intervention. The question is not whether AI agents w play a role in HEDIS improvement but rather which companies will successfully deploy these capabilities and capture the associated value. The window of oppor is open now for a well-funded startup with strong founding team credentials to establish market leadership before incumbents wake up to the opportunity or w capitalized competitors enter the space. For investors seeking exposure to the intersection of AI and healthcare with clear paths to revenue and compelling un economics, this represents a rare opportunity to back a company addressing a m billion-dollar market pain point with technology that actually works. For entrepreneurs passionate about using AI to solve real problems in healthcare del building this company would be both intellectually rewarding and potentially transformative for the industry. The HEDIS insurgency awaits its leaders.

3 Likes • 1 Restack

← Previous

Nex

Discussion about this post

Comments Restacks

Write a comment...