# Building an AI-First Safety Net: Pair Team's Vision for Precision Benefits i Medicaid

DEC 03, 2025



DISCLAIMER: The views and opinions expressed in this essay are my own and d reflect the views of my employer, or any of its affiliates.

If you are interested in joining my generalist healthcare angel syndicate, reach o trey@onhealthcare.tech or send me a DM. Accredited investors only.

## TABLE OF CONTENTS

# ABSTRACT

Medicaid serves over 70 million Americans and represents roughly $700 billion i annual spending, yet remains largely untouched by the AI revolution reshaping commercial healthcare. In a new policy brief published in NPJ Digital Medicine, Team and researchers from Stanford, Berkeley, and Carnegie Mellon introduce "Precision Benefits," a framework for using AI-powered digital assistants to deli targeted support at the exact moment it can prevent health deterioration and soc crisis.

The core insight: predictive analytics can identify people at risk, but without rea time automated intervention, those insights sit unused in dashboards. Pair Team building the infrastructure to close this gap.

Original Research Paper by Pair Team:
https://link.springer.com/article/10.1038/s41746-025-02195-9

## Key components of the vision:

- AI-powered digital assistants operating 24/7 in multiple languages for care navigation, eligibility verification, and behavioral health support

- Real-time intervention triggered by predictive signals like missed appointment prescription refills

- Integration across fragmented systems connecting healthcare providers, comm organizations, and benefits agencies

- Evidence-based approaches with trials showing 51% reduction in depression an 31% reduction in anxiety from AI chatbots

The economic case is substantial: AI could save Medicaid $37-67 billion annuall through administrative efficiency alone, before accounting for clinical and opera

improvements. But Pair Team's work goes beyond cost savings to reimagine how technology can deliver more human, proactive, and accessible care to vulnerable populations. This represents a fundamental shift from reactive crisis manageme proactive support that catches people before they fall.

# Introduction: Reimagining Medicaid Through AI

Medicaid faces a fundamental tension. It serves over 70 million Americans, many complex medical and social needs, yet operates with chronic workforce shortage fragmented service delivery, and administrative burden that consumes 15% of ev healthcare dollar. Beneficiaries struggle with language barriers, behavioral healt access, and unmet social needs. Meanwhile, the technology reshaping how peopl navigate every other part of their lives has barely touched Medicaid.

Pair Team, working with researchers at Stanford, Berkeley, and Carnegie Mellon published a policy brief in NPJ Digital Medicine that lays out a different path forward. The paper introduces "Precision Benefits," a framework for using AI to deliver the right support to the right person at the right time, before avoidable h and social deterioration occurs. This isn't about incremental efficiency gains. It's about fundamentally reimagining how a safety net program can function when it AI-powered infrastructure at its core.

The timing matters. AI's ability to process language, reason, and respond has ma rapidly. Large language models can now conduct natural conversations across languages and contexts. Frameworks like retrieval-augmented generation (RAG) ground AI responses in verified data rather than hallucinations. The Model Con Protocol (MCP) lets AI systems connect to external databases and take real-worl actions through APIs. Together, these technologies form the foundation for digi assistants that can actually help people navigate complex systems rather than fru them with rigid decision trees.

Pair Team's work comes at a policy moment when both efficiency mandates and priorities align. H.R.1 requires states to improve Medicaid verification and redu

improper payments while fostering innovation. The economic pressure is real: w potential savings of $37-67 billion annually just from administrative automation states have both budget incentive and political cover to try new approaches. The question isn't whether AI will come to Medicaid, but whether it will be built responsibly and deployed in ways that actually serve beneficiaries rather than ju cutting costs.

## The Scale of the Opportunity

The numbers tell a stark story. Medicaid spending approaches $700 billion annu serving over 70 million Americans. Administrative overhead consumes roughly 1 every healthcare dollar. Apply AI automation to that administrative burden and looking at $37-67 billion in potential annual savings, scaling proportionally fron estimates that AI could reduce overall U.S. healthcare spending by 5-10% or up t $360 billion.

But Pair Team's brief makes clear that administrative efficiency is just the starti point. The real opportunity comes from clinical and operational improvements: reducing redundant testing, optimizing resource allocation, improving care coordination, preventing avoidable ER visits and hospitalizations. When you cat someone before they fall into crisis, the savings compound exponentially. A few hundred dollars of timely intervention can prevent thousands in emergency care

The access gaps are massive and well-documented. About 25% of Medicaid beneficiaries speak a language other than English at home. Wait times for behav health appointments can exceed 30 days in many areas, with some regions facing longer delays. Community health workers and care coordinators are overwhelme under-resourced. Transportation to appointments remains a persistent barrier. T aren't problems you solve by hiring more staff because the workforce shortages a structural, especially in long-term services and supports (LTSS) and home and community-based services (HCBS).

What makes Medicaid particularly compelling as a testing ground is that it serve populations with the highest needs and the most fragmented care. Multiple chro

conditions are the norm rather than the exception. Social determinants of health housing instability, food insecurity, and lack of transportation directly impact he outcomes. Language barriers complicate every interaction. Eligibility requireme are complex and vary state by state. If you can build AI systems that work in this environment, delivering personalized support to people navigating all these challenges simultaneously, then those systems will work anywhere.

## What Digital Assistants Actually Mean

Pair Team uses the term "AI-powered digital assistants" deliberately to distingu what they're building from basic chatbots. These are agentic systems designed to accomplish tasks on behalf of users with a degree of autonomy. They can initiate workflows, interact with multiple systems, and support care navigation based on input or internal reasoning. The difference between a rules-based chatbot and w Pair Team envisions is the difference between a phone tree and a human assistan actually understands your situation and can help you solve problems.

The technical foundation is what makes this possible now when it wasn't a few y ago. Large language models provide natural conversation and general-purpose reasoning that can handle ambiguity and context. RAG connects those models to verified knowledge sources and maintains memory across interactions so the ass remembers previous conversations. MCP provides the connective tissue to exter systems, letting AI assistants check eligibility, schedule appointments, coordina transportation, or trigger alerts to care teams by calling the appropriate APIs. Together, these technologies enable digital assistants that don't just answer ques but actually help people navigate complex systems.

Picture what this looks like in practice. A Medicaid member texts in Spanish ask about their upcoming cardiology appointment. The digital assistant confirms the appointment in the provider's EHR, checks whether transportation is arranged, that it's not, coordinates with the state's non-emergency medical transportation broker to schedule a ride, sends confirmation back to the member with pickup d and adds a reminder to their calendar. All automated, all in Spanish, no human

intervention unless something breaks. The member gets a seamless experience e
though the backend involves integration across multiple disconnected systems.

Or consider a more complex scenario that Pair Team highlights. A community h
worker at a food bank is helping someone who can't afford their diabetes medica
The worker asks the digital assistant (via voice) to check Medicaid enrollment st
discovers the person is eligible but not enrolled, walks through a streamlined
application right there, schedules a primary care appointment for next week, arr
transportation, and flags the case for pharmacist outreach about medication assi
programs. The assistant handles all the backend integration across eligibility sys
provider scheduling, transportation coordination, and care management platforr
What would normally take multiple phone calls, long hold times, and several day
back and forth happens in one conversation.

This is what Pair Team calls the "embedded staff" model. Instead of digital assis
being a member-facing portal that nobody uses, they become infrastructure that
enhances everyone who touches Medicaid. The food bank employee, the 211 cris
operator, the school nurse, the housing case manager. They all get augmented by
that knows how to navigate the complexity of healthcare and social services on b
of the people they're trying to help.

The multilingual, 24/7 availability aspect can't be overstated. An AI that can con
fluent conversations in Spanish, Mandarin, Vietnamese, Arabic, and dozens of o
languages removes barriers that interpretation services can't fully solve. It's not
translation, it's native conversation that understands cultural context and adapts
appropriately. And because it's always available, members can get help at 2am w
they remember a question or 6pm after work when they finally have time to deal
health stuff.

## Precision Benefits: The Core Innovatio

Here's where Pair Team's vision gets really interesting. The concept of Precision
Benefits represents a fundamental shift in how safety net programs can function

Healthcare has had predictive analytics for years. Every health plan has some kir risk stratification model that identifies members likely to have an ER visit, hospi admission, or chronic disease exacerbation. The problem is these predictions mo just sit in reports. Maybe a care manager gets a list of high-risk members to call, they're overwhelmed, under-resourced, and calling people who may or may not answer. By the time human outreach happens, if it happens at all, the window fo intervention has often passed. The insight was there, but nothing happened with

Precision Benefits flips this model completely. Instead of prediction leading to a manual workflow that may or may not execute, prediction triggers automated, personalized intervention at the exact moment it can make a difference. The AI doesn't wait for a care manager to review a dashboard and add a task to their qu acts immediately because it can, it's software running 24/7 with zero marginal co interaction. This is the key insight that makes Pair Team's approach different fr traditional care management.

The example in the paper is perfect because it's so common. Single mother, work hourly job, juggling caregiving responsibilities, managing diabetes and hyperten Her own mother gets sick, she misses work shifts to help, income drops, she star skipping her own medical appointments and rationing medication to make rent. Blood sugar climbs, blood pressure spikes, nobody notices the warning signs un lands in the ER. Health destabilized, job at risk, eviction looming. This cascade happens all the time in Medicaid.

But here's the thing. That missed prescription refill? That no-show appointment Those are observable signals that predictive models can flag in real time. Withou automated intervention, they're invisible until the crisis hits. With Precision Be the AI catches the missed refill immediately, reaches out via the member's prefe channel (text, voice, whatever), understands through conversation that money's t offers to connect them with emergency assistance programs, schedules a teleheal visit to adjust medications if needed, and follows up to make sure they picked up prescription. Total interaction cost is basically nothing. Potential savings from preventing an ER visit and all the downstream consequences? Easily thousands o dollars, plus the immeasurable human impact of keeping someone stable and ho

This is what Pair Team means by delivering the right resource, in the right dose, the right time, before breakdowns become emergencies. It's precision medicine l applied to the entire continuum of health and social needs. Just as clinicians adj medications based on patient response, Medicaid can use AI to titrate support, shifting from reactive crisis management to responsive, proactive care.

For prevention, these tools can flag early risks like missed screenings, rising BM trends, or social isolation patterns, then offer nudges or referrals to local progra before problems develop. For chronic disease management, AI can support daily stability: reminding someone with asthma to refill inhalers before they run out, t symptoms to identify patterns, stay indoors during high pollution days, and con with their doctor if control is worsening. The system adapts to individual needs circumstances, providing more intensive support when someone's struggling and lighter touch when they're stable.

## Understanding the Barriers to AI Adoption

Pair Team's brief doesn't sugarcoat the challenges, and it's worth understanding AI adoption in Medicaid has been slow despite obvious opportunities. The barri are real and solving them requires more than just good technology.

First, Medicaid has historically been hesitant to fund emerging innovations ever cost-saving potential is clear. States are risk-averse, budgets are perpetually tigh political incentives favor avoiding mistakes over taking chances on new approac Even when pilots show clear ROI, scaling is slow because procurement is fragme and every state wants to customize everything. The lack of clear reimbursement pathways and financial incentives discourages investment in AI tools. This is compounded by Medicaid's regulatory complexity and state-by-state variation in delivery models, from managed care to fee-for-service, requiring coordination ac multiple levels of government.

Second, fairness and safety concerns are paramount. AI models trained on comm insurance data will not generalize well to Medicaid's socioeconomically and raci

diverse populations. The demographics are different, the social determinants are different, the disease burden is different. Using training data that doesn't repres the population you're serving risks biased predictions and unequal care. Genera AI introduces additional risks like hallucination (making up information), inappropriate tone, and overgeneralization if not properly constrained.

Pair Team addresses these concerns head-on with specific mitigation strategies. grounds outputs in verified data rather than allowing the model to generate unconstrained responses. Synthetic data can augment underrepresented populat in training sets without privacy risks. Reinforcement learning from human or ex feedback aligns model behavior with fairness and empathy standards. Rigorous p deployment testing using LLM-as-a-judge evaluation frameworks can ensure transparency, reliability, and fairness before integration into care. These aren't theoretical solutions, they're practical approaches that can be implemented toda

Third, workforce displacement concerns deserve attention even though the fram often wrong. Medicaid already faces a structural mismatch between service dema and workforce supply, especially in LTSS and HCBS where workforce shortages severe and getting worse as the population ages. AI should be viewed not as repl human workers but as augmenting human capacity. By automating routine tasks triaging cases, and freeing professionals to focus on hands-on, relationship-base care, AI can actually improve job quality while expanding access. A care coordin who spends half their day on paperwork and phone tag can spend that time buil relationships with members instead if the administrative tasks are automated.

Fourth, technology capacity poses a real constraint. Many Medicaid agencies lac in-house expertise and infrastructure to deploy and govern AI tools effectively. Fragmented data systems and limited interoperability hinder integration with healthcare and social service networks. Different states use different eligibility systems, different managed care organizations use different care management platforms, providers use different EHRs. Getting these systems to talk to each ot hard, expensive, and often requires negotiating data sharing agreements that rai privacy concerns. Trust is also a barrier, with justified fears that data sharing co lead to surveillance or be used in ways that harm beneficiaries.

Addressing these challenges requires investment in secure, interoperable system alongside transparent governance and genuine community engagement. You can deploy AI tools and hope people trust them. You have to build trust through transparency about how the systems work, clear policies about data use and priv meaningful oversight mechanisms, and demonstrated commitment to fairness ar accountability.

## The Regulatory Reality

One of the most valuable contributions in Pair Team's brief is clarity about FDA regulation, which often creates anxiety for healthcare AI developers. The reality more favorable than many assume, especially for the use cases Pair Team is focu on.

The key distinction is whether the AI is making patient-specific diagnostic or treatment decisions that replace clinical judgment, or whether it's supporting navigation, coordination, and logistics. Digital assistants that handle appointme scheduling, eligibility verification, care navigation, reminders, and basic membe education fall outside FDA jurisdiction under the 21st Century Cures Act becau they don't diagnose or treat medical conditions.

Even more sophisticated tools can qualify as non-device Clinical Decision Suppo (CDS) if they meet four criteria: they're healthcare provider-facing (not direct to consumer making medical claims), they show their inputs and reasoning transparently, they support rather than replace clinical judgment, and they allow clinicians to independently review and verify recommendations. This is an impo safe harbor that Pair Team's brief highlights.

So if you're building an AI that reminds a member with asthma to refill their inh tracks symptoms through simple questions, and suggests they message their doc things are getting worse, that's not a medical device. If you're building an AI tha reviews home monitoring data and tells a clinician "this patient's asthma control poor based on these usage patterns, consider stepping up therapy," that could be device CDS if designed to meet the four criteria. If you're building an AI that

automatically adjusts medication dosing without physician review, that's definite
device and requires premarket review.

The smart strategy Pair Team implicitly recommends is to intentionally scope
products to non-device functions initially. Get to market, prove value, build trus
Then layer on clinical decision support features designed to meet CDS criteria. I
future versions need patient-specific diagnostic or therapeutic capabilities that c
into device territory, FDA has Predetermined Change Control Plans (PCCPs) tha
allow companies to define in advance how models will evolve, what validation m
will be used, and how performance will be monitored. Submit the PCCP with in
review and you can iterate the model without repeated submissions, which is cru
for AI systems that need to adapt and improve over time.

This regulatory pathway is actually quite workable for the vision Pair Team lays
Most of what Precision Benefits requires (care navigation, eligibility support,
transportation coordination, connection to community resources, appointment
reminders, medication refill prompts) falls clearly outside device regulation. The
that touch clinical care can be designed as non-device CDS with appropriate
safeguards. This is a much clearer and faster path than trying to get device appr
for every feature.

There are also state-level AI governance considerations. Some states are passing
requiring bias audits, transparency reports, and human review for automated
decisions that affect people's access to services. These requirements are generall
aligned with best practices anyway. If you're building responsibly with bias testi
human oversight, and clear explanations, compliance shouldn't be a heavy lift. In
Pair Team's brief proposes that Medicaid AI systems should include clear fallba
human support pathways, transparent explanation of reasoning, user feedback lo
(like "I'm not satisfied, take me to a human"), and ongoing trust monitoring thro
surveys and metrics.

## Pair Team's Implementation Strategy

The brief lays out a comprehensive strategy for how AI adoption in Medicaid sh actually happen, with specific recommendations for federal and state action. Thi where Pair Team moves from vision to practical implementation.

At the federal level, the Centers for Medicare and Medicaid Services (CMS) need take a leadership role. Pair Team proposes CMS launch an AI Innovation Challe fund pilot projects that enhance administrative efficiency, streamline eligibility, improve care coordination. This would provide the financial resources and polit legitimacy for states to experiment without bearing all the risk themselves. A St Medicaid Director's Letter outlining best practices would offer implementation guidance, addressing practical questions like whether payments for digital assis should count under Medical Loss Ratio versus Administrative Loss Ratio, how t structure contracts with AI vendors, and how to evaluate effectiveness.

The brief calls for states to create AI Centers of Excellence, modeled in part afte Centers of Excellence in Regulatory Science and Innovation (CERSI) that exist fo medical devices and pharmaceuticals. These would serve as dedicated hubs brin together policymakers, technology developers, healthcare providers, and health to evaluate AI in real-world Medicaid settings. Priorities could include predictiv analytics for early detection of health and social risk, AI-assisted care navigation outreach, and public-private partnerships for scaling innovation. These centers also serve the crucial function of developing shared standards and evaluation frameworks so every state doesn't have to reinvent the wheel.

Pair Team emphasizes that AI tools should undergo rigorous bias audits and benchmarking to ensure fairness in access and outcomes. Training models on Medicaid-representative data improves accuracy, but for large language models, fairness also depends on safeguards like human-in-the-loop review for sensitive decisions, output constraints that prevent harmful or biased responses, and prospective monitoring for differential impacts across populations. The brief proposes pilot user testing that emphasizes acceptability and trust measures bef scaling, recognizing that technical performance alone doesn't guarantee real-wo success.

Beyond pilot programs and operational improvements, Pair Team argues that CI should signal a bold ambition: to establish Medicaid as a leading example of an A first public service. This isn't just about making the current system more efficie about demonstrating what's possible when public programs embrace intelligent technologies to better meet individual needs. The implications extend far beyon Medicaid to Medicare, other safety net programs, and public services globally.

## The Path Forward

What Pair Team is building represents more than incremental improvement to a existing system. It's a fundamental reimagining of how safety net programs can function when designed around AI infrastructure from the ground up.

The Precision Benefits framework addresses a core limitation of traditional care management: the gap between identifying risk and actually intervening before cr occurs. Predictive models have been around for years but without automated intervention capabilities, those predictions sit unused. By combining prediction real-time, personalized, automated outreach and support, Pair Team's approach that gap. The AI can act on signals immediately, at scale, with zero marginal cos interaction.

The embedded staff model solves another fundamental problem, which is that Medicaid beneficiaries interact with dozens of different touchpoints (clinics, pharmacies, community organizations, benefits agencies, schools, housing progr and none of those systems talk to each other effectively. By making AI assistants available to everyone who touches Medicaid, not just as a member-facing portal, Team creates connective tissue across a fragmented ecosystem. The food bank w the school nurse, the 211 operator all get augmented with tools that know how t navigate complexity on behalf of the people they serve.

The focus on 24/7 multilingual support addresses access barriers that workforce expansion alone can't solve. You can't hire enough bilingual care coordinators to every language, every time zone, every moment when someone needs help. But yo build AI that operates around the clock in any language with cultural competen

built in. This isn't about replacing human connection, it's about making sure eve has access to basic support regardless of when they need it or what language the speak.

The emphasis on behavioral health is particularly important given the massive a gaps in this area. The Therabot study that Pair Team cites showed remarkable re 51% reduction in depression scores, 31% reduction in anxiety, 19% reduction in e disorder risk from a generative AI mental health chatbot. Participants reported feeling meaningful connections with the AI. This doesn't replace therapy or psychiatric care, but it provides support that many Medicaid beneficiaries simpl can't access today because of provider shortages and wait times.

What makes Pair Team's work distinctive is the combination of technical sophistication, practical implementation focus, and genuine commitment to responsible development. The brief doesn't handwave away concerns about bias, fairness, privacy, and safety. It addresses them directly with specific mitigation strategies. It acknowledges workforce concerns while reframing AI as augmenta rather than replacement. It provides regulatory clarity that makes the path forw navigable rather than daunting.

The economic case is compelling with potential savings of $37-67 billion annual from administrative efficiency alone. But the human case is what really matters. Medicaid serves people who face multiple overlapping challenges: chronic healt conditions, social determinants of health, language barriers, transportation issue housing instability, food insecurity. The current system is designed for these challenges to be addressed separately, if at all, by disconnected programs and providers who don't communicate effectively. Pair Team's vision is a system whe coordinates across all these domains to catch people before they fall, providing t right support at the right time to keep everyone on an upward trajectory.

The work that Pair Team is doing extends beyond building a product. They're establishing a model for how public programs can leverage AI to better serve vulnerable populations. That model has implications for Medicare's 65 million beneficiaries, for safety net programs serving hundreds of millions globally, and

public services far beyond healthcare. This is infrastructure work in the truest s
Not the flashy stuff that makes headlines but the foundational systems that dete
whether society functions equitably or leaves people behind.

Pair Team is building the rails for an AI-first safety net, and if they succeed, the
implications will reshape how we think about the relationship between technolo
and social support. The question isn't whether AI will transform public progran
whether that transformation will be done thoughtfully, responsibly, and in ways
actually serve the people who need help most. Pair Team's work in Medicaid sho
what's possible when you get that approach right.