personify HEALTH

eBook

The Self-Funded Employer's Guide to AI: What to Embrace and What to Avoid

Three Use Cases in Employee Benefits Where AI Fills Unmet Needs—and Why Humans Are Still at the Center of a Tech-and-Touch Approach

Executive Summary

When business leaders engage artificial intelligence (AI), they're usually after two things: the ability to scale and the ability to predict. In benefits management, AI delivers value in both ways—and by and large, people agree. One survey found that roughly 7 in 10 benefits leaders think AI will positively affect benefits administration.

73%

of benefits leaders think AI will improve benefits administration.

At the same time, artificial intelligence isn't human intelligence. And in a world as uniquely people-centric as employee benefits, there has to be some governance between "high-tech" and "high-touch" as AI expands its footprint.

At Personify Health, we've given a lot of thought to this—having coupled AI enablement with a people-powered approach. And what have we learned? It's nuanced.

In some contexts, AI is the apparatus that can help self-funded employers jointly solve two of their biggest concerns: rising healthcare costs (nearly 8% higher in 2025) and a fraught employee experience (52% of people need help navigating the healthcare system). Automation unlocks cost-efficient scale and predictive insights as it addresses those concerns, supporting user personalization; payment integrity and fraud, waste, and abuse (FWA) detection; and basic member support.

AI can help self-funded employers solve two of their biggest concerns:

high healthcare costs and a fraught employee experience.

On the other hand, business leaders and their employees have valid concerns about AI's security risks and accuracy. Further, many worry about the technology's inherent biases, bringing a new term to the lexicon known as ethical AI that's aimed at eliminating systemic concerns, including racial prejudice.

These are all reasons why certain areas of benefits administration—like content expertise, complex navigation, and denials management—are best kept people-driven, with strict protocols and monitoring in place. This requires a unique blend of tech and touch, where humans tackle some functions and technology addresses others.

In this eBook, we'll explore what we consider the "go" and "no-go" zones for AI in benefits management across three key areas: wellbeing, navigation and advocacy, and third-party administration (TPA). Learn the use cases where it makes sense to adopt AI, where it may not, and how to balance it all for an efficient and seamless experience that serves employers and employees.

Use Case in Wellbeing: Algorithm-driven personalization

People expect personalized experiences in all areas of health and wellbeing, even (and especially) from their employers. AI meets those needs at scale.

People have grown accustomed to tailored experiences in everyday life, especially for topics as personal as healthcare. One survey found that 71% of consumers expect personalization from the brands they interact with, for example.

7 in 10

consumers expect personalization from brands.

And yet, consumers aren't getting personalization for their most personal of needs, their healthcare journey: A Personify Health survey found that only 23% of employees say health and wellbeing programs are initiated in a way that gives them personalized recommendations that matter to their health journey. It's hard for employers to do this with existing resources, especially with a modern melting pot of diverse workforce geographics and psychographics.

23%

of employees say health and wellbeing programs provide personalized recommendations.

This is a classic use case for AI in benefits administration: personalization at scale. The technology can monitor an employee's behaviors, preferences, and evolving health status by analyzing data from multiple sources, including claims, self-reported information, and third-party consumer data.

In the case of Personify Health, this data combines to inform the types of content employees see—from their daily tasks, such as scheduling reminders, to educational content about relevant topics like GLP-1s and heart health. One factor in this personalization engine is "data twinning," which Kerry Hudgins, group product manager at Personify Health, explains this way:

"Collaborative filtering is one example of data twinning where we can look at somebody similar to you in age, gender, employer, geography, etc.," she said. "We couple that with interaction data, such as a person's unique habits and journeys. So if you and I are of similar demographics and we engage in similar habits and journeys, but there is one activity that I've done that you haven't, the platform might recommend that action to you."

We have also expanded this capability a step beyond with predictive modeling. Built in-house, these models can help predict the likelihood of having a health risk or condition.

"A key use case for a predictive model that identifies members at high risk for pre-diabetes is to support early intervention—before a formal diagnosis is made. By leveraging patterns in clinical, behavioral, and demographic data, the model can surface individuals who may not yet show overt symptoms but are trending toward metabolic risk. This enables targeted, preventive outreach—such as nutrition counseling, physical activity programs, or lifestyle coaching—aimed at reducing progression to diabetes and lowering long-term healthcare costs."

- **Data twinning:** Using similarities in demographics to influence recommendations.
- **Predictive modeling:** Addressing data gaps to predict existing diagnoses that can inform recommendations.

That's the tech. Here's the touch: While AI may be driving content recommendations at a workforce-wide scale, it can't (and shouldn't) do everything. Some functions remain human-driven.

Expertise from clinical and behavioral experts is one example. Our educational content is human-authored and human-reviewed, eschewing some of the most egregious concerns with AI-generated medical content. The boom of "AI wellness" furthers risks overpromising what AI tools can (or should) do for employees. Some leaders fear that the technology is not mature enough for many of these use cases.

This is why having a comprehensive layer of human-to-human support is critical. After all, as important as those personalized recommendations are to users, human care management is even more critical to give that final nudge toward action. An article may introduce the importance of specialty care, but the care manager is the one who'll facilitate booking the visit or navigating authorizations, for example.

"Recommending and even initiating human connections where appropriate helps support people directly when needed," said Jeff Yoshimura, chief product officer at Personify Health. "While technology—and AI specifically—can enable personalization at scale, the human element will absolutely and always be the most important one."

Use Case in Navigation and Advocacy: Intelligently guided self-service

AI gives employees a balanced blend of both—the ability to solve their own problems while also having an easy path toward human support.

More than 70% of people feel like the healthcare system fails them. That's unsurprising, given the complex and confusing web of networks, authorizations, claims, cost sharing, and access. As they navigate these challenges, most people—81%—try to find their own solutions before contacting customer service.

8 in 10

people try to self-resolve issues before calling customer support.

Having a straightforward self-service path to answer these questions can mean everything to a frustrated employee. That's why many benefits stakeholders are investing in AI-optimized solutions that incorporate member chatbots.

Two such resources are available to employees using the Personify Health platform: member services chatbots and technical support chatbots. Trained on frequently asked questions provided to the models, these solutions use natural language processing to interpret members' questions and then match them with the right educational content.

For example, someone might ask the member services chatbot a question about how to file a claim, or the technical support chatbot about how to reset a password. The technology matches the question with a relevant resource, helping members feel supported as they navigate self-service.

Chatbots use natural language processing to interpret questions and match members to relevant support.

These solutions are always evolving too, notes Hudgins. While there is an opportunity to escalate questions to human support in complex cases—through a fully staffed call center and email, for example—we at Personify Health recognize the extent to which humans inherently want to self-solve before asking for help (even when that help is readily available).

Again, it's that mix of tech and touch.

"We are always exploring new iterations on these chatbots, including our current focus, which is to weave in resources that are available on the platform into the chat itself," Hudgins said. "But at the same time, we realize that some problems require live customer support. We want to continually focus on tech advancement and expansion without losing sight of having those essential channels for human connections."

"It's always going to be a blend of technology and people."

Use Case in Third-Party Administration: Fraud, waste, and abuse reduction

By identifying wasteful spending that disproportionately affects socially vulnerable members, AI-optimized platforms can support both cost containment and health equity.

Fraud, waste, and abuse (FWA) is a significant problem that can acutely hit self-funded employers. In all, there are nearly $1 trillion in wasted claims or services each year in the U.S.—but many commercial plans have the mechanisms and might to identify those wasteful services. Self-funded employers, on the other hand, often have to manually flag and fend off such costs on their own.

$1,000,000,000,000

There is nearly $1 trillion in wasted claims or services each year in the U.S.

AI is helping meet these unmet needs by enabling FWA identification at scale, particularly when deployed through a TPA like Personify Health.

A recent study analyzed this use case in a joint effort between Personify Health, Health at Scale, MIT, and the University of Michigan. Peer-reviewed and published in NEJM Catalyst Innovations in Care Delivery, the study observed how AI could help identify FWA paid claims of $4,000 and higher. By using AI to escalate probable areas of waste to human reviewers, the initiative saved $11.8 million over eight months, equating to 1.2% of total spending.

AI-enabled FWA saved $11.8 million over eight months. More than half of the claims sent for human review were reduced

Of all claims sent for human review, 54% were reduced, resulting in an average savings of $3,916 per claim. Socially vulnerable members captured even more, with an average of $3,973 saved per claim—an important finding, given that FWA occurs more frequently in these populations. The highest savings came from cardiovascular and critical care services, each yielding an average reduction of more than $10,000.

While AI is automating FWA identification, AI should not be automating denials, however. At Personify Health, we recognize that this is where that critical element of touch comes in: Humans are still the ones reviewing escalations for final adjudication, an important nuance in the blend of platforms and people.

"There is a lot of concern about the level of AI involvement in FWA identification, but we're not in the business of letting platforms decide a person's access to healthcare services," Yoshimura said. "Human reviewers are the last stand. AI doesn't replace them—it augments them as the identifier and not adjudicator around probable FWA."

AI-enabled FWA engines escalate questionable claims to human reviewers. They should not automate denials.

Where AI works—and where it doesn't

Green Light—Go	Red Light—Stop
Content recommendations: Personalize the member experience with algorithm-driven content feeds tailored to each employee's unique needs.	**Content expertise:** AI can help inform who gets served what content, but leave the actual clinical and behavioral expertise to the humans.
Basic navigation: Provide members with a self-service opportunity to resolve common issues, such as searching for in-network providers or finding benefits.	**Complex navigation:** AI can nudge members in the right direction for basic concerns, but leave the complex topics to call centers, email, and other channels for live human support.
FWA identification: Flag and escalate potentially incorrect claim amounts.	**FWA adjudication:** AI can identify areas of potential waste, but leave the actual review of those escalated claims to the humans.

Capture opportunities, control for risk

Employers and employees have concerns about AI, creating a need to capitalize on the opportunities of automation while addressing and mitigating the risks.

As the technology matures (particularly with conversational apps), AI is gradually stepping into consumers' routines—with many people regularly using the tool at work, home, or both. OpenAI reports that there are 400 million ChatGPT users (and growing), for example.

It's also becoming a pivotal part of business: More than 60% of leaders have reported investments in AI, according to a 2024 report. At Personify Health, we too have seen gains from AI involvement, reducing time spent on model development from months and weeks to hours, according to Amit Jain, chief technology officer at Personify Health. Jain is also an inaugural member of Bloomberg Tech's Innovators Circle.

"Personify Health is accelerating critical innovation in healthcare and wellbeing for millions of our members through AI, machine learning, and large language models—all with a strong commitment to data privacy, ethical standards, and responsible AI," Jain recently said in a post announcing his participation in the group.

But despite the fact that individuals are warming up to the technology's presence in everyday life and business operations, worries still linger—especially in health and benefits. One survey among benefits leaders found that the biggest concerns were data security and accuracy.

Top concerns with AI in benefits administration:
Data security (34 %) and accuracy (22%)

As they turn to partners like Personify Health, one way for self-funded employers to build confidence in these technologies is by requesting documentation. That should include an outlined governance process of how data will be:

- **Unbiased, including free of racism or stereotypes**
- **Accurate and free of hallucinations**
- **Complete and wholly representative**
- **Secure with rigorous privacy controls**

Mitchell Carr, project manager at Personify Health, explains:

"Employers should expect a documented process about how AI is used and in what capacities, now and in the planned product roadmaps for the future," he said. "It's that documentation and governance that provides that reassurance that your data is safe in the platform, it's accurate, and it won't be sold or sent anywhere else. "

"And most importantly, that governance helps form the line between where technology exists and where humans step in—people should always be kept involved in a human-in-the-loop model."

These recommendations align with others in the industry. According to a report from Deloitte, top challenges of generative AI include problems with data strategies, systems, processes, and hallucinations . Keeping humans in the loop and continuously learning are among the key opportunities to foster responsible AI use for improved ethics, accuracy, security, and compliance.

Businesses are also increasingly interested in AI assistants that execute administrative tasks. We embrace these technologies internally to capture notes during coaching sessions, as well as processing data into holistic, compatible formats.

Enabling the dual strengths of tech and touch

CAI enables scale and predictive power at an unprecedented rate. As self-funded employers have to do more with less, that dual advantage is proving to be a significant one in benefits administration.

Three use cases where AI offers immediate value include wellbeing through personalized recommendations, navigation and support through member self-service, and third-party administration through FWA detection.

But keep in mind that in those applications and others, the technology is an enabler: delivering expertise but not generating it , navigating basic issues but not complex ones, and automating waste detection but not waste denials.

"AI opens up new worlds for self-funded employers, but as those use cases play out, we as an industry have to be intentional about keeping humans involved," Yoshimura said. "People are still the arbiters of technology, and Personify Health is responsibly leading the way. We'll always need a high-tech, high-touch approach."

If your organization is interested in these AI use cases for optimal benefits administration, learn how we can help at PersonifyHealth.com.

About Personify Health

By bringing industry-leading third party administration, holistic wellbeing, and navigation solutions together, all in one place, we have created the industry's first and only personalized health platform. With decades of experience and global operations, we empower diverse and unique businesses – and diverse and unique people – to engage more deeply in health at a lower cost. Through our proprietary combination of data-driven personalization, science-backed methodology, and concierge-level clinical expertise, our end-to-end platform makes it easier to proactively address people's needs across their lives. With a personalized, holistic, and powerfully simple experience, we are redefining industry expectations and what it means to manage health.

Let's connect →

personify™ HEALTH