The AI Catalyst Pulse

February 20th, 2025

Upcoming AI Catalyst Events

Mark your calendars

- **Mar 26: Executive Deep Dive: How AI is Turbocharging Patient Flow and Throughput** | *Register here*
  - Join our Executive Deep Dive to learn how health systems are utilizing AI to drive significant ROI within patient flow and throughput. We will share how AI Catalyst members are approaching ROI calculations for AI solutions in the patient flow space, as well as implementation tips to ensure your AI investments are improving ClinOps efficiency and care coordination.
- **Apr 10: Candid Vendor Conversations - Perspectives on Palantir** | *Register here*
  - Join a candid discussion with executive peers who have implemented or evaluated Palantir's healthcare AI solutions. Get unfiltered insights about implementation challenges, cost considerations, and realized benefits. You'll hear direct feedback about integration complexity, staff adoption experiences, and how Palantir's solutions compare to alternative vendors in the market.
- **Apr 15: Epic Systems and the Future of AI-Assisted Nursing Documentation** | *Register here*
  - Learn about Epic's evolving AI nursing documentation solutions, their impact on clinical workflows, and real-world implementation guidance. We'll be joined by Epic's development experts for a candid discussion of Epic's development roadmap, as well as how to address key issues relating to workforce buy-in and adoption of these solutions.

Anthropic Just (Accidentally) Released the Biggest Survey Yet of 'Shadow AI' in Healthcare

What happened: The AI company Anthropic has released what it calls the Anthropic Economic Index, "the first large-scale empirical measurement of which tasks are seeing AI use across the economy." Anthropic researchers examined millions of conversations with their AI assistant, Claude, and mapped each one to job-related tasks defined by the U.S. Department of Labor.

Their findings show that while jobs across the economy are using conversational AI every day — with 36% of occupations using AI for at least a quarter of their tasks — healthcare lags significantly behind other sectors, including other regulated industries like financial services and law.

Why it matters: Anthropic's goal with their "economic index" was to track AI adoption across industries. But by releasing their underlying data, they accidentally gave healthcare leaders something far more valuable: the best evidence to date of how healthcare workers are using AI behind their organizations' backs.

The key is that, for reasons related to Anthropic's data agreements, the study *excluded* activity from business customers. As such, it largely shows how healthcare workers are using AI *outside of official contexts* – giving us what may be our best audit ever of so-called "shadow AI."

Here's what the AI Catalyst team found when we crunched Anthropic's healthcare data.

What healthcare workers seem to be doing in AI's shadows

Before we dive in, an important note about limitations. To protect user privacy, Anthropic didn't associate the conversations in their dataset with specific users. Instead, they used AI to group conversations into job-related tasks defined by the U.S. Department of Labor.

In other words, while Anthropic could determine that a given conversation related to *a task typically conducted by a healthcare professional*, they couldn't say for sure that the user was, in fact, a healthcare worker. That said, Anthropic said the vast majority of conversations appeared work-related.

With that disclaimer in mind, we analyzed the healthcare-related queries in Anthropic's data. Broadly speaking, they fell into three categories:

- **Backend administrative work (~44%):** This is all the "behind-the-scenes" stuff required to make healthcare run. Think about internal documentation, reporting, and workflow optimization.
- **Patient communication (~36%):** Everything related to what we say to patients. Education materials, procedural explanations, and counseling content.
- **Clinical applications (~20%):** The meat and potatoes of clinical care: diagnosing, interpreting tests, and developing treatment plans. Examples include AI support for diagnostics, analysis, and treatment planning.

Digging deeper: The most worrying uses of 'shadow AI' in healthcare

Some of these uses seem innocuous. But based on our analysis, about one-third of the queries in Anthropic's data raise what we'd consider "red flags" – suggesting risks of patient harm, legal violations, or privacy concerns. To give a few examples:

- "Interpret[ing] laboratory results and communicat[ing] findings to patients or physicians"
- "Assess[ing] the identity, strength, or purity of medications"
- "Interpret[ing] the outcomes of diagnostic imaging procedures"
- "Develop[ing] individualized treatment plans for patients, considering patient preferences, clinical data, or the risks and benefits of therapies"

We hardly need to explain how troubling it is if healthcare professionals are truly using Claude – a consumer-facing, non-specialist AI chatbot – for these tasks at work without authorization. For one thing, AI makes mistakes that could threaten patient care. Just as worrying, if clinicians are sharing enough context for AI to respond meaningfully to these requests, they may be disclosing protected health information.

To be clear, this data *doesn't* prove that any particular red-flag conversation involved a healthcare professional acting inappropriately – but because "shadow AI" by definition happens in unapproved contexts, it's almost impossible to definitively measure such misuse.

Even so, I think it's reasonable to say this data suggests that perhaps one-third of healthcare's "shadow AI" use poses serious risks. That's plenty enough to be scary.

What AI use can you bring in from the 'shadows'?

Still, there's a more optimistic view of the data: Perhaps two-thirds of healthcare's "shadow AI" use is potentially *acceptable*, if done through proper channels. This includes tasks such as:

- "Writ[ing] research reports and other publications to document and communicate research findings"
- Reviewing reports "for spelling, grammar, clarity, consistency, and proper medical terminology"
- "Prepar[ing] statistical reports, narrative reports, or graphic presentations of information, such as tumor registry data"

These patterns reveal what your healthcare workers truly want from AI: help reducing administrative burden without compromising clinical judgment. They're just currently getting that help through unauthorized channels.

For health system leaders, this data provides both warning and roadmap. Your staff will use AI — if not through official channels, then in the shadows. The question is how to bring the best AI uses into the light while shutting down the dangerous ones.

Questions to consider:

1. Which specific AI use cases from this data could our organization safely implement through official channels?
2. What immediate steps should we take to detect and prevent high-risk clinical applications, particularly those involving protected health information?
3. How might we adapt our IT procurement processes to better accommodate the rapid pace of AI tool development while maintaining necessary safeguards?

Value Capture vs. Value Creation: Your First Critical AI Decision in 2025

First in a four-part series examining the critical strategic dilemmas healthcare executives must navigate in 2025.

Every dollar your health system invests in AI has to serve one of two goals: securing a larger share of today's healthcare dollars, or transforming how care gets delivered over time. This is healthcare's first fundamental AI dilemma in 2025: Do you prioritize shorter-term value capture or longer-term value creation?

(A small disclaimer: All of the dilemmas we'll explore in this series are, to some extent, false binaries; you *can* do some value capture while also pursuing value creation. That said, we've heard repeatedly from AI Catalyst members that budgets are tight and that they must make hard trade-offs. Our goal is to provide frameworks that help you weigh the pros and cons of each approach and give clear guidance to your team.)

The case for value capture: You need immediate wins in the 'AI arms race' with your payers

The revenue cycle arms race is accelerating. One survey found that final denial rates for inpatient care surged 51% between 2021 and 2023, and denials have only increased since then as payers have

deployed increasingly sophisticated AI tools. It's debatable whether payers' tools are terribly accurate — in one recent case, UnitedHealth was accused of using an AI model with a 90% error rate to deny Medicare Advantage claims — but their impact is undeniable.

Fortunately, you're not helpless in this fight. Health systems are deploying their own AI tools and seeing dramatic results:

- One AI Catalyst member boosted revenue by $15 million in year one using AI-powered CDI.
- Emory Healthcare cut appeals letter writing time by 90%, with 75% of AI-drafted appeals succeeding.
- Mass General Brigham's AI tools identified 90% more patients with adverse social determinants of health, potentially improving the likelihood of prior authorizations.

And you can achieve at least some of these gains *fast*. Some organizations report that they've been able to begin drafting AI-powered appeals letters within hours of launching a vendor tool.

But there's a catch: Because your payers are also deploying AI to advance their own bottom line, your early gains may not last. The health system that saw a $15 million boost in first-year CDI revenue, for instance, saw the gains drop to $7 million in year two as payers adapted.

As Dr. Robert Wachter of UCSF pithily framed the challenge to the *New York Times*, "[Payers'] AI will deny our AI, and we'll go back and forth."

The case for value creation: You could transform healthcare – and see compounding gains over time

Some health systems are making a longer-term bet: focusing their discretionary AI investments on transforming care delivery. In other words, rather than engaging in a zero-sum battle with payers over every dollar currently spent on healthcare, they're trying to use AI to provide better, more efficient, or just plain *new* care.

The early results are compelling. For instance, physicians at Texas Health Resources who use DAX CoPilot to automate their clinical documentation are saving more than five hours weekly — time they're using to see more patients while actually increasing face-to-face interaction. At WellSpan Health, AI virtual assistants are handling routine patient outreach and education around colonoscopies that once consumed valuable nursing hours.

Unlike revenue cycle gains, which evaporate as payers adapt, these operational improvements compound over time. Staff get better at working with AI tools. Patients experience better care. The benefits multiply.

Which should you choose: Value capture vs. value creation?

While you'll undoubtedly need to do both *some* value capture and *some* value creation, trying to do both aggressively risks muddling your message and confusing your staff. You need a clear strategic emphasis.

Our guidance: Consider prioritizing value capture if you're under immediate financial pressure and need rapid ROI, or if you've seen sharp recent *increases* in your denial rates (which might suggest that your payers are already beating you in the AI "arms race"). In that case:

- Identify the “pain points” in your revenue cycle (for example, prior authorizations or CDI) and choose AI partners who target those specific areas.
- If in doubt, start with appeals letter automation, which typically shows the fastest returns.
- Use today’s “value capture” efforts to build the AI governance and technology infrastructure that will be required for tomorrow’s “value creation” efforts.

On the other hand, consider prioritizing value creation if you have short-term financial breathing room and your payer mix (for instance, a high proportion of traditional Medicare) makes you less vulnerable to AI-driven denials. This path also makes sense if your organization is already excited about AI's potential, since you'll need broad buy-in for transformational efforts. In that case:

- Focus first on proven use cases like clinical documentation and patient flow optimization, where other systems have already demonstrated lasting impact.
- Build AI literacy across your organization deliberately, as large-scale transformation requires staff who understand both the potential and limitations of AI tools.
- As you demonstrate success with initial projects, measure both financial and operational impacts. You'll need this data to justify continued investment in transformation over short-term revenue gains.

Finally, here’s an intriguing data point: When we posed the “value creation vs. value capture” dilemma to health system CEOs and board chairs at The Health Management Academy's Trustee Summit, 82% favored value creation. That's a contrast to what we’ve seen from CIOs and COOs, who are more evenly split between the approaches.

The gap suggests that while healthcare’s top leaders envision AI primarily as a transformative force, the immediate pressures of revenue capture are harder to ignore on the ground.

Questions to consider:

1. Which specific revenue cycle processes are losing you the most money to denials? That's where you need AI defense first.
2. If you choose value creation, how will you explain to your board why you're accepting short-term revenue losses for long-term transformation?
3. How will you create organization-wide understanding of your strategic choice and its implications?

AI Strategy Quick Hits

Noteworthy moves from peers to implement AI technologies

Catch up on health systems who have established new leadership roles for AI development:

- **AdventHealth** appoints Rob Purinton as its first Chief Artificial Intelligence Officer
- **The University of Utah** appoints its first-ever Chief of Artificial Intelligence Officer, Manish Parashar

- **Providence** CEO announces new leaders focused on responsible AI adoption and innovation
- **Mount Sinai Health System's** new Chief Digital Officer discusses digital transformation plans

Here's the latest on health system partnerships with AI developers:

- **St. Luke's Health System** negotiates deal with **Ambience Healthcare** that is reportedly 10x less expensive than other market leaders
- **Cedars-Sinai** pilots AI-powered mobile app, Aiva Nurse Assistant, to help nurses reduce documentation burdens
- **Sutter Health** and **GE Healthcare** establish seven-year strategic partnership to enhance diagnostic imaging services
- **UNC Health** partners with **Abridge** after successful pilot with 100 clinicians
- **AdventHealth** partners with **ParkourSC** to enhance supply chain operations
- **Mass General Brigham** collaborated with **IBM** to develop an AI tool for predicting extreme heat events
- **Mayo Clinic** pilots **VoiceCare** AI to automate back-office processes

Other strategy quick hits:

- How **Mayo Clinic Health System** plans to enhance its AI initiatives
- **Vanderbilt University Medical Center** establishes an AI Patient and Family Advisory Group
- **UPMC Enterprises** develops Ahavi, a virtual environment for testing and refining AI models

Emerging Use Cases

New capabilities that indicate AI's potential

Discover how AI technology is transforming mental health care:

- A new study on AI's effectiveness in detecting suicide risk among patients in neurology clinics
- Healthcare professionals in Wisconsin explore AI in mental health care

AI-enabled technology continues to support early detection of conditions:

- **Mass General Brigham** is using LLMs to enhance cancer treatment
- **Providence Saint Joseph Hospital** and **Providence Redwood Memorial Hospital** adopt Ceribell for seizure detection
- **Hackensack Meridian Mountainside Medical Center** integrates two AI-powered technologies to enhance diagnosis and treatment of coronary artery disease (CAD)
- **FDA** grants breakthrough device designation to the Serial CTRS AI model for non-small cell lung cancer patients

Other emerging use cases:

- A **Duke University Health System** study found GPT-4 significantly improved patient comprehension of discharge summaries
- **Vanderbilt Medical Laboratories** transitions to digital pathways for tissue sample imaging
- How the emergence of 'AI agents' may impact healthcare operations
- Researchers experiment with AI's effectiveness in the treatment of persistent atrial fibrillation
- Researchers propose using AI-cloned voices to help patients with chronic pain
- **The University of Michigan** is leading a project to develop AI-powered mobile clinics in rural areas

Market Moves

A round-up of AI company announcements and stories

- **Rad AI** completes a $60M Series C funding round
- **Teladoc Health** plans to acquire **Catapult Health** for $65M to enhance its virtual preventative care services
- **Quibim**, a startup focused on AI-driven medical imaging analysis, raises $50M
- **Oscar Health** builds on its AI use cases in response to its expanding membership

Policy Updates

Understanding the evolving AI regulatory and legislative landscape

- Vice President Vance pushes for AI deregulation
- Former FDA commissioner advocates for FDA to revert to its previous regulations regarding AI in clinical decision support software
- Judge allows AI-related lawsuit against **UnitedHealth** to proceed
- **FDA** releases draft guidance addressing the use of AI data in drug submissions

The News in Numbers

Data points that caught our eye

40% of health systems say their organizations have already experienced a significant-to-moderate return on their generative AI investments, according to a recent Deloitte survey

66% of US adults in a recent JAMA study reported low trust in their health care system to use AI responsibly

66% of surveyed physicians in a recent AMA study indicate that they currently use AI in their practices (up 28% from the previous year)