


# Labor Market Disruption from AI in Healthcare: Where the Real Money Is

MAR 07, 2026



---

## Table of Contents

Abstract

Section 1: What the Anthropic Research Actually Says

Section 2: The Gap Between Capability and Deployment Is the Story

Section 3: Healthcare Is Not One Labor Market

Section 4: Why Insurance Gets All the Press But Misses the Point

Section 5: Care Delivery Is Where the Leverage Lives

Section 6: What the Hiring Slowdown Signals for Health Tech Founders

Section 7: How to Think About This as an Investor

## Abstract

This essay uses the March 2026 Anthropic labor market report as a launching point to think through AI's real impact on healthcare employment. Key data points and arguments include:

- Anthropic's "observed exposure" framework shows a massive gap between theoretical AI capability and actual deployment across all industries

- Healthcare practitioners rank lower on observed exposure than most expect, with medical record specialists (66.7%) and customer service reps in health settings (7 rank very high)
- No measurable increase in unemployment among highly exposed workers as of 2026, but a 14% drop in hiring of 22-25 year olds into exposed roles
- Healthcare labor costs represent 55-65% of total operating expenses for most health systems, far exceeding the 20-30% typical of health insurers
- The case that care delivery, not insurance, is the highest-leverage target for AI displacement
- Implications for founders and investors thinking about where to build

## **Section 1: What the Anthropic Research Actually Says**

There is a new Anthropic paper worth reading carefully before making any strong claims about AI eating jobs. Published in March 2026 and authored by Maxim Massenkoff and Peter McCrory, it introduces something called “observed exposure” which is a more grounded way of measuring how much AI has actually penetrated given occupation versus how much it theoretically could. The distinction matters enormously, and the healthcare implications are significant enough that anyone writing checks into digital health or building companies in the space should internalize the core findings before making assumptions about where the displacement story is headed.

The paper builds on earlier work, specifically the Eloundou et al. framework from 2023 that rated tasks on a simple scale: fully automatable by an LLM alone, automatable with additional tools, or not automatable at all. That framework was useful but purely theoretical. What Massenkoff and McCrory layer on top is actual usage data from Anthropic’s own platform, the Economic Index, which captures where people are using Claude in professional settings. They cross-reference that against

O\*NET database of around 800 occupations and their component tasks, then adjust for whether the observed use is genuinely automated versus just augmentative, meaning a human is still meaningfully in the loop.

The result is an exposure score by occupation that reflects real-world deployment just capability on paper. And the headline finding is that actual deployment lags theoretical capability by a wide margin across almost every occupational category. The paper is pretty direct about this being the central finding. Business and finance occupations have the highest theoretical exposure, but even there the observed coverage is a fraction of what's possible. Computer and math occupations show 97% theoretical exposure from Eloundou et al., but only 33% observed coverage from actual Claude usage. That 61-point gap is not a rounding error. It represents the friction between what AI can do and what employers and workers are actually deploying it to do, whether that friction comes from regulatory constraints, workforce integration challenges, risk aversion, or just the normal pace of technology diffusion.

For the ten most exposed occupations in the observed framework, computer programmers top the list at 74.5%, followed by customer service representatives at 70.1% and data entry keyers at 67.1%. Medical record specialists appear at 66.7%. Financial and investment analysts show up at 57.2%. The automation happening is real. But it is concentrated in administrative and information-processing work, not in tasks requiring clinical judgment or physical care. That distinction becomes the foundation for a bigger argument about where healthcare labor displacement is most likely to accelerate.

## **Section 2: The Gap Between Capability and Deployment Is the Story**

Before getting to healthcare specifically, it is worth sitting with the deployment data for a minute because it has real implications for how to think about investment timelines and market sizing. The Anthropic data shows that 97% of observed Claude usage falls into tasks that Eloundou et al. rated as theoretically automatable, either fully or with tools. So users are not asking AI to do things it cannot do. The problem

is that a huge portion of what AI could theoretically do is simply not being asked yet.

The authors walk through a specific example that is instructive. Eloundou et al. the task “authorize drug refills and provide prescription information to pharmacists” as fully exposed, meaning an LLM could theoretically handle it twice as fast with human involvement. But that task does not show up in actual Claude usage data. The reasons are obvious to anyone who has spent time in health tech: regulatory expertise, liability concerns, DEA considerations, and the fact that the software layer connecting pharmacy workflows to any AI system simply does not exist at scale yet. The capability is there. The deployment infrastructure, regulatory permission structure, and integration work are not.

This is a pattern that repeats across healthcare specifically, and it is why healthcare shows a consistently smaller red area relative to the blue on Anthropic’s radar chart comparing theoretical versus observed exposure by occupational category. Healthcare practitioners as a category do not appear in the top ten most exposed lists at all, despite the fact that a significant portion of clinical documentation, coding, and administrative work is theoretically highly automatable. The friction is not technical capability. It is the deployment layer.

For investors, this gap is actually where the opportunity lives. The markets that fill this gap first will generate outsized returns. That is true in any industry, but it is especially true in healthcare because the stakes of getting it wrong are higher, the regulatory moats are deeper, and the labor cost savings on the other side are substantially larger than in most other sectors.

## **Section 3: Healthcare Is Not One Labor Market**

One of the conceptual errors that keeps showing up in AI-and-healthcare conversations is treating healthcare as a monolithic labor market. It is not. It is at least a minimum three distinct labor markets with very different cost structures, regulatory environments, and AI exposure profiles: health insurance and managed care, hospital

and health system operations, and ambulatory or outpatient care delivery. Lumping these together produces analysis that is directionally wrong on the most important questions.

Health insurance is mostly an information-processing business with some customer service, sales, clinical review, and regulatory compliance layered on top. The labor that goes into running a major payer is heavily weighted toward knowledge work doing tasks that the Anthropic framework would rate as highly exposed. Utilization management reviewers are essentially doing a form of clinical decision support on paper. Claims adjusters are doing document-heavy pattern matching. Prior authorization coordinators are navigating rule-based workflows that are, in principle, almost entirely automatable. These are exactly the kinds of tasks that show up in observed exposure occupations.

Hospital and health system operations are something entirely different. Labor here includes registered nurses, physicians, surgical techs, imaging techs, physical therapists, housekeeping, dietary, security, transport, and every other function required to run what is essentially a 24-hour manufacturing operation for human bodies. The distribution of labor across clinical versus administrative functions varies by system, but as a general rule, direct care delivery jobs, the ones involving human patient contact, represent the majority of FTEs and the majority of labor expenses. These jobs do not show up anywhere near the top of the Anthropic exposure ranking. Registered nurses have a theoretical exposure score that is moderate at best, and observed exposure is quite low because the tasks that constitute nursing care are information-processing tasks. They are judgment-heavy, physically present, and deeply relational.

The ambulatory care market, meaning physician practices, outpatient clinics, urgent care centers, and the like, sits somewhere between these two extremes. There is significant administrative labor involved, including front-desk scheduling, billing, coding, prior auth coordination, referral management, and patient communication, all of which is highly automatable. But the clinical labor, the actual visits and procedure and care coordination, involves the same human-centered work that makes health systems hard to automate at the clinical layer.

Understanding these distinctions is not academic. It determines where a founder should build, where an investor should allocate, and what the realistic labor displacement curve looks like over a five to ten year horizon.

## **Section 4: Why Insurance Gets All the Press But Misses the Point**

The prior auth story has dominated AI-in-healthcare coverage for the last two years and for understandable reasons. The workflow is rule-based, the documentation burden is absurd, the human cost is enormous, and the political pressure to automate it is growing. CMS has been pushing payers toward faster decision timelines. Physicians hate it. Patients hate it. The administrative overhead is visible and quantifiable. So it is not surprising that this is where most of the AI-in-payer narrative has concentrated.

The same goes for claims processing, fraud detection, appeals management, and member services. These are all legitimate automation targets. The leading payers are already deploying AI in all of these areas, and the efficiency gains are real. UnitedHealth, Elevance, and Cigna have all discussed AI-driven cost improvements in their operational earnings commentary. The academic work on AI exposure in administrative roles supports the thesis that this category is genuinely high on the automation curve.

But here is the problem with anchoring the healthcare AI thesis to the insurance sector: the labor leverage is smaller than people think relative to the total cost structure of the industry. A major commercial payer running at 15-20% administrative costs on a medical loss ratio framework is a different kind of target than a health system running 60% of its operating budget on labor. The insurance sector is also a smaller employer than people realize relative to the hospital sector. The American Hospital Association reports roughly 6.5 million hospital employees in the US. The health insurance industry employs roughly 500,000 to 600,000 people. The labor surface area is just not comparable.

Put another way, if AI eliminates 30% of payer administrative jobs over the next decade, that is a significant workforce event for the people involved, but it is not a macro-scale labor market disruption story for healthcare. If AI eliminates 15% of health system labor costs through a combination of automation, augmentation, and workflow redesign, that is one of the largest efficiency gains in the history of American industry, full stop.

The Anthropic data actually hints at this dynamic even though the paper does not address healthcare specifically in the cost-structure framing. The occupations with the highest observed exposure are concentrated in knowledge work and administrative processing. Customer service representatives, data entry keyers, and medical record specialists are all in the top four. These are exactly the job categories that exist in abundance at health systems, not just payers. The difference is that at a hospital, those administrative workers are a minority of total labor. The nursing techs, the allied health professionals, they are the majority. And the AI story for those workers is a longer and more complicated one.

## **Section 5: Care Delivery Is Where the Leverage Lives**

Here is the core argument: the highest-value target for AI-driven labor reduction in healthcare is not prior auth or claims processing. It is the operational layer of care delivery, specifically the documentation, coordination, and decision-support burden that sits on top of clinical workers and prevents them from operating at the top of their license. The financial case is straightforward once you actually look at health system cost structures.

For a typical large academic medical center or integrated health system, labor represents between 55% and 65% of total operating expense. For community hospitals the number is similar. That labor is disproportionately concentrated in nursing, which is both the largest clinical workforce category and one of the most expensive to recruit, train, and retain in the current market. Agency and travel nursing costs ballooned during and after COVID and have not fully normalized. The American

Nurses Association has documented persistent vacancy rates in the 10-15% range in many health systems. Turnover costs for a single RN are frequently cited in the \$40,000 to \$60,000 range when factoring in recruitment, onboarding, and production ramp. The labor problem in care delivery is not abstract. It is an acute financial challenge that hospital CFOs are managing quarter to quarter.

Now layer in what AI can actually do for clinical workers right now, not in some speculative future state but in deployed products that exist today. Ambient clinical documentation tools, the Nuance DAX category, the Abridge category, can reduce documentation burden on a physician or advanced practice provider by 50% or more per patient encounter. That is not theoretical. Those systems are in production at major health systems with published data behind them. A physician spending two hours per day on documentation who gets that back to one hour is not getting laid off. But a health system deploying that tool across 500 physicians is getting a capacity equivalent of 250 physician-hours per day without adding headcount. That is enormous.

The same logic extends to nursing documentation, which is even more fragmented and time-consuming than physician documentation. Nurses spend an estimated 25-30% of their time on documentation tasks depending on the study and setting. Bringing that down by even 30% through intelligent EHR integration and ambient capture tools does not eliminate nursing jobs. It changes what nurses spend their time doing, which means health systems can serve more patients with existing staff. Given the vacancy rates and the cost pressure, that is more financially valuable than cutting headcount.

Then there is care coordination and transitions of care, which is one of the most labor-intensive and manually driven processes in health system operations. Discharge planning, post-acute placement, follow-up call centers, chronic disease management outreach, all of this involves significant human labor doing tasks that are substantial information-processing in nature. The AI exposure for these roles is higher than bedside clinical care. And the financial stakes are not trivial either, since readmission penalties under CMS programs create direct revenue exposure for every patient who bounces back after a preventable discharge event.

The clinical decision support layer is where things get more speculative but also interesting. Tools like OpenEvidence are already changing how clinicians access evidence at the point of care. If the AI layer can help a nurse practitioner work to full scope of their license more confidently, you get leverage on physician labor costs. If it helps a specialist see more patients per day by reducing cognitive overhead on routine cases, you get capacity expansion without FTE growth. None of this is about replacing clinicians. It is about making existing clinicians more productive, which is a key lever for a sector running at negative operating margins at many institutions, is the most useful financial lever available.

The Anthropic paper does not make this specific argument, but the underlying framework supports it. Occupations with high observed AI exposure are projected by the BLS to grow less through 2034. Customer service representatives, which include a large category of health system call center and patient access workers, have a minus six to minus eight percent projected employment change alongside a 70% observed exposure score. That is not a coincidence. The market is already anticipating this for the nursing workforce and the clinical operations workforce are not showing the same trajectory, because the AI penetration there is still in early innings.

The investment implication is that the AI companies building for care delivery operations, not just for payer workflows, are chasing a much larger labor cost pool. The total hospital labor expense in the US is somewhere in the range of 700 to 900 billion dollars annually depending on how you scope it. The insurance administration labor pool is a fraction of that. If you build a product that saves health systems 5% of labor through productivity improvement, you are addressing a tens-of-billions-dollar market. The payer automation story is real but the addressable market for labor efficiency tools is materially smaller.

## **Section 6: What the Hiring Slowdown Signals for Health Tech Founders**

One of the more interesting findings in the Anthropic paper that tends to get lost in the headline unemployment result is the hiring signal for younger workers. The

finds no measurable increase in unemployment among highly exposed workers of which they interpret as limited evidence of AI-driven labor displacement to date nested inside that finding is a different and arguably more important signal: a 14% drop in the job-entry rate for workers aged 22 to 25 into highly exposed occupations relative to 2022. This finding is just barely statistically significant, but it mirrors Brynjolfsson et al. result showing a 6 to 16% fall in employment for young workers in highly exposed occupations.

What does this mean for healthcare? If the pattern holds, health systems and payers are quietly reducing entry-level hiring into administrative and information-processing roles. They are not laying off existing workers, partly because that creates legal and reputational risk and partly because those workers are genuinely hard to replace as AI tools underperform. But they are not backfilling attrition in the same ways. Front desk staff turns over. Medical records clerks retire. Call center positions open up. And instead of hiring one-for-one, organizations are evaluating whether AI tools can absorb some of that workload.

This is actually how most technology-driven labor transitions happen in practice: not a layoff event. It is an attrition story playing out over five to ten years. For founders building in this space, the implication is that the ROI case for health system buyers is increasingly framed around reduced hiring dependency rather than headcount reduction. That is a softer sell but it is a durable one, because hiring in healthcare is genuinely astronomical given the certification requirements, the training time, and the competitive market for even entry-level clinical support workers.

For health tech investors, the hiring slowdown signal is worth monitoring as a leading indicator. If the 22-25 age cohort is being absorbed into exposed occupations at a lower rate, it suggests employer-side anticipation of AI substitution even where AI tools are not fully deployed yet. That is a real demand signal for the companies building those tools. Health systems are not buying ambient documentation software because they want to give nurses a better experience, even if that is how it is marketed. They are buying it because they are trying to close 150 to 200 basis point operating margin gaps and labor is the biggest lever they have.

# Section 7: How to Think About This as an Investor

The Anthropic framework gives investors a more rigorous lens than most of what has been circulating in the market. The theoretical exposure metrics have been around since the original Eloundou paper in 2023 and they have generated a lot of misplaced conviction about which sectors are going to see rapid AI displacement. The observed exposure data adds a meaningful corrective: markets where the gap between theoretical and observed exposure is largest are markets where the deployment infrastructure, not the AI capability itself, is the binding constraint.

In healthcare, that deployment gap is enormous and it is almost entirely explained by regulatory friction, EHR integration complexity, liability architecture, and the general risk aversion of clinical organizations. None of those are permanent barriers. They are addressable with the right combination of regulatory strategy, technical integration work, and clinical evidence generation. The companies that successfully close the deployment gap in care delivery settings are building on top of the largest labor pool in the American economy.

The BLS data cited in the paper adds a useful grounding mechanism. Jobs with low observed AI exposure are projected to grow less through 2034. That projection embeds assumptions about technology adoption that are probably conservative given the pace of capability improvement. But it also reflects something real: labor market participants, including employers and workers, are already incorporating AI expectations into their behavior. Health systems doing workforce planning today are not projecting the same ratio of FTEs to patient volume that they would have projected in 2019. That behavioral shift is a demand signal for the tools that actually deliver the promised productivity.

The most defensible investment thesis in this space right now is the one that focuses on care delivery productivity rather than payer automation. Not because payer automation is a bad business, but because the labor cost pool is smaller, the market is more consolidated into a handful of large incumbents with internal AI developmen

capacity, and the regulatory pathway for AI in claims and prior auth is being shaped by CMS in ways that commoditize the workflow faster than the moats can be built.

Care delivery is messier, more fragmented, and harder to sell into. But it is also a much larger and more durable opportunity. The health systems that figure out how to run with 10% less labor through AI-driven productivity, not through layoffs but through attrition management, scope expansion, and documentation automation have a structural cost advantage over competitors. The vendors that enable those outcomes will have sticky contracts, real clinical evidence, and the kind of integration depth that is genuinely hard to replicate.

The Anthropic paper ends with a note of intellectual honesty that is worth quoting in spirit if not in letter: this is early evidence, the effects are small and in some cases statistically marginal, and the framework is most useful before the effects become obvious. That is actually a useful framing for the investment opportunity too. The healthcare labor displacement story is not yet visible in aggregate unemployment. The hiring slowdown signal for young workers is just barely there. The BLS projections are a gentle negative slope, not a cliff. Which means there is still time to build and invest ahead of the curve rather than chasing something that has already played out. That window will not stay open indefinitely.



3 Likes • 2 Restacks

← Previous

Next

## Discussion about this post

Comments

Restacks



Write a comment...

---

