

Minutes Are Not Money

Why AI looks cost-effective in the model but elusive on the ledger, and a practical framework for turning AI output into realized return in health systems

Mona Veiseh
Founder & CEO, KinRidge Advisors

September 29, 2026 · Version 1.0
Published by KinRidge Advisors · kinridgeadvisors.com

AI does not create value. Health systems create value when they turn AI output into decisions, action, and measurable results. Most AI business cases underestimate that conversion — and overstate the return.

Executive summary

AI adoption in health care is no longer the question. In 2026, 81 percent of physicians surveyed by the American Medical Association reported using AI in their practices, more than double the 2023 rate.[1] Seventy-one percent of nonfederal acute care hospitals reported predictive AI integrated with their EHR in 2024.[2] Yet chief financial officers still struggle to find the return. As one 2026 analysis put it, “AI investments are rising, but proven financial impact is elusive.”[3]

This paper argues that the gap is structural and predictable. Every AI output must pass through a chain before it becomes value: the signal must hold up locally, someone must act on it, that person must have the authority to act, the released capacity must be redeployed, and the result must be measured. Value leaks at every link, and the losses multiply. A tool can perform exactly as advertised and still deliver a fraction of its modeled return.

Five implications follow for health-system leaders:

- Evaluate the workflow, not the tool. The unit of analysis is a decision or task: who receives the output, what they can now do, who approves it, and what measurable result follows.
- Treat freed time as a claim, not a return. Minutes released become money only when the organization decides, in advance, what those minutes will buy and then shows that it happened.
- Fix the weakest link before buying a better model. A more accurate model adds little if most of its outputs are ignored or cannot be acted on.
- Start where the ledger is largest. Labor accounts for about 60 percent of hospital expenses.[12] Workforce deployment is where conversion matters most and where the chain from forecast to paid hours can be traced most clearly.
- Choose paths deliberately and protect what matters. Screen every efficiency path for size, conversion, evidence, readiness, and values fit. Automate tasks, never relationships, and write patient and employee guardrails into the decision rules so they can veto an attractive efficiency.

We propose a single portfolio measure for executives, the conversion rate: realized, finance-validated value divided by the value modeled in the business case. It shows where AI value is being lost and makes the conversation between clinical, operational, and finance leaders concrete.

1. The paradox: cost-effective on paper, elusive on the ledger

The published economic literature on clinical AI is strikingly positive. A 2025 systematic review of 52 economic evaluations from 15 countries found that 51 of them, or 98 percent, concluded that AI-based strategies were cost-effective, cost-beneficial, or cost-saving.[4] A separate 2025 review of 19 studies in npj Digital Medicine reached a similar conclusion, finding clinical AI “highly cost-effective or even cost-saving across many applications.”[5]

Look more closely, however, and the evidence rests on models rather than ledgers. Most of the 52 evaluations used decision-analytic models, whose results depend on assumptions that the authors warn “may not hold in practice.” AI-related costs were inconsistently reported.[4] In the 19-study review, only four studies included indirect costs.[5] When researchers looked instead at prospective, real-world deployments of deep-learning systems, the picture changed. A 2026 systematic review found only 20 such studies. Just one evaluated implementation cost, and none evaluated sustainability.[6]

Figure 1. The evidence for AI’s economic value is strongest where it is least tested. Sources: Wu et al., systematic review of 52 economic evaluations[4]; npj Digital Medicine systematic review of 20 prospective real-world deep-learning implementations.[6] The two reviews cover different study sets and are shown together to contrast modeled and observed evidence.

The broader economy shows the same pattern. MIT’s NANDA initiative reported in 2025 that about 95 percent of companies in its dataset saw little or no measurable profit-and-loss impact from generative AI pilots. It attributed the gap less to model quality than to a “learning gap”: tools that “don’t learn from or adapt to workflows.”[7] In health care, Oliver Wyman’s 2026 assessment concluded that “the technology often works as advertised but does not deliver value when placed on top of fragmented operating models, processes, and policies.”[3]

The model says yes. The ledger says maybe. The difference is not a flaw in the model. It is everything that happens between the model’s output and the organization’s income statement.

2. Minutes are not money

Most AI business cases in health care promise time: fewer minutes documenting, faster reads, less manual scheduling. Time is a real benefit. But freed time is capacity released, not value realized. It turns into financial return only when the organization converts it into something it can measure: more visits, fewer paid hours, less agency spend, lower turnover, or better quality.

Ambient documentation shows the distinction clearly.

The time savings are real but modest per clinician. In a randomized trial of 238 UCLA Health physicians, one ambient scribe reduced time-in-note by 9.5 percent relative to usual care; another showed no significant reduction. In the second month, the better-performing tool cut average time-in-note from 4 minutes 30 seconds to 3 minutes 49 seconds.[8] Across a full clinic day, that matters. But it does not, by itself, change what a physician is paid or how many patients are scheduled.

Large totals can hide small units. The Permanente Medical Group reported that ambient AI saved nearly 16,000 hours of documentation across 7,260 physicians and about 2.6 million encounters over 15 months.[9] Averaged across all participating physicians, that is roughly two hours each over the period. It is a meaningful gain for burnout and experience. Whether it is a financial return depends entirely on what happened next.

The financial case remains unproven at the system level. The Peterson Health Technology Institute found that ambient scribes likely reduce burnout and cognitive load, but that “the financial impact is unclear.”[10]

Oliver Wyman stated the conversion problem directly: “If physicians continue seeing the same number of patients supported by the same staffing model, and documentation gains are simply absorbed as marginal time savings, the bottom-line impact remains limited.”[3] Mayo Clinic’s chief AI implementation officer made a similar point, noting that on the care delivery side “it’s really hard to apply line of sight to what we might think of as ROI.”[11]

None of this argues against documentation AI. Reducing burnout may be reason enough to invest, and retention has real financial value. It argues for honesty in the business case. If the value is clinician experience, say so and measure it. If the value is financial, name the conversion mechanism before launch.

3. The conversion chain

Every AI use case, clinical or operational, passes through the same five links between output and value. Each link is a point where value can leak. The table below is a KinRidge framework; the evidence column shows why each link matters.

Table 1. The five links of the AI conversion chain. KinRidge framework.

Link	The question	How value leaks	Evidence
1. Signal	Does the output hold up in our data, population, and workflow?	Performance falls short of what the developer reported	A widely deployed sepsis model reached an AUC of 0.63 on external validation, well below developer-reported performance.[13]
2. Decision	Does the person who receives it act on it?	Alerts are ignored, overridden, or arrive at the wrong moment	A meta-analysis found physicians overrode 90 percent of drug-interaction alerts.[14]
3. Action	Does that person have the authority and options to act?	Policy, budget, contracts, or staffing rules block the response	Oliver Wyman: financial benefit “requires changes in clinical pathways and workforce roles.”[3]
4. Capacity	Is the released time or resource redeployed?	Savings are absorbed as slack	Documentation gains “simply absorbed as marginal time savings” limit impact.[3]
5. Measurement	Can finance see and attribute the result?	No baseline, no comparison, costs omitted	Only 1 of 20 real-world deployments evaluated implementation cost.[6]

The multiplication problem

The losses at each link do not add up; they multiply. If a use case keeps 85 percent of its value at the signal link, 75 percent at the decision link, 80 percent at the action link, 60 percent at the capacity link, and 80 percent at the measurement link, the organization realizes about 24 percent of the modeled value. No single link looks like a failure. Together, they turn a strong business case into a disappointing one.

Figure 2. Illustrative: modest losses at each link compound into a large conversion gap. Illustrative values chosen to demonstrate compounding, not empirical estimates. Actual retention at each link varies by use case and must be measured locally.

The implication for investment is counterintuitive. Improving the weakest link usually yields more than improving the model. In the illustration, raising model performance from 85 to 95 percent lifts realized value from about 24 to 27. Raising capacity redeployment from 60 to 90 percent lifts it to about 37. Health systems often spend their evaluation energy on the first link, model accuracy, because it is the most technical and the easiest to benchmark. The larger losses usually sit further down the chain, in management decisions that no vendor controls.

4. Start with a decision, not a tool

The unit of analysis should be a complete workflow: who receives the AI output, what action becomes possible, who approves that action, and which measurable result follows. A model that predicts demand without changing a staffing decision has produced information, not labor savings. A documentation assistant that drafts a note but increases correction time may not reduce burden at all. The relevant comparison is always the current process, including its workarounds and hidden costs.

For each opportunity, leaders should identify the conversion point: the specific place where released capacity or better information turns into a result the organization can bank.

Table 2. AI opportunities and their conversion points. Adapted from the KinRidge working draft.

Opportunity	Plausible value	Conversion point	What must be verified
Clinical documentation	Less writing and after-hours work	Added visit capacity, reduced scribe spend, or retention	Total time including review; note quality; clinician experience
Diagnostic triage and imaging	Better prioritization; released specialist capacity	Changed reading protocol or reduced backlog	Patient outcomes; false positives; work redistributed
Workforce deployment	Better match between demand, skills, availability, and location	Change in paid hours, agency, and overtime at comparable volume	Actual labor spend; fill rate; safety; staff experience
Access and patient flow	Earlier action on discharge, referrals, and appointments	Freed beds or slots actually filled	Throughput, waits, cancellations, downstream bottlenecks
Revenue operations	Fewer corrections, delays, and preventable denials	Net collections and reduced rework staffing	Net collections after licensing and integration costs

5. Choosing the path: what to automate, what to augment, what to keep human

Once leaders see AI as a conversion problem, a harder question follows. There are endless paths to efficiency, and every vendor offers a new one. How does a health system decide which paths are right for it, and how does it make those choices without losing sight of its values, its patients, and its people? We propose answering in three steps: choose the right paths, decide the right role for AI within each one, and hold every choice to explicit values guardrails.

Step one: choose the paths worth taking

Not every efficiency is worth pursuing. A path deserves investment when it passes five screens, applied in order. The first two screens do most of the work: a path that is small or has no clear conversion point should not advance, however impressive the technology.

Table 3. Five screens for choosing efficiency paths. KinRidge framework.

Screen	The question	What a strong answer looks like
1. Size	How large is the cost, risk, or burden this addresses?	A major ledger line, such as labor, or a burden that drives turnover or harm
2. Conversion	Is there a clear conversion point and a named owner?	Leaders can state in advance what released capacity will buy
3. Evidence	Has this configuration worked in comparable settings?	Prospective or real-world evidence, not only modeled estimates
4. Readiness	Do we have the data, workflow clarity, and change capacity?	Data is reliable; the people who do the work are ready to help redesign it
5. Values fit	Does it strengthen what we promise patients and staff?	It returns time to care, preserves trust, and treats staff fairly

Applied honestly, these screens tend to point in the same direction: toward high-volume administrative and operational work with a clear owner, and toward workforce deployment, where the ledger is largest. They also tend to reject paths that save little, cannot be measured, or ask patients and staff to absorb the cost of someone else's efficiency.

Step two: decide the role of AI in each task

Within any chosen path, the decision is rarely all or nothing. It is made task by task, and two questions settle most cases: How serious and reversible is an error? and How much does the task depend on human judgment, relationship, or trust? The answers place each task in one of four roles.

Table 4. Deciding the role of AI in a task. KinRidge framework; placement of specific tasks should be confirmed locally.

	Low dependence on human relationship or judgment	High dependence on human relationship or judgment
Lower consequence, easily reversed	Automate, with monitoring. Eligibility checks, prior authorization packet assembly, routine scheduling logistics, claim status follow-up.	Assist, and keep the human in front. Draft patient messages for staff review; ambient documentation that returns attention to the patient.
Higher consequence, hard to reverse	Augment, with human verification. Imaging triage, deterioration alerts, staffing forecasts that a manager approves.	Keep human. Delivering a diagnosis, consent and goals-of-care conversations, decisions about an individual employee's role, pay, or discipline.

Two principles sit behind the grid. First, automate tasks, never relationships. Patients do not experience a health system as a set of workflows; they experience it through people. Second, favor augmentation where judgment matters. Economist Erik Brynjolfsson has argued that AI which complements people creates more value than AI which merely imitates them, and that “when AI is focused on augmenting humans rather than mimicking them, then humans retain the power to insist on a share of the value created.”[17] The MASAI mammography trial is a good example: AI decided how many radiologists should read an image, but radiologists still made the call.[16]

Augmentation carries its own risk, which leaders should manage rather than ignore. In an observational study nested within a colonoscopy trial in Poland, endoscopists’ adenoma detection rate in standard, non-AI procedures fell from 28 percent to 22 percent after they became used to routine AI support. The authors present this as a hypothesis requiring confirmation, not a settled finding.[18] Physicians share the concern: 88 percent in the AMA’s 2026 survey said they worried about potential skill loss.[1] Where AI supports clinical judgment, organizations should monitor human performance as well as model performance, and preserve regular practice without the tool.

Step three: hold every choice to values guardrails

Values are easiest to protect when they are written into the decision rules rather than recited afterward. We recommend four guardrails, each with a measure that can veto an otherwise attractive efficiency.

- **Patients:** be open, and give them a say. Patients want to know. In a Pew Research Center survey published in August 2026, 72 percent of U.S. adults said it was extremely or very important that a provider tell them if AI is used in their care. Fifty-three percent said they had little or no say over whether it is used, and 46 percent were not sure whether it had been used in their care at all.[19] An earlier Pew survey found that 60 percent of Americans would be uncomfortable if their own provider relied on AI in their care.[20] Health systems should disclose AI use in plain language, explain what stays in

human hands, and track patient experience as a guardrail for every patient-facing use case.

- **Employees:** design with them, not around them. The people who do the work know where the real toil is and where the risks hide. In McKinsey's 2026 nursing survey, the leading conditions nurses named for wider AI adoption were strong data security and privacy protections and evidence that AI improves quality and patient safety, each cited by 54 percent.[21] In American Nurse's 2025–2026 trends survey, the most common concern about AI was lack of human connection.[22] Frontline staff should help choose the tasks to automate, test the tools, and have a respected way to override them.
- **Released capacity:** make a pledge before launch. The conversion point is also a values decision. Before an initiative starts, leaders should state how freed time will be used: how much goes to patient time, how much to staff well-being or workload relief, and how much to added capacity or cost reduction. Publishing that pledge to the teams involved builds trust, and measuring against it keeps the organization honest.
- **Fairness:** check who carries the cost. Efficiency that shifts burden to night shifts, junior staff, or particular patient groups is not efficiency; it is relocation. Scheduling tools, for example, should be judged on perceived fairness as well as cost, as the Korean AI scheduling study did.[23]

Five questions before scaling

A values test for any AI efficiency. Would we be comfortable explaining it to the patient in the room? Does it give time back to care or to the people who provide it? Did the people who do the work help design it? Can a human override it, and is that override respected? Do we know who benefits from the capacity it releases?

6. What the strongest evidence has in common

The best results in health care AI share a feature that is easy to overlook: the workflow itself changed, so conversion was designed in rather than hoped for.

Imaging: when AI changes the rule

The Swedish MASAI trial randomized about 106,000 women to AI-supported mammography screening or standard double reading. AI was used to triage examinations to single or double reading based on risk scores and to mark suspicious findings for radiologists. The AI-supported pathway reduced screen-reading workload by 44 percent, detected 29 percent more cancers, and produced 12 percent fewer interval cancers, 82 compared with 93.[15, 16]

What made MASAI convert is that it changed a rule, not just the information available. Low-risk examinations moved from two readers to one. The released radiologist capacity was built into the protocol instead of left for individuals to absorb. It is strong evidence for

that clinical configuration, not a general estimate of AI performance in other specialties, but it is the clearest model of what designed conversion looks like.

Documentation: benefit without a conversion mechanism

Ambient documentation is the opposite case. The benefit to clinicians appears real, adoption has been rapid, and products vary in performance.[8, 10] But the conversion point is usually undefined. Time is freed, and the operating model stays the same. Health systems should measure full encounter documentation time, accuracy, clinician correction effort, burnout, and patient experience before assigning financial value to time saved. Where financial return is expected, they should decide in advance whether it will come from visit capacity, reduced scribe spending, or retention.

Decision support: signal without a decision

Clinical decision support illustrates the second link. When an external validation found that a widely deployed sepsis model alerted on 18 percent of hospitalized patients while missing 67 percent of patients with sepsis, the problem was not only accuracy. It was the volume of signals that clinicians had to triage.[13] Across drug-interaction alerting more broadly, a meta-analysis found that physicians overrode 90 percent of alerts.[14] An output that is overridden nine times in ten has almost no chance to convert, however accurate it becomes.

7. Workforce: the largest line on the ledger

If conversion is the problem, workforce deployment is where solving it matters most. Labor remains the largest single category of hospital spending. The American Hospital Association reports that about 60 percent of hospital expenses in 2025 went to the workforce, that workforce costs rose 5.6 percent, and that total hospital expenses grew 7.5 percent, more than twice the rate of growth in hospital prices.[12] A small improvement in how labor is deployed is worth more than a large improvement in most other categories.

Workforce is also where the conversion chain is easiest to trace. A demand forecast leads to a staffing recommendation, which leads to a manager's decision, which leads to paid hours, which leads to cost and service outcomes. Every link can be observed in existing time-and-attendance, scheduling, and finance data. The AHA's 2026 workforce scan makes the same connection, noting that "AI's value relies on thoughtful implementation" and that aligning AI with redesigned workflows is "the strongest link to improved financial performance." It highlights float pools, internal agencies, and predictive scheduling as ways to reduce overtime, turnover, and incentive spending.[24]

Published evidence for AI scheduling is encouraging but context-specific. A study at a Korean tertiary hospital found that schedules generated by an AI system reduced difficult shift patterns, such as night-off-evening rotations, gave nurses more consecutive days off,

and were associated with job satisfaction.[23] That is useful implementation evidence. It does not establish a universal labor-cost return for U.S. health systems or for managed service provider and vendor management system environments.

The KinRidge hypothesis

Our working hypothesis is that the constraint in workforce AI is rarely the forecasting model. It is the data and decision layer beneath it: verified skills, credential status, availability, preferences, and scheduling constraints, brought together so that a manager can choose among internal, float, overtime, and contingent options with confidence. A reliable workforce data layer would strengthen every link in the chain for the scheduling, staffing, and care operations applications that health systems already own.

Testing this hypothesis means tracing the full chain from forecast to recommendation to manager action to paid hours and service outcomes. A pilot should record every override and its reason, so the team learns whether the constraint is the model, the data, the policy, or local trust. Overrides are the most valuable data a workforce AI pilot produces, because they show exactly which link is leaking.

8. What counts as realized ROI

A credible business case separates financial return from clinical and workforce value, and it states the conversion mechanism for each. We recommend three definitions, agreed with finance before launch:

- Net financial benefit equals verified savings and incremental contribution attributable to the intervention, minus licensing, integration, training, monitoring, and change-management costs.
- ROI equals net financial benefit divided by total investment over a defined period.
- Conversion rate equals realized, finance-validated value divided by the value modeled in the original business case. It is the measure that shows whether a portfolio is learning to turn AI output into return.

Freed clinician minutes count as financial return only when the system can show what became possible because of them, such as measurable additional capacity or avoided paid work. Safety, access, and staff experience belong in the same decision record, even when they cannot responsibly be converted to dollars. A tool that improves retention and burnout with no direct savings can still earn its place, but the decision should be made on those terms, not on an inflated financial case.[4, 5]

An illustrative staffing pilot

Consider a system that forecasts unit demand 48 hours ahead and matches that forecast with eligible internal capacity. The primary financial outcome is the change in total labor cost at comparable patient volume and acuity, measured against a credible control or adjusted baseline. Guardrails include unfilled shifts, overtime, agency hours, nurse-reported fairness, and safety events. The pilot tracks each link: forecast accuracy, the share of recommendations accepted, the share that could be acted on under current policy, the paid hours that actually changed, and the savings finance can confirm.

If the system merely shifts cost from agency contracts to unsustainable overtime, it has not demonstrated operating value. If it lowers agency spend while raising burnout or unfilled shifts, the guardrails should veto it. And if recommendations are accurate but rarely accepted, the next investment belongs in decision design and trust, not in a better model.

9. A practical adoption sequence

The sequence below builds conversion into each stage rather than checking for it at the end.

Table 5. Adoption sequence with conversion gates. KinRidge operating design, expanded from the working draft.

Stage	What happens	Conversion question to answer
1. Select	Identify a costly or consequential decision and establish its baseline	Where is the conversion point, and who owns it?
2. Map	Map data sources, authority, exceptions, and the current workflow with the people who do the work	Can the recipient act on the output under current policy?
3. Validate	Assess performance and failure modes locally, across sites and populations	Does the signal hold up here?
4. Pilot	Run a limited implementation with an explicit comparison and documented human oversight; log overrides	Which link is leaking, and why?
5. Gate	Review outcomes and costs with finance before scaling	Is the conversion rate high enough to justify scale?
6. Sustain	Monitor whether benefits persist as data, workflows, and versions change	Is the value still being realized now?

Published guidance supports this emphasis on observing actual use. The DECIDE-AI reporting guideline focuses on early-stage clinical evaluation of AI-driven decision support in real settings, and the 2026 review of prospective deep-learning deployments found

real-world implementation evidence still scarce.[6, 26] Accreditation is moving in the same direction. The Joint Commission's 2026 review guide for its Responsible Use of AI in Health Care certification asks organizations to show pre-deployment testing, lifecycle monitoring, and how they "validate vendor-reported performance."[25]

The final stage connects this paper to KinRidge's companion white paper, The Half-Life of a Best Practice. Conversion is not a one-time achievement. A use case that converts well today can leak again after a model update, an EHR upgrade, or a change in staffing policy. The review cadence for each AI use case should match how fast the organization changes the systems and workflows it depends on.

10. Implications for health-system leaders

- Fund the whole change, not just the software. Budget for workflow redesign, data work, training, monitoring, and the management time needed to convert released capacity. Bain and KLAS report that executives now favor AI with "tangible results and short time-to-value windows."[27] Those results come from the change around the tool as much as the tool itself.
- Name a conversion owner for every initiative. Each AI initiative needs an operational sponsor, a clinician or workforce owner, an evaluation plan, and one person accountable for the conversion point. That person decides, before launch, what released capacity will be used for.
- Manage a portfolio, not a collection. The highest-value portfolio may pair modest tools that reliably remove work with more sophisticated tools whose value is still unproven. Report the conversion rate across the portfolio, and make explicit decisions about when to expand, revise, or stop.
- Govern the human-AI workflow, not just the model. Governance should examine the performance of the combined workflow and its risks throughout use, as the NIST AI Risk Management Framework recommends.[28]

Five questions for the executive team

- For our three largest AI investments, what is the conversion point, and who owns it?
- What share of modeled value have we actually realized, and could finance confirm it?
- Which link in the chain is leaking most, and are we investing there or in the model?
- Where have we released capacity that was never redeployed?
- Are we applying AI to our largest cost line, the workforce, with the same rigor as to documentation and revenue cycle?

11. How KinRidge can help

A practical engagement begins with a conversion diagnostic of the current AI portfolio: for each use case, the modeled value, the conversion point, the owner, the evidence at each link, and the realized value finance can confirm. KinRidge and the health system then select one or two high-value decisions, often in workforce deployment, and redesign the workflow, data layer, and decision rights around them.

The teams run those initiatives through the full adoption sequence, logging overrides, gating scale on the conversion rate, and training internal leaders to repeat the method. The work is complete when the health system can independently identify a valuable decision, redesign the work around it, and show a sustained improvement that clinicians and finance leaders both recognize.

Closing

The strategic question is no longer whether a health system can adopt AI. It is whether it can turn what AI produces into value it can prove. That is the standard by which AI should earn its place in the operating model.

References

Sources are linked for verification. This paper distinguishes published evidence from KinRidge recommendations and illustrative operating models. Frameworks labeled as KinRidge designs are proposals, not published standards.

[1] American Medical Association. AMA: AI usage among doctors doubles as confidence in technology grows. March 12, 2026. <https://www.ama-assn.org/press-center/ama-press-releases/ama-ai-usage-among-doctors-doubles-confidence-technology-grows>

[2] Office of the National Coordinator for Health Information Technology. Hospital Trends in the Use, Evaluation, and Governance of Predictive AI, 2023–2024. Data Brief 80. <https://healthit.gov/data/data-briefs/hospital-trends-use-evaluation-and-governance-predictive-ai-2023-2024/>

[3] Oliver Wyman. Turning health system AI spending into bottom-line results. March 2026. <https://www.oliverwyman.com/our-expertise/perspectives/health/2026/mar/turning-health-system-ai-spending-into-bottom-line-results.html>

[4] Wu WT et al. Economic evaluation of AI-assisted technologies in healthcare: A systematic review. Journal of Food and Drug Analysis. 2025. <https://pmc.ncbi.nlm.nih.gov/articles/PMC12795340/>

[5] El Arab RA, Al Moosa OA. Systematic review of cost effectiveness and budget impact of artificial intelligence in healthcare. npj Digital Medicine. 2025. <https://www.nature.com/articles/s41746-025-01722-y>

[6] Prospective real-world implementation of deep learning systems in healthcare: A systematic review guided by implementation science. npj Digital Medicine. 2026. <https://www.nature.com/articles/s41746-026-02358-2>

[7] Fortune. MIT report: 95% of generative AI pilots at companies are failing (on MIT NANDA, The GenAI Divide: State of AI in Business 2025). August 18, 2025. <https://fortune.com/2025/08/18/mit-report-95-percent-generative-ai-pilots-at-companies-failing-cfo/>

[8] Lukac PJ et al. Ambient AI Scribes in Clinical Practice: A Randomized Trial. NEJM AI. 2025. <https://pmc.ncbi.nlm.nih.gov/articles/PMC12768499/>

[9] Kaiser Permanente Division of Research. AI-assisted notetaking gains steady support from Kaiser Permanente physicians. 2025. <https://divisionofresearch.kaiserpermanente.org/ai-assisted-notetaking-gains-steady-support-from-kaiser-permanente-physicians/>

- [10] Peterson Health Technology Institute. Leading Health Systems: AI-Powered Scribes Alleviate Clinician Burnout; Financial Impact Unclear. March 25, 2025. <https://www.phti.org/announcement/ai-scribes-reduce-clinician-burnout/>
- [11] Healthcare Dive. Health systems are racing to adopt AI. But can they prove ROI? October 27, 2025. <https://www.healthcaredive.com/news/healthcare-ai-adoption-return-investment-roi-hlth/803653/>
- [12] American Hospital Association. New AHA report: Hospitals face increased challenges and financial pressures as they care for patients (2025 Cost of Caring). March 11, 2026. <https://www.aha.org/press-releases/2026-03-11-new-aha-report-hospitals-face-increased-challenges-and-financial-pressure-they-care-patients>
- [13] Wong A et al. External Validation of a Widely Implemented Proprietary Sepsis Prediction Model in Hospitalized Patients. JAMA Internal Medicine. 2021. <https://pmc.ncbi.nlm.nih.gov/articles/PMC8218233/>
- [14] Felisberto M et al. Override rate of drug-drug interaction alerts in clinical decision support systems: A brief systematic review and meta-analysis. Health Informatics Journal. 2024. <https://journals.sagepub.com/doi/10.1177/14604582241263242>
- [15] Gommers J et al. Interval cancer, sensitivity, and specificity comparing AI-supported mammography screening with standard double reading in the MASAI study. The Lancet. 2026. <https://pubmed.ncbi.nlm.nih.gov/41620232/>
- [16] Lund University. AI support in breast cancer screening: fewer missed cancer cases. 2026. <https://www.lunduniversity.lu.se/article/ai-support-breast-cancer-screening-fewer-missed-cancer-cases>
- [17] Brynjolfsson E. The Turing Trap: The Promise & Peril of Human-Like Artificial Intelligence. Stanford Digital Economy Lab. 2022. <https://digitaleconomy.stanford.edu/news/the-turing-trap-the-promise-peril-of-human-like-artificial-intelligence/>
- [18] University of Oslo. Continuous exposure to AI might reduce the skills of endoscopists (on Budzyń K, Romańczyk M et al., Lancet Gastroenterology & Hepatology, 2025). August 2025. <https://www.med.uio.no/helsam/english/research/groups/clinical-effectiveness/news/2025/ai-colonoscopy.html>
- [19] Pew Research Center. Americans want transparency when AI is used in their health care. August 25, 2026. <https://www.pewresearch.org/short-reads/2026/08/25/americans-want-transparency-when-ai-is-used-in-their-healthcare/>
- [20] Pew Research Center. 60% of Americans would be uncomfortable with provider relying on AI in their own health care. February 22, 2023. <https://www.pewresearch.org/science/2023/02/22/60-of-americans-would-be-uncomfortable-with-provider-relying-on-ai-in-their-own-health-care/>
- [21] McKinsey & Company. Transparency and trust for AI in nursing (2026 McKinsey Nursing AI Insights Survey). July 2026. <https://www.mckinsey.com/featured-insights/charts/transparency-and-trust-for-ai-in-nursing>
- [22] American Nurse. 2025–2026 Trends Survey. April 2026. <https://www.myamericannurse.com/2025-2026-trends-survey/>
- [23] Inha University Hospital investigators. Shift nurses' work quality and job satisfaction after implementing the IH-NASS AI scheduling system. BMC Nursing. 2025. <https://pubmed.ncbi.nlm.nih.gov/40597282/>
- [24] American Hospital Association. 2026 Health Care Workforce Scan: Executive Summary. <https://www.aha.org/system/files/media/file/2025/12/2026-Health-Care-Workforce-Scan-Executive-Summary.pdf>
- [25] The Joint Commission. Responsible Use of Artificial Intelligence in Health Care Certification: Review Process Guide. June 1, 2026. <https://digitalassets.jointcommission.org/api/public/content/7bd26620a3e54c0ead2138afa37fc167?v=5f349b78>
- [26] Vasey B et al. Reporting guideline for the early-stage clinical evaluation of decision support systems driven by artificial intelligence: DECIDE-AI. Nature Medicine. 2022. <https://www.nature.com/articles/s41591-022-01772-9>
- [27] Bain & Company and KLAS Research. Healthcare AI investment focusing on hard-dollar returns and clinical workflows amid shift from pilots to widespread production. 2025. <https://www.bain.com/about/media-center/press-releases/20252/healthcare-ai-investment-focusing-on-hard-dollar-returns-and-clinical-workflows-amid-shift-from-pilots-to-wide-spread-productionbain--company-and-klas-research/>
- [28] NIST. AI Risk Management Framework Playbook. <https://www.nist.gov/itl/ai-risk-management-framework/nist-ai-rmf-playbook>

To discuss how this applies to your organization, contact KinRidge Advisors at kinridgeadvisors.com/#contact.

© 2026 KinRidge Advisors. Suggested citation: Mona Veiseth. "Minutes Are Not Money." KinRidge Advisors White Paper, September 29, 2026.