

The Half-Life of a Best Practice

Why a practice is only “best” if it demonstrates value—and keeps proving it at the speed your organization changes

Mona Veisoh

Founder & CEO, KinRidge Advisors

October 19, 2026 · Version 1.0

Published by KinRidge Advisors · kinridgeadvisors.com

A practice is “best” only while it delivers locally demonstrated value, meets applicable safety and compliance requirements, and remains supported by evidence after material changes. Financial claims require finance validation; clinical and safety claims require appropriate outcome evidence. Its review cadence must match the organization’s clockspeed.

Executive summary

Health systems have long used “best practice” as a finish line: a pilot succeeds, the approach gets the label, and the label stays put. That habit was already risky. Well-designed trials routinely overturn established care, and the commonly cited path from research evidence to routine practice runs to about 17 years.^[1, 2, 3] Artificial intelligence now makes the habit untenable. Epic reports that its software release cadence has shortened from about 18 months to three months; an ARK analysis reported a median 11-day interval between frontier model releases in early 2026.^[4, 5]

This paper argues for a stricter and more useful definition. An organization may call something a best practice only when two conditions hold at the same time:

- It is proven. The practice has met locally agreed value and outcome criteria against a baseline established before launch. Finance validates financial claims; appropriate clinical leaders validate clinical and safety claims. Mandatory compliance practices demonstrate that requirements are met, with the effectiveness and cost of implementation made explicit.
- It is current. The evidence has been re-examined since the last material change to the technology, workflow, people, or patients that the practice depends on.

Neither condition is enough alone. A proven practice that has not been reviewed after a material change is a legacy practice running on old evidence. A reviewed practice that has not yet demonstrated its intended value remains an experiment. The practical consequence is the paper’s central operating rule: the review cycle for a best practice must match the clockspeed of the organization that runs it. A system that changes its technology quarterly but reviews results annually can spend much of the year relying on evidence that predates material changes to its tools or workflows.

We propose that health systems build a project management office, or PMO, that acts as the organization’s clock keeper. It would own the register of practices, record when each was last proven, track the “evidence debt” that builds with every change, and trigger review before that debt turns into harm or wasted spend.

1. The end of the permanent best practice

Medicine has good evidence that its own standards expire. A review of a decade of original articles in the New England Journal of Medicine found 363 studies that tested an established standard of care. Of those, 146, or 40.2 percent, reversed the practice, while 138 reaffirmed it.^[2] A later review of more than 3,000 randomized trials in JAMA, The Lancet, and NEJM identified 396 medical reversals.^[3] These were not fringe practices. They were accepted approaches that a stronger test later showed to be ineffective or inferior.

Even the syntheses that underpin best practices age quickly. A survival analysis of high-quality systematic reviews found that the median time before new evidence signaled a need to update was 5.5 years; 15 percent showed such a signal within one year and 23 percent within two.^[6]

The system is also slow to act on new evidence. The frequently repeated estimate that research takes 17 years to reach clinical practice is itself contested; its authors note that the figure “hides complexities” because different studies measure different stages.^[1] The estimate illustrates a translation challenge, but it is not a universal measure of review cadence. Meanwhile, the volume of knowledge keeps growing. One widely cited estimate projected that the doubling time of medical knowledge would fall from 3.5 years in 2010 to about 73 days by 2020.^[7] That is a projection, not an observed doubling rate or a current measurement.

Put these facts together and a pattern emerges. Best practices get declared quickly and retired slowly, if at all. The label outlives the evidence. AI does not create this problem. It speeds it up to the point where leaders can no longer ignore it.

The question is no longer “Is this a best practice?” It is “When was this last proven, and what has changed since?”

2. Three clocks, running at different speeds

MIT's Charles Fine introduced the idea of “clockspeed” to describe how quickly an industry's products, processes, and organizations change. He observed that different industries, and even different capabilities inside one firm, “change, grow, or obsolesce at different rates.”^[8] Health systems now run on at least three clocks, and the gap between them explains why best practices decay.

The technology clock

This is the pace at which the tools change. At its 2026 users group meeting, Epic described how its release cadence has shifted: “Back in the day, we released new versions about every 18 months,” said Seth Howard, Epic’s executive vice president of research and development. “Now we release new versions every three months. ... Today’s three-month release contains more code than the typical 18-month release of the past.”^[4] Frontier AI moves faster still. Analysis by ARK Investment Management, based on Artificial Analysis data, put the median gap between frontier model releases from leading U.S. labs at 37.5 days in 2023 and 11 days across five labs through April 28, 2026, as reported the next day.^[5]

The organization clock

This is the pace at which a particular health system chooses to absorb change, and it varies widely. Health systems upgrade their EHRs anywhere from annually to quarterly. Memorial Hermann upgrades quarterly and stays no more than one release behind. Northwell and Cleveland Clinic upgrade twice a year. Nebraska Methodist Health System updates its Oracle Health EHR annually; its leader explained, “there has to be a reason why we’re going to introduce change into our environment.”^[9] Buying is speeding up as well: Menlo Ventures reports that health-system purchasing cycles for AI average 6.6 months, compared with 8.0 months for traditional IT.^[10]

The evidence clock

This is the pace at which the organization proves that a practice still works. It can become the slowest clock when business cases are not revisited and benefit reviews remain annual or ad hoc. Federal survey data show that in 2024, 79 percent of hospitals using predictive AI reported post-implementation evaluation or monitoring, but 18 percent answered “don’t know” when asked whether they did so at all.^[11]

Figure 1. Selected change and translation intervals—not directly comparable review cycles



*Contested research–translation estimate; not a local review interval.

Sources: Morris et al. (research–to–practice estimate, contested)^[1]; Menlo Ventures (buying cycles)^[10]; Becker’s Hospital Review (Epic release cadence)^[4]; ARK Investment Management via OfficeChai (frontier model intervals)^[5]. Intervals are shown on a logarithmic scale. Research translation, purchasing, and software release intervals measure different processes; this comparison does not measure a health system’s local evidence–review cadence. The 11-day estimate reflects the reporting period available in April 2026.

The key insight is not that technology moves fast. Leaders already know that. It is that a best-practice claim becomes uncertain when the evidence clock runs slower than material organizational change. Every time the organization changes something a practice depends on, whether a software version, a staffing model, a patient population, or a payer rule, the old proof describes a slightly different organization. The gap between the last proof and the present is what we call evidence debt. Like technical debt, it builds quietly and costs more the longer it is left alone. Cleveland Clinic’s IT leadership described the same dynamic for EHR upgrades: “We have found it’s important not to get behind, which makes for a significant backlog and higher level of effort over time.”^[9]

This also frees leaders from a false choice. An organization does not have to match the vendor’s clock. Nebraska Methodist’s deliberate annual cadence and Memorial Hermann’s quarterly cadence can both be sound. What cannot be sound is changing on one clock and proving on a slower one.

3. Two tests that earn the name

If a best practice must be both proven and current, every practice in a health system falls into one of four categories. The grid below is a KinRidge framework, not a published standard, but it gives executives a shared language for an uncomfortable conversation. “Proven” refers to demonstrated value and applicable outcome requirements, not a positive financial return for every practice. The grid assesses evidence; it does not by itself authorize stopping essential care or mandatory compliance activity.

	Current: re-examined since the last material change	Stale: not re-examined since the last material change
Proven: locally demonstrated value and applicable outcome requirements met	Best practice. Keep it, spread it where context matches, and schedule the next review.	Legacy practice. It worked once. Reassess affected evidence before expansion or renewed investment; manage continuity and safety during review.
Unproven: intended value or required outcomes not yet demonstrated	Experiment. Promising and monitored. Keep discretionary implementation bounded until intended value is shown; essential or mandatory activity requires appropriate safeguards.	Habit. Neither proven nor reviewed. Prioritize assessment, then redesign or retire where appropriate; do not discontinue essential or mandatory activity without a safe alternative.

Table 1. The proven–current grid. Practices move between cells as evidence builds up or decays.

The grid gives a practical meaning to the phrase “in today’s age.” In a slow-changing environment, a proven practice could stay in the top-left cell for years. In an AI-enabled environment, it can slide into the legacy cell within a single release cycle without anyone noticing, because the practice itself was never changed. Only its context was.

This proposal builds on work in evidence currency, living guidelines, continuous AI monitoring, and de-implementation.^[6, 12, 13, 14] Appendix A sets out those connections and distinguishes published findings from the KinRidge operating design.

4. Demonstrated value is the entry ticket

The proven test must fit the purpose of the practice. Financial and operational investments require measured benefits, total costs, and finance validation of financial claims. Clinical and safety practices require locally demonstrated effectiveness and safety, with resource implications made explicit. Mandatory compliance practices must demonstrate that requirements are met; their implementation should be assessed for effectiveness, burden, and cost rather than required to generate positive financial ROI. These are proposed KinRidge decision rules, not a replacement for applicable clinical or regulatory standards. Local assurance may draw on transferable external evidence; not every institution must repeat the original clinical trial. The need for local testing depends on the practice, its dependencies, and risk. Financial evaluation should state incremental costs, the time horizon, attribution, and benefit persistence, without double counting released capacity, converted capacity, and contribution margin.

Enthusiasm is not evidence, and adoption is not return. A widely cited 2025 study of enterprise generative AI from MIT's NANDA initiative has been widely reported as finding that roughly 95 percent of generative AI pilots fail to reach measurable value. It attributed the gap less to model quality than to a "learning gap": tools that "don't learn from or adapt to workflows."^[15] In healthcare, Bain and KLAS report that executives now prioritize AI with "tangible results and short time-to-value windows," even as many providers say it is still early to quantify precise financial returns.^[16]

Ambient clinical documentation shows why financial claims require an ROI test, and why that test has to be local.

- The benefit to clinicians looks real. Peterson Health Technology Institute found that early adopters report reduced burnout and cognitive load, but that "the financial impact is unclear." Its executive director put the challenge plainly: health systems "need to measure the results."^[17]
- Products in the same category performed differently. In a randomized trial of 238 UCLA Health physicians, one ambient scribe reduced time-in-note by 9.5 percent relative to usual care, while another showed no significant reduction.^[18] "Ambient AI" is a category, not a practice. The practice is a specific tool in a specific workflow in a specific organization.
- Scale does not equal return. 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 between October 2023 and December 2024.^[19] The total is impressive. A simple division gives roughly 2.2 hours per participating physician, but that is not an exposure-adjusted estimate: participation and use varied over the period. Those hours may support more visits, shorter days, better retention, or better notes. Each outcome must be measured separately; financial ROI additionally requires validated financial benefits net of full costs.

Health-system leaders acknowledge the difficulty. Mayo Clinic's chief AI implementation officer noted that on the care delivery side "it's really hard to apply line of sight" to ROI, and Ardent Health's chief digital officer cautioned that success metrics must be "clearly defined and measurable alongside a baseline."^[20] The answer is to keep financial claims disciplined while recognizing that clinical, safety, and people outcomes require their own evidence. Do not label a nonfinancial benefit as a financial return. We recommend a value ledger with six lines, agreed with finance before launch:

Ledger line	What it captures	Why it must be separate
Hard savings	Cash costs removed: agency labor, outsourced services, supplies	The only line a CFO can bank without further decisions
Capacity released	Staff time freed, measured against baseline	Not a realized financial return until converted; track capacity separately from cash savings
Capacity converted	Released time turned into visits, throughput, or reduced overtime	Whether released capacity becomes an operational or financial benefit
Revenue and margin	Added contribution, collections, documentation accuracy	Must be net of payer behavior and coding compliance
People and quality	Retention, burnout, safety, patient experience	Outcome evidence in its own right; a gain can justify a practice, and a loss can veto one. Do not relabel it as cash savings.
Costs shifted	New work, review burden, or risk created elsewhere	Include review burden and work transferred elsewhere in full costs

Table 2. The value ledger. Proposed KinRidge method, building on benefits–realization practice in healthcare.^[21]

A practice moving through the framework

Illustrative example—not observed results. The dates and decisions below are hypothetical and demonstrate the proposed method, not a product endorsement or an empirical finding.

January 31: locally proven. A health system completes an ambient-scribe pilot. Against a baseline agreed before launch, the evaluation shows documentation–time benefits, acceptable note quality, and a realized financial benefit net of licensing, integration, training, and review costs. Finance validates the financial claims and clinical leaders validate quality. The practice enters the proven–current grid as a best practice, with January 31 recorded as its last-proven date.

April 15: material change. A vendor model update changes how draft notes are produced. The practice moves to legacy/under review because the January evaluation does not fully cover the new configuration. Its evidence debt consists of specific unresolved questions: whether time benefits persist, correction work has increased, and quality or costs have changed. The PMO records the model dependency, assigns the clinical operations owner, and sets a May 31 review deadline. Leaders defer expansion; continued use remains subject to clinical safety oversight.

May 31: targeted reassessment and decision. The team retests the affected documentation–time, correction–work, quality, and cost measures. If the agreed criteria still hold, the practice returns to best-practice status with a new last-proven date, and leaders can back continued use or expansion. If they do not, leaders authorize bounded redesign, park expansion pending specified evidence, or retire the tool through a safe transition. The review updates the evidence affected by the change rather than automatically repeating the entire original study.

5. A best practice in four acts: the sepsis model

The Epic Sepsis Model illustrates why local clinical validation and evidence currency matter. This case does not establish financial ROI, and predictive discrimination alone does not prove improved patient outcomes. It demonstrates why a widely deployed tool still requires assessment in its local setting.

Act one: adoption. Hospitals across the country implemented the model on the strength of developer-reported performance.

Act two: local evidence. A 2021 external validation at Michigan Medicine covering 38,455 hospitalizations found an area under the curve of 0.63, “substantially worse than the performance reported by its developer.” The model missed 67 percent of patients with sepsis while alerting on 18 percent of all hospitalized patients.^[22]

A related context-shift episode (April 2020). In April 2020, the University of Michigan deactivated the model because of spurious alerting tied to pandemic-era changes in patient characteristics, which “fundamentally altered the relationship between fevers and bacterial sepsis.”^[23] The model had not changed. The organization had.

Act four: the updated model was evaluated locally. Epic released a revised model in 2022 that could be fine-tuned on each site’s historical data. A 2026 prospective study across four health systems found better discrimination, with AUROC between 0.82 and 0.92, but high variability across institutions: the alert threshold needed to reach 60 percent sensitivity ranged from 14 to 37 depending on the site. The authors concluded that health systems “should conduct internal evaluations to validate local performance.”^[24]

Across these episodes, the product name did not establish local validity or continued fitness for use. The evidence depended on the version, population, setting, and time of evaluation. This supports the paper’s currency test; a complete value judgment would also require evidence of clinical effects, workflow burden, safety, and resource use.

6. Review at the speed of your organization

If evidence debt builds with every material change, then the right review interval for a practice is not a fixed calendar date. It is set by the fastest-changing element the practice depends on, and it is shortened by risk. A practice that relies on an AI model updated monthly cannot be governed by an annual review, however disciplined that review is. A stable process with stable inputs does not need monthly scrutiny, and imposing it would burn out the teams doing the work.

The table below translates that principle into an operating design. The intervals are illustrative starting points for calibration, not published standards. Every material change triggers an assessment of which evidence remains applicable. The depth of reassessment depends on the affected dependencies, uncertainty, and potential harm; it need not repeat the original evaluation in full. Review of clinical safety and performance may need to occur sooner than a financial reassessment.

Practice clockspeed	Typical examples	Monitoring	Value and currency review
Fast	AI-enabled workflows; vendor-updated models; generative AI tools; rapidly changing payer rules	Continuous performance and drift signals; frontline exception reporting	Assess affected evidence at every material model or version change; at least quarterly value and currency review
Moderate	EHR-dependent workflows; scheduling and capacity tools; revenue-cycle automation	Monthly adoption and outcome dashboards	Assess affected evidence at each material EHR upgrade the practice touches; at least twice-yearly value and currency review
Slow	Policy and process standards with stable inputs; physical-environment practices	Quarterly indicators and exception review	At least annually
Any practice	Upgrade, acquisition, new site of care, staffing-model change, safety signal, or deterioration in ROI	Triggered immediately	Immediate assessment of affected evidence; scope and urgency of reassessment proportionate to risk

Table 3. Matching review cadence to practice clockspeed. Illustrative KinRidge operating design.

Three design rules keep this approach workable:

1. Choose your clock deliberately. An organization's clockspeed should be a leadership decision, not a side effect of vendor release schedules. Then align the review cadence to the clock you chose.
2. Review the dependency, not just the practice. Tag each practice with the systems, models, roles, and populations it relies on. A material change triggers assessment of the affected evidence; the scope of revalidation depends on what changed and the associated risk.
3. Make "last proven" visible. Every practice in the register should display the date it last passed both tests. A best practice without a date is a belief.

7. Why AI shortens the half-life

Practices can lose effectiveness when their context changes, even when people follow the design. AI-enabled practices add further dependencies: because the model, the data, or the way people use its output has changed. Researchers describe this as dataset shift, which occurs "when a machine-learning system underperforms because of a mismatch between the data set with which it was developed and the data on which it is deployed." Common causes include changes in technology, in population and setting, and in behavior, such as new reimbursement incentives.^[23] The FDA has similarly described how changes in clinical practice, patient populations, data inputs, and infrastructure can degrade the real-world performance of AI-enabled medical devices; that discussion concerns regulated devices, not every workforce tool.^[25]

Recent accreditation and governance guidance supports local validation and continuing monitoring. In September 2025, the Joint Commission and the Coalition for Health AI released

guidance calling for policies, “appropriate local validation,” and monitoring, with Joint Commission’s CEO noting “how quickly AI is changing healthcare – and at a scale I’ve never seen.”^[26] In June 2026, the Joint Commission issued a review guide for its Responsible Use of AI in Health Care certification. The guide asks organizations to show evidence such as “performance/bias dashboards, drift alerts, review cadence, revalidation triggers,” and to explain how they “validate vendor-reported performance.”^[27] NIST’s AI Risk Management Framework likewise treats risk management as continuous across the AI life cycle.^[28]

AI can also speed up the evidence clock itself. Researchers writing in the *Journal of Medical Systems* describe language-model systems that reproduced and updated an entire issue of Cochrane reviews, 12 reviews covering 146,276 citations, in under two days of work, compared with an estimated 12 person-years by traditional methods. They caution that performance is uneven and requires human oversight.^[29] The implication for health systems is significant. The same technology that shortens the half-life of a practice can also shorten the time needed to re-prove it, if the organization builds the capability to use it.

8. The PMO as the organization’s clock keeper

A conventional project office coordinates schedules, budgets, and deliverables. The PMO proposed here would also act as an internal team of experts in process design, implementation, change management, data, finance, clinical operations, and technology. It would not replace service-line or frontline ownership. It would give those leaders a consistent method for moving from idea to evidence to adoption, and then back to evidence. AHRQ describes systematic learning from internal data and external evidence in learning health systems, and published healthcare case studies describe PMO and transition-support teams contributing project, change, and evaluation expertise. Those studies support feasibility, not a universal ROI claim for the model proposed here.^[30, 31, 32]

PMO responsibility	Decision it enables	Evidence to retain
Innovation intake	Which problems merit investment?	Problem, baseline, sponsor, alternatives, practice clockspeed
Design and implementation	What work, data, and roles must change?	Workflow map, dependency tags, integration plan, risk log
Change management	Can teams adopt the change safely?	Training, adoption, feedback, permitted local adaptations
Benefit realization	What value actually occurred?	Six-line value ledger; finance validation of financial claims; appropriate clinical and safety validation
Clock keeping	Is the evidence still current?	Last-proven date, evidence debt, next trigger, review decision
Portfolio decision	Should the practice continue, change, spread, or end?	Proven—current grid position; retirement rationale

Table 4. Responsibilities of a clockspeed-aware PMO. Proposed operating design; the literature does not prescribe a single PMO structure.

How the PMO earns its keep

The PMO is subject to the same value discipline as the practices it oversees. Its full cost includes dedicated staff, analytics and tools, and the clinical, operational, and finance time required for implementation and review. Its own value ledger should distinguish realized financial benefits, capacity released and converted, clinical or safety improvements, and spending ended or redirected through documented portfolio decisions.

Portfolio reporting should identify the PMO's contribution without attributing every initiative's benefit to the PMO. Forecast savings and estimates of avoided loss remain separate from realized financial results. Leaders should agree the PMO's objectives and evaluation method in advance, compare its contribution with its full costs, and adjust the operating model when review burden exceeds its decision value. No universal PMO return is assumed.

From evidence to commitment: Back It or Park It

The PMO should help leaders make explicit choices, not simply keep projects moving. A "Back It or Park It" discipline connects evidence to commitment: back an initiative when its value, readiness, and resource requirements justify proceeding; park it when a specific uncertainty or unmet condition makes commitment premature. Where evidence no longer supports a practice, recommend redesign or retirement rather than allowing continued spending by default. Any change to essential care or mandatory compliance requires a safe, compliant transition.

Backing an initiative is not a permanent endorsement or a declaration that it is already a best practice. Each commitment should identify an accountable owner, expected value, authorized resources, and conditions that trigger reconsideration. Parking is equally deliberate: record why the initiative is waiting, what would need to change, and when it will be reviewed. New evidence—not

renewed enthusiasm—should determine whether it returns to the active portfolio.

The proven–current grid describes the strength and currency of the evidence; Back It or Park It determines the commitment that evidence justifies. Applied to new initiatives and existing practices, this discipline connects “When was this last proven, and what has changed since?” to a decision about what the organization does next. It is presented here as a concept, not a prescribed scoring system.

The living practice register

The PMO’s central tool is a register in which each practice records the problem it addresses, the essential steps that produce the benefit, allowable local adaptations, the sites where it has been tested, its implementation cost, its value ledger, the applicable clinical, safety, or compliance criteria, its authorized resource commitment, its dependencies, its clockspeed tier, its last-proven date, its review owner, its next review trigger, and any reason and condition for parking further investment. A practice advances from candidate to locally proven to transferable only as evidence builds up. It moves to under review when a trigger fires, and to retired when the evidence no longer supports it. Retirement is a success of the system, not a failure of the team.

IHI’s Model for Improvement uses small tests of change to learn how an intervention works locally before broader implementation, and the StaRI reporting standard emphasizes describing both fidelity and adaptation.^[33, 34] Together they support a practical rule: preserve the elements that cause the benefit, and let the implementation approach respond to local conditions. Variation between sites, like the sepsis thresholds that ranged from 14 to 37, should be explained rather than averaged away.^[24]

Three rhythms, one system

The PMO sets three connected rhythms rather than one fixed calendar. Frontline teams monitor critical exceptions as work happens. Initiative owners review adoption, benefit, and unintended effects at the interval set by the practice’s clockspeed. Executives periodically decide which practices to expand, refine, replace, or stop. IHI’s work on sustaining improvement places responsibility in the daily management system, supported by leaders at several levels, and the dynamic sustainability framework argues for continued adaptation in changing settings rather than fixed fidelity to an original design.^[35, 36]

9. Independent implementation, internal ownership

An independent implementation partner can bring a useful outside view at the start. It can map the real process across departments, surface inconvenient baseline data, test assumptions made by software vendors, and align the business case with clinical and financial leaders. That independence matters more as vendors move faster: the Joint Commission’s certification guide explicitly asks how organizations validate vendor-reported performance.^[27] The MIT NANDA research also found that line managers, not only central AI teams, need to drive adoption.^[15]

Evidence on facilitation is nuanced. Trials and comparative studies have examined external and blended internal–external facilitation, with effects that vary by intervention and setting. They do not show that an external team is always better.^[37] The claim here is narrower: an independent partner can help establish credible methods and challenge conflicts of interest, while internal ownership is necessary to embed and adapt the work. The goal is local capability, not permanent dependence on a consultant.

10. Five questions for the executive team

Leaders can test their own position against this thesis with five questions:

1. What is our clockspeed? How often do we change the technology, workflows, and staffing models our most important practices depend on, and did we choose that pace on purpose?
2. Which of our “best practices” have a last-proven date? For those that do not, what evidence would we show a board member who asked today?
3. Where is our evidence clock slower than our organization clock? Start with AI-enabled and EHR-dependent practices.
4. What have we retired in the past year? A portfolio that never retires anything is not learning.
5. Could our own team defend the evidence and value method without the vendor in the room? Financial claims should withstand finance review, and clinical or safety claims should withstand appropriate clinical review. If not, the organization does not yet own its evidence.

11. How KinRidge can help

A practical engagement begins with a clockspeed diagnostic: an inventory of active practices and AI-enabled workflows, their dependencies, their last-proven dates, and the gap between the organization clock and the evidence clock. KinRidge and the health system then design the PMO charter, decision rights, staffing, value ledger, practice register, and review cadence together.

The teams then run a small number of initiatives through the full cycle, including at least one AI-enabled implementation and one process redesign, while training internal leaders to do the analysis and facilitation themselves. The transition is complete when the health-system team can independently evaluate a new idea, lead adoption, defend its evidence and value method, keep its evidence current, and recommend when a practice should change or retire.

The PMO’s purpose is not to protect every process it helped launch. It is to help the health system keep asking the harder question, at the speed it actually moves: Is this still the best way to work, and can we show the value it delivers now?

Appendix A. How this builds on prior work

The idea that knowledge decays is not new. This paper draws on evidence currency, living guidelines, continuous AI monitoring, de-implementation, and learning health systems. Its proposed contribution is to connect those ideas to the individual practice inside one organization, judged by demonstrated value and evidence currency.

Body of work	Prior findings or recommendations	What this paper adds
Evidence currency	High-quality systematic reviews often need updating within a few years; half showed a signal to update within 5.5 years. ^[6]	Applies the same logic to an organization's own operating practices, not only to published research.
Living guidelines	Recommendations can be continuously updated. One leading handbook sets no fixed interval but uses evidence searches every three months as its current standard. ^[12]	Sets the update pace by how fast the organization changes, not only by how fast new research appears, and requires demonstrated value, with finance validation of financial claims.
Continuous AI monitoring	Health systems can monitor deployed models in real time for accuracy, drift, and fairness, and pause them when needed, as in Vanderbilt's open-source VAMOS platform. ^[13]	Extends lifecycle monitoring beyond AI models to any practice, and ties it to value as well as safety.
De-implementation	Structured frameworks exist for reducing or stopping ineffective, harmful, or low-value care. ^[14]	Builds the review system upstream, so low-value practices are caught at the next trigger rather than years later.
Half-life of knowledge	Facts decay at measurable rates across fields. ^[38]	Turns a general idea into an operating rule and a named risk: evidence debt.
AI value and operating models	Health systems are reportedly "short on impact from AI," and the piece argues that value depends on redesigning workflows with continuing measurement and executive review. ^[39]	Replaces one fixed cadence with a cadence matched to each practice's clockspeed, and redefines what earns the name "best practice."
"Next practice" critiques	Business writers have argued that best practices from elsewhere rarely transfer and that organizations should pursue "next practices" instead. ^[40]	Keeps the term but raises the bar: a practice must be proven and current in this organization to carry the label.

Table A1. Prior work and the contribution of this paper.

Four elements form the proposed KinRidge approach: a locally demonstrated value test and a currency test; review cadence matched to material organizational change rather than only a calendar or vendor release schedule; evidence debt applied to operating practices; and the proven-current grid as a portfolio tool. Financial claims require finance validation, while clinical, safety, and compliance practices require evidence appropriate to their purpose. These are propositions offered for testing, not settled findings or a claim of established uniqueness.

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. Morris ZS, Wooding S, Grant J. The answer is 17 years, what is the question: understanding time lags in translational research. *Journal of the Royal Society of Medicine*. 2011. <https://journals.sagepub.com/doi/10.1258/jrsm.2011.110180>
2. Prasad V et al. A decade of reversal: an analysis of 146 contradicted medical practices. *Mayo Clinic Proceedings*. 2013. <https://pubmed.ncbi.nlm.nih.gov/23871230/>
3. Herrera-Perez D et al. A comprehensive review of randomized clinical trials in three medical journals reveals 396 medical reversals. *eLife*. 2019. <https://pmc.ncbi.nlm.nih.gov/articles/PMC6559784/>
4. Becker's Hospital Review. 34 new Epic features coming — and when. August 24, 2026. <https://www.beckershospitalreview.com/healthcare-information-technology/innovation/34-new-epic-features-coming-and-when/>
5. OfficeChai. Frontier labs are releasing new models faster than ever, shows data (ARK Investment Management analysis of Artificial Analysis data). April 29, 2026. <https://officechai.com/ai/frontier-labs-are-releasing-new-models-faster-than-ever-shows-data/>
6. Shojania KG et al. How quickly do systematic reviews go out of date? A survival analysis. *Annals of Internal Medicine*. 2007. <https://pubmed.ncbi.nlm.nih.gov/17638714/>
7. Densen P. Challenges and opportunities facing medical education. *Transactions of the American Clinical and Climatological Association*. 2011. <https://pmc.ncbi.nlm.nih.gov/articles/PMC3116346/>
8. Fine CH. Industry Clockspeed and Competency Chain Design: An Introductory Essay. MIT Sloan working paper. 1996. <https://dspace.mit.edu/bitstream/handle/1721.1/2614/5WP-3893-36119547.pdf?sequence=1>
9. Becker's Hospital Review. Health systems rethink EHR upgrades amid AI shift. April 28, 2026. <https://www.beckershospitalreview.com/healthcare-information-technology/innovation/health-systems-rethink-ehr-upgrades-amid-ai-shift/>
10. Menlo Ventures. 2025: The State of AI in Healthcare. October 2025. <https://menlovc.com/perspective/2025-the-state-of-ai-in-healthcare/>
11. Office of the National Coordinator for Health Information Technology. Hospital Trends in the Use, Evaluation, and Governance of Predictive AI, 2023–2024. <https://healthit.gov/data/data-briefs/hospital-trends-use-evaluation-and-governance-predictive-ai-2023-2024/>
12. Australian Living Evidence Collaboration. Living Guidelines Handbook, version 1.1. August 2025. <https://livingevidence.org.au/wp-content/uploads/Living-Guidelines-Handbook-v1.1-Aug2025.pdf>
13. Healthcare Innovation. VAMOS Collaborative: An Open Source Platform for Trustworthy AI. May 2026. <https://www.hcinnovationgroup.com/analytics/ai/ai-governance/article/55379896/vamos-collaborative-an-open-source-platform-for-trustworthy-ai>
14. De-Implementation of Low-Value Care. 2023. <https://pmc.ncbi.nlm.nih.gov/articles/PMC10842822/>

15. 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/>
16. 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/2025/healthcare-ai-investment-focusing-on-hard-dollar-returns-and-clinical-workflows-amid-shift-from-pilots-to-widespread-productionbain-company-and-klas-research/>
17. 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/>
18. Lukac PJ et al. Ambient AI Scribes in Clinical Practice: A Randomized Trial. NEJM AI. 2025. <https://pmc.ncbi.nlm.nih.gov/articles/PMC12768499/>
19. 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/>
20. 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/>
21. Identifying value in healthcare transformation initiatives: an evaluation of an approach to benefits realisation. BMJ Open Quality. 2023. <https://doi.org/10.1136/bmj-oq-2023-0002349>
22. 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/>
23. Finlayson SG et al. The Clinician and Dataset Shift in Artificial Intelligence. New England Journal of Medicine. 2021. <https://pmc.ncbi.nlm.nih.gov/articles/PMC8665481/>
24. Wong A et al. Multicenter Prospective Validation of an Updated Proprietary Sepsis Prediction Model. JAMA Network Open. 2026. <https://pmc.ncbi.nlm.nih.gov/articles/PMC12949446/>
25. U.S. Food and Drug Administration. Measuring and Evaluating AI-Enabled Medical Device Performance in the Real World: Request for Public Comment. 2025. <https://www.fda.gov/medical-devices/digital-health-center-excellence/request-public-comment-measuring-and-evaluating-artificial-intelligence-enabled-medical-device>
26. Joint Commission and Coalition for Health AI. Joint Commission and CHAI release initial guidance to support responsible AI adoption across U.S. health systems. September 17, 2025. <https://www.chai.org/news/joint-commission-and-coalition-for-health-ai-chai-release-initial-guidance>
27. 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>
28. NIST. AI Risk Management Framework Core. <https://airc.nist.gov/airmf-resources/airmf/5-sec-core/>
29. Peltonen LM, Topaz M, Zhang Z. From Research to Practice in Days, not Decades: Why Leaders Must Act Now. Journal of Medical Systems. 2025. <https://pmc.ncbi.nlm.nih.gov/articles/PMC12672604/>
30. Agency for Healthcare Research and Quality. How Learning Health Systems Learn: Lessons from the Field. <https://www.ahrq.gov/learning-health-systems/how-lhs-learn.html>

31. Implementation of evidence-based practices in the context of a redevelopment project in a Canadian healthcare organization. 2012. <https://pubmed.ncbi.nlm.nih.gov/23121763/>
32. A Health Care Project Management Office's Strategies for Continual Change and Continuous Improvement. 2017. <https://pubmed.ncbi.nlm.nih.gov/28953067/>
33. Institute for Healthcare Improvement. Model for Improvement: Testing Changes. <https://www.ihl.org/library/model-for-improvement/testing-changes>
34. Pinnock H et al. Standards for Reporting Implementation Studies: StaRI statement. BMJ. 2017. <https://www.bmj.com/content/356/bmj.i6795>
35. Institute for Healthcare Improvement. Sustaining Improvement. White paper. <https://www.ihl.org/library/white-papers/sustaining-improvement>
36. Chambers DA et al. The dynamic sustainability framework: addressing the paradox of sustainment amid ongoing change. Implementation Science. 2013. <https://link.springer.com/article/10.1186/1748-5908-8-117>
37. Smith SN et al. Comparative effectiveness of external versus blended facilitation on collaborative care model implementation. Health Services Research. 2020. <https://pmc.ncbi.nlm.nih.gov/articles/PMC7704469/>
38. Arbesman S. The Half-Life of Facts: Why Everything We Know Has an Expiration Date. Current/Penguin. 2012. <https://archive.org/details/half-life-of-facts-sw000-arbe>
39. McKinsey & Company. The health system CEO imperative: Turning AI's promise into performance. June 15, 2026. <https://www.mckinsey.com/industries/healthcare/our-insights/the-health-system-ceo-imperative-turning-ai-promise-into-performance>
40. Arthur D. Little. Moving from "best practice" to "next practice" to drive effective digital transformation. https://www.adlittle.com/sites/default/files/viewpoints/adl_moving_from_best_practice_to_next_practice_to_drive_effective_digital_transformation-min.pdf