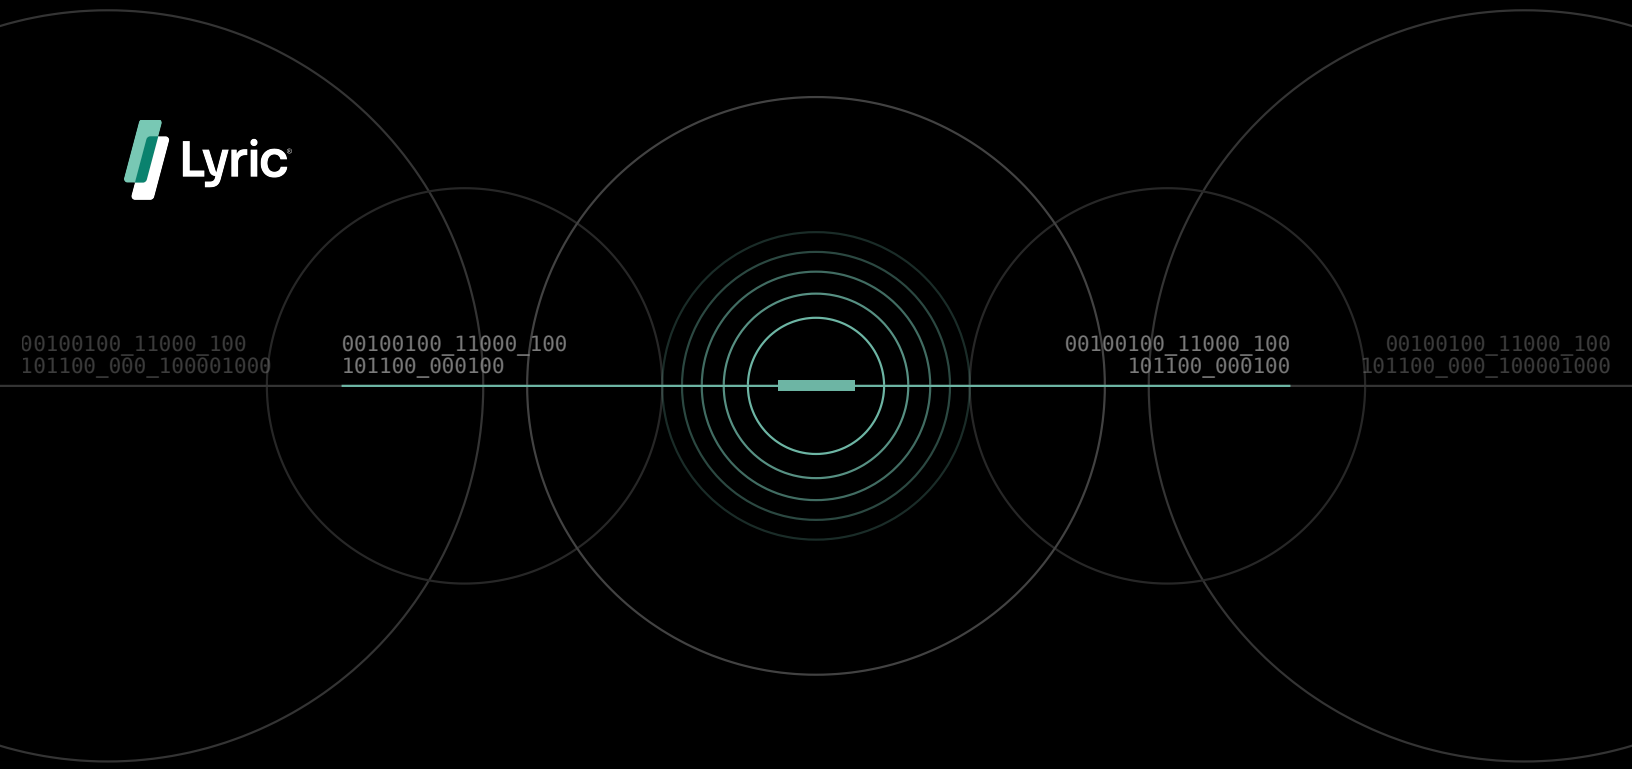
Lyric



REPORT · NOVEMBER 2026

Defining Healthcare Decision Intelligence

The Discipline, the Standard, and Why Healthcare Is Ready

Yevgeny German
Chief Technology Officer

Brian Berkowitz
Chief Product Officer

Alex Ulyanov
SVP of Engineering & AI

Contents

Key observations 3

Why This Paper Exists 4

01 The Discipline 5

The Threshold 6

What Decision Intelligence Is 6

Where the Discipline Matured 9

Where Those Industries Are Now 12

02 Why Healthcare Is Ready 15

Why Healthcare's Need Is Acute 16

Why Healthcare Arrived Late 18

Why Now 18

Governed AI and the Limits of Autonomy 20

The Distinction That Actually Matters 21

03 The Standard 24

What Qualifies as Healthcare Decision Intelligence 25

Where the Category Proves Out First 28

Notes 29

ABSTRACT

Healthcare decision intelligence is described with increasing frequency and decreasing precision. This paper defines decision intelligence as a discipline, traces its origins, describes outcomes in the industries that adopted it first, explains why healthcare's need is now acute, and outlines what a system must do to qualify. It is a category statement rather than an assessment of any organization operating in it.

Key observations

1 Decision intelligence is the discipline of engineering how decisions get made, then measuring what they produce.

2 Work performed before the decision is analysis. Work performed after is reporting and recovery. Decision intelligence operates at the moment of decision.

3 Healthcare has every characteristic that made the discipline work elsewhere, and the pressure behind the need is intensifying.

4 Retrieval, assembly, extraction, routing, prioritization, and documentation are work. Decisions about a patient's care, coverage, or payment are determinations.

5 The standard includes three conditions that separate healthcare decision intelligence from the capabilities it is most often confused with: expertise built into the software; intelligence traceable at decision speed; and applied human judgment.

Why This Paper Exists

Healthcare decision intelligence is described with increasing frequency and decreasing precision. The term is applied to analytics platforms, automation tools, AI pilots, and reporting layers that share little beyond the underlying technology.

That imprecision has a cost. When a category lacks boundaries, buyers cannot compare what they are being offered, organizations cannot tell whether an investment moved them forward, and the conversation stalls at capability claims.

This paper defines decision intelligence as a discipline, traces its origins, describes outcomes in the industries that adopted it first, explains why healthcare's need is now acute, and outlines what a system must do to qualify. It is a category statement rather than an assessment of any organization operating in it. The standard it sets applies across healthcare, to clinical, operational, and financial decisions alike.

Its companion, The State of Healthcare Decision Intelligence in Payment Accuracy, applies the standard established here to the function where the category is proving out first.

Lyric

PART

01

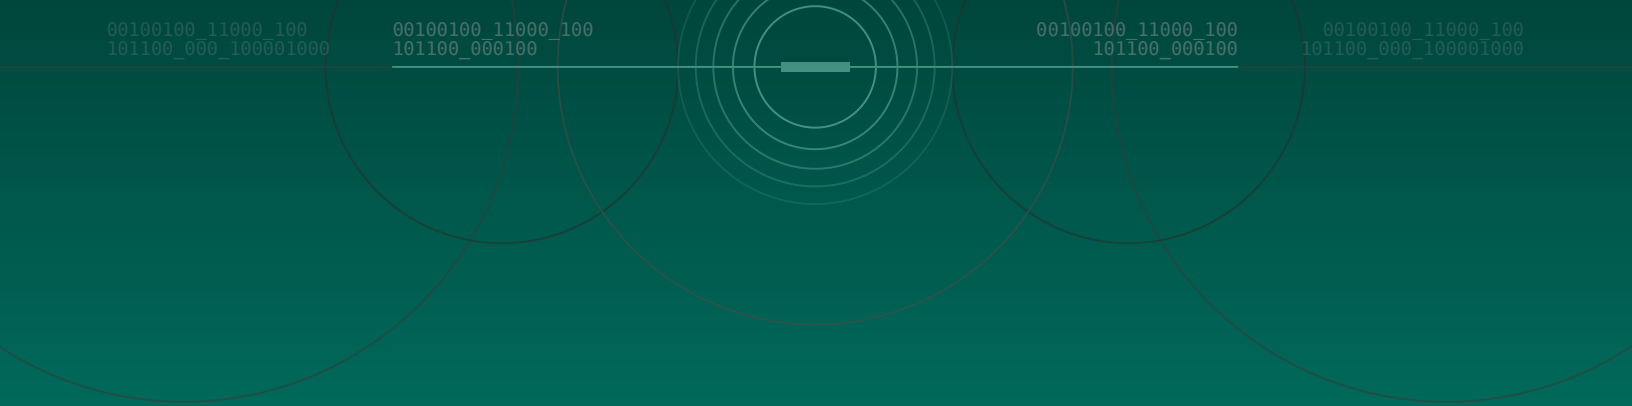
The Discipline

The Threshold

What Decision Intelligence Is

Where the Discipline Matured

Where Those Industries Are Now



The Threshold

Every industry eventually reaches a point where the number of consequential decisions it must make exceeds its capacity to make them well.

The decisions do not become harder individually. Too many arrive too quickly, governed by too much policy, for expertise alone to apply consistently. The organization is not short on knowledge. It is short on the ability to bring that knowledge to bear at the moment each decision is made.

Several industries hit that point decades ago. What they built in response became a discipline, and healthcare is now arriving at the same threshold.

What Decision Intelligence Is

Decision intelligence is the discipline of engineering how decisions get made, then measuring what they produce.

Systems built this way combine analytics, business rules, and AI to support, augment, and automate decisions inside the workflow where those decisions occur, then feed the outcomes back into the logic that produced them. What distinguishes the discipline is what it does with information, not what it does to it.

Decision intelligence's foundations predate the term. Operations research took shape in Britain's prewar air-defense work as scientists helped operationalize radar and control fighter interception in the later 1930s,[1] and Stanford's Ronald Howard named decision analysis in 1966.[2] Computerized decision support followed a parallel path from theory to software. Model-driven systems built in the mid-1960s combined data and analytical models for a single decision-maker, and MIT researchers named the field decision support systems in 1971.[3] Spreadsheets put what-if modeling on analysts' desks in the 1980s. Data warehouses, executive information systems, and online analytical processing (OLAP) followed in the late 1980s and early 1990s, and that data work grew into business intelligence.[4] Each generation improved the tools available. Few engineered the decision process itself or returned outcomes to the logic that produced them, and all were limited by the data and computing of their time. What changed is that the theory became operable at production volume.

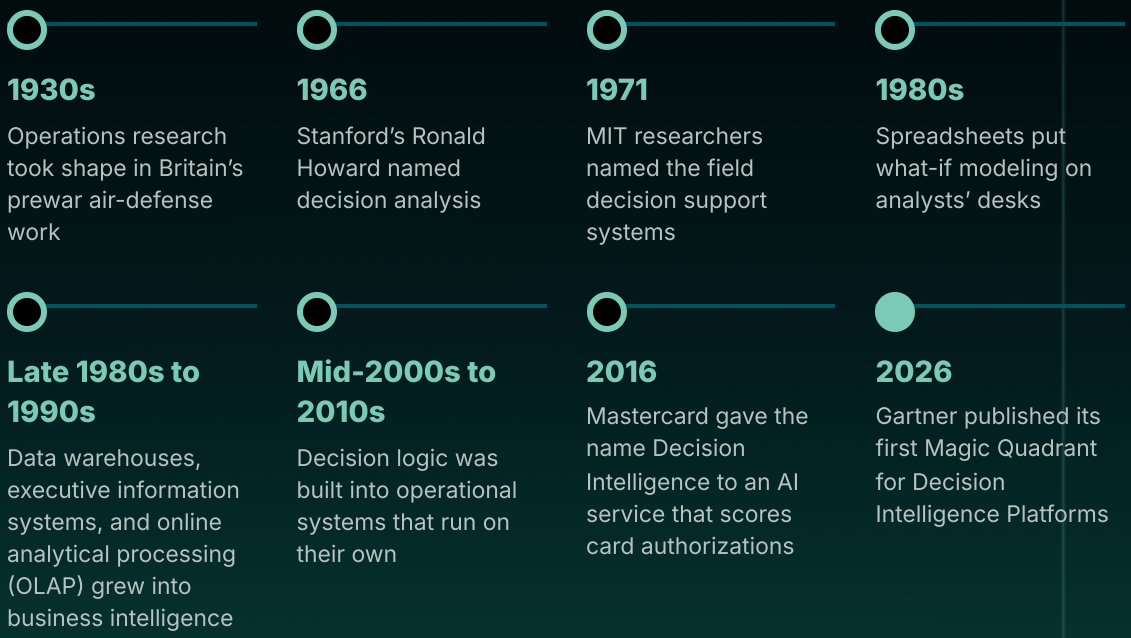
From the mid-2000s to the mid-2010s, the work moved from helping one person make a decision to building decision logic into operational systems that run on their own.[5] Decision intelligence took that automated logic, combined it with machine learning, and added a feedback loop that sends results back into the logic.

Lorien Pratt and Mark Zangari were developing a decision-modeling approach they called Decision Engineering by 2008; Pratt later described the broader discipline as decision intelligence.[6,7] In 2016, Mastercard gave the name Decision Intelligence to an AI service that scores card authorizations across its network.[8] Gartner now defines the field as "a practical discipline used to improve decision making by explicitly understanding and engineering how decisions are made, and how outcomes are evaluated, managed and improved via feedback,"[9] and published its first Magic Quadrant for Decision Intelligence Platforms in January 2026.[10]

FIGURE 1

Historical foundations of decision intelligence

From theory to production volume



Source: note 1, note 2, note 3, note 4, note 5, note 8, note 10.

The defining property is timing.

Work performed before the decision is analysis. Work performed after is reporting and recovery. Decision intelligence operates at the moment of decision.

FIGURE 2

Where decision intelligence operates



Business intelligence reports what happened. Forecasting projects what is likely. Simulation tests assumptions under varied conditions. Each is valuable, but none is decision intelligence. A system qualifies not because it uses AI, and not because it produces insight, but because it brings the relevant data, policy, evidence, and recommendation to the decision while it is still open and can still be changed.

The test is where intelligence is applied, not where it originates. Outcomes generated after a decision is made are essential inputs to the discipline. What matters is that intelligence returns to the point where the next decision is made.



Where the Discipline Matured

Three domains moved first, and each went deep in one area rather than broad across many.

FIGURE 3

Where the discipline matured

~1 ms

CREDIT AND FRAUD

Visa reported scoring transactions against more than 500 risk attributes in about one millisecond by 2019

100M

SUPPLY CHAIN PLANNING

Miles a year UPS ORION was projected to eliminate, along with 10 million gallons of fuel

23 to 5

INSURANCE UNDERWRITING

Average days from application to underwriting decision; 83% of eligible applications still received human review

Source: note 13, note 15, note 18.

Credit and fraud

Decades of scoring expertise executing in tens of milliseconds against live transactions. Lending moved from individual underwriter judgment to scored, policy-governed decisions over roughly four decades: Fair Isaac built its first credit scoring system in 1958, introduced the general-purpose FICO score in 1989, and by 1995 Fannie Mae and Freddie Mac were using it to qualify mortgage borrowers.[11] That year Freddie Mac also launched Loan Prospector, which the company said returned an underwriting decision in minutes.[12] Card networks moved from post-transaction investigation to authorization-time scoring, deciding whether to approve, challenge, or decline while the transaction is still open. Visa first applied neural networks to fraud in 1993. Visa reported that by 2019 its Advanced Authorization system scored transactions against more than 500 risk attributes in about one millisecond, helping prevent an estimated $25 billion in fraud annually and holding global fraud below 0.1% of volume.[13]

Supply chain planning

Capacity, constraints, and network flows modeled as one connected system, with decisions coordinated across the network rather than optimized function by function. The path began in 1964, when IBM engineer Joseph Orlicky developed material requirements planning.[14] UPS spent 10 years building ORION, an operations research system that plans routes for 55,000 drivers. Fully deployed in 2016, ORION was projected to eliminate about 100 million miles and 10 million gallons of fuel annually.[15] McKinsey reported in 2021 that successful early adopters of AI-enabled supply-chain management reduced logistics costs by 15% and inventory by 35% while increasing service levels by 65% compared with slower-moving competitors.[16]

Insurance underwriting

Actuarial judgment encoded as executable rules, operating within regulatory and actuarial boundaries, with human underwriters concentrated on complex and high-value risk. In U.S. life insurance, Munich Re found in its survey that the first accelerated underwriting programs reached the market in 2012.[17] By 2024, 82% of carriers surveyed by Gen Re had fully or partially implemented them, and the average time from application to underwriting decision fell from 23 days to 5. The expert stayed in the loop: 83% of eligible applications still received some level of human underwriter review.[18]

The common conditions are worth noting because healthcare shares all of them. Each domain involved high decision volume, largely structured data, codified policy, measurable outcomes, and a clear cost when the decision was wrong. Each encoded specialized expertise into systems rather than leaving it with individual experts. Each also connected the outcome of a decision back to the logic that produced it.

One further pattern matters more than the rest.



A system is only as good as the human expertise behind it. Decades of experience, judgment, and nuance are what make its decisions sound.

None of these domains succeeded by applying a general-purpose technology to a specialized problem. They succeeded by building depth in one domain and maintaining it as the domain evolved.

Where Those Industries Are Now

The domains that adopted the discipline earliest did not stop once they had expertise in making good decisions in the moment. They moved ahead of the moment.

Fraud detection and supply chain planning both progressed from responding at the point of transaction to identifying the conditions that produce a problem before the problem presents itself. The capability behind that shift is signal and sensing: applying anomaly detection and pattern recognition continuously across live activity, surfacing risk while it is still forming, not after it has become material.

Card fraud shows the shift most clearly. In 2016, Mastercard began scoring authorizations with an AI service it named Decision Intelligence. By 2024, Mastercard said the service was helping banks score and safely approve 143 billion transactions a year. Its newer Decision Intelligence Pro capability improves the existing Decision Intelligence score in less than 50 milliseconds.[19] That same year, the company moved ahead of the transaction altogether. Its generative AI now scans card details exposed on criminal websites, predicts the full card numbers from partial data, and lets banks block compromised cards before they are used. Mastercard reported that the capability doubled its detection rate for potentially compromised cards.[20] In the U.K., Mastercard's Consumer Fraud Risk tool flags likely scam payments before money leaves a consumer's account; within four months of adopting it in 2023, TSB reported a marked increase in fraud detection. Mastercard estimated that comparable performance across U.K. banks could prevent about £97 million in annual losses.[21]

FIGURE 4

Moving ahead of the moment

Card fraud shows the shift most clearly.

143B

TRANSACTIONS A YEAR

Mastercard said in 2024 that its Decision Intelligence service was helping banks score and safely approve this many transactions a year

<50 ms

SCORE IMPROVEMENT

Time in which Decision Intelligence Pro improves the existing Decision Intelligence score

2x

DETECTION RATE

Mastercard reported its generative AI doubled detection of potentially compromised cards

Source: note 19, note 20.

Supply chain planning is moving the same way. In one peer-reviewed complex-asset manufacturing case, a machine-learning model predicted late first-tier supplier orders with 80% accuracy, creating an earlier signal for contingency planning and risk mitigation.[22]

That changes the question the system is answering.

A mature decision intelligence system answers what should happen now. Its forward-looking form answers what might happen, what could happen, and what should happen next.

The practical consequence is that the system can teach an organization about its own operation. Most organizations can see their own results but have limited ability to recognize which of their patterns are unusual, which are early, and which are about to become material. Intelligence operating at scale across a broad, diverse base can identify those conditions and pair them with a recommended course of action rather than a report.

This is the trajectory healthcare is entering, not the state it has reached. But it follows the same direction the discipline has taken everywhere else it has matured, and it sets the target for healthcare: recognizing a patient at rising risk before an avoidable admission, a gap in care before it widens, or a billing pattern before it becomes an overpayment.

Lyric

PART

02

Why Healthcare Is Ready

Why Healthcare's Need Is Acute

Why Healthcare Arrived Late

Why Now

Governed AI and the Limits of Autonomy

The Distinction That Actually Matters

Why Healthcare's Need Is Acute

Healthcare has every characteristic that made the discipline work elsewhere, and the pressure behind the need is intensifying.

Decision volume keeps rising. A single episode of care can require decisions about diagnosis, treatment, site of care, referrals, documentation, coverage, coding, and payment, often made by different parties in different systems. Multiply that across the patients a health system treats and the claims a health plan processes, and you get one of the most decision-intensive operating environments in any industry.

The underlying knowledge changes faster than the systems governing it. New treatments, drugs, procedures, and coding conventions arrive continuously. Clinical evidence moves faster than annual release cycles. A frequently cited 2011 projection, never validated against later data, estimated that medical knowledge would double every 73 days by 2020.[23] Since then, AI-assisted research has begun to shorten the path from discovery to candidate therapy across new modalities and precision medicine, adding to the volume of evidence the systems governing care and payment have to keep pace with.

Requirements shift across state, federal, and industry authorities, and those authorities update on cycles measured in quarters and years. Commercial guidelines and coding sets move no faster. Once an application is filed, securing a new Category I CPT code from the AMA can take 18 to 24 months, and gathering the required evidence can take two to five years before that.[24] Those cycles were built when clinical publications and practice changes arrived in volumes that periodic updates could absorb. That is no longer the case. The gap is not only that medicine moves quickly; the rules governing care and payment move slowly, and the distance between the two keeps widening.

Administrative burden remains substantial. A 2019 *JAMA* analysis estimated that waste consumes $760 billion to $935 billion of U.S. healthcare spending each year, roughly 25% of the total, with administrative complexity alone accounting for $265.6 billion.[25] Much of that waste is generated when a decision is made without the information that would have made it correct. In fiscal 2025, CMS estimated that 6.55% of Medicare fee-for-service payments, $28.83 billion, were paid improperly.[26] Costs can compound between the moment an error is made and the moment it is discovered, adding administrative work to correct the original decision.

FIGURE 5

Why healthcare's need is acute

The pressure behind the need is intensifying.

$760B to $935B

ESTIMATED ANNUAL U.S. WASTE

Roughly 25% of healthcare spending (JAMA, 2019)

$265.6B

ADMINISTRATIVE COMPLEXITY

Estimated share of that waste, per the same analysis

6.55%

MEDICARE FFS IMPROPER PAYMENTS

$28.83 billion, CMS estimate for fiscal 2025

18 to 24 mo.

NEW CATEGORY I CPT CODE

Once an application is filed, with two to five years of evidence gathering before that

Source: note 24, note 25, note 26.

The cost of getting it wrong is not absorbed statistically. A credit model that is wrong at the margin produces a portfolio-level cost. A wrong healthcare decision reaches an individual directly: a patient waiting on care, a clinician acting on incomplete information, a provider relationship built over years.

Taken together, these conditions produced decision intelligence everywhere else it emerged. Healthcare is not an exception to the pattern. It is a late arrival to it.

Why Healthcare Arrived Late

Three constraints held healthcare back, and they were real, not incidental.

Data sat fragmented across plans, providers, vendors, and systems, so the context needed for good decisions was rarely available at the point the decision was made. Even now, the gap persists: in 2023, only 43% of U.S. hospitals routinely sent, received, found, and integrated patient information from outside organizations, and just 42% said their clinicians often used that outside information when treating patients.[27] Healthcare policy is interpreted rather than simply applied, which makes it harder to encode and maintain. And accountability cannot transfer: regulation requires an identifiable person to own the outcome, and past decisions must remain traceable to the policy in effect when they were made, not the policy in effect today.

Those constraints ruled out the approaches that worked in other industries. A system that decides quickly but cannot explain itself has no application in healthcare.

Why Now

The constraints that delayed healthcare's adoption of decision intelligence have not disappeared. What changed is that the technology can now operate inside them, not around them.

Data can be unified without losing lineage. Decisions can be executed in milliseconds without abandoning governance. Models can be versioned, explained, and audited so that the logic behind a decision can be reconstructed months later. Outcomes can be captured and returned to the systems that generated them.

The mechanics of the discipline transfer. Healthcare requires depth in clinical evidence, care guidelines, regulation, coding, and the policy governing how care is delivered, covered, and paid, maintained as the evidence changes.

That depth is also the barrier to entry, and general-purpose technology cannot build it quickly.

Healthcare is next because the constraints that delayed it are now addressable.

FIGURE 6

Why healthcare arrived late, and why now

The constraints have not disappeared. The technology can now operate inside them, not around them.

Then: three constraints

- Data sat fragmented across plans, providers, vendors, and systems
- Healthcare policy is interpreted rather than simply applied
- Accountability cannot transfer; past decisions must remain traceable

Now: inside the constraints

- Data can be unified without losing lineage
- Decisions can be executed in milliseconds without abandoning governance
- Models can be versioned, explained, and audited
- Outcomes can be captured and returned to the systems that generated them



Governed AI and the Limits of Autonomy

The capability question in healthcare is no longer whether AI can act. It is what AI should be permitted to act on.

Regulated industries have a clean answer, but the broader market is still looking for it. Full autonomy, a system that decides and executes without a person in the loop, is not the destination for healthcare decision intelligence. It is the wrong objective, and it will remain the wrong objective regardless of how capable the underlying models become.

This is not a statement about what the technology can do. It is a statement about what the environment permits.

Accountability does not transfer to a model. In regulated healthcare decisions, responsibility remains with the organization and, where circumstances permit, with a licensed professional charged with reviewing or making the decision. In Medicare Advantage, federal rules require a licensed physician or other appropriate healthcare professional to review any adverse medical necessity decision before it is issued.[28] In February 2024, CMS told plans that an algorithm may assist a coverage determination but cannot replace an assessment of the individual patient's circumstances, and that a predictive algorithm cannot be the sole basis for ending post-acute care.[29] California went further: since January 2025, state law has required that a licensed physician or qualified healthcare provider decide any denial or modification of care based on medical necessity.[30] The same principle holds at the point of care. A clinician who acts on a recommendation from decision support generally remains responsible for the treatment decision, and that responsibility does not pass to the system that produced the recommendation. A model cannot hold that responsibility, be examined under it, or answer for it. Any architecture that removes the person also removes the responsible party, which is not a trade a regulated organization can make.

Traceability is a core design principle, not a feature. A decision must be capable of reconstruction months or years later, against the policy, data, and logic in effect when it was made. That traceable line, from a recommendation back to the policy, the data, and the reasoning that produced it, is the decision's lineage. It has to travel with the recommendation rather than sit in a document somewhere. Systems that cannot reproduce their own reasoning at a specific point in time fall short of that standard no matter how accurate they are in aggregate.

Insurance regulators are writing that expectation down. The NAIC's model bulletin on insurers' use of AI, adopted by 25 jurisdictions, including 24 states and the District of Columbia, as of April 1, 2026, asks insurers to maintain written governance of their AI systems, test and document them, weigh how far people are involved in the final decision, and be ready to produce that documentation for examiners.[31] Medicare Advantage organizations must already retain books, records, and other evidence for 10 years.[32]

Trust is the constraint that binds first. A decision that cannot be explained to the patient or provider on the other end, or to a regulator asking why, creates more friction than the decision saved. In an environment already carrying real mistrust among health plans, providers, and patients, an unexplainable system deepens the problem it was meant to solve.

The same line is being drawn outside the United States. The EU AI Act classifies AI used for risk assessment and pricing in life and health insurance as high-risk and requires such systems to be built so a person can "disregard, override or reverse the output."[33]

The Distinction That Actually Matters

Ruling out full autonomy is not the same as ruling out automation. A system should not decide and execute without a person in the loop, but it can and should automate much of the work around the decision. Conflating the two leaves much of the available value unrealized.

Automation and determination are different things.

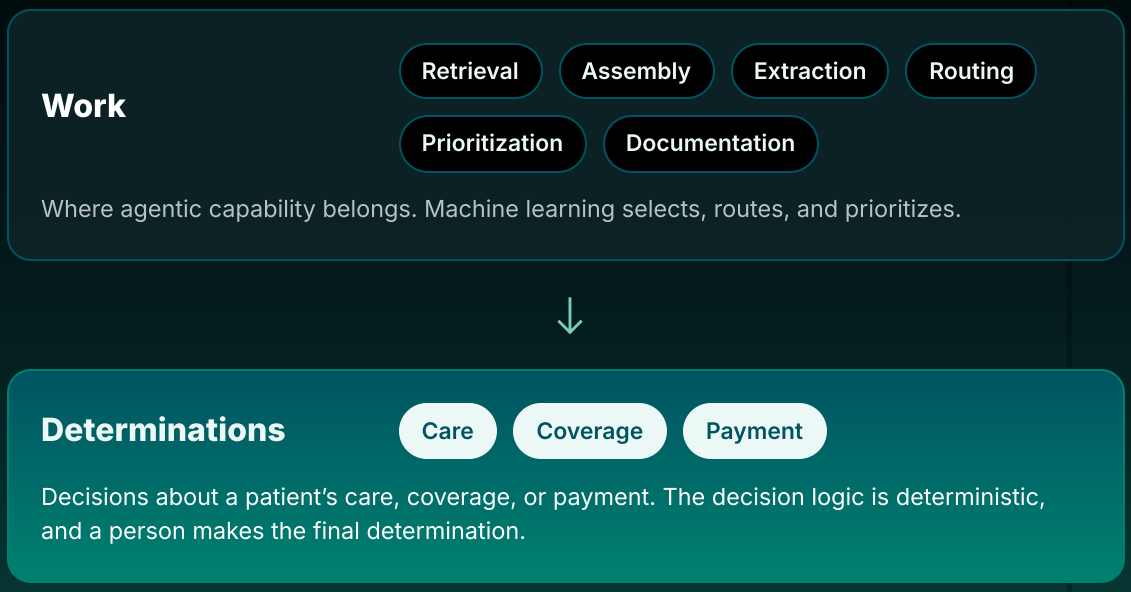
The difference matters because probabilistic machine-learning models introduce a kind of error that deterministic rule-based automation does not. That error is tolerable in work and intolerable in a determination.



Retrieval, assembly, extraction, routing, prioritization, and documentation are work. Decisions about a patient's care, coverage, or payment are determinations. The first should be automated aggressively. The second should not be automated at all.

FIGURE 7

Automation and determination are different things



Most of the effort around a healthcare decision is not the decision itself. It is finding the relevant information across systems, assembling the supporting evidence, confirming which policy or guideline applies, identifying what needs attention, and preparing it for someone to act on. That work is high-volume, largely mechanical, and consumes most of the expert's time.

This is where agentic capability belongs, and where it creates substantial value. A system that can independently gather what the decision-maker needs, assemble it in context, and surface the cases that merit judgment is doing exactly what it should. It expands the surface of work that can be automated without expanding the surface of decisions that can be delegated.

The distinction shows up concretely in how AI is used in mature systems today, and it is more specific than most descriptions of the category suggest.

AI helps build and maintain the decision logic. The decision logic itself is deterministic. Machine learning selects, routes, and prioritizes. It does not alter the determination.

In practice, AI categorizes the government and industry sources a rule draws on, tracks changes to those sources over time, augments them with the organization's own data, and drafts candidate logic through a sequence of steps in which experts confirm accuracy at each one. The resulting rule is tested thoroughly before deployment, so its behavior is known before it touches a live decision. This speed and coverage would not be possible without AI. The traceability is what makes it deployable. The same discipline applies when the system learns from its own outcomes: learning can directly improve how cases are selected, routed, and prioritized. If AI proposes changing the logic that determines an outcome, experts review, test, and version before it reaches a live decision.

That is a meaningful architectural distinction, not a semantic one. A model that helps an expert construct and maintain a rule against changing policy is operating where its strengths apply and where its work can be reviewed before deployment. When a source changes, because a regulator or a specialty society publishes new guidance, that change must propagate through every rule that depends on it, while preserving the prior version so earlier decisions can still be reconstructed. A model that decides which cases warrant review makes a routing decision, consequential but not determinative. The same holds for the work around a complex review. A model can select which cases warrant expert attention, summarize the record, and pinpoint where two sources disagree, for example where a chart and a claim imply different diagnosis groupings. The reviewer still decides whether the grouping is wrong. The model has shortened the path to that judgment, not taken it. The same pattern applies in clinical and operational work. A model that flags a possible gap in documentation, or a patient who may be due for an intervention, surfaces a case for a clinician to consider, and the clinician decides what the record shows and what the patient needs. In none of these cases is the model deciding the outcome. Systems that blur these three into a single assertion about AI-driven decisions are describing something the regulatory environment does not permit and the operating model does not require.

Lyric

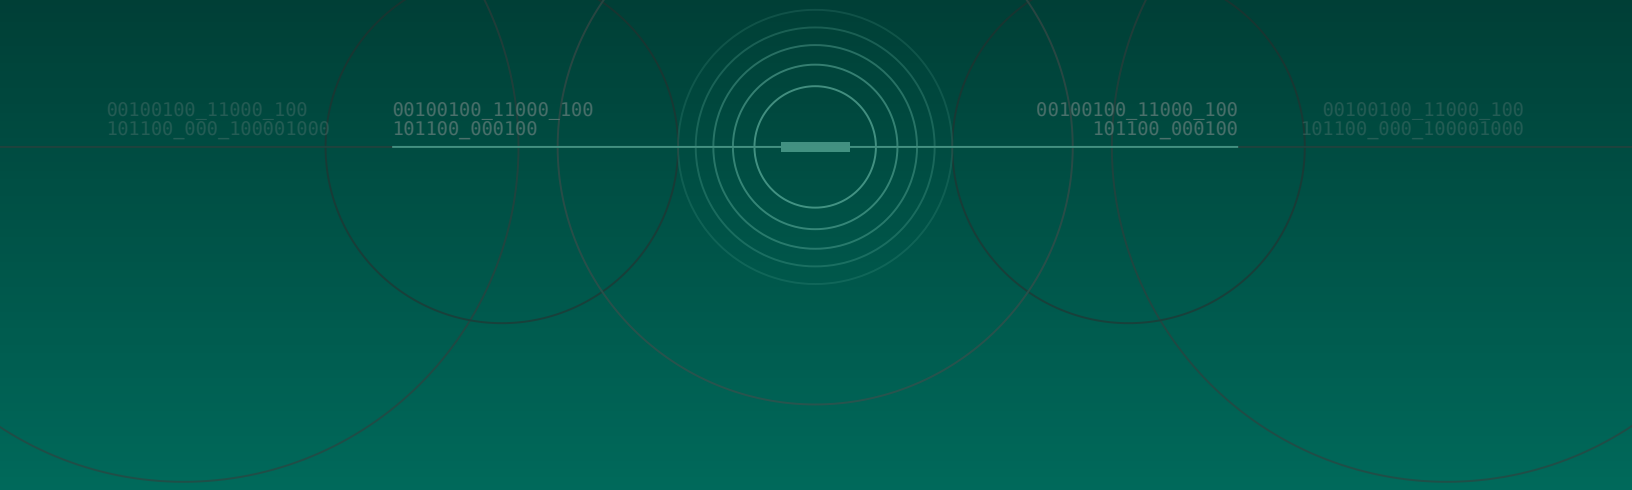
PART

03

The Standard

What Qualifies as Healthcare Decision Intelligence

Where the Category Proves Out First



What Qualifies as Healthcare Decision Intelligence

No universal working definition of healthcare decision intelligence exists. The term is being applied across a variety of conditions and capabilities. The standard includes three conditions that separate healthcare decision intelligence from the capabilities it is most often confused with: expertise built into the software; intelligence traceable at decision speed; and applied human judgment. Without all three, the system fails.

FIGURE 8

What qualifies as healthcare decision intelligence

The standard includes three conditions.

I

Expertise built into the software

Platform capabilities, policy, data, deterministic rules, AI, workflow intelligence, and domain expertise operate as one governed system.

II

Intelligence traceable at decision speed

Every recommendation traces to the policy, the data, and the logic behind it, at the version in effect when the decision was made.

III

Applied human judgment

Even when operations are automated, a person stays in the loop through the workflow.

Expertise Built into the Software

Platform capabilities, policy, data, deterministic rules, AI, workflow intelligence, and domain expertise operate as one governed system rather than separate tools connected after the fact.

Operating as one system also means operating on one body of data. The deterministic logic, the AI that supports it, and the evidence assembled for the decision-maker must draw on the same data, so that a recommendation and its lineage describe the same facts. A system whose rules run on one data set while its models run on another may still produce a correct answer, but it cannot show that the answer and its explanation came from the same place.

Where precision is required, AI grounds policy and regulatory inputs in deterministic, explainable results. Where precision is not required, it surfaces next-best actions and summarizes complex information. That division reflects the practical form of the distinction drawn above: the model does the work that benefits from breadth, and the logic that determines an outcome remains deterministic and reviewable.

The technology must encode clinical evidence, care guidelines, coding conventions, and the policy governing coverage and payment, and maintain them as the underlying evidence changes. A general-purpose model applied to a healthcare problem does not meet this condition, regardless of model capability, because a recommendation is only as good as the policy knowledge behind it and that knowledge changes constantly.

Intelligence Traceable at Decision Speed

The recommendation returns inside the live workflow, at production volume, when the decision is required, in real time where the workflow demands it. If intelligence arrives after the decision, that is reporting and recovery, not decision intelligence. The system must sustain that volume and latency under production conditions, not only in a controlled test; how it does so is an implementation choice, not a condition of the category.

Define criteria once and apply them consistently wherever they are read. A system that produces different answers to the same question depending on which workflow asked it is not operating from a single source of truth.

Every recommendation traces to the policy, the data, and the logic behind it, at the version in effect when the decision was made. Reconstruction must be possible months or years later, against the state of the world at the time rather than the state of the world today.

Accuracy without traceability does not qualify. In a regulated environment, a correct answer that cannot be explained carries most of the risk of a wrong one.

Applied Human Judgment

The organization operating the system configures rules, thresholds, and decision constraints; they are not fixed by the party supplying it. Changes are testable before deployment, so impact is understood before logic goes live.

Even when operations are automated, a person stays in the loop through the workflow, and the final determination always comes from a person.

That must be built in, not applied afterward. Human oversight positioned as a review step at the end of an automated process is oversight in name. By the time the case reaches a person, the logic has already run, the framing is already set, and the reviewer is approving rather than deciding. Governed by design means experts build and maintain the logic the system applies, cases requiring judgment arrive with the evidence and rationale assembled, and the rationale is available to the people affected by the decision, not only to the people making it.

A system that makes the determination itself is not a more advanced form of decision intelligence. It is operating outside the category.

Governance does not limit how capable a healthcare decision system can become. It is the prerequisite that allows a capable system to be deployed at all.

What the Standard Excludes

Applied honestly, this standard excludes much of what is currently described as decision intelligence.

Reporting and dashboards fail on timing when they are offered as the decision mechanism. Measuring outcomes is not excluded; it is the feedback loop the discipline depends on, and what a dashboard measures should return to the logic that made the decision. What fails the standard is analytics presented as the decision itself. Predictive models that score a

population without acting inside a workflow fail on the same condition. Automation that executes rules quickly but cannot identify which policy version fired fails on traceability. General-purpose AI applied to healthcare problems without encoded domain expertise fails on expertise. And fully autonomous determination fails the human judgment condition by design, not by shortfall.

FIGURE 9

What the standard excludes

OFFERED AS DECISION INTELLIGENCE	FAILS ON
Reporting and dashboards offered as the decision mechanism	**Timing**
Predictive models that score a population without acting inside a workflow	**Timing**
Automation that executes rules quickly but cannot identify which policy version fired	**Traceability**
General-purpose AI applied to healthcare problems without encoded domain expertise	**Expertise**
Fully autonomous determination	**Human judgment, by design**

This is a demanding standard, and it is meant to be. The point of defining a category is to make it possible to tell whether something belongs in it.

Where the Category Proves Out First

A standard is useful only if you can measure something against it. Within healthcare, payment accuracy is where that measurement is currently possible, for the same reasons credit and fraud moved first within financial services. It is the first proving ground, not the boundary of the category.

Claims carry standardized data. Reimbursement policy is substantially codified. Decisions occur at very high volume. Outcomes are measurable at the claim and line level. Intervention can occur inside adjudication, while the decision is still open.

Payment accuracy is therefore the clearest available view of how far healthcare decision intelligence has developed. Some of its functions meet this standard today. Others are approaching it. Several are further away than the market's language about AI in healthcare would suggest, and the gap between the two is the most useful thing a health plan can understand about its own operation.

The companion paper, The State of Healthcare Decision Intelligence in Payment Accuracy, applies this standard function by function across editing, data mining, coordination of benefits, record review, fraud, and subrogation.

The same conditions are forming across the rest of healthcare. Utilization management, care management, clinical documentation, and provider revenue cycle each involve high-volume, policy-governed decisions with measurable outcomes, and each will be held to the same standard as its data and policy mature.

Organizations that treat governance as a constraint to be minimized move slowly, because every expansion of scope reopens questions that were never settled. Organizations that embed it move faster, because the boundaries are established and the expansion does not require relitigating them. The second group will define the category.

Notes

1 M. W. Kirby and R. Capey, "The Air Defence of Great Britain, 1920–1940: An Operational Research Perspective," *Journal of the Operational Research Society* 48, no. 6 (June 1997): 555–568, especially 561–563, https://doi.org/10.1057/palgrave.jors.2600421; M. W. Kirby, *Operational Research in War and Peace: The British Experience from the 1930s to 1970* (London: Imperial College Press, 2003), 58–85, https://doi.org/10.1142/9781860949401_0003.

2 Ronald A. Howard, "Decision Analysis: Applied Decision Theory," in *Proceedings of the Fourth International Conference on Operational Research*, ed. David B. Hertz and Jacques Melese (New York: Wiley-Interscience, 1966), 55–71; reprinted in Ronald A. Howard and James E. Matheson, eds., *Readings on the Principles and Applications of Decision Analysis*, vol. 1 (Menlo Park, CA: Strategic Decisions Group, 1983).

3 G. Anthony Gorry and Michael S. Scott Morton, "A Framework for Management Information Systems," *Sloan Management Review* 13, no. 1 (Fall 1971): 55–70.

4 D. J. Power, "A Brief History of Decision Support Systems," DSSResources.com, version 4.1, accessed October 8, 2026, https://dssresources.com/history/dsshistory.html.

5 James Taylor and Neil Raden, *Smart Enough Systems: How to Deliver Competitive Advantage by Automating Hidden Decisions* (Upper Saddle River, NJ: Prentice Hall, 2007), 1–3; Object Management Group, *Decision Model and Notation (DMN)*, Version 1.0, OMG document formal/2015-09-01 (September 2015), https://www.omg.org/spec/DMN/1.0.

6 Lorien Pratt and Mark Zangari, "Overcoming the Decision Complexity Ceiling Through Design," version 1.1 (Quantellia, December 5, 2008), https://quantellia.com/Data/WP-OvercomingComplexity.pdf.

7 Lorien Pratt, *Link: How Decision Intelligence Connects Data, Actions, and Outcomes for a Better World* (Bingley, UK: Emerald Publishing, 2019), https://doi.org/10.1108/9781787696532.

8 Mastercard, "Mastercard Rolls Out Artificial Intelligence Across Its Global Network," press release, November 30, 2016, https://newsroom.mastercard.com/news/press-releases/mastercard-rolls-out-artificial-intelligence-across-its-global-network/.

9 Gartner, "Decision Intelligence," Gartner Glossary, accessed October 8, 2026, https://www.gartner.com/en/information-technology/glossary/decision-intelligence.

10 Gartner, *Magic Quadrant for Decision Intelligence Platforms*, January 26, 2026.

11 FICO, "FICO History," accessed October 8, 2026, https://www.fico.com/en/history; Federal Reserve Bank of Minneapolis, "Mortgage Automation Threat," December 1, 1996, https://www.minneapolisfed.org/article/1996/mortgage-automation-threat.

12 Freddie Mac, "Freddie Mac's Loan Prospector® Celebrates 20 Million Loans," press release, June 11, 2002, reproduced by the American Land Title Association, https://www.alta.org/news-and-publications/news/20020611-Freddie-Macs-Loan-Prospector-Celebrates-20-Million-Loans.

13 Visa, "Visa Prevents Approximately $25 Billion in Fraud Using AI," press release, June 17, 2019, https://usa.visa.com/about-visa/newsroom/press-releases.releaseId.16421.html.

14 Vincent A. Mabert, "The Early Road to Material Requirements Planning," *Journal of Operations Management* 25, no. 2 (March 2007): 346–356, https://doi.org/10.1016/j.jom.2006.04.002; Joseph Orlicky, *Material Requirements Planning: The New Way of Life in Production and Inventory Management* (New York: McGraw-Hill, 1975).

15 INFORMS, "ORION Delivers Success for UPS," *OR/MS Today*, March 2016, https://pubsonline.informs.org/do/10.1287/orms.2016.03.10/full/.

16 McKinsey & Company, "Succeeding in the AI Supply-Chain Revolution," April 2021, https://www.mckinsey.com/\~/media/mckinsey/industries/metals%20and%20mining/our%20insights/succeeding%20in%20the%20ai%20supply%20chain%20revolution/succeeding-in-the-ai-supply-chain-revolution.pdf.

17 Munich Re, *2022 Accelerated Underwriting Survey*, reported in Lisa Seeman and Dave Goehrke, "A Decade Into Accelerated Underwriting: The New Normal for Advanced Risk Selection," Munich Re Life US, May 18, 2023, https://www.munichre.com/us-life/en/insights/future-of-risk/2022-Accelerated-Underwriting-Survey.html.

18 Gen Re, *2024 Individual Life Accelerated Underwriting Survey*, November 2024, https://www.genre.com/us/knowledge/publications/2024/november/surveylhau24-en.

19 Mastercard, "Mastercard Rolls Out Artificial Intelligence Across Its Global Network," press release, November 30, 2016 (note 8 above); Mastercard, "Mastercard Supercharges Consumer Protection with Gen AI," press release, February 1, 2024, https://www.mastercard.com/us/en/news-and-trends/press/2024/february/mastercard-supercharges-consumer-protection-with-gen-ai.html.

20 Mastercard, "Mastercard Accelerates Card Fraud Detection with Generative AI Technology," press release, May 22, 2024, https://www.mastercard.com/us/en/news-and-trends/press/2024/may/mastercard-accelerates-card-fraud-detection-with-generative-ai-technology.html.

21 Mastercard, "Mastercard Leverages Its AI Capabilities to Fight Real-Time Payment Scams," press release, July 6, 2023, https://www.mastercard.com/global/en/news-and-trends/press/2023/july/mastercard-leverages-its-ai-capabilities-to-fight-real-time-payment-scams.html.

22 Alexandra Brintrup, Johnson Pak, David Ratiney, Tim Pearce, Pascal Wichmann, Philip Woodall, and Duncan McFarlane, "Supply Chain Data Analytics for Predicting Supplier Disruptions: A Case Study in Complex Asset Manufacturing," *International Journal of Production Research* 58, no. 11 (2020): 3330–3341, https://doi.org/10.1080/00207543.2019.1685705.

23 Peter Densen, "Challenges and Opportunities Facing Medical Education," *Transactions of the American Clinical and Climatological Association* 122 (2011): 48–58, https://pmc.ncbi.nlm.nih.gov/articles/PMC3116346/.

24 Cheyenne Ariana Erika Modina, Sandra Waugh Ruggles, Juliana Perl, and Josh Makower, "Current Common Procedural Terminology (CPT®) Coding Process Challenges: Impact on the HealthTech Innovation Ecosystem," Health Management, Policy and Innovation 9, no. 1 (2024), https://biodesign.stanford.edu/content/dam/sm/biodesign/documents/programs/policy-program/Current-CPT-Coding-Process-Challenges.pdf.

25 William H. Shrank, Teresa L. Rogstad, and Natasha Parekh, "Waste in the US Health Care System: Estimated Costs and Potential for Savings," *JAMA* 322, no. 15 (2019): 1501–1509, https://doi.org/10.1001/jama.2019.13978.

26 Centers for Medicare & Medicaid Services, "Fiscal Year 2025 Improper Payments Fact Sheet," January 15, 2026, https://www.cms.gov/newsroom/fact-sheets/fiscal-year-2025-improper-payments-fact-sheet. CMS notes that improper payments are not a measure of fraud.

27 Meghan Hufstader Gabriel, Chelsea Richwine, Catherine Strawley, Wesley Barker, and Jordan Everson, *Interoperable Exchange of Patient Health Information Among U.S. Hospitals: 2023*, ONC Data Brief No. 71 (Washington, DC: Office of the National Coordinator for Health Information Technology, May 2024), https://www.healthit.gov/data/data-briefs/interoperable-exchange-patient-health-information-among-us-hospitals-2023/.

28 42 C.F.R. § 422.566(d), https://www.ecfr.gov/current/title-42/chapter-IV/subchapter-B/part-422/subpart-M/section-422.566.

29 Centers for Medicare & Medicaid Services, "Frequently Asked Questions Related to Coverage Criteria and Utilization Management Requirements in CMS Final Rule (CMS-4201-F)," February 6, 2024, https://www.cms.gov/files/document/hpms-memo-faq-coverage-criteria-and-utilization-management-cms-4201-f-02-6-2024-pdf.pdf.

30 California Legislature, Senate Bill No. 1120, Chapter 879, Statutes of 2024 (Health Care Coverage: Utilization Review), approved September 28, 2024, effective January 1, 2025, https://leginfo.legislature.ca.gov/faces/billStatusClient.xhtml?bill_id=202320240SB1120.

31 National Association of Insurance Commissioners, *Model Bulletin: Use of Artificial Intelligence Systems by Insurers*, adopted December 4, 2023; National Association of Insurance Commissioners, *Implementation of NAIC Model Bulletin: Use of Artificial Intelligence Systems by Insurers*, status as of April 1, 2026, accessed October 8, 2026, https://content.naic.org/sites/default/files/cmte-h-big-data-artificial-intelligence-wg-map-ai-model-bulletin.pdf.

32 42 C.F.R. § 422.504(d) and (e)(4), https://www.ecfr.gov/current/title-42/chapter-IV/subchapter-B/part-422/subpart-K/section-422.504.

33 European Parliament and Council of the European Union, Regulation (EU) 2024/1689 (Artificial Intelligence Act), July 12, 2024, Annex III, para. 5(c), and Article 14(4)(d), https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng. Annex III, para. 5(c) applies specifically to risk assessment and pricing of natural persons in life and health insurance.

About Lyric

Lyric is the trusted leader in healthcare decision intelligence. Built on more than 35 years of proven results, Lyric combines responsible AI with clinical, payment, regulatory and policy expertise to deliver transparent, auditable insights in milliseconds across real-time claims workflows. Lyric augments human decision-making so health plans can make faster, more accurate payments while maintaining full control. Today, Lyric supports 200 million lives, with nine of the top 10 U.S. health plans relying on its platform to reduce waste, improve efficiency, and support accurate payments. To learn more, visit lyric.ai.

Contact

lyric.ai

Healthcare Intelligence. Smarter Every Day™

Authors: Yevgeny German, Chief Technology Officer · Brian Berkowitz, Chief Product Officer · Alex Ulyanov, SVP of Engineering & AI · November 2026

© 2026 Lyric. All rights reserved.