

HRN Group

The Miracle You're Asking For

Why Your RCM Team Can't Do What You're Expecting Them To Do

*The math behind underpaid claims, overwhelmed billers,
and why the solution isn't more staff*

A White Paper by HRN Group | May 2026
For Hospital CEOs, CFOs, and Healthcare Policy Leaders

Executive Summary

Your revenue cycle team isn't underperforming. They're doing something that shouldn't be possible - and you're measuring them as if it should be easy.

Consider what you're actually asking a billing manager at a rural hospital to do: review a denied claim, determine which of **nine different payer portals** it belongs to, remember whether that specific payer requires modifier documentation on the original submission or as a separate attachment, recall whether the appeal deadline is 30, 60, 63, or 180 days, cross-reference the denial code against that payer's specific resubmission rules, verify coordination of benefits, check eligibility status, and compose a clinically appropriate appeal letter - for one claim. Then do it again. And again. Hundreds of times.

This paper makes the case that the gap between what hospitals expect from revenue cycle management and what is humanly possible has grown beyond what staffing, training, or effort can close. The problem isn't your people. It's the math.

Section 1: The Job Description That Doesn't Exist

No hospital would post a job listing that reads: "Must simultaneously hold in working memory the billing rules, appeal deadlines, modifier requirements, and submission preferences of nine different payer organizations, across 74 financial class codes, 87 denial reason codes, 780 procedure codes, and 1,748 diagnosis codes. Must make zero errors."

Yet that is precisely what rural hospital revenue cycle teams do every day. Not because anyone designed it that way, but because the complexity of the payer landscape grew incrementally - one new MCO contract at a time, one policy update at a time, one portal change at a time - until the job became something no human being was built to do.

The Complexity No One Measured

At one rural Kansas hospital, we mapped every variable that goes into processing a single underpaid or denied claim. The number was 58.

Not 58 steps. Fifty-eight distinct data points that must be evaluated, cross-referenced, and acted upon correctly for a single claim to have the best chance of recovery. They include:

- **Payer identity** - not just 'UHC' or 'BCBS,' but which specific plan variant, which portal, which submission method, which mailing address
- **Denial code interpretation** - the same CARC code (e.g., CO-16) means different things and requires different responses depending on the payer
- **Regulatory context** - which OIG audit findings, CMS policy clarifications, or state administrative code sections support the hospital's position
- **Contract rate verification** - comparing what was paid against what should have been paid under the specific payer contract, including fee schedule lookups across hundreds of CPT codes
- **Coordination of benefits** - determining primary vs. secondary payer, verifying eligibility dates, checking for retroactive coverage activation
- **Deadline management** - each payer has different appeal windows, and missing one by a single day makes the claim permanently unrecoverable

A billing manager processing these claims isn't doing data entry. She's performing a simultaneous multi-variable optimization across regulatory, contractual, clinical, and financial dimensions - for every single claim on her desk.

The Kansas Revenue Lane Matrix

To understand why this is structurally impossible to manage manually, consider what we call the revenue lane matrix - the number of unique decision pathways that exist in a single state's underpaid claims environment.

In Kansas alone:

- **895** payer name variants mapped to canonical carriers
- **74** financial class codes determining routing logic
- **87** CARC denial reason codes requiring different responses
- **29** denial sub-classifications driving action strategy
- **780** CPT procedure codes with payer-specific rules
- **1,748** diagnosis codes affecting coverage determinations

That's over 27 million possible routing combinations for a single state. No billing team of any size was designed to navigate this.

This isn't a staffing problem. Doubling the team doubles the hands but doesn't reduce the complexity each pair of hands must manage. Every biller still faces the same 27 million pathways. You cannot solve an exponential problem with linear resources.

Section 2: The Arms Race You Didn't Sign Up For

There's a second dimension to this problem that most hospitals haven't fully reckoned with: the payers aren't standing still.

What the Payer Side Looks Like

Major payers - UnitedHealthcare, Anthem, Centene, and their subsidiaries - have invested hundreds of millions of dollars in AI-powered claims adjudication systems. These systems don't process claims the way a human reviewer does. They apply machine learning models trained on millions of claims to identify patterns, flag outliers, and apply denial logic at a speed and scale that no human team can match.

When your billing team submits a claim, it isn't reviewed by a person. It's processed by an algorithm that has been specifically optimized to identify reasons to deny, downcode, or underpay. The payer's AI is looking at the same 58 data points your biller is - except it's doing it simultaneously, across every claim, in milliseconds.

This creates a fundamental asymmetry. **The payer side uses AI to process and deny claims. The hospital side uses a person with a spreadsheet to appeal them.**

The Denial Factory

The result is predictable. Denial rates have increased industry-wide, not because hospitals are billing worse, but because payer adjudication systems have gotten better at finding reasons to deny. Every policy update, every new modifier requirement, every change to prepayment validation criteria is another input to the payer's denial algorithm.

Your biller learns about these changes when claims start getting denied. The payer's system knew about them the day they took effect. That timing gap - the payer knows the rules changed before the hospital does - is where a significant portion of recoverable revenue disappears.

At the rural Kansas hospital we analyzed, we found over \$700,000 in charges across 1,172 individual claims that were either denied, underpaid, or never appealed. These weren't billing errors. They were claims that fell into the gap between what the payer's system required and what the hospital's team had the capacity to track.

You don't bring a spreadsheet to an AI fight. But that's exactly what most hospitals are doing.

Section 3: The Compound Cost of Serial Processing

Human cognition is fundamentally serial when it comes to complex regulatory information. Cognitive load research - from Sweller's foundational work through modern healthcare decision-making studies - consistently demonstrates that decision quality degrades as the number of simultaneous variables increases.

What Serial Processing Costs

A skilled biller can hold perhaps 3-4 payer rule sets in working memory simultaneously. She knows that Sunflower's appeal deadline is 30 days and that UHC uses different modifier requirements. But she cannot simultaneously cross-reference those rules against 6,000 active claims, 45 policy manual sections, four different filing deadline calendars, and nine separate payer portals.

This isn't a criticism. It's a description of how human cognition works. The issue isn't competence - it's capacity. Rural revenue cycle teams are structurally under-resourced relative to the complexity they manage.

The result is that each individual limitation compounds:

- One missed appeal deadline on one claim: \$500
- The same miss pattern across 40 similar claims over 6 months: \$20,000
- Add the claims that were never identified as appealable in the first place: \$60,000
- Add the eligibility gaps, coordination of benefits errors, and contract rate underpayments: \$200,000+

These aren't hypothetical numbers. They're representative of what we consistently find when we analyze 12 months of remittance data from a single rural facility.

The Collection Rate Variance Problem

At one rural hospital, collection rates on Medicaid managed care claims varied between 9.6% and 14.1% - depending on which MCO was being billed. The variance wasn't explained by payer generosity or claim complexity. It was explained by which MCO's rules the billing team happened to know best.

The MCO they'd worked with longest had the highest collection rate. The MCO with the most recent policy changes had the lowest. The team's knowledge hadn't kept up with one payer's rule changes, and the revenue impact was measurable.

This is what it looks like when complexity exceeds human capacity. Not catastrophic failure - gradual, invisible leakage that compounds until it becomes material.

Section 4: What the Data Actually Shows

When we performed a comprehensive revenue analysis at a rural Kansas hospital - a Rural Emergency Hospital serving 65,000 people across 8 counties - we found something that reframed the entire conversation about their financial position.

The Discovery

Using 12 months of 835 EDI remittance files and encounter data, analyzed against state-specific Medicaid rules, payer contract rates, and regulatory requirements, the analysis revealed categories of recoverable revenue that had been structurally invisible:

Revenue Gap Category	Scale
Denied claims never appealed - across all portals	1,172 claims identified
Total charges in active recovery pipeline	\$700K+
Payer portals requiring simultaneous management	9 distinct portals
Payer name variants requiring canonical mapping	895 aliases
Data points evaluated per claim for recovery action	58 per claim

These numbers weren't hiding. They were sitting in the hospital's own billing system. The data had been there for months. What was missing was the capacity to read it all at once.

The Pattern Behind the Numbers

The critical finding wasn't the dollar amount. It was the pattern. The claims that fell through weren't random. They clustered around specific combinations of payer, denial code, and service type that the billing team didn't have the bandwidth to systematically track.

The same CO-16 denial from UHC requires a fundamentally different response than a CO-16 from Sunflower - different documentation, different submission method, different deadline, different appeal address. The billing team knew this in theory. But executing it consistently across hundreds of claims, from nine different payers, with different portals and different rules, is the miracle that hospitals keep asking for.

Section 5: Why the Answer Isn't What You Think

The conventional responses to revenue cycle performance gaps are hiring, training, and technology. All three miss the point.

Hiring Doesn't Scale the Problem

Adding staff adds hands, not intelligence. A second biller facing the same 27 million routing combinations has the same cognitive limitations as the first. You've doubled your labor cost without reducing the complexity each person manages. In a rural hospital operating at negative margins, the math doesn't work.

Training Doesn't Survive Policy Changes

You can train your team on every MCO's current rules. By the time the training is complete, three of those MCOs will have updated their policies. Payer rule changes are continuous, and each change creates a new gap between what your team knows and what the payer's system expects. Training is essential, but it's a treadmill, not a solution.

Enterprise AI Wasn't Built for This

The healthcare AI market is crowded with excellent products - AKASA, Waystar, CodaMetrix, and dozens of others. They are almost entirely irrelevant to the rural hospital Medicaid managed care problem, for four reasons:

- **Scale mismatch:** Built for health systems processing millions of claims. A 25-bed CAH processing 30,000 encounters doesn't generate enough volume for their models.
- **Integration burden:** Require EHR integration, API connections, and IT infrastructure that rural hospitals don't have. Implementation timelines of 6-12 months.
- **State-specificity gap:** Trained on national claims data. They don't know that Sunflower's L2 appeal deadline differs from UHC's, or that Kansas Administrative Code § 28-15-1040 governs timely filing differently than Nebraska Title 471.
- **Wrong output:** They produce dashboards and analytics. Rural hospital billing teams need claim-level worklists with specific action steps, sorted by deadline urgency.

The rural hospital problem doesn't need better analytics. It needs better intelligence - state-specific, payer-specific, and delivered in formats a 2-person billing team can act on Monday morning.

Section 6: A Different Kind of Problem Requires a Different Kind of Solution

The gap between what hospitals need from revenue recovery and what any human team can deliver is structural. It will not close with effort, overtime, or determination. It can only be closed with intelligence that operates at the same scale and specificity as the problem itself.

That means:

- **Processing all 58 data points simultaneously**, across every claim, every time - not serially, not from memory, not approximately
- **Mapping every claim to its exact position** in the 27-million-combination routing matrix - the right payer, the right portal, the right deadline, the right appeal format
- **Generating submission-ready appeal packets** - not reports that require interpretation, but documents that can be sent to the payer with the regulatory citations, clinical arguments, and contract references already composed
- **Updating continuously** as payer policies change, denial patterns shift, and new regulatory guidance is issued - without requiring the billing team to track those changes themselves

This is what purpose-built revenue intelligence looks like. Not a dashboard. Not a report. A system that does the work your team can't humanly do, delivered in the format they already use, requiring no new software, no IT project, and no change to how they work.

We don't change how your team works. We change what your team knows.

What This Means for Hospital Leadership

If you're a CEO or CFO reading this, the question isn't whether your revenue cycle team is doing a good job. They almost certainly are - given what you're asking them to manage with the resources they have.

The question is whether you've measured the gap between what they can humanly accomplish and what the complexity of the payer landscape actually requires. In our experience, that gap ranges from 5-15% of net revenue at rural hospitals. Most of it is recoverable. Almost none of it is visible without the right kind of analysis.

The hospitals that close this gap don't do it by working harder. They do it by acknowledging that the problem has outgrown the solution, and that a new kind of intelligence is required to match it.

Ready to see what the gap looks like in your data?

HRN Group offers a no-cost initial assessment on 3-6 months of your remittance data to show what's recoverable. No software to install. No IT project. No commitment beyond a 30-minute conversation.

Contact: David Thorne | david@highvaluechange.com | hrn.highvaluechange.com