

After after coding: imagining risk adjustment when accuracy breaks the business model

JAN 14, 2026 • PAID



Share

Abstract

This essay speculates on a future version of risk adjustment in which overcoding is structurally prohibited and undercoding is no longer a meaningful source of underpayment. The premise is not that compliance finally wins, but that measurement improves enough to eliminate diagnosis capture as a primary economic lever. In this world, value based care does not become easier. It becomes more exposed. The essay explores what kinds of regulation could realistically produce this state, what business models would survive it, how capital and care would likely be misallocated during transition, and why the end state is less morally satisfying but more economically honest than the current system.

Key themes

- Risk adjustment as a measurement problem, not a coding problem
- Regulatory designs that reward symmetry and evidence rather than documentation effort
- Business models that emerge when diagnosis arbitrage disappears
- Capital allocation failures inside VBC once RAF growth flattens
- Why accuracy creates fragility before it creates efficiency

Table of Contents

1. The End of Coding as Strategy
2. When Everyone Is Accurately Sick
3. Separating Discovery From Monetization
4. Turning Risk Scores Into Regulated Systems
5. Redesigning the HCC Instead of Policing It
6. Provider Behavior Without Documentation Leverage
7. What Breaks Inside VBC First
8. Where Founders Should and Should Not Build
9. The Political Economy of Accuracy
10. Why This Is Still Worth Doing

The End of Coding as Strategy

Risk adjustment has spent the last twenty years pretending it was about fairness it was really about feasibility. Paying capitated plans without some proxy for disease burden was never going to work, so diagnosis codes stepped in as a stand-in for expected cost. That decision solved one problem and created another. Once diagnosis became a payment primitive, capturing them became strategy.

The system did not drift into this. It behaved exactly as designed. When payment is tied to documented disease states, organizations will invest in finding, documenting and defending those states. They will build workforces, workflows, analytics, incentives, and software to do it at scale. They will focus on diagnoses with elast

thresholds, ambiguous criteria, or reflexive logic. They will get very good at adding and much worse at removing.

The pathology people now complain about is not abuse in the cartoon sense. It is industrial optimization in a system where documentation is monetized. This may be because banning overcoding without redesigning the substrate does not fix the problem. It just relocates it. The moment one set of diagnoses is excluded or downweighted, attention shifts to the next most profitable surface. The model tells the market what to care about, and the market listens carefully.

The future imagined here starts when that lesson stops being economically useful. Not because everyone becomes ethical, but because measurement improves enough that diagnosis capture loses its leverage. In that world, coding is no longer a growth strategy. It becomes table stakes at best and liability at worst.

When Everyone Is Accurately Sick

One of the comforting myths in health policy is that better data automatically makes payment fairer. In reality, better data mostly makes mispricing harder to hide. When undercoding is no longer a problem, average disease burden goes up. Not because people suddenly get sicker, but because the system finally notices what was already there.

This creates a distribution problem. Historically underdocumented populations drift upward toward their actuarial reality. Populations that benefited from aggressive documentation drift downward. Risk scores converge. Variance collapses. From a modeling perspective, this looks like success. From an operational perspective, it creates confusion.

When everyone is accurately sick, nobody is obviously sicker. Risk stratification becomes less discriminative. Care management programs that relied on diagnosis-derived tiers find half their population clustered in the same band. The instinctive response is to spread resources thinner. That feels safer than making sharper cuts; it is also how capital gets misallocated quietly.

This is where value based care starts to wobble. Many care models assumed that improved documentation would align revenue with cost. What they discover instead is that revenue aligns with predicted cost, not realized cost. Prediction error becomes the dominant source of variance. The documentation subsidy that quietly covered operational inefficiency disappears. Care models have to stand on their own economics, and many cannot.

Separating Discovery From Monetization

The most productive regulatory move in this future is not banning diagnosis capability but splitting it into two explicitly different phases. Discovery and monetization should not be the same act.

Discovery is cheap. It should be continuous, aggressive, and encouraged. Plans and providers already infer disease burden from medications, labs, utilization, and occurrence patterns. They already generate suspect lists. The problem is not finding conditions. The problem is paying for them too easily.

Monetization, by contrast, should be gated. It should be probabilistic. It should be evidence-weighted. Instead of asking whether a patient has a condition, payment systems should ask how confident the system is that the condition meaningfully contributes to expected cost, and what evidence supports that belief.

This changes incentives immediately. A one-time assessment can surface a suspicion but it cannot drive full payment without corroboration. Chart reviews matter, but only if they align with longitudinal signals. Diagnoses with fuzzy thresholds still exist but their contribution is bounded by uncertainty rather than amplified by ambiguity.

Most importantly, this design removes the asymmetry that dominates today. Adding becomes easy. Paying becomes hard. Removing becomes just as valuable as adding. Saying no becomes a reimbursable act.

Turning Risk Scores Into Regulated Systems

Today risk scores are treated as outputs. Regulators inspect what comes out and about whether it should have. That approach fails at scale because modern risk adjustment is optimized at the process level, not the code level.

A future-proof regulatory framework treats the risk scoring pipeline itself as a regulated object. Plans would be required to maintain documented internal control over how risk is inferred, reconciled, corrected, and submitted. Not vague compliance programs, but named controls with measurable performance.

Symmetry becomes mandatory. If a system looks for unsupported additions, it must look just as hard for unsupported deletions. Lineage becomes explicit. Every payment relevant condition carries an evidence trail showing where it came from, what supported it, and how contradictory signals were resolved.

Accuracy is rewarded not with praise, but with liquidity. Plans with low variance, reversal rates, and balanced correction get faster reconciliation and fewer holdbacks. Plans that drift get haircuts automatically. This is boring, operational pressure, and is exactly why it works.

Redesigning the HCC Instead of Policing It

Much of the current dysfunction comes from treating HCCs as static buckets triggered by codes. That architecture invites gaming because codes are cheap and thresholds are malleable.

A more durable design treats HCCs as composite burden constructs. Payment is activated not by a label, but by a pattern. Medication persistence, lab trends, prior admissions, functional decline. Diagnosis becomes one signal among many, not the master key.

This also shortens the arbitrage window. When weights are tied dynamically to observed utilization patterns rather than historical averages frozen years ago, the

payoff from strategic targeting shrinks. Innovation shifts away from finding the elastic diagnosis and toward improving measurement fidelity.

Provider Behavior Without Documentation Leverage

When documentation stops driving revenue directly, provider behavior changes ways that are both healthy and uncomfortable. The incentive to chase specificity fades. The incentive to stabilize patients rises.

Care becomes smoother rather than simply cheaper. Avoiding volatility matters. Avoiding unexplained spikes matters. That is good medicine most of the time.

But providers also lose agency. They no longer control how their patients are translated into payment through documentation effort. Algorithms and evidence engines take that role. Expect backlash framed as transparency and autonomy concerns. Some of that is real. Much of it is economic.

Compensation models have to adjust. Paying clinicians per condition makes no sense when conditions are no longer currency. Paying for evidence reconciliation, care continuity, and closure of unsupported conditions starts to matter. In a symmetric system, saying no is just as valuable as saying yes.

What Breaks Inside VBC First

The first thing to break is the quiet assumption of year over year RAF growth. Most VBC contracts depend on it implicitly. Flatten that curve and margins evaporate.

Benchmarks do not reprice quickly. Shared savings pools shrink. Downside risk triggers more often. The response cycle is predictable. Renegotiation, exit, consolidation.

The second thing to break is care management focus. When stratification becomes less sharp, resources get diluted. High need patients still exist, but they are hard

isolate. The result is average care for many instead of excellent care for the few who need it most.

Where Founders Should and Should Not Build

Anything that depends on discovering new documentation opportunity is living borrowed time. The future belongs to boring infrastructure.

Risk controllers. Evidence engines. Reconciliation platforms. Audit readiness systems. Products that reduce variance rather than inflate point estimates.

There is also room for services that look unscalable until regulation forces standardization. Risk integrity advisory. Pipeline certification. Evidence protocol design.

The best positioning avoids the language of optimization entirely. Sell predictability. Sell confidence intervals. Sell fewer surprises.

The Political Economy of Accuracy

There is no version of this future where everyone is happy. Plans lose leverage. Providers lose control. Vendors lose categories. CMS inherits fragility.

Accuracy concentrates pain. It exposes mispricing that documentation once hid. It reallocates capital toward populations that were previously subsidized by opacity. That costs money, and someone has to own that politically.

Incrementalism fails here. Partial reform invites arbitrage. Hybrid models create loopholes. The only stable equilibrium is one where accuracy itself is monetized and inaccuracy is operationally expensive.

Why This Is Still Worth Doing

The end state is not perfect. Measurement will always lag reality. Models will always be wrong at the edges.

But a world where overcoding is prohibited and undercoding is irrelevant is still better than the current one. Not because it is morally cleaner, but because it forces honesty. Care models have to justify themselves on outcomes and utilization, not narrative success. Capital allocation becomes explicit instead of accidental.

Risk adjustment stops being a debate about intent and becomes a discipline of measurement. That shift is uncomfortable. It slows some kinds of innovation. It accelerates others.

For health tech, the radical outcome is not a new winner, but a quieter market with fewer miracles and fewer scandals. Accuracy does not feel like progress at first. It feels like constraint. In the long run, it is the only thing that makes the rest of the system legible.



2 Likes • 1 Restack

← Previous

Next

Discussion about this post

Comments

Restacks



Write a comment...