§ f.03 · field note

Two practices, identical financials, different deals.

A note on behavioral patterns in physician practice claims data and what standard diligence methodology cannot see.

H. Chaudhry · Crescent Intelligence
crescentintel.com/research/two-practices-identical-financials

9 May 2026
3418 words

summary

Two physician practices with identical financial profiles — same revenue, same EBITDA, same payer mix, same provider count — can have materially different deal risk profiles. The difference is not visible in the trailing twelve-month financials that anchor standard diligence. It is visible in the structure of the behavioral residual: the patterned gap between what each practice should have been paid under contract and what each was actually paid. The residual contains three kinds of information that the standard methodology cannot extract because the standard methodology aggregates rather than decomposes. First, payer-specific concentration: the residual is heavily concentrated in one or two specific payers in some practices, dispersed across many in others. Second, procedure-specific patterns: the residual is concentrated in specific high-volume procedures with identifiable mechanism, or dispersed across many procedures with no clear pattern. Third, temporal drift: the residual is widening or narrowing over time, reflecting active changes in payer behavior. Each of these patterns has direct implications for deal pricing, escrow structure, post-close risk, and exit timing. None of them is visible to a diligence methodology that examines aggregate revenue and aggregate variance. The corrections are computable for any deal where the target's historical claims data is available. This note develops the analytical framework, demonstrates the divergence on a worked example, and traces the implications for deal mechanics.

1 the central claim

Standard physician practice diligence examines aggregate financials. Net patient revenue, EBITDA, payer mix at the level of percentage of revenue, denial rate as a single number, days in accounts receivable as a single number. These metrics are sufficient to characterize the practice at the level of summary statistics. They are not sufficient to characterize the practice at the level of structural risk.

Two practices with identical summary statistics can have very different structural risk profiles. The difference lives in the patterned gap between what each practice should have been paid under its contracts and what each was actually paid. We have argued elsewhere that this gap — the behavioral residual B in the decomposition $R = C + B$ — is not random. It contains patterns. The patterns are concentrated by payer, by procedure, and by time.[1]

The patterns are computable for any practice where claims data is available. They are also invisible to any methodology that examines only the aggregate.

The implications for deal underwriting are direct. A practice whose residual is concentrated in a single payer is structurally sensitive to that payer's behavior changing — a contract renegotiation, a policy update, a switch to a different plan structure can shift the residual materially. A practice whose residual is dispersed across the payer mix is structurally more stable. The two practices may have identical historical

[1] marginalia

See § M.01, On the deterministic decomposition of healthcare payment. *The decomposition $R = C + B$ establishes that realized revenue is the sum of a calculable component (C, the contractually owed amount) and a residual component (B, the behavioral variance). The structure of B is the subject of this note.*

performance. They have different forward risk profiles. The standard methodology cannot distinguish them; the deterministic methodology can.

This note develops the analytical framework. Section 2 describes what the behavioral residual is and how the standard methodology treats it. Section 3 develops the three patterns the residual contains. Section 4 describes the deterministic methodology. Section 5 demonstrates the divergence on a worked example. Section 6 traces the implications for deal mechanics.

2 what behavioral residual is, and how standard methodology treats it

For any physician practice, the realized payment on a given claim is the sum of two things. The contractually owed amount under the applicable fee schedule, modifiers, place-of-service rules, and edit logic — this is C, and it is calculable exactly from the claim and the contract. And the difference between what was paid and what was owed — this is B, and it captures the totality of behavioral influence on the payment. Downcoding, denial, prior authorization friction, contract noncompliance, payer adjudication error, slow payment that affects cash collection — all of it lives in B.

The decomposition is exact. Aggregated across the practice's full claims history, R (realized revenue) equals C plus B by construction. The interesting analytical question is not whether B is positive or negative — that is determined by the practice's specific operational performance and payer behavior. The interesting question is the structure of B. Where it concentrates, in which procedures, with which payers, and how it changes over time.

Standard diligence methodology, applied to acquired practice claims data, does not perform this decomposition. The methodology aggregates. It examines net collection rate as a single number — the share of expected revenue actually collected, computed at the aggregate level. It examines denial rate as a single number — first-pass denial across all payers, all procedures, all dates of service. It examines days in accounts receivable as a single number — average days to cash across the practice's full revenue base.

Each of these aggregate metrics has a single value that summarizes the practice. The QofE report cites the value, benchmarks it against industry standards, flags it if it falls outside acceptable ranges. The methodology treats deviation from benchmark as a binary signal — concerning or not concerning — without examining the underlying structure of the deviation.

This treatment loses information. Two practices with the same aggregate net collection rate can have very different distributions of underpayment across payers, procedures, and time. Two practices with the same first-pass denial rate can have very different denial concentration profiles. Two practices with the same days in AR can have very different payer-specific cash flow dynamics. The aggregate numbers are identical. The underlying structures, which determine forward risk, are not.

The aggregation problem is structural to the methodology, not a deficiency in any individual diligence firm's execution. The diligence industry built its methods at a time when claim-level computation under both contracts was infeasible — the same data, computation, and rule-encoding constraints we have addressed in earlier work.[2]

The constraints have lifted; the methods have not updated. The result is that diligence continues to characterize practices at the level of aggregate summary statistics, even though the structural information that matters for deal risk lives one level deeper.

3 the three patterns

The behavioral residual contains three kinds of structure. Each is computable from the claims data once the decomposition is performed. Each is directly relevant to deal risk. Each is invisible to aggregate analysis.

Pattern one: payer-specific concentration.

Denial rates vary substantially across payers. In ACA Marketplace data for 2023, UnitedHealth Group denied 33% of in-network claims, Anthem 23%, Aetna 22%, Cigna 21%, and Kaiser Permanente 6%.[3]

The range is meaningful — UHC's denial rate is 5x Kaiser's, and the major commercial payers vary by 12 percentage points among themselves. The pattern is not specific to denial rates. Underpayment rates, downcoding aggressiveness, prior authorization friction, and days-to-cash all vary substantially across payers. Different payers have systematically different behavioral profiles.

The practice's exposure to specific payers is the relevant variable. Two practices with identical payer mix at the percentage-of-revenue level can have very different residual concentration profiles. Practice A has 40% commercial revenue split as 25% UHC, 10% Anthem, 5% other. Practice B has 40% commercial revenue split as 10% UHC, 10% Anthem, 10% Cigna, 10% Kaiser. The aggregate denial rate, computed across both practices, can land at the same number. The structural risk is very different. Practice A is sensitive to UHC behavior — a UHC policy update, a contract renegotiation, a denial-rate change can shift the residual materially. Practice B is dispersed across payers with different behavioral profiles, so no single payer's behavior change has comparable impact.

The pattern is computable claim by claim. For each payer, the analysis identifies the realized payment, the contractually owed payment, the gap, and the structure of the gap (denial, downcoding, slow payment, contract noncompliance). The result is a payer-specific residual profile that distinguishes payer-concentrated practices from payer-dispersed practices at a level of granularity the aggregate denial rate cannot reach.

Pattern two: procedure-specific patterns.

Payer behavior is not uniform across procedures. Specific high-volume codes attract systematic downcoding — the AMA's payer E&M downcoding resource specifically identifies CPT 99204, 99205, 99214, 99215, and modifier-25 claims as the most commonly downcoded.[4]

[2] marginalia

See § M.02, On the obsolescence of sampling methodology in healthcare diligence. *The structural reasons that diligence methodology continues to use aggregate analysis instead of decomposed analysis are the same reasons that sampling methodology persists despite the lifting of its original constraints.*

[3] marginalia

Centers for Medicare & Medicaid Services Health Insurance Marketplace Quality Rating System, 2023 plan year data; compiled across reporting insurers in seven or more states.

[4] marginalia

American Medical Association, "Payer Evaluation and Management (E/M) Downcoding Programs: What You Need to Know," 2024.

Specific procedures attract concentrated prior authorization friction — UnitedHealthcare's prior authorization list has been modified at least eight times in 2024–2026 across cardiology, radiology, spine surgery, arthroplasty, and gastroenterology codes.[5]

Specific procedures attract concentrated underpayment — orthopedic shoulder arthroscopy codes 29827 versus 29822 differ by over $800 per case, and payers enforce documentation hierarchies that drop higher-specificity codes to lower-specificity codes algorithmically when documentation is incomplete.[6]

A practice whose residual is concentrated in a small number of high-volume procedures has very different forward risk than a practice whose residual is dispersed across many procedures. Concentrated residual is often actionable — once identified, the practice can address the specific issue through coding adjustments, documentation improvements, or contract renegotiation on the affected procedures. Dispersed residual is harder to address and represents structural drag rather than addressable drag.

The pattern matters for deal pricing in two distinct ways. The actionability of the residual determines the realistic upside from post-close operational improvements. The concentration of the residual determines the structural exposure to procedure-specific payer behavior changes. Aggregate analysis sees neither.

Pattern three: temporal drift.

Payer behavior changes over time. The most dramatic recent example: Medicare Advantage denial rates rose 4.8% from 2023 to 2024 alone — outpacing every other payer category in MDaudit's network analysis of 1.2 million providers.[7]

The average denied amount per MA claim rose 22.4% to approximately $1,000 in 2025. KFF analysis of MA prior authorization data found that insurers made nearly 53 million PA determinations in 2024, with 80.7% of appealed denials overturned but less than 12% of denials ever challenged.[8]

The MA behavioral shift is large enough that a practice's economics under the 2022 MA regime are not a reliable guide to its economics under the 2024–2026 regime.

The pattern is not specific to MA. UnitedHealthcare's prior authorization requirements have undergone substantive changes at least eight times in 2024–2026 across cardiology, radiology, spine surgery, and other specialty codes. Cigna and Aetna implemented automated E&M downcoding policies in 2025, triggering a formal AMA resolution opposing the practice and an AAFP letter to DOJ, FTC, and CMS urging investigation.[9]

BCBS Massachusetts implemented expanded claims review for high-complexity E&M overcoding effective November 3, 2025.[10]

Initial commercial claim denial rates climbed from 10.2% to 11.8% across the industry in the same period.[11]

The implication for deal underwriting: trailing twelve-month financials capture the practice's economics under recent payer behavior. They do not capture whether that behavior has been stable or actively deteriorating. A practice whose residual has been widening over the trailing 24 months under increasing payer aggressiveness

[5] marginalia

UnitedHealthcare provider portal prior authorization bulletin history, 2024–2026; available at uhcprovider.com.

[6] marginalia

Industry data on orthopedic shoulder arthroscopy reimbursement structure; documented at the CPT code level in payer fee schedules.

[7] marginalia

MDaudit, 2024 network analysis covering 1.2 million providers.

[8] marginalia

Kaiser Family Foundation analysis of CMS Medicare Advantage data, 2024.

[9] marginalia

American Academy of Family Physicians letter to the Department of Justice, Federal Trade Commission, and Centers for Medicare & Medicaid Services, November 21, 2025.

[10] marginalia

Blue Cross Blue Shield of Massachusetts provider notice, November 2025.

[11] marginalia

has a different forward trajectory than a practice whose residual has been stable or narrowing under decreasing payer aggressiveness. The two practices can have identical TTM financials and very different forward run-rates.

The temporal pattern is computable from the historical claims data once the residual is decomposed. The aggregate TTM financials do not contain it. The standard methodology, which uses TTM as the baseline for forward projection, cannot see whether the projection should be increased, held constant, or decreased based on the trajectory of the underlying behavioral component.

4 the analytical methodology

The deterministic methodology decomposes the realized payment on each claim into the contractually owed amount and the residual. The decomposition is mechanical once the apparatus exists — a rule system that encodes the contract's fee schedule, modifiers, edits, and place-of-service logic; a claim record; and the realized payment from the remittance.[12]

For each historical claim, the analysis computes C exactly, observes R from the remittance, and derives B by subtraction. Aggregated across the practice's full claims history, the decomposition produces the structural representation of the behavioral residual.

The structural representation supports three kinds of analysis the aggregate cannot.

For payer concentration, the analysis computes the residual at the payer level. For each payer with which the practice has a contractual relationship, the analysis produces a payer-specific residual profile — the per-claim gap, the aggregate gap, the structure of the gap (denial, downcoding, slow payment, contract noncompliance). The practice's exposure to specific payers becomes visible at a level the aggregate denial rate cannot reach.

For procedure concentration, the analysis computes the residual at the CPT-code level. For each code in the practice's high-volume mix, the analysis produces a procedure-specific residual profile. Codes where the residual is concentrated become visible. Codes where the residual is consistent with industry norms become visible. The practice's procedure-specific risk and addressable opportunity become visible.

For temporal drift, the analysis computes the residual at the time-series level. The per-month behavioral residual, decomposed by payer and procedure, produces the trajectory of the practice's behavioral economics over time. Drift is detectable as systematic change in the residual against stable C — which can only be observed if C is computed exactly, because aggregate analysis cannot distinguish whether changes in R reflect changes in C, changes in B, or both.

The methodology generalizes across specialties. The three patterns appear in any physician practice with sufficient claims volume to support the analysis. The specific payers, procedures, and temporal dynamics vary by specialty and geography. The analytical structure is the same.

[12] marginalia

The computational requirements are described in detail in § M.01. The methodology assumes a rule system capable of computing C at the resolution required for exact derivation, the acquired practice's historical claims data, and the relevant contracts. All three are available in any competent diligence engagement.

5 a worked example

We work through two synthetic ortho practices. Both have identical headline financials — $50M net patient revenue, 22% EBITDA margin ($11M), 28 physicians, payer mix of 35% Medicare, 40% commercial, 25% other. Both have identical aggregate diligence metrics: net collection rate of 94%, first-pass denial rate of 8%, days in AR of 36. The standard QofE report would produce essentially identical risk assessments for these two practices.

The deterministic decomposition reveals two structurally different deals.

table 1 · behavioral residual decomposition, two synthetic ortho practices

DIMENSION	PRACTICE A	PRACTICE B
Net patient revenue	$50M	$50M
EBITDA	$11M (22%)	$11M (22%)
Aggregate net collection rate	94%	94%
Aggregate first-pass denial rate	8%	8%
Aggregate days in AR	36	36
DECOMPOSED BEHAVIORAL RESIDUAL		
UnitedHealth share of commercial book	60%	25%
UHC-specific denial rate	18%	12%
Residual concentration in top 1 payer	67%	24%
Procedure concentration of residual (top 3 codes share)	71%	32%
Trailing 24-mo residual trajectory	+14% (widening)	–3% (stable)
MA exposure (% of total revenue)	18%	6%
MA-specific days-to-cash	41	32

Synthetic example for illustration. Numbers chosen to reflect realistic ranges observed across the literature; not drawn from any specific deal.

Practice A's residual is concentrated. Sixty percent of its commercial book is UHC. UHC's denial rate against this practice is 18%, materially above the industry average and above the practice's overall first-pass denial rate of 8% — meaning that the 8% aggregate denial rate is masking a much higher payer-specific denial rate balanced by lower denial rates elsewhere in the book. The residual is also procedure-concentrated; 71% of the residual variance is generated by the practice's three highest-volume procedures (which in an ortho practice are typically TKA, THA, and arthroscopic procedures). The trailing 24-month residual trajectory shows the residual widening

14% — payer behavior against this specific practice has been actively deteriorating over the trailing two years.

Practice B's residual is dispersed. UHC accounts for only 25% of the commercial book, and the UHC denial rate of 12% is closer to the practice's aggregate. The residual is dispersed across procedures (32% concentration in top three codes versus 71% for Practice A). The trailing 24-month trajectory shows the residual stable or marginally narrowing. The practice has lower MA exposure (6% versus 18%), and its MA-specific days-to-cash is materially better.

The two practices have identical headline metrics. Their structural risk profiles are very different.

Implications for Practice A:

The deal is structurally exposed to UHC behavior. Any change in UHC's policies, prior authorization protocols, or contract terms affects 60% of the commercial book directly. The 18% UHC-specific denial rate is materially above industry norms, which suggests that the practice is either subject to UHC prepayment review (a payer-specific risk factor) or has process issues specific to the UHC workflow. The widening 24-month trajectory indicates that whatever is producing the elevated UHC residual has been accelerating — there is no reason to expect TTM financials to project forward without further deterioration. The 18% MA exposure with 41-day days-to-cash creates structural working capital drag that the EBITDA margin of 22% does not capture.

A buyer underwriting Practice A on TTM financials projects steady forward EBITDA at $11M. A buyer underwriting Practice A on decomposed financials projects EBITDA declining 2–4% per year under continued UHC behavior trajectory, with potential single-event shocks if UHC introduces specific changes affecting the practice's procedural mix. The forward run-rate at exit (year 5) under the decomposed analysis is approximately $9–10M, not $11M.

Implications for Practice B:

The deal is dispersed across payers and procedures. The aggregate metrics reasonably represent the practice's structural performance. Forward TTM projection is reasonable. The 6% MA exposure with 32-day days-to-cash creates manageable working capital characteristics.

A buyer underwriting Practice B on TTM financials projects steady forward EBITDA at $11M. A buyer underwriting Practice B on decomposed financials projects approximately the same.

The deal-pricing implication:

At identical headline financials, the two practices warrant materially different deal prices. Applied to a standard 10x EBITDA entry multiple, Practice A is worth approximately $95M (reflecting forward EBITDA degradation to $9.5M, capitalized at the same multiple). Practice B is worth approximately $110M. The standard methodology, examining only aggregate metrics, prices both at $110M.

The mispricing on Practice A under the standard methodology is approximately 15%. Across a portfolio of similar deals where some are Practice A-like and some are

Practice B-like, the standard methodology produces systematic overpayment on the Practice A-like deals and accurate pricing on the Practice B-like deals. The buyer who computes the decomposition correctly distinguishes the two and adjusts pricing accordingly.

6 implications for deal mechanics

The decomposition has direct implications for how deals should be structured beyond purchase price.

Escrow structure. Standard escrow provisions hold back a percentage of purchase price against representations and warranties, with release tied to specific time periods. The decomposed analysis allows escrow to be sized and released against specific named risks — the UHC-concentrated practice should have escrow sized against UHC behavioral risk and released against specific UHC migration milestones; the dispersed practice does not require the same escrow structure. Standard escrow practice does not vary by risk profile because the standard methodology cannot identify the risk profile.

Earnout structure. Where earnouts are used, they should be tied to specific synergy realization milestones — UHC contract migration completed by month 12, MA payment timeline improved to platform benchmark by month 9. The decomposed analysis produces the specific milestones; aggregate analysis produces only general goals that are difficult to enforce.

Representations and warranties. The deal should include payer-specific reps covering the historical UHC behavior, the trailing 24-month trajectory, and the specific procedure-level patterns identified in diligence. Standard reps cover aggregate performance ("net collections of X%") and miss the structural risks that drive forward economics.

Post-close monitoring. The decomposed analysis identifies the specific patterns to monitor post-close. Practice A's monitoring program should track UHC-specific denial rate by month, UHC-specific underpayment rate on top three procedure codes, and the 24-month residual trajectory. Practice B's monitoring program can use less granular tracking. The standard methodology produces general monitoring KPIs that miss the practice-specific risk profile.

Exit timing and exit story. The decomposed analysis informs when the asset is ready to sell and what the exit story should be. Practice A's exit story should address the UHC concentration and the mitigation actions taken to reduce it; Practice B's exit story does not require the same treatment. The standard methodology produces exit stories that mirror entry stories regardless of what changed during the hold period.

The corrections are computable for any deal where the target's historical claims data is available. The data is acquired in any competent diligence engagement. The decomposition is engineering work, not research work. The buyer who performs it sees deals the standard methodology cannot distinguish. The buyer who does not performs the standard analysis and prices both practices at the same number.

The cost of getting this wrong is structural to the methodology. Across a wave of physician practice rollups currently in market — with first-cycle PE exits accelerating, second-bite transactions building, and strategic acquirers entering the same competitive set — the aggregate mispricing across the sector from inability to decompose behavioral residual is substantial. The corrections are available to any buyer willing to compute the decomposition under the target's actual claim history.

see also

§ M.01 On the deterministic decomposition of healthcare payment.
§ M.03 On the architecture of decomposable targets.
§ F.01 MPFS 2026: the headline narrative on ortho/spine reimbursement is wrong.
§ F.02 Contract migration synergy is computable. Most deal models don't compute it.