

Twenty Five Times the Surgeon's Pay: How Surgical Assistants Are Using No Surprises Act Arbitration Loopholes to Out Earn the Doctors They Assist in Operating Rooms Nationwide

JUL 04, 2026



Share



Part I Podcast free on Spotify.



Thoughts on Healthcare Markets & Technology



Part I: 25x the Surgeon's Pay: How Surgical Assistants Are Using No Surprises Act Arbitration Loopholes to Out Earn the Doctors They Assist in Operating Rooms Nationwide

A scoliosis surgery. Surgeon earned \$8,016. The assistant standing next to him got \$196,566. That is a 24x multiple in favor of the person holding the retractor...

▶ Listen now

30 minutes ago · Thoughts on Healthcare



Part II Podcast episode for paid subscribers only. Also available on Spotify.



Thoughts on Healthcare Markets & Technology



Part II: 25x the Surgeon's Pay: How Surgical Assistants Are Using No Surprises Act Arbitration Loopholes to Out Earn the Doctors They Assist in Operating Rooms Nationwide

A scoliosis surgery. Surgeon earned \$8,016. The assistant standing next to him got \$196,566. That is a 24x multiple in favor of the person holding the retractor...

▶ Listen now

28 minutes ago · Thoughts on Healthcare

To listen to paid episodes in Apple or Spotify, link settings on those platforms (instructions inside the Podcast).

Table of Contents

The Setup

The Sixteen Percent Rule and How It Got Blown

The Wisconsin Fusion Trick, or How to Turn

Who's Actually Cashing These Checks

Why Arbitrators Keep Picking the High Nurse

The Industry Response, or Everyone Points at Everyone Else

What This Means for Payment Integrity as an Asset Class

Where This Probably Goes From Here

Abstract

This piece breaks down the New York Times investigation published June 29, 2026 into surgical assistant billing under No Surprises Act arbitration, where assistants are pulling in payments up to 24–25 times what the operating surgeon collects for the same case. The Times' examples include a scoliosis surgery where the assistant's arbitration award was about 24 times the surgeon's payment, a Texas prostatectomy with a roughly 27-fold differential, facial feminization in New York where the assistant (who is married to the surgeon) received about 16 times the surgeon's pay, and a New Jersey breast reconstruction with an assistant payment more than 40 times the surgeon's. Key data points covered include the standard 16 percent benchmark for assist fees in commercial fee schedules, the Wisconsin spinal fusion case that was split into 11 separate arbitration filings, the Texas prostatectomy where the assistant out-earned the surgeon by more than 27x, and reactions from AHIP, Elevance,



Discover more from Thoughts on Healthcare Markets & Technology

Expert analysis of healthcare and life sciences markets, technology, investment, entrepreneurship policy, and AI — for investors, entrepreneurs,...

Over 4,000 subscribers

Enter your email...

Subscribe

By subscribing, you agree Substack's [Terms of Use](#), and acknowledge its [Information Collection Notice](#) and [Privacy Policy](#).

Already have an account? [Sign in](#)

Surgical Assistants, and CBO's own walkback on its cost projections for the law. The back half gets into why this matters for anyone building payment integrity, fraud detection, or provider risk models, and how this exact pattern (facility-driven out-of-network staffing plus arbitrator information asymmetry) is a preview of where a lot of qui tam and SIU attention is headed next.

Thoughts on Healthcare Markets &
Technology is a reader-supported
publication. To receive new posts and
support my work, consider becoming a
free or paid subscriber.

The Setup

So the Times ran a piece this week that's basically a gift to anyone who spends time thinking about provider payment mechanics, and the headline number is the kind of thing that makes you do a double take. A scoliosis surgery in New York: surgeon gets \$8,016, the assistant standing next to him gets \$196,566. Not a typo, not a decimal error, just a straight up 24x multiple in favor of the person holding the retractor instead of the person doing the fusion. Same pattern shows up in a Texas prostatectomy: according to the Times, the surgeon pulled \$1,843 while the assistant's arbitration award was \$50,456, a differential of more than 27x. Facial feminization case in New York, surgeon gets \$12,767, the assistant (who also happens to be the surgeon's spouse) gets \$210,000. Breast reconstruction in New Jersey, \$2,707 versus \$111,000. You start to notice the pattern isn't a glitch, it's a strategy, and it's one that a very specific set of actors figured out how to run at scale.

The mechanism here is the federal No Surprises Act, the federal law that took effect January 1, 2022 and was supposed to protect patients from getting blindsided by out-of-network bills when, say, an anesthesiologist who doesn't take their insurance shows up during an in-network surgery. The law set up an independent dispute resolution process, IDR for short, where an out-of-network provider and a health plan each submit a number and a neutral arbitrator picks one or the other, no splitting the baby allowed. It's baseball-style arbitration, basically, borrowed straight from labor negotiations. The idea was solid on paper: stop balance billing, let providers and payers hash out fair value through a structured process instead of sticking patients with the tab. What nobody quite modeled out is what happens when the "providers" gaming the system aren't the marquee surgeons everyone assumes are driving costs, but the second and third chair folks nobody really thinks about.

The Sixteen Percent Rule and How It Got Blown Up

Here's the baseline everyone in the reimbursement world already knows: assist-at-surgery fees are typically set as a percentage of the primary surgeon's allowed amount, and 16 percent is the number that shows up over and over in commercial fee schedules. That's not arbitrary, it reflects the actual labor differential. The assistant is positioning robotic arms, managing suction, handing off instruments, doing real and sometimes clinically important work, but it's fundamentally a supporting role relative to the person making the actual surgical decisions and doing the cutting. Insurers built their fee schedules around that ratio for decades without much controversy.

What the Times investigation surfaces is that when these claims go through arbitration instead of a standard contracted or benchmark payment, that 16 percent ratio doesn't just bend, it inverts. According to the data reviewed, assistants are hitting payments up to roughly 24–25 times the surgeon's rate in some cases, not 16 percent of it, 24–25 times it. The mechanics of why this happens are worth sitting with because they're not really about clinical value at all, they're about how arbitrators are instructed to weigh evidence and what evidence actually gets put in front of them. Arbitrators are told to consider the median in-network rate as an anchor (the qualifying payment amount, or QPA), along with provider experience, case complexity, and prior contracted rates. Surgeons, because they're the recognizable, credentialed, high-profile party in the room, often end up accepting the patient's insurance and billing in-network, meaning their payment is anchored to that boring, transparent, contracted number. Assistants, meanwhile, frequently operate out of network on purpose, sometimes through staffing arrangements specifically structured that way, which means their claim goes to arbitration completely divorced from any surgeon comparison.

The arbitrator evaluating the assistant's claim in isolation has no idea the surgeon on the same case got paid a tenth or less as much, because in a lot of these situations they're literally different arbitrators looking at different case files with no visibility into each other's decisions. And when you zoom out from the anecdotal cases to the broader data, you see the same directional pattern: empirical work on NSA arbitration finds median awards 2–3 times median in-network commercial rates and well above QPA, with providers winning over 80 percent of disputes.

The Wisconsin Fusion Trick, or How to Turn One Surgery Into Eleven Paydays

The single best illustration in the whole piece, and the one that ought to get circulated in every payment integrity Slack channel, is the Wisconsin case involving neurosurgeon Arvind Ahuja and a 2024 spinal fusion. The total arbitration payout for that one operation came to \$196,215, split as \$125,508 to Ahuja as the primary surgeon and \$70,707 to his assistant. Compare that to what a normal contracted rate would have produced: the Times reports that the insurer estimated roughly \$9,310 for the surgeon and \$1,562 for the assistant under typical in-network payment, meaning the arbitration outcome represents something like a 21x multiple on the total expected payment.

The way they got there is the part that should really catch the attention of anyone who thinks about claims structuring for a living. Instead of billing the fusion as a single procedure, the practice broke it into pieces and brought 11 separate cases to arbitration. Eleven separate filings, eleven separate arbitrator assignments, eleven separate opportunities to win, and they won every single one. That's not an assistant capitalizing on an anomaly, that's a repeatable, scalable operating procedure, the kind of thing you'd build a standard operating checklist around if you were running this as a business line, which, functionally, someone clearly was.

This is where the loophole framing in the headline actually earns its keep. The No Surprises Act's arbitration "cooling-off" provision requires a 90-day gap before the same parties can re-arbitrate the same or similar services, explicitly designed to prevent exactly this kind of repeated gaming. But if you fragment one surgery into 11 discrete billable components, each one is technically a different service, so the lockout rule never triggers. Nobody wrote the CPT coding rules or the arbitration eligibility rules with this specific combination in mind, and that's exactly the kind of seam that shows up whenever you build a rules-based adjudication system without adversarial testing.

Anyone who's spent time in claims editing or SIU work will recognize this instantly: it's unbundling, dressed up in arbitration clothing instead of billing clothing, and it works because the arbitration side of claims processing has nowhere near the fraud, waste, and abuse tooling that the standard adjudication pipeline has had built up over 20 years.

Who's Actually Cashing These Checks

One detail that's easy to skim past but is actually pretty important for anyone modeling provider risk is who these high-earning assistants tend to be. Per the Times reporting, the assistants pulling these outsized arbitration awards are sometimes physicians but more often nurses or physician assistants, meaning a huge chunk of

this economic upside is flowing to mid-level clinicians who, under any normal fee schedule logic, would never come close to touching surgeon-level compensation for a single case. That's not a knock on NPs or PAs, plenty of them are doing legitimate, clinically valuable work in these roles. But it does mean the incentive structure being created here is completely disconnected from clinical training, risk, or decision-making authority, and entirely a function of who happened to bill out of network and who happened to get assigned a favorable arbitrator on a given day.

There's also the spousal detail buried in the New York facial feminization example, where the assistant collecting \$210,000 against the surgeon's \$12,767 is married to the surgeon. That's a related-party transaction hiding inside what looks, on paper, like an arms-length arbitration outcome, and it's exactly the kind of pattern that would jump off the page in any decent claims analytics pipeline if anyone were running related-party detection against IDR award data. Nobody currently is, at scale, because IDR award data isn't structured or centralized the way standard claims data is. That gap is basically an open invitation.

Illinois-based Surgikal Assistants, one of the staffing firms operating in this space, doesn't even really dispute the characterization. Its CEO, Luis Aragon, told the Times these six-figure assist awards are "signs of a system gone awry," calling amounts like \$100,000 or \$50,000 "way out of line" and flatly saying it's "not sustainable for health care costs." That's a fairly remarkable quote to get from someone whose business model presumably benefits from exactly this dynamic, and it tells you something about how obviously unsustainable the current trajectory is even to people inside the ecosystem profiting from it.

Why Arbitrators Keep Picking the High Number

The structural reason this keeps happening comes down to how IDR arbitrators are set up to evaluate a claim in a vacuum. Under the federal process, the arbitrator cannot split the difference or invent their own number; they have to pick one of the two competing offers in full. That baseball arbitration format was chosen deliberately to discourage extreme anchoring, the theory being that if you know the arbitrator has to pick a side, you'll submit something reasonable because a wildly inflated ask looks obviously unreasonable next to the payer's number.

The problem is that theory only holds if the arbitrator has good comparative information to judge reasonableness against, and for assistant-at-surgery claims specifically, that comparative information mostly doesn't exist in any usable form. There's no standardized public benchmark for what an assist fee "should" be in

arbitration the way there sort of is for primary surgeon fees, and the qualifying payment amount that's supposed to anchor everything gets undermined the moment providers strategically avoid network participation specifically so their claim has no in-network comparison to be measured against.

Add to that the fact that a huge share of these procedures are scheduled, elective, non-emergency operations, not the ER surprise bill scenario Congress was actually picturing when it wrote the statute. Brady Connaughton, a New Jersey attorney who advises union health plans, put it about as bluntly as a lawyer is going to put anything on the record, calling the pattern “a blatant disregard of both the spirit of the law and what the law says.” Legislators were focused on protecting patients from getting stuck between an ambulance ride and an out-of-network ER doc, not on a scheduled scoliosis correction booked six weeks in advance where the facility chose to staff the case with out-of-network assistants on purpose. That's the part that turns this from an unfortunate side effect into something that looks a lot more like designed exploitation.

When you zoom back out beyond assistants, the broader arbitration statistics reinforce the sense that arbitrators are defaulting to the higher number much more often than not: research and policy discussions now routinely cite provider win rates above 80 percent, with arbitration awards that are 2–3x median in-network rates and far above QPA and Medicare benchmarks.

The Industry Response, or Everyone Points at Everyone Else

AHIP's public statement on this, through spokesperson David Allen in comments cited by the Times, characterizes the pattern as “outrageous provider-driven abuse” that's adding billions in wasteful spending and driving up costs for everyone, and pushes for policy fixes to the flawed IDR incentive structure. That's a predictable position from the payer trade group, but it's backed by something a little more interesting: the Congressional Budget Office itself has started walking back its own projections.

CBO acknowledged this month that the emerging evidence suggests the law “might not have the savings effects” originally anticipated, and flagged risk of higher premiums, bigger federal deficits, and increased anti-competitive consolidation among providers who use arbitration leverage as a negotiating cudgel to demand better in-network rates in the first place. That last part is the sleeper detail. Some provider groups are reportedly using the mere threat of inflated arbitration outcomes as leverage to extract higher contracted rates from payers who'd rather just avoid the

arbitration exposure entirely. So the loophole isn't just generating direct overpayments, it's distorting the underlying in-network contracting market too, which is a much bigger and stickier problem than any single scoliosis case.

Elevance has been more concrete about actually doing something. The insurer is rolling out a facility administrative policy across affiliated commercial plans in 11 states that specifically targets facility-driven use of nonparticipating providers for planned, nonemergency care, trying to push participating hospitals and surgical centers toward staffing cases with in-network assistants. That policy work is happening alongside broader federal rule changes: in May 2026, the Departments of Health and Human Services, Labor, and the Treasury finalized new IDR operations rules that, among other things, cut the per-party administrative fee from \$115 to \$15 for disputes initiated on or after June 11, 2026, doubled the maximum number of line items that can be batched in a single dispute from 25 to 50, and established an IDR Gateway platform to standardize and track disputes.

What This Means for Payment Integrity as an Asset Class

If you're building anything in the payment integrity, fraud detection, or provider risk scoring space, this is one of those moments where a very specific operational pattern suddenly jumps from "edge case you hear about at conferences" to "front-page New York Times story that regulators and general counsels are going to read." The underlying mechanics here—facility-driven out-of-network staffing, fragmented billing of a single surgical episode into multiple arbitration-eligible line items, and arbitrators making decisions in claim-level silos without cross-case visibility—are not exotic.

They look a lot like the same playbooks you already see in ordinary claims fraud and abuse work: unbundling, use of related parties, and exploitation of information asymmetries between different parts of the payment pipeline. What's different is that they're happening in a part of the system—federal IDR—that historically hasn't been wired into the same analytics and oversight stack as regular adjudication.

For payment integrity specifically, there are three immediate implications:

IDR data has to be treated as core, not peripheral. If arbitration outcomes are routinely delivering 2–3x in-network rates and occasionally 20x+ multiples for

assistants on single cases, then any true-up, reserve, or trend model that ignores IDR exposure is going to be wrong in ways that matter.

Facility behavior belongs firmly inside provider risk models. In almost every Times example, it's the facility making conscious staffing choices that set up the out-of-network assistant exposure. If your risk scoring is still primarily physician-centric, you're going to miss the locus of control for these arbitrage plays.

Where This Probably Goes From Here

On the regulatory side, the combination of the Times piece, CBO's walkback, and growing empirical evidence that arbitration awards systematically overshoot negotiated rates makes some form of tightening almost inevitable. The May 2026 IDR rule changes are a first step focused on operations—fees, batching, standardized codes, and a gateway platform—but they don't directly attack the assistant-specific loopholes yet.

Expect at least three trajectories:

- More aggressive payer-side policy moves, like Elevance's facility administrative policy, aimed at cutting off the out-of-network assistant pipeline before it ever hits IDR.
- Pressure to incorporate comparative benchmarks for assists into arbitrator guidance, whether through sub-regulatory instructions or future rulemaking, so that a six-figure assist award has to be explicitly justified against something other than the provider's narrative.
- Expansion of SIU and audit focus into IDR decisions themselves, including arbitrator-level pattern analysis, as researchers and payers continue to document provider win rates above 80 percent and awards well above market medians.

For people underwriting payment integrity as an asset class—VCs, private equity, and strategically minded payers—the takeaway isn't that the No Surprises Act failed. The law clearly did what it was supposed to do for patients in emergency and unplanned care settings. The takeaway is that every time you carve out a parallel payment track to solve one problem, you create a new surface area for arbitrage somewhere else.

Assistants turning one spinal fusion into 11 arbitration paydays and walking away with 24x surgeon-level multiples are just the most photogenic edge of that surface. The more prosaic reality—quiet leverage in contract negotiations, facility policies that normalize out-of-network staffing, and arbitrators making systematically generous

decisions in data-poor environments—is where most of the dollars, and most of the future enforcement, are going to live.

Subscribe to Thoughts on Healthcare Markets & Technology

By Thoughts On Healthcare Markets & Technology · Hundreds of paid subscribers

Expert analysis of healthcare and life sciences markets, technology, investment, entrepreneurship, policy, and AI — for investors, entrepreneurs, hospital and insurance executives, and physicians navigating the business of healthcare.

By subscribing, you agree Substack's [Terms of Use](#), and acknowledge its [Information Collection Notice](#) and [Privacy Policy](#).

← Previous

Discussion about this post

Comments Restacks



Write a comment...