Prior Auth & Denials Are Healthcare's Most Hated Processes But Medicare and Medicaid Lose $100-300B a Year Fraud While Commercial Plans Lose 1 3% and the Difference Is Largely That Commercial Plan

APR 14, 2026

Table of Contents

1. Abstract

2. The Great Prior Auth and Denials Paradox

3. Fraud by the Numbers: Government Programs vs. Commercial Plans

4. How Prior Auth and Denials Actually Work as Fraud Prevention

5. The Medicare and Medicaid Fraud Landscape: A Quick Tour of the Wreckage

6. Why Private Payers Don't Have This Problem (Or At Least Not Nearly as Bad)

7. The Uncomfortable Tradeoff Nobody Wants to Talk About

8. Where the Opportunities Are

9. The Bottom Line

Abstract

- Prior authorization and claims denials are universally despised across the healt ecosystem, with bipartisan legislative efforts aimed at curtailing their use in commercial insurance.

- Meanwhile, Medicare and Medicaid lose an estimated $100-300B+ annually to improper payments and outright fraud, numbers that dwarf fraud losses in commercial plans.

- Commercial payers deploy prior auth, claims denials, utilization management, sophisticated analytics that function as a de facto fraud and abuse prevention lay one that government programs largely lack.

- The thesis here: the very mechanisms providers and patients hate most in commercial insurance may be the primary reason private plans don't hemorrhag money to fraud at anything close to the rate of public programs.

- Prior auth blocks fraud prospectively by forcing review before payment. Denia block fraud retrospectively by catching suspicious claims after submission but b or shortly after payment. Together they create a two-layer defense that governme programs have historically lacked.

- For health tech entrepreneurs and operators, this creates a massive opportunity space around fraud prevention, payment integrity, and automation tools targetin government programs, while also raising hard questions about the real cost of dismantling prior auth and denial controls in commercial plans.

- Key opportunity themes include AI-driven prior auth automation, intelligent d management, predictive fraud analytics for government payers, payment integrit platforms, and provider workflow tools that balance access with accountability.

The Great Prior Auth and Denials Paradox

There is a strange disconnect happening in healthcare policy right now, and it is worth pausing on because it has real implications for anyone building, operating

investing in health tech. On one side of the conversation, you have providers, pa advocacy groups, and frankly most of Congress united in their desire to gut prio authorization requirements and rein in claims denials in commercial insurance. complaints are legitimate and well-documented. Prior auth delays care. Denials appeals processes that consume enormous provider resources and sometimes res patients simply going without needed treatment. Together, these mechanisms cr administrative burden that costs the system billions, burns out physicians, and occasionally kills people. The horror stories are real, and nobody is arguing othe

But on the other side of the ledger, something far less discussed is happening. Government insurance programs, specifically Medicare and Medicaid, are bleed money at a rate that should make anyone with a finance background physically uncomfortable. We are talking about improper payment rates that, depending o whose numbers you trust, run somewhere between $100 billion and $300 billion year. Some estimates go higher. A meaningful chunk of that is straight-up fraud. billing errors. Not coding mistakes. Fraud. Fake patients, phantom clinics, orga crime rings, the whole nine yards.

And here is the part that nobody seems to want to connect: Medicare and Medic have historically operated with far less rigorous prior authorization and far lowe denial rates than commercial plans. When government programs do deny claims often retrospective, meaning the money already moved and recovering it become lengthy, expensive process that frequently fails. The correlation between "fewer prospective controls and less aggressive denial practices" and "dramatically mor fraud" is not subtle. It is staring everyone in the face. Yet the policy conversatio treats these as two completely separate issues, as if the people screaming about p auth and denials in commercial insurance and the people screaming about fraud government programs are living on different planets.

> They are not. They are describing two sides of the same coin.

Fraud by the Numbers: Government Programs vs. Commercial Plans

Let's get into the actual numbers because this is where the argument gets hard t ignore. CMS publishes improper payment rates annually, and while the methodc has shifted over the years, the directional story is consistent and grim. Medicare for-service has run improper payment rates in the 6-8% range in recent years, wh on a base of roughly $450 billion in annual FFS spending translates to somewhe around $30-40 billion per year just in traditional Medicare FFS. Medicare Adva adds another layer of complexity with risk adjustment coding issues that GAO a OIG have estimated cost taxpayers tens of billions more annually. HHS OIG has the MA overpayment figure in the range of $12-25 billion depending on the year methodology.

Medicaid is arguably worse on a percentage basis. The most recent CMS data pe the Medicaid improper payment rate above 20% in some years, though that numl bounced around due to measurement changes. On a program that spends over $7 billion annually including both federal and state share, even a conservative 10% improper payment rate means $70 billion walking out the door incorrectly. And 20%+ figure, when it showed up, implied something north of $140 billion.

Now compare that to commercial insurance. The National Health Care Anti-Fra Association has historically estimated that fraud accounts for roughly 3-10% of t health spending, but that is an aggregate number across all payers. When you isc commercial plans specifically, the fraud and improper payment rates tend to run dramatically lower than government programs. The big commercial payers, your UnitedHealthcare, Anthem, Aetna, Cigna, Humana on the commercial side, typi report fraud loss ratios well under 3%, and most internal estimates from payer executives suggest the real number is closer to 1-2% for well-managed commerci books of business.

So you have got government programs losing somewhere in the range of 8-20% t improper payments and fraud, and commercial plans losing maybe 1-3%. That is rounding error. That is an order of magnitude difference. And the single biggest structural distinction between these two payer types, besides the obvious scale a population differences, is the intensity of prospective utilization management ar

willingness to deny claims that do not meet clinical or billing criteria. Which is fancy way of saying prior auth and the associated denial machinery.

How Prior Auth and Denials Actually Work as Fraud Prevention

Most of the public conversation about prior auth and denials focuses on their ro clinical gatekeeping. Does the patient really need that MRI? Is that brand-name medically necessary when a generic exists? Should this surgery happen at an outpatient center instead of an inpatient facility? These are the utilization management questions that drive providers crazy, and rightfully so in many case where the clinical answer is obvious and the auth process just adds friction and Similarly, the denials conversation tends to focus on legitimate claims getting re for technicalities or documentation gaps, forcing providers into costly appeal cy

But there is a second function of both prior auth and denials that gets almost ze airtime: fraud prevention. These two mechanisms work as complementary layers defense. Prior auth works as a prospective fraud deterrent because it forces revi the service before it happens. Denials work as a concurrent and retrospective fra barrier because they catch suspicious claims that make it past the front door, rej payment before or shortly after the money moves. Together, they create a two-lay system that is fundamentally different from the pay-and-chase model that Medic has historically relied on, where claims get paid first and audited later, sometime much later, sometimes never.

Think about it from the perspective of someone running a fraudulent billing operation. If you are billing a commercial plan for, say, a series of expensive gene tests on patients who never actually received them, you have a problem at two le First, the plan is likely going to require prior authorization for those tests. Some going to review the clinical documentation before approving the service. The pa primary care physician may get a notification. The plan may require the test to b performed at a credentialed lab. Second, even if you somehow get past prior autl claims adjudication system is going to run those claims through editing logic, m

policy rules, and increasingly sophisticated analytics before payment. Claims tha trigger flags get denied, and the denial creates a paper trail that feeds into the pl special investigations unit. There are multiple checkpoints that a fraudulent cla to clear before money changes hands, and each checkpoint generates data that m the next fraudulent claim harder to push through.

Now try the same scheme against Medicare FFS. Submit the claim. Get paid in 1 days. Maybe get audited in two years. Maybe not. If a claim does get denied in Medicare, it is usually for a coding or documentation technicality, not because someone prospectively reviewed the clinical scenario. The structural vulnerabili enormous, and organized fraud rings know it. The DOJ has prosecuted cases inv hundreds of millions of dollars in Medicare fraud perpetrated by operations tha for years before anyone caught on. Some of the most notorious cases, like the $1. billion home health fraud takedown in 2022 or the $1.4 billion telemedicine frau schemes that DOJ rolled up during and after COVID, specifically exploited the absence of prospective controls and the low denial rates in Medicare.

Commercial prior auth and denials are not perfect. They are often clunky, slow, applied in situations where they add cost without clinical value. But as structura anti-fraud mechanisms, they work remarkably well. Prior auth as the front gate denials as the back gate create a system where it is genuinely difficult for frauds operate at scale.

The Medicare and Medicaid Fraud Landscape: A Quick Tour of the Wrecka

For anyone not tracking the government program fraud space closely, a brief tou recent enforcement actions is instructive. DOJ's Health Care Fraud Strike Force which operates in major metro areas across the country, has been running coord takedowns for over fifteen years. The annual totals are staggering. In June 2023, announced charges against 78 defendants across the country for approximately $ billion in alleged fraud. The 2022 action tagged $1.7 billion. These are annual ev now, recurring like clockwork, and the amounts keep growing.

The schemes are diverse and increasingly sophisticated. Durable medical equipr fraud remains a perennial favorite, with operations billing for wheelchairs, brac and orthotics that patients never received. Home health fraud is massive, particu in states like Texas and Florida, where fake home health agencies have been cau billing for services on patients who were either dead, not homebound, or never visited. Compounding pharmacy fraud exploded a few years back, with pharmac billing government programs for expensive custom compounds that were either dispensed or therapeutically unnecessary.

And then there is the telemedicine fraud wave that came out of COVID. When C loosened telehealth restrictions during the public health emergency, including dropping prior auth requirements, expanding the types of services eligible for telehealth billing, and effectively lowering the bar for claim denials, fraud opera moved in almost immediately. The DOJ has since prosecuted telemedicine fraud totaling multiple billions, with schemes typically involving call centers that wou cold-call Medicare beneficiaries, conduct sham telehealth visits, and then bill fo expensive genetic tests, durable medical equipment, or pain creams. The benefic often had no idea their Medicare numbers were being used.

The common thread in almost all of these schemes is the absence of prospective review and the low rate of claim denial. The fraudsters specifically target Medic and Medicaid because these programs pay first and investigate later, and becaus odds of any given fraudulent claim being denied before payment are relatively lc compared to commercial plans. The pay-and-chase model combined with low de rates is not just inefficient. It is an invitation.

Medicaid fraud has its own special flavor, often involving providers in long-term behavioral health, personal care services, and substance abuse treatment. The behavioral health and substance abuse categories have been particularly problen with “Florida shuffle” style operations cycling patients through sham treatment programs and billing Medicaid for services that were either not rendered or gro substandard. California’s Medicaid program alone has estimated fraud losses in billions annually.

Compare any of this to what happens in commercial insurance and the contrast i sharp. Commercial fraud certainly exists, but it tends to be smaller in scale, dete faster, and harder to sustain because the prospective controls and more aggressiv denial practices catch anomalies before large sums move.

Why Private Payers Don't Have This Problem (Or At Least Not Nearly as Bad

The question worth asking is: what exactly are commercial plans doing different The answer is not one thing but a layered system of controls, and prior auth and denials sit near the top of that stack.

First, commercial plans maintain provider credentialing processes that are more rigorous than Medicare's enrollment system. To bill a commercial plan, a provid typically needs to be credentialed through a process that verifies licensure, malpractice history, practice location, and specialty qualifications. Medicare has own enrollment process, obviously, but it has historically been more permissive slower to remove bad actors. CMS has made improvements here, particularly thr the Affordable Care Act's enhanced screening provisions, but the commercial credentialing process remains tighter in practice.

Second, commercial plans use prior authorization as a prospective control on hig cost services. This means that before expensive imaging, surgeries, specialty dru genetic tests, and durable medical equipment are authorized, someone at the pla its utilization management vendor reviews the request. This review serves a dual purpose. It assesses medical necessity, which is the clinical gatekeeping function everyone complains about, and it validates that the requesting provider, the pati and the proposed service all check out. It is very hard to bill a commercial plan f service on a patient who does not exist or from a facility that is not real when son is reviewing the request prospectively.

Third, commercial plans deploy denial logic as a second line of defense. Claims make it past prior auth still run through automated editing systems, medical pol engines, and payment integrity algorithms at the point of adjudication. Claims t

not match the authorization, that come from non-credentialed providers, that cc coding anomalies, or that trigger fraud indicators get denied. These denials serv different function than the prior auth layer. Where prior auth prevents fraud frc entering the system, denials catch it at the point of payment and stop the money moving. The combination of prospective and concurrent controls is what makes commercial system so much harder for fraudsters to exploit than government programs, which tend to rely more heavily on post-payment audits.

Fourth, commercial plans invest heavily in analytics and special investigation un The big payers run sophisticated data science operations that flag billing anoma provider outliers, and suspicious patterns in near-real time. When these analytic generate flags, the response is often a targeted increase in prior auth requiremer the flagged provider or an increase in the denial rate for specific claim types, cre a feedback loop that tightens controls around suspicious actors. Medicare has it analytics capabilities through CMS's Center for Program Integrity and contracto but the commercial payer analytics infrastructure is generally more advanced an more aggressively deployed.

Fifth, commercial plans have a direct financial incentive to prevent fraud and de improper claims. Every dollar lost to fraud comes directly off the bottom line. M loss ratio regulations under the ACA mean that plans need to keep admin costs v bounds, but fraud losses hit the medical cost side of the equation and directly in profitability. This creates a strong alignment between the payer's economic inter and aggressive use of prior auth and denials. Medicare and Medicaid, by contras spending taxpayer money, and while CMS certainly has fraud prevention prograı the institutional urgency is different. Bureaucratic processes, interagency coordination challenges, and political dynamics all slow the government's respor relative to what a commercial plan with direct profit exposure can do.

Sixth, commercial plans benefit from smaller, more defined networks. A commer plan knows its providers. It has contracts with them. It knows their billing patte their patient panels, their historical utilization. When something looks off, the s is easier to spot against a known baseline, and the plan can respond by increasin prior auth requirements or denial rates for that specific provider. Medicare, whi

essentially an open network where any enrolled provider can bill the program, h
much harder signal-to-noise problem. The sheer scale and openness of the Medi
provider base makes it structurally more vulnerable to bad actors blending in.

The Uncomfortable Tradeoff Nobody Wants to Talk About

Here is where this gets politically uncomfortable, which is probably why the
conversation rarely happens in mixed company. The prior auth and denials mach
that everyone hates in commercial insurance is performing a fraud prevention
function that is saving enormous amounts of money. The government programs
lack equivalent controls are hemorrhaging cash to fraudsters at a rate that, if it
happening in any other sector, would be considered a national scandal.

This does not mean that prior auth and denials as currently implemented are op
They are not. Prior auth is too slow, too manual, too often applied to routine ser
that do not need prospective review. Denials are too often applied to legitimate c
for technical reasons, creating enormous appeal volumes and delaying payment t
honest providers. The Gold Card programs that some states have implemented,
high-performing providers get exempted from prior auth requirements, are a sm
step in the right direction. So are the CMS interoperability rules requiring elect
prior auth by 2027, and the various health tech solutions automating both the pr
auth and denial management workflows.

But the conversation about reforming these mechanisms needs to happen with c
eyes about what happens when you remove prospective controls and reduce deni
rates entirely. The Medicare and Medicaid experience provides a natural experin
and the results are not encouraging. When you pay first and chase later, you lose
of money. When you do not require prospective justification for services, fraud s
easily. When your denial rates are low and your provider enrollment and credent
processes are permissive, bad actors get in and stay in.

The advocacy community, and frankly a lot of the health policy commentariat, ta
about prior auth and denial reform as if the only variable is access to care. And a

matters enormously, nobody is disputing that. But the fraud prevention functior matters too, and pretending it does not exist leads to policy proposals that could very expensive unintended consequences.

Consider what would happen if commercial plans were required to eliminate pri auth and dramatically reduce denial rates, as some legislative proposals have suggested. Based on the differential fraud rates between commercial and govern programs, you would expect a meaningful increase in fraudulent billing. How mu Hard to say precisely, but even a 2-3 percentage point increase in the commercia fraud rate would represent tens of billions of dollars annually, costs that would ultimately flow through to employers and consumers in the form of higher prem The actuaries at the big plans have modeled these scenarios, and the numbers ar pretty.

Where the Opportunities Are

For anyone building or operating in health tech, this tension between prior auth denial reform on one hand and fraud prevention on the other creates several mas opportunity areas.

The most obvious is prior auth automation. The market for solutions that make auth faster, less burdensome, and more clinically intelligent is large and growing Platforms that can automate the prior auth submission process, use clinical data pre-populate authorization requests, and reduce turnaround times from days to minutes are addressing a real pain point without eliminating the prospective rev function. The value proposition is straightforward: keep the fraud prevention be of prior auth while removing the administrative friction that delays care and bur out clinicians.

Denial management and optimization is equally compelling, and it cuts both way the provider side, tools that help practices and health systems manage denials m efficiently, automate appeals, and reduce denial rates for legitimate claims have massive addressable market. On the payer side, solutions helping plans deploy smarter denial logic that catches fraud and coding errors without generating the

massive false positive rates that plague current systems are equally valuable. The outcome is a denial system that is more accurate in both directions: fewer denial legitimate claims and more denials of fraudulent ones.

Government program fraud prevention is arguably the biggest greenfield opport in the bunch. Given the scale of improper payments in Medicare and Medicaid, t is an enormous market for solutions that bring commercial-grade prior auth, de logic, and fraud analytics to government programs. CMS has been investing in th area, and the Fraud Prevention System that CMS operates has identified and prevented billions in improper payments. But the gap between government and commercial capabilities remains wide. Predictive models, provider risk scoring platforms, and real-time claims surveillance tools specifically built for the Medi and Medicaid context address a gap that costs taxpayers hundreds of billions an

Payment integrity platforms represent a related but distinct opportunity. Payme integrity goes beyond fraud to include coding accuracy, clinical validation, and b compliance. The payment integrity market for commercial payers is already well established with large incumbents. But the government program payment integr space is less mature and arguably more impactful given the higher improper pay rates.

Provider-side workflow tools that help legitimate providers navigate the prior au and denials landscape while ensuring compliance represent another substantial market. Think of this as the provider workflow layer that sits between the clinic decision and the payer authorization or claim adjudication. Tools that can predi whether a prior auth will be required, pre-check the likely approval criteria, ass the necessary documentation automatically, submit electronically, and proactive address the most common denial triggers before submission are valuable to prov and do not threaten the fraud prevention function that payers rely on.

And then there is intelligent utilization management, which is more speculative potentially very large. This means moving beyond binary approve/deny logic to n nuanced, risk-stratified approaches. A system that applies intensive prospective review and higher denial thresholds to high-risk providers, new-to-network enti

and unusual service patterns, while fast-tracking authorizations and reducing de friction for established providers with clean billing histories. This is essentially Gold Card programs do at a coarse level, but there is room for much more granu data-driven approaches that could dramatically reduce administrative burden fo vast majority of legitimate providers while actually increasing scrutiny on the sn percentage of actors who account for most of the fraud.

The Bottom Line

The prior auth and denials debate in healthcare is one of those situations where loudest voices in the room are not necessarily wrong, but they are definitely incomplete. Yes, prior auth as currently implemented is often terrible. Yes, deni legitimate claims waste enormous resources and delay necessary care. Yes, both to be reformed, automated, and made more intelligent.

But the argument that prior auth should simply be eliminated and denial rates s across the board, that these mechanisms serve no useful purpose, that they are p tools for payers to deny care and boost profits, that argument does not survive co with the data. The differential fraud rates between commercial plans (which use auth and denials aggressively) and government programs (which historically have tell a clear story. These mechanisms, for all their flaws, are functioning as critica fraud prevention layers. And the scale of fraud in programs that lack them is genuinely breathtaking.

For health tech entrepreneurs and operators, this creates a rare situation where sides of the equation present opportunity. The reform side needs better technolo make prior auth less painful and denials more accurate. The fraud prevention si needs better technology to bring government programs up to something approac commercial-grade integrity. And the sweet spot is solutions that can do both simultaneously, reducing the burden on legitimate providers while increasing th detection of fraudulent actors.

The regulatory trajectory supports this view. CMS is moving toward electronic p auth requirements. Multiple states are passing Gold Card laws. Congress contin

advance bipartisan prior auth and denial reform legislation. And at the same tim
CMS is investing in enhanced program integrity tools and the DOJ continues to
up healthcare fraud enforcement. These parallel tracks are not contradictory. Th
complementary. The future of utilization management is not “fewer controls.” It
“smarter controls.” And that is a technology problem, which means it is exactly t
kind of problem that health tech companies can solve.

The fraud numbers in government programs are not going down on their own. T
political pressure to reform prior auth and denials in commercial plans is not go
away. Both problems need technology solutions. Both represent enormous mark
And both are underfunded relative to their scale.

2 Likes • 1 Restack

← Previous

Discussion about this post

Comments Restacks

Write a comment...