# The Home Health Arbitrage: How CMS Just Created a Billion-Dollar Market Opportunity

NOV 30, 2025 • PAID



DISCLAIMER: The views and opinions expressed in this essay are solely my own do not reflect the views, opinions, or positions of my employer, Datavant, or any affiliates.

If you are interested in joining my generalist healthcare angel syndicate, reach o trey@onhealthcare.tech or send me a DM. Accredited investors only.

## ABSTRACT

CMS's 2026 Home Health Prospective Payment System final rule introduces significant payment reforms including a mandatory cap on therapy service utiliz expanded rural add-on payments, and new health equity adjustments. This essay proposes a revolutionary business model that exploits the gap between tradition high-intensity therapy-focused home health and the new payment incentives fav comprehensive care coordination. The proposed venture would function as a technology-enabled care orchestration layer that helps home health agencies opt their mix of services while capturing value through improved outcomes and redu therapy costs. Key elements include predictive analytics for therapy utilization, automated care planning that balances medical and social needs, and a novel rev

model that shares savings generated from reduced therapy overutilization. The addressable market exceeds 3.5 million Medicare beneficiaries receiving home h annually, with potential to capture 200-400 basis points of margin improvement participating agencies.

# TABLE OF CONTENTS



# What CMS Actually Changed and Why I Matters

Let me start by saying that most people reading CMS rules fall into two categori policy wonks who live for this stuff and operators who skim for the payment rat changes and move on. I am asking you to be neither. Instead, think like an invest looking for arbitrage opportunities created by regulatory change, because that is exactly what this rule creates.

The 2026 home health final rule does three things that matter. First, it implemer what CMS calls a therapy utilization cap, which is bureaucratic speak for “we ar going to penalize you if too many of your patients get a ton of physical therapy v

Specifically, if more than 58% of your thirty-day periods involve ten or more ther visits, you start losing money on every period above that threshold. The penalty i trivial at 5% of the payment, which on a $2,000 episode means you are giving bac $100 every time you cross that line.

Second, the rule expands rural add-on payments from 3% to 5% and makes them available to a wider swath of counties. This matters because rural home health h always operated on terrible unit economics due to drive time and patient density issues. The expanded add-on does not fix the problem but it makes rural market slightly less punishing.

Third, and this is where things get interesting, CMS introduced a health equity adjustment that provides bonus payments for agencies serving high proportions dually eligible beneficiaries or patients in underserved areas. The adjustment is around 2-3% of base payment, but it signals where CMS is heading. They want agencies focused on medically complex, socially complicated patients rather thar traditional home health model of cherry-picking the easiest cases.

Now, if you are running a home health agency today, this rule creates a massive operational headache. The typical agency generates maybe 4-6% operating margi good year. Most agencies have built their entire care model around maximizing therapy visits because that is what the old payment system rewarded. A typical t day home health period might include twelve physical therapy visits, eight skille nursing visits, and maybe an occupational therapy evaluation. Under the new rul that profile is suddenly underwater because you are in the penalty zone.

The immediate response from most agencies will be to arbitrarily cap therapy vi nine per period to stay under the threshold. This is the wrong answer but it is w most operators will do because they are not thinking about the underlying care r just the payment mechanics. And that creates the opportunity.

## The Therapy Utilization Problem Noboc Talks About

Here is a dirty secret about home health that anyone who has spent time in the industry knows but nobody says out loud: a huge percentage of therapy visits are clinically unnecessary. I am not saying they are fraudulent or that therapists are providing real services during those visits. I am saying that the marginal value of eight versus visit six is often negligible, and the decision to authorize that eighth has more to do with payment optimization than patient outcomes.

Why does this happen? Because the old home health payment model was based on therapy thresholds. If you provided six to nine therapy visits in a period, you got X. If you provided ten to thirteen visits, you got paid X plus 20%. This created an obvious incentive to get every patient who could tolerate it up to that ten-visit threshold. Clinical directors learned to write care plans that justified the visits, therapists learned to find things to work on that would support the medical nece documentation.

The result is that American home health patients get way more therapy than pat in other countries with similar outcomes. If you look at Medicare data, the avera home health patient receives about nine therapy visits per sixty-day episode. In Canada, the comparable number is around four visits. Outcomes are basically identical. That five-visit difference represents pure waste, but it is waste that ha baked into the business model for two decades.

CMS knows this, which is why they are implementing the cap. But here is the problem: agencies do not have the tools or expertise to figure out which patients actually need ten therapy visits versus which patients would do fine with six. Th current approach is one-size-fits-all care planning where the therapist does an in evaluation and then basically everyone gets the same visit schedule regardless of clinical presentation or functional trajectory.

What agencies need is a way to predict, at the time of admission, which patients require high-intensity therapy and which patients can achieve their goals with a lower-intensity approach. They also need help redesigning care plans to substitu other interventions, like home health aide services or remote monitoring, for ma therapy visits. And they need this to happen without sacrificing outcomes, becau

readmission rates go up or functional improvement scores go down, they lose mo through the value-based purchasing program.

This is not a problem you can solve by hiring smarter clinicians or implementing better training. This is a data problem and an orchestration problem. And that is where the opportunity lives.

## The Business Model: Care Orchestratio as a Service

The company I am describing would function as a software-enabled service that between the home health agency and the patient, helping to optimize the care pl and service mix to maximize outcomes under the new payment rules. Think of it Cohere Health of home health, but instead of prior authorization we are doing c orchestration and therapy utilization management.

The core product has three components. First, a predictive model that ingests re data at the point of admission and outputs a recommended care plan with visit frequency by discipline. The model would be trained on millions of historical ho health episodes with outcomes data, learning which patient characteristics predi high therapy needs versus low therapy needs. Variables would include diagnosis utilization, functional status scores, living situation, availability of informal caregivers, and about fifty other features you can extract from the OASIS assessm

The model would not just predict therapy needs but would also identify opportu to substitute lower-cost interventions. For example, a patient recovering from a replacement might traditionally get twelve PT visits over sixty days. The model m identify that this patient has a supportive family environment and good baseline functional status, and recommend six PT visits plus a home health aide to assist activities of daily living plus a remote monitoring program to track mobility and scores. The total cost is lower, the agency avoids the therapy cap penalty, and the patient outcome is equivalent or better because they are getting more holistic su

Second component is the care plan automation and documentation engine. Once model generates a recommended care plan, the software would auto-populate th relevant sections of the clinician's documentation, flag any medical necessity gap need to be addressed, and create visit schedules that balance clinical needs with operational efficiency. This is important because one of the barriers to changing patterns is that it creates more work for clinicians. If you make it easier rather t harder, you get adoption.

Third component is the ongoing monitoring and adjustment layer. As the episod progresses, the platform would track actual visit utilization against the plan, mo patient progress through whatever data sources are available like remote monito devices or patient-reported outcomes, and recommend adjustments in real time. patient is not hitting their functional milestones, the system might recommend adding visits. If a patient is progressing faster than expected, it might recommer stepping down intensity earlier.

The revenue model is what makes this interesting. Rather than charging a SaaS f patient, which is hard to sell into home health agencies operating on tight margi the company would operate on a shared savings model. The agency pays a small platform fee, maybe $20 per episode, and then the company captures 40-50% of t savings generated from reduced therapy utilization and avoided penalties.

Let me show you why this works. Average home health episode pays around $2,0 Under current patterns, an agency might have 65% of episodes in the high thera utilization bucket, which means they are paying penalties on 7% of their volume per episode. That is $7 per episode across their whole book, or about $24 millior annually for a mid-sized agency doing 350,000 episodes per year.

If the platform can reduce high-utilization episodes from 65% to 55%, the agenc saves the penalty on 10% of volume. That is $10 per episode times 350,000 episo $3.5 million in penalty avoidance. Add in another $5-8 per episode in direct ther cost reduction from lower visit intensity, and you are looking at total savings of 25 per episode. The company captures $10 per episode in performance fees, the agency nets $10-15 per episode in margin improvement, and everyone is happy.

The beauty of this model is that it aligns incentives perfectly. The company only makes money if it actually reduces utilization and improves economics for the ag There is no incentive to undertreat patients because if outcomes suffer, the agen gets hammered on their quality scores and stops using the platform. And the sha savings model means you can sell into the CFO rather than having to convince a clinical leader to change workflows, which is always a harder conversation.

## Why This Works Now and Not Three Years Ago

Timing is everything in healthcare investing. This exact business model would n have worked in 2022 for three reasons. First, the payment incentives were differ Before this rule, agencies were rewarded for high therapy utilization, so there w economic reason to reduce visits. A platform that optimized for lower utilizatior would have cost agencies money.

Second, the data infrastructure was not mature enough. Building the kind of predictive model I am describing requires access to large-scale claims data, OAS assessments, and outcomes measures. Three years ago, getting clean, linked data across those sources was nearly impossible. Today, with CMS data sharing progr and the maturation of the all-payer claims database ecosystem, you can actually assemble the training data you need.

Third, and most importantly, agencies were not ready to change their care mode Healthcare providers are deeply conservative organizations that resist operationa change unless forced by regulation or competitive pressure. The new rule create forcing function. Every agency in America is now scrambling to figure out how t under the therapy cap without destroying their margins. They are desperate for solutions, which means they will actually adopt new technology rather than runr six-month pilots that go nowhere.

The other timing element is competitive. There are no incumbent players solvin problem. The major home health software vendors like Homecare Homebase anc Axxess are focused on operational workflow and billing, not clinical decision sup

or utilization management. The prior authorization platforms like Cohere and Utilization Review have not moved into post-acute, so there is no direct competi with a similar model.

That gives a new entrant maybe eighteen to twenty-four months to build market position before the big players wake up and either build competing products or acquire you. In healthcare IT, that window is everything. If you can sign up fifty seventy-five agencies in that window and prove out the economics, you become t category leader and everyone else is playing catch-up.

## The Technical Architecture and Data Moat

Let me talk about how you would actually build this thing because the technical details matter for understanding defensibility. The core asset is the predictive m and building a good model requires three things: training data, features, and a feedback loop.

Training data comes from two sources. The easy source is Medicare claims data, you can access through the CMS data sharing programs or by partnering with a aggregator. Claims give you diagnosis codes, prior utilization, demographic info payment amounts. The harder source is clinical data from OASIS assessments, w contain functional status measures, wound characteristics, medication lists, and bunch of other stuff that actually predicts care needs. OASIS data is not centrall available, so you have to get it directly from agencies.

This creates a classic cold start problem. Agencies will not give you their data u you can prove value, but you cannot build a good model without their data. The around this is to build a minimum viable model using only claims data and publ available OASIS benchmarks, then improve the model iteratively as you add age partners and accumulate proprietary clinical data. The first version of the model not be great but it will be good enough to generate some savings, and every episo you process makes the model better.

Features are where the real IP lives. A naive approach would just use diagnosis a age to predict therapy needs, which would produce a mediocre model. A sophist approach incorporates things like prior home health utilization patterns, specifi OASIS items that correlate with therapy response, social determinants data from external sources, and interaction effects between different variables. You also wa segment by geography because therapy utilization norms vary wildly by region.

The feedback loop is critical for moat-building. Every time an agency uses the platform and completes an episode, you get an outcome observation. Did the pat hit their functional goals? Did they have an adverse event? Did they get readmit This lets you continuously retrain the model and improve accuracy. After proces few hundred thousand episodes, your model is materially better than what any n entrant could build, because they do not have access to your proprietary outcom data.

The other technical component worth mentioning is the integration layer. Home health agencies use a dozen different software systems for clinical documentatio scheduling, billing, and so on. Your platform needs to ingest data from those sys generate recommendations, and push outputs back into clinician workflows. Bu those integrations is tedious work but it creates meaningful switching costs onc agency is live.

One important architectural decision is whether to build the remote monitoring home health aide coordination components in-house or partner with existing ve My instinct is to partner initially and build selectively over time. Remote monitc is a commodity at this point with dozens of vendors offering similar products, so is no reason to reinvent that wheel. Home health aide coordination is more inter and potentially defensible, but it is also operationally complex and distracts fror core value proposition. Start with partnerships and build only what you need to control the customer relationship and capture the economics.

## Unit Economics and Market Sizing

Let me walk through the math on this business because healthcare investors car about unit economics and path to profitability. The core cost structure is softwa development, data acquisition, and customer success. Revenue is performance fe from agencies.

Development costs are front-loaded. You probably need a team of fifteen to twer people in year one: a few data scientists, several software engineers, a clinical tea validate the care protocols, and some customer success folks to manage agency relationships. Fully loaded cost is maybe $200-250K per person, so $3-4 million year one burn. Add in data acquisition costs, cloud infrastructure, and general overhead and you are looking at $5-6 million to get to market.

Customer acquisition is relatively capital-efficient because you are selling into a concentrated market. There are about 11,000 Medicare-certified home health ag in the US, but the top 500 agencies represent over 60% of total volume. You can those 500 agencies through direct sales and industry conferences without buildi massive sales org. Assume you need a VP of sales and three or four account execu plus some marketing spend. Maybe $2 million annually in fully loaded sales and marketing costs.

On the revenue side, a typical mid-sized agency does 300,000 to 500,000 episode year across their footprint. If you are capturing $10 per episode in performance that is $3-5 million in annual revenue per customer. Get to twenty customers in two and you are at $60-100 million in revenue. Gross margins should be 75-80% scale because the product is software-driven with minimal variable costs.

The key question is how fast you can sign up agencies and what percentage of th volume you can capture. In the best case, you land a large national chain like Amedisys or Encompass and they roll you out across their entire footprint in twe months. That is a single sale that could be worth $50 million in annual revenue. realistically, you start with regional agencies and mid-sized chains, sign up ten t fifteen customers in year one, and expand from there.

Market sizing depends on how you define the addressable opportunity. If you lir to traditional Medicare home health, you are looking at about 3.5 million benefi receiving home health annually, generating around 7 million episodes at an aver payment of $2,000 per episode. Total market is $14 billion in Medicare payment you can drive $15-20 in savings per episode and capture half of that value, the ar revenue opportunity is $50-70 million per percentage point of market share.

But the real opportunity is bigger than traditional home health. Medicare Advar plans are increasingly pushing members into home-based care as a way to manag costs. MA plans have much stronger incentives to reduce unnecessary utilizatior which means they would be highly motivated to mandate or incentivize use of a platform like this. If you can get a few large MA plans to make your platform a preferred partner for home health utilization management, you could potentially capture 20-30% of the MA home health market, which is growing at 15-20% ann

There is also a secondary market in analytics and benchmarking. Once you have processed millions of episodes and built a proprietary dataset on therapy utiliza patterns and outcomes, you can sell insights back to agencies, payers, and CMS i This is not huge revenue initially but it becomes a nice margin business at scale.

## Competitive Landscape and Defensibili

The competitive question comes down to who else could build this and why they not already. The obvious potential competitors are the existing home health soft vendors, the prior authorization platforms, and the big health IT companies like or Oracle.

Home health software vendors like Axxess, WellSky, and Homecare Homebase h relationships with thousands of agencies and could theoretically add utilization management features to their platforms. But these companies are fundamentally workflow and billing businesses, not clinical decision support businesses. Their competency is making it easy for agencies to document visits and submit claims, optimizing care plans. They also have a conflict of interest because many of then

make money on a per-visit basis through their billing products. Recommending visits would reduce their revenue.

Prior authorization platforms like Cohere Health have the technical and operati playbook for utilization management but they have not moved into post-acute ca They focus on high-cost procedures and specialty drugs where the savings per ca much larger. Home health is a lower-margin, higher-volume opportunity that re a different go-to-market motion. Could they expand into this space? Absolutely. they would be starting from scratch in terms of agency relationships and clinical protocols.

The big health IT vendors could build competing products but they move slowly tend to focus on large hospital systems rather than post-acute providers. Epic ha home health module but it is basically an afterthought compared to their inpatie and ambulatory products. Oracle bought Cerner which has basically no presence home health. This is not a market that the big strategic players are focused on.

The more interesting competitive threat is agencies building internal solutions. chains like Amedisys or LHC Group have the scale to hire data scientists and bu their own predictive models. But most agencies are focused on operational execu not product development. Building a sophisticated machine learning platform requires a totally different skill set and culture than running a home health ager My guess is that a few large players will try to build internally and most will con it is not worth the distraction.

Defensibility comes from four sources. First is the data moat I described earlier. you have processed millions of episodes with outcomes data, you have a dataset t nobody else can replicate. Second is integration and switching costs. Once an ag has integrated your platform into their clinical workflows and trained their staff moving to a competitor is painful. Third is regulatory alignment. As CMS contir to evolve payment models and quality measures, you can stay ahead of those cha and maintain product-market fit while competitors play catch-up. Fourth is netv effects. As you add more agencies, you can benchmark performance across organizations and provide better insights to everyone in the network.

None of these moats are permanent or unassailable. Healthcare is full of exampl companies with strong market positions that got disrupted by new entrants or regulatory changes. But if you execute well, you can build a business that is defe for five to seven years, which is plenty of time to generate significant value for investors.

## The Path to Exit

Exit strategy matters for angel investors, so let me talk through the most likely p The boring but realistic exit is acquisition by a strategic buyer in three to five ye The most obvious acquirers are the large home health chains that want to bring utilization management in-house or the existing software vendors looking to add clinical decision support capabilities. Amedisys, Encompass, and LHC Group ha been acquisitive over the past few years. Homecare Homebase sold to Advent International for over $1 billion in 2021, showing that private equity is willing t for market-leading home health tech assets.

Valuation in a strategic exit would probably be based on a revenue multiple give you would still be in growth mode and potentially not profitable. Comparable healthcare IT deals in the utilization management space have traded at 8-12 time forward revenue. If you can get to $80-100 million in revenue with strong growth retention metrics, a $600 million to $1 billion exit is plausible. That would be a nice outcome for early investors.

The more exciting but less likely path is staying independent and building towa IPO. The public market comps for healthcare IT services businesses like Health or Cotiviti are in the 4-6 times revenue range, which is lower than private marke valuations. You would need to get to $300-400 million in revenue with a path to profitability before an IPO makes sense. That is probably a seven to ten year tim

There is also a potential private equity recap scenario where you bring in a PE p after the business is mature and take some chips off the table while continuing t grow. This is attractive if you think the market opportunity is large enough to su building a much bigger company but you want to de-risk your personal position

The least likely but most transformative outcome is that CMS decides they want mandate utilization management for home health the way they mandate prior authorization for certain services. If that happens and your platform becomes th category leader, you could have a $5-10 billion outcome on your hands. I put tha probability at under 5%, but it is worth keeping in the back of your mind.

One important point on exit timing: healthcare regulatory cycles matter a lot. Th therapy cap rule takes effect in 2026, which means agencies will start seeing the financial impact in late 2026 and really feel the pain in 2027. That is when dema solutions will peak. You want to be in market by early 2026, scale aggressively through 2027 and 2028, and be positioned for an exit in 2029 or 2030 when you h demonstrated sustained value creation and locked in major customers.

The worst possible outcome is being too early or too late. Too early means you a trying to sell a solution before agencies feel the pain, which leads to long sales c and low conversion. Too late means competitors have already established marke position and you are fighting for scraps. The window for this opportunity is prob 2025 through 2027. After that, the market will have consolidated around a few winners.

Healthcare investing is all about finding regulatory arbitrage opportunities and building businesses that help providers navigate the gap between old incentives new requirements. The 2026 home health rule creates exactly that kind of gap. Agencies know they need to change their care models but they do not have the te expertise to do it effectively. That is a massive opportunity for a well-designed, thoughtfully executed technology platform that aligns incentives and delivers re value. The companies that win in healthcare are the ones that make providers' li easier while improving economics and outcomes. This business model does all th

← Previous

Nex

## Discussion about this post

Comments Restacks

Write a comment...