The Unfair Advantage Nobody Talks About: How Skipping BAAs Unlocks Venture-Scale Growth in Health Tech

MAR 29, 2026 • PAID

Abstract

This essay argues that the most fundable, scalable health tech companies of the decade will be built around business models that structurally avoid HIPAA's bus associate agreement requirement during their high-growth phases. OpenEviden the case study – it went from zero to $50M ARR, then $150M ARR, then a $12B valuation in roughly three years by operating as a physician-facing knowledge to rather than a PHI handler, letting it grow like consumer software while healthca enterprise companies were stuck in procurement hell. But OpenEvidence isn't tl story. The deeper thesis is that the highest-upside opportunities remaining in he tech share a common architecture: you're reallocating dollars, not documenting This essay covers three specific categories – employer healthcare control planes, patient-controlled health graph infrastructure, and healthcare financial rails – a explains the product architecture, go-to-market physics, and defensibility dynan each. Key data points and framing: OpenEvidence hit 40% of US physicians and monthly consultations in roughly 24 months; self-insured employer spend excee $1T annually; US healthcare payments exceed $4T/yr; and none of the three cate described require real-time PHI ingestion during their core operating motion.

Table of Contents

- The Structural Thesis: Reallocating Dollars vs. Documenting Care

- Category 1: Employer Healthcare Control Planes

- Category 2: Patient-Controlled Health Graph Infrastructure

- Category 3: Healthcare Financial Infrastructure

- Where the Real Moats Are

- How to Bet

The BAA Is a Moat in Reverse

Most people in health tech treat the business associate agreement as just anothe contract to get through. Sign it, check the box, move on. What they're missing is the BAA isn't just a legal document – it's a gravitational force that slows everyth down by roughly 12 to 24 months and turns your sales cycle into a compliance negotiation. Every health system that needs to sign one has a legal team, an infos team, a vendor assessment process, and a procurement function that exists speci to delay your deal. Epic's App Orchard review alone takes six to nine months jus get listed. Then each health system has to individually enable your product. The legal review. Then BAA signature. Then IT integration. By the time you're in production, you've burned cash, diluted your cap table, and watched a better-capitalized competitor catch up. The BAA is healthcare's version of a moat – exc protects the incumbents, not the new entrant.

This isn't some niche edge case. It's the fundamental reason health tech has historically underperformed on a risk-adjusted basis relative to other enterprise software verticals. The capital required to survive a health system procurement c is enormous. The revenue that comes out the other end is lumpy and contract-dependent. And the whole thing is brittle – one renewal risk can crater your AR Investors know this. Operators know this. Yet the field keeps producing compar that walk straight into the BAA trap because they think the clinical data access worth it.

Sometimes it is. EHR integrations, clinical AI tools that need real-time patient d anything touching identified health records at the point of care – those compani have to play the game. The question is whether your idea actually requires that, whether you've been pattern-matching to existing health tech archetypes withou thinking hard about what data you truly need to build value. A surprising numb genuinely large opportunities in healthcare turn out to require no PHI at all. Th companies that figure this out early get to grow like software companies instead healthcare companies. The difference in outcomes is dramatic.

OpenEvidence and the Permission Expl

OpenEvidence is the clearest recent proof point. The company launched in 2023 built a physician-facing clinical knowledge tool – essentially an AI search engine grounded in peer-reviewed medical literature. The key structural decision, whet intentional from day one or emergent, was to build a tool that physicians query u clinical questions rather than patient records. No PHI flows through the system during the core use case. A physician types "what's the current evidence on tirzepatide for HFpEF?" and gets a synthesized answer drawn from NEJM, JAM the broader literature. The physician never inputs a patient name, a DOB, an MI diagnosis code. The system never touches a covered entity's PHI as part of its sta operating motion.

That distinction sounds technical and boring. The growth implications were any but. Because no PHI was flowing, no BAA was required for the base product. Individual physicians could sign up, verify their NPI, and start using the tool the day. There was no procurement cycle. No legal review. No infosec questionnaire. enterprise negotiation. The product grew word-of-mouth between doctors the sa way a B2C app grows – doctor tells colleague, colleague signs up, colleague tells more. This is the distribution dynamic that virtually every health tech entrepren envies and almost none of them achieve because their product architecture force them out of it.

The numbers are hard to argue with. By mid-2024, roughly 358,000 verified US physician consultations were happening per month. A year later, that number wa million monthly consultations and climbing, representing over 2,000% year-over growth. The company was registering over 65,000 new verified clinicians per mo reached 40% of US physicians. Sacra pegged annualized revenue at $150M in 202 from $7.9M in 2024, with 90% gross margins. The monetization model is pharmaceutical and medical device advertising at CPMs that can reach $1,000 or more, compared to $5 to $15 for typical social media – a function of the audienc quality and the decision context. In January 2026, OpenEvidence raised $250M a $12B valuation co-led by Thrive and DST. That's a full consumer-scale growth cu in healthcare, achieved without a single enterprise procurement cycle during the hyper-growth phase.

The irony is that OpenEvidence does offer BAA-based access for covered entities want to input PHI as part of their workflow. But that's the expansion layer, not t acquisition motion. The core growth engine was the no-PHI, no-BAA, word-of-m viral loop. That's the thing to study. Most health tech founders build the BAA-required product first and then try to add virality on top. OpenEvidence inverted Start viral, then add the enterprise layer when you have leverage.

The Structural Thesis: Reallocating Dollars vs. Documenting Care

OpenEvidence is a specific case in a broader pattern. The clearest articulation of this: health tech companies that document, coordinate, or improve care delivery almost always require PHI. Health tech companies that reallocate, optimize, or arbitrage the financial flows underneath care delivery often don't. The first categ includes EHR-adjacent tools, care management platforms, clinical decision supp that needs real-time patient context, remote monitoring, and most of what peopl traditionally think of as "health IT." The second category includes employer ben infrastructure, healthcare financial rails, market transparency tools, data infrastructure that works on de-identified or aggregated datasets, and anything

plugs into the financial layer of healthcare without needing to know that it's Mr
Jones in bed 4.

This matters because healthcare's financial layer is both enormous and almost comically inefficient. US healthcare spend is over $4T annually. Self-insured employers alone are responsible for roughly $1T of that. The payment system is with latency – anywhere from 30 to 120 days from service to settlement in many Provider cash flow is volatile because of denial rates. Price discovery is broken. A the data infrastructure that sits underneath all of it was mostly built in the 1990 These are fintech-scale problems sitting inside a healthcare wrapper. And unlike clinical care delivery, they don't fundamentally require patient-level PHI to solv

The three categories below are the most fundable expressions of this thesis right Each has specific architectural patterns that support the no-BAA growth motion specific buyer dynamics that differ from traditional health IT sales, and specific building properties that separate the defensible from the commodity.

Category 1: Employer Healthcare Contr Planes

Self-insured employers collectively spend somewhere north of $1T per year on healthcare benefits, and most of them have essentially no real-time visibility or control over where that money goes. The TPA processes claims post-hoc. The PE optimizes for its own margin. Point solutions multiply. By the time a CFO sees a quarterly claims report, the economic outcome is already locked in – the utilizat happened, the procedures were done, the unit prices were whatever the carrier's network rates happened to be. There's no mechanism to intercept that spend in time and redirect it to better economic outcomes.

That gap is the product opportunity. The right frame isn't "benefits administrat software" or even "benefits navigation." It's a real-time control plane that sits be the employer's benefit design and the moment a member enters the healthcare sy At the architectural level, this requires three layers. The first is an economic gra normalized pricing across providers built on allowed amounts rather than

chargemaster fiction, procedure-level bundled cost estimates built on CPT clust logic, and longitudinal provider performance data that blends cost and outcome proxies. The second is a decisioning engine – real-time routing logic that can ide when a member is about to have an MRI at a $3,200 facility when an equivalent c available for $400, and surface that decision with a member-facing incentive to c behavior. The third is a control surface for the employer – scenario simulation ("what's the financial impact if we steer 20% of orthopedic volume to these three facilities?"), budget allocation tools, and contract intelligence that identifies wh the carrier is extracting margin.

None of that architecture requires touching identified clinical records in real tir The economic graph can be built on claims-derived datasets, which are de-ident or used in aggregate, plus public pricing data, plus provider performance data fr CMS and similar sources. The routing and incentive layer operates on administr data about what procedure is being ordered and where. The employer UI works entirely on financial and utilization analytics. The system can run almost entirel outside HIPAA's PHI definitions during its core operating motion. The BAA que may arise at the edges – particularly if the platform integrates with a TPA and st processing individual member claims in identifiable form – but the core product launch and growth phase can happen without it.

The go-to-market dynamics here are fundamentally different from clinical healt The buyer is the CFO or the head of benefits, not the CMO or the clinical inforr team. The conversation is about financial outcomes, not care quality metrics. Th sales pitch is "we will reduce your healthcare spend by 8 to 15 percent in year or with a performance-based fee structure, not "we will improve clinical outcomes 36-month outcomes study period." This is an ROI sell that can close in one budg cycle, not a clinical outcomes sell that requires multi-year data and a hospital committee vote. The expansion motion is also financially logical: start with steer and navigation, expand into network redesign and direct contracting enablemen eventually own the full benefits orchestration layer. Each step adds value to the employer and adds switching cost.

The companies most frequently cited in this space – Collective Health, Transcar and their cohort – have built pieces of this but haven't assembled a true real-tim control plane architecture. Most of the existing navigation tools are static recommendation engines, not programmable routing infrastructure. The moat i strong version isn't the navigation UI – that's a commodity. The moat is the proprietary provider pricing graph built from actual transaction data, the reinforcement learning loop on steerage effectiveness across a large member population, and the embedded financial workflows that process incentive payme and premium offsets. When your platform becomes the system that determines v employer healthcare dollars flow, you have a defensible position. Until then, you one of thirty navigation apps in the market.

Category 2: Patient-Controlled Health Graph Infrastructure

This one is consistently misunderstood because founders keep trying to build it consumer app rather than as infrastructure. The consumer app version – "give patients control of their health records, maybe add a wellness dashboard" – has tried roughly 200 times since 2010. It doesn't work, for reasons that are pretty w understood: patients don't have enough intrinsic motivation to maintain a perso health record, the data acquisition story requires deep EHR integration that reintroduces all the enterprise friction you were trying to avoid, and there's no obvious monetization. The infrastructure version of this same idea is a complete different business.

The infrastructure version starts from a different problem statement. Healthcar no canonical, persistent identity layer. MRNs are provider-scoped and don't por Payer IDs churn. Device data from wearables is siloed in consumer cloud platfor Lab results are trapped in portal UXs that nobody uses. The result is that longit health data – the kind that would actually power useful clinical AI, real trial recruitment, accurate risk stratification, and genuine cross-provider coordinatio doesn't exist in usable form for most patients. The identity fragmentation probl means that even if you could pull all the data together, you couldn't reliably stit

into a single patient record without probabilistic matching logic that most healt companies haven't invested in.

Building the infrastructure layer means solving four things in sequence. First, identity resolution across claims, labs, device data, and pharmacy records using deterministic and probabilistic matching – this is the hardest part technically ar thing most teams underestimate. Second, data normalization into a unified sche where labs map to LOINC, medications map to RxNorm, and imaging metadata handled correctly, with temporal alignment solved properly. Third, a permission engine that lets the patient grant fine-grained, tokenized access to specific data elements for specific use cases – not a blanket records release, but a surgical acc grant that lets a pharma company query for trial eligibility without receiving rav records. Fourth, an API layer that enterprise buyers actually pay for: trial recrui endpoints, risk scoring feeds for employers and payers, data access for clinical A developers.

The reason this business can avoid BAAs in its growth motion is structural. The is user-authorized and user-controlled. The company isn't acting on behalf of a covered entity – it's acting as an agent of the patient, under the patient's explicit authorization. That's a different regulatory posture than a health IT vendor that handles PHI on behalf of a hospital. It's the architecture that Plaid uses for finar data, and it's the reason Plaid was able to grow without being a licensed financia institution. The analogy isn't perfect – health data has more sensitivity than ban transaction data – but the regulatory structure supports a similar go-to-market l if you build it correctly.

The counterintuitive GTM move here is to start on the demand side, not the sup side. Don't start by acquiring patients. Start by acquiring enterprise buyers who a desperate need for the data: pharma companies running trial recruitment, digi health companies that need longitudinal patient data for their AI models, emplo doing risk stratification. Lock in the demand-side contracts first, then use those contracts as both economic justification and incentive design for patient onboar Patients have a much clearer reason to join a platform when there's a tangible be attached – trial eligibility screening with a cash incentive, for instance – than wh

the pitch is generic health record portability. The supply-demand sequencing m enormously here and is where most teams get the strategy backwards.

If you get both sides working, this is one of the few true winner-take-most dyna in health tech. Network effects are real – more users produce more valuable longitudinal datasets which attract more enterprise buyers which fund better incentives for users. Switching costs on the enterprise side are deep – once your is integrated into a pharma company's trial recruitment pipeline or an AI compa training process, replacing it is expensive and disruptive. And the data gravity compounds over time in a way that's hard for a late entrant to replicate, because longitudinal data is inherently time-dependent. A company that has three years longitudinal data on a patient cohort is not replicable by a company that starts t

Category 3: Healthcare Financial Infrastructure

Healthcare is a $4T-plus annual payment system that processes transactions on infrastructure designed in roughly 1994. Payment latency runs 30 to 120 days in specialty scenarios. Provider cash flow is volatile because payer denial rates can to 20 percent on first submission. Patient responsibility has grown as a percenta total collections as deductibles have risen, creating a receivables management problem that most providers lack the technology to handle. Price certainty at th of service is essentially nonexistent outside of a handful of cash-pay contexts. Th are not care delivery problems. They are financial infrastructure problems that happen to occur inside a healthcare context.

Building the financial infrastructure layer means tackling this in three modular components. The first is payments – real-time or near-real-time adjudication an payment rails, provider-side collections optimization, and patient financing emb at the point of care. The second is credit and liquidity – advance payments again expected reimbursements, working capital lines for providers that are calibrated their specific payer mix and denial rate history, and risk-adjusted lending using utilization patterns as underwriting inputs. The third is pricing and risk – the a

to predict episode cost before the episode occurs with enough accuracy to under a bundled payment or offer a fixed-price surgical package outside the payer cont framework.

None of this requires PHI. The payments layer operates on financial transaction and billing metadata. The credit and liquidity layer underwrites based on aggreg utilization patterns, historical payer behavior, and denial rate data – all of which be handled at a practice or facility level rather than a patient level. The pricing a risk engine needs claims data to build actuarial models, but claims can be de-identified or aggregated for that purpose. The core operating motion of this bus runs entirely outside HIPAA's PHI definitions. There are edge cases – patient financing that involves individual billing records, for instance – where BAA que surface, but those are scoped to specific product modules rather than the whole business.

The GTM wedge is specialty providers with acute cash flow pain: ambulatory su centers, oncology practices, dental groups. These buyers have high revenue concentration, significant denial exposure, and limited access to traditional worl capital products because bank lenders don't understand healthcare revenue cycl enough to underwrite it correctly. Landing there with a combination of faster payments plus working capital unlocks a very fast payback cycle – providers can financial impact within 60 to 90 days, which drives strong referral dynamics and word-of-mouth within specialty networks. The expansion path runs upmarket through the MSO and PE-backed rollup layer, where the financial infrastructure opportunity scales with the consolidated provider group, and eventually into the and employer side where risk-sharing contract enablement creates new product surface area.

The incumbents here are surprisingly weak given the market size. Legacy clearinghouse infrastructure is fragmented and built on antiquated rails. RCM software companies have the workflow but not the financial product layer. No or built the full-stack financial platform for healthcare that is analogous to what m fintech did for SMB payments and lending. The moat in the strong version come from embedded payment rails that are difficult to rip out once integrated,

underwriting models trained on healthcare-specific transaction data that improv with scale, and the two-sided network effects that emerge when you have enough providers and enough capital partners to create a real marketplace dynamic. The analogy to Stripe, Brex, or AvidXchange is apt – none of those companies won be they were the cheapest option. They won because they became the infrastructure that everything else depended on.

Where the Real Moats Are

Across all three categories, the pattern for durable moat-building is the same. Commodity versions of each idea exist already – navigation tools, personal health record apps, RCM software. Those are built on static data, lack reinforcement lo and can be replicated or commoditized by a better-funded competitor. The dura versions are built on proprietary data assets that compound over time, network e that make the platform more valuable with each additional participant, and emb workflows that create switching costs independent of product quality. The empl control plane moat is the provider pricing graph and the steerage effectiveness learning loop. The patient health graph moat is the longitudinal identity resolut dataset and the enterprise integration depth. The financial infrastructure moat i underwriting model trained on healthcare-specific transaction history and the embedded payment rails.

The regulatory architecture matters enormously for how these moats develop. A company that avoids BAAs during its high-growth phase gets to scale its data as and its network before the compliance burden kicks in. By the time BAA requirements become relevant – typically when the enterprise expansion motion begins and the product starts handling identified PHI – the company already has leverage. It has the data, the distribution, the brand, and the customer relationsl The BAA is still an annoying contract, but it's not a growth constraint anymore. OpenEvidence's playbook is to grow to 40% physician penetration on the no-BA product, then offer BAA-based access as the enterprise upsell. That sequencing right model.

How to Bet

For investors putting capital into this thesis, the ranking on risk-adjusted convic looks roughly like this. Employer control planes have the fastest path to revenue because the buyer is a CFO making an ROI decision, the deal sizes are significar the product development risk is moderate rather than high. You're mostly build top of data that already exists – claims, provider pricing, utilization patterns – ar packaging it into a control interface that displaces the incumbent TPA analytics nobody uses. Financial infrastructure is the most underserved category relative t market size, with the strongest fintech analogs and the highest confidence that durable moats are achievable if execution is strong. Patient-controlled health gra infrastructure is the highest-upside and highest-execution-risk bet – the networ effects are real and the TAM is enormous, but solving the cold start problem on the user and enterprise side simultaneously is genuinely hard, and the identity resolution technical challenge is regularly underestimated by teams entering the space.

The throughline across all three is the same structural shift happening in paralle healthcare is moving from a documentation-and-reimbursement architecture tov programmable economic infrastructure architecture. The companies that get to meaningful scale without getting trapped in BAA-driven enterprise sales cycles be positioned to build the financial control layers, data substrates, and payment that everything else runs on. That's not a point solution outcome. That's an infrastructure outcome. And infrastructure, when done right, compounds.

2 Likes • 1 Restack

← Previous

Discussion about this post

Comments Restacks

Write a comment...