


The Peptide Economy vs the Healthcare AI Economy: Which Side of the Trade Matters More

APR 05, 2026



Share

Abstract

This essay examines the emerging competitive dynamics between two massive healthcare market forces, GLP-1 receptor agonists and adjacent peptide therapies on one side and healthcare-focused artificial intelligence on the other, and argues framing them as separate markets fundamentally misunderstands where value will concentrate over the next decade. Key questions addressed include:

- Which market, peptides or healthcare AI, will generate more absolute economic value by 2035
- GDP impact modeling for both categories across the US and China
- Regulatory divergence between US and Chinese approaches to both peptides and healthcare AI
- The oral peptide transition and its downstream effects on distribution, pricing, and market access
- Labor market implications for clinical and administrative healthcare workers
- Commodity risk analysis across both categories
- Profit center mapping for hospitals, payers, employers, pharma, and the direct-consumer wellness segment

- Actionable opportunity identification for healthcare entrepreneurs and angel investors

Table of Contents

Two Trades Walk Into a Bar

Market Sizing: The Numbers That Actually Matter

GDP Math: Peptides Think Bigger Than You Think

USA vs China: Different Playbooks Same Endgame

Regulatory Velocity: Who Is Moving Faster and Why

Oral Peptides Change Everything

Labor Impacts: Who Loses Their Job and When

Commodity vs Moat: What Survives the Next Five Years

Where Profit Centers Actually Land

The Entrepreneur Opportunity Map

Convergence Is the Real Trade

Two Trades Walk Into a Bar

There is a conversation happening right now in every health tech investor group every LP meeting, every pitch session, and it goes something like this. Someone says GLP-1s are the biggest thing to happen to healthcare since antibiotics. Someone says no, AI is going to restructure the entire cost basis of medicine. Then everyone argues for forty five minutes and nobody changes their mind. The problem is not one side is wrong. The problem is that almost everyone is analyzing these as independent phenomena when they are deeply, structurally entangled. The pepti

economy and the healthcare AI economy are not competing markets. They are co-dependent systems that will amplify or constrain each other in ways that most investment theses have not yet accounted for. But since people love a horse race, the essay will run the numbers on both before explaining why the real alpha is in the intersection.

Market Sizing: The Numbers That Actually Matter

Start with peptides. The GLP-1 receptor agonist market alone hit roughly \$50 billion in global revenue in 2024, and most credible forecasts have it north of \$130 billion by 2030. That is just the GLP-1 class. Layer in GIP/GLP-1 dual agonists like tirzepatide, the amylin analogs in development, the glucagon receptor agonists, the emerging obesity-adjacent peptide candidates targeting MASH and cardiovascular outcomes, and you are looking at a therapeutic category that could plausibly exceed \$200 billion in annual revenue within seven years. Novo Nordisk and Eli Lilly between them represent something like \$900 billion in combined market cap, and a meaningful chunk of that valuation is the market pricing in decades of peptide dominance. In context, the entire global oncology market was about \$220 billion in 2023. Peptides are on track to rival cancer drugs in total market size. That sentence would have gotten you laughed out of a room five years ago.

Now healthcare AI. This is harder to size because the category is genuinely messy. Morgan Stanley pegged the healthcare AI market at around \$10 billion in 2024 and projected it to \$45-50 billion by 2030. McKinsey has thrown around figures suggesting AI could create \$200-360 billion in annual value across US healthcare through productivity gains and waste reduction. But there is an important distinction between revenue captured by AI companies and value created by AI deployed in healthcare settings. The revenue numbers are modest relative to peptides. The value creation numbers are enormous but diffuse. An AI system that saves a health system \$40 million a year in denied claims recovery does not show up as AI market revenue the same way that semaglutide scripts show up as pharma revenue. This measurement problem makes healthcare AI look smaller than it actually is, which is both a pro-

and an opportunity for investors who know how to look past top-line market size reports.

The honest answer on pure market size is that peptides win on direct revenue and are not particularly close. A single blockbuster drug generates more revenue than most companies will see in a decade. But revenue is not the same as economic impact, that distinction matters enormously for where entrepreneurs should be building

GDP Math: Peptides Think Bigger Than You Think

The GDP impact question is where things get genuinely interesting. Obesity costs the US economy somewhere between \$1.7 and \$2 trillion annually when you account for direct medical spending, lost productivity, disability, and related comorbidities. GLP-1s and next-generation peptides can reduce obesity prevalence by even 15-20 percent over the next decade, you are talking about hundreds of billions in GDP just from the productivity gains alone. Goldman Sachs published research suggesting GLP-1 adoption could boost US GDP by 0.4 percent annually by the early 2030s which sounds small in percentage terms but translates to roughly \$100 billion per year in additional economic output. And that estimate may be conservative because it does not fully model the second-order effects on things like workforce participation rates among previously disabled individuals, reduced caregiver burden, and lower disability insurance payouts.

Healthcare AI has a different GDP transmission mechanism. It does not create economic value primarily through therapeutic outcomes. It creates value through deflation and labor productivity. If AI can reduce administrative costs in US healthcare by 20-30 percent, that is roughly \$200-300 billion per year in savings from a roughly \$1 trillion administrative spend base. But healthcare cost reduction has a complicated relationship with GDP because healthcare spending itself is a component of GDP. Reducing it makes people healthier and frees up capital for other uses, but the accounting gets weird. The productivity channel is cleaner. If AI enables clinical

to see 20 percent more patients or reduces diagnostic error rates, those gains flow more directly into economic output.

On net, peptides probably have the larger near-term GDP impact because the burden is so massive and the therapeutic effect is so direct. AI has the larger long-term GDP impact because it compounds. A drug treats one condition. An AI system that reduces diagnostic latency or optimizes care pathways creates value across multiple conditions simultaneously. But “long-term” in this context might mean 15-20 years which is longer than most fund cycles.

USA vs China: Different Playbooks Same Endgame

The US and China are approaching both of these markets with fundamentally different strategies, and understanding those differences is critical for anyone making capital allocation decisions.

In peptides, the US has a massive head start. Novo Nordisk is Danish but the US is its largest market, and Eli Lilly is American. The US regulatory infrastructure, payer landscape, and provider distribution channels are all built to support branded specialty pharmaceuticals at scale. China is playing catch-up but doing it aggressively. There are over a dozen Chinese biotechs with GLP-1 candidates in clinical development, and several are already in Phase 3 trials. Companies like Innovent Biologics and Hengrui Medicine are developing both injectable and oral formulations. China's strategy is classic: let the US companies prove the biology, then compete on manufacturing cost and domestic market access. The NMPA has been accelerating review timelines for obesity drugs, partly because China's obesity rate has roughly tripled over the past two decades and is now a genuine public health crisis. But Chinese peptide companies face real barriers in Western markets, including regulatory complexity, quality perception issues, and the formidable patent estate that Novo and Lilly have constructed.

In healthcare AI, the competitive landscape looks different. China is not just catching up, it is genuinely competitive. Chinese healthcare AI companies have advantages

American companies do not, starting with data access. China's hospital systems generate enormous volumes of structured clinical data and the regulatory barriers using that data for model training are dramatically lower than in the US. The lack of a HIPAA equivalent (or rather, the existence of data protection rules that are enforced very differently in practice) means Chinese AI companies can train on datasets that American companies cannot legally touch. China has also been more aggressive in deploying AI in clinical settings. NMPA has approved AI-assisted diagnostic tools at a pace that makes the FDA look glacial. Infervision, for example, had AI diagnostic tools deployed across hundreds of Chinese hospitals before most American health AI companies had finished their first clinical validation study.

The scorecard looks roughly like this. In peptides, the US leads on innovation and market capture, China competes on cost and domestic scale. In healthcare AI, China leads on data access and deployment velocity, the US leads on foundational model capability and premium market monetization. Neither country has a decisive advantage across both categories, which creates interesting dynamics for cross-border investment strategies.

Regulatory Velocity: Who Is Moving Faster and Why

US regulation on peptides is actually moving faster than most people expected. The FDA approved tirzepatide for obesity on a relatively aggressive timeline, and the agency's willingness to consider cardiovascular outcome data as a basis for expansion indications has been notable. CMS coverage decisions have been more mixed, with Medicare still not covering anti-obesity medications under most Part D plans as early as 2025, though legislative efforts to change that have gained bipartisan traction. The real regulatory bottleneck in the US is not FDA approval, it is payer coverage and formulary access. Getting a drug approved is one problem. Getting it paid for at scale is a completely different problem, and the US has not solved the second one yet.

On healthcare AI, the FDA has cleared over 900 AI-enabled medical devices as of early 2025, but most of those are relatively narrow applications like radiology triage or

interpretation. The regulatory framework for more ambitious AI applications, like autonomous clinical decision support or AI-driven treatment planning, remains underdeveloped. The FDA's proposed regulatory framework for AI/ML-based software treats these tools more like traditional medical devices than like the continuously learning systems they actually are, which creates friction for companies trying to deploy models that improve over time. There are signs this is changing, the agency signaled interest in a "predetermined change control plan" approach that would allow certain types of model updates without requiring new submissions, but the implementation details are still being worked out.

China's NMPA has taken a more permissive stance on both fronts. For peptides, the agency has created accelerated pathways for drugs that address unmet domestic needs, and obesity now firmly qualifies. For AI, China's approach has been to regulate outputs rather than processes, meaning the agency cares more about whether a diagnostic AI gets the right answer than about how the model was trained. This is philosophically different from the FDA's approach, which tends to scrutinize the entire development lifecycle. The Chinese approach enables faster deployment but arguably creates more risk of poorly validated tools reaching patients. Which regulatory philosophy produces better long-term outcomes is genuinely unknown; anyone who tells you they know the answer is selling something.

Oral Peptides Change Everything

The transition from injectable to oral peptide formulations deserves its own section because the market dynamics shift dramatically. Oral semaglutide already exists in the form of Rybelsus, but its bioavailability is low (roughly 1 percent) and it requires fasting and specific dosing protocols that reduce adherence. The next generation of oral peptide delivery technologies, including things like permeation enhancers, prodrug reformulations, and novel capsule designs, could push oral bioavailability into ranges that make the injectable vs oral efficacy gap much smaller.

When that happens, several things change simultaneously. First, the addressable market expands dramatically because the psychological barrier to treatment drops

Many patients who would never self-inject will take a pill. Analyst estimates suggest oral formulations could expand the total addressable patient population by 40-60 percent. Second, the distribution channel shifts. Injectable GLP-1s flow through specialty pharmacy and require cold chain logistics. Oral peptides can move through standard pharmacy distribution, which reduces costs and complexity. Third, and is the part most people miss, oral formulations commoditize faster. It is harder to differentiate a pill than an injection device. The moment you are competing on a tablet form factor, generic and biosimilar competition becomes much more straightforward, and pricing power erodes faster.

For entrepreneurs, the oral peptide transition creates opportunities in formulation technology, last-mile delivery, adherence monitoring, and companion diagnostics: companies that figure out how to identify optimal responders, titrate doses using time biomarker data, and manage the side effect profile through AI-driven personalization will capture value that the peptide manufacturers themselves can't easily internalize. This is one of the clearest examples of the peptide-AI convergence thesis.

Labor Impacts: Who Loses Their Job and When

Both peptides and healthcare AI have significant labor market implications, but operate through different mechanisms and on different timescales.

Peptides affect healthcare labor through demand reduction. If GLP-1s and related therapeutics meaningfully reduce obesity prevalence, the downstream effect on surgical volumes is substantial. Bariatric surgery is the obvious one, that market shrinks by 30-40 percent, but the cascade extends to orthopedic procedures (fewer knee and hip replacements in younger patients), cardiac interventions (fewer stents, fewer bypasses), and even sleep medicine (reduced CPAP utilization as sleep apnea prevalence drops). Each of these represents a labor pool. Fewer bariatric surgeries means fewer bariatric surgeons, fewer anesthesiologists dedicated to those cases, fewer surgical techs, fewer post-op nurses. The healthcare workforce does not ac

quickly to demand shifts like this. Medical training pipelines have 7-15 year lag. Hospitals that built service lines around surgical volumes that are now declining are at real strategic risk.

Healthcare AI affects labor through task automation. The clearest near-term impact is on administrative roles. Revenue cycle management, prior authorization, coding, claims processing, and scheduling are all workflows where AI can reduce headcount or at minimum reduce headcount growth. A large health system might employ 50,000-800 people in revenue cycle roles. If AI can automate 40-50 percent of those tasks over the next five years, that is 200-400 positions that do not get backfilled when people leave. Clinical roles are more insulated in the near term, but the medium-term picture is more complex. AI scribes are already reducing the need for human medical scribes. AI-assisted diagnostic tools will not eliminate radiologists but they will change the ratio of radiologists to imaging studies, which means fewer new radiology positions than historical growth rates would suggest.

The combined effect of both trends is a healthcare labor market that gets structurally smaller relative to population over the next decade. That is a big deal in an industry that employs roughly 17 million Americans and has been one of the most reliable sources of middle-class job growth for decades.

Commodity vs Moat: What Survives the Next Five Years

In peptides, the drug molecule itself is increasingly a commodity. There are dozens of GLP-1 receptor agonists in development globally. The biology is well understood and manufacturing is complex but not proprietary in a way that prevents competition once patents expire. What retains pricing power is clinical data (the outcomes trials that support specific indications), device/delivery innovation (autoinjectors, oral formulations with differentiated absorption profiles), and brand equity with prescribers and patients. Novo Nordisk's moat is not semaglutide the molecule. It's the STEP trial program, the Ozempic brand, the FlexPen device, and the manufacturing scale they built over decades. But even those moats are time-limited.

Semaglutide biosimilars will arrive, probably in the 2031-2033 timeframe depending on patent litigation outcomes, and when they do, pricing will compress meaningfully.

In healthcare AI, the commodity risk is arguably even more acute. Foundation models are rapidly converging in capability. The difference between a GPT-4 class model and a Claude class model on most healthcare NLP tasks is marginal and shrinking. Vertical AI applications that are essentially wrappers around foundation models with some healthcare-specific prompt engineering are already facing pricing pressure. What retains value is proprietary data assets (clinical datasets that cannot be easily replicated), workflow integration (deep embedding in EHR systems and clinical processes that creates switching costs), and regulatory clearance (FDA-cleared algorithms have a moat even if the underlying technology is replicable, because the clearance process itself takes 12-24 months and significant capital). Companies building healthcare AI that relies primarily on model sophistication are in trouble. Companies building healthcare AI that relies on data network effects and workflow lock-in are in a much stronger position.

The meta-lesson here is that in both categories, the thing that looks like the product (the drug, the model) is not where durable value accretes. Value accretes in the systems around the product, the data, the distribution, the regulatory position, and the clinical evidence. Entrepreneurs who understand this build very different companies than entrepreneurs who think they are in the drug business or the AI business.

Where Profit Centers Actually Land

Hospitals are going to struggle with both trends in the near term. Peptides reduce surgical volumes which are hospitals' highest-margin service lines. A hospital that generates 15-20 percent of its operating margin from bariatric, orthopedic, and cardiac surgical cases is looking at real margin compression if those volumes decrease by even 10-15 percent. Some hospitals will pivot to becoming peptide prescribing and monitoring centers, but the revenue per patient on chronic medication management is a fraction of what a surgical episode generates. On the AI side, hospitals can capture value through operational efficiency, but most health systems are bad at actually

realizing labor cost savings from technology. They tend to redeploy rather than reduce headcount, which means the savings show up as throughput increases rather than as cost reductions. That is fine if you have unmet demand, which most systems do, but it does not flow to the bottom line in the way that a CFO would like.

Payers are in a complicated position on peptides specifically. If GLP-1s genuinely reduce downstream medical costs, then covering them is actuarially rational even at current pricing. The problem is the timing mismatch. You pay for the drug now but don't realize the savings over 5-10 years, but employer-sponsored insurance turns over annually and Medicare has a different budget cycle than a private insurer. Some payers are already experimenting with outcomes-based contracts where the peptide manufacturer shares risk on whether the drug actually reduces total cost of care. These are interesting but operationally complex. On AI, payers are natural buyers because anything that reduces claims cost or improves utilization management directly benefits their economics. The prior authorization automation use case alone is potentially worth billions in aggregate payer savings.

Employers might actually be the biggest winners from both trends. They bear the ultimate cost of both medical spending and productivity losses from chronic disease. If peptides make their workforce healthier and AI makes their healthcare spend more efficient, employers capture value on both sides. The smart self-insured employers are already thinking about this, offering GLP-1 coverage as a benefit while simultaneously deploying AI-driven plan design optimization to manage total cost of care. This is a real whitespace for entrepreneurs building tools that help employers navigate both trends simultaneously.

Pharma's position is obvious on peptides (they make the drugs) but their AI strategy is underdeveloped. Most large pharma companies are using AI for drug discovery and clinical trial optimization, which is valuable but incremental. The bigger opportunity is using AI for real-world evidence generation and outcomes-based commercialization, essentially using AI to prove that their drugs work in real-world populations and then using that evidence to negotiate better coverage and pricing. Few pharma companies are doing this well yet.

And then there are the wellness bros. The direct-to-consumer peptide market, including telehealth platforms prescribing compounded semaglutide, men's health clinics offering peptide stacks, and influencers promoting BPC-157 and other research peptides, represents a genuinely significant market that most institutional investors dismiss as noise. It is not noise. The DTC peptide market was probably \$1 billion in 2024 if you include compounding pharmacies, and it is growing faster than the branded market in percentage terms. The FDA crackdown on compounded semaglutide will reshape but not eliminate this channel. The wellness consumer who wants peptides is not going away. They will just shift to whatever the next access formulation is. Entrepreneurs building compliant, clinically grounded DTC peptide platforms are positioning for a market that institutional capital has largely ignored.

The Entrepreneur Opportunity Map

For healthcare entrepreneurs looking at this landscape today, there are several zones of opportunity that merit serious attention.

First, the data layer between peptides and AI. Whoever builds the longitudinal data connecting peptide utilization to long-term health outcomes across large populations will have an extraordinarily valuable asset. This does not exist today in a clean, accessible form. Claims data gives you part of the picture, EHR data gives you another part, patient-reported outcomes data gives you a third part, but nobody has stitched them together at scale for the peptide population specifically. The company that does this becomes the intelligence layer for payers making coverage decisions, pharmaceutical companies running real-world evidence studies, and employers evaluating ROI on GLP-1 benefits.

Second, AI-enabled peptide prescribing and management. The current prescribing model for GLP-1s is crude. Patient comes in, gets a prescription, titrates up on a standard schedule, maybe gets some dietary counseling. There is essentially no personalization based on genomic, metabolic, or behavioral data. The opportunity is to build an AI-driven precision prescribing platform that optimizes peptide selecti

dosing, and combination therapy based on individual patient characteristics is w open. This is a classic convergence play.

Third, the workforce transition infrastructure. If both trends reduce healthcare demand in specific categories, someone needs to build the retraining and redeployment platforms. A bariatric surgical nurse whose case volume is declining needs a pathway to another role. A medical coder whose job is being automated reskilling options. The companies that build workforce transition platforms specifically for healthcare will find a receptive market among health systems, un and government agencies.

Fourth, the employer benefits intelligence layer. Self-insured employers are mak million-dollar decisions about GLP-1 coverage, AI-driven plan optimization, and workforce health strategy with remarkably little analytical support. Most rely on benefits consultant, who is often conflicted. The opportunity for an independent powered employer health strategy platform is significant and largely untapped.

Convergence Is the Real Trade

The thesis that matters is not peptides vs AI. It is peptides times AI. The most valuable companies in healthcare over the next decade will be the ones that sit at intersection, using AI to make peptide therapies more effective, more personalized and more efficiently delivered, while using peptide-driven health improvements generate the outcomes data that makes healthcare AI actually useful. Neither reaches its full potential without the other. Peptides without AI remain blunt instruments prescribed on crude protocols. AI without therapeutic intervention actually change patient trajectories remains an optimization tool for a broken system. The entrepreneurs who see this convergence clearly, and build for it specifically, playing a different game than the ones picking sides. And in the long run, playing a different game is the only reliable way to generate outsized returns.



4 Likes • 1 Restack

[← Previous](#)

Discussion about this post

Comments

Restacks



Write a comment...