Building for the post transparency economy: Trump’s great healthcare plan

JAN 17, 2026 • PAID

2

Share

Abstract

Trump’s Great Healthcare Plan doesn’t just mandate disclosure, it creates the conditions for entirely new market structures that can’t exist under current opa This analysis examines five forward-looking business opportunities that emerge from compliance requirements but from the new economic behaviors transparer enables: real-time pharmaceutical arbitrage networks, dynamic facility routing optimization, predictive MLR trading markets, consumer health financing rebui around known prices, and labor-verified health access models. Each represents a business that’s structurally impossible today but becomes viable when informati asymmetries collapse.

Key opportunities:

- Pharmaceutical spot markets enabling real-time drug price arbitrage across ch

- AI-driven patient routing engines optimizing site-of-care decisions against live pricing

- MLR futures markets letting insurers hedge medical cost volatility

- Consumer credit products underwritten against transparent procedure pricing

- Blockchain-verified work credentials creating portable Medicaid eligibility

The highest-return plays involve building market infrastructure for behaviors th transparency makes rational but currently impossible.

Table of Contents

When Opacity Dies, Markets Reorganize

Pharmaceutical Spot Markets: The Drug Price Bloomberg

Intelligent Care Routing: Uber for Site Selection

MLR Derivatives: Hedging the Medical Cost Lottery

Known-Price Lending: Credit Products for Transparent Healthcare

Portable Work Verification: Identity Layer for Benefits

Second-Order Effects and Timing

When Opacity Dies, Markets Reorganiz

Healthcare's current infrastructure assumes prices stay hidden. Payment system route transactions without price discovery. Referral patterns ignore cost differer Insurance products can't differentiate on actual medical cost management becau nobody knows what things cost until after they happen. Supply chains operate o negotiated contracts rather than spot pricing. The entire stack gets rebuilt when transparency breaks these assumptions.

The Great Healthcare Plan's mandates matter less for the compliance burden th making new business models economically viable. Once PBM spreads become pu pharmaceutical arbitrage becomes possible. Once hospital prices get published i machine-readable formats, algorithmic site selection becomes feasible. Once MI calculations face real scrutiny, hedging instruments become valuable. The regul

don’t create these opportunities directly, they remove the barriers that made the impossible.

This requires thinking past the obvious infrastructure plays. Everyone sees the n for compliance dashboards and reporting tools. Fewer people see the entirely ne markets that become possible when information flows freely. The analogy runs t financial market structure. When stock prices moved from phone calls to electro tickers, the obvious play was building ticker infrastructure. The bigger play was building for decimalization, algorithmic trading, and market making behaviors t ticker transparency enabled but didn’t directly mandate.

Healthcare’s moving from phone-call pricing to electronic tickers. The question becomes what market structures emerge when every participant can see everyon else’s prices in real-time and route accordingly. That’s a different infrastructure than healthcare currently runs.

Pharmaceutical Spot Markets: The Dru Price Bloomberg

PBM spread transparency doesn’t just force disclosure, it creates conditions for pharmaceutical spot markets that trade drugs like commodities. Right now, drug pricing happens through annual contracts between manufacturers, PBMs, and pharmacies. Spreads stay hidden so price discovery never happens. Arbitrage opportunities exist everywhere but nobody can execute because transaction costs exceed potential gains without infrastructure.

The future company builds a pharmaceutical trading platform where drugs get p dynamically based on real supply and demand across channels. Self-insured emp could buy commonly prescribed drugs directly from manufacturers through spo auctions, cutting out PBM spreads entirely. Pharmacies could arbitrage between wholesalers based on real-time price feeds. Manufacturers could dump excess inventory into secondary markets when production runs over forecasted demand None of this works without transparent pricing and neutral exchange infrastruc

The technical architecture resembles commodity exchanges more than healthcar platforms. You need order matching engines that can handle thousands of SKUs different dosage forms and package sizes. You need clearing and settlement infrastructure that interfaces with existing pharmaceutical distribution network You need market data feeds that normalize pricing across different contract structures and rebate arrangements. You need fraud prevention since drug diver carries serious regulatory risk.

The business model compounds in interesting ways. Start by building the transp pricing layer, basically scraping and normalizing all the newly disclosed PBM sp data. That creates the information substrate for market making. Then introduce exchange functionality, letting sophisticated buyers and sellers transact directly. Eventually the platform becomes the price-setting mechanism for generic drugs, similar to how commodity exchanges set wheat or copper prices.

The regulatory angle gets tricky since FDA oversees drug distribution chains an wholesaler licensing requirements limit who can transact. But the plan's deregul bent suggests openness to secondary markets if they demonstrably lower costs. T early partnerships matter, you need a manufacturer willing to experiment with s sales, a large self-insured employer wanting to test direct purchasing, and a phar chain that sees arbitrage value. String together proof points before scaling.

Market sizing starts with generic drugs since brand pricing remains constrained patent monopolies and manufacturer price controls. Generic spend runs around billion annually with 20-40% in spreads and inefficient distribution markups. Capturing even 5% of that volume through more efficient spot markets creates a billion dollar gross merchandise value platform charging basis points on transac

Intelligent Care Routing: Uber for Site Selection

Site-neutral payment policies eliminate the facility fee advantages hospitals enjo procedures that could happen in ambulatory settings. This doesn't just shift volu it makes optimal site selection a solvable algorithmic problem for the first time.

Currently, referrals happen through physician relationships, insurance networks geographic convenience. Nobody optimizes for price because prices stay hidden EOBs arrive months later. That changes completely when hospitals must publish procedure costs in machine-readable formats and site-neutral payments equalize reimbursement.

The company worth building creates an intelligent routing layer that recommen optimal care sites based on price, quality, availability, and patient preferences. P needs knee imaging, the platform instantly evaluates every facility within reason distance, pulls their published prices for that specific procedure code, factors qu metrics and patient reviews, checks real-time appointment availability, and pres ranked options with total out-of-pocket cost estimates. The patient books direct through the platform which handles insurance verification and payment process

This only works with comprehensive price data and real-time facility integration price transparency mandates create the data substrate, requiring machine-reada files with negotiated rates for every payer-facility combination. The platform ing these feeds continuously, normalizes pricing across different file formats, and maintains a live database of what every procedure costs at every facility for every insurance plan. That's the foundation layer.

The routing optimization gets more sophisticated over time. Early versions just s by price and show quality ratings. Later versions incorporate patient preference learning, predicting which factors matter most for individual users based on pas selections. Eventually the platform runs predictive models on facility capacity, recommending appointments at times when providers have openings and might discounts to fill schedules. It becomes Kayak meets OpenTable meets ZocDoc bu actually useful because pricing data makes the recommendations economically meaningful.

The network effects compound as volume grows. More patient bookings give fac incentive to offer preferential pricing through the platform. Better facility pricin attracts more patients. Higher volumes justify building deeper integrations with scheduling systems and EMRs. Eventually the platform becomes the dominant p

acquisition channel for ambulatory facilities, capturing transaction fees on both of the marketplace.

The self-insured employer angle creates the initial wedge. These plans already tr steering patients to lower-cost facilities through reference-based pricing and cer of excellence programs, but execution remains clunky without good tooling. The platform gives benefits teams a turnkey solution, reducing their claims costs by 30% on steerable procedures through better site selection. That value propositio funds initial development and proves the model before expanding to individual consumers.

Revenue models stack multiple layers. Transaction fees on bookings, maybe 3-5% procedure cost. Subscription fees from employers for white-labeled routing tool Lead generation fees from facilities wanting patient volume. Data licensing to pa and providers who want market intelligence on competitive pricing. The unit economics work because the platform's eliminating real waste, not just redistrib existing dollars.

MLR Derivatives: Hedging the Medical Cost Lottery

Stricter medical loss ratio enforcement on Medicare Advantage plans creates ma financial volatility that nobody can currently hedge. MLR rules require insurers spend at least 85% of premium revenue on medical care or rebate the difference members. Plans hit this threshold through a combination of enrollment manage utilization controls, and accounting optimization. When CMS tightens enforcem and closes loopholes, plans face much higher variance in their ability to meet requirements year-to-year.

This creates conditions for an MLR derivatives market where insurers can hedge medical cost risk and investors can take exposure to healthcare utilization trend Think about it structurally: plans know their premium revenue with reasonable certainty based on CMS benchmarks and enrollment. They have much less certa about medical costs due to flu severity, new expensive drugs launching, utilizatio

pattern shifts, and provider contract negotiations. That's a hedgeable risk if you the right instruments and exchange infrastructure.

The basic product resembles weather derivatives or catastrophe bonds. An MA p projects 84% MLR for the year based on historical utilization. They buy MLR protection at 86%, paying a premium to a counterparty who agrees to cover medi costs exceeding that threshold. If actual MLR hits 88%, the counterparty pays th difference. If MLR comes in at 82%, the plan keeps the savings and the counterp keeps the premium. Both sides benefit from reducing uncertainty.

Building this market requires several components that don't exist. You need standardized MLR calculation methodologies that both sides trust, probably tie regulatory reporting definitions but with faster settlement cycles. You need neu claims data feeds that provide real-time MLR tracking without revealing compet information. You need credit infrastructure ensuring counterparties can cover potential payouts. You need regulatory approval since these instruments probab qualify as insurance or derivatives requiring oversight.

The customer base starts with smaller MA plans that can't absorb MLR volatilit easily as United or Humana. A regional plan with 50,000 Medicare members mig face existential risk if MLR swings 5 points in the wrong direction. They'll pay meaningful premiums to cap that exposure. As the market matures, even large p participate to smooth earnings and manage capital requirements.

The counterparty side attracts sophisticated investors wanting healthcare exposu uncorrelated with equity markets. Pension funds, endowments, and hedge funds already trade catastrophe bonds and weather derivatives. MLR instruments offer similar risk-return profiles with potentially better data and modeling capabilitie since healthcare utilization follows more predictable patterns than natural disas

Market sizing follows from MA premium volume around 500 billion annually. If 10% of plans hedge 20% of their MLR risk, that's 10 billion in notional value cre hundreds of millions in transaction fees and market making spreads for the exch operator.

The regulatory pathway requires convincing CMS that MLR hedging instrument improve market stability rather than enabling gaming. The pitch emphasizes tha hedging encourages smaller plan participation in MA by reducing capital requirements, promoting competition that the plan's broader goals support. Stat insurance regulators need comfort that the instruments don't constitute unlicen insurance. The derivatives framework potentially fits better, treating MLR swap other financial hedges rather than insurance products.

Known-Price Lending: Credit Products Transparent Healthcare

Price transparency mandates combined with site-neutral payments create somet that's never existed: known healthcare prices before care delivery. This enables entirely new consumer financing products underwritten against actual procedur costs rather than historical charge estimates. Currently, medical loans and paym plans get priced based on rough estimates of what procedures might cost, creati high default rates and expensive credit. That inefficiency disappears when exact become available upfront.

The new lending product works like this: patient needs hip replacement, the pla pulls exact negotiated rates for that procedure at available facilities given their insurance coverage. Say it's 15,000 out-of-pocket after insurance. The patient ap for financing through the platform which underwrites based on known procedu cost, their insurance coverage details, and standard credit factors. Because the lo amount matches real liability rather than an estimate, default risk drops substar That lets the lender offer better rates than traditional medical credit cards or pa plans.

The underwriting sophistication increases over time. Early versions just use pub prices and insurance verification. Later versions incorporate predictive models c complications and additional care needs based on patient health data and facility quality metrics. Eventually the platform's building custom risk scores for medic procedures similar to how auto lenders evaluate vehicle values and borrower pro

The facility partnerships create interesting dynamics. Providers have strong incentives to work with the lending platform since it removes payment friction f patients. Currently, patients facing large out-of-pocket costs defer care or haggl payment plans, creating collections headaches for facilities. The platform conve that into guaranteed payment at time of service, eliminating accounts receivable for facilities. That value justifies facility fees or revenue sharing arrangements.

The insurance integration matters for the product's viability. The platform need time benefit verification showing exact patient cost-sharing for specific procedu specific facilities. That requires deep integration with payer systems and eligibil databases. It also requires handling edge cases where patients hit deductibles mi year or have multiple insurance sources with coordination of benefits complexity

Revenue models layer multiple streams. Interest on loans, obviously. Originatior paid by patients or facilities. Risk-adjusted pricing that captures excess spread c lower-risk borrowers. Late fees and servicing revenue on the loan portfolio. Data monetization selling anonymized pricing and utilization intelligence to market participants.

The capital requirements scale with loan volume but remain manageable throug securitization. The platform originates loans, pools them into tranches with diff risk profiles, and sells them to institutional investors who want exposure to cons medical credit. The known-price underwriting creates better performance metri than traditional medical debt, making the securities more attractive to buyers.

Market sizing follows from the 50-70 billion in annual out-of-pocket medical spending on procedures where financing makes sense. Capturing 10% of that vol through better underwritten loans creates a 5-7 billion loan book generating 20(million in annual interest income and fees.

Portable Work Verification: Identity La for Benefits

Medicaid work requirement implementations create a technical problem states c solve with existing infrastructure: continuous verification that beneficiaries mee employment thresholds without creating administrative nightmares. Current eligibility systems rely on periodic attestations and manual documentation revie That doesn't work when requirements mandate 20 hours weekly employment ac 52 weeks. States need real-time work verification that doesn't require constant paperwork from beneficiaries or employers.

The solution involves building a portable work verification credential that beneficiaries control but employers validate, essentially creating a decentralized identity layer for employment status. The technical architecture uses blockchair similar distributed ledger technology to create tamper-proof work records. Emp submit hour attestations to the platform which writes them to the beneficiary's credential. State Medicaid systems query the credential to verify ongoing eligibi without intermediating every transaction.

This requires solving several hard problems. You need employer participation w creating new compliance burdens. The platform must integrate with existing pa systems to automatically submit hour attestations rather than requiring manual entry. You need privacy protection so beneficiaries control who sees their work history and states only get eligibility verification without detailed employment information. You need fraud prevention since the stakes involve both benefits ac and program integrity.

The go-to-market starts with states implementing work requirements who face impossible administrative challenges using legacy systems. The pitch emphasize reduced overhead costs compared to manual verification processes and better beneficiary experience through automated eligibility. Early adopter states becom reference customers proving the model works at scale.

The employer side requires careful positioning. Large employers with significan hourly workforces already run complex payroll systems and don't want additiona reporting requirements. The platform needs to offer clear value beyond Medicai verification to justify integration effort. That might include portable work crede

that employees use for credit applications, rental verification, or other identity r
It might include workforce analytics that help employers optimize scheduling ar
retention.

The beneficiary experience determines adoption since they control credential sh
The platform must feel like it protects their privacy and gives them agency over
data rather than creating another surveillance mechanism. That probably means
consumer-friendly mobile apps with granular sharing controls and clear explana
of who sees what information. It means fraud protection so stolen credentials ca
used inappropriately. It means dispute resolution when employers submit incorr
hour data.

Revenue comes from multiple sides. States pay per-verification fees lower than t
current administrative costs for manual eligibility checks. Employers might pay
premium features or workforce analytics. Beneficiaries might pay for extended
credential functionality beyond Medicaid verification. The platform could mone
anonymized workforce data to researchers and policymakers studying employme
patterns.

The network effects compound as coverage expands. More participating employe
make credentials more valuable to beneficiaries. More beneficiaries using creder
give employers incentive to integrate. More states adopting the system create pr
on employers with multi-state workforces to participate. Eventually the credenti
becomes standard infrastructure for employment verification beyond just Medic

The long-term vision extends to other means-tested benefits requiring work
verification like SNAP or housing assistance. It potentially becomes the identity
for how Americans prove employment status across contexts, replacing paystubs
employer letters with cryptographically verified credentials.

Second-Order Effects and Timing

These opportunities share common characteristics that make them different fro
typical healthcare venture bets. They're structurally dependent on regulatory ch

eliminating current information asymmetries. They're building market infrastru
rather than delivering care. They capture value from making new behaviors
economically rational rather than changing clinical outcomes. They scale throug
network effects and data accumulation rather than care delivery capacity.

The timing dynamics favor early action more than typical healthcare plays. Most
digital health companies face long validation cycles proving clinical efficacy and
willingness to reimburse. These companies face validation cycles proving that
transparency creates the predicted behaviors and that their infrastructure captu
the resulting value. That's faster to test and cheaper to prove wrong, making the
innovation cycles more capital efficient.

The regulatory uncertainty cuts both ways. The plan might face implementation
delays or legal challenges that slow timelines. Specific provisions might get wate
down through rulemaking or congressional revision. But the directional bet rem
sound: transparency increases and information asymmetries decrease. The exact
mechanisms matter less than positioning to capture value from more efficient
markets.

The exit landscape looks different too. These aren't clinical companies that get
acquired by health systems or pharma companies. They're infrastructure plays tl
might get acquired by existing healthcare technology vendors, financial services
companies entering healthcare, or data platforms wanting healthcare exposure. T
might IPO as vertical SaaS or fintech companies rather than healthcare compani
accessing better multiples through different comps.

The capital strategy probably involves less total funding than typical healthcare
ventures but faster deployment cycles. You can build MVPs and test core hypotl
within 12-18 months rather than requiring five-year clinical validation studies. T
means seed and Series A matter more than later-stage growth capital. It means f
on product-market fit and unit economics rather than long-term strategic positi

The talent requirements emphasize different skills than typical healthcare startu
You need people who understand market microstructure and exchange design, n

clinical workflows. You need regulatory strategists who can navigate CMS rulem and state Medicaid policy, not hospital sales experts. You need data engineers w normalize messy healthcare datasets, not population health clinicians. The team more like a fintech or marketplace startup than a care delivery company.

Building for post-transparency healthcare means betting that markets reorganiz around information efficiency once opacity disappears. That's not a crazy bet. It' every other industry evolved when pricing became transparent and transaction c dropped. Healthcare's just getting started.

2 Likes

← Previous

Nex

Discussion about this post

Comments Restacks

Write a comment...

© 2026 Thoughts on Healthcare · Privacy • Terms • Collection notice
Substack is the home for great culture