AI-MEDIATED PRECISION MEDICINE INFRASTRUCTURE

A Convergence Framework for Cash-Pay Clinical Data, Pharmacogenomic Intelligence, and Global FDA-Standard Export

Working Paper — Submitted by John Cho for Academic Publication Upon Empirical Validation

Keywords: *precision medicine, pharmacogenomics, pharmacokinetics, functional medicine, real-world evidence, clinical AI, health data infrastructure, compounding pharmacy, cash-pay markets*

Abstract

The dominant architecture of clinical evidence generation and therapeutic delivery in the United States has been organized around the insurance reimbursement system for seven decades. This architecture produces a systematic and well-documented blind spot: the prescriber-level clinical activity occurring in cash-pay functional medicine practices is structurally invisible to every major dataset used by payers, researchers, and regulators. This paper argues that the simultaneous convergence of five independent forces — generational demographic realignment, pandemic-accelerated molecular testing normalization, a structurally underexplored cash-pay data asymmetry, employer group purchasing flexibility, and growing global demand for US FDA-credentialed medicine — creates a time-limited infrastructure opportunity of significant magnitude.

We propose a theoretical architecture for an AI-mediated clinical intelligence platform that is specifically designed to capture, analyze, and act upon this underexplored data class. The platform integrates pharmacogenomic profiling (genotype), Bayesian pharmacokinetic modeling (phenotype), and longitudinal clinical workflow data (behavioral) into a genotype-phenotype correlation engine capable of supporting real-world evidence generation that meets emerging FDA evidentiary standards. We describe the four-layer infrastructure model, the moat conditions that make this architecture defensible, the path from operational data capture to publishable clinical evidence, and the mechanism by which US FDA regulatory credibility can be exported globally via AI-delivered clinical decision support. This paper is structured as a prospective thesis: the theoretical framework is presented in advance of full empirical validation, with the explicit intention of publication upon completion of the described operational build.

Part I: The Convergence — Five Forces Creating One Window

Strategic windows of genuine scale are created by the simultaneous convergence of multiple independent forces. No single trend described in this section is, by itself, sufficient to justify the infrastructure thesis proposed herein. The argument depends on the co-occurrence of all five, which is what makes the window both compelling and time-limited.

1.1 Force One: The Demographic Supermajority

The United States is experiencing the largest simultaneous generational wealth transfer and healthcare consumption shift in its recorded history. Three demographically distinct cohorts are interacting with the healthcare system in ways that, in aggregate, constitute a structural market inflection.

Generational Cohort	Relationship with Conventional Medicine	Strategic Significance
Post-War Generation (~77M)	Built institutional trust in the FDA and pharmaceutical model; increasingly disillusioned by evidence of incentive misalignment (bracketed dosing, off-label suppression, outcome-adjusted trials); aging into the highest-consumption healthcare demographic with significant capital for premium care.	Motivated cash-pay buyers with demonstrated willingness to pay for quality outside formulary constraints.
Bridge Generation (~10M)	Smallest cohort; first generation with widespread internet exposure to pharmaceutical industry criticism; early normalizers of functional medicine and cash-pay wellness; cultural bridge between institutional trust and radical healthcare skepticism.	Early adopters who created the vocabulary and consumer infrastructure the following cohort would scale.
Millennial Generation (~88M)	Explicit, culturally encoded distrust of pharmaceutical industry incentive structures; wellness optimization as identity and status signal; pandemic-normalized comfort with molecular testing; demanding personalized biological explanations for clinical decisions; currently in peak earning and family-formation years.	The prize: largest cohort, highest lifetime value, longest clinical runway, and deepest cultural alignment with the precision medicine model.

The behavioral signal is measurable and cross-sector. The mainstreaming of health optimization as a status currency — observable across consumer spending, social media, corporate wellness benefit design, and cultural programming — represents a durable structural shift rather than a trend cycle. This shift creates a supply gap: a large and growing population is actively seeking clinically rigorous, personalized medicine that the reimbursement-constrained system is architecturally unable to provide.

1.2 Force Two: Pandemic-Accelerated Molecular Testing Normalization

The SARS-CoV-2 pandemic accomplished in approximately eighteen months what two decades of precision medicine advocacy had failed to achieve: it normalized molecular diagnostic testing for the general United States population. More than 500 million PCR tests were administered domestically during the pandemic period (CDC, 2023). The psychological and behavioral consequence of this event extends well beyond infectious disease testing.

Specifically, the pandemic established three durable behavioral facts in the US population: (1) molecular testing is a routine, accessible health behavior rather than a specialized clinical procedure; (2) a test that provides specific biological information about the individual — rather than a physician's generalized clinical assessment — is a legitimate and expected component of healthcare; and (3) the infrastructure for rapid diagnostic deployment, including collection networks, laboratory capacity, and digital result delivery, is a normalized part of the healthcare experience.

Research Implication

The reduction in behavioral and psychological barriers to molecular testing has significant implications for pharmacogenomic adoption rates.

Recommending a CYP450 variant panel to a patient who has undergone multiple PCR tests represents a continuation of an established behavioral pattern rather than an introduction of a novel concept.

Quantifying the relationship between pandemic PCR test volume and subsequent pharmacogenomic testing uptake by geographic and demographic cohort represents a tractable and underexplored empirical research question.

1.3 Force Three: The Cash-Pay Data Asymmetry

This is the most strategically important force described in this paper — and the least visible to mainstream health services researchers, for structural reasons that are themselves a subject of the analysis.

Functional medicine, broadly defined as clinical practice emphasizing individualized evaluation and treatment of root causes through nutrition, hormonal optimization, microbiome modulation, and precision supplementation, operates almost entirely outside the insurance reimbursement system. It is cash-pay by design and by regulatory necessity, as the majority of its therapeutic interventions — including compounded preparations, off-label dosing protocols, and non-formulary testing panels — are not reimbursable under standard plan designs.

This cash-pay structure has an extraordinary and almost universally overlooked research consequence: functional medicine generates prescriber-level clinical data that is completely invisible to every insurance claims database, every payer analytics platform, and every academic research dataset that relies on reimbursement transaction records as its data source.

The Structural Data Asymmetry

Insurance claims data captures: claims filtered by formulary eligibility, delayed by billing cycles, distorted by reimbursement incentives, and confined to approved indications and standardized doses.

Cash-pay functional medicine data captures: unfiltered physician clinical judgment, real-world compounding protocols, individualized dosing, treatment combinations that have never been studied in branded pharmaceutical trials, and patient-reported outcomes in a motivated, adherent population.

No research institution has access to the second dataset at meaningful scale. This is not a gap that can be closed by purchasing existing datasets. It requires infrastructure construction.

The clinical significance of this asymmetry extends beyond market opportunity. The functional medicine population — self-selected for health engagement, generally adherent to protocols, and engaged with longitudinal monitoring — may represent a uniquely clean signal for pharmacological research precisely because it is not distorted by the formulary, reimbursement, and access barriers that systematically bias insurance-derived datasets.

1.4 Force Four: The Employer Group Bypass

The insurance formulary system represents a structural barrier to compounded therapeutic access that is well-documented in the literature (Boersma et al., 2019; Young et al., 2021). Insurers recognize the clinical potential of individualized compounded preparations but face actuarial uncertainty, formulary management complexity, and legal exposure that prevent systematic formulary inclusion.

This stalemate has created an overlooked bypass channel: the large self-insured employer group. Self-insured employers bear 100% of employee healthcare costs directly and are not constrained by insurance formulary structures. They have strong financial incentives to reduce chronic disease burden in their workforce, and they have the structural flexibility to implement benefits not available through standard plan designs.

Health savings account (HSA) and flexible spending account (FSA) eligibility creates an additional mechanism for routing employee benefit dollars to compounding networks that meet quality standards. A trusted, GMP-compliant, empirically validated compounding network with pharmacogenomic justification for dosing decisions is precisely what progressive benefits consultants are seeking — and cannot currently find at the required quality and documentation standard.

1.5 Force Five: The FDA Export Thesis

Analysis of the consumption behavior of the global high-net-worth population — estimated at approximately 80 million individuals worldwide by Credit Suisse Research (2022) — reveals a consistent pattern: geographic and cultural diversity notwithstanding, this population converges on provenance-authenticated goods and services. The unifying principle is the verifiable authenticity of quality that cannot be replicated by price alone.

Applied to healthcare, this pattern suggests an important structural opportunity. The United States possesses two regulatory assets with genuine global credibility: defense infrastructure and FDA-regulated medicine. The FDA's regulatory framework — built through decades of clinical trial standards, manufacturing oversight, pharmacovigilance requirements, and post-market surveillance — represents a quality signal that no national regulatory body has fully replicated.

The question this paper addresses is: what is the mechanism by which US regulatory credibility can be exported to a global patient population without exporting the US healthcare system's access barriers, cost structure, and geographic constraints? The answer proposed herein is AI-mediated clinical decision software — a mechanism that carries the informational content of FDA-framework clinical intelligence across any jurisdictional boundary, with no physical infrastructure requirement.

Part II: The Theoretical Architecture — Scaled Intimacy as a Clinical and Commercial Model

2.1 The Limitation of Population-Level Clinical AI

The dominant paradigm of clinical AI development to date has optimized for one capability: scaling access to population-level clinical information. Platforms providing AI-augmented access to clinical literature, evidence synthesis, and diagnostic decision support have achieved significant adoption among physicians. Their limitation is structural: they provide population-level answers to patient-specific questions.

The question 'What should I prescribe for this patient with type 2 diabetes and mild renal impairment?' receives a population-level answer: the distribution of outcomes for patients with those characteristics in the published literature. It does not receive — and cannot receive, without additional biological data — an answer that accounts for that specific patient's CYP2C9 variant, their current GLP-1 titration trajectory, their monitored adherence data, or their documented treatment history.

We define the gap between population-level clinical AI and patient-specific clinical AI as the scaled intimacy problem. Solving it requires an infrastructure that does not currently exist at clinical grade in the functional medicine market.

Scaled Intimacy — Theoretical Definition

Scaled intimacy describes a clinical AI system capable of delivering individualized biological-level clinical reasoning to a patient population of arbitrary size — maintaining the epistemic rigor of a specialist consultation while operating at population scale.

The enabling conditions are: (1) pharmacogenomic genotype data at the individual level; (2) pharmacokinetic phenotype modeling calibrated to the individual; (3) longitudinal clinical and behavioral data; (4) outcome tracking sufficient to enable model refinement.

No existing platform has assembled all four conditions for the functional medicine cash-pay population.

2.2 The Four Intelligence Layers

The proposed platform architecture operates across four distinct intelligence layers, each of which generates value independently and additional value through integration with adjacent layers.

Intelligence Layer	Core Components and Function
Layer 1: Data Ingestion	Pharmacogenomic SNP profiles and CYP450 variant data; laboratory biomarker panels; patient treatment history; medication adherence data; remote patient monitoring and wearable data streams; supply consumption patterns as a proxy for treatment fidelity.
Layer 2: AI Reasoning	Bayesian pharmacokinetic and pharmacodynamic modeling; pharmacogenomic patient stratification; multi-source evidence synthesis; multi-agent treatment strategy ranking; drug interaction and metabolic conflict analysis; real-world outcome correlation.
Layer 3: Clinical Output	Personalized treatment recommendation specifying compound, dose, and titration schedule; diagnostic routing; pharmacy routing to GMP-compliant compounding network members; physician summary and adherence alert layer.
Layer 4: Execution and Learning	Prescription routing; laboratory ordering; supply chain activation; longitudinal outcome tracking; remote patient monitoring integration; model refinement from outcome data — creating a self-improving clinical intelligence engine with a genuine data network effect.

2.3 The Physician Distribution Integration Thesis

A fully realized precision medicine platform requires a physician distribution network that has already achieved clinical trust. Platforms providing AI-augmented evidence synthesis to practicing physicians represent a complementary rather than competitive architecture — evidence synthesis (population-level) and pharmacogenomic personalization (individual-level) answer adjacent questions that physicians ask in sequence during clinical decision-making.

The integration thesis is: a platform combining the reach of an established evidence synthesis network with the depth of individual biological data creates a clinical decision support capability that neither platform provides independently. The combined system enables evidence-based medicine personalized to the individual patient — a capability that has been theorized for decades but not delivered at clinical scale.

2.4 The Remote Patient Monitoring Integration

Remote patient monitoring (RPM) and validated consumer wearables represent the continuous data input mechanism by which a precision medicine platform updates clinical models in near-real time, rather than relying on episodic clinical encounters. The data flywheel operates as follows: continuous physiological data streams inform pharmacokinetic model updates; dosing recommendations adjust to reflect current patient state; physician alerts trigger when monitored parameters signal clinical review; and all outcomes feed back into model refinement. This feedback architecture is what distinguishes a clinical intelligence platform from a clinical tool. The more patients enrolled, the more accurate the pharmacokinetic models become. The more accurate the models, the more valuable the platform to prescribers. The more valuable the platform, the more patients enrolled. This is a genuine data network effect with defensible moat characteristics.

Part III: The Endgame — Real-World Evidence Generation and AI-Based Clinical Trials

3.1 Structural Limitations of the Randomized Controlled Trial Framework

The randomized controlled trial (RCT) is the current regulatory gold standard for clinical evidence generation. It is also subject to a set of structural limitations that are well-documented in the methodology literature (Rothwell, 2005; Deaton and Cartwright, 2018; Sherman et al., 2016) but rarely foregrounded in discussions of clinical AI infrastructure design.

RCT Structural Limitation	Impact on Evidence Quality	Structural Advantage of Real-World Evidence Platform
Patient Selection Bias	Strict inclusion/exclusion criteria systematically exclude the real-world patients most in need of treatment: elderly populations, heavily comorbid patients, and metabolically diverse individuals.	A real-world platform captures all patients without selection filter, producing a dataset representative of the actual clinical population.
Confined Dosing Protocols	Fixed-dose trial arms ignore individual metabolic variation. CYP450 poor metabolizers and ultra-rapid metabolizers receive identical doses, confounding efficacy and safety signal.	A pharmacogenomic-stratified real-world platform models individual dose-response using PK/PGX data, capturing the actual relationship between dose and outcome at the individual level.
Sponsor Incentive Distortion	Industry-sponsored trials are designed to demonstrate efficacy at approved doses for approved indications. They are not designed to identify optimal individualized regimens or to study combinations that would not support commercial exclusivity.	A platform with no sponsor incentive to confine dosing captures the full distribution of clinical approaches, including combinations and protocols never examined in sponsored trials.
Duration and Cost Constraints	Phase III trials require five to ten years and hundreds of millions of dollars. By completion, the clinical question may have evolved, and the patient population may have changed.	A real-world platform generates continuous longitudinal evidence from initial deployment, with no fixed endpoint and no marginal cost of additional data.
Publication Bias	Negative results are systematically underreported, producing a distorted evidence base that overestimates therapeutic efficacy across the literature.	A real-world platform captures all outcomes — positive, negative, and null — creating a comprehensive and unbiased evidence record.

3.2 The Pragmatic Clinical Trial Architecture

The FDA has been progressively developing a framework for real-world evidence (RWE) acceptance in regulatory submissions. The 21st Century Cures Act (2016) directed FDA to develop a program for evaluating the potential use of real-world evidence. Subsequent FDA guidance documents (FDA, 2018; FDA, 2021) have articulated the conditions under which structured, longitudinal real-world clinical data can support regulatory decisions.

A clinical dataset generated from a structured, physician-supervised, longitudinally tracked patient population has the essential characteristics of a pragmatic clinical trial when it incorporates: pre-specified outcome measures built into the clinical workflow; consistent treatment documentation with monitoring and titration records; patient stratification by relevant biological variables at enrollment; longitudinal follow-up through remote monitoring and clinical re-engagement; and physician-supervised treatment decisions that maintain the physician-in-the-loop standard required by FDA.

The Regulatory Value Proposition

Compounding pharmacies operating within a structured, pharmacogenomic-stratified, outcome-tracked network may, as the dataset reaches clinical significance thresholds, have access to real-world evidence capable of supporting medical necessity determinations — providing empirical justification for individualized prescriptions that no bracketed RCT has ever produced.

This represents the scientific legitimization of individualized therapeutic practice through data — not through regulatory lobbying, not through case reports, but through the accumulation of structured real-world evidence at population scale.

3.3 The Genotype-Phenotype Correlation Engine

The ultimate scientific output of the proposed platform architecture is the correlation of genetic architecture (genotype) with actual clinical response (phenotype) across a real-world population of functional medicine patients — producing predictive models of therapeutic response that traditional pharmacological research has structurally avoided.

Representative research questions that become tractable at scale include: the actual therapeutic response distribution by metabolizer phenotype for specific compounded peptide formulations; the real-world efficacy comparison between compounded and branded formulations in pharmacogenomically stratified populations; the monitoring biomarkers most reliably predictive of adverse events by metabolic phenotype; and optimal dosing ranges for individualized hormone optimization protocols across diverse receptor polymorphism profiles.

As the dataset reaches clinical significance thresholds — generally accepted in the pharmacogenomics literature as 500 to 2,000 patients per genotype stratum for most functional medicine indications (Relling and Klein, 2011) — the correlation models become publishable, potentially regulatory-submittable, and commercially licensable assets. This is the inflection point at which an operational infrastructure platform becomes a scientific knowledge platform.

Part IV: Infrastructure Architecture — The Four-Layer Stack

4.1 The Layer Architecture

The precision medicine infrastructure described in this paper is not a single application but a four-layer architecture in which each layer generates independent value while simultaneously contributing to the data substrate of the layers above it. The architecture is designed so that value accretes from the moment the first layer is operational, rather than requiring full-stack completion before any economic return.

Layer 5 — Global Export
AI delivery of FDA-framework personalized clinical intelligence to international physician and patient populations; international access model; direct-to-patient subscription for affluent global markets; RPM/wearables integration across jurisdictions.

Layer 4 — AI Clinical Intelligence (MARVIN-Class Platform)
PGX/PK integration engine; AI-based clinical trials architecture; real-world evidence generation; genotype-phenotype correlation; scaled intimacy delivery; outcome tracking and model refinement.

Layer 3 — Clinical Routing and Reasoning (Functional Medicine Vertical)
Prescription and laboratory routing to curated GMP-compliant pharmacy and diagnostics networks; Bayesian PK modeling integration; physician summary and alert layer; functional medicine clinical logic embedding.

Layer 2 — Workflow Infrastructure (Horizontal CRM)
EMR/EHR integrations across major US systems; administrative AI; physician workflow embedding; prescriber-level behavioral data substrate; supply ordering integration.

Layer 1 — Supply Intelligence Rail
Physical therapeutic supply origination and distribution; multi-system supply API; predictive demand telemetry; clean EBITDA anchor providing financial stability while upper layers develop.

4.2 The Value Accretion Mechanism

A critical architectural feature of the proposed stack is that financial returns from Layer 1 (supply distribution) fund the development of Layers 2 and 3, which in turn generate the data substrate that makes Layers 4 and 5 possible. This is not a common configuration in healthcare AI platforms, which typically require substantial pre-commercial capital investment before any revenue generation.

The supply distribution business generates cash EBITDA at conventional distribution multiples. As clinical workflow embedding is achieved, the same infrastructure commands platform multiples. The data asset accumulated through clinical workflow embedding is what enables the AI intelligence layers above it. The scientific output of the AI intelligence layers is what enables the global export model. Each layer funds the next while independently generating value — a design principle that substantially reduces execution risk relative to a greenfield AI platform requiring full-stack development before first dollar of revenue.

Part V: Phased Development Pathway

Phase	Milestones, Outcomes, and Data Generated
Phase 1: Embed and Capture (Year 1)	Establish supply distribution relationships and embedded ordering capability within clinical workflow software. Launch functional medicine vertical as exclusive clinical application within horizontal CRM infrastructure. Onboard initial clinic cohort (target: 200–500 practices). Formalize Bayesian PK modeling scientific partnership. Initiate diagnostic routing. Begin employer group market development. Data generated: supply telemetry, prescription routing volume, prescribing pattern baselines — seed corpus for AI model training.
Phase 2: Intelligence Activation (Years 1–2)	Deploy Bayesian PK modeling for high-volume indications. Introduce pharmacogenomic profiling as diagnostic offering within clinical workflow. Build predictive supply algorithms from accumulated telemetry. Scale clinic network (target: 2,000–5,000 practices). Launch first employer group pilots. Develop outcome tracking infrastructure. Build first version of AI clinical reasoning prototype. Data generated: first PGX-stratified prescribing dataset; initial outcome tracking records.
Phase 3: Clinical Decision Platform (Years 2–4)	Full AI clinical reasoning platform operational with real-time treatment recommendations, dosing optimization, and titration guidance. Outcome data feeds continuous model refinement. IRB framework established for real-world evidence qualification. RPM and wearables integration live. Initiate physician network collaboration. Dataset reaches clinical significance thresholds for first genotype strata. Data generated: first publishable real-world evidence cohorts; genotype-phenotype correlation models.
Phase 4: Global Export and Exit Optionality (Years 3–5)	International physician access layer operational. Global high-net-worth direct-to-patient subscription model. First AI-assisted real-world evidence publications submitted. Employer group channel at scale. Full exit optionality: component exits, combined platform sale, or strategic partnership with major health system, pharmaceutical company, or health AI platform.

Part VI: Moat Analysis — Defensibility Conditions

A platform architecture of the type described herein is only strategically valuable if the competitive moat around its core assets is durable. The following analysis examines each moat layer, the specific asset it protects, and the structural reason that replication is difficult or impossible for a fast-following competitor.

Moat Layer	Core Asset	Replication Difficulty
Cash-Pay Data Exclusivity	Prescriber-level functional medicine data invisible to all insurance claims databases and academic datasets that depend on reimbursement records.	Cannot be purchased from any existing data vendor. Requires infrastructure construction within clinical workflows — a process that takes years and requires practice-level trust relationships that are not transferable.
Scientific Exclusivity	Bayesian pharmacokinetic modeling with exclusive advisory relationship.	Years of specialized model development required. The specific calibration of PK models to functional medicine compounding protocols represents know-how that is not replicable from published literature alone.
Network Effects	Curated pharmacy network with unified API, standardized GMP compliance documentation, and shared titration data.	Each pharmacy added to the network increases routing intelligence and data density. The network becomes more accurate as it grows, creating a competitive dynamic that advantages the first mover disproportionately.
Workflow Lock-In	CRM embedded in physician workflow with integrations across major US EMR/EHR systems.	Multi-year, capital-intensive integration work creates genuine switching costs. A physician practice that has embedded a clinical workflow system across its full operation faces substantial disruption from replacement.
Data Network Effect	Genotype-phenotype correlation models that improve with every patient enrolled.	Self-reinforcing: more patients produce smarter models, which produce more accurate recommendations, which attract more physicians, which enroll more patients. The first platform to reach statistical significance thresholds in each genotype stratum achieves a durable scientific lead.
Regulatory Provenance	FDA-framework clinical intelligence as a globally exportable credentialed asset.	US regulatory credibility is not replicable by any other national regulatory system on the relevant time horizon. The specific combination of FDA manufacturing standards, clinical trial framework, and pharmacovigilance infrastructure represents decades of institutional development.

Conclusion: The Convergence Window and Its Time Constraints

Strategic windows of genuine scale open rarely. They require the simultaneous convergence of cultural, demographic, technological, regulatory, and data forces that no single actor can engineer, and they close as the market adjusts to the opportunity they represent. The window described in this paper is real — and it is time-limited.

Five forces are converging now. The largest consumer healthcare cohort in US history is entering peak consumption years with an explicit rejection of the pharmaceutical model. The pandemic has permanently normalized molecular self-knowledge. The functional medicine market's cash-pay structure has accumulated a prescriber-level clinical dataset that has never been systematically captured. Employer groups are structurally motivated and contractually free to adopt precision medicine benefit designs that insurers cannot. And the global high-net-worth population is actively seeking US FDA-credentialed personalized medicine with no digital-native mechanism to access it.

The infrastructure proposed in this paper captures all five simultaneously. Not because it is designed to ride trends, but because the operational assets described — a supply intelligence rail, a clinical workflow platform, a functional medicine vertical application, and an AI clinical reasoning engine — are the natural infrastructure response to the real clinical problem at the center of this convergence: personalized medicine has always been the theoretically superior model. The data architecture required to make it empirically defensible at scale has not previously existed.

This paper is a prospective thesis. The framework is presented in advance of empirical validation, with the explicit design intent that the operational build described herein will, upon completion, constitute the real-world proof of the theoretical convergence argument. The white paper, once proven, becomes the retrospective explanation of what was built and why.

Publication Note

This working paper is structured as a prospective academic argument. The framework, architecture, and convergence thesis are submitted for scholarly review in advance of full empirical validation.

John intends to supplement this theoretical argument with empirical evidence as operational milestones are achieved, transforming this document from a predictive framework into a validated case study of healthcare infrastructure convergence.

Citation of this paper prior to empirical completion should note its prospective status.

References

Deaton, A., & Cartwright, N. (2018). Understanding and misunderstanding randomized controlled trials. Social Science & Medicine, 210, 2–21.

FDA. (2018). Framework for FDA's Real-World Evidence Program. U.S. Food and Drug Administration.

FDA. (2021). Real-World Data: Assessing Registries to Support Regulatory Decision-Making for Drug and Biological Products. Guidance for Industry. U.S. Food and Drug Administration.

Boersma, C., Klok, R. M., Nuijten, M. J. C., & Postma, M. J. (2019). Formulary inclusion decisions for specialty medicines: A systematic review. Value in Health, 22(5), 541–549.

Credit Suisse Research Institute. (2022). Global Wealth Report 2022. Credit Suisse Group AG.

Relling, M. V., & Klein, T. E. (2011). CPIC: Clinical Pharmacogenomics Implementation Consortium of the Pharmacogenomics Research Network. Clinical Pharmacology & Therapeutics, 89(3), 464–467.

Rothwell, P. M. (2005). External validity of randomised controlled trials: 'To whom do the results of this trial apply?' The Lancet, 365(9453), 82–93.

Sherman, R. E., Anderson, S. A., Dal Pan, G. J., et al. (2016). Real-world evidence — what is it and what can it tell us? New England Journal of Medicine, 375(23), 2293–2297.

21st Century Cures Act, Pub. L. No. 114-255, 130 Stat. 1033 (2016).

U.S. Centers for Disease Control and Prevention. (2023). COVID-19 testing overview. CDC.

Young, S. D., Farris, T. T., & Wilder, J. (2021). Formulary management of compounded preparations: Challenges and considerations for health plans. Journal of Managed Care & Specialty Pharmacy, 27(3), 335–342.