

GNQ INSILICO

WHITE PAPER

Reimagining Pharmacovigilance: How Causal Biological AI and Digital Twins Are Transforming Drug Safety in Oncology

Confidential | For Authorized Use Only | **June 2026**
Sudhir Saxena | C.T.O | gnq.ai

Executive Summary

Pharmacovigilance — the science of detecting, assessing, and preventing drug adverse effects — is at an inflection point. Despite decades of investment in post-market surveillance systems, the industry continues to face a fundamental limitation: current frameworks are reactive, population-level, and blind to the biological individuality that drives the majority of serious adverse drug reactions (ADRs).

In oncology, where treatment regimens are aggressive and patient populations are biologically heterogeneous, this gap is most consequential. Serious ADRs remain a leading cause of treatment discontinuation, hospitalization, and mortality — even for approved, well-studied agents.

GNQ InSilico's Quantum Biological Reasoning Model (Q-BRM) represents a paradigm shift. By combining causal AI, mechanistic biological pathway reasoning, and patient-specific digital twins, Q-BRM enables a new class of pharmacovigilance: predictive, personalized, and mechanistically grounded. Rather than waiting for ADR signals to emerge in population data, Q-BRM reasons from an individual patient's biology — before, during, and after treatment — to anticipate harm, explain its mechanism, and guide safer therapeutic decisions.

Key Finding

Q-BRM identified a clinically significant off-target tissue effect in dexamethasone therapy that population-level pharmacovigilance systems are structurally incapable of detecting — illustrating the platform's potential to transform post-market surveillance in oncology.

The Limits of Traditional Pharmacovigilance

Traditional pharmacovigilance relies on spontaneous adverse event reporting, electronic health record (EHR) mining, and population-level statistical signal detection. These methods share a common weakness: they are inherently retrospective. A signal only emerges once enough patients have already been harmed.

In oncology, this limitation is compounded by several factors:

- **Biological heterogeneity:** Cancer patients carry unique genomic, transcriptomic, and proteomic profiles that determine how drugs behave at the individual level. Population averages obscure — rather than illuminate — the mechanisms that drive harm in specific subpopulations.
- **Polypharmacy complexity:** Oncology regimens typically combine multiple agents, each interacting with the patient's biology in ways that are difficult to disentangle in post-market data.
- **Mechanism opacity:** Standard surveillance identifies that an ADR occurred, but not why — leaving clinicians without actionable guidance for the individual patient in front of them.
- **Latency of harm:** Off-target tissue effects, cumulative toxicities, and variant-driven metabolic failures often manifest months or years after exposure, making causal attribution nearly impossible with conventional tools.

The result is a system that protects populations in aggregate but consistently fails individual patients — particularly those whose biology places them outside the statistical mainstream.

Dimension	Traditional Pharmacovigilance	Q-BRM Approach
Detection timing	Post-hoc, signal-based	Prospective, pre-exposure
Biological scope	Population-level statistics	Individual patient biology
Mechanistic depth	Observational correlation	Causal pathway reasoning
Off-target detection	Limited — requires signal volume	Structural — variant-pathway modeling
Post-market utility	Passive surveillance	Active digital twin monitoring
Actionability	Label updates (months–years)	Real-time clinical guidance

The Q-BRM Platform: Architecture and Capabilities

Q-BRM is a three-layer causal reasoning engine operating at the intersection of genomics, biological pathway science, and quantum-inspired optimization. It is designed to simulate how a specific patient's biology will respond to a specific therapeutic intervention — at a mechanistic, tissue-specific level.

Layer 1: Relational Dependency Graph (RDG) + Bio Avatar

The RDG encodes the causal architecture of human biology — genes, proteins, enzymes, metabolic pathways, and their dependencies — as a structured knowledge graph. For each patient (Bio Avatar), the RDG is instantiated with their specific variant profile, expression data, and clinical history, creating a personalized biological map.

Layer 2: Quantum Reasoning Causal Model (Quantum RCM)

Operating at levels 2 and 3 of causal inference — intervention and counterfactual reasoning — the Quantum RCM simulates the downstream effects of therapeutic interventions across the patient's RDG. It can model not only what will likely happen, but what would have happened under alternative treatment choices, enabling genuine counterfactual safety analysis.

Layer 3: BRM Agentic Orchestration

The orchestration layer coordinates multi-pathway reasoning across clinical contexts, integrating real-world evidence, GNQ RDG, and patient longitudinal data. It surfaces insights through BioLens — GNQ's clinical workbench interface — in formats designed for clinician decision support.

Case Study : Dexamethasone — Tissue-Specific Therapeutic Inversion via HSD11B1

Clinical Context

Dexamethasone is a foundational component of multiple myeloma treatment regimens, valued for its immunosuppressive and anti-inflammatory effects within the bone marrow tumor microenvironment. Despite its widespread use, a clinically significant subset of MM patients exhibits both reduced therapeutic response and disproportionate systemic toxicity — including chronic kidney disease (CKD), metabolic suppression, and infection vulnerability.

The HSD11B1 Variant Mechanism

HSD11B1 (11 β -hydroxysteroid dehydrogenase type 1) is an enzyme responsible for intracellular reactivation of glucocorticoids from their inactive ketone forms — converting 11-oxo-dexamethasone back to active dexamethasone within specific tissues. This reactivation step is critical for dexamethasone to exert its intended immunosuppressive effect at the site of disease.

In patients carrying a loss-of-function HSD11B1 variant, Q-BRM's tissue-level pathway modeling reveals a clinically consequential asymmetry:

- In the bone marrow tumor microenvironment — where MM cells reside — HSD11B1-mediated reactivation of dexamethasone is impaired. Active drug concentrations at the tumor site are insufficient to suppress MM cell growth and the pro-tumor immune environment. The result is treatment resistance.
- Systemically, circulating active dexamethasone — which does not require HSD11B1 reactivation — continues to reach peripheral organs including the liver, kidney, and adipose tissue. These organs experience ongoing glucocorticoid-mediated immunosuppression and metabolic stress despite the drug's failure to work at the intended site.

Key Insight

This is not a conventional off-target receptor binding story. It is a tissue-specific therapeutic inversion: the HSD11B1 variant redirects dexamethasone's effective bioavailability away from the tumor microenvironment toward peripheral organs, flipping the drug's risk-benefit ratio at the individual patient level.

Post-Market Surveillance Implications

From a pharmacovigilance perspective, the CKD and metabolic complications observed in this patient population would appear in adverse event databases as standard dexamethasone toxicities — indistinguishable from effects in the broader MM population. There is no population-level signal that would reveal the underlying HSD11B1 mechanism, because the variant's frequency and the confounding of polypharmacy would obscure the association for years.

Q-BRM detects this pattern prospectively, by simulating dexamethasone's tissue-level distribution and reactivation dynamics against the patient's specific HSD11B1 genotype. It surfaces not only the risk of harm, but the mechanistic reason — enabling the clinician to adjust the regimen before toxicity accumulates.

Digital Twins and Continuous Post-Market Surveillance

The two case studies above illustrate Q-BRM's capacity for pre-treatment risk stratification. The platform's pharmacovigilance potential extends further: into continuous, longitudinal monitoring through patient-specific digital twins.

GNQ's Bio Avatar Platform maintains a living computational model of each patient — updated as new clinical, genomic, and real-world data becomes available. This enables a post-market surveillance paradigm that is fundamentally different from anything currently in regulatory practice:

- **Prospective harm anticipation:** Rather than flagging an ADR after it has occurred, the digital twin can signal emerging biological stress patterns — rising pathway perturbation scores, enzymatic bottleneck accumulation — before clinical symptoms manifest.
- **Counterfactual safety analysis:** For any observed adverse event, the digital twin can model what would have happened under alternative treatment choices, providing mechanistic attribution that traditional surveillance cannot.
- **Subpopulation discovery:** Across a population of digital twins, Q-BRM can identify previously unknown biological subgroups who share elevated risk for specific ADRs — enabling targeted label updates, risk minimization programs, and companion diagnostic development.
- **Real-world evidence integration:** As treatment outcomes accumulate, the digital twin population continuously refines Q-BRM's predictive models, creating a self-improving pharmacovigilance loop that strengthens over time.

This architecture aligns directly with the vision articulated by leading voices in healthcare AI: that transformative insights will emerge from models that individualize biological reasoning and simulate it at the patient level. GNQ's platform is built to deliver exactly that.

Regulatory Positioning and FDA Alignment

GNQ's approach is designed to complement and enhance existing regulatory frameworks — not to circumvent them. Q-BRM is in alignment with FDA's Innovative Science and Technology Approaches for New Drugs (ISTAND) program, which specifically invites novel analytical methods for drug development and post-market applications.

The platform's causal reasoning architecture is well-suited to several evolving regulatory priorities:

- **Real-World Evidence (RWE):** FDA's increasing reliance on RWE for post-approval safety decisions creates a natural fit for Q-BRM's ability to contextualize real-world observations within mechanistic biological models.
- **Pharmacogenomics:** The agency's expanding framework for genomic biomarker labeling aligns directly with Q-BRM's variant-level pathway scoring — providing the mechanistic evidence base that current genomic label updates often lack.
- **Predictive Safety Testing:** Q-BRM's digital twin simulation capability maps to FDA's interest in in silico trial methodologies as complements to traditional clinical pharmacology studies.

The Strategic Opportunity for Pharma

For pharmaceutical organizations managing large post-market drug portfolios, Q-BRM offers a fundamentally new capability layer — one that transforms pharmacovigilance from a compliance cost center into a strategic asset for label defense, product differentiation, and precision medicine positioning.

Label Defense and Risk Management

As regulatory scrutiny of post-market safety intensifies, the ability to demonstrate mechanistic understanding of ADR subpopulations — rather than statistical correlation alone — strengthens label defense and supports risk evaluation and mitigation strategy (REMS) programs with biological specificity that traditional surveillance cannot provide.

Precision Oncology Differentiation

In an increasingly competitive oncology market, the ability to identify which patients will benefit from a therapy and which face elevated harm risk — before treatment — creates meaningful differentiation. Q-BRM provides the mechanistic substrate for companion diagnostic development, patient selection algorithms, and precision dosing protocols.

Post-Market Lifecycle Value

Approved oncology agents with manageable safety profiles in the right patients can achieve significant market expansion if the at-risk subpopulations are identified and protected. Q-BRM enables this segmentation with biological precision — converting a broad adverse effect into a manageable, predictable, and avoidable risk for the majority of patients.

Pipeline Safety Enrichment

For agents in late-stage development, Q-BRM's digital twin simulations can prospectively model off-target tissue effects in patient subpopulations defined by genomic and pathway characteristics — enabling safety-enriched trial design that reduces late-stage failures and post-approval surprises.

Conclusion

The future of pharmacovigilance is not more data — it is deeper biological reasoning applied to individual patients. The two case studies presented in this paper illustrate a simple but profound truth: the mechanisms that determine whether a drug harms a patient are written in that patient's biology, long before the drug is prescribed. The challenge has always been reading and reasoning from those signals at the speed and scale that clinical practice demands.

Q-BRM is designed to meet that challenge. By combining causal AI, mechanistic pathway science, and patient-specific digital twin infrastructure, GNQ InSilico has built a platform that can transform pharmacovigilance from a retrospective surveillance function into a prospective, personalized, and mechanistically grounded layer of clinical intelligence.

For pharmaceutical organizations committed to leading the next era of drug safety — one defined not by population statistics, but by individual biological truth — Q-BRM represents a foundational capability worth building toward, together.

*GNQ InSilico, Inc. | gnqinsilico.com | Confidential
For partnership inquiries, contact: business@gnqinsilico.com*