GNQ INSILICO

Silicon-to-Patient™ | Q-BRM – The Power behind GNQ Precision Medicine Platforms

CLINICAL PROVIDER EDITION

WHITE PAPER

Q-BRM for Clinical Decision Support

How GNQ Insilico's proprietary Reasoning Model enables precision medicine at personalized level

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

Executive Summary

Clinical decision support today offers clinicians one of two things: a correlation-based probability derived from patients who looked like this patient in the past, or a guideline recommendation derived from the average patient in a published trial. Neither answers the question that matters in the consultation room: ***what will happen to this specific patient under each available treatment option?***

GNQ Insilico's **Q-BRM — Quantum Biomedical Reasoning Model** — answers this question. Powered by a patient-specific Relational Dynamic Graph (RDG), quantum-computed molecular interaction energies, and Neural Stochastic Differential Equations, Q-BRM's Digital Twin Platform (DTP) generates individual-level causal treatment predictions with mechanistic rationale and calibrated uncertainty bounds — operationalizing the third rung of Pearl's Causal Ladder at the bedside.[1]

Q-BRM does not rely only on population averages or trial statistics as its primary reasoning source. Causal intelligence is generated from first principles — biology, chemistry, and quantum physics. Real-world evidence, such as the MMRF CoMMpass cohort, **validates** Q-BRM's predictions. This means Q-BRM delivers reliable, mechanistically grounded guidance even for patients with rare genotypes, novel therapy combinations, or profiles underrepresented in published trials.

> **The Clinical Question Q-BRM Answers**
> Standard AI: 'Patients with similar profiles responded to Drug A 68% of the time.'
> Q-BRM: 'This patient's BCL2 pathway is overactive due to t(11;14) translocation — venetoclax will restore apoptotic sensitivity. The PSMB5 A49T variant predicts bortezomib resistance at a median of 7 months. Starting with venetoclax and sequencing to daratumumab before month 6 maximizes durable response probability.'

1. The Limits of Current Clinical AI

Large language models and correlation-based clinical AI tools have proliferated rapidly. Rigorous independent evaluation reveals a consistent pattern: these systems perform well on medical knowledge tests but fail systematically on realistic clinical decision tasks involving data integration, guideline adherence, and treatment planning.[3]

These limitations are not resolvable with more data or larger models. They reflect **architectural constraints** in correlation-based and pattern-matching systems — fundamental reasoning failures that persist regardless of model scale.[4]

For oncology centers and precision medicine programs specifically, three gaps are most clinically consequential:

1.1 Population Statistics Cannot Guide Individual Patients

A 68% response rate means nothing to the patient sitting in front of you. What matters is whether their specific molecular state — their exact genomic variants, pathway activation levels, and treatment history — predicts response or resistance. This requires a causal reasoning model analyzing the individual, not a population average.

1.2 Resistance Is Predictable, Not Inevitable

The conventional model treats resistance as an unpredictable event discovered at progression. Q-BRM treats resistance as an emergent causal process — one whose trajectory can be simulated before it manifests clinically. Proactive therapy sequencing before resistance emerges is categorically different from reactive switching after progression.

1.3 Novel Patients Need First-Principles Reasoning

Patients with rare mutations, novel combination therapies, or genomic profiles underrepresented in published trials receive the least reliable guidance from data-driven platforms — precisely because these platforms need historical data to function. Q-BRM's first-principles causal approach delivers the same quality of mechanistic reasoning regardless of whether prior data exists for the patient's specific profile.

2. Q-BRM's Clinical Intelligence Engine

2.1 Patient-Specific Relational Dynamic Graph (RDG)

For each patient, Q-BRM instantiates a **patient-specific Relational Dynamic Graph** — a live causal network encoding their individual genomic variant profile, pathway activation states, treatment history, and longitudinal biomarker trajectories. The RDG is continuously updated as new data arrives via FHIR R4/R5 integration with the clinic's EHR system.

The RDG encodes curated pathway relationships and 2M+ causal links with typed, directional biological semantics — not just associations, but mechanistic cause-and-effect chains from gene variants to pathway activity to clinical outcomes. This is the substrate on which all clinical reasoning operates.

2.2 Quantum-Accurate Drug-Target Modeling

Drug-protein binding is a quantum mechanical phenomenon. Classical statistical models cannot accurately capture electron correlation, orbital hybridization, or allosteric effects in binding pockets. Q-BRM employs **DMET+SQD** — IBM's near-term quantum chemistry framework — to

compute binding free energies at quantum accuracy for each drug-target interaction relevant to the patient's therapy options.[2]

This matters clinically when patients carry mutations that alter binding site geometry — such as PSMB5 A49T in proteasome inhibitor resistance — where the difference between predicted sensitivity and predicted resistance hinges on quantum-level changes in molecular interaction that classical models cannot resolve.

2.3 Digital Twin Counterfactual Simulation

Q-BRM's **Digital Twin Platform (DTP)** deploys Neural Stochastic Differential Equations over the patient-specific RDG to simulate continuous biomarker trajectories under each treatment option being considered. The output is not a point prediction — it is a full trajectory with calibrated uncertainty bounds, answering:

- How will M-protein, serum free light chains, and key biomarkers evolve under Treatment A vs Treatment B over the next 12 months?
- At what point is resistance most likely to emerge, and through which pathway mechanism?
- What is the optimal therapy sequencing strategy to maximize durable response while preserving future treatment options?
- What would have happened if this patient had received a different first-line therapy? (counterfactual retrospective analysis for treatment-experienced patients)

This is Pearl's Rung 3 causal reasoning — counterfactual inference — operationalized at the individual patient level.[1] It is the highest level of causal intelligence available, and it is only possible with a fully specified patient-specific causal model.

2.4 Pearl's Causal Ladder in the Clinic

Pearl's Rung	Formal Query	Question Type	Q-BRM Implementation
Rung 1 — Seeing	P(Y | X)	Association	RDG instantiation from multi-omics. IVS scores 681+ typed causal pathways. Relational structure encodes causal semantics — not text embeddings.
Rung 2 — Doing	P(Y | do(X))	Intervention	Q-BRM's Quantum RCM: DMET+SQD computes quantum-accurate causal edge weights. Neural ODEs propagate interventional states continuously through the relational causal graph.

Pearl's Rung	Formal Query	Question Type	Q-BRM Implementation
Rung 3 — Imagining	P(Yx | X'=x', Y'=y')	Counterfactual	Digital Twin Neural SDEs simulate diverging biomarker trajectories under alternative treatment paths — patient-specific counterfactual inference at the individual level.

3. Clinical Workflow Integration

3.1 FHIR R4/R5 and EHR Integration

GNQ's Digital Twin Platform is built natively on FHIR R4/R5 standards, enabling direct integration with modern EHR systems. Multi-omics data (RNA-seq, variant calls, methylation profiles), clinical records, laboratory values, and imaging biomarkers are ingested through standardized FHIR APIs. The patient-specific RDG is updated in near-real-time as new data arrives, ensuring clinical recommendations reflect the most current patient state.

3.2 Federated Architecture and HIPAA Compliance

GNQ's federated scoring architecture maintains patient data privacy by running AI models locally at each clinical site to generate pathway scores. Only anonymized causal scores — not raw patient data — are transferred centrally for platform learning. The DTP is fully HIPAA-scoped with data governance frameworks appropriate for multi-site oncology deployments.

3.3 Clinical Decision Support Output Format

Q-BRM delivers clinical reasoning insights in formats designed for integration into existing workflows:

- **Treatment Recommendation Report:** Ranked treatment options with causal pathway rationale, predicted response trajectories, resistance timeline projections, and uncertainty bounds.
- **Molecular Support:** Structured causal analysis of genomic findings, pathway implications, and mechanistic treatment rationale for board review.
- **Monitoring Alert System:** Real-time tracking of biomarker trajectories against predicted values — flagging deviations that indicate emerging resistance or unexpected response patterns.
- **Retrospective Counterfactual Analysis:** For treatment-experienced patients, Q-BRM models alternative treatment history pathways to inform optimal sequencing of future lines.

4. Clinical Validation and Deployments

4.1 MMRF CoMMpass: Multiple Myeloma Validation

GNQ's clinical validation has been conducted in oncology. One such validation is in Multiple Myeloma using the MMRF CoMMpass cohort — 1,100+ multiple myeloma patients with longitudinal genomic profiling, detailed treatment records, and outcome data. Q-BRM's causal predictions such as bortezomib resistance (PSMB5 A49T), venetoclax sensitivity (BCL2 t(11;14) pathway), and sequential therapy optimization have been validated against CoMMpass observed outcomes.

Key validated capabilities in MM include proteasome inhibitor resistance prediction, BCMA pathway bispecific antibody response modeling, and sequential therapy counterfactual simulation across treatment lines.

4.2 KAIRA Health: Multi-Site Clinical Deployment

Q-BRM's Digital Twin Platform is deployed under a multi-year agreement with KAIRA Health across clinics in Toronto, New York, Miami, Chicago, and others. This deployment represents the first at-scale clinical operation of patient-specific RDG instantiation and counterfactual treatment simulation in a live clinical setting, generating real-world operational validation of Q-BRM's clinical workflow integration.

4.3 Therapeutic Breadth

Q-BRM's RDG-Driven Quantum RCM architecture is indication-agnostic. Current active clinical applications include NSCLC (EGFR pathway combination therapy), ALS (multi-pathway neurodegenerative combination approach), Cognitive Warrior Program for high-performance cognitive optimization — demonstrating platform generalizability across oncology, neurodegenerative disease, preventive medicine and other areas.

5. Competitive Differentiation

Clinical Capability	Standard CDSS	Correlation AI	GNQ Q-BRM DTP
Treatment Recommendation Basis	⚠ Guideline + population stats	✗ Pattern matching	✓ Patient-specific causal pathway
Individual Counterfactuals	✗ Not available	✗ Not available	✓ Neural SDE Digital Twin
Resistance Foresight	✗ Reactive (post-progression)	✗ Not available	✓ Proactive — pre-progression
Mechanistic Rationale	⚠ Guideline reference	✗ Black box output	✓ Pathway-level causal explanation

Clinical Capability	Standard CDSS	Correlation AI	GNQ Q-BRM DTP
Multi-Omics Integration	⚠ Genomics only (typically)	⚠ Limited	✅ Full multi-omics + FHIR
Novel / Rare Genotypes	✗ No prior data	✗ No prior data	✅ First-principles causal model
FHIR R4/R5 Integration	⚠ Variable	⚠ Variable	✅ Native FHIR R4/R5
HIPAA Compliance	✅ Standard	✅ Standard	✅ Federated + HIPAA-scoped
Uncertainty Quantification	✗ Point estimates	✗ Point estimates	✅ Calibrated credible intervals
Combination Therapy Guidance	⚠ Guideline-based	✗ Limited	✅ Causal pathway complementarity

6. Clinical Engagement Model

GNQ engages clinical providers under a subscription model designed for the operational realities of oncology programs and precision medicine centers

Implementation includes dedicated clinical integration support, EHR connectivity configuration, clinical staff training, and ongoing platform updates as Q-BRM's knowledge graph and quantum simulation capabilities expand.

7. Conclusion

The question that matters in precision medicine is not 'what do patients like this usually respond to?' — it is 'what will this patient respond to, and why?'

Q-BRM answers the second question through individual-patient causal reasoning grounded in quantum physics, systems biology, and formal causal inference.

For oncology centers, neurodegenerative programs, and precision medicine practices, Q-BRM delivers a categorically different quality of clinical decision support: not a probability derived from the past, but a mechanistic prediction grounded in the biology of the patient in front of you — with the causal rationale to explain it, and the calibrated uncertainty to trust it.

Partner with GNQ Insilico

GNQ Insilico, Inc. | gnq.ai
To arrange a clinical platform demonstration, integration assessment, or pilot program discussion, contact the GNQ clinical partnerships team.

Glossary of Key Terms

Q-BRM — Quantum Biomedical Reasoning Model. GNQ's proprietary intelligence engine powering DAP, DSP, and DTP. Composed of the Relational Dynamic Graph (RDG), Quantum Relational Causal Model (Quantum RCM), and agentic BRM orchestration layer.

RDG — Relational Dynamic Graph. GNQ's proprietary knowledge graph substrate: 681+ pathway relationships, 20,000+ genes, 2M+ causal links. Typed, weighted, directional biological relationships with continuous temporal dynamics.

Quantum RCM — Quantum Relational Causal Model. The causal inference framework within Q-BRM: DMET+SQD quantum edge weights + Neural ODE/SDE propagation over the RDG. Extends Pearl's SCMs to relational heterogeneous graphs with feedback loops.

DAP — Drug Assessment Platform. GNQ's causal pathway scoring and drug evaluation engine using Integral Variant Scoring (IVS).

DSP — Drug Simulation Platform. GNQ's dynamic causal simulation engine using Neural ODEs and Neural SDEs.

DTP — Digital Twin Platform. GNQ's HIPAA-compliant patient-specific causal modeling and clinical decision support layer.

BRM — Biomedical Reasoning Model. GNQ's agentic orchestration layer built on LangGraph and MCP. Handles natural language intent parsing and orchestrates calls to DAP, DSP, and DTP.

DMET+SQD — Density Matrix Embedding Theory + Sample-based Quantum Diagonalization. IBM's near-term quantum chemistry framework for computing molecular ground states on real quantum hardware without fault-tolerant qubits.

IVS — Integral Variant Scoring. GNQ's proprietary methodology for scoring pathway-level causal impact from patient genomic variant profiles.

Neural ODE — Neural Ordinary Differential Equation. Continuous-time neural architecture for modeling causal state evolution (interventional inference — Pearl Rung 2).

Neural SDE — Neural Stochastic Differential Equation. Extends Neural ODEs with principled stochasticity for uncertainty quantification; enables counterfactual simulation (Pearl Rung 3).

SCM — Structural Causal Model. Judea Pearl's formal framework for causal reasoning using DAGs and structural equations.

DAG — Directed Acyclic Graph. Required by standard SCMs. Cannot represent biological feedback loops — a key limitation addressed by GNQ's RDG.

rNPV — Risk-adjusted Net Present Value. Standard pharma valuation metric for drug programs; GNQ improves rNPV inputs by raising probability of success and accelerating go/no-go timing.

References

[1] **Pearl, J.** "Causality: Models, Reasoning, and Inference." Cambridge University Press, 2nd ed., 2009. ISBN: 978-0-521-89560-6. Foundational framework for Structural Causal Models, do-calculus, and the three-rung Causal Ladder.

[2] Robledo-Moreno, J., Motta, M., et al. "Chemistry beyond the scale of exact diagonalization on a quantum-centric supercomputer." *Science Advances* 11(25), eadu9991, June 2025. DOI: 10.1126/sciadv.adu9991. IBM Heron + Fugaku; 45–77 qubits; validates DMET+SQD on real near-term hardware.

[3] Hager, P., Jungmann, F., et al. "Evaluation and mitigation of the limitations of large language models in clinical decision-making." *Nature Medicine* 30(9), 2613–2622, 2024. DOI: 10.1038/s41591-024-03097-1. Across 2,400 real cases, LLMs performed significantly worse than physicians on realistic clinical tasks.

[4] Song, P., et al. "Large Language Model Reasoning Failures." *arXiv*:2602.06176, February 2026. First comprehensive survey of systemic LLM reasoning failure modes; directly relevant to reliance on LLMs for clinical inference.

[5] Maier, M., Marazopoulou, K., Arbour, D., Jensen, D. "A Sound and Complete Algorithm for Learning Causal Models from Relational Data." *Proceedings UAI 2013*, arXiv:1309.6843. Introduces Relational Causal Models (RCM) as a formally sound extension of Pearl's SCMs to heterogeneous relational graphs.

[6] Chen, R.T.Q., Rubanova, Y., Bettencourt, J., Duvenaud, D. "Neural Ordinary Differential Equations." *NeurIPS 2018*, pp. 6572–6583. arXiv:1806.07366. Foundational Neural ODE paper; underpins GNQ's continuous-time causal trajectory simulation.