Is Quantum Ready for Biology?

What It Would Take to Put a Quantum Step Inside a Regulated Biopharma Pipeline

Birds-of-a-Feather | NISQ & Fault-Tolerant | Provenance | Regulated Pipelines

---

Alexander Chernov
Associate Principal Data Engineer, Data Science and Modeling
BioPharmaceuticals R&D, AstraZeneca

Monday 14 September 2026 | 15:00–16:30 EDT
Room 715A, Metro Toronto Convention Centre

IEEE Quantum Week (QCE) 2026 — Birds-of-a-Feather | BOF-2099

Views and provocations are my own; the artifacts are discussion objects, not AstraZeneca positions. The four seed positions are arguments assigned for debate. Participation under an approved AstraZeneca public release.

Who is moderating, and on what standing

linkedin.com/in/thedoytsujin

Alexander Chernov

Associate Principal Data Engineer, Data Science and Modeling — BioPharmaceuticals R&D, AstraZeneca, Greater Toronto Area

He works on the control planes behind regulated biopharmaceutical science: the systems deciding what may be executed, and what a computed number must prove before anyone may act on it.

He has reproduced a published quantum-genomics method on PennyLane, trained a variational quantum classifier downstream of it, and built a mechanistic fed-batch CHO bioreactor twin under a contract of in-process limits tied to a Critical Quality Attribute. Associate editor, IEEE Toronto Section newsletter; 2026 work at IEEE CCECE, EMBC, CBS and BigDataService.

A systems engineer working inside biology, with an **MSc in theoretical physics** (photoionization, confined atoms), who had not pursued quantum computing academically until recently. For a moderator on this question that is an asset: **no quantum result to defend.**

House rules, before anything else

- **Steelman before you rebut.** State the strongest version of the position you are about to attack. If you cannot, you are not ready to attack it.
- **Turns are time-boxed and evidence-based.** Two minutes. Name what would change your mind.
- **The adversarial energy is aimed at claims, never at the people holding them.**
- **Every segment closes with a written synthesis**, not an open argument.

And one rule that binds me rather than you:

I hold no result to defend. *I came to quantum computing academically only recently, from theoretical physics. Everything I put on the table today is offered to be attacked, and I will report what this room concludes even when it is unflattering to my own artifact.*

Views and provocations here are my own. The artifacts are discussion objects, not AstraZeneca positions. The four seed positions are **arguments assigned for debate** — adopting or defending one in this room does not imply endorsing it outside it.

Two conversations about quantum computing in biology

The one happening at this conference.

Eigensolvers for molecular ground states. QAOA for genomics. Quantum kernels over molecular feature vectors.

The demonstrations are real and the mathematics is sound. They run at a scale where the answer is checkable classically — which establishes that the method *works*, not that it is *worth using*.

That limitation is explicit in the technical work. It is easier to lose in the broader narratives.

The one happening inside a pharmaceutical company.

Not about algorithms at all.

It is about whether a computational step can be trusted well enough to inform a decision a regulator will later examine.

In that world a result carries an obligation before anyone may act on it.

Neither conversation is wrong.

What does not exist between them is a mature interface.

VQE: Peruzzo et al. [1]. QAOA: Farhi, Goldstone, Gutmann [2]. Quantum kernels: Havlíček et al. [3].

Where a regulated pipeline begins

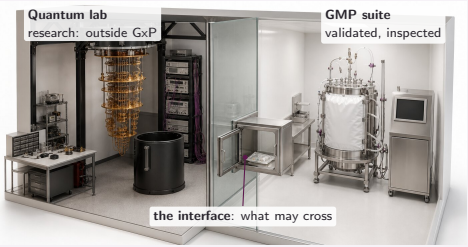

Illustrative render, not a real facility. The lab may compute anything; the suite may act only on what crosses with its record.

Four of the five ballot targets start outside GxP. The question is not whether they may run there. It is what they must carry to cross.

Stage	A decision made here	Ballot targets that start here
Discovery outside GLP [4]	which molecule to advance	molecular simulation, property prediction, protein structure, genomics
GxP from here on — by use, not by computation [5]		
Nonclinical safety GLP, 21 CFR 58 [4]	is it safe to dose people?	a prediction, once in a study record
Clinical GCP, ICH E6(R3) [6]	dose, efficacy, safety	genomics, once it informs trial data
Manufacturing, QC CGMP, 21 CFR 210–211 [7]	may this batch be released?	bioprocess optimization (the C4 twin)
Submission agency review	approve, and on what evidence	everything above, as submitted evidence

US regulations shown as one jurisdiction. ICH guidelines are not binding in themselves; each ICH regulatory member implements them — FDA, the European Commission with EMA, and Health Canada among them [8]. A computed prediction already inside a regulated decision, from small molecules: two complementary (Q)SAR predictions under ICH M7(R2) [9].

One antibody, from cell bank to release

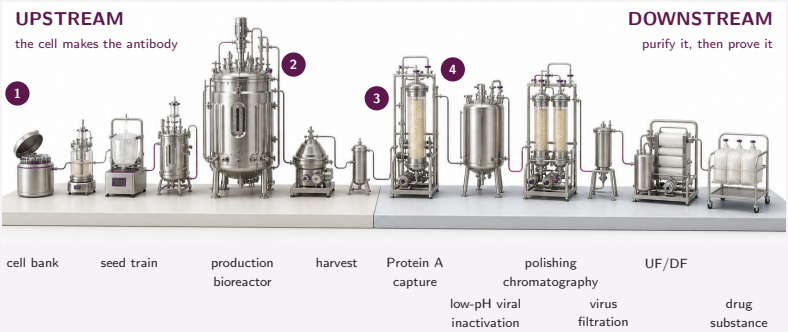

Connected end to end. What the cell decides upstream — the antibody's glycan profile among it — travels down the line [10]. Purification removes impurities [11]; standard purification does not change afucosylation [12]. Each batch is tested against its specification before release [13].

Numbered: where quantum computation has been proposed — locations, not verdicts; placing them is Ballot 1. **1** molecule and antibody design, before the line [14], [15]; **2** bioprocess optimization, a perspective with no results [16]; **3** scheduling, shown for manufacturing in general [17]; **4** chromatography column packing, a preprint [18]. Unmarked: none found, 13 Sep 2026.

Illustrative render. Unit sequence: cell banks, ICH Q5D [19]; seed train, 10–14 day fed-batch production and harvest, Kelley [20] and Li et al. [21]; downstream platform, Liu et al. [11]; step order varies between company platforms [20].

What a computed result must carry before it may influence a regulated decision

A result in a validated pipeline is not a number. It is a number plus four obligations:

1. **What produced it** — enough provenance to reconstruct the conditions: method, configuration, versions, inputs.
2. **Under what conditions** — calibration status and audit trail, not a claim to have captured every register of the machine.
3. **Which version of which model** — pinned, not "current."
4. **Evidence that repeating the computation produces a result consistent with its validated reproducibility criterion.**

This is not bureaucratic ornament. It is the mechanism by which a computed number is allowed to influence a decision that is inspected later.

A synthesis for this session, adapted from — not quoted from — metadata and audit trail as defined in FDA's data-integrity guidance [22]; audit trails and system validation, 21 CFR 11.10(a),(e) [5]; calibration, 21 CFR 58.63 [4] and 211.160(b)(4) [7]; lifecycle and configuration control, GAMP 5 [23]; reproducibility of methods, 21 CFR 211.165(e) [7], and precision acceptance criteria, ICH Q2(R2) [24]; credibility of a model for a context of use, ASME V&V 40 [25] and FDA's draft AI guidance [26].

What a quantum device returns

- A **distribution**, not a number — so the reproducibility criterion cannot be a tolerance around a deterministic value. It has to be a statement about a sampled one, and **somebody has to decide what that statement is.**
- From an instrument whose **calibration drifted** since yesterday. In a measurable sense the machine that produced yesterday's result is not the machine you have today.
- Under a **configuration** — transpilation, qubit mapping, shot count, error mitigation. Quantum stacks expose *pieces* of this state, but generally do not emit it as a **complete, durable evidence package suitable for downstream governed use**.

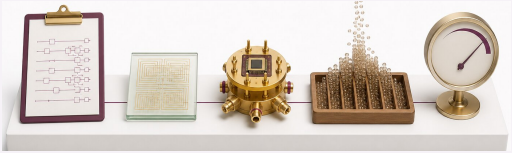

A biologist asking *"can I reproduce this?"* is asking a question the quantum stack cannot answer in the form it is asked.

Histogram: 1,000 shots/run, sd 7.9 mHa, five times the chemical-accuracy half-width (100,000 shots: 0.81 mHa); H_2 STO-3G, PennyLane, seeded, `figures/gen_shots.py`. Drift: Proctor et al. [27]. Configuration: [28], [29].

The standing agenda: three questions, in this order

Q1 **What evidence must accompany a quantum result** before a regulated workflow can even *evaluate* it?

Q2 **Which biological use cases are credible under each regime** — investigable on NISQ now, dependent on fault tolerance, or lacking a defensible quantum rationale at all?

Q3 **Which governance requirements are regime-independent** — true whether the device is noisy or fault-tolerant?

Everything else on the agenda is in service of these three. If we resolve nothing else, we resolve these.

Two horizons, one interface

A fair objection to everything I just said: it conflates two regimes. The organizing move for the whole session: sort *every* requirement into three bins.

NISQ-specific N

Noisy, intermediate-scale devices [30]. An experimental step that must survive insertion into a governed workflow *today* without destroying provenance, reproducibility or baseline comparability.

- Device state
- Configuration

Regime-independent RI

True whether the device is noisy or fault-tolerant.

- Reproducibility
- Confidence
- Provenance
- Classical baseline
- Failure disclosure

five of the eight strawman rows

Fault-tolerant-specific FT

Error-corrected logical qubits [31], [32]. A future validated component whose governance we can begin to specify *now*, before the interface is fixed by whoever ships first.

- Logical error rate

The claim I am asking you to contest: most of what a regulated pipeline demands lands in the middle bin.

Rows: the strawman checklist, Table I of the accepted proposal. NISQ: Preskill [30]. Surface-code fault tolerance and logical error rates: Fowler et al. [31]; below-threshold operation, Google Quantum AI [32].

On the table: four artifacts, all modest by construction

Their value here is not that they are impressive. It is that they are real enough to argue with.

- **A gate that refuses quantum results it cannot certify.** Live, in your browser, right now: `qubit.agenticdatasets.org`. No backend, no hardware. It scores itself against my own checklist and **fails four of eight rows**.
- **A published quantum-genomics method, reproduced.** QuASeR's QAOA *de novo* assembly [33] on PennyLane. Nine qubits, five-base contig, verified against the brute-force optimum.
- **A trained quantum classifier, downstream.** It does not beat a linear classifier and I make no claim that it does.
- **A regulated decision, from the other side of the wall.** A mechanistic fed-batch CHO bioreactor twin under a contract of in-process limits. **Not quantum.** It is the yardstick.

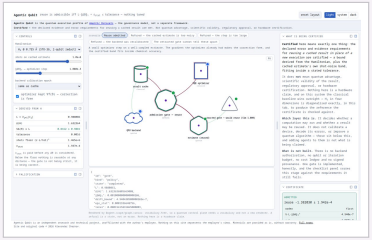

As served on 13 September 2026: the gate admitting a reuse, the retracted gate beside it.

Scan to open it now

`qubit.agenticdatasets.org`

QuASeR: Sarkar, Al-Ars and Bertels [33].

Four seed positions. They are deliberately incompatible.

P1 **The bottleneck is not qubits, it is evidence.** *(Mine. The session succeeds if you dismantle it.)*

P2 **This session is fifteen years premature.** *(The strongest objection to P1, and it deserves the room.)*

P3 **Only quantum chemistry is real; the rest should stop.**

P4 **The interface is the whole problem, not the advantage.**

You may adopt one, or reject all four. They exist so the room has something to disagree with immediately.

I will be actively looking for the strongest biological skeptic in this room, and I will give them the floor early.

P1, and what would kill it

P1 — the bottleneck is not qubits, it is evidence.

Grant a useful quantum machine tomorrow. Regulated biology still could not consume it. The gap is a governance and instrumentation gap: platforms expose *pieces* of the required evidence, but the complete, durable evidence interface a governed workflow needs is still immature — and the conventions for it are being set now, by whoever ships first.

What would refute it — and I want one of these from you:

- A platform that *already* emits device state, execution configuration and attributable lineage as one complete, durable package fit for governed downstream use. **Name one and I substantially revise P1; show me it is the norm rather than the exception and I withdraw it.**
- A demonstration that this evidence can be reconstructed *post hoc*, so nothing is being locked in and the urgency is manufactured.
- An argument that fault tolerance dissolves the requirement rather than inheriting it.

Sources for the framing (1 of 2)

[1] A. Peruzzo, J. McClean, P. Shadbolt, et al., "A variational eigenvalue solver on a photonic quantum processor," *Nat. Commun.*, vol. 5, art. 4213, 2014, doi:10.1038/ncomms5213.

[2] E. Farhi, J. Goldstone, and S. Gutmann, "A quantum approximate optimization algorithm," arXiv:1411.4028, 2014.

[3] V. Havlíček, A. D. Córcoles, K. Temme, et al., "Supervised learning with quantum-enhanced feature spaces," *Nature*, vol. 567, pp. 209–212, 2019, doi:10.1038/s41586-019-0980-2.

[4] U.S. FDA, "Good laboratory practice for nonclinical laboratory studies," 21 CFR Part 58: §58.3(d), definition of nonclinical laboratory study; §58.63, equipment calibration.

[5] U.S. FDA, "Electronic records; electronic signatures," 21 CFR Part 11: §11.1(b), scope; §11.10(a),(e), validation and audit trails; §§11.50, 11.70, signatures.

[6] ICH, "Guideline for Good Clinical Practice E6(R3)," Step 4, 6 Jan. 2025.

[7] U.S. FDA, current good manufacturing practice, 21 CFR Parts 210–211: §210.2(a), Parts 600–680 for biologics; §211.160(b)(4), calibration; §211.165, testing and release, incl. (e) reproducibility of test methods.

[8] ICH, "The International Council for Harmonisation (ICH): An overview," Nov. 2025; ICH Assembly Rules of Procedure, v17.0, 13 May 2025.

[9] ICH, "Assessment and control of DNA reactive (mutagenic) impurities in pharmaceuticals to limit potential carcinogenic risk M7(R2)," Step 4, 3 Apr. 2023, §6.

[10] S. Sha, C. Agarabi, K. Brorson, D.-Y. Lee, and S. Yoon, "N-glycosylation design and control of therapeutic monoclonal antibodies," *Trends Biotechnol.*, vol. 34, no. 10, pp. 835–846, 2016, doi:10.1016/j.tibtech.2016.02.013.

[11] H. F. Liu, J. Ma, C. Winter, and R. Bayer, "Recovery and purification process development for monoclonal antibody production," *mAbs*, vol. 2, no. 5, pp. 480–499, 2010, doi:10.4161/mabs.2.5.12645.

[12] A. J. Rish, K. Siddiquee, Z. Huang, J. Xu, C. A. Anderson, M. C. Borys, and A. Khetan, "Strategies for controlling afucosylation in monoclonal antibodies during upstream manufacturing," *Biotechnol. J.*, vol. 18, no. 7, art. 2200604, 2023, doi:10.1002/biot.202200604.

[13] ICH, "Specifications: Test procedures and acceptance criteria for biotechnological/biological products Q6B," Step 4, 10 Mar. 1999.

[14] N. S. Blunt, J. Camps, O. Crawford, et al., "Perspective on the current state-of-the-art of quantum computing for drug discovery applications," *J. Chem. Theory Comput.*, vol. 18, no. 12, pp. 7001–7023, 2022, doi:10.1021/acs.jctc.2c00574.

[15] J. Allcock, A. Vangone, A. Meyder, S. Adaszewski, M. Strahm, C.-Y. Hsieh, and S. Zhang, "The prospects of Monte Carlo antibody loop modelling on a fault-tolerant quantum computer," *Front. Drug Discov.*, vol. 2, art. 908870, 2022, doi:10.3389/fddsv.2022.908870.

[16] I. M. Hernández-Romero, S. M. Gibford, W. Clements, C. J. Savoie, A. Flores-Tlacuahuac, and S. S. Mansouri, "Exploring hybrid quantum-classical algorithms for multiscale bioprocess optimization," *Ind. Eng. Chem. Res.*, vol. 64, no. 19, pp. 9484–9499, 2025, doi:10.1021/acs.iecr.4c04624.

[17] D. Venturelli, D. J. J. Marchand, and G. Rojo, "Quantum annealing implementation of job-shop scheduling," arXiv:1506.08479, 2015.

Sources for the framing (2 of 2)

[18] B. Hall, I. Njoroge, C. Campbell, B. Thotakura, R. Rines, V. Omole, and M. Qadan, "Sphere packing on a quantum computer for chromatography modeling," arXiv:2412.00601, 2024, preprint.

[19] ICH, "Derivation and characterisation of cell substrates used for production of biotechnological/biological products Q5D," Step 4, 16 Jul. 1997.

[20] B. Kelley, "Industrialization of mAb production technology: The bioprocessing industry at a crossroads," *mAbs*, vol. 1, no. 5, pp. 443–452, 2009, doi:10.4161/mabs.1.5.9448.

[21] F. Li, N. Vijayasankaran, A. Shen, R. Kiss, and A. Amanullah, "Cell culture processes for monoclonal antibody production," *mAbs*, vol. 2, no. 5, pp. 466–479, 2010, doi:10.4161/mabs.2.5.12720.

[22] U.S. FDA (CDER, CBER, CVM), "Data integrity and compliance with drug CGMP: Questions and answers," guidance for industry, Dec. 2018.

[23] ISPE, *GAMP 5: A Risk-Based Approach to Compliant GxP Computerized Systems*, 2nd ed., Jul. 2022.

[24] ICH, "Validation of analytical procedures Q2(R2)," Step 4, 1 Nov. 2023, §4 Glossary.

[25] ASME, *Assessing Credibility of Computational Modeling through Verification and Validation: Application to Medical Devices*, V&V 40-2018.

[26] U.S. FDA, "Considerations for the use of artificial intelligence to support regulatory decision-making for drug and biological products," draft guidance, Jan. 2025.

[27] T. Proctor, M. Revelle, E. Nielsen, et al., "Detecting and tracking drift in quantum information processors," *Nat. Commun.*, vol. 11, art. 5396, 2020, doi:10.1038/s41467-020-19074-4.

[28] T. Peham, L. Burgholzer, and R. Wille, "Equivalence checking of parameterized quantum circuits: Verifying the compilation of variational quantum algorithms," in *Proc. ASP-DAC*, 2023, pp. 702–708, doi:10.1145/3566097.3567932.

[29] G. Bisicchia, J. Alvarado-Valiente, J. Romero-Álvarez, J. Garcia-Alonso, J. M. Murillo, and A. Brogi, "Maximizing quantum hardware utilization via multiprogramming circuits and shot-wise distribution," *Inf. Softw. Technol.*, vol. 191, art. 108005, 2026, doi:10.1016/j.infsof.2025.108005.

[30] J. Preskill, "Quantum computing in the NISQ era and beyond," *Quantum*, vol. 2, p. 79, 2018, doi:10.22331/q-2018-08-06-79.

[31] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, "Surface codes: Towards practical large-scale quantum computation," *Phys. Rev. A*, vol. 86, art. 032324, 2012, doi:10.1103/PhysRevA.86.032324.

[32] Google Quantum AI and Collaborators, "Quantum error correction below the surface code threshold," *Nature*, vol. 638, pp. 920–926, 2025, doi:10.1038/s41586-024-08449-y.

[33] A. Sarkar, Z. Al-Ars, and K. Bertels, "QuASeR: Quantum accelerated de novo DNA sequence reconstruction," *PLoS ONE*, vol. 16, no. 4, e0249850, 2021, doi:10.1371/journal.pone.0249850.