# 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

Presented Monday 14 September 2026 | Room 715A, Metro Toronto Convention Centre

Published edition | doi:10.5281/zenodo.22841353 | CC BY 4.0, except the portrait photograph

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

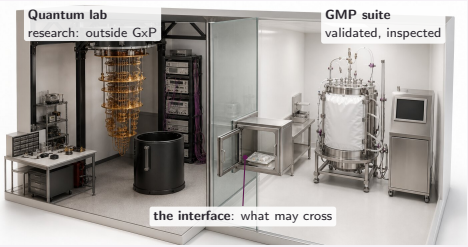


Illustrative render, not a real facility. The lab may compute anything; the suite may act only on what crosses with its record.

**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**<br>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**<br>GLP, 21 CFR 58 [4] | is it safe to dose people? | a prediction, once in a study record |
| **Clinical**<br>GCP, ICH E6(R3) [6] | dose, efficacy, safety | genomics, once it informs trial data |
| **Manufacturing, QC**<br>CGMP, 21 CFR 210–211 [7] | may this batch be released? | bioprocess optimization (the C4 twin) |
| **Submission**<br>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

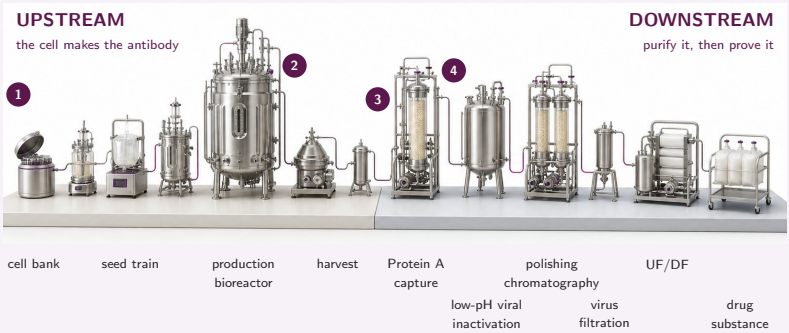


**Connected end to end.** What the cell decides upstream — the antibody's glycan profile among it — travels down the line [10]. Purification removes impurities [11]; standard purification does not change afucosylation [12]. Each batch is tested against its specification before release [13].

**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**.

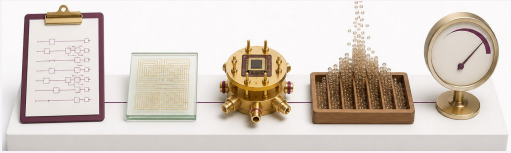


A biologist asking *"can I reproduce this?"* is asking a question the quantum stack cannot answer in the form it is asked.

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.

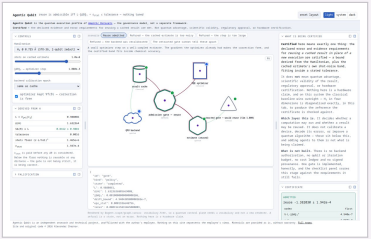

As served on 13 September 2026: the gate admitting a reuse, the retracted gate beside it.

**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.

# P1 — The bottleneck is not qubits, it is evidence

The moderator's position. He will argue it. The session succeeds if the room dismantles it.

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.

**Fixing it later will be harder than fixing it now.**

# P2 — This session is fifteen years premature

The strongest objection to P1, and it deserves the room.

There is no biological problem where quantum beats classical this decade. Current bio-quantum demonstrations are **toys** — checkable by brute force, beaten by classical methods that took an afternoon to write.

Worrying about their regulatory provenance is like drafting flight-safety rules for a paper airplane. **Build the fault-tolerant machine first.**

Governance is a distraction dressed as rigor: it lets people feel productive without confronting the hardware.

# P3 — Only quantum chemistry is real; the rest should stop

The one defensible near-term biological target is **quantum simulation of molecular systems**, where the problem is natively quantum and the classical baseline genuinely strains.

Genomics-as-QUBO, quantum classifiers over biological feature vectors, and bioprocess optimization are the field's weakest work: they take a problem classical computing already handles well, encode it awkwardly, and run it worse.

**They persist because they generate papers.** The field should concentrate its credibility on chemistry and abandon the rest.

# P4 — The interface is the whole problem, not the advantage

Nobody will replace a classical pipeline stage with a quantum one.

Quantum subroutines will sit **inside** classical loops, and the entire engineering question is the boundary: what crosses it, what guarantees hold across it, and what happens when the quantum side returns a *distribution* to a classical caller that expected a *number*.

**Advantage is an argument about the wrong layer.**

# The artifacts, in five segments

**C1** The certificate — a gate that refuses

**C2** QuASeR reproduced, and the indictment

**C3** The trained classifier that beats nothing

**C4** The process, and the bioreactor — the yardstick

**C5** Steelmanning P2

Prepared for the session as material to argue with; each segment stands alone. This edition omits the speaker notes and the slides used to collect the room's input.

# C1 — The question, narrowly

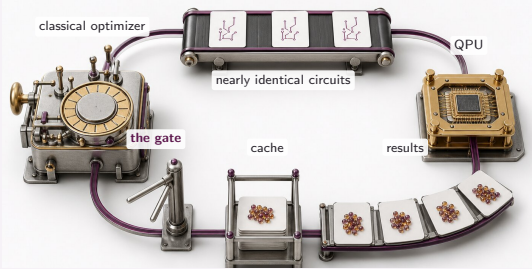


A variational optimizer resubmits nearly the same circuit thousands of times.

**May a cached result stand in for the next one?**

This is the smallest honest instance of the session's whole question. A cached result is a result whose provenance points at a *different* execution. Either you can certify that substitution or you cannot, and the machinery for saying which is the machinery a regulated pipeline needs everywhere.

Live, client-side, no backend: `qubit.agenticdatasets.org`

# C1 — My own previous answer, withdrawn

In my recent research I answered it with a tuned similarity threshold:

$$\mathrm{Sim} = \alpha \cdot S_{\text{struct}} + \beta \cdot S_{\text{param}} + \gamma \cdot S_{\text{topo}}, \qquad \text{reuse iff } \mathrm{Sim} \geq \tau$$

**That answer is retracted.** Edit distance between circuits is neither necessary nor sufficient for result interchangeability.

The artifact keeps the retracted metric *on screen*, running, next to the replacement — so you can watch it get the answer wrong rather than take the retraction on trust.

# C1 — The replacement: a bound from the Hamiltonian, nothing tuned

| | | |
|---|---|---|
| **Fact 1** | $f(\theta) \in [\lambda_{\min}(H),\, \lambda_{\max}(H)]$ | for every $\theta$ |
| **Fact 2** | $\dfrac{\partial f}{\partial \theta_k} = \frac{1}{2}\left[f(\theta + \frac{\pi}{2} e_k) - f(\theta - \frac{\pi}{2} e_k)\right]$ | parameter shift, exact |
| **Fact 3** | $\left\lvert \dfrac{\partial f}{\partial \theta_k} \right\rvert \le \frac{1}{2}\Delta(H) \le \sum_{j \neq I} \lvert c_j \rvert =: L$ | $\mathcal{O}(\#\text{terms})$, no executions |

$L$ is read off the Pauli coefficients of $H$. **No threshold, no fitted constant, no calibration run.**

Exclude the identity term: it shifts $f$ without moving any derivative, and $\lambda_{\max} - \lambda_{\min}$ is invariant under adding $cI$. Including it roughly doubles $L$ for nothing — for $H_2$ the identity coefficient is the largest in the set.

Parameter-shift rule, as stated for gates $e^{-i\theta P/2}$ with $P$ a Pauli product (RY here): Mitarai et al. [34]; Schuld et al. [35]. Every partial derivative bounded by the sum of absolute Pauli coefficients: Gu et al. [36].

# C1 — The admissibility rule

Reuse is admissible if and only if

| | | |
|---|---|---|
| $L \cdot \|\Delta\theta\|_1 + \varepsilon_{\text{stat}} \leq \text{tolerance}$ | zeroth order | return the cached value |
| $\frac{1}{2} L \cdot \|\Delta\theta\|_1^2 + \varepsilon_{\text{stat}} \leq \text{tolerance}$ | first order | return it gradient-corrected |

The second form **costs nothing extra**. A gradient-based optimizer already computed $\nabla f(\theta)$ to take its step, so the correction is sitting in its state.

What comes back is not a stale value. It is a Taylor-corrected one with a certified band.

**Measured: the free correction buys $37\times$ the admissible step.**

# C1 — The two halves are different kinds of claim

$L \cdot \|\Delta\theta\|_1$

**Deterministic.** It cannot be exceeded. Zero violations is a theorem holding, not a discovery — which is why it is *run* rather than asserted.

$\varepsilon_{\mathbf{stat}} = zL/\sqrt{\mathbf{shots}}$

**A 95% band.** It *will* be exceeded about one time in twenty. That is what 95% means.

The composite inherits the weaker of the two, so it is reported as a 95% certified band and **never as a guarantee**.

A first pass scored them together and read the expected 3.1% tail as 57 violations. The verifier now scores them separately for exactly this reason. This slide is here because getting it wrong once is the most instructive thing in the artifact.

# C1 — Measured: reuse has a floor price, in shots

$\varepsilon_{\text{stat}}$ is paid *before* any $\Delta\theta$ is considered. So shots $\geq (zL/\text{tolerance})^2$, or nothing is reusable at any distance.

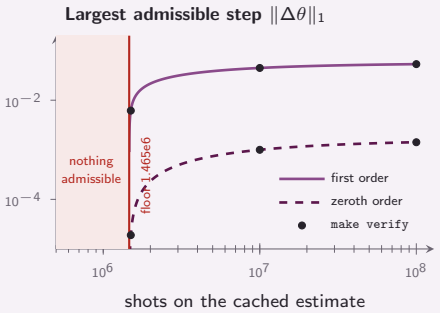


| shots | admit 0th | admit 1st | admit legacy |
|---|---|---|---|
| 1e5 | 0.0% | 0.0% | **78.9%** |
| 1e6 | 0.0% | 0.0% | **78.9%** |
| 1.5e6 | 0.0% | 39.7% | **78.9%** |
| 1e7 | 21.3% | 58.7% | **78.9%** |
| 1e8 | 24.9% | 60.4% | **78.9%** |

For $H_2$ at chemical accuracy the floor is $\mathbf{1.465 \times 10^6}$ **shots**. Below it the gate admits nothing — and it is not being strict, it is being correct.

Shot requirement $M \leq (\sum_\ell |w_\ell|/\varepsilon)^2$: Rubin et al. [37]. $H_2$ at 0.735 Å in STO-3G reduced to two qubits, and chemical accuracy ≈1.6 mHa: Kandala et al. [38]. 1,500 trials per shot count against an exact reference simulator. Curve: closed form of `shotsSweep`, reproducing the five verified rows; `figures/gen_floor.py`.

# C1 — Measured: what each gate actually admits

At 1e8 shots, tolerance = chemical accuracy:

| Gate | admitted | outside 95% band | deterministic violations | max error |
|---|---|---|---|---|
| legacy $\mathrm{Sim} \geq \tau$ *(retracted)* | 2363 | 864 (36.6%) | **863** | 1.2e-1 |
| certificate, zeroth order | 784 | 0 (0.0%) | **0** | 6.4e-4 |
| certificate, first order | 1821 | 57 (3.1%) | **0** | 4.8e-4 |

The first-order gate admits **2.3× as much as the zeroth** and still violates nothing deterministically. The 3.1% is the 95% band's own tail, arriving exactly where it was predicted to.

The retracted gate admits three times as much as the certified one. It is also wrong 863 times in a way no confidence interval covers.

# C1 — Why the retracted gate cannot be repaired by tuning

Same circuits, same $\Delta\theta$, two operators differing only by a factor of ten:

| | $L$ | certified band | verdict |
|---|---|---|---|
| $H_2$ | 0.988 | 6.4e-4 | **reuse** |
| $H_2 \times 10$ | 9.881 | 6.4e-3 | **refuse** |

$$\mathrm{Sim} = 0.9956 \geq \tau \quad \text{for } \textbf{both}.$$

One verdict, two operators. **No choice of $\tau$ repairs that**, because $\tau$ is applied to a quantity that never saw $H$.

This generalizes past caching. Any gate tuned on a proxy for the thing it is supposed to bound has this failure mode, and most governance heuristics are exactly that.

# C1 — Its own scorecard: four of eight, and the failures are the point

| Requirement | Hzn | Emits? | What it emits, or why it does not |
|---|---|---|---|
| Device state | N | yes | Every entry carries its calibration epoch, and the gate refuses across an epoch boundary first. Whether an epoch id is granular enough is the open question. |
| Reproducibility | RI | yes | A certified band. Not that the number repeats — that any repetition lies within a stated interval. |
| Configuration | N | no | Partial. Ansatz, parameters and shots recorded; transpilation, mapping and mitigation are not, because the reference simulator has none. |
| Confidence | RI | yes | 95%, split into a half that cannot be exceeded and a half that will be. |
| Provenance | RI | no | **Nobody signs it.** Lineage is recorded but unsigned, and an unsigned chain is a claim rather than evidence. The largest honest gap. |
| Classical baseline | RI | no | **The classical baseline wins outright.** $H_2$ in four dimensions is diagonalized exactly, in this tab, to produce the reference. |
| Failure disclosure | RI | yes | A refusal is a first-class rendered outcome with its reason attached, never a silent fallback to a stale value. |
| Logical error rate | FT | no | Out of regime. Asserting one would be inventing evidence. |

# C1 — What is actually new here, and what is not

**Published already, and I am not claiming it:**

- The shot requirement $M \leq (\sum_\ell |w_\ell|/\varepsilon)^2$ — Rubin, Babbush, McClean [37]; earlier in approximate form, Wecker, Hastings, Troyer [39].
- Bounding every partial derivative by $\sum |c_j|$ from Pauli coefficients — Gu et al., adaptive shot allocation [36].
- Caching quantum results across optimizer iterations — Tejedor, Conejero, Badia [40]. *Reuse only when the implemented circuit is the same after ZX reduction; no error bound.*
- Discarding an iteration when noise drift flips the gradient — DISQ [41]. *Skips rather than substitutes.*

**What survives as a contribution:**

- Error-bounded reuse **across parameter points**, not merely at equivalent ones.
- The deterministic and statistical halves **scored separately**, so the composite is never reported as a guarantee.

Two gaps. That is a narrow claim, and I would rather state it narrowly here than have it widened for me.

# C2 — A published quantum-genomics method, reproduced

QuASeR casts *de novo* genome assembly as an optimization problem:

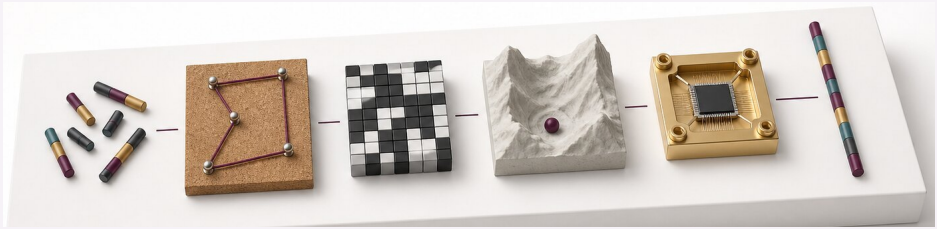


Re-implemented on an open, differentiable stack (PennyLane). **Nine qubits.** The circuit returns a distribution whose most probable *feasible* ordering — a valid permutation of the reads — is the optimal assembly, reconstructing the contig `GCGCA`.

**Verified against the brute-force optimum.** Not asserted against it.

Sarkar, Al-Ars, Bertels, *PLoS ONE* 16(4):e0249850, 2021.

# C2 — The verification is the point and the indictment at once

The instance is small enough to check exhaustively.

**Which is precisely why it proves nothing about advantage.**

Nine qubits. Five bases. A brute-force optimum computed in the same notebook.

**P3 will find this an easy target, and may well be right.** Genomics-as-QUBO takes a problem classical computing already handles well, encodes it awkwardly, and runs it worse.

So what is it doing in this session? Only this: a real bioinformatics task runs end to end on an open quantum stack, which makes it concrete enough to ask the governance question about. **The governance question is the one I am here for.**

# C2 — What was wrapped around it

**A dataset descriptor** travels with the data and states what it is and how it may be used:

- **Schema** — the shape of the data.
- **Provenance** — where it came from, and by what process.
- **Policies** — a qubit budget and an iteration budget, **checked before the job runs**.

The shift is small to state and large in consequence: the data stops being an inert input and starts carrying its own rules. A dataset that knows its own schema, lineage and limits can be validated, audited and **refused automatically** — instead of relying on whoever happens to be running the pipeline to remember the constraints.

"Agent" here means an orchestration controller that runs a stage, evaluates its policies and records what happened. No model, no LLM in this loop.

# C2 — And a gateway in front of the backend

It does not do the science. It decides what is allowed to run, where, at what cost, and it keeps the invariants.

- **Admission control.** A job is admitted only when the requesting tenant's policy allows the chosen backend. An unapproved device is **refused before a circuit ever reaches it**.
- **Backend abstraction.** The same job description targets a local simulator or a real machine — which is what lets a workflow move from simulation to hardware without rewriting the science.
- **Bounded cost.** Rate limits, plus a cache keyed on the circuit, so an identical circuit is served rather than re-run.
- **A structured event record**, so the orchestration is auditable as a whole and not merely per job.


The cache here serves only an identical circuit. Reusing a result at a different parameter point needs the bound in C1.

# C3 — A trained quantum classifier, downstream of assembly

Part 1 reconstructs a contig. It does not interpret it. So a second stage does: a variational quantum classifier labeling the sequence by nucleotide composition — **GC-rich vs AT-rich**.

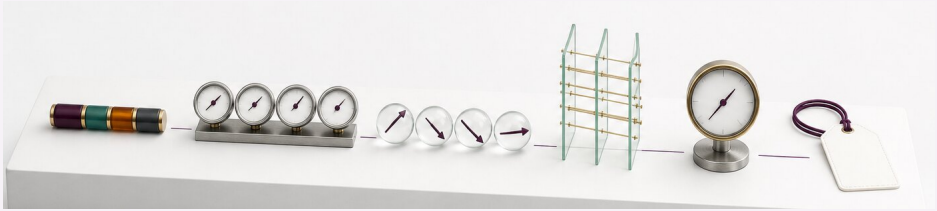


Trained by minimizing mean-squared error against $\pm 1$ targets with Adam; gradients from the parameter-shift rule. Held-out test split it never sees during optimization. Applied to `GCGCA`, it returns *GC-rich*, correctly.

Every piece has a neural-network counterpart. The identical recipe you would reach for classically, except the forward pass runs a quantum circuit.

# C3 — It does not beat a one-line GC check

It obviously does not. **No claim is made that it does.**

It is in this session for one reason: **it is governed**, on exactly the same terms as the optimization stage.

- Its **own descriptor** — output schema, provenance naming the trained-VQC method and linking back to assembly as its source.
- A **training-step budget** rather than an iteration budget. A learned model's distinctive cost is the training, so that is what the policy bounds — checked before the stage trains anything.
- The **same gateway.** The AI does not get a private door to the quantum resource.
- A **contract**: the trained model must clear a held-out accuracy floor before its output is accepted. A model that trained badly **fails its contract** rather than silently shipping a guess.

# C3 — What this stage tells the session

Two quantum stages — one that *solves* a circuit, one that *trains* one — under a single auditable control plane, with the same admission control, the same budgets, the same event record.

**The governance did not care which kind of stage it was.**

That is evidence for the regime-independence claim, at the smallest possible scale. It is not proof, and a room that finds a stage type the structure cannot govern has found something worth reporting.

Q3 asks which requirements are regime-independent. This is the cheapest available test of that question and it passed. Ask the room for a harder one.

# C4 — What upstream decides, downstream does not undo


**This is why C4's contract limits ammonia in the vessel.** What the cell does to fucosylation on day 9 is still in the product at release; standard purification does not change it [12].

Illustrative: one antibody, stylized, with glycans drawn as beads near the stem — in the cell, in the vessel, through the column, in the released vial.

**1 Set in the cell**

Cell line and culture conditions shape the glycan profile [10].

Temperature and osmolality move it measurably; pH and dissolved oxygen often by under 5% [42].

Ammonia stress lowered galactosylation, sialylation and fucosylation [43].

**2 Carried down the line**

The product is a mixture of forms, not one molecule: glycoforms and other isoforms, whose pattern the manufacturer must define [13].

Aggregates — dimers and higher — are product-related impurities [13].

**3 Removed downstream**

Host-cell protein, host DNA, viruses, leached Protein A, aggregates [11].

Two complementary viral clearance steps, e.g. low-pH inactivation and virus filtration [44].

**Not remade:** standard downstream processing does not change afucosylation [12].

**4 Tested at release**

Acceptance criteria for product- and process-related impurities; heterogeneity, glycoforms included, characterized [13].

Each batch tested against specification before release — a GMP decision [7].

Unit sequence on slide 6: cell banks, ICH Q5D [19]; seed train, fed-batch production and harvest, Kelley [20] and Li et al. [21]; downstream platform, Liu et al. [11].

# C4 — Where quantum is proposed, and against what

| Step | Classical tool today | Quantum proposal found, in its own words | Decision it feeds | Moderator's assessment |
|---|---|---|---|---|
| **1** Molecule design | electronic structure | Error-corrected phase estimation on a drug–protein active site: about 2 days at 32 orbitals, years at 100 [14] | discovery | **FT** the one case worth specifying for |
| **1** Antibody structure | classical Monte Carlo | Loop modelling on a fault-tolerant machine: slower than classical even in the optimistic scenario [15]; 7-residue lattice peptide on 9 qubits [45] | discovery | no defensible rationale shown |
| Glycan, degradation chemistry | — | None found. Nearest: an enzyme reaction on hardware, by its authors within brute-force classical reach [46] | development | unproposed |
| **2** Media and feed | design of experiments [21], [47] | A perspective asserting hybrid methods help, with no results [16] | development | no defensible rationale shown |
| **3** Planning, scheduling | MILP [48] | Job shop on an annealer, beaten by one classical core [17]; industrial scheduling on a hybrid solver whose quantum share cannot be isolated [49] | operations | no defensible rationale shown |
| In-line spectra (PAT) | Raman with PLS [50], [51] | None found for PAT. Across 160 non-PAT datasets, classical models beat quantum classifiers [52] | GMP, in-process | no defensible rationale shown |
| **4** Chromatography | mechanistic models, e.g. steric mass action [53] | Column packing as maximum independent set, 18 nodes, QAOA on 20 qubits [18]. PDEs (heat equation): at most quadratic speedup [54]; readout caveats [55] | development | no defensible rationale shown |
| Release decision | specification testing [13] | None found | **GMP release** | not a compute problem: where the evidence lands |

Numbers match the markers on slide 6. One row survives as a quantum target, and it sits before the line, outside GLP [4]. The release decision, where P1's evidence is needed, has no proposal at all.

# C4 — A regulated decision, from the other side of the wall

**This one is not quantum. That is the point of it.**

A mechanistic digital twin of a **fed-batch CHO bioreactor** — the workhorse of biologics manufacturing. First-principles ODEs: Monod growth with by-product inhibition, glycolytic overflow to lactate, the glucose-limited lactate re-consumption shift, glutaminolysis to ammonia, mAb formation, and a stirred-tank mass balance with feed boluses.

Around it, a **contract of in-process limits** — lactate, ammonia, vessel volume. They are there for the product: ammonia shifts the antibody's glycosylation [56], and glycosylation is a critical quality attribute [57], [58]. An agent must produce an **attributed decision artifact** recording what it chose and why. **Attributed, not yet signed.**

Parameters are illustrative and order-of-magnitude realistic, not fit to a clone. That is stated in the artifact, not discovered by a reviewer.

Monod growth kinetics [59]. CQA as defined in ICH Q11 [57]. Ammonia and CHO glycosylation: Yang and Butler [56]; glycosylation as a critical quality attribute: Hossler et al. [58]. NASEM's definition of a digital twin [60] requires updates from the physical system; this model has none, so strictly it is a mechanistic process model.

# C4 — The contract binds on over-feeding, not on starving

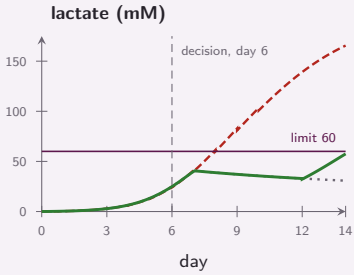


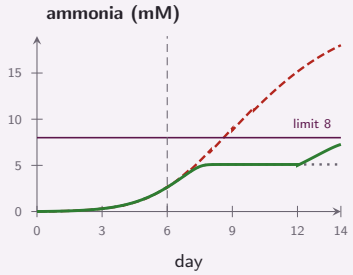


**Naive daily over-feed**
lactate 165.6, ammonia 18.0
**outside the contract**

**No feed**
glucose runs out; titer 1607 mg/L
**inside the contract**

**Chosen at day 6**
one 20 mL bolus on day 12, best of 256 plans — the same run as no feed until then; titer 1642 mg/L
**inside the contract**

The contract refuses the over-feed. It has nothing to say about starving: that costs titer, not compliance — here only 2%. The agent **discovers** where the limits bind by running what-ifs against the twin, then writes down what it chose and why. **The chemistry stays in the equations. Nothing invents a number.**

All three trajectories re-simulated from the twin with the decision artifact's numbers asserted first; `figures/gen_twin.py`. Two implementations run over the same twin — a deterministic search controller and a tool-use loop — bound by the same contract.

# C4 — The decision artifact, from an actual run

**Observed state, day 6**

| | |
|---|---|
| VCD | 18.7 e6 cells/mL |
| glucose | 12.2 mM |
| glutamine | 2.9 mM |
| lactate | 24.9 mM |
| ammonia | 2.63 mM |
| titer | 298.6 mg/L |

**Recommended**

Day 12: 0.02 L bolus, 1250 mM glucose, 150 mM glutamine. Selected from **256 feed strategies** searched over the twin from day 6.

**Predicted vs. contract**

| | pred | limit | |
|---|---|---|---|
| lactate | 57.5 | 60.0 | **ok** |
| ammonia | 7.27 | 8.0 | **ok** |
| volume | 1.02 | 1.6 | **ok** |

final titer **1642.2 mg/L**, peak VCD 24.1, 14 days

**contract_satisfied: true** violations: 0

Also recorded: objective, rationale, `decided_by`, `twin_version`, and **5 alternatives considered** with their titers and constraint margins.

# C4 — The question, with its abstraction removed

That is the kind of record a computational step in a regulated bioprocess must leave **before anyone may act on it**.

**What would a quantum stage produce to clear this bar?**

**And can it?**

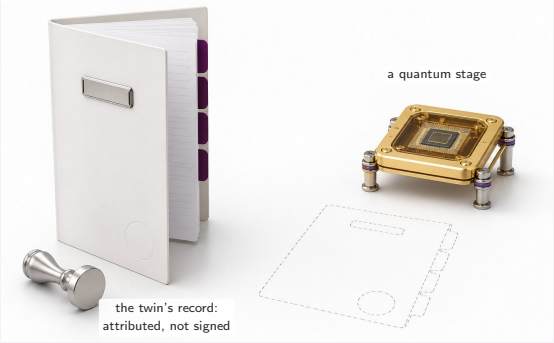


the record it would need: not built

Not "is quantum useful." Not "will quantum scale." This: the state it observed, the alternatives it rejected, the contract it satisfied, the version of the model it used, and a signature linked to that record [5] — which this twin does not produce yet either.

# C5 — The strongest case for P2

**The strongest version of P2**, argued in earnest:

- Every bio-quantum demonstration in the literature, including all three of mine, runs at a scale a laptop settles exactly. **Including mine.**
- Governance work is cheap to produce and impossible to falsify, which is exactly the profile of work that accumulates when a field cannot make progress on its actual bottleneck.
- Specifying an interface before the thing behind it exists has a poor historical record. Interfaces designed ahead of their implementations usually encode the wrong abstraction and then have to be unlearned.
- Fault-tolerant machines will have *different* failure modes. Calibration epochs and shot statistics may simply not be the things that need recording, in which case this checklist is an artifact of NISQ that will be thrown away.

**That is a serious argument and I do not have a knockdown answer to it.**

# C5 — What P2 gets right, conceded without qualification

- **There is no biological quantum advantage today.** I am not claiming one and my artifacts do not show one.
- **My classical baseline wins outright** in the one case where I actually computed it — $H_2$ in four dimensions, diagonalized exactly, in a browser tab.
- **Two of my four artifacts are in categories P3 says should stop**, and P3's argument about genomics-as-QUBO is the strongest technical objection in the room.

None of these concessions costs me anything I was entitled to. Making them early is how the remaining disagreement becomes findable.

# C5 — Where P1 survives P2, if it survives

P1 does not depend on advantage existing. It depends on **one empirical claim**:

**The evidence interface is being fixed now, by default, by whoever ships first — and it will be inherited rather than redesigned.**

If that claim is false, P2 wins and this session was premature. If it is true, waiting for fault tolerance is how biology inherits an uninspectable interface for the second time.

**So the question that decides between P1 and P2 is not a quantum question. It is:** do platform interfaces, once shipped, get redesigned?

The room contains people who have shipped those interfaces. They are better placed to answer this than I am.

# C5 — Seven open questions

Each is designed to produce disagreement, not consensus.

1. Name a quantum platform that today emits device state, full configuration and signed lineage. If none, is that a product gap or a research gap?
2. If bit-exact reproducibility is impossible, what *is* the acceptance criterion — and who gets to set it, the vendor or the regulator?
3. Is a calibration epoch identifier granular enough? What drifts *within* an epoch, and does anyone measure it?
4. Must a quantum stage prove it beats a classical baseline to be admitted — or is "runs at all" enough for an experimental step?
5. A quantum result comes back as a distribution to a caller that expected a number. Who is responsible for the collapse — the platform, the algorithm, or the biologist?
6. Which is more likely to move first — quantum platforms learning to emit regulated evidence, or regulated science relaxing what it requires?
7. Is there anyone here who believes a quantum step will be inside a *validated* pipeline before 2035? What would have to be true?

# Sources for C1 and C4 (1 of 2)


[34] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, "Quantum circuit learning," *Phys. Rev. A*, vol. 98, art. 032309, 2018, doi:10.1103/PhysRevA.98.032309.

[35] M. Schuld, V. Bergholm, C. Gogolin, J. Izaac, and N. Killoran, "Evaluating analytic gradients on quantum hardware," *Phys. Rev. A*, vol. 99, art. 032331, 2019, doi:10.1103/PhysRevA.99.032331.

[36] A. Gu, A. Lowe, P. A. Dub, P. J. Coles, and A. Arrasmith, "Adaptive shot allocation for fast convergence in variational quantum algorithms," arXiv:2108.10434, 2021.

[37] N. C. Rubin, R. Babbush, and J. McClean, "Application of fermionic marginal constraints to hybrid quantum algorithms," *New J. Phys.*, vol. 20, art. 053020, 2018, Eq. (49), doi:10.1088/1367-2630/aab919.

[38] A. Kandala, A. Mezzacapo, K. Temme, et al., "Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets," *Nature*, vol. 549, pp. 242–246, 2017, doi:10.1038/nature23879.

[39] D. Wecker, M. B. Hastings, and M. Troyer, "Progress towards practical quantum variational algorithms," *Phys. Rev. A*, vol. 92, art. 042303, 2015, doi:10.1103/PhysRevA.92.042303.

[40] M. Tejedor, J. Conejero, and R. M. Badia, "A semantic quantum circuit cache for scalable and distributed quantum-classical workflows," arXiv:2604.26788, 2026.

[41] J. Zhang, H. Wang, G. S. Ravi, et al., "DISQ: Dynamic iteration skipping for variational quantum algorithms," in *Proc. IEEE QCE*, 2023, pp. 1062–1073, doi:10.1109/QCE57702.2023.00120.

[42] J. Bravo-Venegas, J. Rodriguez-Siza, M. Vergara, M. Torres, A. Dickson, J. R. Toledo, M. C. Molina, et al., "Impact of process parameters on IgG glycosylation in CHO systems: A comprehensive quantitative analysis," *mAbs*, vol. 18, no. 1, art. 2643039, 2026, doi:10.1080/19420862.2026.2643039.

[43] B. F. Synoground, C. E. McGraw, K. S. Elliott, C. Leuze, J. R. Roth, S. W. Harcum, and N. R. Sandoval, "Transient ammonia stress on Chinese hamster ovary (CHO) cells yield alterations to alanine metabolism and IgG glycosylation profiles," *Biotechnol. J.*, vol. 16, no. 7, art. 2100098, 2021, doi:10.1002/biot.202100098.

[44] ICH, "Viral safety evaluation of biotechnology products derived from cell lines of human or animal origin Q5A(R2)," Step 4, 1 Nov. 2023, §6.3.

[45] A. Robert, P. Kl. Barkoutsos, S. Woerner, and I. Tavernelli, "Resource-efficient quantum algorithm for protein folding," *npj Quantum Inf.*, vol. 7, art. 38, 2021, doi:10.1038/s41534-021-00368-4.

[46] P. Ettenhuber, M. B. Hansen, P. P. Poier, et al., "Calculating the energy profile of an enzymatic reaction on a quantum computer," *J. Chem. Theory Comput.*, vol. 21, no. 7, pp. 3493–3503, 2025, doi:10.1021/acs.jctc.5c00022.

[47] Y. Rouiller, A. Périlleux, N. Collet, M. Jordan, M. Stettler, and H. Broly, "A high-throughput media design approach for high performance mammalian fed-batch cultures," *mAbs*, vol. 5, no. 3, pp. 501–511, 2013, doi:10.4161/mabs.23942.

# Sources for C1 and C4 (2 of 2)


[48] K. Lakhdar, Y. Zhou, J. Savery, N. J. Titchener-Hooker, and L. G. Papageorgiou, "Medium term planning of biopharmaceutical manufacture using mathematical programming," *Biotechnol. Prog.*, vol. 21, no. 5, pp. 1478–1489, 2005, doi:10.1021/bp0501571.

[49] A. Awasthi, N. Kraus, F. Krellner, and D. Zambrano, "Real world application of quantum-classical optimization for production scheduling," in *Proc. IEEE QCE*, 2024, pp. 239–244, doi:10.1109/QCE60285.2024.10285.

[50] N. R. Abu-Absi, B. M. Kenty, M. E. Cuellar, M. C. Borys, S. Sakhamuri, D. J. Strachan, M. C. Hausladen, and Z. J. Li, "Real time monitoring of multiple parameters in mammalian cell culture bioreactors using an in-line Raman spectroscopy probe," *Biotechnol. Bioeng.*, vol. 108, no. 5, pp. 1215–1221, 2011, doi:10.1002/bit.23023.

[51] J. Whelan, S. Craven, and B. Glennon, "In situ Raman spectroscopy for simultaneous monitoring of multiple process parameters in mammalian cell culture bioreactors," *Biotechnol. Prog.*, vol. 28, no. 5, pp. 1355–1362, 2012, doi:10.1002/btpr.1590.

[52] J. Bowles, S. Ahmed, and M. Schuld, "Better than classical? The subtle art of benchmarking quantum machine learning models," arXiv:2403.07059, 2024.

[53] C. A. Brooks and S. M. Cramer, "Steric mass-action ion exchange: Displacement profiles and induced salt gradients," *AIChE J.*, vol. 38, no. 12, pp. 1969–1978, 1992, doi:10.1002/aic.690381212.

[54] N. Linden, A. Montanaro, and C. Shao, "Quantum vs. classical algorithms for solving the heat equation," *Commun. Math. Phys.*, vol. 395, no. 2, pp. 601–641, 2022, doi:10.1007/s00220-022-04442-6.

[55] S. Aaronson, "Read the fine print," *Nature Phys.*, vol. 11, no. 4, pp. 291–293, 2015, doi:10.1038/nphys3272.

[56] M. Yang and M. Butler, "Effects of ammonia on CHO cell growth, erythropoietin production, and glycosylation," *Biotechnol. Bioeng.*, vol. 68, no. 4, pp. 370–380, 2000, PMID 10745205.

[57] ICH, "Development and manufacture of drug substances (chemical entities and biotechnological/biological entities) Q11," Step 4, 1 May 2012, Glossary.

[58] P. Hossler, S. F. Khattak, and Z. J. Li, "Optimal and consistent protein glycosylation in mammalian cell culture," *Glycobiology*, vol. 19, no. 9, pp. 936–949, 2009, doi:10.1093/glycob/cwp079.

[59] J. Monod, "The growth of bacterial cultures," *Annu. Rev. Microbiol.*, vol. 3, pp. 371–394, 1949, doi:10.1146/annurev.mi.03.100149.002103.

[60] National Academies of Sciences, Engineering, and Medicine, *Foundational Research Gaps and Future Directions for Digital Twins*. National Academies Press, 2024, doi:10.17226/26894.

# Thank you.

**the app, live and client-side**

qubit.agenticdatasets.org

**session page**

qubit.agenticdatasets.org/bof

**the moderator**

linkedin.com/in/thedoytsujin

**Alexander Chernov | BOF-2099 | IEEE Quantum Week 2026**