A Few Research Projects I Have Worked On

Some work in experiment, fault tolerance, control, and foundations

Jianshuo Gao

Report for Prof. Zhendong Zhang's group
The University of Hong Kong

2026

Experimental AMO • Fault-tolerant quantum computing • Quantum control • Foundations

A quick timeline

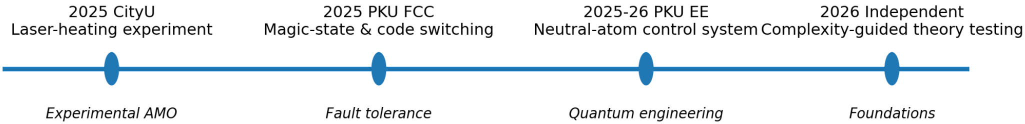

- I am usually interested in problems where **hardware**, **information**, and **verification** show up together.
- These projects look different on the surface, but for me they are connected by the same habit: build something, check what really limits it, and then try to say clearly what was learned.

What connects these four projects?

Prepare

How can we prepare useful quantum resources when the hardware and noise are both non-ideal?

Certify

How do we tell whether a claimed quantum resource is really better than the ingredients used to build it?

Control

How do we turn abstract pulse sequences into deterministic lab timing?

Choose

When is a more complicated explanation really worth it?

Main point

For me the common pattern is simple: first make the system controllable, then find the main limitation, and then decide how to judge success quantitatively.

Roadmap

Project I

Laser heating of atoms at City University of Hong Kong (2025)

Project I: question, setup, and what I did

- **Scientific question:** can I quantify heating induced by a modulated laser beam through changes in rubidium absorption linewidths?
- **Experimental goal:** build the optical path, design the AOM modulation scheme, acquire spectra, and fit linewidth–temperature curves.
- **What I did personally:** this was a fairly complete experimental project for me: I set up the optics, took the data, and did the fitting myself.
- **Why it matters:** spectroscopy turns a qualitative “the atoms are hotter” statement into a quantitative thermal diagnostic.

Outputs

- optical layout
- AOM scan scheme
- spectral dataset
- linewidth model

Internal CityU project data and notebook, 2025.

Spectroscopy as thermometry

For an atomic transition of optical frequency ν_0, thermal motion gives Doppler broadening

$$\Delta\nu_D = \nu_0\sqrt{\frac{8\ln 2\, k_B T}{mc^2}}.$$

- T enters through the velocity distribution; broader velocity spread means broader absorption lines.
- In practice, the measured linewidth is not purely Doppler-limited.
- Collisions, power broadening, and instrumental effects can all contribute.

Main point

The core logic of the experiment is simple: heating changes the line shape, and the line shape gives back quantitative information about temperature and broadening mechanisms.

Why a Voigt profile is the natural physical model

A useful decomposition is

$$\text{Gaussian part} \;\leftrightarrow\; \text{Doppler broadening}$$

$$\text{Lorentzian part} \;\leftrightarrow\; \text{collisional / homogeneous broadening.}$$

For the full-width at half maximum (FWHM), a good approximation is

$$\omega_V(T) \approx 0.5346\,\Gamma(T) + \sqrt{0.2166\,\Gamma(T)^2 + \Delta\nu_D(T)^2}.$$

- $\Delta\nu_D(T)$ grows only as $\sqrt{T}$.
- $\Gamma(T)$ can grow much faster because vapor density and collisional rate increase rapidly with temperature.
- This gives a natural explanation for an apparent crossover from weak to strong linewidth growth.

Olivero–Longbothum-type approximation used in my notebook-based analysis.

Optical setup used in the experiment

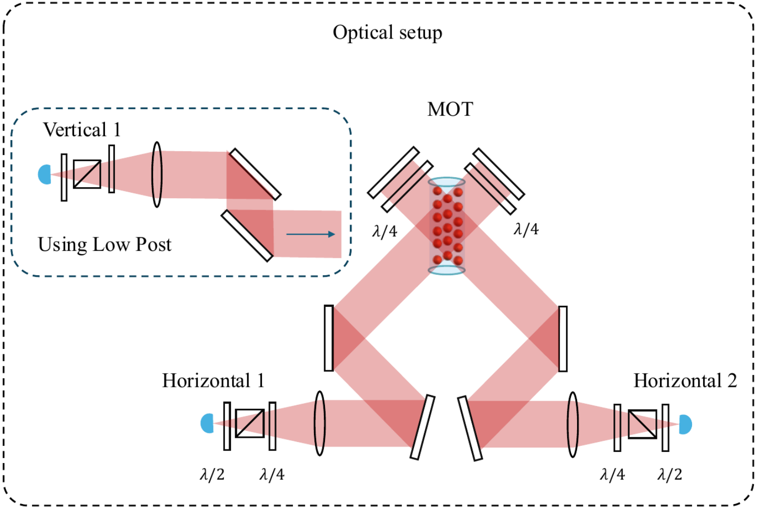

PBS:
Waveplate:
Mirror:
Lens:
Flip mirror
frameless mirror

Frequency scan and heating actuation

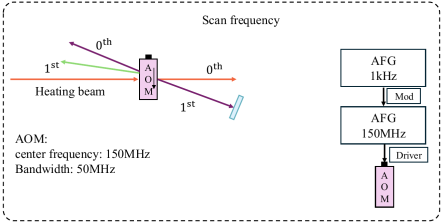

- AOM center frequency: about 150 MHz.
- Reported bandwidth in my project note: about 50 MHz.
- A lower-frequency function generator modulates the RF drive and acts as a heating control knob.
- Frequency and amplitude therefore become experimentally tunable heating parameters.

Main point

AOM modulation is not only a control tool; here it also becomes the physical mechanism that drives heating and spectral broadening.

Data workflow I actually used

1. Sweep the spectroscopy coordinate (piezo / scan voltage) across several rubidium features.
2. Record transmission traces for room temperature and heated conditions.
3. Fit local windows of the dips to extract FWHM for each feature.
4. Repeat across different heating settings or controlled cell temperatures.
5. Compare empirical trends with a physics-motivated broadening model.

Why this matters

This part was good training for me because the lab work, fitting, and physical interpretation all had to be done together.

Representative room-temperature spectrum

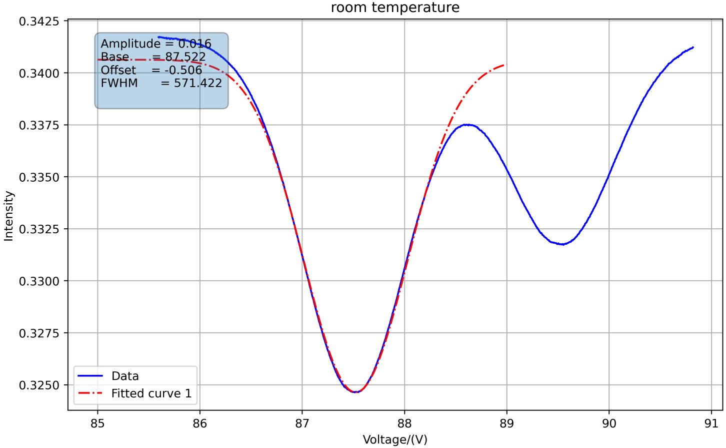

- The main dip is well fit by a local model.
- The fitted FWHM on this trace is about 571 MHz.
- A second dip is visible at larger scan voltage, so local fitting windows are essential.

Tracking multiple spectral features

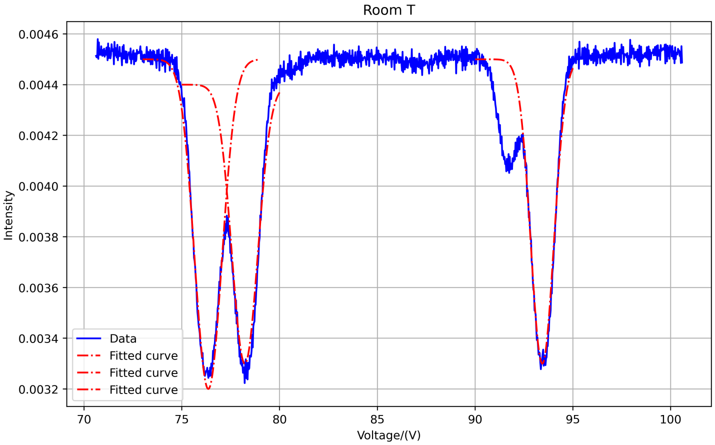

- I tracked **three distinct features** instead of collapsing the spectrum to a single number.
- This is useful because different features can respond differently to heating, overlap, or fitting window bias.
- The later temperature sweep kept these three features separate throughout the analysis.

Heating conditions visibly reshape the spectra

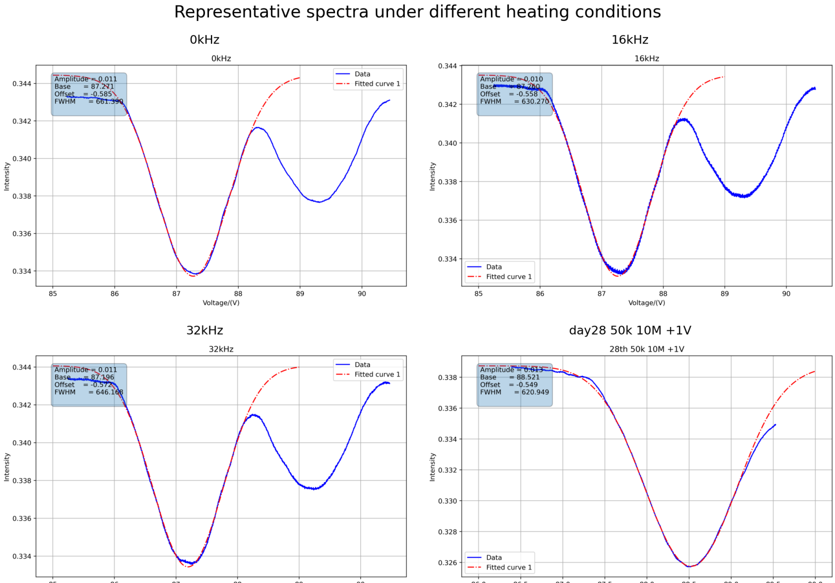

A direct comparison: room temperature versus heated cell

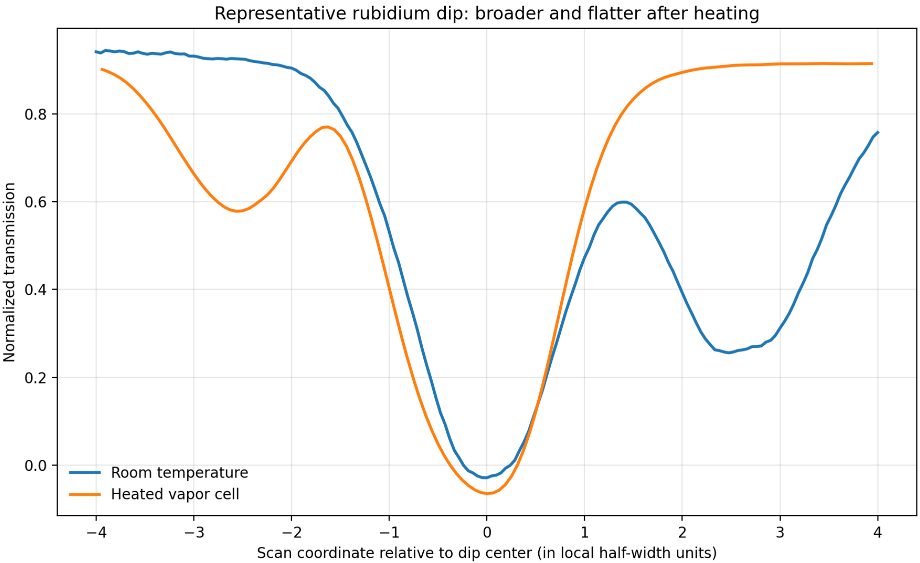

Temperature sweep: all three linewidths grow strongly

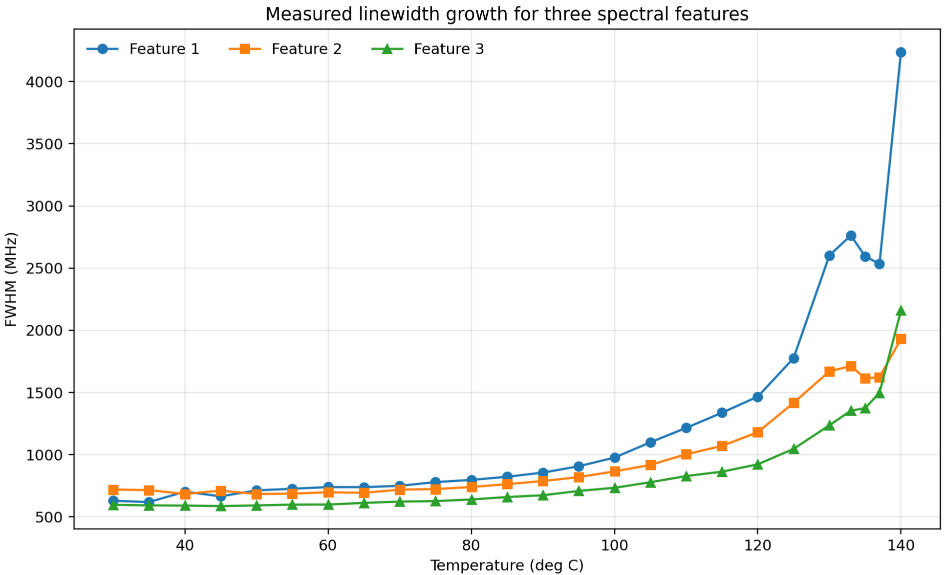

Empirical fit versus a physical model

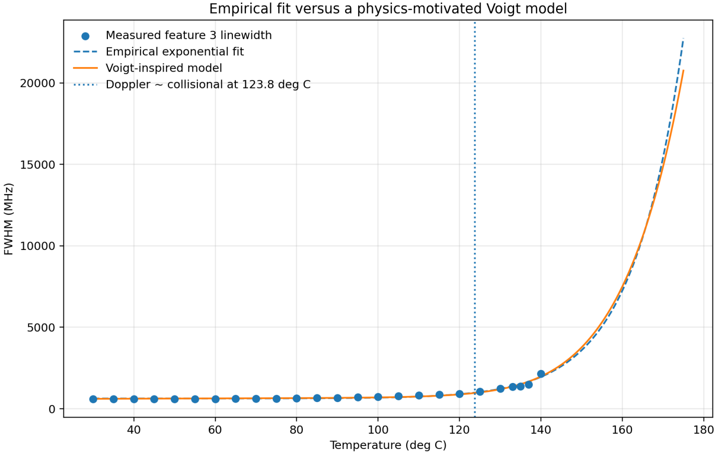

- An empirical fit of the form

$$\omega(T) = \omega_0 + ab^T$$

  captures the measured trend over the observed range.
- But a purely empirical fit says *what* the curve does, not *why* it does it.
- That motivated a Voigt-inspired interpretation.

Physical picture: a crossover of broadening mechanisms

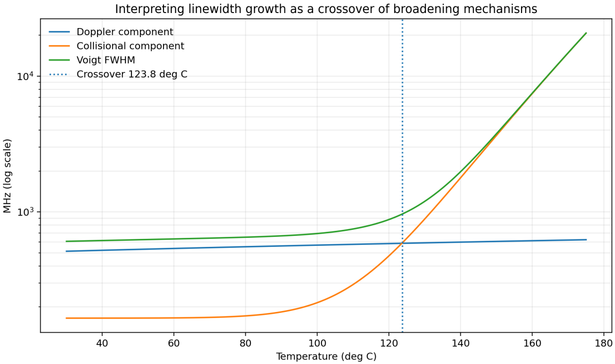

- I modeled the homogeneous component as

$$\Gamma(T) = \Gamma_0 + \beta e^{A - B/T_K}.$$

- In the fitted model, Doppler and collisional contributions become comparable near 124 °C.
- Above that regime, the collisional contribution dominates the linewidth growth.

Main point

What I took from this part is that an empirical fit is only the start; the more important thing is to say what mechanism is plausible and where the model is still only suggestive.

AOM diffraction geometry and order identification

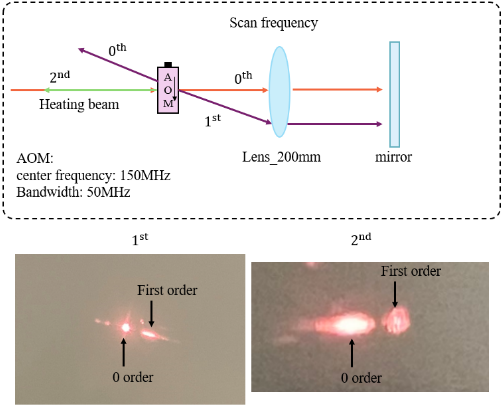

- I used the AOM as both a **frequency shifter** and a **heating actuator**.
- The top panel shows the scan geometry: a 150 MHz center frequency and about 50 MHz usable bandwidth.
- The lower images identify the **0th** and **1st** diffraction orders in practice, which was important for alignment and power delivery.
- This check confirmed that the intended beam path was experimentally realized before I interpreted any spectroscopy data.

Heating reshaped the absorption peak rather than increasing atom number

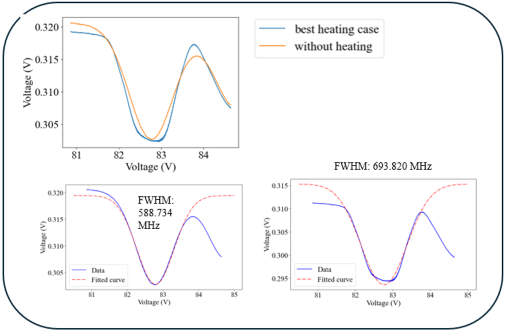

- After turning on the heating beam, the measured background became higher, consistent with extra scattered light or weak fluorescence.
- In the processed comparison (upper-left panel), the heated trace therefore has to be shifted upward before comparing line shapes.
- The key observation is that the **peak shape changes**: the dip becomes distorted and is no longer well described by a simple Gaussian.
- Our interpretation was that this method mainly modifies the **velocity distribution**, while the total absorbing atom number changes much less.

Why I stopped here for the moment

Reason

This means the technique may be better for studying **non-thermal velocity-distribution effects** than for the original motivation of directly changing atom number. After discussing it with Prof. Qu, I decided to pause the project and leave it for future continuation if a more suitable question appears.

What I learned from project I

Experiment

This project gave me practice with the whole chain: optics, scan design, data acquisition, and checking the raw traces.

Analysis

It also made me much more comfortable extracting quantitative information from imperfect spectra and comparing fitting models.

Interpretation

I also learned not to stop at a descriptive plot if there is a reasonable physical interpretation to discuss.

Project II

Magic-state preparation, code switching, and gauge-fixing ideas at PKU CFCS

Why non-Clifford resources are the expensive part

- In most fault-tolerant architectures, Clifford gates such as H, S, and CNOT are relatively cheap.
- A standard non-Clifford resource is the T gate,

$$T = \text{diag}(1, e^{i\pi/4}).$$

- Clifford circuits with Pauli measurements and stabilizer-state preparation remain classically efficiently simulable (Gottesman–Knill).
- Therefore, large-scale universal quantum computation requires an additional non-Clifford resource — typically magic states.

Background summarized from my 2025 group report on magic states and code switching.

Magic states as the usual non-Clifford resource

A single-qubit magic state can be written as

$$|T\rangle = T|+\rangle = \frac{|0\rangle + e^{i\pi/4}|1\rangle}{\sqrt{2}}.$$

- Once high-fidelity $|T\rangle$ states are available, one can inject T gates into an otherwise Clifford circuit.
- This is why **magic-state preparation** is one of the main overhead bottlenecks in fault-tolerant quantum computing.
- My work in this area has two parts: understanding the protocol level in detail, and then asking how to compare or improve entire families of protocols.

How I got into this topic

Expository foundation

I wrote a short note comparing **surface codes** and **color codes**, and a longer group-report-style note on **magic states and code switching**. Writing those notes forced me to sort out the basic logic for myself instead of just remembering protocol names.

Research direction

Under Prof. Yuan Xiao, After that, I started asking a more specific question of my own: can the postselection idea used in code switching also help gauge fixing, which is a more general way to convert between codes?

My reports: *Surface Code vs Color Code* (2025) and *Magic States and Code Switching* (2025).

The stabilizer-code language I keep using

For a code with stabilizer generators $S_1, \ldots, S_m$, the projection onto the code space is

$$\Pi_{\mathcal{C}} = \frac{1}{2^m}(I + S_1)(I + S_2)\cdots(I + S_m).$$

From this one gets the standard counting relation

$$\mathrm{dim}(\mathcal{C}) = 2^{n-m}, \qquad k = n - m.$$

- This is the language I use to move between expository work, protocol analysis, and current research questions.
- It also makes it natural to think of code switching and gauge fixing as **changing the constraint set while preserving logical information**.

Surface codes versus color codes: why color codes matter here

	Surface code	Color code
Basic geometry	usually square-lattice-like	trivalent, face-colored lattice
Transversal Clifford gates	limited	richer Clifford structure
Transversal non-Clifford gate	typically no direct transversal T	3D color codes can support transversal T
Best role in this story	excellent main computational code	useful for *preparing* magic states by code switching

Main point

The attraction of code switching is exactly this: prepare the non-Clifford resource in a code that naturally supports it, then move the state into the code you actually want to compute in.

The key idea of code switching

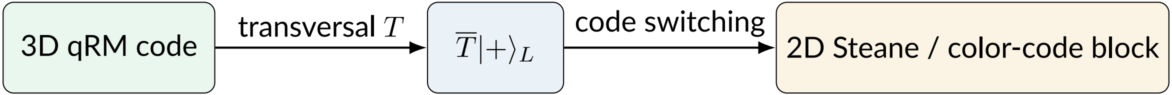

- Use a code with a transversal non-Clifford gate to create a logical magic state.
- Transfer the encoded state into a code better suited to the rest of the computation.

The codes in the protocol I studied most closely

Steane code

$$[[7,1,3]], \qquad \overline{Z} = Z_1 Z_2 Z_3, \quad \overline{X} = X_1 X_2 X_3.$$

A 2D color code / CSS code that is very convenient as a small logical destination block.

Quantum Reed–Muller (qRM) code

$$[[15,1,3]], \qquad d_Z = 3,\ d_X = 7.$$

This 3D color-code instance supports a transversal logical T gate, which is the crucial resource.

Daguerre & Kim, *Code switching revisited*, 2025.

A transversal logical T gate on qRM

The qRM code admits the transversal pattern

$$\overline{T} = T_1 T_2^\dagger T_3 T_4^\dagger \cdots T_{15},$$

with alternating T and $T^\dagger$ on odd and even qubits.

- This is powerful because the expensive non-Clifford operation is implemented *transversally* at the logical level.
- The later protocol then converts this encoded magic state into a 2D code block through a logical teleportation step.
- This is an elegant way to partly evade the usual magic-state-distillation overhead at small distance.

Logical flow of the code-switching protocol

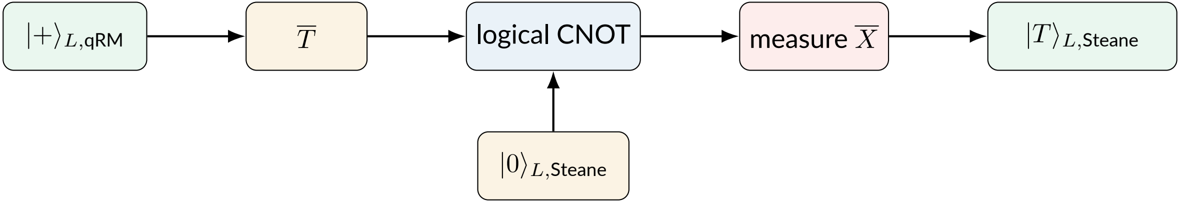

- The state transfer is essentially a one-bit teleportation protocol at the logical level.
- What makes it interesting is the fault-tolerant implementation underneath that logical cartoon.

A one-way transversal CNOT is the enabling bridge

With the qubit labeling used in the protocol, the logical CNOT is realized as

$$\mathrm{CNOT}_{\mathrm{qRM}\to\mathrm{Steane}} = \bigotimes_{i=1}^{7} \mathrm{CNOT}_i .$$

- This is a **one-way** transversal bridge: the control is in the qRM block and the target is in the Steane block.
- That asymmetry is not a nuisance; it is precisely what matches the state-transfer protocol.
- It is one reason code switching is naturally attractive on architectures with nonlocal connectivity, such as trapped ions and potentially neutral atoms.

Where the main danger comes from

- Both qRM and Steane are distance-3 codes in this low-overhead setting.
- A single two-qubit fault during syndrome extraction can spread into a weight-2 or higher error.
- Such hook errors can exceed the correctable set and become logical failures after the code-switching step.

Therefore

The protocol lives or dies by **how cleverly syndrome extraction is designed and filtered**.

Flag qubits plus postselection: the core protection idea

Accept only events in a benign set $\mathcal{A}$,

$$\mathcal{A} = \{(f, s) : f = 0,\ s \in \mathcal{S}_{\text{allowed}}\},$$

where f is the flag outcome and s is the measured syndrome.

$$P_{\text{fail}|\mathcal{A}} \ll P_{\text{fail}}.$$

- Flags detect dangerous fault propagation patterns.
- Postselection discards suspicious rounds.
- The remaining accepted rounds have much better conditional fidelity.

Main point

This is the main conceptual move that later motivated my gauge-fixing question: can the same “flag + reject” logic protect other code-conversion procedures?

Published benchmark numbers I used as a reference

Protocol variant	Logical infidelity	Acceptance	Comment
Code switching + EC	4.6×10^{-5}	84%	exact Monte Carlo estimate
Code switching + postselection	$\sim 5.1 \times 10^{-7}$	$\sim 84\%$	polynomial extrapolation
Estimated space-time volume		≈ 1335 qubit-steps	

- These numbers caught my attention because they already look quite competitive in a small-code, near-term setting.
- For me, this was a sign that code switching is not only a nice idea on paper, but something worth simulating carefully.

Daguerre & Kim, 2025; see also my own benchmark note on logical $|T\rangle$ preparation.

What I simulated after reading the code-switching paper

Part A: reproduction

- Reconstructed the qRM $\to$ Steane logical-$|T\rangle$ protocol in **Stim**.
- Sampled the pre-T qRM preparation with `FlipSimulator`.
- Implemented the logical T step using the paper's **Clifford-frame** trick.
- Evaluated the final accepted branch with exact stabilizer-projector postselection.

Part B: gauge-fixing extension

- Started from direct qRM $\to$ Steane gauge fixing.
- Replaced the generic qRM encoder by the paper's Fig. 5 $|+\rangle_L$ preparation plus logical-X verification.
- Added a new **qRM X-cell postselection** layer immediately before gauge fixing.

Main point

I can now reproduce the published protocol and test gauge-fixing variants in the same simulation framework.

Reproducing the published code-switching postselection curves

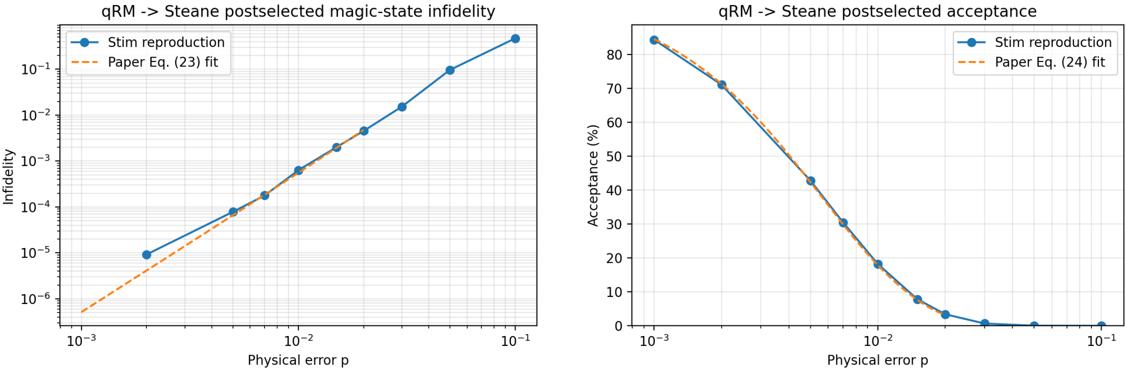

- In the low-p regime, my Stim reconstruction tracks the paper's Eq. (23) logical-infidelity fit and Eq. (24) acceptance fit.
- At $p = 10^{-3}$, the accepted branch stays near the **reported $\sim 84\%$ yield**, which was an important sanity check before I moved on to new protocol variants.

How I modified the gauge-fixing simulation

1. Use the paper's Fig. 5 qRM $|+\rangle_L$ preparation instead of a generic synthesized encoder.
2. Keep logical-X verification and the full optimized Z-syndrome / flag-postselection layer.
3. Before direct qRM $\to$ Steane gauge fixing, measure the **four qRM X-cell checks**.
4. Accept only the **trivial X-cell syndrome** branch, then continue the conversion.

Noiseless checks passed

- Fig. 5 preparation + postselection accepts deterministically.
- The upgraded conversion maps $|+\rangle, |-\rangle, |Y\pm\rangle$ into the expected Steane logical states.

Main point

The new ingredient is precise: I am not just saying "gauge fixing might benefit from postselection"; I inserted a concrete qRM-side filter and tested it.

Targeted comparison at $p = 10^{-3}$

Protocol variant	$1 - F$	P_{success}
Original direct gauge fix	6.49×10^{-3}	84.4%
Fig. 5 prep + logical-X verification + full Z/flag postselection	7.10×10^{-3}	87.3%
+ qRM X-cell postselection before gauge fixing	3.57×10^{-3}	83.5%

- The extra X-cell filter is the first clear win: about $1.8\times$ **to** $2\times$ **lower** accepted-branch infidelity at nearly the same yield.

Targeted $p = 10^{-3}$ comparison from my March 2026 gauge-fixing runs.

Uniform scan of the upgraded gauge-fixing protocol

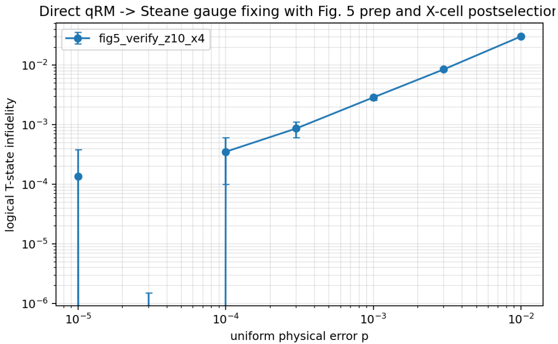

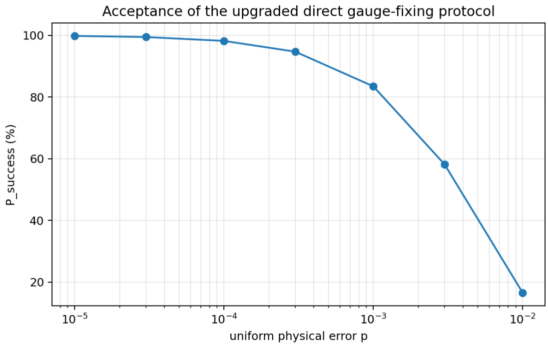

- The upgraded protocol continues to improve the accepted branch over a broad p range, while the success probability degrades in the expected postselection-like way as noise increases.
- A log-log fit over $p \geq 10^{-4}$ gives slope ≈ 0.97, so the present version is **better** but still not yet in the clean $O(p^2)$ regime that I ultimately want.

What I learned from section II

Reproduce

A lot of my time here went into reconstructing the published protocol carefully enough that I could trust my own simulation code.

Modify

Once the baseline looked reliable, I tried one concrete modification: qRM X-cell postselection before gauge fixing.

Judge honestly

I also try to be careful about whether a new idea really changes the scaling, not just whether one point happens to improve.

Project III

Neutral-atom quantum-computing control system development at PKU EE

The control-system family I am working on

Two deployment scales

- **WYS8U:** 19-inch 8U chassis, 12 functional slots.
- **WYS2U:** 2U agile platform for ADC / DAC / DDS / GPIO prototyping.
- Common control style: Python front-end, 4 ns timing resolution.

WYS8U / WYS2U user manuals, 2025.

RWG is the module most directly relevant to neutral-atom control

- **4-channel RF source** with parameterized real-time waveform generation.
- Up to **128 simultaneous tones** in a tunable 100 MHz band (0.1–400 MHz class).
- Frequency / amplitude ramping, nonlinear modulation, and amplitude compensation.
- 4 mark outputs, upgradable to GPIO.
- Representative RWG-loop delay in the WYS8U test manual: about 560 ns.

Main point

For neutral-atom work, this is a pretty useful combination: fast RF output, synchronization, and room for feedback.

RWG module docs and WYS8U test manual, 2025.

What I am doing in this system now

- Turning high-level pulse descriptions into **deterministic RTMQ schedules** that know about carriers, ramps, timing dependencies, and trigger relations.
- Thinking about how a neutral-atom experiment should be decomposed across **master**, **RWG**, and **auxiliary analog / digital modules** instead of being hard-coded instrument by instrument.
- Keeping the interface **calibration-friendly**: parameter scans, feedback hooks, and later system growth should be natural rather than patched on.

Why this project matters to me

For me this is the point where abstract ideas meet real lab constraints: clocking, latency, routing, waveform generation, and hardware abstraction all have to work together.

Project IV

Why simpler theories are more likely to be correct — but only quantitatively

Why I started thinking about simplicity more carefully

- Physicists often say that simpler theories are preferable.
- But that slogan hides at least three different questions:
  1. What do we mean by the “length” or “complexity” of a theory?
  2. How much evidence is needed to justify a more complicated model?
  3. If experiments are expensive, which measurements repay the complexity penalty fastest?
- In this side project, I try to put these three questions into one operational framework.

Predictive description length instead of a vague idea of elegance

Let E be a fixed experimental language and let a theory T be represented by a computable predictor on that language.

$$K_{\text{Pred}}(T) = \min\{|p| : \ p \text{ computes the theory's predictions on } E\}.$$

Because this ideal quantity is not computable in general, I also introduced an operational upper bound: a **complexity certificate** of length $L_{\text{cert}}(T)$.

$$K_{\text{Pred}}(T) \leq L_{\text{cert}}(T) + O(1).$$

- The point is not to talk about elegance in a vague way.
- I want to attach a concrete and checkable complexity penalty to a predictive theory.

Posterior odds = Occam factor × Bayes factor

Assign prior weights proportional to 2^{-L_i} for theories T_i with description lengths L_i. Then

$$\frac{P(T_i \mid x)}{P(T_j \mid x)} = 2^{-(L_i - L_j)}\,\frac{\nu_i(x)}{\nu_j(x)}.$$

- The first factor is the **Occam penalty**.
- The second factor is the usual **likelihood reward** from data.
- A more complicated theory may still win — but it must pay back its extra description length with enough information gain.

Main point

In this picture, “simple theories are more likely to be correct” is not a slogan. It just means more complicated theories start with a penalty, and strong enough data can overcome it.

When should a more complicated theory eventually win?

Using an MDL / BIC-style score,

$$L_{\mathsf{BIC}}(M; x_{1:n}) = -\log_2 \hat{L}_M(x_{1:n}) + \frac{d_M}{2} \log_2 n,$$

I show the following logic:

- if the true data-generating distribution is KL-separated from a simpler competitor class by at least δ,
- then the cumulative information gain grows like $n\delta$,
- while the complexity penalty grows only like $\frac{d}{2} \log_2 n$.

Therefore a sufficiently informative complex model must eventually win.

Bell witnesses make the information gap computable

In a CHSH scenario, let $\omega(P)$ be the Bell-game win probability. Then for any local model Q,

$$D_{\mathsf{KL}}(P\|Q) \geq D_{\mathsf{KL}}(\mathsf{Bern}(\omega(P))\|\mathsf{Bern}(\omega(Q))) \geq D_{\mathsf{KL}}(\mathsf{Bern}(\omega(P))\|\mathsf{Bern}(3/4)) .$$

Also,

$$\omega(P) = \frac{1}{2} + \frac{S(P)}{8},$$

where $S(P)$ is the CHSH value for the appropriate sign convention.

- A Bell witness therefore certifies a strictly positive lower bound on how much worse every local model must be.

Witness-certified crossover scale in CHSH tests

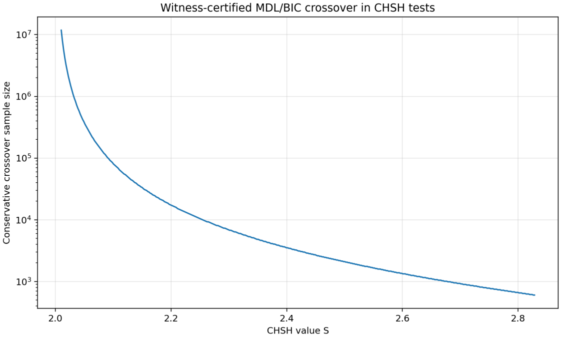

Define the witness-certified information rate

$$\delta_{\mathsf{CHSH}}(\omega) = D_{\mathsf{KL}}(\mathsf{Bern}(\omega) \| \mathsf{Bern}(3/4)) \,/\, \log 2.$$

A conservative MDL / BIC crossover estimate then solves

$$n\,\delta_{\mathsf{CHSH}}(\omega) \gtrsim \frac{d_L - d_Q}{2} \log_2 n.$$

- Larger Bell violations repay the complexity penalty faster.

Empirical case study: CHSH data and BIC comparison

From the four-photon dataset reanalyzed in my note, I obtained

$$\hat{S} = 2.274548, \qquad \sigma_{\hat{S}} \approx 0.0567.$$

Then I compared:

- a compact two-parameter cosine quantum model, and
- the best-fit local hidden-variable model in the CHSH local polytope.

Result:

$$L_{\mathsf{BIC}}(\text{local}) - L_{\mathsf{BIC}}(\text{quantum}) \approx 21.1 \text{ bits}.$$

Interpretation

The local model can achieve a slightly better raw likelihood because it is more flexible, but the compact quantum model wins once model complexity is charged.

My 2026 note using Wang et al.'s CHSH dataset.

Experiment design should maximize information per unit cost

If an experiment menu is available, the right question is not only "which theory wins?" but also "which next experiment is cheapest per bit of evidence?"

$$\text{Choose } e^\star \in \arg\max_e \frac{\delta(e)}{c(e)}.$$

In Bell scenarios, if $\pi(x, y)$ is the setting distribution, this becomes

$$\max_\pi \frac{\delta(\pi)}{\bar{c}(\pi)}.$$

- This is where the theory becomes experimentally actionable.
- Complexity penalties, KL information gain, and actual resource costs are all in the same equation.

A simple Bell-design lesson from the paper

- Suppose one CHSH setting is much more expensive than the others.
- Uniform sampling is no longer optimal if the objective is evidence per resource.
- In the example in my note, a cost-aware schedule improved the expected evidence-per-cost by about a factor of 2.5 relative to uniform sampling.

Why I like this result

It turns "Occam's razor" from a philosophical slogan into an **experiment-design principle**: spend your next unit of laboratory cost where it buys the most discriminating information.

Why this matters for logical magic-state certification

Recent logical magic-state experiments certify a target state using both single-copy logical measurements and a two-copy logical Bell-basis witness. In the high-fidelity regime with $\eta = 1 - F \ll 1$,

$$\sigma_{\hat{\epsilon}} \sim \sqrt{\eta/N},$$

so resolving infidelity η requires only

$$N_{\text{2-copy}} = O(\eta^{-1})$$

accepted shots, whereas direct single-copy estimation costs

$$N_{\text{tomo}} = O(\eta^{-2}).$$

- This is exactly the same information-per-cost principle, now in a fault-tolerant setting.

Certification cost in space–time volume

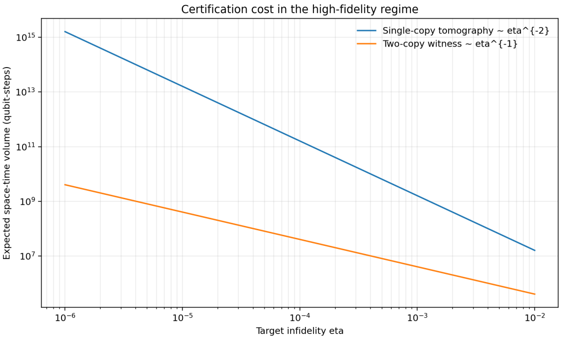

Using representative acceptance rates and a code-switching preparation cost V, one gets

$$V_{\text{tomo}}(\eta) \sim \frac{V}{p_1}\eta^{-2}, \qquad V_{\text{2copy}}(\eta) \sim \frac{2V}{p_2}\eta^{-1}.$$

- In the high-fidelity regime, the two-copy witness becomes dramatically cheaper.
- The asymptotic advantage grows as the target infidelity decreases.

Scaling analysis from my 2026 note; representative V, p_1, p_2 from recent code-switching work.

A recent side project: deriving Galilean one-particle kinematics

Question

Can Galilean kinematics be recovered from a **smooth family of reference states** instead of being assumed from the start?

Current result

$$\mathcal{H} \simeq L^2(\mathbb{R}^3) \otimes \mathbb{C}^{2s+1}$$

$$K_i = mX_i, \qquad H = \frac{P^2}{2m} + E_0$$

$$J = X \times P + S$$

- I extend Giannelli–Chiribella's observable–generator duality from a single reference state to a local manifold with **time, translation, rotation, and boost directions**.
- Then I combine that with covariant localization and a central boost–translation holonomy condition.
- So my newest foundations project is trying to recover standard Galilean one-particle kinematics from operational / information-theoretic ingredients.

Current draft: *Galilean One-Particle Kinematics from a Smooth Family of Reference States* (2026).

What I actually want to say here

Simplicity

A shorter predictive description should get a prior advantage.

Evidence

A more complicated theory can still win if the data bring enough KL information to pay back that complexity penalty.

Design

Experiments should be chosen to maximize information gain per unit cost, not just per shot.

Main point

So for me, “simpler theories are more likely to be correct” only makes sense as a quantitative statement that can in principle be tested.

Why I think project IV still fits with the others

- Project I asked how to extract a reliable thermal signal from an optical experiment.
- Project II asked how to prepare and compare high-value logical resources under realistic assumptions.
- Project III is about building the controller that makes such experiments real.
- Project IV asks a meta-question: **when do the data justify the extra complexity of the model or protocol we are using?**

My answer

I do not really separate foundations from engineering. If we want trustworthy protocols, we need good resource metrics, and that usually means being clear about complexity and evidence.

Conclusion

What I think these projects add up to

Very short summary

Project	Main technical content	What I took from it
Laser heating of atoms	optics, AOM control, spectroscopy, linewidth fitting, mechanism-level interpretation	I can run experiments and analyze the data physically
Magic states and code switching	stabilizer / color-code theory, low-overhead magic-state preparation, fair benchmarking, gauge-fixing ideas	I can connect expository mastery to current research questions
Neutral-atom control system	hardware abstraction, deterministic timing, multi-instrument synchronization	I care about how theory becomes executable in the lab
Complexity-guided theory testing	MDL / BIC, KL bounds, Bell witnesses, cost-aware certification	I like building quantitative frameworks for trustworthy scientific claims

What I hope I can contribute in future collaborations

- I am most comfortable on projects that move back and forth between **experiment**, **control**, **fault tolerance**, and **conceptual methodology**.
- Strong interest in architectures where nonlocal connectivity makes low-overhead logical-resource preparation realistic.
- I do not mind the less glamorous work of turning broad claims into concrete comparisons and executable protocols.
- I am especially interested in projects where a clean theoretical idea can later be tested, benchmarked, or compiled into real hardware control.

Selected references

- L. Daguerre and I. H. Kim, *Code switching revisited: low-overhead magic state preparation using color codes*, 2025.
- L. Daguerre *et al.*, *Experimental demonstration of high-fidelity logical magic states from code switching*, PRX 2025.
- D. Litinski, *Magic State Distillation: Not as Costly as You Think*, Quantum 2019.
- C. Gidney, N. Shutty, and C. Jones, *Magic state cultivation*, 2024.
- H. Bombin, *Dimensional jump in quantum error correction*, New J. Phys. 2016.
- J. Gao, *A Reproducible, Coverage-Aware Benchmark of Auditable Operating Points for Logical $|T\rangle$ Preparation*, 2026 draft.
- J. Gao, *Complexity-Guided Theory Testing: From Bell Witnesses to Logical Magic-State Certification*, 2026 draft.
- J. Gao, *Galilean One-Particle Kinematics from a Smooth Family of Reference States*, 2026 draft.
- K. Wang *et al.*, *Violation of Bell inequality with unentangled photons*, Science Advances 2025.

Thank you!

Questions

Jianshuo Gao