Prospective Structural Classification of Discrete Time-Crystalline Recurrence

Integer-Depth Closure, Blind Validation, and a Quasi-Periodic Domain Boundary

UNNS Research Collective

UNNS Substrate Research Program

August 20, 2026

Abstract

Periodic or subharmonic recurrence is not, by itself, sufficient to distinguish structurally meaningful time-crystalline organization from ordinary classical recurrence. We introduce a frozen recurrence-depth structural metric $C(q)$, a recurrence-family contrast F, and a shuffle-null significance test p, embedded in a hierarchical chamber (TIME-CRYSTAL-I v1.1.0) that separates temporal recurrence from higher-order rigidity, collective, and spectral evidence. The metric, its search range, and its thresholds were frozen against a ten-record known-physics validation corpus *before* a prospective campaign began; candidate identity was withheld from analysis and a cryptographic lock was written before any ground-truth reveal. Three candidates not used in metric calibration were then tested prospectively under the frozen chamber. A prospective large-period digital-quantum-computer discrete-time-crystal (DTC) candidate (C001), not used in metric calibration, returned $q_0 = 4$, $C = 0.5483$, $F = 0.3617$, $p = 0.0149$ (TEMPORAL_RECURRENCE), while its no-recompilation control returned NO_TEMPORAL_ORDER. A prospective numerical prethermal-DTC candidate (C002), not used in metric calibration, returned $q_0 = 2$, $C = 0.5191$, $F = 0.5064$, $p = 0.0050$ (TEMPORAL_RECURRENCE), while its matched $\epsilon = 0$ control returned NO_TEMPORAL_ORDER despite a *higher* raw closure ($C = 0.6250$). A third candidate (C003), posthoc-revealed as a robust experimental $\mathbb{Z}_2$ discrete time quasi-crystal (DTQC), returned $q_0 = 6$, $F = 0.0934$, $p = 0.9602$ (NO_TEMPORAL_ORDER); its matched breakdown-regime control also returned NO_TEMPORAL_ORDER. We conclude that the frozen temporal grammar prospectively identifies a class of supported integer-depth temporal recurrence, and that this same grammar encounters a domain boundary in the tested quasi-periodic multi-clock regime. Non-admission under this frozen grammar is not treated as evidence of the absence of physical temporal order in the DTQC candidate. The results support distinguishing the demonstrated integer-depth recurrence class from the tested quasi-periodic multi-clock regime, motivating a separate future closure formalism rather than a retrofit of the present chamber.

Contents

1 Introduction — From Periodicity to Structural Recurrence

The opening question of this work is methodological before it is physical: *how can one distinguish structurally meaningful discrete time-crystalline recurrence from ordinary periodic or subharmonic behavior?*

Time crystals were first proposed as a possible phase of matter breaking continuous time-translation symmetry [1]; the periodically driven (Floquet) discrete-time-crystal concept that followed established rigidity and stability criteria against local perturbations and disorder [2, 3, 4], and has since been surveyed in several reviews [5, 6]. The present work does not revisit that microscopic theory; it asks a structural-classification question that sits downstream of it.

Three facts motivate the question and rule out the simplest candidate answers.

1. Classical systems can exhibit exact period-2 ($2T$) and period-4 ($4T$) recurrence with no many-body content whatsoever. An exact classical two-cycle and a logistic-map period-4 attractor both reproduce themselves with arbitrarily high fidelity (§3).
2. Consequently, the mere detection of $q_0 = 2$ or $q_0 = 4$ — a fundamental recurrence depth — cannot by itself establish a many-body time-crystalline state.
3. Likewise, a large recurrence amplitude does not automatically imply structural admissibility: we show below (§7) a prospective control record whose raw closure *exceeds* that of its matched true candidate, yet which fails every other criterion.

These facts motivate the central hypothesis of this work:

> *Temporal recurrence can be treated as a structural family whose membership depends not only on recurrence strength, but on recurrence depth, family organization, and separation from appropriate null alternatives.*

This leads to the research question addressed here: *can such a structural recurrence grammar be frozen and tested prospectively on time-crystalline systems excluded from its calibration?*

We answer this question with a three-candidate prospective design, each candidate probing a distinct axis of the claim:

- **C001** tests *recurrence-depth generalization* — whether a grammar developed predominantly in a $2T$ lineage can independently recover a different fundamental depth ($q_0 = 4$) on data excluded from calibration, without being told the target period.
- **C002** tests *candidate/control specificity* in an independent physical setting — whether the grammar can be fooled by a control with higher raw recurrence than its matched true candidate.
- **C003** tests *domain validity* — what the frozen grammar does when confronted with a physically real but structurally different temporal-order regime (quasi-periodic, multi-clock) rather than another integer-period example.

We report that C001 and C002 succeed as designed, and that C003 does not enter the supported class — a result we argue is more scientifically informative than a third positive would have been, because it identifies where the present formalism's domain ends rather than merely extending a string of confirmations.

2 Structural Recurrence Framework

2.1 Recurrence-depth closure

For a candidate recurrence depth q and a temporal state sequence $\{x_t\}$, pairwise closure between states separated by q steps is defined as

$$\text{closure}(x, y) = \text{similarity}(x, y) \times \text{support}(x, y), \tag{1}$$

$$\text{similarity}(x, y) = 1 - \frac{\|x - y\|}{\|x\| + \|y\|}, \tag{2}$$

$$\text{support}(x, y) = \text{clip}\left(\frac{\min(\text{RMS}(x), \text{RMS}(y))}{\text{RMS}_{\text{early}}}, 0, 1\right), \tag{3}$$

where $\mathrm{RMS}_{\mathrm{early}}$ is the median root-mean-square amplitude over the first several time points. The support factor exists specifically to prevent a thermalized, near-zero state from being scored as "closed" merely because $0 \approx 0$: without it, a fully depolarized trajectory would trivially maximize similarity.

Averaging closure over all valid pairs at separation q yields the depth spectrum $C(q)$, $q = 1, \ldots, q_{\max}$ (frozen at $q_{\max} = 10$; Table 2). Here q denotes an integer separation within the ordered temporal sequence, not a period pre-assigned by the experimenter.

2.2 Fundamental recurrence depth

The fundamental recurrence depth q_0 is selected automatically as the value of $q \in \{2, \ldots, q_{\max} - 1\}$ with the largest local spectral contrast against its immediate neighbors,

$$q_0 = \arg\max_q \Big[C(q) - \tfrac{1}{2}\big(C(q-1) + C(q+1)\big) \Big]. \tag{4}$$

Crucially, $q_0 = 2$ is *not* hard-coded anywhere in this selection rule. This becomes important in §6, where a prospective large-period candidate, excluded from calibration, independently produces $q_0 = 4$.

2.3 Recurrence-family contrast

Once q_0 is fixed, the recurrence-family contrast F compares the closure of q_0's multiples (the "family") against all remaining, non-family depths:

$$F = \overline{C(q)}\Big|_{q \bmod q_0 = 0} - \overline{C(q)}\Big|_{q \bmod q_0 \neq 0}. \tag{5}$$

Conceptually, $C(q_0)$ asks *how much recurrence exists*; F asks whether that recurrence belongs to a distinguishable structural family rather than being an isolated coincidence at one depth.

2.4 Shuffle-null significance

A shuffle-null test estimates the probability that the observed family contrast could have arisen from a temporally shuffled surrogate of the same data (200 surrogates, seed 20260819). The frozen temporal gate therefore requires all three conditions jointly:

$$C(q_0) \geq 0.15, \qquad F \geq 0.10, \qquad p \leq 0.10. \tag{6}$$

Figure 1 illustrates the closure spectrum concept on a representative near-ideal trajectory from the raw ingestion corpus (§4), showing the alternating even/odd structure that the local-contrast rule (Eq. 4) is designed to detect.

3 Why Temporal Recurrence Is Not Yet a Time Crystal

Before presenting the prospective results, we use the ten-record validation corpus to establish *why* a four-sector hierarchy (temporal → rigidity → collective → spectral) is necessary at all, rather than a single temporal detector.

The most striking corpus comparison is between an exact classical two-cycle and the genuine many-body MBL-DTC record (Mi et al. dataset [9, 10]). The classical two-cycle achieves $C = 1.000$, $F = 1.000$ — essentially perfect scores on the temporal axis — yet the chamber places it at

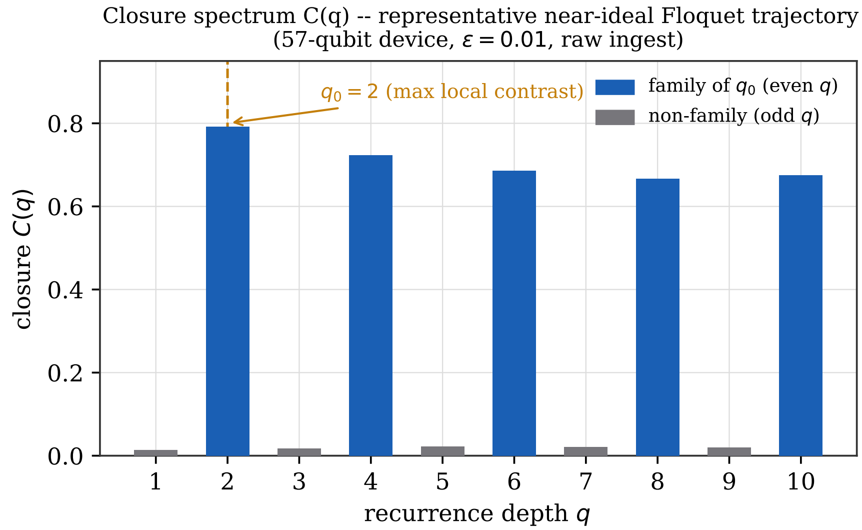

Figure 1: Closure spectrum $C(q)$ for a representative near-ideal Floquet trajectory (57-qubit device, $\epsilon = 0.01$, raw ingestion corpus). The alternating even/odd structure is exactly what the local-contrast rule (Eq. 4) is built to detect; here it correctly selects $q_0 = 2$. This illustrates the concept metric only and is not one of the three prospective candidates.

RIGID_RECURRENCE (Level 2), not Level 4, because it has no rigidity evidence beyond an exact sign-flip family and no collective or spectral content (both correctly marked N/A for a classical record, not failed). Similarly, a classical logistic-map period-2 attractor reaches $C = 1.000$ but only $F = 0.213$ and remains at TEMPORAL_RECURRENCE (Level 1).

By contrast, the Mi 20-qubit MBL-DTC record reaches MANY_BODY_TIME_CRYSTAL_ADMISSIBLE (Level 4) specifically because all four sectors are independently supported: temporal ($q_0 = 2$, $C = 0.954$, $F = 0.947$), rigidity (3/3 significant initial-state pairs, universality $U = 0.986$), collective (finite-size order growth, $\alpha = 0.374$; broad initial-state order, $n = 500$), and spectral (quantum-typicality retention 0.979). Table 1 and Figure 2 summarize all ten records.

Table 1: Validation corpus, ten records. Verdicts consistent with pre-specified expected physics class in all ten cases; N/A marks a quantum-only sector correctly not imposed on a classical/generic-domain record (chamber rule, not a workaround).

ID	Domain	Expected class	Verdict	Lv	q_0	C	F	p
frey57_mbl_dtc	quantum	MBL-DTC	INSUFFICIENT_DOMAIN_EVIDENCE	2	2	0.682	0.618	0.0050
mi20_mbl_dtc	quantum	MBL-DTC	MANY_BODY_TIME_CRYSTAL_ADMISSIBLE	4	2	0.954	0.947	0.0050
mi20_prethermal	quantum	prethermal DTC-like	TEMPORAL_RECURRENCE	1	2	0.293	0.231	0.0846
mi20_thermal	quantum	thermal	NO_TEMPORAL_ORDER	0	2	0.290	0.245	0.2985
classical_exact_2T	classical	ordinary period-2	RIGID_RECURRENCE	2	2	1.000	1.000	0.0050
classical_logistic_2T	classical	nonlinear period-2	TEMPORAL_RECURRENCE	1	2	1.000	0.213	0.0050
classical_logistic_4T	classical	nonlinear period-4	TEMPORAL_RECURRENCE	1	4	0.987	0.222	0.0050
classical_damped_2T	classical	transient period-2	TEMPORAL_RECURRENCE	1	2	0.192	0.118	0.0547
classical_chaos	classical	chaotic	NO_TEMPORAL_ORDER	0	8	0.535	0.020	0.3930
iid_random	generic	random	NO_TEMPORAL_ORDER	0	8	0.260	0.009	0.4975

This corpus-level result licenses the central principle of the framework:

recurrence $\neq$ many-body time-crystalline admissibility.

The prospective campaign reported below (§§6–8) concerns the temporal gate specifically. We state explicitly, and repeat in §12, that C001 and C002 are *not* presented as prospective confirmations of the full Level-4 hierarchy: their rigidity, collective, and spectral sectors remained untested because no standardized evidence for those sectors was supplied in the corresponding bundles.

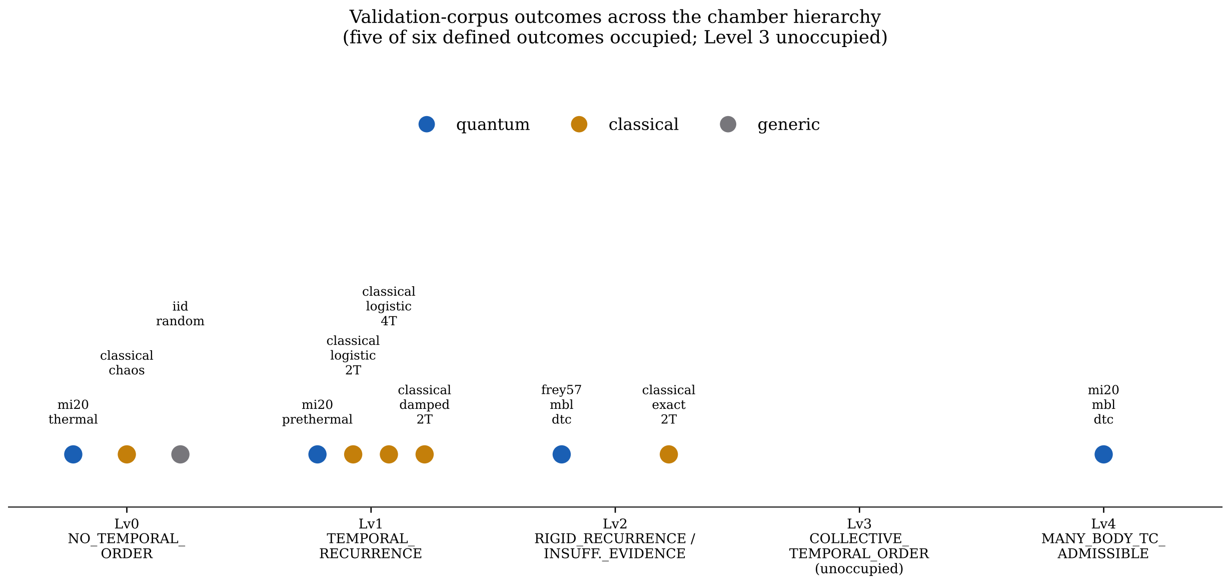

Figure 2: The ten-record validation corpus positioned across the chamber's defined outcomes. Five of the six defined outcomes are occupied; Level 3 (COLLECTIVE_TEMPORAL_ORDER) has no corpus example. The exact classical two-cycle (Lv2) and the Mi MBL-DTC record (Lv4) are the anchor comparison discussed in §3.

4 Independent Physics Reconstruction and Metric Freeze

Methodological credibility for the prospective claims below rests on one design choice: the underlying physics was reconstructed independently *before* the temporal-closure metric was defined, and the metric was calibrated once against the validation corpus and then frozen before any prospective candidate existed to the chamber.

Raw ingestion of the 57-qubit digital-quantum-computer dataset [7, 8] (84 `.dat` files, $5 \times 51 \times 57$ float32 arrays, 30 ϵ values, Brooklyn and Manhattan devices) passed a coarse sanity gate: DTC/thermal half-frequency amplitude ratio $11.729\times$ ($> 5\times$ required). This raw-ingest ratio is deliberately coarse and is distinct from the subsequent methods-faithful reconstruction.

The methods-faithful reconstruction (baseline thresholds $W_0 = 0.15$, $W_f/W_0 = 2/3$) gives three reconstruction indicators together with a two-indicator consensus:

Indicator	Value
Raw discrete variance maximum	$\epsilon = 0.090$
Grid-derived smoothed variance peak	$\epsilon = 0.0755$
Segmented depolarization break	$\epsilon = 0.080$
Two-indicator consensus	$\epsilon = 0.07775$
Published reference ϵ_c	≈ 0.075
Absolute consensus difference	0.00275
Reconstructed half-frequency response ratio	$31.532\times$

The raw discrete critical-fluctuation maximum sits at $\epsilon = 0.090$; the value near 0.075 emerges only after light smoothing whose bandwidth is derived from the ϵ grid itself, not set to match the published answer. The decay-break location is comparatively threshold-sensitive; the subharmonic-variance peak is comparatively stable. We report all three raw, smoothed, and sensitivity results rather than only the favorable number.

The firewall enforced throughout is:

physics reconstruction $\rightarrow$ metric definition $\rightarrow$ validation $\rightarrow$ metric freeze $\rightarrow$ prospective testing.

This is more than a technical detail: it is part of the argument that C001–C003 were not used, even implicitly, to redefine the metric. The frozen temporal metric carries SHA-256 checksum

`06344c8641aa14e6663cc3cba65d86fcdb9e37857909e85da6bc0d689e5b2553`

and the frozen external protocol (thresholds, q-range, shuffle parameters; Table 2) carries

`2e85f96bb66f1d689e89eaa2916bb26639f608311735f450eb2822a64f92c449`.

Table 2: Frozen temporal protocol (`external_quantum_dtc.json`, v1.1.0). Calibrated once against the validation corpus and applied unchanged to C001–C003.

Parameter	Value
q search range	1–10
Minimum closure $C(q_0)$	0.15
Minimum family contrast F	0.10
Maximum shuffle p	0.10
Shuffle surrogates	200
Shuffle seed	20260819

5 Prospective Validation Protocol

The prospective campaign (TC_PROSPECTIVE_01) follows a blind-lock-reveal workflow:

candidate $\rightarrow$ blind analysis $\rightarrow$ evidence audit $\rightarrow$ analysis lock $\rightarrow$ identity reveal $\rightarrow$ posthoc comparison.

We refer to this workflow throughout as the *prospective firewall* (Figure 3), and enforce four rules:

1. **Identity separation.** Physical interpretation and expected class are withheld from the blind analysis package; ground-truth files are stored outside the blind candidate bundle.

2. **Metric immutability.** No threshold or q-range change occurs after the campaign starts. Table 2's values are identical across all three candidates and their controls.

3. **Adapter neutrality.** Adapters may re-express raw source measurements in the chamber's input form (e.g. stroboscopic reduction, dual-clock recombination) but may not inject an expected period, a sign correction chosen to favor a target frequency, a Fourier filter toward a desired period, or synthetic evidence.

4. **Lock before interpretation.** The quantitative result and verdict are cryptographically committed (SHA-256 analysis lock) before any reveal step. All six locks for this campaign are listed in Table 3.

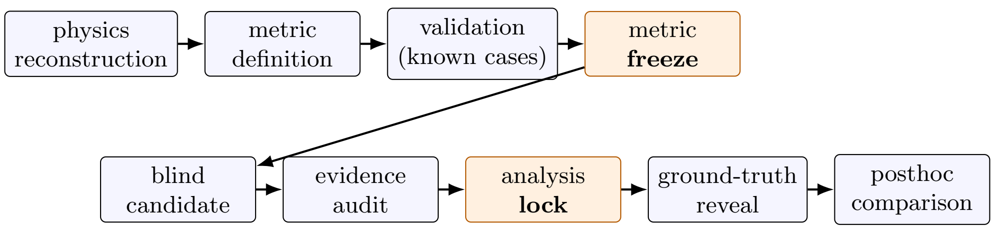

Figure 3: The prospective firewall. Top row: the development-corpus firewall (§4), ending in a metric freeze. Bottom row: the per-candidate prospective protocol (§5), in which the quantitative result and verdict are cryptographically locked *before* identity reveal.

Table 3: Authoritative analysis-lock SHA-256 checksums (truncated), all six locked runs of TC_PROSPECTIVE_01, written before ground-truth reveal.

Run	`analysis_lock_sha256` (truncated)
C001	`5deef733c2f71683868...65f035c16`
NRCTRL	`3dc18024e8310ff0766...512f995a03a4e`
C002	`1404bdcf53eab116713...00cc245e61`
C002_CTRL	`7a9b544919219cea0d4...820737297f3b`
C003	`2cce593284bf341633...7182c1e0c066`
C003_CTRL	`561ed97513e009cff58...d099741e4d18e`

6 Prospective Result I — Generalization to $q_0 = 4$

Candidate C001 was constructed from eight experimental IBM `Z_*` channels of a reported large-period digital-quantum-computer DTC signature [11, 12], stroboscopically combined with no target-period leakage: no sign alignment, no Fourier filter toward period 4, no smoothing, no target period, no $q_{\max}$ change, and no mixing of mitigation variants into the primary state vector. The blind analysis returned

$$q_0 = 4, \quad C(4) = 0.5483, \quad F = 0.3617, \quad p = 0.0149 \;\Rightarrow\; \text{TEMPORAL\_RECURRENCE}.$$

The independent **no-recompilation control** (NRCTRL) — not a matched-parameter control, but a bundle built without the recompilation step applied to the candidate — returned

$$q_0 = 6, \quad F \approx -0.0035, \quad p \approx 0.9950 \;\Rightarrow\; \text{NO\_TEMPORAL\_ORDER}.$$

[demonstrated prospectively] A recurrence grammar developed predominantly in a $2T$ lineage independently selected $q_0 = 4$ on a prospective large-period candidate excluded from calibration, without target-period information, and the no-recompilation control was rejected under the same frozen gate. Higher sectors (rigidity, collective, spectral) remained NOT_TESTED for C001 because no standardized evidence for those sectors was supplied in this bundle; the chamber reports Level 1 rather than inventing evidence for untested sectors.

Figure 4 (left panel) shows C, F, and $1-p$ for C001 against NRCTRL relative to the frozen gate thresholds.

7 Prospective Result II — Closure Is Not Enough

Candidate C002 is an exact 14-spin numerical prethermal-DTC simulation ($N = 300$ slow-kick protocol, near $\epsilon = \pi$), a physically distinct setting from C001. The adapter performed only the documented stroboscopic reduction $(x_t, y_t, z_t) \to (x_{300n}, y_{300n}, z_{300n})$; no alternating sign, target period, Fourier filter, smoothing, or threshold change was introduced. The blind analysis returned

$$q_0 = 2, \quad C = 0.5191, \quad F = 0.5064, \quad p = 0.00498 \;\Rightarrow\; \text{TEMPORAL\_RECURRENCE}.$$

The **matched control** (C002_CTRL, same $N = 300$ protocol at $\epsilon = 0$) returned

$$q_0 = 6, \quad C = 0.6250, \quad F = 0.00124, \quad p = 0.7960 \;\Rightarrow\; \text{NO\_TEMPORAL\_ORDER}.$$

Note that $C_{\text{control}} = 0.6250 > C_{\text{candidate}} = 0.5191$, yet the candidate is classified TEMPORAL_RECURRENCE while the control is classified NO_TEMPORAL_ORDER. We box this result because it is, in our view, the clearest single demonstration in the campaign of why the temporal gate is a three-condition conjunction rather than a closure threshold alone:

> $C(q_0)$ **alone is not the temporal classification rule.** The matched control's closure is real but structurally undifferentiated ($F \approx 0$, $p \approx 0.80$): its recurrence-family contrast is statistically consistent with temporally shuffled surrogates. The candidate's lower closure sits inside a well-separated recurrence family ($F = 0.5064$, $p = 0.005$), whose family contrast is not consistent with the shuffled null.

[demonstrated prospectively] C002 demonstrates candidate/control specificity in a second, physically independent quantum-many-body setting, and shows directly that recurrence-family structure and the null model — not raw closure magnitude — govern the temporal verdict.

8 Prospective Result III — A Quasi-Periodic Domain Boundary

Candidate C003 probes an experimental $\mathbb{Z}_2$ discrete time quasi-crystal with two incommensurate drive clocks [13, 14]. A direct one-observable ingestion was rejected outright because the closure metric requires at least two accepted coordinates; rather than duplicate the signal, use error bars as a fake second coordinate, or invent a measurement, a dual-clock adapter was constructed from the experiment's actual two-clock structure using only stored measured samples, $X_n = [S_x(n\tau_1), S_x(n\tau_2)]$, with no interpolation and no target quasi-crystal frequency introduced.

The blind analysis of the candidate returned

$$q_0 = 6, \quad C = 0.1988, \quad F = 0.0934, \quad p = 0.9602 \;\Rightarrow\; \text{NO\_TEMPORAL\_ORDER},$$

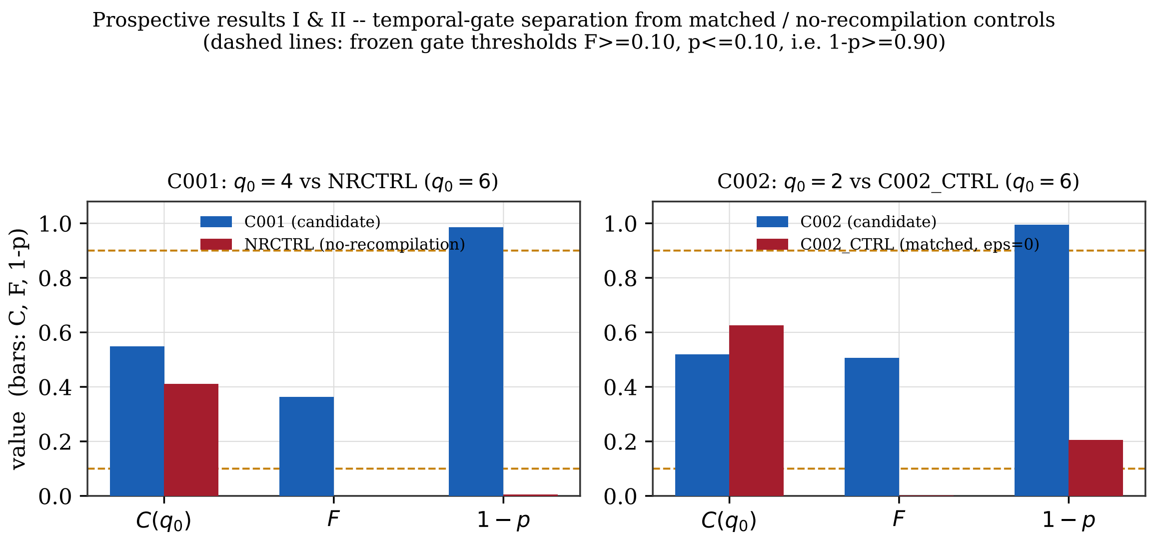

Figure 4: Prospective results I and II. **Left:** C001 ($q_0 = 4$) against its no-recompilation control (NRCTRL, $q_0 = 6$). **Right:** C002 ($q_0 = 2$) against its matched $\epsilon = 0$ control (C002_CTRL, $q_0 = 6$); note the control's higher raw closure C despite failing F and p. Dashed lines mark the frozen gate thresholds ($F \geq 0.10$; $p \leq 0.10$, plotted as $1 - p \geq 0.90$).

and the matched control returned

$$q_0 = 6, \quad C = 0.0946, \quad F = 0.0309, \quad p = 0.9950 \;\Rightarrow\; \text{NO\_TEMPORAL\_ORDER}.$$

Both records fail the frozen temporal gate. Only after both blind analyses were locked were the identities revealed: the candidate is a *robust, long-interaction-regime* experimental DTQC ($\tau_1 = 2.00\,\mu$s, $\tau_2 = 3.236\,\mu$s, rotation factor $1 - \epsilon = 0.933$), and the control is a *matched short-interaction breakdown regime* with split subharmonic responses ($\tau_1 = 0.25\,\mu$s, $\tau_2 = 0.404\,\mu$s).

We emphasize the quantitative — not merely categorical — distinction between candidate and control: the candidate's family contrast ($F = 0.0934$) is roughly $3\times$ the control's ($F = 0.0309$), a larger, but still sub-threshold, recurrence-family contrast in the physically ordered candidate that nonetheless never crosses the frozen $F \geq 0.10$ line (Figure 5).

Interpretation — stated precisely

We do *not* claim "the DTQC has no temporal order." The correct, bounded statement is:

> *Under the frozen dual-clock adapter and the integer-depth ($q \leq 10$) recurrence grammar, the tested robust DTQC regime did not enter the supported temporal-recurrence class.*

NO_TEMPORAL_ORDER is the chamber's verdict *inside this grammar*. It does not mean absence of physical temporal order, and it does not establish that all DTQCs, all multi-clock structures, or UNNS closure formalisms in general lack a representation for this regime — only that this frozen adapter and grammar did not supply one. This distinction is restated in §10 and §12. We do not alter the frozen v1.1.0 record, which correctly reports the machine verdict as NO_TEMPORAL_ORDER; but the more precise reading of that verdict for C003 is that the candidate fell outside the supported integer-depth temporal-recurrence grammar, not that temporal order is absent from the underlying physical system.

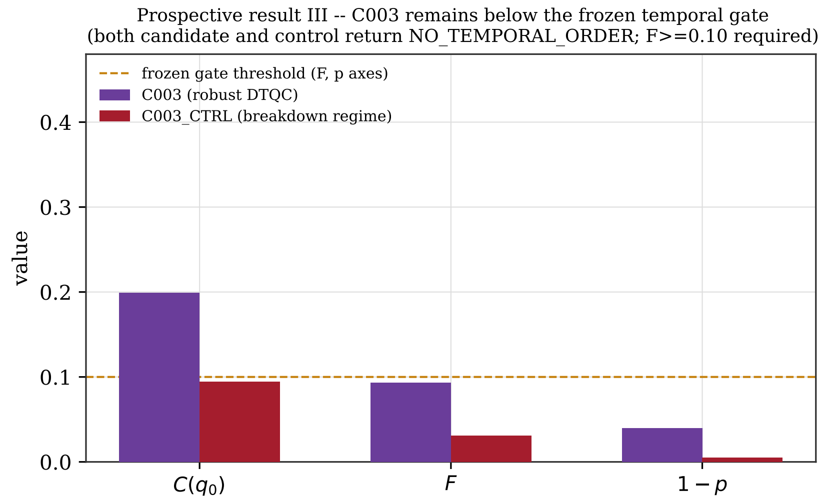

Figure 5: Prospective result III. Both records fail the joint frozen temporal gate. C003 exceeds the closure threshold ($C = 0.1988 \geq 0.15$) but fails both the family-contrast and shuffle-null criteria; C003_CTRL fails all three criteria. The candidate's F is $\sim 3\times$ the control's, but neither crosses the $F \geq 0.10$ threshold (dashed line).

9 Cross-Candidate Structural Classification

Table 4 collects all six locked prospective runs. Bringing C001–C003 together produces an empirical structural map rather than three independent anecdotes (Figure 6):

$$\underbrace{\text{C001 } (q_0 = 4),\ \text{C002 } (q_0 = 2)}_{\text{supported integer-depth recurrence}} \qquad \text{vs.} \qquad \underbrace{\text{C003 (robust DTQC)}}_{\text{unsupported integer-depth recurrence}} \quad .$$

C001 and C002 arise in physically unrelated settings (a digital-quantum-computer trajectory; an exact numerical spin simulation) yet both map into the same supported integer-depth recurrence class under the frozen grammar. C003 shares the broader "time-crystalline" label with C001/C002 in the physics literature, yet lies outside the same structural grammar. This motivates the central conceptual move of the paper: *physical taxonomy ≠ structural taxonomy.* We do not yet formally define a second grammar for the quasi-periodic regime ($R_{\text{multi-clock}}$); C003's role is precisely to motivate, rather than pre-empt, that future development (§10).

Table 4: Prospective campaign, TC_PROSPECTIVE_01. All values from cryptographically locked blind analyses (Table 3); physical class revealed only after locking.

Run	q_0	C	F	p	Temporal verdict	Posthoc physical class
C001	4	0.5483	0.3617	0.0149	TEMPORAL_RECURRENCE	reported large-period DTC candidate
NRCTRL	6	0.4109	−0.0035	0.9950	NO_TEMPORAL_ORDER	no-recompilation control
C002	2	0.5191	0.5064	0.0050	TEMPORAL_RECURRENCE	prethermal DTC (N=300 numerical)
C002_CTRL	6	0.6250	0.0012	0.7960	NO_TEMPORAL_ORDER	matched control, $\epsilon = 0$
C003	6	0.1988	0.0934	0.9602	NO_TEMPORAL_ORDER	robust experimental DTQC
C003_CTRL	6	0.0946	0.0309	0.9950	NO_TEMPORAL_ORDER	DTQC breakdown-regime control

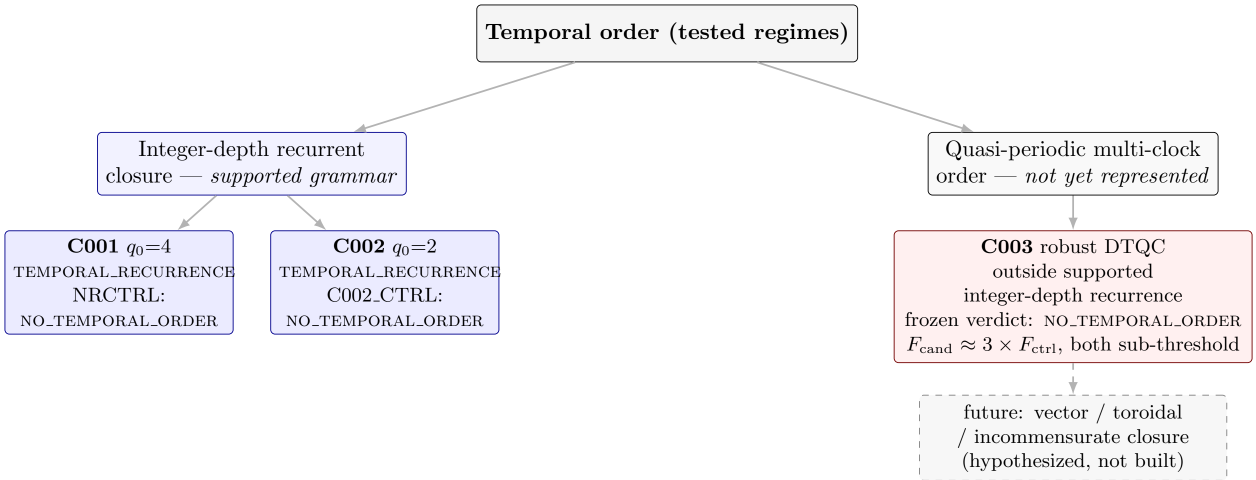

Figure 6: Temporal structural classification map. C001 and C002 (physically unrelated systems) both enter the supported integer-depth recurrence class under the frozen grammar; their controls do not. C003 (a physically real, robust DTQC) remains outside that same class under the identical frozen grammar — the first documented domain boundary of TIME-CRYSTAL-I v1.1.0. The quasi-periodic branch and its hypothesized future closure coordinates are not developed in this work.

10 Discussion — What Has Actually Been Classified?

10.1 Structural class rather than physical label

C001 and C002 arise in different physical contexts (digital-quantum-computer hardware; exact numerical spin simulation) but map into the same supported integer-depth recurrence class under the frozen grammar. C003 shares the broader time-crystalline label but lies outside the same structural grammar. Physical taxonomy and structural taxonomy are therefore not interchangeable within this framework, and treating them as identical would overstate what the chamber has shown.

10.2 Non-admission is grammar-relative

We formalize the lesson of §8 explicitly: non-admission under a frozen grammar G is not equivalent to physical non-existence of the corresponding order. It means only that the tested system, under the specific adapter and closure coordinates comprising G, does not satisfy G's admissibility conditions. This is a necessary maturation of the UNNS terminology used in earlier chamber-development work, where "admissible" and "exists" were not always kept as distinct concepts.

10.3 Domain discovery is part of validation

A successful prospective instrument need not return positives indefinitely. C003 is scientifically valuable precisely because the metric remained frozen and encountered a boundary rather than being retuned to accommodate it. A chamber capable of producing a bounded, reproducible negative result under a frozen protocol is more useful than one that can always be adjusted post hoc until the answer becomes positive; the latter would not have been falsifiable in any meaningful sense.

10.4 Toward temporal admissibility families

We introduce, but explicitly do not develop, candidate coordinates for a future quasi-periodic closure formalism: vector recurrence depth (q_1, q_2), toroidal phase closure on $\mathbb{T}^2$, irrational/incommensurate recurrence ratios, and multi-frequency closure families more generally. These are hypotheses motivated by C003, not results of the present work, and any such branch requires its own unseen validation data rather than retrofitting into TIME-CRYSTAL-I v1.1.0.

10.5 Structural grammars have empirical domains

Here, *empirical domain* refers to the range of physical regimes for which the grammar has demonstrated structural validity, not merely the set of inputs on which the algorithm can be evaluated.

The domain boundary encountered in §8 supports a more general reading than a single negative result. A structural grammar is fixed by an adapter, a set of closure coordinates, and a set of frozen thresholds; nothing guarantees in advance that it extends to every physically ordered system that might be presented to it. The prospective campaign shows that whether a given physical system falls inside or outside a frozen grammar's domain is not always decidable by inspection of the physics alone — it is an empirical question, answerable only by locked, prospective testing of the kind performed here.

Two logically distinct explanations are available whenever a candidate returns a non-admission verdict, and a single chamber run cannot by itself distinguish between them: the candidate may genuinely lack the structural organization the grammar is built to detect, or the candidate may possess that organization in a form the grammar's representation does not preserve. C003's non-admission is naturally read as the second case, since its physical DTQC order is independently documented in the source literature; the parsimonious conclusion is that the present dual-clock, integer-depth adapter does not represent that order, not that the order is absent. Adjudicating between the two explanations in general requires an independently specified, frozen alternative representation, tested prospectively on new candidates — not a retrospective adjustment of the grammar that already produced the negative result (§10).

Read this way, C003 is not simply a result to be set aside as inconclusive. It marks a discovered edge of the present grammar's empirical domain, and it motivates future closure formalisms explicitly: any successor grammar built to admit quasi-periodic order should itself be frozen and tested on new, unseen candidates under the same prospective discipline used in this work, rather than validated against C003 itself.

11 Relation to the UNNS Substrate: Temporal Structure as a Distinct Structural Chart

The relevant substrate-level implication of this work is narrower than a general claim about recursive structure in physics: UNNS is useful here not as an alternative microscopic theory of time crystals, but as a structural coordinate system for classifying recurrence organization. The time-crystal chamber provides an unusually clean instance of this idea because the prospective campaign demonstrates *both* directions of the underlying claim: different physical settings (C001, C002) map to the same structural class, while a system carrying the same broad physical umbrella label ("time-crystalline," including C003) does not map into the same supported structural class under the present grammar. We draw out three specific consequences of this below, deliberately restricted to what the prospective results themselves establish.

11.1 Order-preserving temporal representation

Structural analysis elsewhere in the UNNS Substrate program is built primarily on ordered numerical ladders $L = (x_1 \leq x_2 \leq \cdots \leq x_n)$, where the object of study is the gap geometry of a sorted sequence and sorting is itself a legitimate, structure-preserving step. The present chamber operates on a different kind of object: a temporal sequence $X_0, X_1, X_2, \ldots$ in which order in time carries the structural information under test. Sorting the values would destroy exactly the property the chamber is designed to detect. TIME-CRYSTAL-I therefore instantiates a sequence-order-preserving structural grammar, distinct from ladder-based admissibility, standing alongside the Substrate's existing ladder formalism rather than as a special case of it.

11.2 From a single structural space to multiple charts

Read together with the ladder formalism, this suggests that structural organization within the UNNS framework is better described through multiple structural charts than through a single master coordinate: a ladder chart for ordered numerical structure, and a temporal-recurrence chart, of the kind developed here, for order-preserving sequences. We do not attempt, in this manuscript, to place these charts in formal correspondence. In particular, we do not establish a mapping between the temporal closure coordinates (C, F, p) used throughout this work and any established UNNS realizability coordinate; any such correspondence, if one exists, is left to future work with its own frozen protocol and its own prospective test, rather than asserted here.

11.3 Grammar-relative structural visibility

The C003 result illustrates a general point about structural grammars: a physically real organization may be fully present in a system yet fail to appear as an admissible structure within a particular chart when that chart's representation and coordinates do not preserve the relevant organization — a form of grammar-relative structural visibility (grammar-selective structural observability). C003 is the motivating empirical case for this reading, not an independent demonstration of the exact representational mechanism responsible for it. This does not privilege one chart as more fundamental than another; it means that admissibility, as reported by any one chamber, is relative to the grammar-representation pair actually used, and that a non-admission verdict inside one chart leaves the question open in charts not yet built or tested.

12 Limitations and Falsifiability

We state limitations plainly rather than treat them as caveats to be minimized.

- Prospective validation in this work applies only to the *temporal* sector. C001 and C002 do not prospectively validate rigidity, collective, or spectral generalization; those sectors remained NOT_TESTED because no standardized evidence for them was supplied in the corresponding bundles.
- C003 constitutes one experimental DTQC candidate/control pair. It establishes a domain boundary for this frozen adapter and grammar, not a general statement about discrete time quasi-crystals.
- The dual-clock adapter used for C003 is physically grounded in the source experiment's actual two-clock structure, but it is not established as the unique possible representation of DTQC dynamics for closure analysis.

- The recurrence-depth search is finite ($q \leq 10$); a physically real recurrence at depth $q > 10$ would not be found by the current implementation.
- The shuffle-null formulation and its frozen thresholds (Table 2) were calibrated once against a ten-record corpus and require broader independent testing across more physical domains.
- The Level-4 hierarchy (§3) has validation-corpus support (one Level-4 record: Mi MBL-DTC) but no prospective confirmation to date.
- The exact generating code and external provenance for the C002 numerical source trajectories are not resolved in the project's internal records (Data and Code Availability); C002's blind/lock/reveal chain is intact, but this open provenance item should be closed before external submission.

We list explicit ways the central conclusion could fail under future testing, in the interest of falsifiability:

- Independent DTC candidates, analyzed under the same frozen protocol, may fail to reproduce the C001/C002 pattern of candidate/control separation.
- A future matched control, constructed under the same neutrality rules, may satisfy the frozen temporal gate where none has so far.
- An independently preregistered, neutral representation of DTQC data — constructed without knowledge of this result — may enter the supported integer-recurrence class, which would narrow or overturn the domain-boundary interpretation of §8.
- Threshold-perturbation analysis (varying $F_{\min}$, $p_{\max}$, or $q_{\max}$ within physically reasonable ranges) may reveal that the C001–C003 separations are less stable than the point estimates in Table 4 suggest.

Hierarchy of claims. We maintain three explicit levels of claim strength throughout this manuscript:

- **[demonstrated prospectively]** $q_0 = 4$ support in C001; $q_0 = 2$ support in C002; candidate/control temporal separation in both cases; the C003 domain boundary under the frozen adapter and grammar.
- **[validation-corpus evidence]** The four-sector hierarchy distinguishes ordinary classical recurrence (Level 2 ceiling for an exact classical two-cycle) from a fully evidenced many-body DTC case (Level 4 for the Mi record).
- **[hypothesized]** Quasi-periodic systems may require a separate admissibility family — multi-clock, toroidal, or vector-depth closure coordinates.

Keeping these three levels visibly distinct is intended to make the manuscript's claims difficult to overinterpret in either direction.

13 Conclusion

A frozen structural recurrence grammar was prospectively tested against three time-crystalline candidates excluded from metric calibration, under blind analysis and post-analysis identity reveal. Two DTC-related candidates independently entered the supported integer-depth recurrence class at $q_0 = 4$ and $q_0 = 2$, while their controls did not. A robust experimental discrete time quasi-crystal remained outside the same support grammar, establishing the first observed domain boundary of the frozen temporal metric. These results support the interpretation of temporal recurrence as a structural admissibility family rather than a universal proxy for time-crystalline order, and motivate a separate prospective treatment of multi-clock quasi-periodic organization in future work. More generally, the prospective boundary encountered at C003 shows that the domain of a structural grammar can itself be discovered empirically: physically real temporal organization need not be structurally visible in every admissibility chart.

Data and Code Availability

The frozen temporal-closure implementation, the validation corpus, and all six locked prospective runs (analysis locks, evidence audits, and posthoc-comparison records) are maintained under the TIME-CRYSTAL-I v1.1.0 and TC_PROSPECTIVE_01 project trees. This internal archive is not yet deposited at a public, citable repository; a versioned public archive (code, frozen metric implementation, validation corpus, and the six locked prospective run bundles) will be deposited at a permanent repository (e.g. Zenodo) at the time of external submission, and this section will be updated with the corresponding DOI. The frozen metric SHA-256 (§4) and the six analysis-lock SHA-256 checksums (Table 3) allow integrity verification of the reported quantities against the preserved locked records; a hash alone does not externally certify chronology, and lock chronology for this campaign is documented by the campaign logs and blind/reveal records rather than by the hashes themselves. A known hash-portability issue (non-canonical byte-level serialization across operating systems) affects derived provenance hashes but not the reproduced scientific quantities (q_0, C, F, p, verdict), which were confirmed identical across Linux and Windows runs; the locally generated Windows analysis locks are treated as authoritative for this campaign (Appendix A.2).

External-data provenance. This manuscript is explicitly prospective and reproducibility-oriented, so the bibliography (References, below) is organized to let an external reader retrace the full provenance chain — *published physical claim* → *exact source archive* → *our neutral transformation* → *locked chamber result* — rather than only the surrounding literature. Foundational/context references [1, 2, 3, 4, 5, 6] are kept compact; the physical papers and exact public data archives for every system actually analyzed (Frey–Rachel/Zenodo 5807219 for §4; Mi et al./Zenodo 5570676 for the Level-4 validation record of §3; Chen et al./Zenodo 15421652, source of `4T-DTC_upload.tar`, for C001 in §6; He et al./Zenodo 14791637, source of `rawdata.xls`, for C003 in §8) are listed in full, since these — not the general background — constitute the auditable trail behind the reported numbers.

The C002 numerical trajectories (14-spin, $N = 300$ slow-kick protocol at $\epsilon = 3.1416$ and $\epsilon = 0.0000$; §7) were preserved as source files with recorded SHA-256 provenance hashes (`C002_SOURCE_HASHES.json`) prior to adapter construction, but the project's internal records do not include the generating simulation code or an external literature/data-archive citation for these two files. Their origin — in-house numerics versus an uncited external source — is therefore *not resolved* by the materials available to this manuscript. We flag this explicitly as an open provenance item rather than asserting either possibility: closing it, before external submission, requires either recovering and archiving the exact generator script and parameters, or identifying and citing an external source, and until then C002 should be read as methodologically sound (its blind/lock/reveal chain is intact) but provenance-incomplete relative to C001, C003, and the validation corpus.

A Supplementary Material (summary)

The following material is maintained alongside the chamber and campaign artifacts rather than reproduced in full here: exact metric implementation and source; threshold calibration history against the validation corpus; individual $C(q)$ spectra for all ten validation records and all six prospective runs; shuffle-null surrogate distributions; the C001 stroboscopic adapter specification; the C002 stroboscopic reduction adapter; the C003 dual-clock adapter and its rejection of a direct one-observable ingestion; the evidence-audit schema; the full blind/reveal protocol document; all cryptographic hashes referenced in this manuscript; the Windows/Linux hash-portability note (Appendix A.2); the complete Level-4 chamber sector definitions; the TC_RIGIDITY_v001 negative-result details (below); and the perturbation-basin retention analyses.

A.1 Rigidity-sector negative result

A deliberate negative result from TC_RIGIDITY_V001 motivated the four-sector architecture used throughout this work. The site-resolved Mi MBL-DTC closure spectrum and the exact ordinary classical $2T$ ladder are geometrically almost the same object: cosine similarity 0.999988952, normalized RMS distance 0.005819. On the monotone recurrence-rigidity axes (C, F, S) (with $S = 1 - p$), the ordinary classical two-cycle dominates or equals the genuine MBL-DTC on every axis tested; the same holds for (C, F, S, U) (with U an initial-state universality score) comparing an ordinary sign-flip family against the Mi initial-state result. Consequently, no monotone scalar built only from C, F, S, U can be justified as a DTC-specific order parameter. The same rigidity vector does still separate the Mi MBL-DTC from its thermal and prethermal controls. Perturbation-basin retention areas: Frey MBL-DTC 0.982; classical $2T$+noise 0.923; classical $2T$+damping 0.670.

A.2 Hash portability note

Scientific outputs (q_0, closure, family contrast, shuffle p, sector status, verdict) reproduced identically across Linux-side and Windows-side runs of the prospective campaign. Serialized evidence hashes and derived analysis-lock hashes differed across platforms, most likely due to non-canonical byte-level serialization or line-ending differences rather than any change in the underlying computation. The campaign treats the locally generated Windows analysis locks (Table 3) as authoritative. This affects provenance portability only; a future maintenance release should canonicalize serialization without altering the frozen metric.

References

[1] F. Wilczek, "Quantum Time Crystals," *Phys. Rev. Lett.* **109**, 160401 (2012).

[2] D. V. Else, B. Bauer, and C. Nayak, "Floquet Time Crystals," *Phys. Rev. Lett.* **117**, 090402 (2016).

[3] N. Y. Yao, A. C. Potter, I.-D. Potirniche, and A. Vishwanath, "Discrete Time Crystals: Rigidity, Criticality, and Realizations," *Phys. Rev. Lett.* **118**, 030401 (2017).

[4] C. W. von Keyserlingk, V. Khemani, and S. L. Sondhi, "Absolute Stability and Spatiotemporal Long-Range Order in Floquet Systems," *Phys. Rev. B* **94**, 085112 (2016).

[5] K. Sacha and J. Zakrzewski, "Time Crystals: A Review," *Rep. Prog. Phys.* **81**, 016401 (2018).

[6] M. P. Zaletel, M. Lukin, C. Monroe, C. Nayak, F. Wilczek, and N. Y. Yao, "Colloquium: Quantum and Classical Discrete Time Crystals," *Rev. Mod. Phys.* **95**, 031001 (2023).

[7] P. Frey and S. Rachel, "Realization of a Discrete Time Crystal on 57 Qubits of a Quantum Computer," *Sci. Adv.* **8**, eabm7652 (2022).

[8] P. Frey, *Realization of a Discrete Time Crystal on 57 Qubits of a Quantum Computer* [dataset], Zenodo (2021). `https://doi.org/10.5281/zenodo.5807219`

[9] X. Mi et al., "Time-Crystalline Eigenstate Order on a Quantum Processor," *Nature* **601**, 531–536 (2022).

[10] X. Mi et al., *Data and Code for "Time-Crystalline Eigenstate Order on a Quantum Processor"* [dataset], Zenodo (2021). `https://doi.org/10.5281/zenodo.5570676`

[11] T. Chen, R. Shen, C. H. Lee, B. Yang, and R. W. Bomantara, "A Robust Large-Period Discrete Time Crystal and its Signature in a Digital Quantum Computer," *Quantum Sci. Technol.* **11**, 025030 (2026).

[12] T. Chen et al., *Code and Essential Data for "A Robust Large-Period Discrete Time Crystal and its Signature in a Digital Quantum Computer"* [dataset], Zenodo (2025). `https://doi.org/10.5281/zenodo.15421652`

[13] G. He, B. Ye, R. Gong, C. Yao, Z. Liu, K. W. Murch, N. Y. Yao, and C. Zu, "Experimental Realization of Discrete Time Quasicrystals," *Phys. Rev. X* **15**, 011055 (2025).

[14] G. He, B. Ye, R. Gong, C. Yao, Z. Liu, K. W. Murch, N. Y. Yao, and C. Zu, *Raw Data: Experimental Realization of Discrete Time Quasi-Crystals* [dataset], Zenodo record 14791637 (2024). Source file `rawdata.xls`; at time of writing the Zenodo record's own metadata exposes the associated arXiv DOI (10.48550/arXiv.2403.17842) rather than a distinct Zenodo-minted DOI, so no separate Zenodo DOI is asserted here.