

Research Article

Measuring Dimensional Collapse: A Two-Domain Framework for Quantifying Response Space Constriction

Daniel Andreas Borschel¹¹. Independent researcher

Trauma-related and chronic stress-related conditions are widely described as involving a narrowing of the responses available to an organism. This idea — recently formalised as *dimensional collapse of response space* — has clinical face validity but no operational definition, and therefore cannot be tested, falsified, or compared across studies.

This paper proposes a measurement framework that renders the construct computable. The central claim is that response space constriction is a property of *distributions*, not of *magnitudes*, and can therefore be quantified using a single family of diversity statistics — Hill numbers — applied in two separable domains: the physiological domain, where the effective dimensionality of autonomic reorganisation across a standardised multi-state protocol is estimated; and the repertoire domain, where the effective number of regulatory strategies is estimated from the subscale profile of an established multidimensional self-report instrument.

Both domains yield indices on the same scale — an *effective number of dimensions* — permitting direct comparison, a coherence statistic, and a discrepancy statistic with potential clinical meaning. The framework is instrument-agnostic in the repertoire domain and requires only stated minimum sensor conditions in the physiological domain.

We report a synthetic validation of the computational pipeline. Simulated organisms constructed to have identical resting RMSSD and identical total excursion but differing configuration geometry were separated by the physiological index (1.55 vs 1.08, non-overlapping 95% intervals), confirming that the index carries information unavailable to conventional resting-state HRV measures. The same simulations revealed two problems not anticipated in the framework's initial specification: the originally proposed six-feature set is substantially redundant, and a whitening correction that

appeared theoretically motivated in fact destroys discrimination. Both findings are reported and the specification revised accordingly.

We derive a falsifiable double-dissociation prediction distinguishing post-traumatic stress disorder (phasic collapse) from myalgic encephalomyelitis/chronic fatigue syndrome (tonic collapse), and specify the empirical validation required. No efficacy claim of any kind is made or implied.

Correspondence: papers@team.qeios.com — Qeios will forward to the authors

1. Introduction

1.1. *The construct and its measurement gap*

Clinical description of trauma sequelae converges on language of narrowing. Patients are described as having fewer options, a smaller world; as reacting identically to situations that differ; as being unable to modulate. In theoretical writing this appears as rigidity, loss of flexibility, defensive fixation, or constriction. A recent formalisation proposed treating the narrowing literally, as a **dimensional collapse of response space**: the organism's accessible state space loses degrees of freedom, so that situations which are objectively distinct come to elicit increasingly similar responses^[1].

The construct is attractive because it is domain-general. It applies to autonomic regulation, affect, interpersonal behaviour, and cognition, and it suggests that superficially unrelated presentations may share a formal structure. It is also, at present, **untestable** — because nobody has said what number to compute.

This is not a minor omission. A construct without an operationalisation cannot accumulate evidence. It can only accumulate advocacy.

1.2. *Why existing measures do not capture it*

Heart rate variability magnitude indices. Resting RMSSD, SDNN, and high-frequency power are widely used as markers of autonomic regulatory capacity, and reduced resting HRV is reliably associated with a broad range of psychopathology^{[2][3]}. But these are *scalar* quantities derived from a *single* state, and a scalar from one state cannot in principle express dimensionality. Two individuals may have identical resting RMSSD while one reorganises markedly across rest, paced breathing, orthostatic challenge and

recovery, and the other does not. On the construct proposed here these are qualitatively different organisms; on resting RMSSD they are indistinguishable. Section 8 demonstrates this concretely.

This is not a criticism of HRV magnitude measures, which do what they were designed to do. It is a statement that they answer a different question: *how much vagal tone is present now, not in how many independent directions can this system move.*

Symptom severity scales. Instruments such as the PCL-5, PHQ-9 or PSS-10 quantify reported distress. Distress and constriction are empirically correlated but conceptually distinct: a person may report high symptom burden while retaining a wide behavioural repertoire, or — clinically more familiar and more concerning — report low burden precisely because the repertoire has narrowed to a small set of well-defended, ego-syntonic responses. A severity total cannot distinguish these.

A third literature is directly relevant. **Regulatory flexibility**^[4], **psychological flexibility**^[5] and **emotional granularity**^[6] all address something adjacent. These are typically operationalised either through experimental paradigms of limited field feasibility, or through self-report scales that again yield a severity-like total. None produces a domain-general dimensionality estimate computable from routinely collected data. The present framework is intended to complement this work, and Section 9.2 specifies the discriminant tests required to show that it does more than rename it.

1.3. What this paper does and does not claim

This paper specifies **how to measure**. It reports no human data; the empirical validation is designed here and will be reported separately. It makes no claim that any intervention alters the proposed indices.

Indeed, a prior pre-post analysis conducted by the author using resting-state HRV endpoints returned a null result (Hedges' $g = -0.31$). That null is part of the motivation for this paper: if an endpoint cannot in principle express the construct, a null on that endpoint carries no information about the construct. The appropriate response to such a finding is not a larger sample. It is a better endpoint — or the abandonment of the construct. This paper attempts the former while specifying, in Section 9.1, exactly what would compel the latter.

2. Theoretical background

2.1. Response space

Let an organism's momentary state be a point in a multidimensional space whose axes are independently varying regulatory parameters. Across time and situations the organism traces a trajectory. Two properties of that trajectory matter here:

- **Excursion** — how far the organism moves. A system responding strongly to challenge has large excursion.
- **Dimensionality** — in how many *independent directions* it moves.

These are separable, and the separation is the entire point. A system can move a great deal along a single axis — large excursion, low dimensionality. This is the phenomenology of a person who has exactly one response to everything and deploys it vigorously. A system can also occupy many axes while barely moving along any — flattening without rigidity.




Figure 1. Excursion and dimensionality are independent properties of a state configuration. Each panel shows five hypothetical states in a two-dimensional regulatory space. The second panel — large excursion confined to a single direction — is the pattern that magnitude-based measures cannot distinguish from the first.

Dimensional collapse is defined as reduction in the effective dimensionality of the accessible state space: the loss of *independence between axes*, such that parameters which should vary separately become yoked and the organism's response to any perturbation approaches a scalar multiple of one stereotyped pattern.

2.2. Two forms of collapse

The framework distinguishes two forms, not assumed to co-occur:

Phasic collapse. The resting configuration may be within normal limits, but the system cannot *reorganise* under challenge — it either fails to shift or shifts stereotypically regardless of the nature of the challenge. Excursion may even be exaggerated. Hypothesised for post-traumatic conditions.

Tonic collapse. The accessible space is chronically reduced across all conditions including rest. Excursion is uniformly low. Hypothesised for conditions of chronic energetic limitation.

This distinction generates the prediction in Section 9.1. It is stated here because it imposes a protocol requirement: **a protocol measuring only rest cannot distinguish the two forms and cannot detect phasic collapse at all.**

2.3. The domain-general claim

The construct is claimed to apply in more than one domain. The framework takes this seriously by requiring that the *same* mathematics apply in each. If dimensional collapse is a general property, a general statistic should capture it; if separate ad-hoc formulas were needed per domain, the generality claim would be doing no work.

This yields a testable consequence: indices computed in different domains with the same formula should converge. Section 6 makes this the central validity test, and Section 11.7 states plainly what a failure would mean.

3. The unifying statistic

3.1. Effective dimensionality via Hill numbers

Both domains reduce to one problem: given a non-negative vector describing how a quantity is distributed across components, compute the *effective number* of components meaningfully occupied.

Ecology solved this with **Hill numbers**^[7] (see ^[8], for the modern treatment). For a normalised non-negative vector $\mathbf{p}$ with J components, where $\sum_j p_j = 1$, the Hill number of order q is

$$D^{\wedge q} = (\sum_j p_j^q)^{1/q} \text{ summed over } j = 1 \dots J$$

with the limiting case at $q = 1$ given by the exponential of Shannon entropy. Two orders are relevant:

- $D^1 = \exp(H)$, $H = -\sum p_j \ln p_j$ — effective number of components, sensitive to rare components.
- $D^2 = 1 / \sum p_j^2$ — inverse Simpson index, identical to the **participation ratio**, dominated by strongly occupied components and considerably more robust to measurement noise in small ones.

Hill numbers behave intuitively: if a quantity is spread evenly across exactly k components and absent from the rest, $D^q = k$ for every q . Output is therefore directly interpretable as “this organism is effectively operating in 2.8 dimensions”.

D^2 is specified as the **primary statistic in both domains** on grounds of noise robustness, with D^1 as a pre-registered sensitivity analysis. This choice must be fixed in advance, not selected post hoc.

3.2. *Why this family*

Three properties recommend Hill numbers over raw entropy, variance ratios, or count-based indices:

1. **Scale.** Output is in units of effective components, not nats or arbitrary index units — the precondition for the cross-domain comparison of Section 2.3.
2. **Replication principle.** Doubling the number of equally occupied components doubles the index. Entropy lacks this property, which is why entropy values are hard to compare across instruments with differing component counts.
3. **Tunable sensitivity.** The order q provides a principled rather than arbitrary way to vary sensitivity to weakly occupied components.

4. Domain D: physiological dimensionality (RRI-D)

4.1. *Principle*

The organism is measured in K standardised states. For each state a vector of M autonomic features is computed. The $K \times M$ configuration describes where the organism goes when perturbed in defined ways. RRI-D estimates the effective dimensionality of that configuration.

The key inferential move: **the signal is not in the values within any state, but in the geometry of the differences between states.**

4.2. Feature set

The feature set is the framework's most consequential technical choice, and simulation (Section 8.3) showed it to be the binding constraint on index performance — more so than the number of states.

Core set (recommended, $M = 4$):

Feature	Axis represented
ln RMSSD	Short-term vagal
ln LF power	Baroreflex / mixed
Mean RR	Chronotropic level
SD1/SD2	Short- vs long-term balance

Features deliberately excluded from the core set:

ln HF power correlated at $r = 0.92$ with *ln RMSSD* in the simulated reference sample. The two are near-redundant: both index respiratory-linked vagal modulation. Including both inflates the apparent occupancy of a single direction and **reduced discrimination between archetypes by roughly 40%** (Section 8.3). HF power should be reported descriptively but not entered into the feature matrix.

Sample entropy proved noise-dominated at the segment lengths available (3–5 min) and likewise reduced discrimination. It is retained as an optional fifth feature only where segment durations exceed 10 minutes.

LF/HF ratio is never entered into the feature matrix, given well-documented interpretive problems^[9]. It may be reported descriptively.

Frequency-domain features are computed by Welch periodogram on evenly resampled (4 Hz, cubic spline) tachograms. Log transformation is applied to all power and dispersion measures.

Users adding features should note that adding a feature correlated with one already present will lower, not raise, measured dimensionality. This is counter-intuitive and is the error most likely to be made by secondary users.

4.3. Computation

Step 1 — Feature extraction per state, as above.

Step 2 — Reference standardisation. Each feature is z-transformed using reference means and standard deviations derived from a **healthy comparison sample measured under the identical protocol**, not from within-person values. Within-person standardisation would remove precisely the between-state variance the index quantifies.

Published HRV reference values^{[10][11]} are **not suitable** for this purpose, for two reasons. They cover resting conditions only, and four of the five states in the protocol of Section 4.6 have no published norms at all. And even for the resting state they are methodologically heterogeneous — differing in recording duration, breathing condition, and spectral estimation method — with frequency-domain measures in particular varying substantially between source studies. Substituting them would introduce more error than it removes.

The reference sample must therefore be collected alongside the study, using the same sensor, protocol, and extraction pipeline. Approximately 50–80 healthy participants suffice for stable feature means and standard deviations; in a case-control design the control group serves this function directly.

The consequence must be stated explicitly in any report: RRI-D is sample-relative, not absolute. Values are comparable within a study but not across studies until a multi-site reference dataset exists. Establishing such a reference is the natural second-generation contribution and is not attempted here.

Step 3 — Centring. The person's $K \times M$ matrix $\mathbf{X}$ is column-centred *within person*, yielding $\tilde{\mathbf{X}}$. Each row now expresses that state's deviation from the person's own average state.

Step 4 — Eigendecomposition. Compute the eigenvalues of the $M \times M$ covariance of $\tilde{\mathbf{X}}$, ordered from largest to smallest, of which r are non-zero, with $r \leq \min(M, K-1)$.

Step 5 — Indices.

Excursion: $E = \text{tr}(\Sigma) = \sum_i \lambda_i$

RRI-D: $D^2(\lambda) = (\sum_i \lambda_i)^2 / (\sum_i \lambda_i^2)$

RRI-D is the effective number of independent directions in which the organism reorganises. **E is reported alongside it and never merged into it:** the two are conceptually distinct (Section 2.1), and their dissociation is the substance of Sections 8 and 9.

4.4. Whitening: tested and rejected

Because features are correlated in the population, it appeared theoretically motivated to whiten each person's state matrix against the normative covariance before eigendecomposition, thereby removing population-level collinearity and allowing state excursions to occupy independent directions.

This was implemented and tested, and it fails. Under whitening, the three simulated archetypes converged to RRI-D values of 1.09, 1.14 and 1.19 — indistinguishable, and with the intact archetype now scoring *lowest*. The mechanism is straightforward in retrospect: whitening amplifies the low-variance directions of the normative covariance, and those directions are dominated by measurement noise rather than physiological signal. The correction inverts the signal-to-noise ratio.

This negative result is reported because it is a natural thing for a subsequent user to try, and because the framework's credibility depends on its failures being visible. **The unwhitened pipeline of Section 4.3 is the specification.**

4.5. Ceilings and the binding constraint

RRI-D cannot exceed $\min(M, K-1)$. This ceiling must be reported in every application, and comparisons across studies with differing M or K are invalid without explicit adjustment.

However, simulation showed that **the nominal ceiling is not the operative constraint**. With $K = 5$ (nominal ceiling 4.0) and the original six-feature set, the intact archetype reached only 1.55, and adding states barely moved it ($K = 3$: 1.59; $K = 4$: 1.59; $K = 5$: 1.64). The limiting factor was redundancy within the feature set, whose own effective dimensionality was 3.44 of 6 nominal features. Reducing the feature set raised the intact archetype to 1.96 — a *larger* gain than any increase in K produced.

The practical implication reverses the framework's initial emphasis: **effort is better spent on feature independence than on adding states**. $K \geq 4$ is sufficient; $K \geq 5$ is preferable but not the priority.

4.6. Acquisition protocol

Sensor requirement. A validated chest-strap ECG-based sensor (e.g. Polar H10) with raw RR export. Optical PPG wrist sensors are **not acceptable**. The author's comparison work found optical output unsuitable for non-linear endpoints; beat detection may suffice for mean rate and gross time-domain measures, but the beat-to-beat timing precision required for SD1/SD2 is not attained. Applications using PPG input must restrict themselves to the time- and frequency-domain subset and report the restriction.

State sequence. Fixed order; order effects are accepted in exchange for standardisation and cross-site comparability.

#	State	Duration	Axis probed
1	Seated rest, spontaneous breathing	5 min	Baseline configuration
2	Paced breathing, 6 breaths/min	5 min	Vagal responsivity
3	Active standing (orthostasis)	3 min	Sympathetic activation
4	Standardised cognitive load	3 min	Central-autonomic coupling
5	Seated recovery	5 min	Return capacity

Total $\approx$ 21 min plus transitions. Where session length is constrained, states 1, 2, 3 and 5 constitute the **minimum viable protocol**. State 2 is never omitted: it is the primary probe of phasic responsivity and, in simulation, the single largest contributor to state separation in the intact archetype.

Phase markers are mandatory. Every recording must carry timestamped transition markers. A continuous recording without markers cannot be segmented post hoc without circularity — inferring state boundaries from the signal and then evaluating state separation in that same signal is not a valid procedure. Recordings lacking markers are unusable regardless of length or quality. This is stated emphatically because the resulting data loss is total and irreversible.

Artefact handling. Malik rule filtering (successive RR differences exceeding 20% flagged); cubic spline interpolation of isolated corrections; exclusion of any state segment with >5% corrected beats; exclusion of any session losing more than one state.

Covariate documentation. Respiratory rate in spontaneous states, time of day, time since last meal, caffeine and nicotine within 4 h, medication (particularly beta-blockers, anticholinergics, SSRIs), habitual physical activity, BMI, age, sex.

5. Domain R: repertoire dimensionality (RRI-R)

5.1. Principle and design decision

RRI-R estimates the effective number of regulatory strategies a person draws on. A deliberate decision was made **not to construct a new self-report instrument**, but to define RRI-R as a *derived index computable on existing multidimensional instruments*.

The rationale is threefold. It eliminates item development and novel-scale validation, the most expensive component of instrument work. It permits immediate use of existing normative data. And — decisively for a construct seeking adoption — it allows any researcher holding a suitable dataset to compute the index retrospectively. A construct testable on data people already hold will be tested; one requiring a proprietary questionnaire will not.

5.2. Instrument requirements

Any instrument may serve as input provided it satisfies:

1. **$J \geq 6$ subscales** representing distinct strategies or modes, not severity gradations of one construct.
2. **Subscales scored as usage or endorsement**, not as pathology. A profile of symptoms is not a repertoire.
3. **Comparable subscale metrics** — equal item counts and response ranges, or documented normalisation.
4. **Published subscale reliabilities**, reported alongside any index computed.

Reference instruments. The Brief COPE^[12] — 28 items, 14 subscales, freely available, German validation available — and the CERQ^[13] — 9 subscales — both satisfy these requirements and serve as reference implementations. The DERS and circumplex interpersonal instruments of the IIP family are plausible additional inputs subject to licensing.

Instruments that do not qualify, and why: symptom inventories such as the PCL-5, PHQ-9 or IDAS-II are severity profiles, not strategy profiles. Using one would render RRI-R operationally indistinguishable from symptom severity, destroying the discriminant validity on which the framework depends.

5.3. Computation

Step 1 — Norm referencing. Each subscale score is z-transformed against published norms.

Step 2 — Non-negativity transform. Hill numbers require non-negative input:

$$u_j = \max(0, z_j + c)$$

with $c = 3$ fixed a priori, a floor three standard deviations below the normative mean. The choice of c is consequential — larger values compress between-person differences, smaller values discard information — and must be fixed in advance and reported, never tuned. Sensitivity analysis across $c \in \{2, 3, 4\}$ should accompany any first application.

Step 3 — Normalisation. $p_j = u_j / \sum u_j$.

Step 4 — Index.

$$\text{RRI-R} = D^2(p) = 1 / (\sum_j p_j^2), \text{ bounded above by } J$$

Step 5 — Total engagement. Report $R_{\text{total}} = \sum u_j$ alongside RRI-R as the repertoire-domain analogue of excursion. As in domain D , breadth and magnitude are reported separately.

5.4 Implementation note

Most clinical assessment platforms compute and store instrument *totals* with severity bands. RRI-R requires the **subscale vector**. Any system intending to support this framework must expose subscale-level output — a one-off but non-trivial engineering requirement, noted here because it is a foreseeable obstacle to adoption and because subscale profiles are clinically more informative than totals in any case.

6. Cross-domain indices

With both indices on a common scale, two further quantities follow. Because the domains have different ceilings, each is first expressed as a proportion of its own maximum:

$$d = \text{RRI-D} / \min(M, K-1) \text{ and } r = \text{RRI-R} / J$$

Coherence: $\text{RRI-C} = 1 - |d - r|$, approaching 1 when physiological and self-reported dimensionality agree.

Discrepancy: the signed difference $\delta = d - r$ carries potential clinical information neither index carries alone:

- $\delta \gg 0$ — physiologically flexible, subjectively constricted. Possible over-report of rigidity, or a self-model lagging behind recovered capacity.
- $\delta \ll 0$ — physiologically collapsed, subjectively flexible. Possible alexithymic or dissociative presentation, or defensive self-description.

These interpretations are **hypotheses, not findings**. They are stated because they constitute the framework's most clinically distinctive prediction and must therefore be exposed to test.

7. Reporting standard

Any application should report, minimally: M, K, and the resulting ceiling; the state sequence and durations; sensor model and sampling rate; artefact criteria and proportion of segments excluded; the feature set and any deviation from Section 4.2; the repertoire instrument, its J, and its subscale reliabilities; the value of c; the Hill order q; and both dimensionality and magnitude indices in each domain.

Indices reported without their ceilings are uninterpretable.

8. Synthetic validation of the pipeline

8.1. Purpose and limits

The demonstration below establishes that the computational pipeline is sound and that RRI-D carries information unavailable to resting-state HRV. It does **not** establish that the index measures dimensional collapse in human beings. The data are generated by a model the author designed, and any generative model will tend to be legible to the analysis built alongside it. This section is a proof of computability, not evidence of validity. Section 10 specifies what evidence of validity would require.

8.2. Method

RR interval series were synthesised from a generative model with controlled mean interval, low-frequency oscillation amplitude, respiratory oscillation amplitude and frequency, and broadband beat-to-beat noise. Features were then extracted from the resulting synthetic signal by the same pipeline that would process real recordings — not assigned directly — so the demonstration tests the extraction and estimation chain end to end. A reference sample of 80 simulated recordings provided normative means and standard deviations. Three archetypes were constructed over $K = 5$ states, each evaluated across 30 independent realisations.

The phasic-collapse archetype was calibrated so that its state parameters lie along a **single direction** in parameter space, and its excursion scale was tuned to match the intact archetype. This produces the

strong form of the claim: two organisms identical on resting HRV *and* on total excursion, differing only in the geometry of their reorganisation.

8.3. Results

Principal comparison (full six-feature set, $K = 5$, nominal ceiling 4.0; 30 realisations):

Archetype	Resting RMSSD	Excursion E	RRI-D	95% interval
A — intact	29.6 ms	8.91	1.55	[1.44, 1.68]
B — phasic collapse	29.6 ms	8.90	1.08	[1.06, 1.11]
C — tonic collapse	10.2 ms	2.39	1.29	[1.13, 1.70]

A and B are **identical on resting RMSSD and on total excursion**. B is not hypo-reactive; it would pass any conventional resting-state HRV screen and any measure of reactivity magnitude. The archetypes are nonetheless separated by RRI-D with non-overlapping 95% intervals. This is the specific empirical gap the framework claims to fill, and the pipeline does fill it under controlled conditions.

C is distinguishable from B on magnitude but shows *higher* dimensionality per unit excursion — the pattern the two-form distinction of Section 2.2 predicts.

Figure 3. Intact and phasic-collapse archetypes are identical on resting HRV and on excursion, and are separated only by dimensionality




Figure 2. Resting RMSSD, excursion, and RRI-D across the three archetypes. The intact and phasic-collapse archetypes are indistinguishable on the first two quantities and separated only by the third. Error bars are 95% intervals across 30 realisations; the dashed line marks the nominal ceiling.

Figure 2. State configurations in principal-component space (single realisation, $K = 5$)




Figure 3. State configurations projected onto their first two principal components, single realisation. The intact archetype occupies a two-dimensional region; the phasic-collapse archetype traces a line, reorganising along one direction only; the tonic-collapse archetype is compressed in all directions. Shaded regions are 2 SD covariance ellipses.

Independence from resting HRV. Across a mixed synthetic sample ($n = 45$), RRI-D correlated with resting RMSSD at $r = 0.21$, well below the $|r| > 0.7$ threshold at which the index would carry no independent information (Section 9.1).

Feature set comparison (intact vs excursion-matched phasic; larger difference is better):

Feature set	D (intact)	D (phasic)	Difference
All six	1.55	1.08	0.47
Drop ln HF	1.65	1.09	0.56
Drop SampEn	1.51	1.09	0.42
Core four (ln RMSSD, ln LF, mean RR, SD1/SD2)	1.88	1.10	0.79
ln RMSSD, ln LF, mean RR	1.96	1.09	0.87

Removing the two redundant or noise-dominated features nearly doubled discrimination. This is the basis for the revised specification in Section 4.2.

Figure 4. Redundant features reduce measured dimensionality and weaken discrimination




Figure 4. Left: feature correlations in the reference sample; ln RMSSD and ln HF are near-redundant at $r = 0.92$. Right: discrimination between the intact and phasic archetypes by feature set. Removing redundant and noise-dominated features nearly doubles separation.

Feature space structure. The six-feature set had an effective dimensionality of 3.44 in the reference sample, with ln RMSSD and ln HF correlating at $r = 0.92$. The feature set, not the state count, is the

binding constraint (Section 4.5).

8.4. *What the simulation revealed that theory did not*

Three of the findings above were not anticipated when the framework was first specified, and two of them contradict its initial recommendations: the six-feature set was worse than a four-feature subset; the K-1 ceiling was not the operative constraint; and whitening, which appeared well motivated, inverted the signal (Section 4.4).

The absolute values also deserve comment. An intact organism reaching 1.88 against a nominal ceiling of 4.0 means the index operates over a compressed range. Whether this compression is a property of the generative model or of the construct cannot be determined synthetically, and it is recorded as a limitation (Section 11.1) rather than resolved.

9. Falsifiable predictions

9.1. *Double dissociation*

Comparing individuals with PTSD, individuals with ME/CFS, and matched controls:

	Resting excursion	RRI-D	State separation (rest ↔ paced)
Controls	normal	high	high
PTSD	normal to elevated	low	low
ME/CFS	low	intermediate	low

PTSD is predicted to show **preserved excursion with reduced dimensionality**; ME/CFS **reduced excursion** with dimensionality less severely affected relative to its excursion. If both groups depart from controls in the same direction on both quantities, the two-form distinction is not supported.

Falsification conditions. The framework is disconfirmed if: (a) RRI-D fails to discriminate PTSD from controls where resting RMSSD does not; (b) RRI-D correlates with resting RMSSD at $|r| > 0.7$, indicating no independent information; (c) the PTSD and ME/CFS profiles do not dissociate as specified.

9.2. Discriminant requirements

RRI-R must not reproduce symptom severity: $|r|$ with PHQ-9 and PCL-5 should be moderate at most ($|r| < 0.5$), and RRI-D should retain incremental variance in group classification after adjustment for depression severity, trauma load, age, sex, BMI and physical activity. Failure would indicate the framework has produced an expensive proxy for existing measures.

9.3. Reliability requirements

Test-retest over 2–4 weeks in a stable subsample: $ICC \geq 0.70$ for RRI-D. Below this the index is not usable at the individual level regardless of group-level discrimination. Given the compressed range observed in Section 8, this is a demanding threshold and its failure is a realistic outcome.

10. Planned validation

Design. Cross-sectional, three groups (PTSD, ME/CFS, controls), no intervention.

Sample. Target $N \geq 240$ with balanced groups; test-retest subsample $n = 50$.

Diagnostic anchoring. Structured interview (CAPS-5 or equivalent) for PTSD; established consensus criteria with independent clinical confirmation for ME/CFS.

Convergent and discriminant battery. Trauma load (CTQ, ACE), symptom severity (PCL-5, PHQ-9, PSS-10), dissociation (DES-II), affect (PANAS), plus the repertoire instrument.

On the ME/CFS group. This population has legitimate and well-founded reasons for scepticism toward psychological framings of their illness. The framework's position must be unambiguous: *tonic collapse is proposed as a description of a physiological state, not a claim about aetiology, and specifically not a claim that ME/CFS is caused by trauma or is psychological in origin*. The prediction that ME/CFS and PTSD dissociate is precisely a prediction that they are **not** the same thing. Patient community review of all study materials prior to submission is treated as a methodological requirement rather than a courtesy, and critical feedback should be expected to alter the protocol.

Governance. Prospective ethics approval; OSF pre-registration before any data collection; publication of scoring code, extraction pipeline and de-identified index-level data.

11. Limitations

11.1. The index operates over a compressed range. In simulation the intact archetype reached 1.88 against a nominal ceiling of 4.0. If this holds in human data, between-person differences will be small relative to measurement error, with direct consequences for the reliability threshold of Section 9.3.

11.2. Eigenvalue estimation from few observations is unstable. Estimating covariance from $K = 5$ observations is statistically demanding. Shrinkage estimation should be evaluated, and bootstrap confidence intervals for individual RRI-D values are strongly preferred over point estimates.

11.3. The feature set is provisional. The core set of Section 4.2 was selected on synthetic data. Feature correlations in human recordings will differ, and the selection must be repeated empirically before the specification is treated as settled.

11.4. Values are not portable between studies. Because no multi-state reference norms exist, RRI-D must be standardised against a locally collected comparison sample (Section 4.3). Two studies reporting an RRI-D of 1.6 are not thereby reporting the same thing. Until a shared reference dataset is established, the index supports within-study contrasts only, and cross-study synthesis — including meta-analysis — is not possible.

11.5. Respiration confounds the physiological domain. HF power in particular is strongly influenced by respiratory rate and depth. State 2 controls this by design; states 1, 4 and 5 do not. Respiratory rate should be measured, not assumed, wherever feasible.

11.6. Diversity indices on subscale scores have known problems. Subscales differ in reliability, so an apparently even profile may partly reflect differential measurement error. Ceiling and floor effects distort profile shape. Ipsative normalisation induces artificial negative dependencies between components. Existing instruments were not designed for this use. These objections are real and the framework does not dissolve them; it requires that subscale reliabilities accompany every RRI-R value and that sensitivity analyses across c and q be conducted.

11.7. Cross-domain convergence may fail. If RRI-D and RRI-R do not correlate, the domain-generality claim of Section 2.3 is unsupported, and the honest conclusion would be that physiological and psychological constriction are distinct constructs sharing a metaphor rather than a mechanism. This is not a remote possibility and the author regards it as a live outcome. It is stated here so that it cannot later be reframed as an incidental finding.

11.8. Synthetic validation is weak evidence. Section 8 demonstrates computability, not validity. A generative model designed by the same author as the analysis will tend to be legible to it.

11.9. Directionality is unknown. Cross-sectional design cannot establish whether reduced dimensionality precedes, follows, or merely accompanies clinical presentation.

11.10. Order effects are accepted, not eliminated. The fixed state sequence trades internal validity for standardisation. A randomised-order substudy would be informative.

12. Conclusion

Dimensional collapse of response space has until now been a description. This paper proposes what is required to make it a measurement: a single statistic — the Hill number of order 2 — applied to normalised non-negative vectors in two domains, yielding effective dimensionality estimates on a common interpretable scale, together with the acquisition protocol, feature specification and reporting standard needed to make values comparable across sites.

Synthetic testing confirms the pipeline separates organisms that resting HRV and reactivity magnitude cannot distinguish, and simultaneously corrected two of the framework's initial recommendations. That the specification changed under testing is reported as a feature of the process rather than concealed as an embarrassment.

The framework is designed to be adoptable without permission. The repertoire index is computable on instruments already in wide use, so existing datasets can be reanalysed. The physiological index requires only a validated chest-strap sensor and a protocol change of approximately fifteen minutes.

Whether the construct survives contact with human data remains open, and Section 9 specifies in advance what would count as failure. That is the point. A construct that cannot fail cannot be evidence for anything.

Notes

This article is posted as a preprint and has not undergone peer review at the time of posting. Critical feedback is actively sought, in particular on three points: the stability of eigenvalue-based dimensionality estimates from $K = 5$ observations (Sections 4.3 and 11.2); the validity of applying diversity indices to questionnaire subscale profiles (Sections 5.3 and 11.6); and the framing of the ME/CFS prediction, on

which input from ME/CFS researchers and from the patient community is especially welcome (Sections 9.1 and 10).

Statements and Declarations

Funding

No specific funding was received for this work.

Potential Competing Interests

The author develops and teaches a coaching methodology that employs HRV-based assessment, and derives income from training and licensing activities related to it. This constitutes a potential non-financial and financial competing interest with respect to a framework for HRV-based measurement. To mitigate it: no claim regarding the efficacy of that or any other intervention is made in this article; the framework is instrument- and method-agnostic by design; the falsification conditions in Section 9 and the limitations in Section 11 were specified so as to permit disconfirmation of the author's own construct; and all code is provided so that every reported result can be independently reproduced.

Ethics

Not applicable. This article reports no research involving human or animal subjects. All data reported in Section 8 are synthetic, generated by the simulation code provided as supplementary material. The validation study outlined in Section 10 has not been conducted; it will require prospective ethics approval before any data collection.

Data Availability

No human or empirical data are reported. All results in Section 8 derive from synthetic data. The complete simulation code, feature extraction pipeline, RRI-D and RRI-R scoring implementations, and figure generation scripts are provided as supplementary data under a CC0 licence, and are sufficient to reproduce every number and figure in this article. Random seeds are fixed in the code, so results are exactly reproducible.

Author Contributions

Daniel Borschel: Conceptualization, Methodology, Formal analysis, Software, Investigation, Writing – original draft, Writing – review & editing, Visualization.

Use of Generative AI

In accordance with Qeios policy and WAME/ICMJE guidance, the author discloses the following. Claude (Anthropic) was used for: literature retrieval and verification of bibliographic details; implementation and execution of the simulation, diagnostic and figure-generation code reported in Section 8 and supplied as supplementary data; and drafting and language editing of the manuscript. The conceptual framework, the construct of dimensional collapse, the decision to adopt Hill numbers as the unifying statistic, the specification choices, and all interpretation of results are the author's. No figures were generated or altered by generative AI: all figures are plots produced by the supplied Python code from the synthetic data described in Section 8. The author has reviewed and verified all analyses, has checked all references against primary sources, and takes full responsibility for the content of this article. Generative AI is not listed as an author and does not qualify for authorship.

Footnotes

¹ $1/(1-q)$

References

1. [△]Borschel DA (2026). "Trauma as a Dimensional Collapse of Response Space: A Quantum-Like Framework with Testable Predictions." *Qeios*. doi:[10.32388/8udzn6](https://doi.org/10.32388/8udzn6).
2. [△]Thayer JF, Lane RD (2000). "A Model of Neurovisceral Integration in Emotion Regulation and Dysregulation." *J Affect Disord*. 61(3):201–216. doi:[10.1016/s0165-0327\(00\)00338-4](https://doi.org/10.1016/s0165-0327(00)00338-4).
3. [△]Kim HG, Cheon EJ, Bai DS, Lee YH, Koo BH (2018). "Stress and Heart Rate Variability: A Meta-Analysis and Review of the Literature." *Psychiatry Investig*. 15(3):235–245. doi:[10.30773/pi.2017.08.17](https://doi.org/10.30773/pi.2017.08.17).
4. [△]Bonanno GA, Burton CL (2013). "Regulatory Flexibility: An Individual Differences Perspective on Coping and Emotion Regulation." *Perspect Psychol Sci*. 8(6):591–612. doi:[10.1177/1745691613504116](https://doi.org/10.1177/1745691613504116).
5. [△]Kashdan TB, Rottenberg J (2010). "Psychological Flexibility as a Fundamental Aspect of Health." *Clin Psychol Rev*. 30(7):865–878. doi:[10.1016/j.cpr.2010.03.001](https://doi.org/10.1016/j.cpr.2010.03.001).

6. [△]Barrett LF, Gross J, Christensen TC, Benvenuto M (2001). "Knowing What You're Feeling and Knowing What to Do About It: Mapping the Relation Between Emotion Differentiation and Emotion Regulation." *Cogn E mot*. 15(6):713–724. doi:[10.1080/02699930143000239](https://doi.org/10.1080/02699930143000239).
7. [△]Hill MO (1973). "Diversity and Evenness: A Unifying Notation and Its Consequences." *Ecology*. 54(2):427–432. doi:[10.2307/1934352](https://doi.org/10.2307/1934352).
8. [△]Jost L (2006). "Entropy and Diversity." *Oikos*. 113(2):363–375. doi:[10.1111/j.2006.0030-1299.14714.x](https://doi.org/10.1111/j.2006.0030-1299.14714.x).
9. [△]Billman GE (2013). "The LF/HF Ratio Does Not Accurately Measure Cardiac Sympatho-Vagal Balance." *Front Physiol*. 4:26. doi:[10.3389/fphys.2013.00026](https://doi.org/10.3389/fphys.2013.00026).
10. [△]Nunan D, Sandercock GRH, Brodie DA (2010). "A Quantitative Systematic Review of Normal Values for Short-Term Heart Rate Variability in Healthy Adults." *Pacing Clin Electrophysiol*. 33(11):1407–1417. doi:[10.1111/j.1540-8159.2010.02841.x](https://doi.org/10.1111/j.1540-8159.2010.02841.x).
11. [△]Voss A, Schroeder R, Heitmann A, Peters A, Perz S (2015). "Short-Term Heart Rate Variability—Influence of Gender and Age in Healthy Subjects." *PLoS ONE*. 10(3):e0118308. doi:[10.1371/journal.pone.0118308](https://doi.org/10.1371/journal.pone.0118308).
12. [△]Carver CS (1997). "You Want to Measure Coping But Your Protocol's Too Long: Consider the Brief COPE." In *J Behav Med*. 4(1):92–100. doi:[10.1207/s15327558jbm04016](https://doi.org/10.1207/s15327558jbm04016).
13. [△]Garnefski N, Kraaij V, Spinhoven P (2001). "Negative Life Events, Cognitive Emotion Regulation and Emotional Problems." *Pers Individ Differ*. 30(8):1311–1327.

Declarations

Funding: This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors. All work reported here — including the development of the framework, the implementation of the analysis code, and the preparation of the manuscript — was carried out by the author independently and funded from the author's own resources. No external body was involved in the design of the framework, the conduct or interpretation of the simulations, the writing of the article, or the decision to submit it for publication. The author intends to seek funding for the validation study outlined in Section 10; no such funding has been applied for, awarded, or received at the time of writing.

Potential competing interests: The author is an independent researcher and practitioner who develops and teaches Ultimate Code, a coaching methodology that incorporates HRV-based assessment, and who derives income from training, licensing, and the associated software platform used by a network of practitioners. Since this article proposes a measurement framework applicable to that field, this

constitutes a potential financial and non-financial competing interest. The author also has an intellectual stake in the underlying construct, having introduced it in a prior preprint (Qeios, DOI 10.32388/8UDZN6), and anticipates seeking research funding for the validation study outlined in Section 10. The pre-post null result cited in Section 1.3 (Hedges' $g = -0.31$) derives from the author's own practice data, and its interpretation here as a limitation of resting-state endpoints rather than as evidence about any intervention reflects an interpretive position the author has an interest in. Several features of this article are intended to offset these interests: no claim of efficacy for any intervention is made; the framework is instrument- and method-agnostic; falsification conditions are specified in advance (Section 9); the limitations section states explicitly what would count as failure of the construct (Section 11); and the complete analysis code is provided so that every reported result can be independently reproduced. The author has no financial or other relationship with any manufacturer of heart rate sensors, including those named or referenced in this article, and received no support of any kind from any commercial entity. The author holds no institutional affiliation; this work was conducted independently and without external oversight.