The Attralucian Essays:
Exploring the Finite

First Edition

 Typeset in LaTeX

The Attralucian Essays

Finite Process Unfolding
A Method for Recovering Temporal Structure from Static Symbolic Forms

Kevin R. Haylett

This essay applies the method of Finite Process Unfolding (FPU) as developed in *Finite Process Unfolding: A Method for Recovering Temporal Structure from Static Symbolic Forms* (Haylett, 2026). FPU is taken here as an established instrument; Bayesian inference is its specimen.

Overview

Bayesian inference is one of the most powerful frameworks in the history of quantitative reasoning. Yet its standard symbolic presentation systematically conceals the sequential experimental structure from which it is inseparable. The canonical expression

$$P(H \mid D) = \frac{P(D \mid H) \cdot P(H)}{P(D)}$$

appears as a static relation between four probability values. It offers no indication that two distinct finite experiments must be carried out in a determinate order, that a physical record must persist between them, or that the hypothesis space must be bounded before the first experiment begins.

This essay applies *Finite Process Unfolding* (FPU) to Bayesian inference. We show that the pipe symbol |—read conventionally as "given"—is not a logical separator between simultaneous quantities. It is a *temporal arrow with a memory cost*: a marker of sequential experimental commitment. The prior is not a belief; it is the *output of a first finite measurement act*, stored and carried forward. The posterior is not a correction; it is the *result of a second finite measurement act* performed on a state space that the first act has already contracted.

The broader claim is that Bayesian inference, properly unfolded, is a trajectory in a finite hypothesis manifold. Each updating step is a geometric contraction. Convergence is basin attraction. The notation hides this. FPU recovers it.

Why Bayes Is Hard to Understand

The pedagogical difficulty of Bayesian statistics is well documented. Students and practitioners trained in frequentist methods routinely misinterpret posterior probabilities. The posterior is read as a frequency, or the prior is treated as arbitrary, or the conditionality is interpreted synchronically rather than sequentially. These errors are not failures of intelligence. They are *failures induced by the notation.*

The standard presentation offers the Bayes equation as a relation between probabilities, all of which appear to coexist at the same moment. The word "conditional" is introduced as a modifier—"the probability of H, *given* D"—without indicating that "given" encodes a temporal commitment: D must have been *produced, recorded, and stored* before the conditioned quantity is defined.

This is a failure of symbolic compression. The equation compresses two sequential experiments and the memory operation between them into a single static expression. What is lost in that compression is precisely what makes Bayesian reasoning difficult to grasp: its inherently diachronic structure. FPU provides the diagnostic instrument. We proceed to apply it.

The Symbolic Frame

We follow the FPU convention: all analysis takes place within the symbolic frame. No appeal is made to entities beyond the boundary at which symbols are formed. Each probability value is treated as the outcome of a finite measurement act—a generonic event—embedded within a bounded hypothesis container.

The hypothesis space $\mathcal{H} = \{H_1, H_2, \ldots, H_n\}$ must be enumerated and bounded *before* the first experiment. This is not a technical convenience. It is an admissibility condition: without a finite, closed container, the first measurement act has no definite output, and the second experiment has no defined state to inherit.

FPU Applied: Unfolding the Bayes Equation

Step 1: Identify the Compression Boundary

The symbolic object under analysis is:

$$S \; = \; P(H \mid D) = \frac{P(D \mid H) \cdot P(H)}{P(D)}$$

This expression is the compression boundary. Behind it lie two experiments, a storage operation, and a normalisation procedure. We treat S as a compressed trace and reconstruct the admissible generating procedure $\mathcal{P}(S)$.

Step 2: Classify the Compression Type

S is a *sequential composition* compressed into a static ratio. Its components are:

- $P(H)$: output of Experiment 1 (prior measurement)
- $P(D \mid H)$: output of Experiment 2, conditioned on stored intermediate state
- $P(D)$: marginalisation over all admissible hypotheses
- $P(H \mid D)$: final contracted state

The pipe symbol $|$ appears twice in S with different meanings, both suppressed by the notation: once as a *conditioning on a stored result* (in $P(D \mid H)$), and once as a *temporal sequencing marker* (in $P(H \mid D)$). FPU requires these to be made explicit.

Step 3: Reconstruct the Sequential Procedure

The admissible generating procedure $\mathcal{P}(S)$ is as follows:

Experiment 1 — Prior Measurement Act

1. Bound and enumerate the hypothesis space: $\mathcal{H} = \{H_1, \ldots, H_n\}$. This act has a cost: each distinction $H_i \neq H_j$ requires a minimum measurement operation $\Delta M > 0$.

2. Assign a prior probability $P(H_i)$ to each hypothesis through a finite measurement or estimation procedure. This may be empirical, theoretical, or conventional, but it must be *finite and completable.* The output is a distribution over $\mathcal{H}$.

3. Store the result. The prior distribution is a physical record: it must persist between experiments. This is the intermediate state σ_1.

Memory Commitment — Storage of σ_1

The intermediate state $\sigma_1 = \{P(H_i)\}_{i=1}^{n}$ must be retained with sufficient fidelity to condition Experiment 2. This is not a passive act. It imposes a minimum storage cost proportional to the cardinality of $\mathcal{H}$ and the precision of each $P(H_i)$. Without this storage, the second experiment cannot proceed.

Experiment 2 — Likelihood Measurement Act

1. Retrieve the stored state σ_1.

2. Conduct the second experiment, generating data D. For each hypothesis $H_i \in \mathcal{H}$, evaluate $P(D \mid H_i)$: the probability of observing D if H_i were true. This is a conditional evaluation *tethered to* σ_1: the hypothesis space was fixed in Experiment 1 and cannot be revised here.

3. Compute the marginal $P(D) = \sum_i P(D \mid H_i)\, P(H_i)$. This is a normalisation over the same bounded container established in Experiment 1.

4. Apply Bayes' rule to contract the state:

$$P(H_i \mid D) = \frac{P(D \mid H_i)\, P(H_i)}{P(D)}$$

5. The output is the posterior distribution $\sigma_2 = \{P(H_i \mid D)\}_{i=1}^{n}$: the contracted state of the hypothesis manifold following both experiments.

Step 4: Identify State Requirements

The procedure $\mathcal{P}(S)$ requires retention of:

- The hypothesis enumeration $\mathcal{H}$ (established before Experiment 1, inherited by Experiment 2)
- The prior distribution σ_1 (output of Experiment 1, input to Experiment 2)
- The intermediate likelihood values $P(D \mid H_i)$ during Experiment 2
- The marginal $P(D)$ as a normalisation constant

None of these requirements appear in the compressed expression S.

Step 5: Estimate Process Cost

A lower bound on the process cost is:

$$\text{Cost}(\mathcal{P}(S)) \;\geq\; (n + 1 + n + 1) \cdot \Delta M \;=\; (2n + 2) \cdot \Delta M$$

where $n = |\mathcal{H}|$ is the cardinality of the hypothesis space, and ΔM is the minimum cost per distinction. The four terms correspond to: bounding $\mathcal{H}$ (n distinctions), completing Experiment 1 (1 act), evaluating likelihoods (n acts), and computing the marginal (1 act). This cost scales linearly with the hypothesis space and is invisible in S.

Step 6: Test Admissibility

The procedure $\mathcal{P}(S)$ is admissible under finite constraints provided:

- $\mathcal{H}$ is finite and enumerable before Experiment 1
- $P(H_i) > 0$ for at least one H_i (non-vacuous container)
- Storage of σ_1 is feasible within the available symbolic container
- Experiment 2 is independent of Experiment 1 except through the inherited container $\mathcal{H}$ and stored state σ_1

A key inadmissibility condition is the *improper prior*: a prior distribution over an unbounded hypothesis space (e.g., a uniform distribution over all real numbers). FPU identifies this not as a technical inconvenience but as a *container violation*: the first experiment has no finite boundary and therefore no admissible output. Experiment 2 cannot proceed.

Step 7: Restate the Expression

$P(H \mid D)$ is a compressed trace of the following:

> A finite hypothesis container was established and measured (*Experiment 1*); the result was stored as an intermediate state (*memory com-*

> *mitment*); a second independent experiment was conducted within the same container, tethered to the stored state (*Experiment 2*); the hypothesis manifold was contracted by the joint evidence of both experiments (*posterior*).

The pipe | in $P(H \mid D)$ does not mean "given D simultaneously." It means: *after D was produced and the intermediate state was consulted.* It is a temporal arrow.

The Geometric Picture

The FPU analysis suggests a natural geometric interpretation that the algebraic notation suppresses entirely.

Let the hypothesis space $\mathcal{H}$ define a finite manifold $\mathcal{M}$, where each point corresponds to a hypothesis and position within the manifold is weighted by probability mass. At the outset, $\mathcal{M}$ is a bounded but relatively diffuse region: the prior distribution σ_1 occupies the manifold with broad support.

Experiment 2 applies a *geometric contraction*: the likelihood function $P(D \mid H)$ selects against regions of $\mathcal{M}$ inconsistent with the observed data D. The posterior σ_2 is a contracted state of the same manifold. The process is not a movement from one number to another; it is a *deformation of a bounded space.*

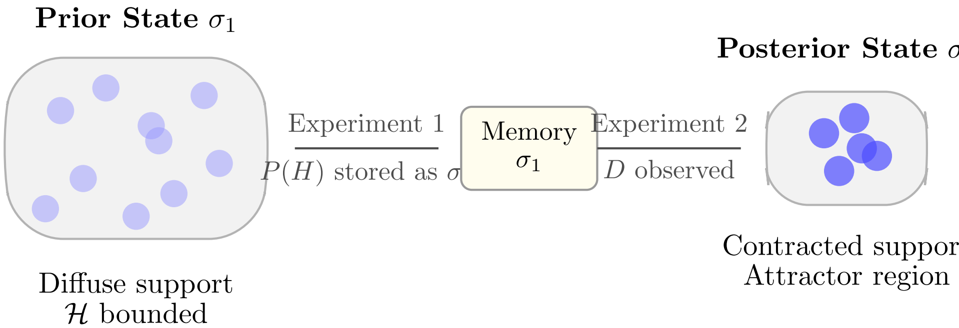

Figure 1: Two sequential experimental acts contracting the hypothesis manifold $\mathcal{M}$. The prior state σ_1 (diffuse support, left) is stored as an intermediate physical record and inherited by Experiment 2. The posterior state σ_2 (contracted support, right) is the geometric result of applying the likelihood $P(D \mid H)$ to the bounded space established by Experiment 1. The compression in the Bayes equation suppresses the entire middle column.

This geometric picture immediately reveals what the algebraic notation conceals:

- The two experiments are *not interchangeable.* Swapping their order does not yield the same posterior, because the container $\mathcal{H}$ and the stored state σ_1 would be different. The sequential structure is not a convenience; it is constitutive.

- The posterior is *not merely a rescaled prior.* It is a geometrically contracted state of the same manifold, produced by a genuinely distinct finite act.

- Improper priors correspond to *unbounded manifolds*: spaces without a finite container. The contraction operation of Experiment 2 has no admissible domain to act upon.

Sequential Updating as Manifold Traversal

The full power of the geometric picture emerges when we consider sequential Bayesian updating: the procedure by which each posterior becomes the prior for a subsequent experiment.

Let σ_0 be the initial prior. After k experiments, the state of the hypothesis manifold is σ_k, related to its predecessor by:

$$\sigma_k(H) \;\propto\; P(D_k \mid H) \cdot \sigma_{k-1}(H)$$

Under FPU, this is not a sequence of algebraic rescalings. It is a *trajectory* in the space of probability distributions over $\mathcal{H}$. Each experiment D_k applies a contraction, deforming the manifold further toward regions consistent with the accumulated evidence.

The trajectory has the following properties that are invisible in the algebraic notation:

- **Ordering matters.** The trajectory $\sigma_0 \to \sigma_1 \to$

$\sigma_2 \to \cdots$ is path-dependent in general when the experiments interact. The intermediate states are not mere computational artifacts; they are the actual history of the process.

- **Convergence is basin attraction.** Under mild conditions, the trajectory converges to a stable region of the manifold regardless of the initial prior σ_0. This is Bayesian consistency, but the geometric picture reveals it as a basin structure: the attractor is the region of $\mathcal{H}$ supported by the data, and the prior determines only the initial position in the basin, not the destination.

- **Memory is structural.** Each σ_k is a compression of the full experimental history $\{D_1, \ldots, D_k\}$. This compression is only admissible if the hypothesis container $\mathcal{H}$ was established at the outset and held fixed throughout the trajectory. Revising $\mathcal{H}$ mid-trajectory is not a Bayesian operation; it is a re-initialisation of the manifold.

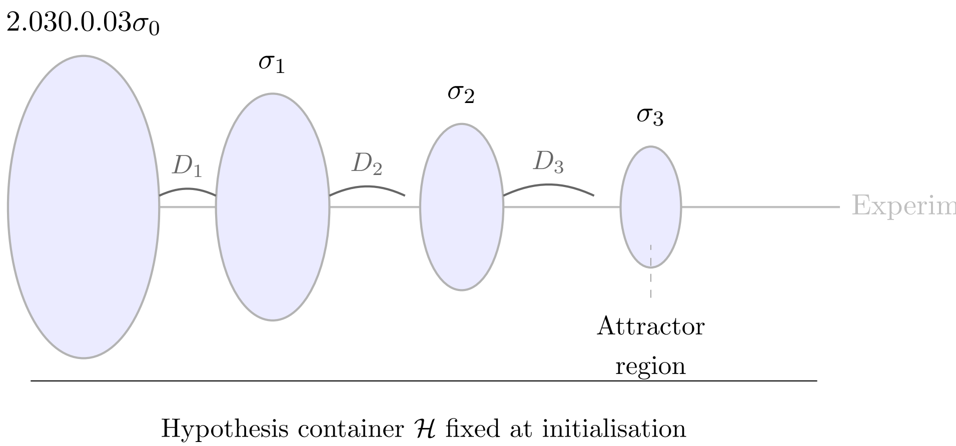

Figure 2: Sequential Bayesian updating as manifold traversal. Each experiment D_k contracts the hypothesis manifold $\mathcal{M}$ further. The trajectory $\sigma_0 \to \sigma_1 \to \sigma_2 \to \sigma_3$ converges toward an attractor region determined by the accumulated evidence. The container $\mathcal{H}$ is fixed throughout; only the distribution over it changes. The algebraic presentation suppresses the trajectory entirely, presenting only the local update rule.

The Pipe Symbol Reconsidered

The FPU analysis yields a precise reinterpretation of the conditional notation. In standard probability theory, $P(A \mid B)$ is defined as:

$$P(A \mid B) = \frac{P(A \cap B)}{P(B)}$$

This definition is synchronic: both A and B are treated as events in the same probability space, and the conditioning is a ratio. There is no time.

FPU reveals that in the Bayesian inference context, this synchronic reading is *a notational fiction*. The quantities $P(H)$ and $P(D \mid H)$ are not co-present values in a static space. They are outputs of distinct finite experimental acts, separated by a memory commitment, with a determinate temporal ordering.

The pipe symbol $\mid$ therefore carries different meanings depending on whether it is read synchronically (as in the ratio definition) or diachronically (as in the experimental procedure). FPU insists on the diachronic reading as the primary one, with the synchronic definition as a legitimate but *process-suppressing* compression.

This has a direct implication for how Bayesian updating should be taught and communicated. The pipe is not a conditioning operator on a static space. It is a *tethering symbol*: it marks the point at which a second experiment is tethered to the stored output of a first.

Connection to Finite Symbolic Mechanics

The FPU analysis of Bayesian inference connects naturally to the broader framework of Finite Symbolic Mechanics (FSM) in the following ways:

Admissibility prior to logic. The hypothesis space $\mathcal{H}$ must be admissible—finite, bounded, enumerable—before any probabilistic reasoning over it can proceed. This is not a consequence of Bayes' theorem but a precondition for its application. FSM places admissibility before logic; the Bayesian case is an illustration of this priority.

Geometric Container Space. The hypothesis manifold $\mathcal{M}$ is a geometric container in the FSM sense: a bounded space within which distinctions are made and measurements are stored. The contraction of $\mathcal{M}$ through updating is a manifestation of Dynamic Flow within a finite container.

Approximations and Measurements. Each probability value in the Bayesian procedure is a finite measurement outcome, not an ideal real number. The precision of $P(H_i)$ is bounded by the measurement process that produced it. FSM requires that all quantities be treated as approximations with finite cost; the Bayesian formalism typically suppresses this.

Finite Reality. The FPU procedure makes explicit that Bayesian inference, properly construed, is a finite process with a finite cost, operating within a bounded container. Its extension to continuous hypothesis spaces, improper priors, and limiting distributions should be treated as approximations to this finite reality, not as the fundamental case.

Discussion

The difficulty of Bayesian statistics is not an accident of pedagogy. It is a consequence of notation that presents a sequential experimental procedure as a static algebraic relation. FPU provides the instrument to diagnose this compression and reverse it.

The central contributions of this analysis are:

1. The pipe symbol $|$ in Bayesian notation is a temporal arrow with a memory cost, not a synchronic conditioning operator.

2. The prior $P(H)$ is the output of a first finite experiment, not a pre-existing quantity. It must be stored before the second experiment can begin.

3. The Bayes equation is a compressed trace of two sequential experimental acts and the memory operation between them.

4. Sequential Bayesian updating is a trajectory in a finite hypothesis manifold, converging toward an attractor region determined by accumulated evidence.

5. Improper priors are inadmissible under FSM: they

correspond to unbounded manifolds without a finite container.

None of these claims conflict with the mathematics of Bayes' theorem. They enrich it by making explicit the process structure that the theorem compresses.

Conclusion

Bayesian inference is not a static relation between four probability values. It is a finite sequential process: a first experiment that bounds and measures a hypothesis space; a memory commitment that stores that measurement; a second experiment that inherits the stored state and contracts the manifold further; and a posterior that is the geometric residue of both acts.

The Bayes equation compresses this process into a single line. FPU unfolds it. What is revealed is not a complication of Bayes but its natural geometry: a trajectory in a finite manifold, driven by evidence, converging to a basin that the data has determined. The symbolic form is not primary. It is a compression of process. And the pipe | is not a given. It is a *then*.

Appendix: FPU Summary Table for Bayesian Inference

FPU Step	Bayesian Interpretation
Compression boundary	$P(H \mid D) = P(D \mid H)\, P(H) \,/\, P(D)$
Compression type	Sequential composition of two experiments
Sequential procedure	Exp. 1 (prior) $\to$ memory $\to$ Exp. 2 (likelihood) $\to$ posterior
State requirements	$\mathcal{H}$, σ_1, $P(D \mid H_i)$, $P(D)$
Process cost	$\geq (2n + 2) \cdot \Delta M$ where $n = \lvert\mathcal{H}\rvert$
Admissibility condition	$\mathcal{H}$ finite; σ_1 stored; Exp. 2 tethered to σ_1
Restated expression	Compressed trace of two sequential contractions of $\mathcal{M}$
Geometric picture	Trajectory in hypothesis manifold; posterior as attractor

Table 1: FPU analysis of Bayesian inference: summary of steps and interpretations.