Contraction Dynamics in Constraint-Shaped Acceptability Distributions

Jeremy McEntire
Cage & Mirror Press
jmc@cageandmirror.com

Abstract

We study contraction dynamics in constraint-shaped acceptability distributions. A constraint type system with four primitives (schema, threshold, scoring, topological) is defined and shown to be decidable and compositionally type-preserving. For Gaussian acceptability distributions under Gaussian scoring with lifecycle variance injection, the Q-dynamics operator is proved to be a global contraction with unique fixed point Q^* and explicit contraction parameter

$$\lambda_* = \frac{\sigma_\kappa^2}{\sigma_*^2 + \sigma_\kappa^2},$$

where σ_κ^2 is the scoring function variance and σ_*^2 is the fixed-point variance of the acceptability distribution. We establish compositional well-formedness: the composition of well-formed constraint architectures is well-formed under type compatibility and energy separability. The cage—the attractor under convergent selection—is characterized as the limit $Q^* \to p_F$ as scoring sharpens, and conditions on selection balance are given under which Q^* preserves generative capacity. Extensions to $d > 1$ via product measures are proved for the commuting case; the non-commuting case and non-Gaussian extensions are catalogued as open problems with conjectured results and bridging strategies.

Keywords: contraction mapping, acceptability distribution, constraint type system, Wasserstein convergence, fixed-point dynamics, compositional well-formedness

1 Introduction

We study the convergence properties of acceptability distributions shaped by constraint architectures. The setting is a system in which agents produce outputs scored by constraint functions, and the distribution of acceptable outputs evolves over time through scoring-selection and lifecycle operations.

Three formal questions arise:

1. When can constraints be composed without circular dependencies?
2. When does composition of well-formed architectures preserve well-formedness?
3. Under what conditions does the acceptability distribution converge to a unique fixed point, and at what rate?

These questions are unified by the observation that constraint typing determines valid compositions, valid compositions preserve well-formedness conditions, and the scoring function's curvature combined with lifecycle variance injection determines whether the dynamics operator contracts.

For the Gaussian case, we derive the contraction parameter in closed form from the scoring variance σ_κ^2 and lifecycle variance σ_L^2. The fixed point, convergence rate, and compositional properties are all explicit. The cage characterization and generative preservation theorems connect the contraction dynamics to organizational phenomena studied in [7, 6]; the systems context is developed there, while the present paper treats the pure mathematics.

Section 2 fixes notation. Section 3 defines the constraint type system. Section 4 proves the main contraction theorem. Section 5 extends to $d > 1$. Section 6 establishes compositional well-formedness. Section 7 characterizes the cage and generative preservation. Section 8 catalogues honest gaps and conjectured extensions.

2 Preliminaries

2.1 Notation

Let $(\Omega, \mathcal{F}, P)$ be a probability space. We write $\mathcal{P}(X)$ for the set of Borel probability measures on a metric space (X, d), and $\mathcal{P}_1(X)$ for those with finite first moment. The Wasserstein-1 distance between $\mu, \nu \in \mathcal{P}_1(X)$ is

$$W_1(\mu, \nu) = \inf_{\gamma \in \Gamma(\mu,\nu)} \int_{X \times X} d(x, y)\, d\gamma(x, y),$$

where $\Gamma(\mu, \nu)$ denotes the set of couplings with marginals μ and ν. The space $(\mathcal{P}_1(X), W_1)$ is complete when (X, d) is complete and separable (Polish) [11]. Transportation-cost inequalities relating W_1 to relative entropy provide concentration bounds that complement the contraction results below [10].

We write $\mathcal{N}(\mu, \sigma^2)$ for the Gaussian distribution with mean μ and variance σ^2, and $\mathcal{N}(\boldsymbol{\mu}, \Sigma)$ for the multivariate Gaussian with mean vector $\boldsymbol{\mu}$ and covariance matrix Σ.

For a positive definite matrix A, we write $\lambda_{\min}(A)$ and $\lambda_{\max}(A)$ for its smallest and largest eigenvalues. The notation $A \succ 0$ means A is positive definite.

A map $T : (M, d) \to (M, d)$ on a complete metric space is a *Banach contraction* with parameter $\lambda \in [0, 1)$ if $d(Tx, Ty) \le \lambda\, d(x, y)$ for all $x, y \in M$. By the Banach fixed-point theorem [2], such T has a unique fixed point x^* and $T^n(x_0) \to x^*$ for any $x_0 \in M$.

2.2 State Space

A constraint architecture $\mathcal{A}$ operates on a state space

$$S = (\text{Agents, Signals, Topology, Energy}).$$

We do not impose further structure on S until the Gaussian specialization in Section 4.

3 Constraint Type System

3.1 Primitive Types

Definition 3.1 (Schema Constraint)**.** A *schema constraint* $\sigma : \text{Signals} \to \{\text{valid}, \text{invalid}\}$ is a total, deterministic, stateless function that partitions the signal space. Schema constraints compose conjunctively: $\sigma_1 \wedge \sigma_2$ is a schema constraint whenever σ_1 and σ_2 are.

Definition 3.2 (Threshold Constraint)**.** A *threshold constraint* $\tau = (f, \theta)$ consists of a scoring function $f : S \to \mathbb{R}$ and a threshold $\theta \in \mathbb{R}$. The constraint gates behavior: an action is permitted when $f(s) \ge \theta$. Threshold constraints are parameterized and stateless.

Definition 3.3 (Scoring Constraint)**.** A *scoring constraint* $\kappa\colon \text{Outputs} \to \mathbb{R}$ assigns a real-valued score to every agent output. A scoring constraint is total (every output receives a score), deterministic (the same output receives the same score), and bounded ($\kappa(o) \in [0, M]$ for some finite M). Scores are used for selection, not gating.

Definition 3.4 (Topological Constraint)**.** A *topological constraint* $\rho\colon \text{Topology} \times \text{Signals} \to \text{Topology}$ defines how the agent topology evolves in response to signals. Topological constraints are the only stateful primitive: they modify the system's structure. A topological constraint is *well-typed* when it satisfies connectivity preservation: for any connected topology T, $\rho(T, s)$ is connected for all signals s.

3.2 Dependency and Compatibility

Definition 3.5 (Dependency Graph)**.** For a set of constraints $C = \{c_1, \ldots, c_n\}$, the *dependency graph* $G(C)$ has vertex set C and a directed edge $c_i \to c_j$ when c_j reads a state component that c_i writes.

Definition 3.6 (Type Compatibility)**.** A constraint set C is *type-compatible* when $G(C)$ is a directed acyclic graph (DAG).

Proposition 3.7 (Decidability)**.** *Type compatibility is decidable in polynomial time. Given n constraints with declared read/write sets, $G(C)$ can be constructed in $O(n^2)$ and cycle-checked in $O(n + |E|)$.*

Proof. Construction of $G(C)$ requires checking each pair of constraints for shared state components: $O(n^2)$ pairs, each checkable in time proportional to the size of the read/write declarations. Cycle detection on the resulting directed graph is linear in the graph size by topological sort. □

3.3 Compound Types

Definition 3.8 (Sequential Composition)**.** For constraints $c_1\colon S \to S$ and $c_2\colon S \to S$ where c_1 writes what c_2 reads, the *sequential composition* $c_1 \,;\, c_2$ applies c_1 first, then c_2 to the resulting state.

Definition 3.9 (Parallel Composition)**.** For constraints c_1 and c_2 that read and write disjoint state components, the *parallel composition* $c_1 \| c_2$ applies both simultaneously. The result is independent of evaluation order.

Definition 3.10 (Conditional Composition)**.** For a threshold constraint τ and constraints c_{true}, c_{false}, the *conditional composition* $\tau \;?\; c_{\text{true}} : c_{\text{false}}$ applies c_{true} when τ is satisfied and c_{false} otherwise.

Proposition 3.11 (Type Preservation under Composition)**.** *Let c_1, c_2 be type-compatible constraints.*

(i) *$c_1 \,;\, c_2$ and $c_1 \| c_2$ are type-compatible with any constraint that was type-compatible with both c_1 and c_2.*

(ii) *For the conditional composition $\tau \;?\; c_1 : c_2$, type compatibility is preserved when additionally τ does not read state that either branch writes.*

Proof. (i) Sequential composition adds an edge $c_1 \to c_2$ to $G(C)$, preserving acyclicity when the original graph plus this edge has no cycle (guaranteed by type compatibility of c_1, c_2). Parallel composition adds no edges (disjoint state).

(ii) Conditional composition adds edges from τ to both branches. The additional hypothesis (no reads from τ on state written by branches) ensures no backward edges from branches to τ, preserving acyclicity. □

4 Gaussian Contraction

4.1 Well-Formedness Conditions

We require four conditions on a constraint architecture $\mathcal{A}$.

Definition 4.1 (Well-Formedness)**.** A constraint architecture $\mathcal{A} = (C, L, E)$ with constraint set C, lifecycle rules L, and energy function E is *well-formed* if it satisfies:

(WF1) **Completeness.** For every agent a and every output o produced by a, there exists at least one scoring constraint $\kappa \in C$ such that $\kappa(o)$ is defined.

(WF2) **Liveness.** The lifecycle rules $L = \{\text{fork}, \text{decay}, \text{merge}\}$ are non-trivially reachable: for each $l \in L$, there exists a reachable state $s \in S$ such that l is triggered at s.

(WF3) **Dissipation.** There exists an energy function $E \colon S \to \mathbb{R}_+$ such that E is bounded, monotonically non-increasing between external input events, and strictly decreasing whenever an agent processes a signal.

(WF4) **Contraction.** The operator $\Phi_{\mathcal{A}}$ (Definition 4.3) is a contraction in W_1:

$$W_1(\Phi_{\mathcal{A}}(Q),\, \Phi_{\mathcal{A}}(Q')) \leq \lambda\, W_1(Q, Q')$$

for some $\lambda \in [0, 1)$ and all Q, Q' in the feasible set.

Definition 4.2 (Acceptability Distribution)**.** For a constraint architecture $\mathcal{A}$ on state space S, the *acceptability distribution* Q_t at time t is

$$Q_t = \frac{1}{|O_t|} \sum_{o \in O_t} \delta_o,$$

where $O_t = \{o : \kappa(o) \geq \theta_{\min} \text{ for all active scoring constraints } \kappa\}$.

Definition 4.3 (Q-Dynamics)**.** The evolution of Q is governed by

$$Q_{t+1} = \Phi_{\mathcal{A}}(Q_t),$$

where $\Phi_{\mathcal{A}}$ is the operator induced by $\mathcal{A}$: given the current acceptability distribution, $\Phi_{\mathcal{A}}$ produces the next-period distribution through scoring-selection and lifecycle operations.

Remark 4.4. Conditions WF1–WF3 can be verified by inspection of the constraint architecture. WF4 requires either analytical derivation of λ from the architecture's properties or empirical estimation. The remainder of this section provides the analytical derivation for the Gaussian case.

4.2 Setup

Let the output space be $\mathbb{R}$ (extension to $\mathbb{R}^d$ in Section 5). Model the acceptability distribution as Gaussian:

$$Q_t = \mathcal{N}(\mu_t, \sigma_t^2).$$

The dynamics have two components:

1. **Scoring-selection step:** reweight Q_t by the scoring function κ, renormalize.
2. **Lifecycle step:** inject variance from fork events, external signal arrival, and agent non-determinism.

4.3 The Scoring-Selection Operator

Model the scoring function κ as Gaussian with precision $1/\sigma_\kappa^2$ centered at μ_*:

$$\kappa(o) = C \exp\left(-\frac{1}{2}\frac{(o-\mu_*)^2}{\sigma_\kappa^2}\right).$$

The reweighted distribution is

$$Q_{\text{score}}(o) = \frac{\kappa(o)\, Q_t(o)}{Z_t}.$$

Since both factors are Gaussian, the product is Gaussian:

$$Q_{\text{score}} = \mathcal{N}(\mu_{\text{score}}, \sigma^2_{\text{score}})$$

where

$$\sigma^2_{\text{score}} = \frac{\sigma_t^2\, \sigma_\kappa^2}{\sigma_t^2 + \sigma_\kappa^2}, \tag{1}$$

$$\mu_{\text{score}} = \frac{\mu_t\, \sigma_\kappa^2 + \mu_*\, \sigma_t^2}{\sigma_t^2 + \sigma_\kappa^2}. \tag{2}$$

4.4 The Lifecycle Operator

Fork, external signals, and agent non-determinism inject variance. Model as additive:

$$\sigma^2_{t+1} = \sigma^2_{\text{score}} + \sigma_L^2, \tag{3}$$

$$\mu_{t+1} = \mu_{\text{score}}, \tag{4}$$

with $\sigma_L^2 > 0$ representing lifecycle variance injection, bounded by dissipation (WF3).

4.5 The Composed Map

Define $\delta = \mu - \mu_*$ (deviation from scoring center). The full update $T\colon (\delta, \sigma^2) \mapsto (\delta', \sigma'^2)$ is:

$$\delta' = \delta \cdot r(\sigma^2), \qquad \text{where } r(\sigma^2) = \frac{\sigma_\kappa^2}{\sigma^2 + \sigma_\kappa^2}, \tag{5}$$

$$\sigma'^2 = h(\sigma^2), \qquad \text{where } h(\sigma^2) = \frac{\sigma^2\, \sigma_\kappa^2}{\sigma^2 + \sigma_\kappa^2} + \sigma_L^2. \tag{6}$$

4.6 The Contraction Theorem

Theorem 4.5 (1D Gaussian Contraction)**.** *For the Q-dynamics map T with Gaussian scoring (variance $\sigma_\kappa^2 > 0$) and lifecycle variance injection ($\sigma_L^2 > 0$):*

(i) *The variance map $h\colon (0,\infty) \to (0,\infty)$ is a global contraction. For all $\sigma^2 > 0$: $h'(\sigma^2) = r(\sigma^2)^2 < 1$.*

(ii) *T has a unique fixed point $(\delta_* = 0,\ \sigma_*^2)$ where*

$$\sigma_*^2 = \frac{\sigma_L^2 + \sqrt{(\sigma_L^2)^2 + 4\,\sigma_L^2\,\sigma_\kappa^2}}{2}.$$

(iii) *For any initial* $\sigma_0^2 > 0$, $\sigma_t^2 \to \sigma_*^2$. *For any initial* δ_0, $\delta_t \to 0$ *at geometric rate*

$$\lambda_* = \frac{\sigma_\kappa^2}{\sigma_*^2 + \sigma_\kappa^2}.$$

(iv) $\lambda_* < 1$ *whenever* $\sigma_\kappa^2 > 0$ *and* $\sigma_L^2 > 0$.

(v) *Explicitly:*

$$\lambda_* = \frac{2\,\sigma_\kappa^2}{\sigma_L^2 + \sqrt{(\sigma_L^2)^2 + 4\,\sigma_L^2\,\sigma_\kappa^2} + 2\,\sigma_\kappa^2}.$$

Proof. **(i)** Compute

$$h'(\sigma^2) = \frac{d}{d\sigma^2}\left[\frac{\sigma^2\,\sigma_\kappa^2}{\sigma^2 + \sigma_\kappa^2} + \sigma_L^2\right] = \frac{(\sigma_\kappa^2)^2}{(\sigma^2 + \sigma_\kappa^2)^2} = r(\sigma^2)^2.$$

Since $\sigma^2 > 0$ and $\sigma_\kappa^2 > 0$, we have $0 < r(\sigma^2) < 1$, so $0 < h'(\sigma^2) < 1$ for all $\sigma^2 > 0$.

By the mean value theorem, for any $\sigma_a^2, \sigma_b^2 > 0$:

$$|h(\sigma_a^2) - h(\sigma_b^2)| = |h'(\xi)|\,|\sigma_a^2 - \sigma_b^2| < |\sigma_a^2 - \sigma_b^2|.$$

Moreover, h maps $(0, \infty)$ into $[\sigma_L^2,\, \sigma_\kappa^2 + \sigma_L^2)$ (since $\sigma^2\sigma_\kappa^2/(\sigma^2 + \sigma_\kappa^2) < \sigma_\kappa^2$ for all $\sigma^2 > 0$). So h maps the compact interval $[\sigma_L^2,\, \sigma_\kappa^2 + \sigma_L^2]$ into itself.

By the Banach fixed-point theorem [2], h has a unique fixed point σ_*^2 in this interval, and $h^n(\sigma_0^2) \to \sigma_*^2$ for any σ_0^2 in the interval. Since h maps all of $(0, \infty)$ into the invariant interval after one step, convergence is global.

(ii) The fixed-point equation $h(\sigma_*^2) = \sigma_*^2$ gives

$$\sigma_*^2 = \frac{\sigma_*^2\,\sigma_\kappa^2}{\sigma_*^2 + \sigma_\kappa^2} + \sigma_L^2.$$

Rearranging: $(\sigma_*^2)^2 - \sigma_L^2\,\sigma_*^2 - \sigma_L^2\,\sigma_\kappa^2 = 0$. The unique positive root is as stated.

(iii) For the mean deviation: $|\delta_t| = |\delta_0| \prod_{k=0}^{t-1} r(\sigma_k^2)$. For $k \geq 1$, $\sigma_k^2 \geq \sigma_L^2$ (since h maps into $[\sigma_L^2,\, \sigma_\kappa^2 + \sigma_L^2)$). Therefore $r(\sigma_k^2) \leq \sigma_\kappa^2/(\sigma_L^2 + \sigma_\kappa^2) < 1$ for all $k \geq 1$. This uniform bound guarantees $|\delta_t| \to 0$ at a rate that converges to $\lambda_* = r(\sigma_*^2)$ as $\sigma_t^2 \to \sigma_*^2$. Combined convergence of $(\delta_t, \sigma_t^2) \to (0, \sigma_*^2)$ follows.

(iv) $\lambda_* = \sigma_\kappa^2/(\sigma_*^2 + \sigma_\kappa^2)$. Since $\sigma_*^2 > 0$ (from the quadratic formula with $\sigma_L^2 > 0$), we have $\sigma_*^2 + \sigma_\kappa^2 > \sigma_\kappa^2$, so $\lambda_* < 1$.

(v) Substituting the expression for σ_*^2 from (ii) into the formula for λ_* from (iii) and simplifying yields the explicit expression. □

Remark 4.6 (Global convergence vs. local stability). The variance map h is a Banach contraction on the invariant interval $[\sigma_L^2,\, \sigma_\kappa^2 + \sigma_L^2]$, proved from $h'(s) < 1$ for all $s > 0$. The full map T is globally convergent but is not a uniform Banach contraction in any single product metric, because the mean contraction rate $r(\sigma_t^2)$ varies with the current variance. The composable contraction parameter is the worst-case rate on the invariant interval, $\lambda_{\text{uniform}} = \sigma_\kappa^2/(\sigma_L^2 + \sigma_\kappa^2)$, which is larger (worse) than the fixed-point rate λ_*.

Table 1: Effect of parameter changes on the contraction rate λ_*.

Parameter change	Effect on λ_*	Interpretation
$\sigma_\kappa^2 \to 0$ (sharper scoring)	$\lambda_* \to 0$	Stronger selection = faster convergence
$\sigma_\kappa^2 \to \infty$ (flatter scoring)	$\lambda_* \to 1$	Weaker selection = slower convergence
$\sigma_L^2 \to 0$ (less lifecycle noise)	$\lambda_* \to 1$	Fixed point approaches degeneracy
$\sigma_L^2 \to \infty$ (more lifecycle noise)	$\lambda_* \to 0$	Faster contraction; Q^* has high variance

4.7 Parameter Dependence

4.8 Convergence of Q-Dynamics

Theorem 4.7 (Convergence—Gaussian). *Let $\mathcal{A}$ be a constraint architecture operating on Gaussian distributions with scoring variance $\sigma_\kappa^2 > 0$ and lifecycle variance injection $\sigma_L^2 > 0$. Then:*

(i) *The Q-dynamics operator $\Phi_\mathcal{A}$ is a global contraction with explicit parameter λ_* given by Theorem 4.5(v).*

(ii) *There exists a unique fixed-point distribution $Q^* = \mathcal{N}(\mu_*, \sigma_*^2)$.*

(iii) *For any initial Q_0, the sequence $Q_t = \Phi_\mathcal{A}^t(Q_0)$ satisfies $W_1(Q_t, Q^*) \to 0$ at geometric rate λ_*.*

Proof. Direct consequence of Theorem 4.5. The map T sends any (μ, σ^2) with $\sigma^2 > 0$ into the invariant set $\{(\mu', \sigma'^2) : \sigma'^2 \in [\sigma_L^2,\, \sigma_\kappa^2 + \sigma_L^2]\}$. This invariant set, equipped with the product metric on $\mathbb{R} \times [\sigma_L^2,\, \sigma_\kappa^2 + \sigma_L^2]$, is complete. The variance map h is a Banach contraction on the invariant interval (Theorem 4.5(i)). The mean map contracts non-uniformly, with rate converging to λ_*. The combined map converges globally by Theorem 4.5(iii). □

Theorem 4.8 (Convergence—General). *Let $\mathcal{A}$ be a constraint architecture satisfying WF1–WF4 with contraction parameter $\lambda < 1$. Then:*

(i) *There exists a unique fixed-point distribution Q^* such that $\Phi_\mathcal{A}(Q^*) = Q^*$.*

(ii) *For any initial distribution Q_0, $W_1(Q_t, Q^*) \le \lambda^t\, W_1(Q_0, Q^*)$.*

(iii) *Convergence is geometric at rate λ.*

Proof. Banach fixed-point theorem on $(\mathcal{P}_1(\text{Output}), W_1)$, which is complete for distributions with finite first moment over a Polish space [11]. □

Remark 4.9. Theorem 4.7 is unconditional: WF4 is derived from the scoring function's curvature and lifecycle's variance injection. Theorem 4.8 is conditional: it requires WF4 as a hypothesis. Theorem 4.7 discharges this hypothesis for the Gaussian case. For non-Gaussian architectures, WF4 must be verified separately.

Corollary 4.10 (Endogenous Q). *The fixed point Q^* resolves the endogenous-Q problem [8]: Q is not exogenous but is determined by the constraint architecture:*

$$Q^* = \lim_{t\to\infty} \Phi_\mathcal{A}^t(Q_0).$$

5 Multi-Dimensional Extension

We extend the contraction result to $\mathbb{R}^d$. Let $Q_t = \mathcal{N}(\boldsymbol{\mu}_t, \Sigma_t)$ with $\Sigma_t \succ 0$, and let the scoring function be $\kappa(\mathbf{o}) \propto \exp(-\frac{1}{2}(\mathbf{o} - \boldsymbol{\mu}_*)^\top \Lambda\, (\mathbf{o} - \boldsymbol{\mu}_*))$ with precision matrix $\Lambda \succ 0$.

5.1 Mean Contraction

The scoring-selection step yields posterior mean

$$\boldsymbol{\mu}_{\text{score}} = (\Sigma_t^{-1} + \Lambda)^{-1}(\Sigma_t^{-1}\boldsymbol{\mu}_t + \Lambda\,\boldsymbol{\mu}_*).$$

Define $\boldsymbol{\delta}_t = \boldsymbol{\mu}_t - \boldsymbol{\mu}_*$. Then $\boldsymbol{\delta}_{t+1} = R(\Sigma_t)\,\boldsymbol{\delta}_t$ where

$$R(\Sigma) = (I + \Sigma\Lambda)^{-1}. \tag{7}$$

The spectral radius satisfies $\rho(R(\Sigma)) = 1/(1+\gamma_{\min})$ where $\gamma_{\min} = \lambda_{\min}(\Sigma^{1/2}\Lambda\,\Sigma^{1/2}) > 0$.

Theorem 5.1 (Mean Contraction in d Dimensions)**.** *For $\Sigma \succ 0$ and $\Lambda \succ 0$, $\|R(\Sigma)\|_2 \leq 1/(1+\gamma_{\min}) < 1$. Hence $\boldsymbol{\delta}_t \to \mathbf{0}$ geometrically.*

Proof. Write $R = (I+\Sigma\Lambda)^{-1}$. Since $\Sigma\Lambda$ has eigenvalues $\gamma_1, \dots, \gamma_d > 0$ (all positive because $\Sigma, \Lambda \succ 0$), R has eigenvalues $1/(1+\gamma_i)$, each in $(0,1)$. The largest eigenvalue of R is $1/(1+\gamma_{\min})$, giving $\|R\|_2 \leq 1/(1+\gamma_{\min}) < 1$. □

5.2 Covariance Contraction

The posterior covariance is $\Sigma_{\text{score}} = (\Sigma_t^{-1} + \Lambda)^{-1}$, and the lifecycle step gives $\Sigma_{t+1} = \Sigma_{\text{score}} + \Sigma_L$ with $\Sigma_L \succ 0$. Define the covariance map

$$H(\Sigma) = (\Sigma^{-1} + \Lambda)^{-1} + \Sigma_L. \tag{8}$$

Theorem 5.2 (Covariance Contraction—Commuting Case)**.** *If Σ and Λ are simultaneously diagonalizable (i.e., $\Sigma\Lambda = \Lambda\Sigma$ at the fixed point), the covariance map H decomposes into d independent scalar maps, each a contraction by Theorem 4.5*(i)*. The fixed point Σ_* exists, is unique, and is globally attracting.*

Proof. In the simultaneous eigenbasis, $\Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_d^2)$ and $\Lambda = \text{diag}(\lambda_1^{-2}, \dots, \lambda_d^{-2})$. The map H acts coordinate-wise: $\sigma_j'^2 = \sigma_j^2\lambda_j^2/(\sigma_j^2 + \lambda_j^2) + (\sigma_L^2)_j$. Each coordinate is a 1D contraction by Theorem 4.5(i). □

Remark 5.3 (Non-commuting case). For general Σ, Λ with $[\Sigma, \Lambda] \neq 0$, the Fréchet derivative of H at Σ_* has spectral radius bounded above by $\lambda_{\text{mean}}^2 < 1$, establishing local stability. Global contraction in the non-commuting case requires monotonicity arguments on the cone of positive definite matrices; this is an open problem (Problem 8.5).

6 Compositional Well-Formedness

Theorem 6.1 (Compositional Well-Formedness)**.** *Let $\mathcal{A}_1 = (C_1, L_1, E_1)$ and $\mathcal{A}_2 = (C_2, L_2, E_2)$ be well-formed constraint architectures satisfying WF1–WF4 with uniform contraction parameters $\lambda_1, \lambda_2 < 1$. Define the composed architecture $\mathcal{A}_{12} = \mathcal{A}_1 \circ \mathcal{A}_2$ with:*

- *Constraint set $C_{12} = C_1 \cup C_2$,*
- *Lifecycle rules $L_{12} = L_1 \cup L_2$,*
- *Energy function $E_{12} = E_1 + E_2$.*

If the constraint sets are type-compatible (Definition 3.6) and the energy functions are additively separable ($E_{12}(s_1, s_2) = E_1(s_1) + E_2(s_2)$), then $\mathcal{A}_{12}$ satisfies WF1–WF4 with contraction parameter $\lambda_{12} = \max(\lambda_1, \lambda_2)$.

Proof. **WF1 (Completeness).** Every output of an agent in $\mathcal{A}_1$ passes through at least one scoring constraint in C_1 (by WF1 for $\mathcal{A}_1$). The union of scoring constraints covers the union of agents.

WF2 (Liveness). Fork, decay, and merge in L_1 are reachable from states in S_1 (by WF2 for $\mathcal{A}_1$). Additive separability implies C_2 does not write to S_1 state components that L_1's triggers depend on, so constraints in C_2 do not block the state transitions that trigger L_1's lifecycle rules. Similarly for L_2.

WF3 (Dissipation). $E_{12} = E_1 + E_2$. Each summand is non-increasing and strictly decreasing on agent action. Additive separability ensures no cross-terms, so E_{12} is strictly decreasing whenever any agent acts.

WF4 (Contraction). For the composed system, the Q-dynamics operator is $\Phi_{\mathcal{A}_2} \circ \Phi_{\mathcal{A}_1}$ (sequential) or $\Phi_{\mathcal{A}_1} \otimes \Phi_{\mathcal{A}_2}$ (parallel, on disjoint output spaces).

Sequential case:

$$\begin{aligned} W_1(\Phi_{\mathcal{A}_2}(\Phi_{\mathcal{A}_1}(Q)),\, \Phi_{\mathcal{A}_2}(\Phi_{\mathcal{A}_1}(Q'))) &\le \lambda_2\, W_1(\Phi_{\mathcal{A}_1}(Q),\, \Phi_{\mathcal{A}_1}(Q')) \\ &\le \lambda_2 \lambda_1\, W_1(Q, Q'). \end{aligned}$$

The parameter $\lambda_2\lambda_1 \le \max(\lambda_1, \lambda_2) < 1$.

Parallel case (disjoint output spaces): W_1 on the product space decomposes. Each component contracts at its own rate. The joint contraction parameter is $\max(\lambda_1, \lambda_2)$. □

7 Cage Characterization

The cage is the regime in which constraints prevent divergent exploration, causing the acceptability distribution to collapse toward the system's internal model. We characterize this mathematically.

7.1 Cage as Attractor

Definition 7.1 (Frame Distribution)**.** For a system with internal model (Frame) F, the *Frame distribution* p_F is the distribution of outputs consistent with F.

Theorem 7.2 (Cage as Attractor—Gaussian)**.** *Under convergent selection—where the scoring function is centered at the Frame mean ($\mu_* = \mu_F$) with scoring variance σ_κ^2—the Gaussian contraction theorem (Theorem 4.5) gives*

$$Q^* = \mathcal{N}(\mu_F, \sigma_*^2).$$

As $\sigma_\kappa^2 \to 0$, $\sigma_*^2 \to \sigma_L^2$ and Q^* concentrates on p_F.

Proof. From Theorem 4.5, the fixed-point mean is $\mu_* = \mu_F$. As $\sigma_\kappa^2 \to 0$,

$$\sigma_*^2 = \frac{\sigma_L^2 + \sqrt{(\sigma_L^2)^2 + 4\,\sigma_L^2\,\sigma_\kappa^2}}{2} \;\to\; \sigma_L^2.$$

The distribution concentrates, with spread determined only by lifecycle variance injection. □

Remark 7.3. The statement "$Q^* = p_F$" is a limiting idealization. Precisely: $W_1(Q^*, p_F)$ decreases monotonically as scoring sharpness increases, approaching zero as $\sigma_\kappa^2 \to 0$. The cage is a continuum, not a discrete state: the degree of caging is determined by scoring sharpness relative to lifecycle variance.

Corollary 7.4 (Cage Properties)**.** *As Q^* approaches p_F:*

(a) *The acceptability distribution is a near-fixed point of the system's own evaluation function, making external disconfirmation structurally invisible.*

(b) *Escape requires either reducing scoring sharpness (relaxing what counts as acceptable) or increasing σ_L^2 (increasing structural exploration).*

7.2 Generative Preservation

Definition 7.5 (Selection Balance)**.** A constraint architecture $\mathcal{A}$ has *selection balance* $\beta \in [0, 1]$ when the scoring function's center is

$$\mu_* = (1 - \beta)\, \mu_F + \beta\, \mu_{\text{divergent}},$$

where μ_F is the Frame center and $\mu_{\text{divergent}}$ is a target reflecting functional novelty.

Theorem 7.6 (Generative Preservation—Gaussian)**.** *Under the Gaussian contraction theorem with selection balance $\beta > 0$:*

The fixed-point mean is $\mu_* = (1 - \beta)\, \mu_F + \beta\, \mu_{\text{divergent}}$, deviating from μ_F by $\|\mu_* - \mu_F\| = \beta\, \|\mu_{\text{divergent}} - \mu_F\|$.

If additionally:

(G1) *the scoring function assigns positive density to outputs outside any ε-neighborhood of p_F,*

(G2) *the threshold $\theta_{\min}$ does not exclude all novel outputs, and*

(G3) *$\lambda_* < 1$ (contraction holds),*

then Q^* has support outside p_F: the system maintains generative capacity.

Proof. With $\beta > 0$, the scoring center $\mu_* \neq \mu_F$. The fixed point $Q^* = \mathcal{N}(\mu_*, \sigma_*^2)$ has its mass centered away from the Frame. By G1 and G2, the tails of Q^* extend into novel territory. By G3 and Theorem 4.5, $Q_t \to Q^*$. The limit distribution has support outside p_F. □

Corollary 7.7 (Design Criterion)**.** *The cage is the limit as $\beta \to 0$. Generative capacity requires $\beta > 0$. The optimal β depends on the domain's cost structure.*

8 Open Problems

The Gaussian contraction theorem is rigorous for its stated assumptions. Extending it requires closing the following gaps. Each is classified by severity and given a bridging strategy.

8.1 Non-Gaussian Contraction

Open Problem 8.1 (Log-Concave Distributions)**.** Let κ be a log-concave scoring function (not necessarily Gaussian) with $-D^2 \log \kappa \succeq \Lambda$ for some $\Lambda \succ 0$, and let Q_t belong to the class of log-concave distributions on $\mathbb{R}^d$. Does the scoring-selection operator preserve log-concavity, and is the composed map a contraction?

Conjecture 8.2 (Log-Concave Contraction)**.** *If κ is strongly log-concave with parameter $\alpha > 0$ and the lifecycle operator adds variance bounded by σ_L^2, then the Q-dynamics operator contracts in W_2 with rate*

$$\lambda_{\mathrm{lc}} \leq \frac{1}{1 + \alpha\,\sigma_L^2}.$$

Bridging strategy. The Bakry-Émery criterion [1] provides contraction rates for log-concave measures under diffusion operators. Otto and Villani [9] established that lower bounds on Ricci curvature imply contraction in Wasserstein distance, providing a geometric framework for the type of contraction studied here. Concentration results for log-concave measures [3] supply the tail bounds needed for W_1 estimates. The scoring-selection step is multiplicative reweighting; the lifecycle step is convolution with a noise kernel. The composed operator may satisfy a Bakry-Émery condition with curvature determined by $-D^2 \log \kappa$. Near the fixed point, the second-order Taylor expansion of $\log \kappa$ recovers an effective Gaussian with precision $\Lambda_{\mathrm{eff}} = -D^2 \log \kappa$, reducing to the proved case for local stability. Global contraction from arbitrary initial conditions remains open.

8.2 Sub-Gaussian Extensions

Open Problem 8.3 (Sub-Gaussian Tails)**.** Let Q_t have sub-Gaussian tails with parameter σ and let κ be bounded and Lipschitz. Under what conditions on κ and the lifecycle operator does the composed map contract in W_1?

Conjecture 8.4 (Sub-Gaussian Contraction)**.** *If κ is L-Lipschitz with unique maximum and the lifecycle operator has sub-Gaussian increments with parameter σ_L, then the Q-dynamics operator contracts in W_1 with rate $\lambda \leq 1 - c\,L\,\sigma_L$ for some universal constant $c > 0$, provided $L\,\sigma_L$ is sufficiently small.*

Bridging strategy. Concentration inequalities for sub-Gaussian random variables, combined with coupling arguments for W_1 contraction. The key step is showing that the scoring-selection operator maps sub-Gaussian distributions to sub-Gaussian distributions with improved parameters.

8.3 Multi-Dimensional via Non-Product Measures

Open Problem 8.5 (Non-Commuting Covariance Contraction)**.** Prove global contraction of the covariance map $H(\Sigma) = (\Sigma^{-1} + \Lambda)^{-1} + \Sigma_L$ on the cone of positive definite matrices when Σ and Λ do not commute.

Status. Local stability is established: the Fréchet derivative of H at Σ_* has spectral radius bounded by $\lambda_{\mathrm{mean}}^2 < 1$. Global contraction requires a monotonicity argument on the Löwner order of the positive definite cone. The 1D monotone argument (used in Theorem 4.5) does not directly generalize because the Löwner order is a partial order rather than a total order.

Conjecture 8.6 (Non-Commuting Global Contraction)**.** *The map H is a contraction in the Thompson metric on the cone of positive definite matrices, with rate bounded by $\lambda_{\mathrm{mean}}^2$.*

Bridging strategy. The Thompson metric $d_T(\Sigma_1, \Sigma_2) = \| \log(\Sigma_1^{-1/2} \Sigma_2 \, \Sigma_1^{-1/2}) \|$ is complete on the positive definite cone and is compatible with the Löwner order. Matrix Riccati-type maps are known to be contractions in the Thompson metric under certain conditions [5]. The map H has the structure of a discrete algebraic Riccati iteration with additive perturbation Σ_L.

8.4 General Contraction

Open Problem 8.7 (General Scoring Functions)**.** Characterize the class of scoring functions κ and lifecycle operators for which WF4 holds. What is the minimal regularity on κ (beyond boundedness and measurability) that guarantees contraction?

Status. Open. The Gaussian case provides a template: contraction arises from the interaction of scoring curvature (which concentrates) and lifecycle variance injection (which spreads). The conjecture is that any scoring function with sufficient curvature near its maximum, combined with sufficient variance injection, produces contraction. Making "sufficient" precise is the problem.

8.5 Mixture Contraction

Open Problem 8.8 (Multi-Component Distributions)**.** When Q_t is a mixture of K Gaussians (each component corresponding to an agent), prove contraction of the overall distribution. The challenge is that fork events change the number of components, and the Wasserstein distance between mixtures does not decompose into per-component distances.

Bridging strategy. Multi-scale contraction argument: (a) each component contracts locally at rate λ_* (Theorem 4.5), (b) routing converges to a Voronoi partition of the input space, (c) fork/merge adjusts the number of components, (d) overall W_1 decreases because each component contracts and the component structure converges. The combined rate is dominated by the slow component (structural convergence). Tools from mean-field theory and adaptive resonance theory [4] are relevant but have not been assembled.

9 Conclusion

For the Gaussian case, the contraction parameter $\lambda_* = \sigma_\kappa^2/(\sigma_*^2 + \sigma_\kappa^2)$ is computed from first principles, not assumed. The fixed point Q^* is unique and globally attracting. Composed architectures preserve well-formedness under type compatibility and energy separability. The cage is characterized as the limit of convergent selection; generative capacity requires nonzero selection balance. The multi-dimensional extension is complete for commuting covariance-precision pairs and locally stable in general.

Five problems remain open: non-Gaussian contraction (log-concave and sub-Gaussian extensions), non-commuting covariance contraction in $d > 1$, general scoring function characterization, and mixture contraction for multi-component systems. For each, we have stated conjectures with bridging strategies. The Gaussian theorem provides the anchor; the extensions are the frontier.

References

[1] Dominique Bakry and Michel Émery. Diffusions hypercontractives. In *Séminaire de Probabilités XIX 1983/84*, volume 1123 of *Lecture Notes in Mathematics*, pages 177–206. Springer, 1985.

[2] Stefan Banach. *Sur les opérations dans les ensembles abstraits et leur application aux équations intégrales*, volume 3. 1922.

[3] Sergey G. Bobkov and Michel Ledoux. Weighted Poincaré-type inequalities for Cauchy and other convex measures. *Annals of Probability*, 37(2):403–427, 2009.

[4] Gail A. Carpenter and Stephen Grossberg. A massively parallel architecture for a self-organizing neural pattern recognition machine. *Computer Vision, Graphics, and Image Processing*, 37(1):54–115, 1987.

[5] Bas Lemmens and Roger Nussbaum. *Nonlinear Perron–Frobenius Theory*. Cambridge University Press, 2012.

[6] Jeremy McEntire. The cage: How fiduciary duty creates organizational incompleteness. Zenodo DOI: 10.5281/zenodo.18828433, 2025.

[7] Jeremy McEntire. Emergence: A programming paradigm for constraint-shaped agency. Working paper, 2025.

[8] Jeremy McEntire. The strategic rate-distortion-perception tradeoff: Bridging blau-michaeli compression theory and crawford-sobel strategic communication. Working paper, 2025.

[9] Felix Otto and Cédric Villani. Generalization of an inequality by Talagrand and links with the logarithmic Sobolev inequality. *Journal of Functional Analysis*, 173(2):361–400, 2000.

[10] Michel Talagrand. Transportation cost for Gaussian and other product measures. *Geometric and Functional Analysis*, 6(3):587–600, 1996.

[11] Cédric Villani. *Optimal Transport: Old and New*. Springer, 2009.