

The Fixed Point of Self-Improvement

Three limits of a learner that trains on its own memory

Michael P. Laing

mllaing@post.harvard.edu

ORCID 0009-0004-5593-5759

Prepared with drafting assistance from Anthropic Claude Fable 5.1

Licensed under Creative Commons Attribution 4.0 (CC BY 4.0)

September 28, 2026 — draft 8

Abstract

A system that improves itself trains each generation on records its previous generation wrote. The companion note [2] shows that a learner trained once on point-valued records of measured quantities converges to an artifact: the collapsed regression function, which under a standard observation model is the classical attenuation by the reliability ratio — in the original coordinate a concave power law absent from the data-generating system. This note asks what the artifact does under iteration. We define the population operator of the loop and prove three limits. If labels are held fixed and only readings are refreshed, the one-generation artifact is a fixed point: no generation improves on the first. If each generation is labeled by the predictions of the last, the learned response contracts to a constant, the unconditional mean of the outcome, at a rate governed by the squared maximal correlation between a reading and the quantity it measures — geometric whenever that correlation is less than one; along the way the training loss goes to zero, so the loop’s own evaluator reports success at the moment the response has been forgotten. If instead the memory stores each reading as a conditioning event on a carried law, the same readings drive the posterior to the truth. The contraction is not model collapse: it has infinite data per generation, samples nothing, and is only sharpened by more data. It is regression to the mean, compounded once per generation, and it is invisible after the fact because the dispersion that would reveal it was discarded when the reading was written down. What is claimed is not the calculation, which is classical, but its location — the memory store — where the loop’s own evaluator cannot see it, and the write convention that removes it. Not claimed: anything about generative models on text or code as such, or the rate at which any deployed system approaches these population limits.

1 The loop

The prospect of AI systems that participate in their own improvement — revising outputs, training on data they generate, patching controllers from streaming telemetry, contributing to the engineering of their successors — is now discussed as an engineering timeline rather than a thought experiment [3, 6]. The debate concerns whether such loops close and how fast. This note concerns a prior question: what a loop converges to, when the memory it trains on stores measured quantities as points.

The setting is the one the core paper [1] and its first companion [2] analyze for a single generation. A quantity X exists; an instrument reports a reading W of it with a known error

mechanism; a downstream outcome Y depends on the quantity through a true response g , $\mathbb{E}[Y | X] = g(X)$, up to outcome noise. The memory stores the reading at face value and the outcome. A learner trained on such records converges to $\mathbb{E}[Y | W = \cdot]$, the collapsed regression function [2, Prop. 1], and under the lognormal observation model of [1, Prop. 4.1] that function is $C w^\lambda$, $\lambda = s^2/(s^2 + \sigma^2)$, with the outcome exact (§4 adds outcome noise, which scales C), though g is linear (Figure 1): on the log scale, the classical attenuation of a regression slope by the reliability ratio λ — regression dilution — carried back to the original coordinate. The learner is not at fault: it answers, with increasing precision, a question the memory answered wrongly at write time.




Figure 1: One generation, the result of [2, Prop. 1]. The population target of any consistent learner on point-valued records is the concave law, not the linear truth. Drawn at the parameters of §4 ($\tau = 0.5$); $C_1 = C e^{\tau^2/2}$.


A self-improving loop runs this once per generation and writes the result back. Two loops are worth distinguishing by what the next generation uses as its label (Figure 2).

Loop A (labels held). The outcome Y stays in the record. Each generation re-reads the quantity and trains to predict Y from the new reading. This is retraining on refreshed telemetry with ground-truth outcomes retained.

Loop B (labels re-derived). Generation k pairs a fresh reading with the previous generation’s prediction on the previous reading. This is pseudo-labeling, self-distillation, and training on self-generated data *on those records for which the training input is a fresh observation of the labeled quantity* (§5); it is the loop that patches a controller from the last controller’s own outputs when the controller is re-run on fresh telemetry.

Loop C (labels mixed). Generation k trains on a fraction $\alpha \in (0, 1]$ of records, chosen independently of the record, that keep the outcome Y and a fraction $1 - \alpha$ labeled as in Loop B.

Loop A: labels held



Loop B: labels re-derived




Figure 2: Loop A refreshes readings and keeps the outcome as label. Loop B labels generation k with generation $k-1$'s predictions on generation $k-1$'s readings, then trains against a fresh reading.

This is a fixed-fraction analogue of the retained-real-data regimes of the model-collapse literature [5, 17]; the accumulating form of [17], which retains every past generation, is not covered (§8).

All three are population statements: every generation is assumed to reach its target exactly. That is the opposite regime from model collapse (§6), and the choice is deliberate.

2 Setup

The assumptions, in one place. (i) The quantity X takes values in a standard Borel space with law π . (ii) Readings $W_1, W_2, \dots$ are generated from X by one Markov kernel Q and are conditionally independent given X . (iii) The outcome Y satisfies $\mathbb{E}[Y | X, W_1, W_2, \dots] = g(X)$ with $\mathbb{E}[Y^2] < \infty$ and g square-integrable under π : outcome noise, if any, is conditionally mean-zero and does not enter any target below. (iv) Every generation reaches its population target over all measurable predictors; the theorems are about targets, not fits. (v) For Theorem 5, Assumption 3, and for its geometric rate a spectral gap, verified in the Gaussian case of §4; for Theorem 8, identifiability of the kernel ($x \mapsto Q(x, \cdot)$ injective). No distributional form is assumed before §4.

Let X take values in a standard Borel space $\mathcal{X}$ with law π . A reading is generated by a Markov kernel Q from $\mathcal{X}$ to a standard Borel space $\mathcal{W}$: given $X = x$, $W \sim Q(x, \cdot)$. Readings $W_1, W_2, \dots$ of the same quantity are conditionally independent given X , each with kernel Q . Write $\nu = \pi Q$ for the marginal law of a reading and $P(w, dx)$ for the posterior of X given $W = w$,

so that $\pi(dx)Q(x, dw) = \nu(dw)P(w, dx)$. The response $g: \mathcal{X} \rightarrow \mathbb{R}$ is square-integrable under π .

Two conditional-expectation operators carry the whole argument:

$$\begin{aligned}(Sh)(x) &= \int Q(x, dw) h(w) = \mathbb{E}[h(W) \mid X = x], \\ (S^*f)(w) &= \int P(w, dx) f(x) = \mathbb{E}[f(X) \mid W = w].\end{aligned}$$

$S: L^2(\nu) \rightarrow L^2(\pi)$ and $S^*: L^2(\pi) \rightarrow L^2(\nu)$ are adjoint, each is a contraction (conditional Jensen), and each fixes constants.

Definition 1 (Loop operators). The one-generation target is $h_A = S^*g$, i.e. $h_A(w) = \mathbb{E}[g(X) \mid W = w]$. The re-observation operator on $L^2(\nu)$ is

$$T = S^*S, \quad (Th)(w) = \mathbb{E}[h(W_1) \mid W_2 = w],$$

where W_1, W_2 are two readings of the same X .

Under Loop A every generation's target is h_A . Under Loop B generation k predicts $h_{k-1}(W_{k-1})$ from W_k , so its target is $h_k = Th_{k-1}$, and with $h_1 = h_A$,

$$h_k = T^{k-1}S^*g = (S^*S)^{k-1}S^*g.$$

Lemma 2. T is self-adjoint and positive on $L^2(\nu)$, with spectrum in $[0, 1]$, $T\mathbf{1} = \mathbf{1}$, and $\int (Th) d\nu = \int h d\nu$. Its fixed space $\{h : Th = h\}$ is exactly the set of $h \in L^2(\nu)$ such that $h(W)$ is almost surely a function of X .

Proof. $T = S^*S$ is self-adjoint and positive by construction, and $\|T\| \leq \|S^*\|\|S\| \leq 1$. $T\mathbf{1} = \mathbf{1}$ since each factor fixes constants; mean preservation is $\langle Th, \mathbf{1} \rangle = \langle Sh, S\mathbf{1} \rangle = \langle h, S^*\mathbf{1} \rangle = \langle h, \mathbf{1} \rangle$. For the fixed space: since $0 \leq T \leq I$, $Th = h$ iff $\langle (I - T)h, h \rangle = 0$ iff $\|Sh\|^2 = \|h\|^2$, and $\|Sh\| = \|h\|$ iff $\mathbb{E}[\text{Var}(h(W) \mid X)] = 0$ iff $h(W) = \mathbb{E}[h(W) \mid X]$ a.s. $\square$

Assumption 3 (Noisy reading). No ν -a.e. nonconstant $h \in L^2(\nu)$ has $h(W)$ almost surely a function of X . Equivalently, the fixed space of T is the constants.

A sufficient condition: there is a reference measure μ on $\mathcal{W}$ with $Q(x, \cdot) \sim \mu$ (mutually absolutely continuous) for π -a.e. x , i.e. every reading is possible for every value of the quantity. Then $h(W)$ a.s. a function of X means $h = c(x) Q(x, \cdot)$ -a.s. for each x , hence $h = c(x) \mu$ -a.e. for each x , hence c is constant. Additive or multiplicative Gaussian error satisfies this; a reading with a deterministic component of X (a reported digit that is always exact) does not, and that component is then preserved by the loop.

3 Three limits

Theorem 4 (Loop A: the artifact is a fixed point). *Under Loop A the target of every generation is $h_A = S^*g$. In particular the transformation error of [2, Prop. 5] — the discrepancy between h_A and g composed with the posterior mean — is the same in every generation.*

Proof. Each generation solves the same population problem $\arg \min_h \mathbb{E}[(Y - h(W_k))^2]$ with $(X, W_k) \sim \pi(dx)Q(x, dw)$; the solution is $\mathbb{E}[Y \mid W_k = \cdot] = S^*g$, since $\mathbb{E}[Y \mid X, W_k] = g(X)$ [2, Prop. 1]. $\square$

Loop A is the honest loop and it still cannot improve: retraining on fresh readings with true outcomes retained reproduces the first generation’s artifact exactly, because the artifact is a property of the record’s population law, not of any sample.

Theorem 5 (Loop B: contraction to the mean). *Under Assumption 3, $h_k \rightarrow \mathbb{E}_\pi[g]$ in $L^2(\nu)$: the learned response converges to the constant equal to the unconditional mean of the outcome. Let*

$$\rho = \sup \{ \|Sh\|^2 / \|h\|^2 : h \in L^2(\nu), \int h d\nu = 0, h \neq 0 \}$$

be the norm of T on mean-zero functions. Then

$$\|h_k - \mathbb{E}_\pi[g]\|_{L^2(\nu)} \leq \rho^{k-1} \|h_A - \mathbb{E}_\pi[g]\|_{L^2(\nu)},$$

and $\rho = R^(X, W)^2$, the square of the Hirschfeld–Gebelein–Rényi maximal correlation between the quantity and a reading. Here $\rho \leq 1$; the rate is geometric exactly when $\rho < 1$, i.e. when T has a spectral gap on mean-zero functions, which Assumption 3 does not by itself guarantee (it holds in the Gaussian case of §4, where the spectrum of T is $\{\lambda^n\}$).*

Proof. By Lemma 2 and the spectral theorem, $T^k h \rightarrow E_{\{1\}} h$, the projection of h onto the fixed space. Under Assumption 3 that space is the constants, and mean preservation identifies the projection: $E_{\{1\}} h = \int h d\nu$. For $h = h_A = S^* g$, $\int h_A d\nu = \mathbb{E}[\mathbb{E}[g(X) | W]] = \mathbb{E}_\pi[g]$. For the rate, T restricted to the closed invariant subspace $L_0^2(\nu)$ of mean-zero functions is self-adjoint with norm $\sup \langle Th, h \rangle / \|h\|^2 = \sup \|Sh\|^2 / \|h\|^2 = \rho$, and $h_k - \mathbb{E}_\pi[g] = T^{k-1}(h_A - \mathbb{E}_\pi[g])$ with $h_A - \mathbb{E}_\pi[g] \in L_0^2(\nu)$. Finally $\|Sh\|^2 / \|h\|^2 = \text{Var}(\mathbb{E}[h(W) | X]) / \text{Var}(h(W))$ for mean-zero h , whose supremum over h is, by the variational characterization of maximal correlation [7] — ρ is the square of the second singular value of the conditional-expectation operator [10, 11] — $\sup_{h,f} \text{corr}(h(W), f(X))^2 = R^*(X, W)^2$. Equivalently $\rho = \eta_{\chi^2}(Q, \pi)$, the χ^2 -contraction coefficient of the reading kernel at the prior, $\sup_{\pi'} \chi^2(\pi' Q \| \pi Q) / \chi^2(\pi' \| \pi)$: a strong data-processing constant [23, 24]. $\square$

Corollary 6 (The evaluator reports success). *Under Loop B the population training loss of generation $k \geq 2$, $\ell_k = \mathbb{E}[(h_k(W_k) - h_{k-1}(W_{k-1}))^2]$, satisfies $\ell_k = \mathbb{E}[\text{Var}(h_{k-1}(W_{k-1}) | W_k)] \leq \text{Var}(h_{k-1}(W)) \leq \rho^{2(k-2)} \text{Var}(h_A(W))$, and $\ell_k \rightarrow 0$ by Theorem 5, while the error against the truth, $\mathbb{E}[(h_k(W_k) - g(X))^2]$, converges to $\text{Var}_\pi(g)$, the smallest error any constant predictor can attain: the loop ends at the best predictor that has forgotten the reading.*

Proof. $h_k(W_k) = \mathbb{E}[h_{k-1}(W_{k-1}) | W_k]$, so ℓ_k is the residual conditional variance; the bound follows from Theorem 5 applied to h_{k-1} . The truth error follows from $h_k \rightarrow \mathbb{E}_\pi[g]$. $\square$

The loop’s evaluator scores generation k against the record, and the record is generation $k-1$. It reports each generation fitting its labels better than the last, and it is right. The quantity that would expose the drift — how much the labels have moved from the outcomes — is not in the record.

Proposition 7 (Loop C: the mixture has its own fixed point). *Under Loop C with retained fraction $\alpha \in (0, 1]$, the targets satisfy $h_k = \alpha h_A + (1 - \alpha)Th_{k-1}$ with $h_1 = h_A$, and*

$$h_k \rightarrow h^* = \alpha (I - (1 - \alpha)T)^{-1} h_A \quad \text{in } L^2(\nu), \quad \|h_k - h^*\| \leq ((1 - \alpha)\rho)^{k-1} \|h_A - h^*\|,$$

with no assumption on ρ : the rate is at most $1 - \alpha$. In the spectral representation of T , $h^ - \mathbb{E}_\pi[g] = \varphi(T)(h_A - \mathbb{E}_\pi[g])$ with $\varphi(t) = \alpha / (1 - (1 - \alpha)t) \in [\alpha, 1]$ on $[0, 1]$; hence $\alpha \|h_A - \mathbb{E}_\pi[g]\| \leq \|h^* - \mathbb{E}_\pi[g]\| \leq \|h_A - \mathbb{E}_\pi[g]\|$, with strict inequalities for $\alpha < 1$ whenever $h_A - \mathbb{E}_\pi[g]$ has spectral mass off $\{1\}$ and off $\{0\}$, and, in the lognormal case of §4, h^* is no longer a pure power law.*

Proof. $F(h) = \alpha h_A + (1 - \alpha)Th$ is a contraction on $L^2(\nu)$ with constant $(1 - \alpha)\|T\| \leq 1 - \alpha < 1$, so it has a unique fixed point h^* , and $\|(1 - \alpha)T\| < 1$ makes $I - (1 - \alpha)T$ invertible by its Neumann series, giving the formula. Then $h_k - h^* = (1 - \alpha)T(h_{k-1} - h^*)$; since T preserves ν -means and $\int h^* d\nu = \int h_A d\nu = \mathbb{E}_\pi[g]$, the difference is mean-zero and the bound follows with ρ the norm of T there. The spectral statement is the Neumann series applied through the spectral measure of T : from $(I - (1 - \alpha)T)h^* = \alpha h_A$ and $T\mathbf{1} = \mathbf{1}$, $(I - (1 - \alpha)T)(h^* - \mathbb{E}_\pi[g]) = \alpha(h_A - \mathbb{E}_\pi[g])$. $\square$

Loop C is the loop that retains some ground truth. It does not return to the truth: it returns to a shrunk copy of the one-generation artifact, and the retained fraction sets how much of the response survives.

Theorem 8 (Conforming memory: the same readings find the truth). *Suppose the memory stores the cell $\mathcal{L}(X \mid W_1, \dots, W_k)$ and treats each new reading as a conditioning event — the operation [1, §3.1] describes and [1, Prob. 4] leaves to be specified, which this note takes as given in the conjugate case of §4 and assumes in general. Then $\mathbb{E}[g(X) \mid W_1, \dots, W_k] \rightarrow g(X)$ almost surely and in L^2 , provided the model is identifiable ($x \mapsto Q(x, \cdot)$ injective).*

Proof. $\mathbb{E}[g(X) \mid W_1, \dots, W_k]$ is a uniformly integrable martingale converging to $\mathbb{E}[g(X) \mid W_1, W_2, \dots]$. By Doob's consistency theorem [8] and [13, Thm. 6.9] for conditionally i.i.d. observations from an identifiable standard Borel model, the posterior $\mathcal{L}(X \mid W_1, \dots, W_k)$ converges weakly to δ_X for π -a.e. value of X , so X is measurable with respect to the tail σ -field of the readings up to a null set, and the limit is $g(X)$. $\square$

The three theorems are one statement about what a memory keeps of a reading. A point-valued memory under Loop A keeps one reading's worth, forever. Under Loop B it keeps, in the limit, none. A measure-valued memory keeps all of them, and k readings conditioned into one cell converge on the quantity. Same instrument, same readings, same learner: the fixed point is decided by the storage convention. Read through the data-processing inequality for the chain $W_{k-1} \text{---} X \text{---} W_k$, the three are: what the default reader of the record computes (Theorem 4), the rate at which the loop loses what the record kept (Theorem 5, ρ its strong-DPI constant), and what reaches the bound (Theorem 8). The title's fixed point is Theorem 4's: the artifact, which the honest loop cannot leave. Theorem 5's fixed point is the constant, which the re-deriving loop cannot avoid.

4 The lognormal case

Take the observation model of [1, Prop. 4.1] extended by outcome noise, as in [2, Prop. 1]: $\log X \sim N(m, s^2)$, $\log W = \log X + \eta$, $\eta \sim N(0, \sigma^2)$, and $Y = \beta X e^\zeta$, $\zeta \sim N(0, \tau^2)$, with X , the reading errors $\eta_1, \eta_2, \dots$ and ζ mutually independent, so that (iii) holds and the response $g(x) = \mathbb{E}[Y \mid X = x] = \beta e^{\tau^2/2} x$ is linear. The outcome noise moves no exponent and scales every constant below by $e^{\tau^2/2}$; it is carried because with $\tau = 0$ the record's law identifies the response by the reverse regression of $\log W$ on $\log Y$, and with $\tau > 0$ it does not [2, §4]. Then $(\log X, \log W)$ is jointly Gaussian with $\text{corr}(\log X, \log W)^2 = s^2/(s^2 + \sigma^2) = \lambda$, and for jointly Gaussian pairs the maximal correlation equals the ordinary correlation [9], so $\rho = \lambda$ in Theorem 5. On the log scale, in the coordinate standardized by ν , T is the Ornstein–Uhlenbeck semigroup with correlation λ (the Mehler kernel), and the exponent λ^k below is the attenuation of §1 iterated: the shrinkage of a regression slope by the reliability ratio, once per generation. Mehler's formula is its name.

Proposition 9 (Exponent compounds). *Under this model $h_k(w) = C_k w^{\lambda^k}$ for every $k \geq 1$, with*

$$C_k = \beta \exp\left(m(1 - \lambda^k) + \frac{1}{2} s^2 (1 - \lambda^{2k-1}) + \frac{1}{2} \tau^2\right), \quad C_k \rightarrow \mathbb{E}[Y] = \beta e^{m+s^2/2+\tau^2/2}.$$

Proof. On the log scale T is the Mehler (Ornstein–Uhlenbeck) operator with parameter λ : given $\log W_2 = z$, $\log W_1$ is Gaussian with mean $m + \lambda(z - m)$ and variance $s^2\sigma^2/(s^2 + \sigma^2) + \sigma^2$. The Mehler operator maps e^{tz} to a constant times $e^{\lambda tz}$ (Gaussian moment generating function); the operator and its Hermite spectrum are standard [12]. $h_1 = C_1 e^{\lambda \log w}$ by [2, Prop. 1], with $C_1 = C e^{\tau^2/2}$, C the constant of [1, Prop. 4.1], the outcome noise being independent of the readings; the step $T(Cw^a) = C \exp(a(1 - \lambda)m + \frac{1}{2} a^2 V) w^{a\lambda}$ with $V = \sigma^2(2s^2 + \sigma^2)/(s^2 + \sigma^2)$ the conditional variance above, applied at $a = \lambda^{k-1}$, sums to the closed form for C_k ; induct. $\square$

The fitted log–log slope of the learned response is therefore λ^k and its constant is C_k : an exhibit can check both to as many places as its sample allows. A C program runs both loops (`iterated_exhibit_20Sep_3/`, seed 20260910) with 2×10^6 records per generation, $m = 0$, $s = 1$, $\sigma = 0.8$, $\tau = 0.5$, $\beta = 2$, so $\lambda = 0.6098$. Loop B slopes for generations 1–6: 0.6098, 0.3720, 0.2268, 0.1383, 0.0845, 0.0515, against $\lambda^k = 0.6098, 0.3718, 0.2267, 0.1382, 0.0843, 0.0514$; fitted constants (least squares in y at the fitted exponent) 2.751, 3.329, 3.575, 3.671, 3.708, 3.722, against $C_k = 2.755, 3.336, 3.582, 3.678, 3.715, 3.728$, converging on $\beta e^{s^2/2+\tau^2/2} = 3.736$. Loop A slopes stay within 0.0006 of λ , at 0.6092–0.6099 (Figures 3 and 4). Under Loop C the fixed point is $h^*(w) = \alpha \sum_{j \geq 0} (1 - \alpha)^j C_{j+1} w^{\lambda^{j+1}}$, a convergent sum of power laws with distinct exponents and positive weights; it is not a power law.




Figure 3: Loop B, six generations of learned response, flattening toward the constant $\mathbb{E}[Y] = \beta e^{m+s^2/2+\tau^2/2}$.

On the conforming side, k readings conditioned into one cell give $\log X \mid W_1, \dots, W_k \sim N(\cdot, (1/s^2 + k/\sigma^2)^{-1})$: the posterior precision grows linearly in k , the pushforward $\mathbb{E}[Y \mid W_1, \dots, W_k] = \beta e^{\tau^2/2} \mathbb{E}[X \mid W_1, \dots, W_k]$ is exact by [1, Thm. 3.2] because the response is affine in X , and no artifact enters at any generation (Figure 5).




Figure 4: Fitted exponent by generation: Loop A at the fixed point λ , Loop B on λ^k . Lines are Proposition 9 (with Theorem 4 for the constant line); markers are the exhibit.

5 Which loops are which

Real pipelines are mixtures, and Loop C is the mixture. Continual fine-tuning on fresh telemetry with ground-truth outcomes retained is Loop A. Theorem 5 needs two things of a pipeline: that the label pass through one point-valued write per generation, and that the reading the next generation trains on be a fresh reading of the same quantity, conditionally independent, given the quantity, of the one the label was computed on. A pipeline that labels a reading and re-trains on that same reading has $T = I$ and sits at Loop A’s fixed point; it is the re-observation — refreshed telemetry, re-measured sensors, input noise or augmentation applied independently at labeling and at training — that makes it Loop B. Distillation, pseudo-labeling, reward models trained on the last policy’s outputs, and controllers patched from their own logs are Loop B *on those records for which the training input is a fresh observation of the labeled quantity*, and Loop A otherwise. The learner’s architecture, the optimizer, and the sample size do not enter, because the theorem is about the target, not the fit.

The rate $\rho = R^*(X, W)^2$ is a property of the instrument — the strong-data-processing constant of its channel at the prior (§3). A reading that determines the quantity nearly exactly has ρ near 1 and forgets slowly; a noisy reading forgets in a few generations. Nothing downstream changes ρ . A better model converges faster to h_k ; it does not move h_k .

6 This is not model collapse

Shumailov et al. [5] show that generative models trained recursively on their own samples degrade: tails are undersampled, the fitted distribution narrows, and after enough generations the model produces a point. Their formal result, and the headline mechanism, is their statistical-approximation error: a finite-sample phenomenon, present with a kernel density estimate and




Figure 5: Same readings, two memories. Left: point-valued, Loop B, the exponent decays to zero. Right: measure-valued, the posterior precision of $\log X$ grows as $1 + ks^2/\sigma^2$.

no learner at all, in which the estimation error of a correctly specified model accumulates generation by generation under repeated resampling and refitting — in the linear-regression analyses [15, 17], linearly in the number of generations under replacement; Gerstgrasser et al. [17] show it is attenuated when real data is retained and samples accumulate. More data per generation weakens it.

Theorem 5 has infinite data per generation, samples nothing, loses no tails, and is only sharpened by more data, which improves each generation’s convergence to its target. Its mechanism is Jensen’s inequality applied to a record that lost its dispersion — regression to the mean [4], compounded once per generation — and the survey literature’s “collapse dynamics” [3] does not contain it. Retaining real data does not help when the real data is itself point-valued, because Theorem 4 says the first generation is already at the artifact. The two effects are distinct in mechanism: one is sampling error compounded by resampling, the other is a population contraction that sampling cannot see. How they combine in a finite-sample Loop B — where the contraction also damps the propagation of sampling error — is not treated here. The nearest prior results are on the other side of each distinction: the regression-setting analyses of model collapse [15] shows that with additive *outcome* noise and infinite data there is no collapse, the complement of the input-side mechanism here and of [2, Rem. 4]; self-distillation as power iteration of a contraction whose iterates degenerate [14] has the shape of Theorem 5 with regularization in the role of re-observation; and iterated learning converging to a prior-determined fixed point [20, 21] is the same contraction with the prior in the role of the record. See also [16, 18, 19, 22].

7 What this says for self-improving systems

The results are stated for a memory layer, not for any deployed system, and their content is a specification, not a warning. None of the calculation is new: the contraction is regression to the mean iterated, and a statistician sees it at once. What is claimed is its location (the memory store), its invisibility to the loop’s own evaluator (Corollary 6), and the write convention that

removes it (Theorem 8). Three consequences are worth stating in the vocabulary of the safety discussion.

The property is not testable from inside the loop. Under Loop B, Corollary 6 says the loop’s own loss goes to zero as the response is forgotten; an evaluator that keeps generation-1 outcomes sees the error rise to $\text{Var}_\pi(g)$, but the record the loop trains from cannot represent it. Under Loop A no evaluation against the record can detect the artifact at all, because the artifact is the record’s own regression function. Whether a memory carries dispersion is a property that must be specified at design time, and it can be checked: the three conformance checks of [2, §5] — Dirac conformance, affine transparency, curvature accounting — test whether an implementation transports the cell, and are checkable on it. Whether the memory also conditions on each reading, which Theorem 8 requires and those checks do not test, is the elicitation question of [1, Prob. 4] and must be specified separately.

Improvement is not the fixed point’s friend. Under each loop a more capable generation is a more faithful copy of the artifact (Theorem 4), of h_k one step nearer the constant (Theorem 5), or of h_k one step nearer h^* (Proposition 7); the step length ρ is the instrument’s. Capability spends itself on distance to target; the target is set by the memory store.

The repair is at the memory store. Theorem 8 needs nothing the instrument does not already know — the error model from its calibration certificate — and in the conjugate case of §4 it costs two numbers per cell. In general the cost is the representation question of [1, Prob. 1] and the kernel question of [1, Prob. 4]; what this note fixes is where the repair must live, not what it costs. A memory that stores the cell and conditions on each reading turns the same telemetry that forgets under Loop B into telemetry that converges on the truth. The difference is the write convention.

8 What is not claimed

No claim is made about the behavior of any particular system. The theorems are population statements about targets, so they bound what learning can achieve and say nothing about how fast a finite learner approaches the bound. Assumption 3 excludes readings with an exactly reported component of the quantity; such a component survives the loop, and the theorem then applies to the noisy remainder. Theorem 8 assumes identifiability and conditionally independent readings; correlated readings — a shared calibration certificate, a shared clock, the couplings the core carries as data — are outside it. A shared unknown offset is not identified by any number of readings, and the posterior then converges to $\mathcal{L}(X \mid X + c)$, not to δ_X : the coupling changes the limit, not only the rate. Carrying the coupling as data is what makes that limit computable; it does not remove it. The response is assumed square-integrable; every generation is assumed to reach its population target over all measurable predictors, so a learner confined to a model class is not covered; and each generation is labeled by the one before it, so mixtures of outcomes and predictions are covered only in the form of Loop C, and predictions from several past generations are not covered. Nothing here is a claim about hallucination in language models, though Corollary 6 — a reported certainty with nothing behind it — is that pathology at the memory layer.

Canonical reference for this note, its companion, and the core: <https://realql.org>.

References

- [1] M. P. Laing, “Measure-Relational Algebra: The Core,” 2026, <https://doi.org/10.5281/zenodo.22999911> (all versions); served at https://realql.org/mra_core.pdf.
- [2] M. P. Laing, “What a Memory Must Preserve: Machine Reasoning over Uncertain Observations,” September 2026, https://realql.org/what_a_memory_must_preserve.pdf.
- [3] M. Chen, L. Wang, B. Qu, “Recursive Self-Improvement in AI: From Bounded Self-Refinement to Autonomous Research Loops,” arXiv:2607.07663v2 (2026).
- [4] N. Sivadas, D. Sibeck, V. Subramanyan, M.-T. Walach, et al., “Regression to the mean can explain saturation of geomagnetic storms,” *Nature* 655, 1143–1147 (2026).
- [5] I. Shumailov, Z. Shumaylov, Y. Zhao, N. Papernot, R. Anderson, Y. Gal, “AI models collapse when trained on recursively generated data,” *Nature* 631, 755–759 (2024), doi:10.1038/s41586-024-07566-y; Author Correction, *Nature* 640, E6 (2025).
- [6] M. Favaro and J. Clark, “When AI Builds Itself,” Anthropic Institute, June 2026, <https://www.anthropic.com/institute/recursive-self-improvement>.
- [7] A. Rényi, “On measures of dependence,” *Acta Math. Acad. Sci. Hungar.* 10, 441–451 (1959).
- [8] J. L. Doob, “Application of the theory of martingales,” *Le Calcul des Probabilités et ses Applications*, Colloq. Internat. CNRS 13, Paris, 23–27 (1949).
- [9] H. O. Lancaster, “Some properties of the bivariate normal distribution considered in the form of a contingency table,” *Biometrika* 44, 289–292 (1957).
- [10] H. S. Witsenhausen, “On sequences of pairs of dependent random variables,” *SIAM J. Appl. Math.* 28, 100–113 (1975).
- [11] A. Buja, “Remarks on functional canonical variates, alternating least squares methods and ACE,” *Ann. Statist.* 18, 1032–1069 (1990).
- [12] D. Bakry, I. Gentil, M. Ledoux, *Analysis and Geometry of Markov Diffusion Operators*, Springer, 2014.
- [13] S. Ghosal and A. van der Vaart, *Fundamentals of Nonparametric Bayesian Inference*, Cambridge University Press, 2017.
- [14] H. Mobahi, M. Farajtabar, P. L. Bartlett, “Self-distillation amplifies regularization in Hilbert space,” *NeurIPS* 2020.
- [15] E. Dohmatob, Y. Feng, J. Kempe, “Model collapse demystified: the case of regression,” arXiv:2402.07712 (2024).
- [16] E. Dohmatob, Y. Feng, P. Yang, F. Charton, J. Kempe, “A tale of tails: model collapse as a change of scaling laws,” *ICML* 2024.
- [17] M. Gerstgrasser et al., “Is model collapse inevitable? Breaking the curse of recursion by accumulating real and synthetic data,” arXiv:2404.01413 (2024).

- [18] S. Alemohammad et al., “Self-consuming generative models go MAD,” ICLR 2024.
- [19] Q. Bertrand, A. J. Bose, A. Duplessis, M. Jiralerspong, G. Gidel, “On the stability of iterative retraining of generative models on their own data,” ICLR 2024.
- [20] T. L. Griffiths and M. L. Kalish, “Language evolution by iterated learning with Bayesian agents,” *Cognitive Science* 31, 441–480 (2007).
- [21] M. L. Kalish, T. L. Griffiths, S. Lewandowsky, “Iterated learning: intergenerational knowledge transmission reveals inductive biases,” *Psychon. Bull. Rev.* 14, 288–294 (2007).
- [22] E. Dohmatob, Y. Feng, A. Subramonian, J. Kempe, “Strong model collapse,” ICLR 2025; arXiv:2410.04840 (2024).
- [23] O. V. Sarmanov, “Maximum correlation coefficient (nonsymmetric case),” *Dokl. Akad. Nauk SSSR* 121, 52–55 (1958).
- [24] Y. Polyanskiy and Y. Wu, “Strong data-processing inequalities for channels and Bayesian networks,” in *Convexity and Concentration*, IMA Vol. Math. Appl. 161, Springer, 2017, 211–249; arXiv:1508.06025.