

SAHILVERSE / FIELD NOTE 026

BUTTERFLY ENGINE.

TWO WORLDS. ONE INITIAL DIFFERENCE. A practical field guide to feedback, sensitivity and intellectual humility.

Creative project author: Sahil Tayal · Sahilverse · October 2026

Companion experiment: butterflyengine.sahiltayal.online

This guide describes a deliberately simplified, computational thought experiment. It does not claim to forecast ecosystems, weather, people, markets or organisational outcomes.

01 / The proposition

If two rule-governed systems begin almost alike, must their futures stay alike? Not necessarily. A small change may vanish, remain local, or be amplified through repeated feedback. The Butterfly Engine explores these possibilities in a discrete toy ecosystem.

The meaningful comparison holds all else constant. World A and World B use the same initial seed, agent placement, resource grid, update order and time-indexed weather. World B differs at the start only in the amount of food on the centre cell.

02 / What you are actually looking at

The grid contains 30 columns by 20 rows: 600 cells. Each cell holds integer food from 0 to 12 and at most one occupant. Each occupant also has a small energy score. Food is coloured by density; agents are represented by pale circles. Heatmap brightness represents absolute difference between paired cells.

A tick is a complete synchronous-ish update of food, followed by ordered movement/feeding and then reproduction. Scan order and neighbour preference are fixed, which makes the program replayable. The rules are invented for this experiment, not calibrated against ecological data.

03 / The rules, precisely

Initial state. A seeded xorshift generator assigns each cell 2–6 food units and gives each cell an approximately 11.5% chance of one agent. Each initial agent receives energy 4–7. At time zero World B centre food is adjusted by the selected amount and clamped between 0 and 12. This clamp matters: the requested change and the realised difference may not match near a boundary.

Weather. At each tick a deterministic derived seed picks dry, mild or wet. Both worlds share that value. Weather grants 0, 1 or 2 growth units respectively. Each cell also gets one nutrient growth unit if at least one orthogonal neighbour had 7+ food in the preceding state. Occupied cells incur a one-unit growth penalty. Food is clamped to 0–12.

04 / Agents and feedback

Agents are processed in row-major order. Each considers staying put or moving north/east/south/west; it chooses the available location with most food, using fixed order to break ties. Locations occupied earlier in the movement pass cannot be selected, and initially occupied target cells are blocked.

An agent loses one energy each tick. If its chosen cell has food, it consumes one food unit and gains two energy. Energy at or below zero causes death. Surviving agents are capped at energy 12. After movement, each agent with energy 8+ can reproduce into the first orthogonally adjacent vacant cell with at least two food units. The offspring starts at energy three and the parent pays three energy.

These interactions create nonlinear thresholds. A one-unit food difference may alter a move, which alters consumption, which alters future food and survival. But that ripple is not inevitable. Constraints, saturation and shared weather can also suppress differences.

05 / Measurements and what they mean

Population is the sum of occupied cells. **Resources** is the sum of food units. **Births and deaths** accumulate from local transition events. **Differing cells** is the count of paired cells with any food, occupancy or energy difference.

Total state distance sums, over every grid cell, $|\text{food A} - \text{food B}| + 3 \times |\text{occupancy A} - \text{occupancy B}| + |\text{energy A} - \text{energy B}|$. Weight three is an editorial design choice, not an established scientific constant. The metric does not measure moral value, wellbeing or real-world impact.

A high distance may come from changed resources rather than changed population. Inspect several measures together. Exported CSV rows represent actual recorded simulation states, including the initial state.

06 / What sensitivity is not

A deterministic system can display changing outcomes without any new randomness. Conversely, a large difference between two runs does not, by itself, prove rigorous mathematical chaos. Formal claims require carefully defined dynamical systems, asymptotic analysis or appropriate Lyapunov evidence. This small bounded, finite-state model does not offer that proof.

07 / Guided experiment A: the sign

Question: Does adding one resource have an effect symmetric to removing one?

1. Set a seed such as 271828. Choose +1 resource and reset. Run 40 ticks, pause, save a snapshot and export CSV.
2. Keep the seed exactly the same. Choose -1 and reset. Run another 40 ticks.
3. Compare distance, population gaps and first divergence. Consider whether clamping changes the realised size of the perturbation.

Interpretation: Any difference you observe belongs to this implementation and seed. Report it as a conditional observation, not a universal law.

08 / Guided experiment B: scale

Question: Is a six-unit nudge guaranteed to have six times the effect of a one-unit nudge?

1. Use the same seed with +1 and +6 initial resource settings, each from tick zero.
2. Read differing-cell count and total distance at ticks 10, 20 and 40.
3. Note threshold effects, saturation and cases in which a larger perturbation produces a smaller eventual gap.

Interpretation: Systems with thresholds can be non-proportional. This is a useful thinking prompt, not evidence for a universal scaling rule.

09 / Guided experiment C: counterfactual replay

Question: When does a divergence become noticeable?

1. Run to tick 60, then drag the recorded timeline backwards. Observe when food differences appear and when agent counts first separate.
2. Rewind to tick 15 and press Run. The recorded future after tick 15 is replaced by a new branch computed from the same state.
3. With unchanged parameters, inspect whether the continuation reproduces the former outcome.

Interpretation: Replayability supports causal inspection inside this program. It does not establish causal effects outside it.

10 / An organisational thought experiment

Consider an imagined onboarding programme. A team provides one additional unit of attention to a newcomer. That contact might create trust, faster learning and peer support; or a bottleneck might absorb it entirely. The point is neither “every detail changes everything” nor “small interventions always succeed.” The question is: which feedback loops, constraints, buffers and thresholds matter?

Use three prompts: (1) What is the “food” variable in the analogy, and what gets lost by using one number? (2) Which choices made by people are not captured by automatic toy agents? (3) What observations would be needed before claiming the metaphor applies to an actual workplace?

11 / How to read failure and sameness

If the worlds become identical, the simulator is still doing something informative: the perturbation was absorbed under those conditions. If they diverge dramatically, avoid narrating the end result as destiny. The model is fully mechanistic and sensitive to chosen rules, scan order, grid size and measurement horizon.

12 / Reflective exercises

Counterfactual discipline. Write one claim about a personal or organisational turning point. Separate observed facts from the story that a minor event caused everything that followed. List plausible alternative pathways.

Resilience audit. Find a system where small disruptions are buffered: redundant processes, clear handoffs, safety margins, or deliberate pauses. How might you model that buffer?

Model humility. Change just one rule mentally: remove reproduction, allow agents to share cells, or reverse tie priority. Which conclusions no longer survive?

13 / Further reading

Edward N. Lorenz (1963), “Deterministic Nonperiodic Flow,” *Journal of the Atmospheric Sciences*, 20(2), 130–141. A foundational paper in deterministic atmospheric dynamics.

Edward N. Lorenz (1972), “Predictability: Does the Flap of a Butterfly’s Wings in Brazil Set Off a Tornado in Texas?” Conference presentation. An influential question about predictability, not a demonstration that a particular butterfly causes a tornado.

John H. Holland (1992), “Complex Adaptive Systems,” *Daedalus*, 121(1), 17–30. A broad entry point to interacting adaptive agents.

14 / Reproducibility note

The site is a no-server HTML/CSS/JavaScript experiment. Use the same seed, perturbation and rules to reproduce the same initial state and sequence. Changing the seed changes both worlds together; changing the perturbation changes only the centre cell in World B at reset. The program uses deterministic fixed update order and deterministic weather; no random draw is made during a tick.

The timeline retains up to 241 states (tick zero through tick 240) to bound memory and CPU. Export CSV before resetting if you need a longer record; no server-side records are kept. Saved snapshots can be exported as JSON. They are not automatically persisted across page reloads.

A caveat about the interface: selecting a saved snapshot starts a fresh replay branch at that recorded tick. The interface cannot import the JSON export; it is a portable inspectable record. The downloadable PDF is this companion guide, not a report automatically generated from your current run.

Creative project author: Sahil Tayal · Sahilverse

Canonical entrance: sahiltayal.com · Contact: sahil@sahiltayal.com

Construct systems that invite inspection. Resist the temptation to confuse a beautiful model with the world.