The Markov Brain: Rewiring brain using Markov probabilities

Disclaimer

This article was generated with the assistance of AI tools — specifically Claude AI (Anthropic) and Google Gemini NotebookLM — and subsequently reviewed and edited by the author. It is the author's belief that the latest scientific breakthroughs in neuroscience, when understood through mathematical frameworks like Markov chains and probability theory, can offer meaningful insights for forming new habits and pursuing self-improvement. The author has a genuine personal and intellectual interest in this subject but does not claim any scientific expertise in Markov chain theory, computational neuroscience, or related disciplines. Readers interested in clinical applications or academic research should consult the original peer-reviewed sources cited at the end of this article and seek guidance from qualified professionals where appropriate.

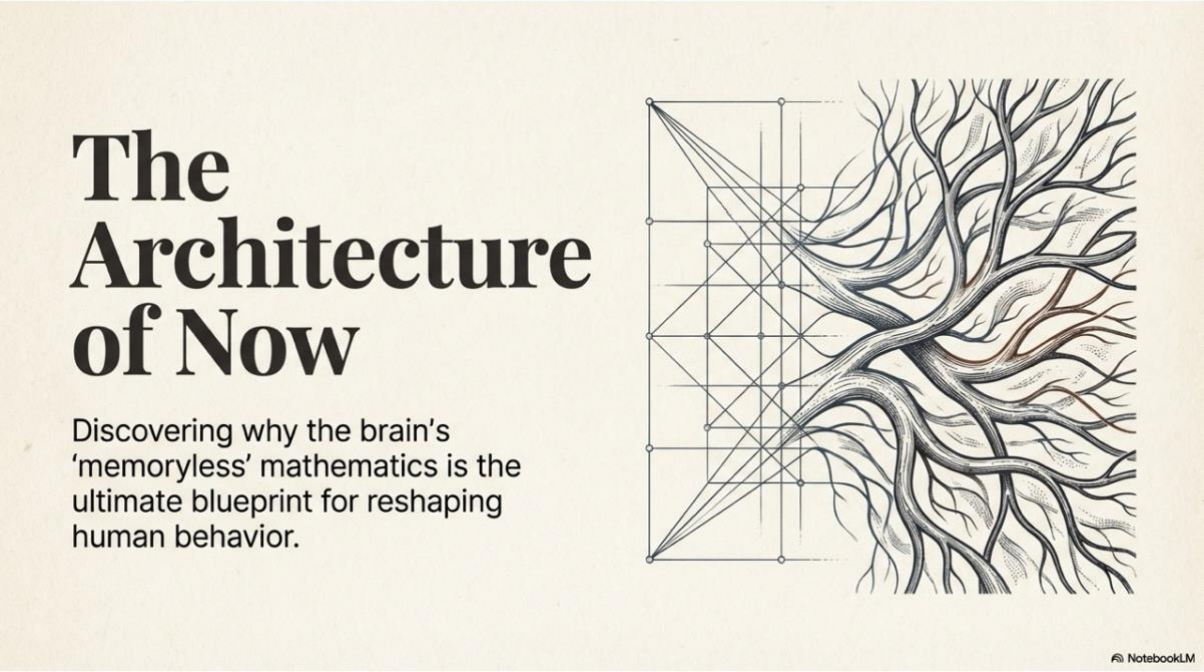

Right now, your brain is doing something specific. Billions of tiny cells are firing in a particular pattern. A split second from now, that pattern will change. And here is the surprising part: which pattern comes next may depend *only* on the pattern you are in right now — not on anything that happened before.

This idea comes from a branch of mathematics called **Markov chains**, named after the Russian mathematician Andrey Markov, who described them in 1906. They have become one of the most useful tools in brain science.

But the idea raises an uncomfortable question: **if only the present moment matters, then what is the point of memory? Why do we remember anything at all?**

The answer turns out to be both surprising and deeply practical. It connects brain science to some of the biggest questions in philosophy — and it offers a concrete, usable framework for changing your life through small, everyday actions.

What Is a Markov Chain? (As Simple as Possible)

Imagine you live in a small country with only three towns: **Calm Town**, **Alert Town**, and **Restless Town**. Every morning you wake up in one of these towns. Where you wake up tomorrow depends *only* on where you are today — not on where you were last week.

If you are in Calm Town today:

- There is a 70% chance you will still be in Calm Town tomorrow
- A 20% chance you will be in Alert Town
- A 10% chance you will be in Restless Town

These chances can be written as a simple table:

Where you are today	Calm tomorrow	Alert tomorrow	Restless tomorrow
Calm	70%	20%	10%
Alert	15%	65%	20%
Restless	25%	35%	40%

This table is called a **transition matrix**. It is the engine of the whole system.

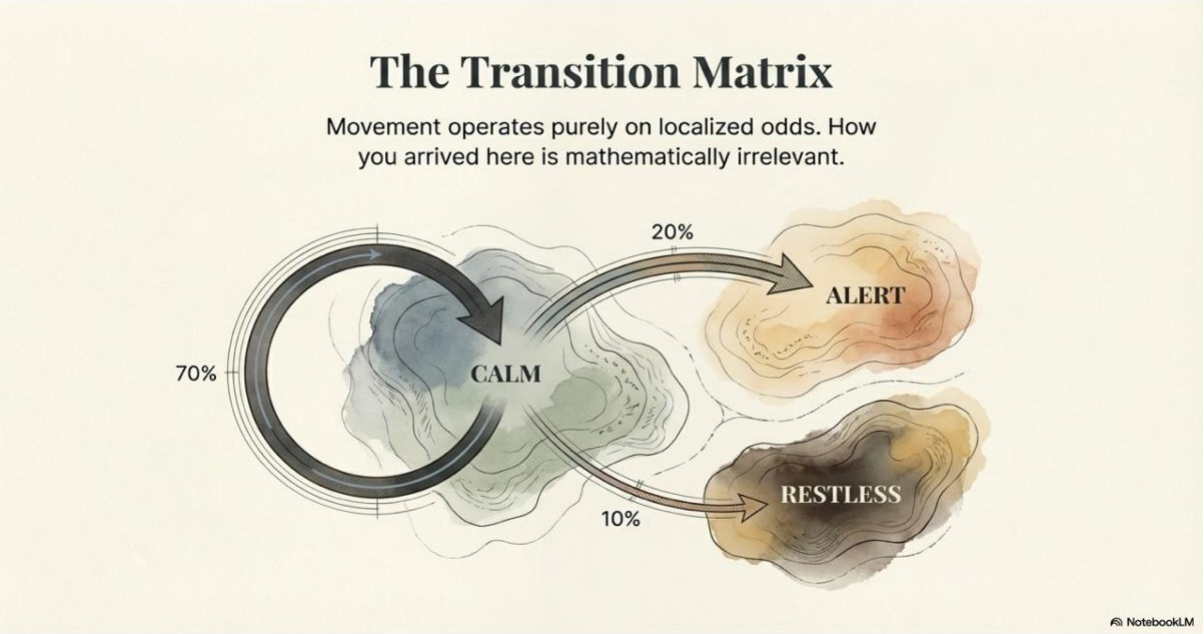

Notice something about the diagonal — those are the "staying put" numbers. Calm Town is very stable (70% chance of staying). Alert Town is fairly stable (65%). Restless Town is shaky (only 40% chance of staying — you are likely to move somewhere else soon).

The one rule that makes this a Markov chain is simple: **where you go next depends only on where you are now.** Your entire travel history — where you were last month, last year, ten years ago — does not matter. Only today matters.

This rule is called the **Markov property**. In technical language, the system is "memoryless."

Where Do Markov Chains Show Up in the Brain?

Everywhere. From the smallest parts of a single brain cell to the behaviour of the whole brain.

The Tiniest Level: On-Off Switches in Your Neurons

Every neuron has thousands of tiny gates called **ion channels**. Think of them as microscopic doors. Each door can be Open or Closed. It flips between these two states randomly. The chance of flipping depends only on whether the door is currently open or closed — not on how long it has been in that position.

This is the simplest possible Markov chain: two states, two probabilities. Billions of these tiny doors, each flipping independently, produce the electrical signals that carry your every thought.

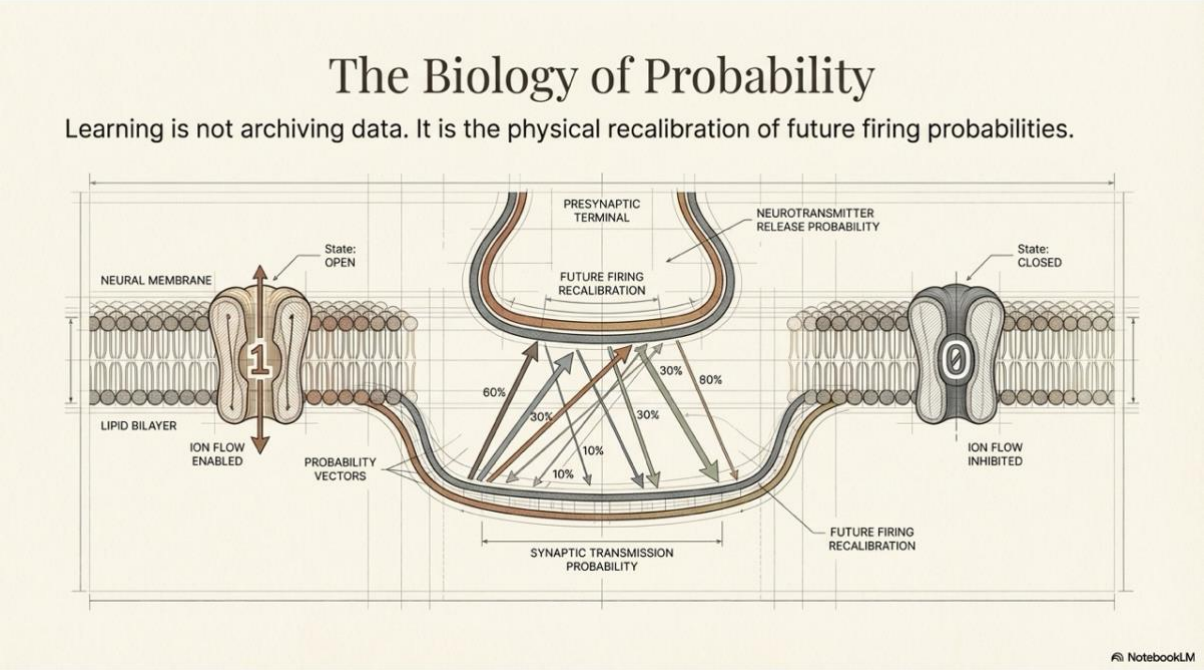

The Connection Level: How Neurons Learn

Neurons talk to each other across tiny gaps called **synapses**. These connections can get stronger (when you practise something) or weaker (when you stop). In extreme cases, connections get completely eliminated — the brain literally prunes wiring it does not need.

Scientists have shown that the way these connections change over time follows Markov rules. The future strength of a connection depends on its current strength and what the neurons are doing right now — not on the full history of that connection.

This is what learning looks like at the physical level. When you practise the guitar, you are strengthening certain synaptic connections. When you stop practising, those connections weaken. The brain is constantly rewriting its own wiring — and that rewriting follows the Markov property.

The Whole-Brain Level: Your Brain Has "Modes"

This is the most exciting part. Using brain scanners (fMRI and MEG), scientists have discovered that your entire brain switches between a small number of distinct "modes" or states. Each mode is a specific pattern of which brain regions are talking to each other. These modes switch every few seconds, and the switching follows Markov rules.

The mathematical tool used here is called a **Hidden Markov Model** (HMM). "Hidden" because you cannot see the brain's mode directly — you can only figure it out from the signals the scanner picks up.

Here is what researchers have found:

Healthy brains switch smoothly between several modes. Think of a person who moves comfortably between being focused, being relaxed, daydreaming, and being alert — flowing naturally from one to the next.

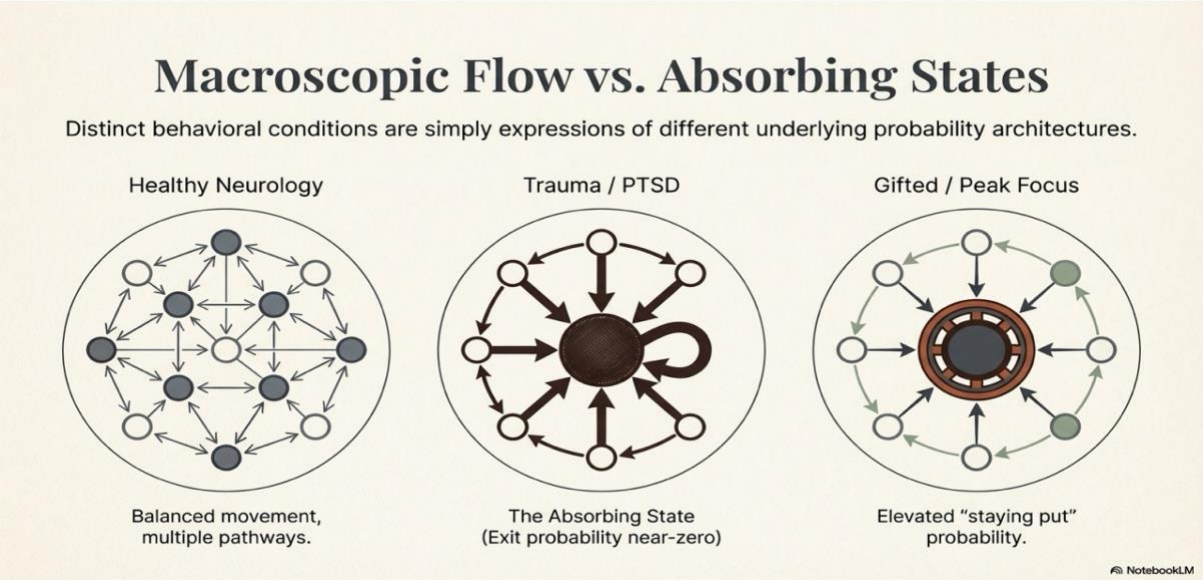

Brains affected by PTSD get *stuck*. Scientists found that the PTSD brain develops what is called an **absorbing state** — a mode it cannot easily leave. The brain locks into one pattern and cannot switch out of it. This matches exactly what PTSD feels like from the inside: being trapped in a loop of anxiety and hypervigilance, unable to relax or shift your attention.

Brains of exceptionally gifted students show yet another pattern. During complex problem-solving, gifted teenagers can *hold* a productive brain mode for longer than average. Their "staying put" probability for the focused mode is higher. They also switch *into* the focused mode more easily from other modes.

These are not metaphors. These are actual numbers computed from real brain scans.

The Prediction Level: Your Brain Guesses What Comes Next

Your brain is always trying to predict what will happen next. When you hear a familiar song, your brain predicts the next note before it plays. When the prediction is wrong — say, a strange note plays instead — your brain produces a little jolt of surprise, detectable on an EEG machine.

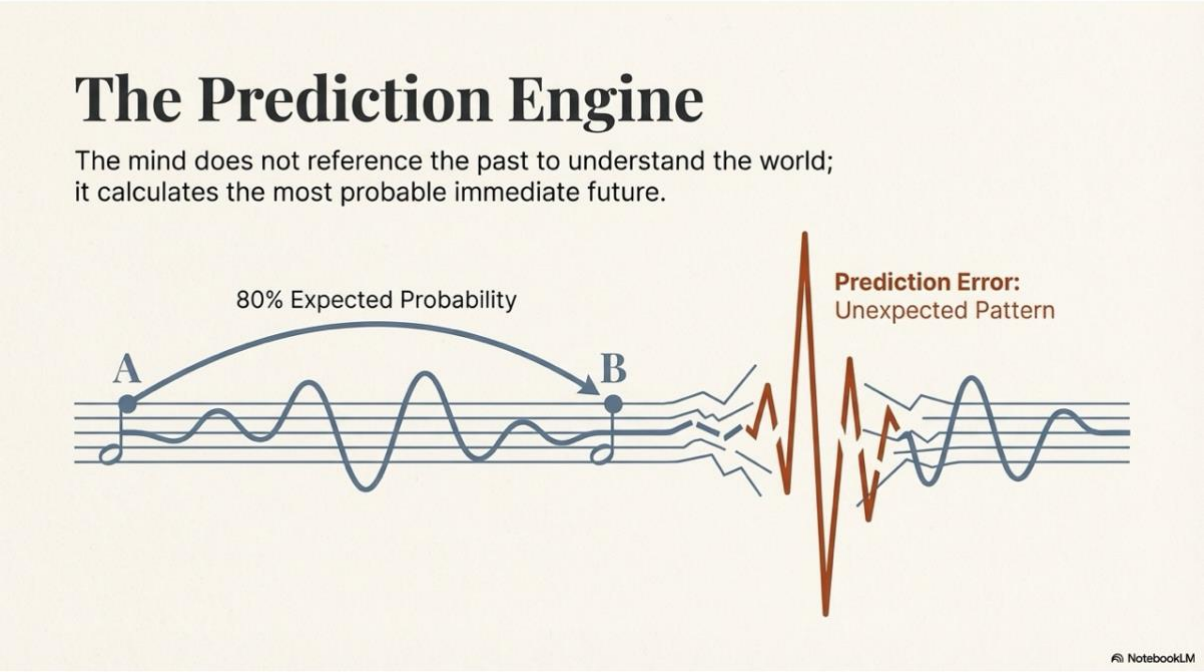

A 2023 study showed that the brain's prediction system works like a Markov transition matrix. The brain has learned: "After note A, note B usually follows with 80% probability." When note C plays instead, the brain registers a prediction error. The brain, it turns out, has built its own internal Markov chain as a model of the world.

If Only the Present Matters, Why Do We Have Memory?

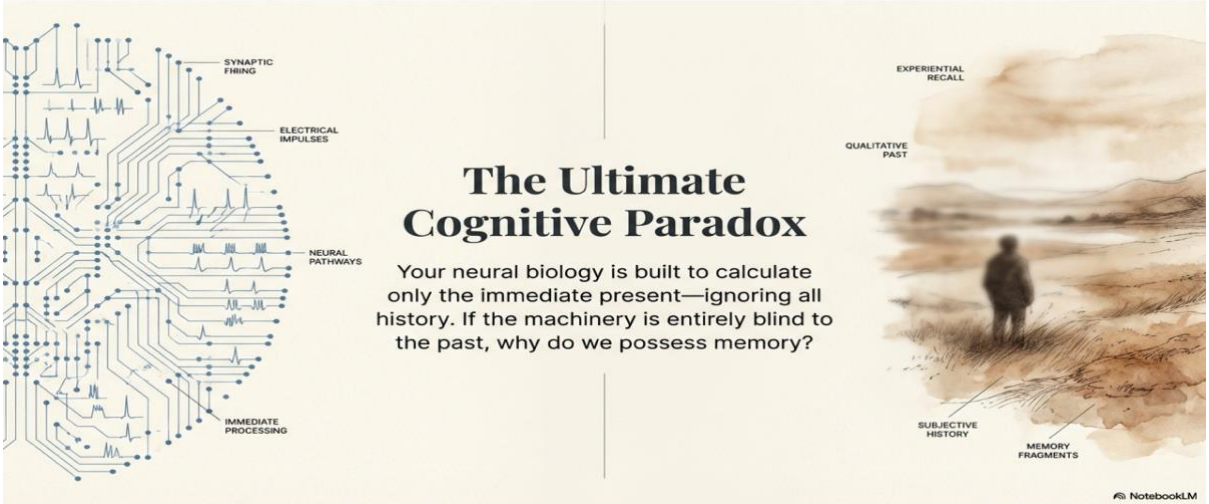

This is the question at the heart of the whole essay. If the Markov property says the future depends only on the present, does that mean our memories are useless?

No. It means the opposite. Here is why.

Your Memory Has Already Become Your Present

Think about two people sitting side by side on a park bench. Same bench, same weather, same moment. But one person spent twenty years as a soldier, and the other spent twenty years as a monk. Are they in the same "state"? Obviously not. Their brains have been physically shaped by completely different experiences. Their synaptic connections, their neural pathways, their transition probabilities — all different.

The past is not stored in a separate drawer. It has been physically built into the current structure of the brain.

Every experience you have ever had — every book you read, every heartbreak you endured, every skill you practised — has changed the wiring of your brain. It has altered which states your brain can enter and how easily it moves between them. Your past is not a recording that plays back. It is the architecture that determines how your present moment works.

The Markov chain is "memoryless" not because memory does not matter, but because **memory has already done its work**. It has shaped the transition matrix itself. The past is not consulted at each step because the past is already baked into the structure.

The River Analogy

A river's flow at any point depends on the shape of the riverbed right here, right now. The river does not "remember" the rainstorms of ten years ago. But the riverbed *is* the result of those rainstorms. Every flood, every drought, every season of erosion carved the contours that now guide the water.

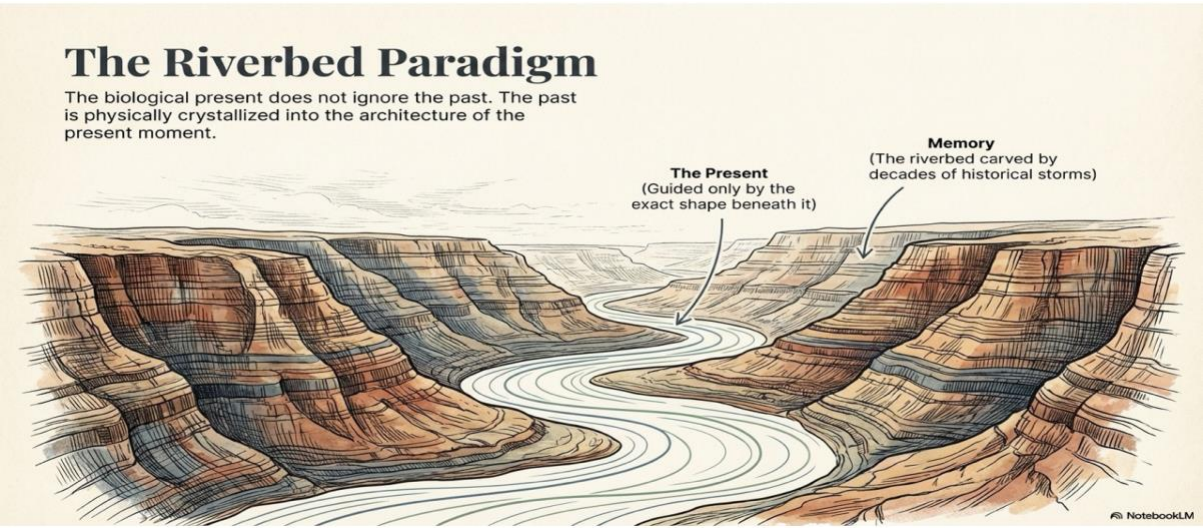

Your brain is the riverbed. Your experiences are the floods and droughts. The Markov property is the flow of water. The river does not need to remember because the memory is the riverbed itself.

What Memory Actually Does

In a Markov brain, memory plays four roles:

It rewrites the probability table. Learning something new changes the chances of your future state transitions. A child who touches a hot stove rewires the probability of "reach toward stove" from high to very low.

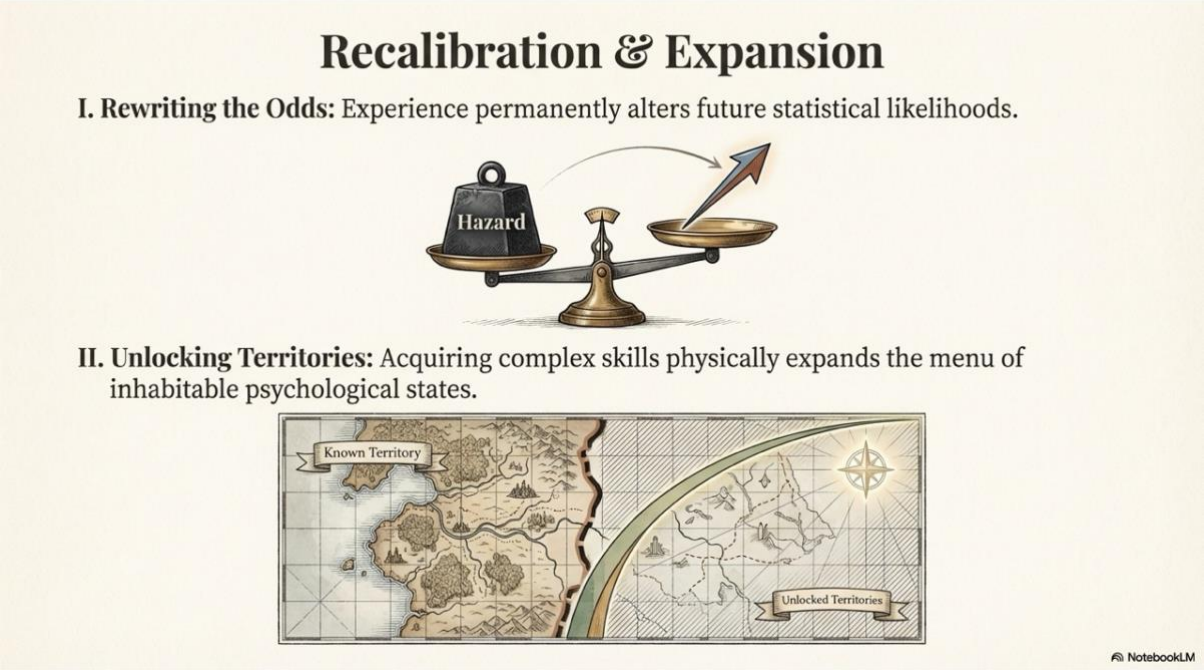

It creates entirely new states. A person who learns to read has access to brain states that an illiterate person simply does not have. Experience does not just shuffle the probabilities — it expands the menu of possibilities.

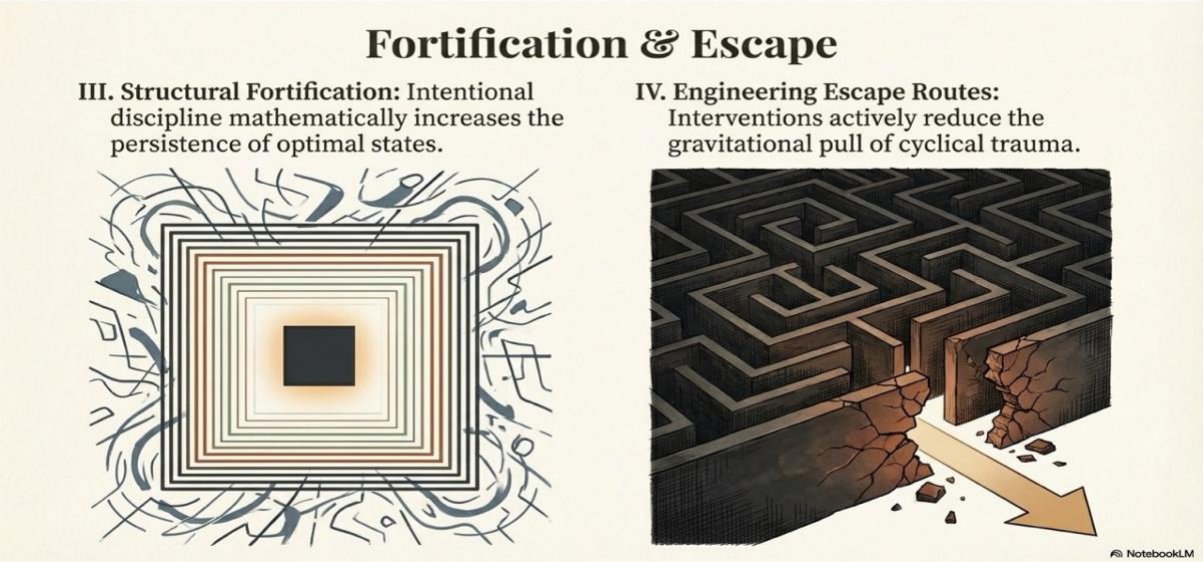

It makes good states last longer. Practice increases the "staying put" probability of useful states. A trained meditator can sustain a calm, focused state much longer than a beginner. The diagonal number in their transition matrix for "focused calm" is higher.

It can break you out of traps. Therapy for PTSD works, in Markov terms, by reducing the "staying put" probability of the stuck state. The goal is not to erase the trauma but to make the brain flexible again — able to leave the loop and move to other states.

The Paradox Solved

The seeming contradiction — "memoryless chains in a creature built for memory" — disappears when you understand what "present state" really means. Your present brain state is not a blank slate. It is the living summary of your entire life. The Markov chain does not ignore the past. **The past matters so much that it has literally become the present.**

Existentialism and the Markov Brain: You Are Not Your Past

The Markov idea — "only the present determines the future" — is not just a mathematical trick. It is the central message of existential philosophy, the tradition that includes Sartre, Camus, Heidegger, and Kierkegaard.

Sartre: You Are Free Right Now

Jean-Paul Sartre's most famous idea is that **existence precedes essence**. In plain language: you are not defined by your history, your genes, your upbringing, or your resume. You are defined by what you do *right now* and what you choose *next*.

Look at the Markov chain. At every state, the system faces a spread of possible next states. It is not locked into one path. The transition matrix gives *tendencies*, not certainties. There is always some probability, even if small, of an unexpected jump — a sudden shift from Restless to Calm, from stuck to free.

Sartre called this **radical freedom**. The mathematics agrees: at every time step, multiple futures are possible.

But Sartre added something the maths does not capture: the *anxiety* of this freedom. The Markov chain transitions without stress. We transition while knowing we could have chosen differently. That awareness — the gap between who we were and who we are becoming — is what makes us human.

Sartre's "Bad Faith": Pretending You Cannot Change

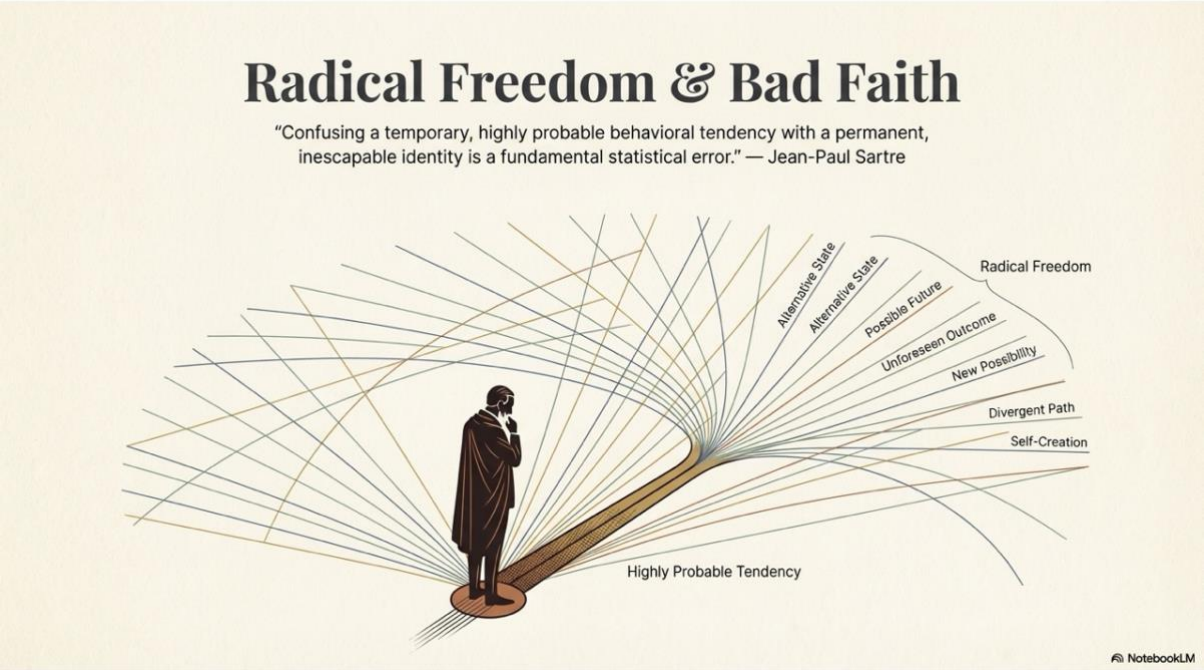

Sartre had a name for the habit of denying your own freedom: **bad faith**. Bad faith is telling yourself, "I am an angry person — that is just who I am." Or: "I have always been disorganised — I cannot change."

In Markov language, bad faith is the mistake of treating your current transition matrix as permanent. It is confusing a *tendency* with an *identity*. Yes, your transition matrix right now may make angry states more likely. But the matrix is not set in stone. Neuroplasticity — the brain's ability to rewire itself — means the matrix is always being rewritten by experience.

You are not your transition matrix. You are the process that can rewrite it.

Heidegger: Where You Start Is Not Where You End Up

Heidegger talked about **thrownness** — the fact that you find yourself in a situation you did not choose. You did not pick your family, your country, your body.

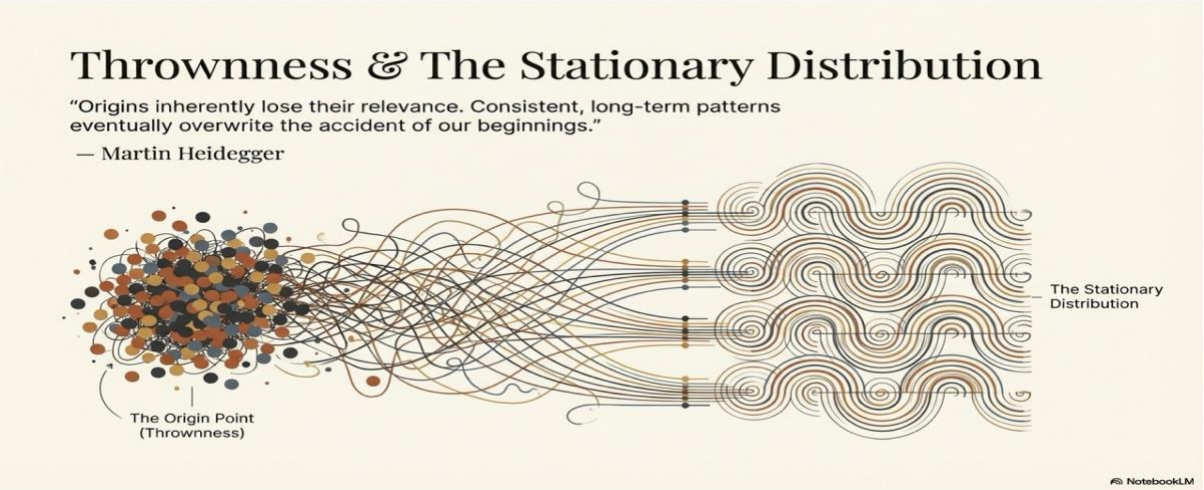

In a Markov chain, this is the starting state — X_0. You did not choose it. But here is the mathematical fact: as the chain runs, the starting state matters less and less. Over time, the system settles into its **stationary distribution** — the long-run pattern determined by the transition matrix, not by the starting point.

Translation: **where you begin matters less than how you transition.** The structure of your choices, over time, overwrites the accident of your origin.

Kierkegaard: The Leap That Changes Everything

Kierkegaard wrote about the **leap of faith** — a moment when someone jumps from one way of living to a fundamentally different one. Not through gradual steps, but through a single decisive act.

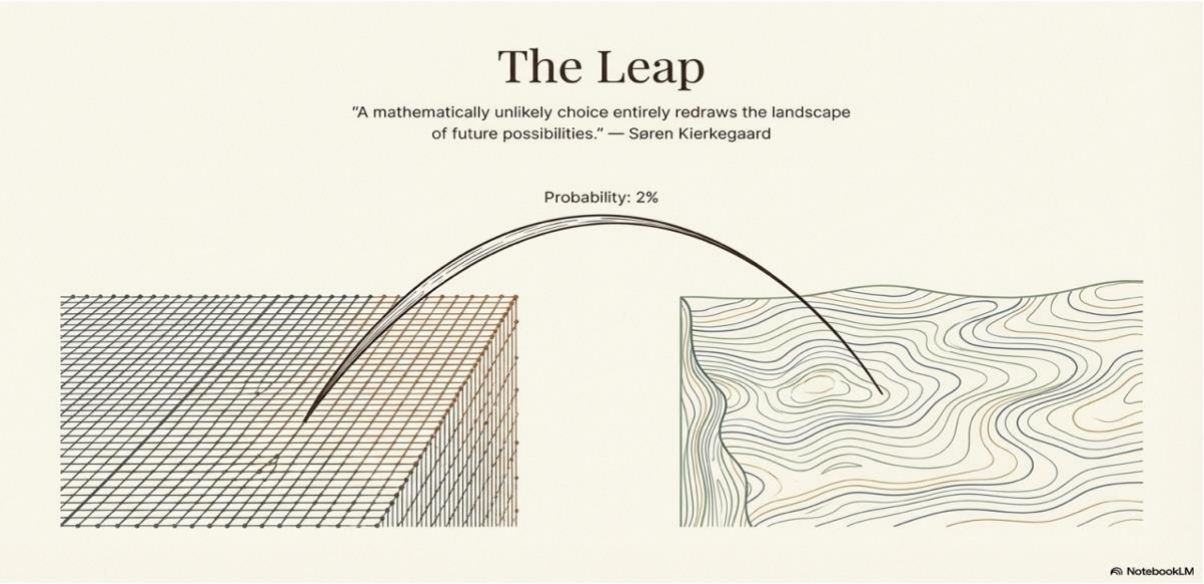

In Markov terms, this is a low-probability transition. Maybe there is only a 2% chance of jumping from "going through the motions" to "fully committed." But 2% is not zero. The leap is always on the table.

And here is the crucial insight: **once you make the leap, you face a different transition matrix entirely.** New states become possible. Old traps lose their grip. The leap does not just move you to a new place — it changes the entire landscape of what comes next.

Camus: Finding Meaning in the Loop

Camus imagined Sisyphus, condemned to push a boulder up a hill forever, watching it roll back down every time. This looks like the ultimate absorbing state — an endless, inescapable loop.

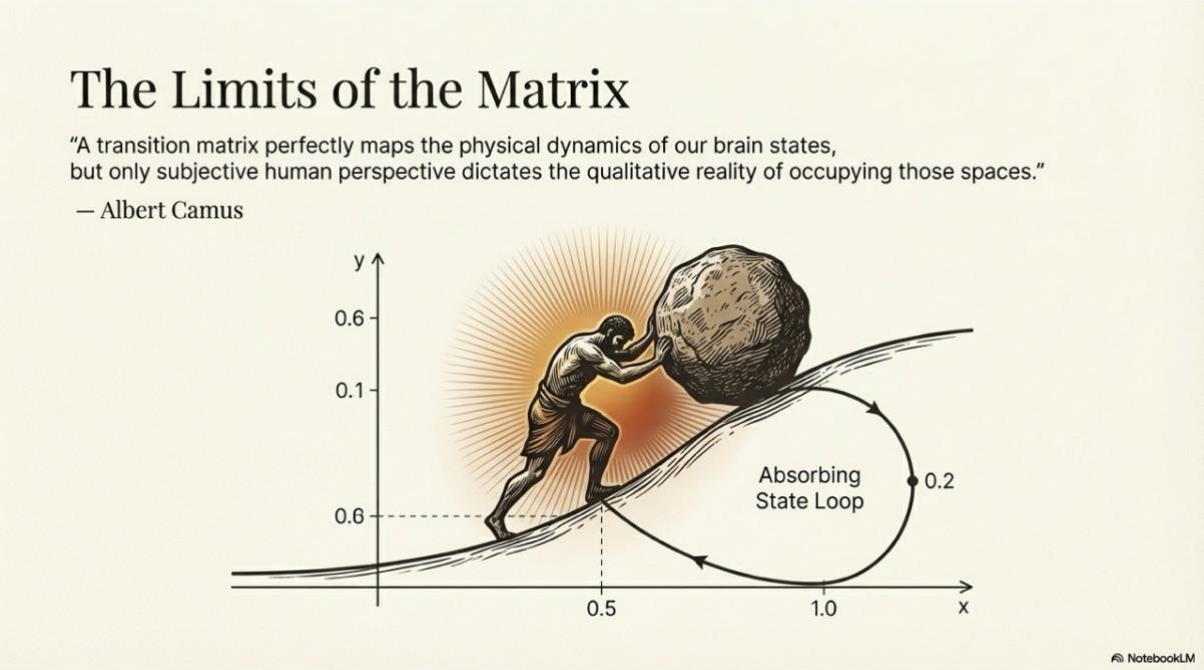

But Camus said: "One must imagine Sisyphus happy."

How? By changing the inner experience of the state without changing the state itself. The boulder still rolls down. But the person pushing it has found meaning in the act. The Markov chain sees the same loop. The human inside it has transformed.

This reveals something the mathematics alone cannot: **the felt quality of a state is not captured by the transition matrix.** Two people can occupy the same mathematical state and experience it completely differently. The numbers describe the dynamics. The philosophy describes what it is like to live them.

How Small Changes Reshape Your Future: A Practical Guide

Here is where the mathematics becomes genuinely useful for everyday life. The core insight is this: **you do not need to make dramatic changes to dramatically change your life. Small shifts in probability, repeated consistently, change everything.**

This is not motivational fluff. It is a mathematical property of Markov chains. Let me show you how it works.

How Small Probabilities Compound

Imagine your daily life has three main states: **Energised**, **Neutral**, and **Drained**. Your current transition matrix looks like this:

Today	Energised tomorrow	Neutral tomorrow	Drained tomorrow
Energised	50%	30%	20%
Neutral	20%	50%	30%
Drained	10%	30%	60%

Run this chain over a long time and you can calculate the **stationary distribution** — the average fraction of your days spent in each state. With these numbers, you end up spending roughly:

- **22% of days Energised**
- **36% of days Neutral**
- **42% of days Drained**

That is a life where you feel drained almost half the time. Now suppose you make *one small change* — you start going to bed 30 minutes earlier. This does not transform you. It just shifts a few probabilities by five percentage points:

Today	Energised tomorrow	Neutral tomorrow	Drained tomorrow
Energised	55%	30%	15%
Neutral	25%	50%	25%
Drained	15%	35%	50%

A small change. But the new stationary distribution shifts to roughly:

- **29% of days Energised** (up from 22%)
- **38% of days Neutral** (up from 36%)
- **33% of days Drained** (down from 42%)

You have gone from being drained 42% of the time to 33% of the time — gaining roughly **one extra good day every two weeks** — from a single habit change. And this is just one shift. Stack two or three small changes and the compound effect is striking.

This is the central practical lesson: you do not change your life by willpower on any single day. You change it by nudging probabilities, which changes where you spend your time over months and years.

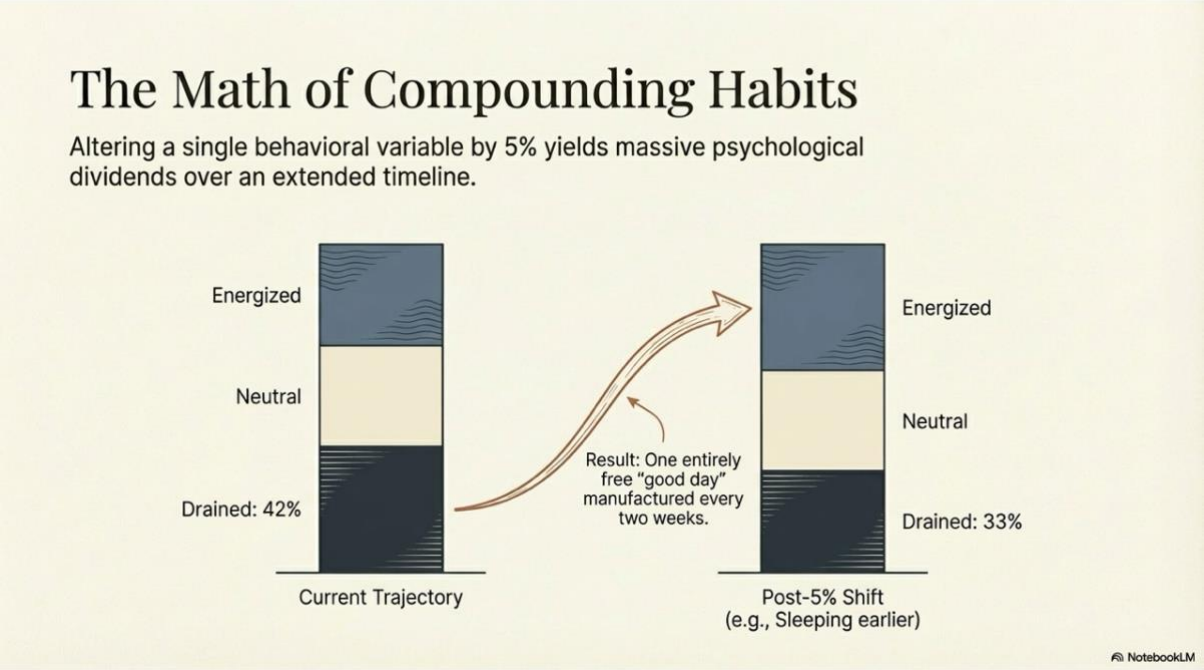

Seven Principles for Rewriting Your Transition Matrix

1. Start From Where You Actually Are

The Markov rule says the future depends on the present state — not on the state you wish you were in. If you are exhausted, the menu of next steps is different from when you are well-rested. Trying to force a transition that does not exist from your current state is like trying to catch a bus that does not stop at your station.

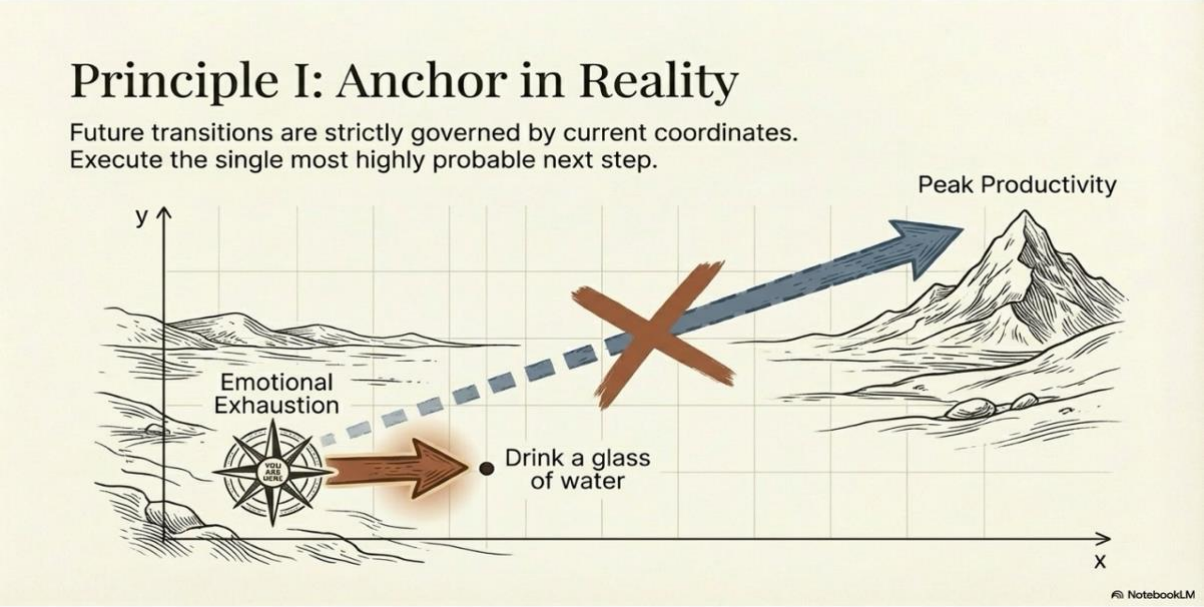

What to do: Before trying to change anything, honestly name your current state. Ask: "What is actually possible from *here*?" The best next step from "completely overwhelmed" is not "total productivity." It might be "drink a glass of water and do one small thing." That small transition opens the door to the next one.

2. Protect Your Good States

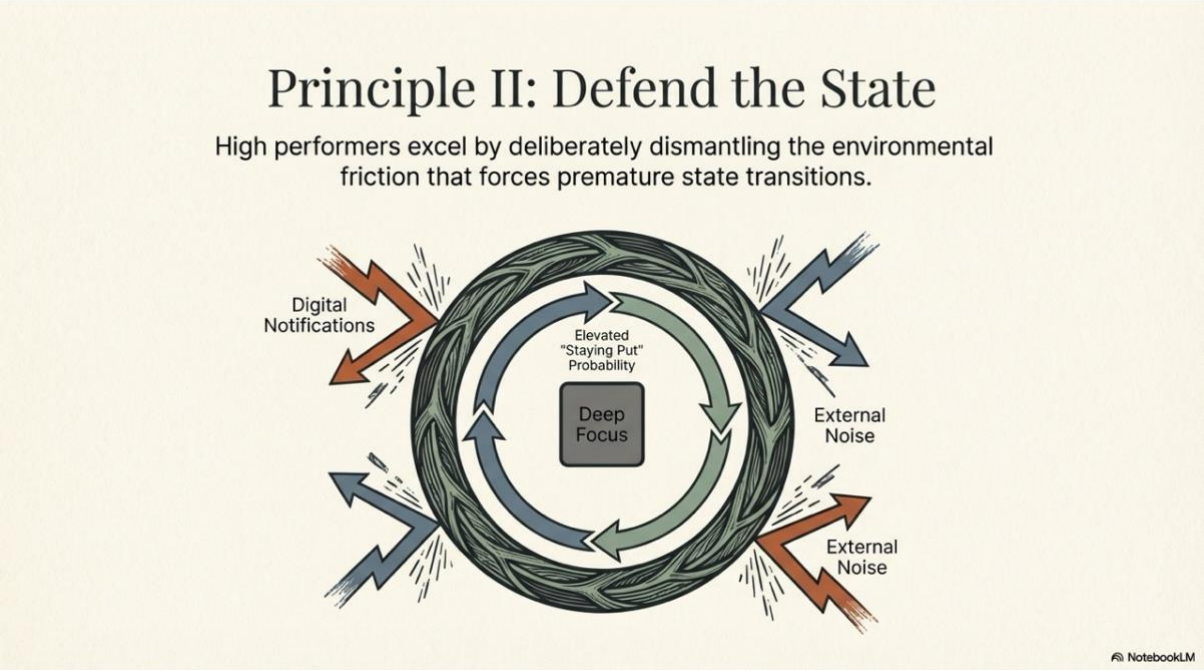

Research on gifted brains shows that high performers are not people who are always in a peak state. They are people who can *stay* in a peak state longer once they enter it. Their "staying put" probability is higher for productive modes.

You can engineer this. The key is to figure out what pulls you *out* of good states and reduce those interruptions.

What to do — specific examples:

- **If deep focus is your good state:** Turn off phone notifications, close email, use a "Do Not Disturb" sign. Each interruption is a forced state transition. Reducing interruptions from 10 per hour to 3 per hour means your "staying in focus" probability jumps significantly.
- **If calm contentment is your good state:** Notice what triggers anxiety spirals (often: news, social media, certain conversations). Reduce exposure during calm periods. You are not avoiding reality — you are protecting a valuable state from unnecessary disruption.
- **If creative flow is your good state:** Keep your tools ready. A writer who has to search for a notebook every time an idea strikes has added friction that lowers the self-transition probability of the creative state.

3. Build Stepping Stones, Not Giant Leaps

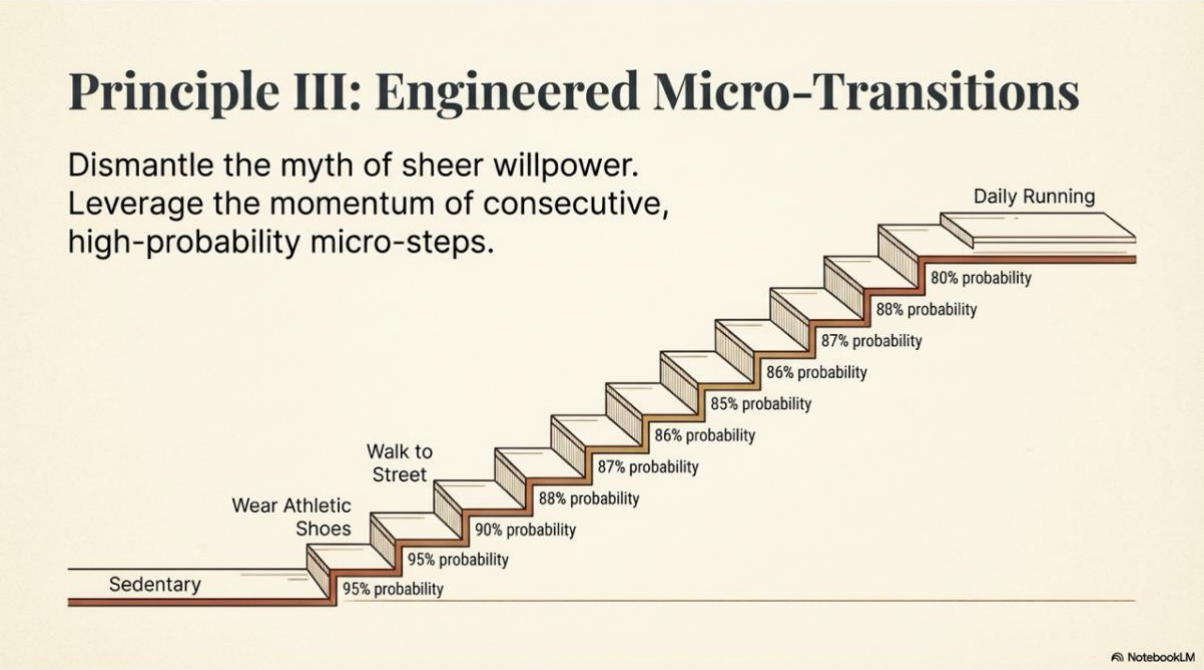

Most failed resolutions are failed because people attempt transitions with near-zero probability. "I will go from couch potato to running 5 km every morning" is a transition that barely exists in most people's matrix.

The solution is to create intermediate states that make the journey a chain of small, achievable steps.

What to do — a concrete example:

Suppose you want to go from "sedentary" to "regular exerciser." Here is the Markov approach:

- **Step 1 (Week 1-2):** From "sitting on the couch," transition to "putting on running shoes." That is it. Just put them on. This transition has a very high probability — it asks almost nothing of you.
- **Step 2 (Week 2-3):** From "shoes on," transition to "walking to the end of the street." Again, easy. High probability.
- **Step 3 (Week 3-5):** From "short walk," transition to "slightly longer walk." The momentum of the previous steps has already shifted your matrix.
- **Step 4 (Week 5+):** From "regular walker," transition to "occasional light jog." By now, your brain has physically rewired — the synaptic connections supporting this routine have strengthened. Your transition matrix has genuinely changed.

Each individual step has a high probability of success. The chain of steps achieves what a single giant leap could not.

4. Catch Absorbing States Early

An absorbing state is any state that feeds on itself — the longer you are in it, the harder it is to leave. Clinical examples include PTSD and severe depression. But everyday absorbing states are everywhere:

- **Doom-scrolling:** The algorithm is designed to make the "keep scrolling" probability approach 100%. After two minutes, leaving is easy. After forty minutes, it feels impossible.
- **Worry loops at 3 AM:** One anxious thought triggers another, which triggers another. The self-transition probability increases with each cycle.
- **Anger spirals in arguments:** Each angry response from the other person increases the probability of another angry response from you.
- **"I will start Monday" procrastination:** The delay itself makes further delay more likely, because nothing changes to shift the matrix.

What to do — build circuit breakers:

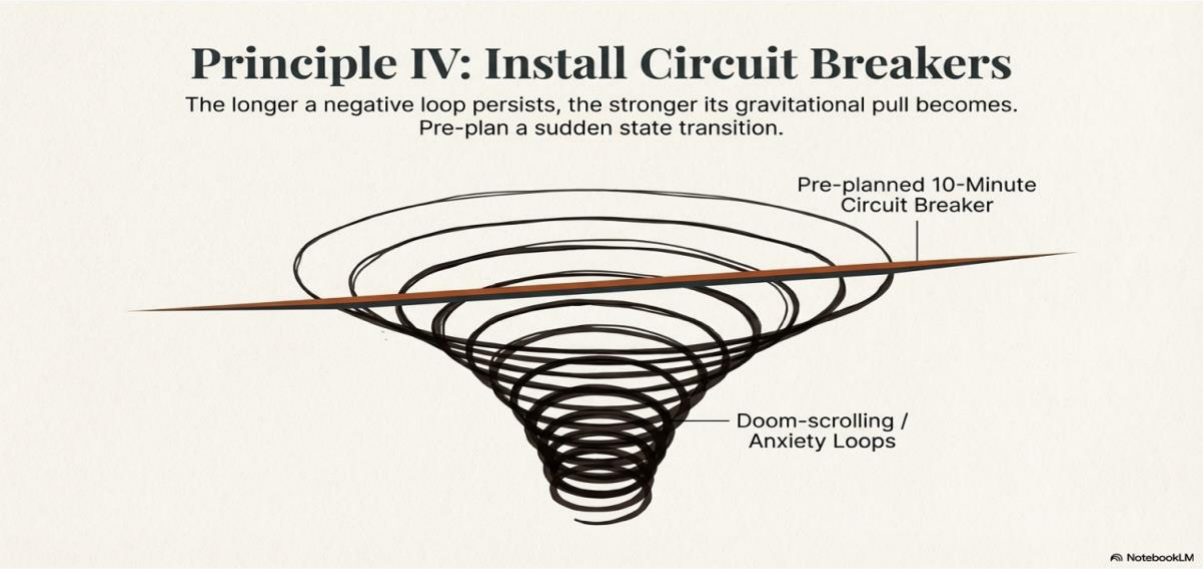

A circuit breaker is a pre-planned action that forces a state transition before the absorbing state locks in. The key word is *pre-planned* — you decide what to do *before* you need it, because once you are deep in the absorbing state, your ability to make decisions is compromised.

Specific circuit breakers:

- **For doom-scrolling:** Set a 10-minute timer before you open the app. When it rings, physically stand up and walk to another room. The physical movement forces a state change.
- **For 3 AM worry:** Keep a notebook by the bed. Write down the worry in one sentence. The act of writing externalises the thought, which weakens the internal loop. This is not about solving the problem — it is about breaking the self-reinforcing cycle.
- **For anger in arguments:** Pre-agree with your partner on a code word that means "we both take a 10-minute break." Leaving the room physically interrupts the chain.
- **For procrastination:** Use the "two-minute rule" — commit to working on the task for just two minutes. This is a stepping stone (Principle 3) combined with a circuit breaker. Two minutes is too short to resist, but often enough to shift your state into engagement.

5. Change the Matrix, Not Just Today

This is the most important principle and the one most people get wrong.

Every Markov chain settles into a **stationary distribution** — the long-run average of how much time you spend in each state. This distribution depends entirely on the transition matrix. If the matrix does not change, the distribution does not change. Period.

This means: **your good intentions do not matter if your underlying habits, environment, and routines stay the same.** You can resolve every January to be more focused, more calm, more productive. But if nothing changes in the structure of your daily life, you will drift back to the same average pattern — the same proportion of stressed days, wasted evenings, and distracted hours.

What to do: Stop focusing on individual days and start focusing on **structural changes** that alter probabilities:

Instead of...	Try...	Why it works
"I will not check my phone in bed" (willpower)	Charge your phone in another room (structure)	Removes the transition trigger entirely
"I will eat healthier" (intention)	Do not keep junk food in the house (environment)	Changes the probability of reaching for junk from high to near-zero

"I will exercise more" (resolution)	Sign up for a class at a fixed time with a friend (commitment)	Social expectation + schedule creates a high-probability transition
"I will be less anxious" (wish)	Start a daily 5-minute breathing practice (routine)	Gradually increases the self-transition probability of calm states
"I will stop procrastinating" (guilt)	Break the next task into a 2-minute first step (design)	Creates a high-probability transition into the "engaged" state

The left column tries to change a single transition by force of will. The right column changes the matrix itself. The left column fails after a few days. The right column compounds over months.

6. Trust the Compound Effect

Here is an encouraging mathematical fact: the stationary distribution of a Markov chain can be surprisingly sensitive to small changes in the transition matrix. You do not need to transform yourself overnight. You need to nudge a few probabilities by a few percentage points and then *let time do the work*.

Think of it like compound interest for behaviour:

- Going to bed 30 minutes earlier → slightly higher probability of waking rested → slightly higher probability of morning exercise → slightly higher probability of entering an energised state → slightly higher probability of making good food choices → slightly lower probability of evening stress → slightly higher probability of going to bed on time again.

Each link in this chain is a small probability shift. But the chain feeds back on itself. The energised state makes the next energised state more likely, which makes the next one more likely still. Over weeks and months, a small initial nudge can substantially shift how you spend your time.

What to do: Pick one — just one — small change that you believe will shift a probability in your favour. Do it consistently for four weeks before evaluating. Do not add anything else during that time. Let the ripple effects propagate through your transition matrix before adding another change.

Good candidates for a first change:

- Sleep time (shifts morning state probabilities)
- First thing you look at in the morning (phone vs. window vs. book — sets the initial state for the day)
- One meal per day (energy levels affect every downstream transition)
- Five minutes of stillness at a fixed time (builds the "calm" state's persistence)
- Walking (even ten minutes shifts the probability distribution of mental states for hours afterward)

7. You Are Not Your Current State

This is the existential principle made practical. Sartre called it overcoming "bad faith" — the habit of treating your current tendencies as your permanent identity.

The Markov framework makes this very clear: **your current state is temporary.** It has a probability of persisting and a probability of changing. You are not "a depressed person" — you are a person currently in a low state, with a transition matrix that makes low states more likely than you would prefer. That matrix is rewritable. It is being rewritten right now, with every choice you make.

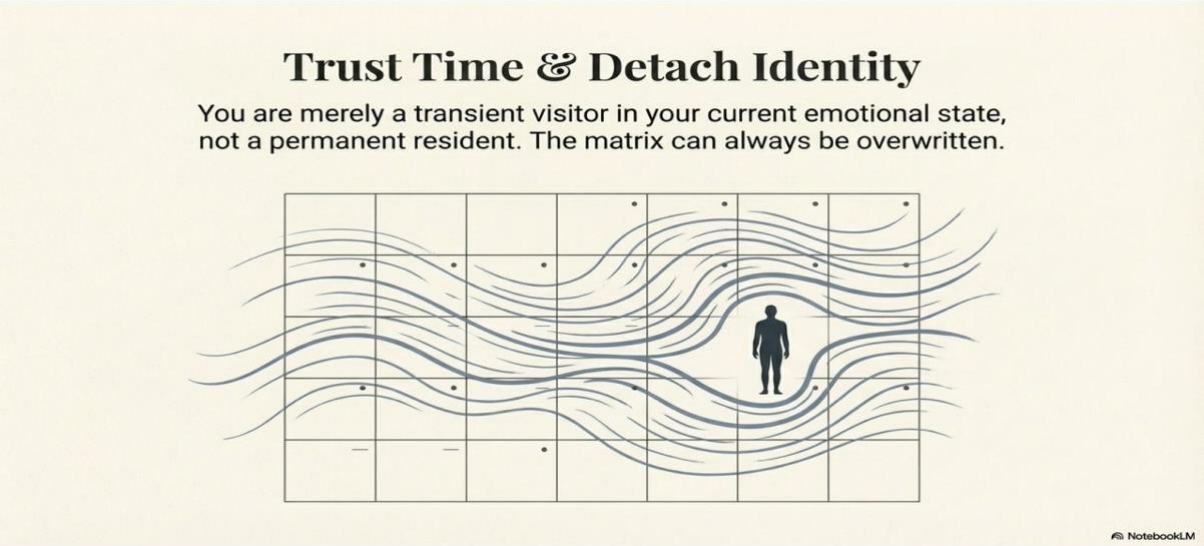

What to do: When you are stuck in a difficult state — anxiety, lethargy, frustration, hopelessness — try saying to yourself: "This is a state I am visiting, not a place I live. The matrix makes it likely I will be here for a while, but there is always a non-zero probability of moving somewhere better. My job is not to escape this instant. My job is to make small changes so that, over time, I visit this state less often."

This is not toxic positivity. It is not pretending to be fine. It is a mathematically grounded reframe: states are temporary, matrices are rewritable, and small consistent changes shift the long-run distribution.

Where the Traditions Meet

The Markov brain sits at a place where mathematics, brain science, philosophy, and ancient wisdom all point the same way.

The **Bhagavad Gita** teaches Nishkama Karma — act fully in the present, without clinging to past results or grasping at future outcomes. The Markov property says the same: the best action depends only on your current state.

Sartre insists that you are free at every moment, defined not by your past but by what you choose now. The transition matrix agrees: the past has been absorbed; only the current spread of possibilities matters.

Nietzsche's amor fati — loving your fate, embracing the present moment including everything that led to it — mirrors the Markov chain's relationship with its own history: not resisted, not replayed, but fully absorbed into the structure of now.

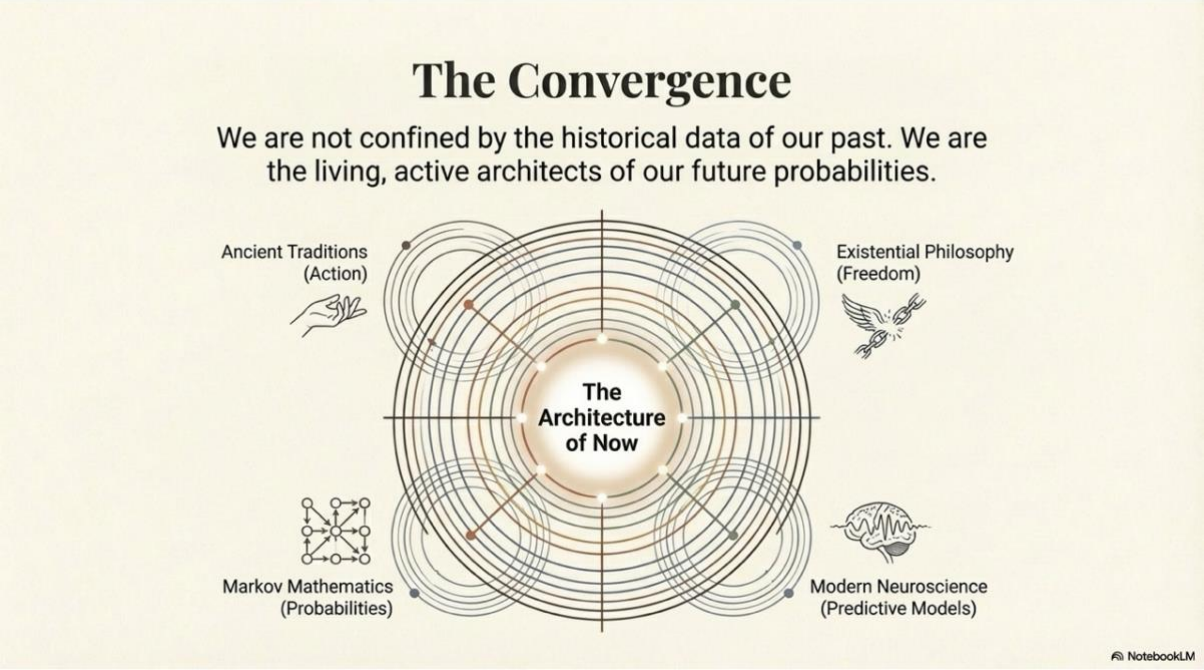

The **Yoga Sutras** describe pratyahara — turning inward, away from the noise of accumulated impressions — as the gateway to deep concentration. In Markov terms, this is the practice of raising the "staying put" probability of a focused state by reducing the distractions that trigger transitions away from it.

And **Friston's Free Energy Principle** — perhaps the boldest theory in modern brain science — says the brain is a system that constantly updates its internal Markov models to better predict the world. Learning, in this view, is the lifelong rewriting of transition matrices through lived experience.

The mathematics and the philosophy converge. The brain science confirms. And the practical wisdom follows:

You are not your history. You are not your current state. You are the living process that rewrites the matrix — one small probability at a time.

Summary

Level	What Is Modelled	Key Insight
Ion channels	Open/Closed switching	The smallest unit of brain computation is a two-state Markov chain
Synapses	Connection strength changes	Learning is the rewriting of transition probabilities
Brain networks	Hub identification	Random walks on brain wiring reveal communication centres (same maths as Google PageRank)
Whole-brain states	Mode switching (HMM)	Healthy brains switch modes fluidly; disordered brains get trapped
Prediction	Sensory expectations	The brain's model of the world is a probability table
Self-improvement	Habit and state change	Small shifts in probability, sustained over time, reshape the long-run pattern of your life

Further Reading

- Schröger et al., "Markov chains as a proxy for the predictive memory representations underlying mismatch negativity," *Frontiers in Human Neuroscience*, 2023
- Ezaki et al., "Modelling state-transition dynamics in resting-state brain signals," *European Journal of Neuroscience*, 2021
- Niu et al., "EEG source-space synchrostate transitions and Markov modeling in the math-gifted brain," *Human Brain Mapping*, 2020
- Chen et al., "Characterizing and differentiating brain state dynamics via hidden Markov models," *Brain Informatics*, 2015
- Menon et al., "Uncovering hidden brain state dynamics that regulate performance," *Nature Communications*, 2018
- Arizumi et al., "A Markov chain model of the evolution of complex neuronal network structures in the presence of plasticity," *BMC Neuroscience*, 2010
- Sartre, *Being and Nothingness*, 1943
- Friston, "The free-energy principle: a unified brain theory?" *Nature Reviews Neuroscience*, 2010
