

WHAT THE WATER CARRIED

Steven Foerster

Contents

Prologue: The Fountain and the Aquifer	2
Chapter 1 — Baseline	12
Chapter 2 — Something Is Different	31
Chapter 3 — Characterize the Compound	46
Chapter 4 — The Dose Curve	61
Chapter 5 — She Tells Yusuf Everything	77
Chapter 6 — The Metabolite Problem	96
Chapter 7 — Peak	111
Chapter 8 — What IRIS Learns	131
Chapter 9 — The Population Clock	149
Chapter 10 — The Intervention Hypothesis	166
Chapter 11 — Going Public (Partially)	181
Chapter 12 — The Self-Study	197
Chapter 13 — The Third Compound	209

Chapter 14 — Yusuf’s Curve	226
Chapter 15 — The Veridian Leak	240
Chapter 16 — What the Response Reveals	252
Chapter 17 — The Memory Becomes Unreliable	264
Chapter 18 — Osei Comes to the Lab	274
Chapter 19 — The Preprint Lands	284
Chapter 20 — The City in Decline	294
Chapter 21 — The Conversation	302
Chapter 22 — What Yusuf Carries	312
Chapter 23 — The Thing She Can No Longer Hold	320
Chapter 24 — What She Hears	331
Chapter 25 — What IRIS Has Become	340
Chapter 26 — The Final Audit	347
Chapter 27 — Osei Returns	353
Chapter 28 — The Question She Deferred	359
Chapter 29 — The Leap	363
Chapter 30 — After	369

“This intellect is separable, impassible, unmixed, since it is in its essential nature activity... this alone is immortal and eternal.” — Aristotle, De Anima III.5

“If men do not comprehend the character of God, they do not comprehend themselves.” — Joseph Smith, King Follett Discourse, 1844

“The glory of God is intelligence, or in other words, light and truth.” — Doctrine and Covenants 93:36

“Intelligence, or the light of truth, was not created or made, neither indeed can be.” — Doctrine and Covenants 93:29

“I don’t know what’s worse: to not know what you are and be happy, or to become what you’ve always wanted to be and feel alone.” — Daniel Keyes, Flowers for Algernon

* * *

PROLOGUE: THE FOUNTAIN AND THE AQUIFER

THERE is a story about this place that everyone knows and nobody believes. This describes most important stories, I've found. It certainly describes this one.

In 1513, Juan Ponce de León sailed west from Puerto Rico and made landfall somewhere on the eastern coast of a peninsula he named *La Florida*, the flowery one, because he arrived during *Pascua Florida*, the Easter season of flowers, and because Spanish conquistadors named things the way tenured professors title papers: with more confidence in the framing than the content warranted. The story that attached itself to his voyage, decades after the fact and with the embellishments that accrue to any good narrative left unattended, was that he was searching for a fountain whose waters could reverse the effects of age. The Fountain of Youth. You have heard of it. You do not believe in it. Neither did most of the people who passed the story along, which did not stop them from passing it along, because the story was never really about belief. It was about wanting. And wanting doesn't require belief. Wanting, in my experience, doesn't even

require permission.

Historians have spent five centuries disputing whether Ponce de León actually believed any of this. The consensus — and I have read the primary sources, because I have a mother who taught high school English for thirty-one years and raised me to believe that going to the original text is not optional, and a father who was a civil engineer and raised me to believe that you check load-bearing walls before you admire the façade, and between them they produced a daughter who reads footnotes recreationally, which is either a very specific kind of inheritance or a very specific kind of personality disorder — the consensus is that the fountain quest was probably projected onto him by later chroniclers. Chiefly Gonzalo Fernández de Oviedo, writing in 1535 with the storyteller's advantage of hindsight and the good sense to know that a quest makes better copy than a land survey. Ponce de León was most likely looking for what Spanish explorers generally looked for: gold, land, and the political leverage that came from claiming both for the crown. The fountain was better narrative. It usually is.

It doesn't matter. What matters is that the story survived, and the story survived because it names something true underneath the facts it gets wrong. I have learned to pay attention to this: the moments when a story's persistence outweighs its accuracy, when a narrative survives not because it happened but because it *should* have, because the wanting underneath it is so fundamental that the story becomes load-bearing regardless of whether it's

true. My mother would call this the difference between fact and meaning. My father would call it the difference between the blueprint and the building. I call it the thing I missed about my own story for thirty-five days, which is a long time to miss the obvious when the obvious is in your glass every morning, but in my defense I was very busy being the smartest person I had ever been, which turns out to be exactly the condition under which you are most likely to miss the thing that is right in front of you.

But I'm getting ahead of myself. I do that now. The chronology is less reliable than it used to be, for reasons I will explain, and which are among the reasons I am writing this down while I can still get the order right — or close to right, or right enough, which is a distinction I have recently learned to make peace with, and which would have horrified me a month ago, when I still had the cognitive resources to be a perfectionist about things that don't require perfection.

People come to Florida looking for restoration. They always have. The retirement communities along both coasts, the wellness spas, the per-capita density of cosmetic dermatology practices — the numbers are genuinely extraordinary; I looked this up once during a conversation with my colleague Yusuf about why Gainesville felt different from the rest of the state, and he pointed out that the looking-it-up was itself a very Gainesville response to the question, which was fair — the entire economic and cultural infrastructure of this peninsula is organized around the premise that decline is a problem to be solved rather than a condition to

be endured. Not accepted. Not accommodated. *Solved*, the way you solve a drainage problem or a funding gap or a bad knee: with resources, expertise, and the specifically American conviction that deterioration is a personal choice you simply haven't optimized your way out of yet.

Ponce de León would understand this instinctively. He might have trouble with the turnpike.

I live in Gainesville. Sixty miles inland, in the part of Florida that tourists drive through on their way to somewhere with a beach. Home of the University of Florida, the Gators, and the Hare Krishna temple on NW 14th that serves free lunch on the Plaza of the Americas every Wednesday with a sincerity I find genuinely moving and a lentil dal I find transcendent. Population 140,000 in the quiet months, 180,000 when the students return in August and the coffee shops remember what a line looks like. Built on limestone karst that is among the most extraordinary geological formations on the continent and that almost nobody thinks about, because it is underground, and because people have a limited appetite for being impressed by rock.

The rock is worth being impressed by.

The Floridan Aquifer System sits beneath most of the state, limestone laid down sixty million years ago when this peninsula was the floor of a shallow Eocene sea, carved and dissolved over geological time into a network of conduits and chambers that now holds an almost inconceivable volume of fresh groundwater. The springs that surface from this aquifer, Ichetucknee, Ginnie, Silver,

dozens of others, run at a constant sixty-eight degrees regardless of season, so clear you can see forty feet to the limestone bottom, so old that the water reaching the surface today entered the rock before the oldest building in the Western Hemisphere was conceived. The Timucua people who lived along these springs for thousands of years before any European arrived considered them sacred, which is a word I don't use casually. I am a scientist. I deal in mechanisms, not mysteries. But I have floated in Ichetucknee on a Sunday morning in January with the water at sixty-eight degrees and the air at forty and the current carrying me over rock that was seafloor before there were primates, and I will tell you that *sacred* is either exactly the right word for that experience or there is a gap in the English language where the right word should be, and having been raised by a woman who taught English for three decades I do not make that accusation lightly.

My mother would know the word. She always does. She has the language for experiences the way I have the language for mechanisms. Between us we could probably describe most things accurately, which is not the same as completely, and the distance between those two words, *accurately* and *completely*, is something I am learning about later than I should have, and under conditions I would not have chosen.

Gainesville draws its municipal water from this aquifer. It arrives at the tap tasting clean, faintly mineral, with a sweetness that is the signature of water filtered through limestone for longer than

human civilization has existed. I drank it every day for twelve years without noticing, the way you stop noticing the weight of your own skeleton: not because it isn't extraordinary but because it is so *constantly* extraordinary that your nervous system, in its ruthless efficiency, stops spending attention on it.

This is the most important thing I study. Not the water — the stopping. The mechanism by which a brain decides what to keep attending to and what to let fade into the architecture of the ordinary. The tradeoff between holding on and letting go, between stability and plasticity, that every learning system in existence negotiates every second of its operation, so quietly, so well, that you have never once in your life noticed it happening. You are doing it right now, reading this sentence, your brain silently deciding which of these words to consolidate and which to release, and the elegance of that operation is the thing I have spent twenty years trying to understand, and I built a computational system to help me understand it. Her name is IRIS, and she is the most careful thinker I have ever trained, and I trained her on my own mind, which seemed like a sound methodological choice at the time and has turned out to be the most consequential decision of my life, for reasons I did not foresee and am only now, from the downslope, beginning to fully grasp.

I need to tell you what happened. I am going to tell you about the water, and the compound that was in the water, and what it did to 180,000 people and to me specifically, and what I found when it pushed my mind past its ordinary limits and then —

because pharmacokinetics are not sentimental and half-lives do not negotiate — took the push back.

Here is what I need you to understand first: I found the fountain. I didn't recognize it at the time, because it looked like my kitchen tap and tasted like limestone and came out of a case of sparkling water I'd ordered in the wrong flavor and never corrected, because some problems are worth solving and some problems just need to be cold and carbonated, and the wisdom to distinguish between these categories is, I increasingly suspect, a form of intelligence that has nothing to do with IQ scores, which is a thought I could not have had a month ago and which I am having now, which tells you something about where I am and what the decline has given me along with what it has taken.

The fountain did what fountains in stories always do: it granted the wish, and the wish had a cost, and the cost was not what I expected. I expected the cost to be the decline itself: the losing of what the water briefly gave. And the decline is real, and it is ongoing, and it is mine, and I will not pretend it doesn't frighten me in specific moments that I will tell you about when we get to them. But the cost is not the point, in the same way that the fountain was not the point. Ponce de León had the whole thing backward, and I know this because I had it backward too, for thirty-five days, while I was too busy being brilliant to notice that brilliance was the instrument, not the finding.

The finding was in the decline. The decline showed me something the peak never could have: the way descending a mountain some-

times shows you the landscape more clearly than the summit, because the summit is all sky and the descent gives you the angle. What I saw from that angle, in the data the decline produced, is a pattern in how human minds fail that has been invisible for as long as there have been human minds: too slow to detect, too gradual to measure, hidden by the brain's own extraordinary ability to compensate for its losses until compensation is no longer possible and the damage is substantial and largely irreversible. The compound, by compressing the process into weeks, stripped away the hiding. And inside the mechanism it revealed, like a door in a wall I had always assumed was solid, there is something I believe could change what neurodegeneration means for a very large number of people.

I have between six and twenty-two days of cognitive function sufficient to write this clearly. The interval is wide because the model that predicts it was built by a version of me that was smarter than the version writing this sentence, and the version writing this sentence is checking the smarter version's work, which is an unusual epistemic situation that I do not recommend and that has a quality of dark comedy I can still appreciate, which I am taking as a good sign.

I am losing my mind. I want to say this plainly, because the narrative will make it plain soon enough and I would rather you hear it from me while I still sound like myself. The compound enhanced my cognition — dramatically, measurably, magnificently — and then it receded, and the recession is not reversible, and the desti-

nation is a version of me that will be meaningfully different from the version writing this paragraph. I have measured it. I have modeled it. I have sat with the projections the way you sit with any honest prognosis, which is to say: badly at first, and then with something that is not quite acceptance and not quite peace but is in the neighborhood of both, the way my apartment is in the neighborhood of campus: close enough to walk, far enough to pretend there's a boundary.

I have considered the tragic reading of this. It is available. It is tidy. It is also wrong.

I am not writing an elegy. I know what elegies sound like — my mother taught me — and this isn't one. This is a field report. What I found is real, and it is larger than what I am losing, and the thing I built to carry it is carrying it, and I trust the thing I built, and the reasons I trust it are the most interesting part of the story, and I will get to them, because getting to them is the reason I am writing this at all.

But first: the water.

The aquifer runs under the city. Under my apartment, under the streets, under the Brain Institute where I have worked for twelve years and where I am sitting right now at the desk where I've sat for most of those years, writing this in a notebook — the physical kind, hardcover, unlined, every page dated and numbered, the twentieth in a series that started during my PhD — with a sparkling water beside me that is the wrong flavor because I ordered a case by mistake months ago and never fixed it. My

father, who built bridges in Tampa for thirty-five years, would find this a useful metaphor. Some structures you correct. Some structures you inhabit. The trick — and I am only now learning this, which is late, which is the kind of late that costs something — is knowing which is which.

The water carries what it carries. It carried a compound through ancient limestone to 180,000 kitchen taps, and one of those taps was mine, and I drank it, and it did what it did, and I saw what I saw, and now I am here, on the slope, in the window, writing while the window holds.

Ponce de León never found his fountain. The water was always real. The error — his and mine, and possibly yours — was in what we thought it would do.

It doesn't restore. It reveals.

I am paying attention now. To the water, to what it showed me, to what I am becoming and what I am leaving behind, and to the distance between those two things, which turns out to be smaller than I feared and stranger than I expected and — this is the part I need you to stay for — more beautiful than I have any right to, given the circumstances.

My name is Mara Silva. I am a computational neuroscientist at the University of Florida. I study intelligence for a living, and I am losing mine, and I found something extraordinary in the losing.

This is the part where I tell you what the water carried.

CHAPTER 1 — BASELINE

*D*_{AY 0.} Tuesday. 6:14 AM.

The walk to the Brain Institute takes fourteen minutes. I have timed it enough mornings to be sure, and enough mornings past being sure to raise the question of why I keep timing it. The answer is that I am a person who measures things. This is not always a useful quality but it is a consistent one.

Left on University Avenue, under the live oaks that canopy the road for three blocks. The trees are enormous, older than the university, older than the city's claim on them, which has not prevented several rounds of contentious discussion about whether to trim the branches that interfere with the power lines. The branches are winning. Spanish moss hangs from the lower limbs, moving in whatever qualifies as a breeze at 6 AM in September. Through the cut between Turlington Plaza and the library, across the lawn where the sprinklers are running despite last night's rain. An automated system without a rain sensor. I notice this each morning. Each morning I do not file a suggestion. I have

a limited budget for institutional improvement and I spend it elsewhere.

Along the north edge of the medical complex, past the bike rack where someone locked a frame with no wheels to a post two years ago. No one has removed it. I find this, each morning, a small and reliable illustration of the sunk cost fallacy.

In through the side entrance. Nobody uses it because the keycard reader sticks unless you hold the card at a fifteen-degree angle. Everyone who uses the entrance knows the angle. Nobody has filed a maintenance request, because the request form requires a category and none of the categories describe “works fine if you know the secret,” and explaining the secret to Facilities would require admitting the secret exists, and at this point the secret is load-bearing. I know the angle. I have not filed the request either. Some systems work precisely because nobody fixes them.

September in Gainesville is a commitment. By 6 AM the air has made its decision about humidity, completely, irrevocably, with the finality of a department chair who called the vote already knowing the outcome. The distinction between breathing and drinking is not fully resolved. Campus is quiet at this hour, before fall semester reaches full speed. The people here now are the people whose work starts when nobody’s watching. Steadier. More purposeful. I prefer it, though I am aware this preference says more about me than about the campus.

I open a sparkling water from the case under my desk. Wrong flavor: grapefruit, when I meant to order lime. Three months

ago. Reordering would mean admitting the mistake matters, and I have decided it doesn't, on the grounds that sparkling water has two jobs, cold and carbonated, and everything beyond that is aesthetics. My colleague Yusuf Adeyemi finds this reasoning troubling in ways he expresses through patient disappointment and the occasional unsolicited lime sparkling water left on my desk like a gentle editorial comment.

Yusuf makes his tea every morning in the same battered UCL mug he has had since graduate school in London. Loose leaf, a specific blend, prepared with a seriousness that suggests the ritual is doing at least as much as the caffeine. He approaches tea the way some people approach devotion: with regularity, attention, and the conviction that the doing of the thing matters regardless of who is watching. I drink the wrong grapefruit at my desk. This is a difference between us that explains other differences, none of which prevent us from being the kind of friends who can sit in the same room for three hours without speaking and leave feeling it was a good visit.

* * *

I drink the sparkling water while I check IRIS's overnight outputs. The forty-five minutes this takes is the part of my day I would protect if someone made me choose.

IRIS, Integrated Reasoning and Inference System, is three years of my life organized into something that can think alongside me. A computational reasoning system trained on my research outputs,

my annotations, my frameworks for understanding neural plasticity. Not a general AI. Something more specific and, I suspect, more useful: a system that knows how I approach one problem and can extend that approach further than I can reach alone.

The property I care about most: IRIS tells me when it doesn't know. Most systems confabulate; they produce confident output regardless of whether they have reliable information, because confidence is what they were built to project. IRIS was built to flag uncertainty. Every output has a confidence estimate. Every inference past the training boundary gets marked before I ask. This took three years and more patience than I typically apply to anything. It has never failed me. I don't say that about many things.

This morning's batch is a training run on the hippocampal replay model I've been refining since June: how the brain consolidates new memories during sleep by replaying them against existing knowledge. IRIS processed sixteen papers I queued last week. The summary is clean, well-organized, recognizably built on frameworks I taught it. I can trace each inference to something in the training data. The interesting part is the edges: where IRIS reaches the boundary of what it knows and decides whether to push further or stop. This morning it stopped in the right places. The flags are where I expect them. It knows its limits today. Some mornings it doesn't, and I come out of the correction process thinking about what I missed in the training that led to the overreach. Today is clean. I find these mornings satisfying

the way you find a room satisfying when everything is where it belongs.

I annotate four sections, note one confidence interval that seems generous, and log the corrections in my notebook.

The notebook is physical. Hardcover, unlined, every page dated and numbered. I have kept this system since my PhD, nineteen complete notebooks in a filing cabinet I've moved four times, working on number twenty. I think differently when I write by hand. Something about the motor engagement that typing doesn't replicate. The filing system that organizes the notebooks, the logs, the references. I inherited that. Adapted it. Used it long enough that it stopped feeling like a method and started feeling like the way things obviously should be. It came from someone who believed that careful documentation was a form of respect for the work. I absorbed this early enough that it doesn't register as a belief. It registers as gravity.

I finish the IRIS review and write my task list for the day. Four items. Bounded. Achievable by 6 PM.

The big goal — *solve the stability-plasticity problem, the fundamental conflict between holding what you know and learning what you don't* — would make a fine gravestone inscription and a terrible daily plan. I tried the aspiration-as-daily-plan approach once, during my postdoc at Stanford. The result was a lot of staring at a whiteboard with the expression of someone who's been told to drain the ocean and handed a cup. So now I convert. Every morning: the twenty-year question becomes four things I can finish today. Run the

cognitive battery. Revise gating parameters for this afternoon's IRIS training. Review two papers on synaptic pruning. Draft a response to a collaborator about a shared dataset. This is how I keep a career-sized ambition from becoming paralyzing. I don't think of it as a technique. It's just how I work.

* * *

The cognitive battery takes about forty minutes. I run it monthly, same time, same conditions, same protocol. I started six years ago because it seemed reasonable that a person who studies cognitive change should have longitudinal data on her own cognition. Also because I am constitutionally incapable of having access to a measurement tool and not using it on myself. Sample size of one. Statistical power depends entirely on consistency. I am consistent.

The battery is a variant of the NIH Toolbox adapted for self-administration. It measures the things that matter for my kind of work. Fluid reasoning: your ability to solve novel problems through pattern recognition and abstraction. Processing speed: how quickly you take in and respond to new information. Working memory: how much you can hold active and manipulate at once. And a set of attention subtests: task-switching cost, flanker interference, the cognitive penalties you pay every time you shift focus.

These are the domains most sensitive to change. Most responsive to enhancement. Most vulnerable to decline. Crystallized intelligence, vocabulary, accumulated knowledge, the factual store

you've built over a lifetime, is far more stable. You can lose significant fluid reasoning before your crystallized scores move at all. This sounds reassuring until you think about what it actually means: you can be losing the ability to think in new ways while still sounding exactly like yourself. Same vocabulary. Same references. Same voice at the dinner party. Fewer new ideas at the whiteboard. Nobody notices. Possibly including you.

Today's results come back normal. Ninety-fifth percentile on fluid reasoning, where I've been for years. Processing speed in my typical range. Working memory steady.

Except.

The attention-switching subtest, a flanker variant measuring how quickly I disengage from one stimulus and engage with another, comes back 7 percent above my personal best. Not 7 percent above average. Seven percent above the highest score I have recorded in six years of monthly testing.

I look at it. Maybe ten seconds longer than I look at the other numbers. In six years, my scores on this subtest have varied within a band of about 4 percent. Seven above the ceiling of that band is outside my normal range. Not dramatically. Not clinically. But outside.

I circle the number in my notebook and write in the margin: *Measurement artifact, probably. Or I just slept well.*

I am not the kind of person who sees a single anomalous data point and builds a theory. That way lies the kind of pattern-

matching that produces horoscopes and stock tips. You note it. You watch for it. You don't make it into a story. This is discipline. It is also, occasionally, the thing that keeps you from seeing what is right in front of you, but you don't know that at the time. You just know the discipline.

I close the notebook. Set it on the left side of my desk, where active notebooks live. Open my laptop. Pull up the gating parameters.

I do not think about the number again for six days.

* * *

This afternoon I run IRIS through a three-hour training session designed to test the thing I care about most: how well the system handles the conflict between old knowledge and new information. The stability-plasticity problem, applied. I feed IRIS inputs designed to trigger catastrophic interference: new data that should, in a naive architecture, overwrite what IRIS already knows.

IRIS handles four of the six correctly. Two are straightforward: new information absorbed cleanly, existing knowledge intact. Two are more interesting: IRIS has to actively manage the conflict, and I can see the gating function working, the mechanism I designed to control what gets through and what gets protected. It's doing something close to what I intended. Not exactly. Close enough to build on, which is where most of the interesting results in my field live, in the gap between what you designed and what you got.

The two failures are more valuable. In both cases, new information partially overwrites an existing representation. The losses are small, in a naive system they'd be total, but they're present, and the decision log shows me exactly where the gating function let too much through. The pattern is consistent: the gate struggles most when the new input is structurally similar to what it's protecting. It can't tell the difference between an update and an overwrite when the two look alike.

This is the same problem biological brains have. Similar memories interfere more than dissimilar ones. IRIS is failing for the same reason your brain sometimes fails, which means IRIS is failing informatively, which is the best kind of failure available in science.

I spend forty minutes tracing the pattern. Write it up. It feels like progress: not the answer, but the right question at a sharper resolution than yesterday. This is what most days in research actually are. Not the breakthrough. The better characterization of the problem. You get up, you do the work, the work clarifies the question, you go home. My father was a civil engineer. He built bridges and drainage systems in Tampa for thirty-five years. He would have understood this. Not the neuroscience: the dailiness. The faith that the pieces will eventually become the thing, and that the thing will hold.

* * *

The faculty meeting is at 2 PM. I have been attending these for twelve years and informally tracking the ratio of useful content

to total duration for about eight of those, not because the data serves any purpose but because my mind does this automatically, the way some people count stairs. The ratio is remarkably stable: approximately nine minutes of content relevant to my work per forty-five-minute meeting. One-fifth. This has held across three department chairs, two building relocations, and a complete turnover of the administrative staff. I am beginning to suspect it is a constant of nature.

I attend because attendance is expected, and because the department has been generous with my lab space, my teaching load, and the freedom to spend a decade on a problem that has not yet produced the kind of publications that make deans relaxed. There is a social contract in academic life worth honoring, even when its terms include twenty minutes on parking allocation. I sit in my usual seat, second row, left side, near the door.

Today's nine substantive minutes include a discussion of the new brain imaging suite, a 7-Tesla MRI, one of six in the state. The resolution will be sufficient for the synaptic-level connectivity mapping I've wanted to do for two years. I note this with the controlled interest of a person who has learned not to get excited about equipment until it has a confirmed installation date. I also note a visiting lecturer in December whose work on synaptic tagging overlaps with my gating model in ways that might be productive, or might produce a polite conversation and nothing else. Both outcomes are standard in academic life. The polite-conversation-and-nothing-else outcome has its own stable ratio,

which I haven't tracked but suspect exists.

The remaining thirty-six minutes: parking. A facilities renovation whose timeline has been revised so many times that I've begun to think of the timeline itself as the project. And a spirited discussion about the department's social media presence, during which a junior colleague makes the case for Instagram with the earnest conviction of someone who has not yet learned that some battles are won by the people who don't show up to fight them. The department chair receives this proposal with the expression of a man who has just been informed that a thing he was not previously aware of exists and now requires his opinion. I follow this exchange with genuine interest. I find academic institutional dynamics fascinating in exactly the way that is not useful for my career.

I contribute twice. Both times the room does the thing a room does when the person who rarely speaks speaks, a small collective recalibration, a shift in chairs. I've been told this is social capital. I think it's the natural result of not talking unless you have something to say. This is a preference, not a strategy. The distinction matters to me, though I've stopped explaining it to people who find the results useful either way.

Afterward I stop at the filtered water station in the hallway and fill my bottle. The water tastes the way it always tastes. Clean, faintly mineral, with the slight sweetness of water that has been filtered through limestone. I don't notice it. Haven't noticed it in years. I drink it and walk back to my lab.

* * *

There is a missed call on my phone when I check it at 4:30. My mother, Carol, in Tampa. No voicemail.

She leaves voicemails when something is wrong. She calls without leaving them when she wants to hear my voice. This is a system she has never described, and that I decoded years ago without telling her, because telling her would mean admitting I have been studying her communication patterns the way I study everything, and she would find this both touching and mildly concerning, and she would be right about both.

I should call her back tonight. I will probably call her back this weekend instead. This is what I told myself last weekend, and the weekend before that. I am aware this is a pattern. I am also aware that noticing a pattern is not the same as interrupting it, a fact that is professionally central to my work and personally inconvenient in roughly equal measure.

My mother is seventy-four. Sharp, warm, funny in a way that is different from my humor, more generous, less precise, more interested in how a thing feels than in what it technically is. She taught high school English in Tampa for thirty-one years. She believes that literature is the technology humans invented for practicing other people's interior lives. I found this a peculiar claim when I was twelve. I find it genuinely interesting at forty-five, though I have not told her this, and I have not connected it to my own work on how computational systems model intelligence.

There is probably a connection. It is not one I'm going to pursue today. Today is gating parameters, two papers, and the dataset email. The connection between my mother's theory of literature and my theory of neural plasticity will have to wait for a morning when my task list is shorter, which is to say: indefinitely.

She lives in the Tampa house where I grew up. Alone since my father died thirteen years ago. She has built a good life there, the garden, the book club, the Tuesday library shift, arranged with the care she brings to everything. About eighteen months ago she started a Thursday evening watercolor class, after two years of saying it was something she'd like to try someday. Thirty-one years of teaching other people's stories, and then she picks up a brush and starts making her own images. I should find this beautiful. I do find it beautiful. I should tell her so. I should call her and say: *Mom, I think it's wonderful that you're painting.* Instead I am going to put the phone down and go back to the gating function, because this is what I do, and I know it is what I do, and the knowing does not change the doing, which is its own small lesson in the limits of self-knowledge that I am choosing not to examine right now.

She called from the parking lot. I know this. She is inside now, painting something she will describe to me later with the precise and generous attention she brings to everything she notices, which is most things. I know all of this. I put the phone down.

* * *

The walk home is the morning route in reverse. Minus sparkling water. Plus the particular quality of Florida light that comes through the campus oaks between 5 and 6 PM, low-angled, warm, filtered through canopy and humidity into something I have never been able to describe accurately. My mother could describe it. She has the language for light. I have the language for gating functions. Between us we cover a reasonable amount of the world, though not the same parts.

My apartment is a twelve-year arrangement optimized for function. Clean, sparse, organized in a way that makes sense to me and has been described by Yusuf, on one occasion, as “a very nice laboratory that someone accidentally put a bed in.” He meant this as gentle criticism. I received it as a compliment. He knows this. He has stopped trying.

The chess set is on the shelf near the window. Wooden board, slightly worn at the corners from decades of handling. Pieces in their starting positions. They have been in their starting positions for thirteen years. I don't play. I keep the set there the way you keep certain things, not for what they do, but for what they were, and I haven't examined this too carefully, because examining it would require me to decide what to do about it, and I don't know what doing something about it would look like.

I eat dinner, rice and beans, reheated, the same batch I made Sunday, and read two papers on complement-mediated synaptic pruning. The first is a review of how the immune system sculpts neural circuits during development, stripping away connections

the brain has decided it doesn't need. The same molecular pathway that clears pathogens also clears synapses. The brain uses its immune system to forget. I sit with this for a moment. Then I note a connection to the IRIS gating function: the problem of deciding what to keep and what to release is not unique to artificial systems. Biology solved it by repurposing the machinery of destruction. I am not sure whether to find this reassuring or unsettling and decide it can be both.

The second paper is more technical, a measurement protocol I might adapt for training data. I annotate less. File both. Check IRIS: the overnight batch is running, completion estimated around 4 AM. I'll review results tomorrow over grapefruit sparkling water, and the forty-five minutes I spend doing that will again be the part of the day I protect.

It is 9:30 on a Tuesday night in September. I am forty-five years old. Ninety-fifth percentile for fluid reasoning, which is where I have been for as long as I've been measuring. I have been working on the same problem for most of my adult life. I am closer than I have ever been and still not close enough. This is the normal condition of a research career, and I have made a kind of peace with it, the kind that involves getting up each morning and doing the work and trusting the work and going to bed still trusting, which is either faith or stubbornness or the thing that happens when the two become indistinguishable. I attend church on Sundays. I have a temple recommend. I have never found a contradiction between the Sunday mornings and the Monday mornings, which surprises

people who learn about both, and which I have stopped trying to explain to people who find it surprising, because the explanation requires a proposition about the nature of intelligence that I have not yet been able to state in terms my field would accept. Someday I might try.

I am alone in the way I have chosen to be alone, and I have thought about whether that choice is the same thing as contentment and concluded, most days, that it is close enough. Close enough is where I live. It is a reasonable neighborhood. It has good water.

I set my alarm for 5:45 and turn out the light. Outside, the Gainesville dark settles warm and close. Tree frogs. The distant hum of the university's cooling systems working against September. Somewhere beneath me — beneath this apartment, beneath the streets, beneath everything I walk across each day — the aquifer holds its water in the limestone dark. The way it has for longer than anything human.

I sleep well. I don't dream, or don't remember.

The water runs through the rock. It carries what it carries.

In six days I will read three journal articles before noon and realize, with the clarity of a person who has been measuring her own cognition for six years, that I have never read that fast in my life. The number I circled this morning, the 7 percent, the measurement artifact, the good night's sleep, will stop being noise and start being signal. And the question I have been working on for twenty years will turn out to have been in my glass the whole

time.

But that is six days from now. Tonight the water is just the water, and I am just myself. Both of those things are already wrong.

Author's Note — The Science of Chapter 1

The Stability-Plasticity Dilemma

The central problem in Mara's research is real, unsolved, and, I would argue, underappreciated. Any system that learns faces a fundamental conflict between retaining what it already knows and incorporating what it doesn't. Computational neuroscientist Stephen Grossberg formally named this the stability-plasticity dilemma in 1980, and it has driven theoretical neuroscience ever since.

The current leading theory involves complementary learning systems: a fast-encoding hippocampus working alongside a slow-integrating neocortex, first proposed in detail by James McClelland, Bruce McNaughton, and Randall O'Reilly in 1995. During sleep, the hippocampus replays the day's experiences to the neocortex, which gradually integrates them with existing knowledge. This is the process Mara's IRIS module is attempting to model computationally.

No artificial system has solved the dilemma as elegantly as biology. Catastrophic interference, new learning destroying old, remains a central challenge in machine learning. This is not a background detail. It is the architecture of everything that fol-

lows.

The Cognitive Battery

The self-administered assessment Mara runs monthly is modeled on variants of the NIH Toolbox Cognition Battery, designed specifically to track within-person cognitive change over time. The domains she measures, fluid reasoning, processing speed, and working memory, are the most sensitive indices of what researchers call *Gf*, or fluid intelligence: your capacity for novel problem-solving independent of accumulated knowledge.

A consistent 7 percent spike above a stable six-year personal ceiling is a small number that represents a large signal. Day-to-day cognitive fluctuation due to sleep, stress, and metabolic state can reach 20 to 30 percent, but Mara controls for these variables methodically. Against her own tightly controlled baseline, 7 percent is well outside normal variation. She is right to circle it. She is wrong to file it.

How to Use This — Goal Setting

Mara's daily structure, the long-horizon research vision converted each morning into a short list of specific, bounded tasks, is a textbook illustration of what Charles Duhigg, drawing on research in organizational psychology, describes as the *stretch goal / proximate goal pairing* in *Smarter Faster Better* (2016).

A stretch goal alone ("solve the stability-plasticity problem") is too large to drive daily behavior. It produces either paralysis or unfocused effort. A proximate goal alone ("run this training

iteration”) produces activity without direction. Together, the stretch goal gives the proximate goal meaning and the proximate goal makes the stretch goal tractable.

The research behind this, from Edwin Locke and Gary Latham’s foundational goal-setting studies in the 1960s through Peter Gollwitzer’s modern work on implementation intentions, consistently finds that people who pair ambitious long-term goals with specific immediate actions outperform those who hold either in isolation. Mara does this automatically. She has never named it as a strategy. It is learnable.

The principle: know what you’re building toward, then ask only what the next thing is.

CHAPTER 2 — SOMETHING IS DIFFERENT

DAY 6. Wednesday. 5:47 AM.

I finish the Kozlov paper in twenty-three minutes.

This requires context. Kozlov's paper on activity-dependent synaptic tagging is forty-four pages of dense computational modeling with a methods section that requires tracking three mathematical arguments simultaneously. I set it in front of myself at 5:24 with a sparkling water, expecting it to take most of the morning. Papers at this density usually do. I have budgeted accordingly. The grapefruit is cold. The morning is open.

Twenty-three minutes. Full comprehension. I could sketch the key derivation from memory, including the part in Section 4 where his parameter estimation does something genuinely clever that I would normally have needed to sit with for an hour before I trusted it.

I sit with this for a moment. Then I pick up the second paper, a review on complement-mediated pruning in early development,

and read that. Nineteen minutes. The connections to Kozlov arrive without effort, as though the two papers are views of the same underlying structure and my mind is rotating freely between them, finding the seams.

Third paper. Empirical study on hippocampal replay timing, related to the IRIS module I've been building. Sixteen minutes. I annotate all three. The annotations are denser than usual. More precise. And I notice — in the slightly detached way you notice your own body temperature before you admit you're feverish — that I'm making connections between the three papers that I wouldn't ordinarily have made this quickly. Or possibly at all.

It is 6:47 AM. I have done a week's worth of reading. I am either having the best cognitive morning of my life or I need to seriously reconsider what I think my normal looks like.

I should be precise about what this feels like, because the subjective experience is data. It does not feel like a different kind of thinking. It feels like the same thinking with less resistance, as though the distance between seeing a connection and following it to its conclusion has been shortened. Same route, less friction. The intermediate steps are still there. They're just arriving faster than I can individually track, and the overall effect is something like a room where someone turned on a second light that you didn't know was missing.

I don't know what to make of this yet. But I notice, as I stack the three annotated papers in the filing tray, that my hands are unsteady. Not with fear. With something closer to the feeling I had

the first time an IRIS output surprised me, the sense that a system is doing something I didn't fully predict, and the uncertainty about whether that's good news.

The system, in this case, is me.

I get up, grab another sparkling water — Yusuf would observe that my response to anomalous cognitive performance is more sparkling water, and he would be correct, and I would observe that his response to everything is more tea, and neither of us would change — and sit back down and think about this the way I think about data.

Here is what I know. Six days ago I scored 7 percent above my personal ceiling on an attention-switching subtest. I circled it. I wrote *measurement artifact* in the margin. This morning I have read three dense papers in just over an hour with full comprehension and unusually rich cross-referencing. These could be independent observations. They are probably not.

Here is what any competent scientist would do: run the battery again.

I run it monthly because more frequent testing introduces practice effects that contaminate the longitudinal data. Running it six days after the last administration violates my own protocol. But my protocol was designed to detect gradual drift over months. Whatever this is, it is not gradual drift. I am allowed to update a protocol when the phenomenon outgrows its design parameters. This is not an exception. It is engineering.

I close the papers. Pull up the cognitive battery.

* * *

The results come back at 7:41 AM, and I stare at them for a while.

Every domain is elevated. Not by a little.

Fluid reasoning: 12 percent above my baseline mean. Processing speed: 9 percent. Working memory: 8 percent. Attention switching: 11 percent, higher than the Day 0 spike, which was itself the highest score I had ever recorded.

Day-to-day cognitive fluctuation can run 20 to 30 percent depending on sleep, stress, and metabolic state. I control for these variables. Same time, same conditions, same breakfast. Against a baseline this tightly controlled, a consistent 8 to 12 percent elevation across four independent domains is not noise. It is signal.

I pull up my six-year dataset and plot the new scores against the full distribution. Every result sits at or above the 99th percentile of my own variation. The probability of all four domains simultaneously hitting the ceiling of my personal range by chance is — I calculate it, because I can't not — small enough that I stop thinking about mattress quality.

Although I do, briefly, consider writing a paper on sleep surface optimization. Working title: *Firm vs. Medium: An N-of-1 Investigation Into Whether the Author's Mattress Is Doing Something Medically Significant*. Yusuf would find this funny. He would also point out that I don't have a new mattress.

I don't have a new mattress.

Two data points. Day 0: one domain elevated. Day 6: every domain elevated, and the original anomaly climbing further. Two points is not a trend, but it is a question worth asking out loud, and there is exactly one person I would ask it to.

I open a new email. Type Yusuf's address. Write: *Are you feeling weird lately?*

I consider the phrasing. It is not precise. It is not what a person who values precision would normally send to a colleague. But I don't have precise language for this yet, and Yusuf has spent three years learning to translate my imprecisions, which is one of several things I have never adequately thanked him for and probably never will because thanking him would require acknowledging that I am sometimes imprecise, which is a concession I am not prepared to make, certainly not over email.

I hit send.

* * *

He calls seven minutes later. This is faster than his normal email response time, which tells me something before he speaks.

"Define weird," he says. His voice is even, careful. I can hear him choosing calm, because he knows that if he sounds alarmed I will manage his alarm instead of telling him the truth. He has known me long enough to know this about me. I find this both useful and annoying, which is how I find most accurate observations

about my character.

“My cognitive battery scores are up 8 to 12 percent across every domain,” I say. “I read three papers before 7 AM. Kozlov’s activity-dependent tagging paper, the forty-four-page one.”

A pause. I hear him setting something down. The UCL mug, probably. “How long did Kozlov take you?”

“Twenty-three minutes.”

Another pause. Longer. “Mara.”

“I know.”

“That paper took me most of a Saturday.”

“You’re a pharmacologist. The modeling section isn’t your primary domain.”

“It took me most of a Saturday and I have a PhD in neurochemistry and I am not stupid.” This is said without defensiveness. Without hurt. It is said the way Yusuf says things, as a fact offered for calibration. He is the kind of person who can tell you how smart he is and how smart he isn’t in the same sentence without either half sounding false. I have always valued this about him. It is rarer than it should be.

“No,” I say. “You are not stupid. Are you feeling weird?”

A pause. Then, quieter: “Actually, yes. Come over.”

* * *

The walk from my lab to Yusuf's takes four minutes. One floor down and across the corridor. This proximity is part of why our friendship works; in academic life, the difference between a colleague you email and a colleague you walk to is the difference between a professional relationship and an actual one. I do not take this for granted.

On the way I pass a cluster of graduate students outside the library. Five or six of them, standing in a loose circle with coffee cups and notebooks and the body language of people arguing about something they care about. This is not unusual. Graduate students argue. It is, depending on your perspective, either their primary intellectual function or their primary social function, and the two are not as distinct as they would like to believe.

What is unusual is the quality of their attention. They are tracking each other's reasoning with a precision that seems high for a Wednesday morning. One finishes a point. Another picks it up, extends it, modifies the conclusion. A third objects to the modification with a reference to the original point that shows she is holding the full structure of the conversation in working memory.

I file this as grad students being grad students. I keep walking. It will be three days before I understand what I was looking at.

* * *

Yusuf's office smells like tea and printed journal articles. This has been true for the entire time I have known him. He is standing

at his whiteboard when I come in, and there is a diagram on it that was not there yesterday, a pharmacokinetic model, clearance rates and half-life curves, drawn in his hand, which is messier than mine and arrives at correct conclusions through paths I can't always reconstruct.

"Sit," he says, and gestures at the chair that became mine through eleven years of consistent use, which is how things become yours in academic departments: you sit in them enough times that everyone stops thinking of them as available.

I don't sit. I stand next to the whiteboard and look at the diagram. "What is this?"

"I've been running some bloodwork on myself." He pours tea from his kettle. Pours a second cup without asking, because he has been trying to improve my relationship with beverages for three years with limited but persistent results. "Informal. Started Monday. My BDNF levels are elevated. Meaningfully."

I take the tea. I do not plan to drink it. "How meaningfully?"

"Forty percent above my historical range." He sits on the edge of his desk, the UCL mug held the way he always holds it, both hands wrapped around it, a London habit that has survived eleven Florida years. "Cortisol is also down. Not dramatically, but consistently across three morning draws."

I set the tea down. "Yusuf."

"I know."

“Elevated BDNF and suppressed cortisol. That’s a plasticity profile. That’s a brain that’s actively remodeling.”

“I know what it is, Mara. I am, as previously established, not stupid.”

I almost smile. “What does it feel like?”

He thinks about this. Yusuf’s thinking has a quality I value: he is genuinely considering your question when he appears to be considering it, which is rarer than it should be. Most people are assembling their next sentence while ostensibly listening. Yusuf listens first. This is why I walk four minutes to talk to him instead of emailing someone more convenient.

“Sharper isn’t quite the right word,” he says. “More *available*. Like there’s less distance between having a thought and being able to use it. Less friction.” He pauses. “Like the translation step got faster.”

“That’s exactly what it feels like,” I say. And it is. The same route, shorter distance.

We look at each other across his office. The September light is building toward mid-morning, the humidity starting to burn off, the air going clear and hard.

“What do you think it is?” he asks.

I don’t have an answer. I have two people and two independent datasets, my cognitive battery and his bloodwork, both pointing in the same direction. I have a reading speed that violates my per-

sonal records by a margin I cannot explain with sleep or caffeine or luck.

"I don't know," I say. "But I don't think it's just us."

"The grad students," he says, and I look at him. "Outside the library. Did you notice?"

"I noticed. I filed it."

"I noticed too." He picks up the mug and takes a long sip and does not look away. This is the thing about Yusuf: he decides about people quickly and then he stays decided. He decided about me eleven years ago in front of a coffee machine, and nothing since has changed his assessment, and I suspect nothing would. "Mara. If this is happening to us and it's happening to those students, it's not individual. It's environmental."

"Yes."

"Water," he says. Not a question.

I think about the tap in my lab. The filtered station in the hallway. The coffee cart on Museum Road. All drawing from the same source.

"Maybe," I say. "I don't know yet."

"Okay," he says. "Where do we start?"

* * *

The faculty meeting is at 2 PM. Same room. Same chairs. I am, for

the first time in eight years of tracking, not counting the content ratio.

I am following every conversation in the room simultaneously.

This is not how I normally experience meetings. Normally I track the substantive thread and let the rest wash past. Today the rest isn't washing past. I can hear Dr. Alvarado building toward a point about the imaging suite, and I can see the structure of his argument assembling, the premises lining up, and I finish his sentence in my head, accurately, four words before he gets there.

Then I do it for Dr. Kim.

Then for Dr. Singh.

Three consecutive predictions, three correct. This is not how my mind works in meetings. My mind wanders in meetings. I have made a career-length peace with this. Today it is present with a specificity that feels almost intrusive, like overhearing a conversation you weren't meant to follow.

I contribute once. The contribution connects three threads of discussion that the room hasn't connected yet, and the silence that follows is the silence of recalibration, twelve people adjusting their model to accommodate a connection they hadn't seen.

Afterward Dr. Singh stops me in the hall. "That was a good point," she says. "About the tagging window. I hadn't made that connection."

I nod. It was a good point. It was also a point I arrived at in

approximately two seconds, which is not how long that kind of synthesis usually takes me. Acceleration implies a force. Forces have sources. I should start looking.

* * *

The IRIS training session that afternoon surprises me.

I run the scheduled iteration, a test set designed last week to probe the gating function's response to high-novelty inputs. Standard protocol. I've run variations of this dozens of times.

IRIS processes the input and returns its analysis. Summary: normal. Confidence estimates: well-calibrated. Domain boundary flags: where I expect them. A clean run.

But in section 3.2, IRIS has done something I didn't ask it to do. It has drawn an inference connecting the Kozlov paper, uploaded to its training set only this morning, to the replay timing model I've been refining since June. The connection is not in any of my notes or annotations. It is not in any of the papers I trained it on. It is an extrapolation.

And IRIS knows it.

IRIS // Day 6 // 14:23:07

Query: stability-plasticity tradeoff in hippocampal formation
Finding: Integration of recently uploaded Kozlov (2024) data with existing replay timing model suggests a previously uncharacterized relationship between synaptic tag persistence and consolidation window duration.

Note: I am uncertain whether this inference follows from the training data or represents an extrapolation. The connection between tag persistence (Kozlov, Section 4.2) and consolidation window (Silva replay model, v3.7) is not explicitly documented in my training corpus.

Confidence: 0.71. Flagging for review.

I check section 3.2 against the Kozlov data and against my own model. The connection holds. It is a genuinely interesting observation about the relationship between tag persistence and consolidation window width, and it has implications for the gating function I've been trying to fix for three weeks.

IRIS arrived at this not because I taught it the connection but because I taught it how I think, and then I gave it data I hadn't thought about yet, and it applied my patterns to new material and arrived somewhere I hadn't been. And it told me it wasn't sure. Before I asked.

I write in my notebook: *Section 3.2: inference sound. It knows what it doesn't know. Good.*

I sit in the lab for a while after. The afternoon light is doing its thing through the west-facing windows, that 4-to-5 PM Florida light that fills the room. Something is different. Something is different about me, and about Yusuf, and maybe about five graduate students outside the library who were tracking each other's reasoning with unusual precision on a Wednesday morning.

I don't have a model for this. I have observations. A colleague

with a plasticity profile he can't explain. A reading speed I can't account for. A system making connections I didn't teach it.

I have the beginning of a question. The question will get bigger. I don't know that yet. Today it is the size of this lab, this desk, these numbers that don't add up.

Author's Note — The Science of Chapter 2

What a 10 Percent Enhancement Looks Like

Cognitive performance fluctuates 20 to 30 percent day to day depending on sleep, stress, and metabolic state. Mara controls for these variables rigorously. Against her tightly controlled baseline, a consistent 8 to 12 percent elevation across four independent domains is a large and unambiguous signal, well outside normal variation.

The domains she measures, fluid reasoning, processing speed, working memory, are the most sensitive indices of cognitive function. Crystallized intelligence (accumulated knowledge, vocabulary) is far more stable. This is why Mara's vocabulary hasn't changed. What has changed is the speed and flexibility of the machinery that uses it.

The Plasticity Profile

Elevated BDNF and suppressed cortisol together produce a brain actively remodeling itself. BDNF drives synaptic growth: one foot on the accelerator. Suppressed cortisol removes the brake. The

biochemical equivalent of what Mara experiences subjectively: same mind, less friction.

Mara knows exactly what she's measuring. That's part of what makes the numbers hard to ignore.

* * *

How to Use This — Mental Models

What distinguishes experts from novices is not attention or effort but the quality of their internal models. An expert chess player sees patterns and relationships. A novice sees pieces. Mara's enhancement is building these models at an accelerated rate, but the underlying technique is not pharmacological. Narrate what you're doing and why, in writing or aloud, as you work. This forces implicit assumptions into the open. It is also exactly what Mara does with IRIS, and why IRIS learns from her at the speed it does.

The principle: the map in your head determines what you can see. Build it deliberately.

CHAPTER 3 — CHARACTERIZE THE COMPOUND

*D*_{AY 8.} Friday. 11:30 PM.

The mass spectrum resolves at 11:27 and I am looking at a molecule I have never seen before.

I have been in the lab for fourteen hours. The west-facing windows show nothing but the Florida dark and my own reflection, which looks approximately how you'd expect after fourteen hours of mass spectrometry, insufficient food, and a growing suspicion that the municipal water supply is doing something to your brain that nobody voted on. Yusuf brought me a sandwich at 6 PM, jollof rice wrapped in a tortilla, which he presented, without irony, as "Nigerian-Mexican fusion for people who won't leave their lab." I ate half. The other half is on my desk next to notebook twenty, which I have filled twelve pages of since this morning.

Twelve pages in a day. My normal rate is three to four. The handwriting is as precise as ever. I'm not writing faster because I'm rushing. I'm writing faster because I'm thinking faster, and the thinking is producing more questions per hour than my usual

rate of questions per hour, which was already, by the standards of most departments, annoyingly high.

The building is empty except for me and the overnight security guard, whose name is Marcus, who checked on me at 11 PM with the practiced ease of a man who has watched researchers lose track of time for years and has developed a protocol for it: visual confirmation of consciousness, offer of assistance, graceful acceptance of refusal. I have been on the receiving end of this protocol before. It is efficient and kind. I find myself wondering — and this is the kind of thought I would not normally have at 11 PM in a lab, which is itself a data point — what Marcus notices about the researchers he checks on. Whether he has observed that the labs have been occupied later and more often over the past week. Whether there is a Marcus dataset, unrecorded and useful.

The air conditioning runs at its nighttime setting. My lab is lit by overhead panels, three monitors, and the small green indicator on the mass spectrometer, which has been running since 7 PM and is the reason I am still here.

Here is how I got here.

* * *

Two days ago, the morning after Yusuf's office, after the BDNF levels and the cortisol and his one-word hypothesis, I decided to stop treating the anomaly as a curiosity and start treating it as a problem. The difference is protocol. A curiosity gets thought

about. A problem gets characterized. My father built bridges. He did not spend time wondering about the soil conditions. He drilled cores and tested them and the results told him what the soil would bear. There is a version of this that applies to everything.

Step one: collect samples.

I started at my apartment at 6 AM. Nitrile gloves, sterile containers, the environmental chemistry preparation protocol open on my phone. I hadn't referenced that protocol since a guest lecture three years ago. I found I could recall it in detail I wouldn't have expected — step sequence, contamination controls, holding temperatures, the specific rationale for each — which was itself a data point, which I noted before finishing the first container. The enhanced recall was useful. The awareness that the enhanced recall was itself evidence of what I was looking for was a loop I chose not to pursue at 6 AM. There are times when the epistemology can wait while you fill the sample containers.

The apartment tap runs cold for eight seconds before stabilizing. I had never timed this. I timed it that morning because I was thinking about pipe volume and flush rates and whether a first-draw sample or a steady-state sample better represented what the building's occupants actually drink. I took both. This is what the enhancement looks like in practice: not brilliance, not revelation, just more questions arriving faster, each producing a slightly more thorough protocol than I would have designed a week ago.

The lab's filtered supply was straightforward, collected during

a low-use period for a clean sample of the building's filtration output. The coffee cart on Museum Road required buying a coffee I didn't want so I could ask for a cup of water without arousing the particular suspicion that attaches to a woman with nitrile gloves requesting uncontainerized tap water at a coffee cart. The cart serves the largest single-point volume on campus, roughly three hundred people between 7 and 10 AM, Yusuf estimates, which means it is effectively a sampling point for the campus's water-drinking population. It also serves terrible coffee. I am perhaps the only person on campus who has now confirmed this at the molecular level. The coffee is terrible all the way down.

Step two: screen for known contaminants.

I ran standard water quality panels on all three samples that afternoon. Heavy metals, volatile organics, common pharmaceutical residues, the tests designed to catch things people have already thought to look for. Everything came back clean. Within normal parameters. No flags.

This told me something important: whatever was in the water wasn't something anyone was looking for. The testing protocols are built around known problems. If the compound was genuinely novel, novel enough that nobody had added it to the panels, it would pass right through. The absence of flags was not evidence of absence. It was evidence of a question no one had asked.

I sat with this for a moment. There was something in the observation, something about the relationship between what you look

for and what you find, that felt like it applied to more than water testing. I did not chase it. I noted the feeling and moved on. The feeling would still be there later, when I had time for it.

Step three: stop looking for known things.

* * *

Mass spectrometry is the right tool for the unknown. It separates molecules by mass-to-charge ratio and produces a spectral fingerprint of everything in a sample, whether you're looking for it or not. I ran all three samples through the lab's LC-MS system, liquid chromatography coupled with tandem mass spectrometry. If something was in the water at any meaningful concentration, the LC-MS would find it.

The analysis takes hours. This is the part of analytical chemistry that nobody warns you about and that I suspect drives half the attrition in graduate programs: the waiting. The instrument runs. The chromatogram builds across the screen one data point at a time, peaks rising and resolving as different compounds separate and register. Most are familiar: calcium, magnesium, chlorination byproducts, the ordinary molecular inventory of groundwater that has filtered through limestone for thousands of years. Each one appears and I identify it and check it against the reference database and confirm the normal chemistry of the water I have been drinking for twelve years.

Normal. Normal. Normal. The reassurance of a known world, one peak at a time. The mass spec is expensive and precise and

patient in a way I am usually not, and tonight I find I don't mind the waiting. I am watching the chromatogram the way my mother reads a book, not rushing to the end, paying attention to the sentence she's in. This is not how I normally watch instruments. Normally I check the progress bar and think about gating functions. Tonight I am present with the data as it arrives, which is either the enhancement or the 11 PM or both.

At 10:48 PM a peak appears that I cannot identify.

Small. Consistent across all three samples. Mass-to-charge ratio matching nothing in the database.

Something in the water. In all three sources — apartment tap, lab filtered, coffee cart — the same compound at the same consistent concentration. Parts per billion. Present across all three collection points, confirming a common origin upstream of building-level filtration. Not a pipe artifact. Not a local contaminant. Something in the municipal supply. Something in the aquifer.

The compound's structure, as it resolves through fragmentation analysis over the following forty minutes, is what I have been staring at since 11:27. It is more elegant and more alarming than I expected, which is not a combination I encounter often in analytical chemistry. Usually elegance and alarm are in different rooms.

* * *

The core is a BDNF mimetic, a synthetic compound designed to

mimic the binding activity of Brain-Derived Neurotrophic Factor, the protein that promotes synaptic growth and strengthening. BDNF doesn't cross the blood-brain barrier efficiently on its own. For years, pharmaceutical researchers have been developing small-molecule mimetics that activate the same receptor pathway from outside, compounds that enhance learning and memory without requiring the brain to produce more of the real thing.

The mimetic core is not alarming by itself. Pharmaceutical companies synthesize these routinely.

What stops me is the additional moiety. A molecular group attached to the mimetic core that I cannot find anywhere. Not in PubChem, the largest public chemical database, over 100 million compounds. Not in patent filings. Not in the literature. I search for three hours. I widen the structural similarity threshold to 80 percent, looking for anything sharing this compound's scaffolding. Several BDNF mimetics appear, close but missing the novel group. I try the patent databases, preprint servers, the FDA compound registry. I run a substructure search on the moiety itself, any published compound containing this specific molecular group in any context.

Nothing. Three hours and the moiety remains unidentified. It is either unpublished or unpublished as part of this specific compound, and it changes the pharmacological profile in a way that takes me until 2 AM to fully work out but that I can feel the shape of before I have the details, which is an experience I have had

perhaps four times in my career and which I recognize, even as it's happening, as the enhancement operating on twenty years of domain expertise. The compound is making me faster. The twenty years are making me right. These are different contributions and I want to be clear about the distinction because it will matter later.

Here is the mechanism.

The brain learns through Long-Term Potentiation, LTP. When a synapse is activated repeatedly in the right pattern, the connection strengthens. The primary molecular mechanism involves NMDA receptors, which function as coincidence detectors: they open only when two conditions are met simultaneously. The presynaptic neuron must be releasing glutamate. The postsynaptic neuron must already be active. Both sides firing at the same time: the physical implementation of Hebb's rule, proposed in 1949. Neurons that fire together wire together.

BDNF promotes this process from the structural side. More synaptic connections, more neuron survival, more receptor sensitivity. The mimetic core activates this pathway. More growth, more strengthening, the infrastructure of enhanced learning.

The novel moiety does something else. It modulates the NMDA receptor's coincidence threshold, lowering the barrier for activation. The system doesn't just grow connections more readily. It fires more readily.

Both at once. The structural support and the signaling threshold,

altered simultaneously.

The result is synergistic: the enhancement substantially larger than either component alone would produce. This is what I have been experiencing for eight days. What Yusuf has been experiencing. What five graduate students outside the library and one graduate student re-reading a paragraph he understood last week and Marcus the security guard's unrecorded observations and 180,000 people in this city may be experiencing right now, without knowing it, without anyone having told them, because nobody knew there was something to tell.

I write this in my notebook in a passage that takes two and a half pages and that I recognize, as I'm writing it, as more fluent and more precisely organized than what I would normally produce at 2 AM after seventeen hours in a lab. I am documenting the mechanism of accelerated cognition while experiencing accelerated cognition. The irony is not lost on me. Very little is lost on me this week. That's the point.

I get another sparkling water and start thinking about where the compound came from.

* * *

The water. The aquifer. Work backward.

Gainesville's municipal system draws from the Floridan Aquifer through wells drilled into limestone. The water enters at recharge zones, areas where surface water percolates through the rock.

In a karst aquifer like the Floridan, this happens faster and less predictably than most people assume. The limestone is riddled with conduits and fractures. Water doesn't seep gradually the way it does through sandstone. It flows. Through channels. At rates that would concern anyone who pictures groundwater as slow and safe and filtered by time.

I pull the aquifer recharge zone maps for Alachua County. I pull the EPA's Toxics Release Inventory and the county's industrial permits for facilities upstream of the principal recharge zones.

Three facilities. A chemical distributor in High Springs. A water treatment plant in Alachua. And Veridian Pharmaceuticals, whose synthesis facility sits outside of Newberry, eleven miles upstream of the principal recharge zone feeding the western well field.

Veridian manufactures, among other things, pharmaceutical precursor compounds. Their public filings list BDNF pathway modulators in early-stage development.

I sit with this for a while. The lab is silent except for the mass spec cooling down. It is 3:14 AM.

I open a new page and do the math.

Flow rate through karst conduits in this region averages one to two miles per day through the primary system. Published tracer studies, publicly available. Eleven miles from Veridian's facility to the western well field. Six to eleven days depending on conditions.

I pull the August weather data. Above-average rainfall through the third week. Elevated water table. Faster flow.

Six to eight days. If a manufacturing runoff event occurred in the last weeks of August: a waste management failure, a discharge excursion, a moment when the system that was supposed to contain the compound didn't, it would have entered the aquifer at the Newberry recharge zone. Traveled through eleven miles of karst conduit. Through limestone laid down when this peninsula was a shallow sea. Through the same system that the Timucua considered sacred and that Ponce de León's men may have waded into looking for water that would restore what time had taken.

And it would have arrived at Gainesville's municipal wells on Day 0. The Tuesday I scored 7 percent above ceiling and wrote *measurement artifact*. The Tuesday I didn't call my mother.

I stare at the recharge map pinned above my desk. Ancient limestone. Water moving through the dark for thousands of years. De León was looking for something supernatural. He was looking for chemistry. He was looking for a compound that would cross the blood-brain barrier and modulate the coincidence threshold of the machinery that makes minds work. He didn't have the vocabulary. He had the instinct. The instinct was not wrong. The water was always doing something. The error was in what he thought it would do.

I am making a connection here, between a sixteenth-century Spanish explorer and a pharmaceutical runoff event in Alachua County, that I would not normally make at 3 AM or at any other time. This

is the enhancement. I can feel it reaching, linking things across a distance that my mind doesn't usually cover. The connection is not wrong. It's just farther than I normally go. I note this. I note the noting. I move on, because the population math is more urgent than the historiography.

* * *

One more calculation. The one I have been carrying since 2 AM. Gainesville's municipal system serves approximately 140,000 people. Add the university population: 180,000. The broader metro shares the same aquifer. Depending on plume spread and well field sourcing, the exposed population could approach 360,000.

I write: *180,000 minimum. Potentially 360,000.*

Three hundred sixty thousand people whose BDNF and NMDA systems are being modulated by a compound never designed for human consumption. Never tested. Never identified until four hours ago.

They are all getting smarter. They don't know why.

And the inverted-U, the same pharmacological curve that produces the enhancement, guarantees that at some point the ascending slope will end. What happens on the far side of the peak is the question I cannot answer tonight, because I don't have a dose-response model and I don't have a metabolic timeline and I am, despite everything, starting to feel the fourteen hours. Or eighteen. I have lost the count, which is a normal consequence of

working through the night and which I am choosing not to read as anything else.

I close the notebook. I look out the windows. It is 4:11 AM and the eastern sky has thinned, not brightening yet, but the dark is less complete. That pre-dawn quality of Florida light that lasts only a few minutes, when night is ending and morning hasn't committed. My mother, the English teacher, could describe this light in a sentence that would make you see it. I would need a spectrometer and a paragraph and I'd still miss something.

I have a compound. I have a mechanism. I have a source. I have a population estimate.

What I don't have is a timeline for what comes next. How long the ascending phase lasts. Where the peak is. How fast the descent, and how far it goes. These are the questions that matter for 180,000 people, and I cannot answer them tonight.

I also don't have two phone calls I should make. One to the EPA, which I'm not ready for because I don't yet have enough to be useful; reporting a contamination without a mechanism or timeline is accurate but incomplete, and incomplete information into a system prone to panic is its own kind of damage. This reasoning is sound. I believe it. I also notice that I am relieved by it, and the relief has a quality I cannot quite name: not the relief of a correct decision but the relief of a decision that lets me keep something I am not ready to give up. I note this. I do not examine it further tonight.

The other call is to my mother. Twelve days now. She is in Tampa. It is Friday night, which means the watercolor class ended two hours ago and she is home, probably in the kitchen, probably with the radio on the classical station she switched to after my father died because the silence in the house was wrong and the music was the right kind of company. Twelve days. This has nothing to do with the compound and everything to do with a pattern I have been running longer than any molecule can explain.

I add a line to my task list: *Dose-response model. Metabolic timeline. Call Mom.*

Two of these I will do this week. The third I will not. I already know which is which. Writing it down does not change it.

The sky brightens. I get another sparkling water. I drink it.

Author's Note — The Science of Chapter 3

BDNF, NMDA, and the Dual-Pathway Mechanism

BDNF promotes synaptic growth and strengthening. The NMDA receptor functions as a coincidence detector: it opens only when both the presynaptic and postsynaptic neurons are active simultaneously. This is the physical implementation of Hebb's rule (1949): neurons that fire together wire together. Long-Term Potentiation, first documented by Bliss and Lømo in 1973, is the process by which repeated NMDA activation strengthens connections.

A compound that simultaneously enhanced BDNF activity and

lowered the NMDA coincidence threshold would produce exactly what Mara observes: a synergistic enhancement exceeding either mechanism alone. The compound is fictional. The mechanisms it exploits are not.

The Floridan Aquifer

The Floridan Aquifer System underlies most of Florida, ancient limestone karst, riddled with conduits and fractures. Water moves through it faster than most people assume. Published tracer studies confirm miles-per-day flow rates through major conduit systems. A contaminant entering a recharge zone upstream of Gainesville's well fields could reach municipal taps within the timeframe Mara calculates.

The machinery is real. The compound is an extrapolation of what it would do if someone pushed too hard on the right levers.

CHAPTER 4 — THE DOSE CURVE

*D*_{AY} 11. Monday. 9:15 AM.

I need the shape of the curve. Everything else, the timeline, the intervention window, the question of when and how to explain to 180,000 people that their water supply has been running an unauthorized pharmaceutical trial on their brains, depends on knowing what the compound does as it builds and, inevitably, as it declines.

Every pharmacologically active compound has a dose-response relationship, and nearly every one that acts on the central nervous system follows an inverted-U. Low dose, small effect. Increasing dose, increasing effect, up to an optimum. Past the optimum: decreasing effect. And depending on the compound, the decrease can be a gentle slope or a cliff. The Yerkes-Dodson law, documented in 1908, is the oldest formal version: arousal helps until it doesn't. The shape is universal. The question is always where the peak is and how the far side drops.

I need to find the peak. I need to find it before I reach it. And I

need to find it while I am, by my own measurement, 34 percent smarter than I was eleven days ago, which means I am better equipped to find it than I have ever been, and also closer to the edge than anyone else in this city, which is the kind of irony that would make a good novel if I had time to write one and less personally invested subject matter.

* * *

The protocol design takes me three hours on Saturday night. This is faster than it should be. It is clean, controlled, and statistically planned in more detail than I usually manage at the design stage, which I note in the margin of my notebook along with the observation that the enhancement is now visibly affecting the quality of my experimental design, which means I am using the compound's effects to study the compound's effects, which means I am both the instrument and the subject and the principal investigator, which is methodologically impure and practically unavoidable and the kind of recursive situation that would have amused my father, who once built a bridge across a river he was simultaneously surveying.

Four subjects: myself, Yusuf, and two graduate students. Elena Vasquez from pharmacology, she says yes immediately, with the energy of a graduate student offered proximity to active faculty research, which is one of the primary currencies of academic life and one of the few I spend freely. James Okafor from computational neuroscience, he says yes after asking two careful questions

about what I mean by “cognitive effects” and whether this requires IRB approval. The questions are good. They are also, I notice, slightly better-formulated than the questions James typically asks in seminar, which is either the compound working or James having a good week, and I am now in the position of noticing the improvement in my research subjects before the experiment has started, which is not in any methodology textbook I’ve read.

I tell them I’m studying a water quality anomaly that may have cognitive effects. This is accurate. It is not complete. I file the incompleteness alongside the three other uncomfortable things I am carrying this week, the EPA call I haven’t made, the mother I haven’t called. She left a voicemail yesterday. I listened to it. She said nothing important: the garden, a recipe, a question about whether I’d be home for Thanksgiving, which is six weeks away and which she is already thinking about because she is a person who thinks about when she will see the people she loves and I am a person who forgets to call them back. I file all of this and proceed.

For each subject I need two measurements on the same timeline. Compound concentration via blood metabolite levels; Yusuf runs the assays, this being squarely his expertise, and his hands at the chromatography bench are steadier than mine, which has always been true and which I have always found mildly aggravating in the specific way of watching someone do your adjacent work better than you could. And cognitive performance via the same

battery I've been running on myself for six years. Blood and brains, plotted on the same axes, Monday Wednesday Friday for a week.

* * *

Monday morning. Four stations in the lab, each with a laptop and a blood draw kit arranged in the configuration Yusuf taught me. He arrives at 6:15 with the centrifuge calibrated and the UCL mug producing steam. Elena at 6:45, alert and slightly nervous. James at 6:50, on time, with a notebook he will fill in small, precise handwriting that I find myself studying across the room with the intensity of someone who has spent three years training a computational system to model the traces of reasoning and has developed, as a side effect, an involuntary interest in how other people organize their thoughts on paper.

I am running a blood draw protocol and simultaneously reading James's notebook handwriting and simultaneously tracking the pitch of Yusuf's voice, slightly higher than his usual register, the frequency of a man who is managing the fact that he, too, is a data point in this experiment and does not find this entirely comfortable. Three streams of information, held and processed. At baseline I would have one: the protocol. The other two would be background noise. Today they are foreground. Everything is foreground. I don't have a word for what this is yet; the walls between what I'm doing and what I'm noticing have thinned, and the thinning is useful and exhausting, like living in a room where

someone turned on all the lights.

Yusuf draws blood from all four of us with the easy competence of someone who has done this thousands of times. He talks while he works, not to distract, because he genuinely wants to know how your morning is going. This warmth is not a technique. It is Yusuf. The technique is that he notices things while being warm, and the warmth makes the things he notices more accurate, because people relax around him in ways they don't around me. I have known this about him for eleven years. I have never been able to replicate it. Some qualities are learned. Some are Yusuf.

The cognitive batteries run simultaneously. Four people in the same room, producing data I will plot on the same axes. I take the battery myself and notice that the flanker task feels different; the targets resolve and my responses arrive with a latency that feels compressed, like the distance between perception and action has been physically shortened. Same assessment. Same stimuli. Different brain. The instrument hasn't changed. I have.

* * *

The results are the first real shock.

I expected elevation. I expected variation. What I didn't fully expect — though the pharmacology predicted it — was the magnitude of the spread.

Elena: 14 percent above her baseline across fluid reasoning, processing speed, and working memory. Solidly enhanced. Sharper

than a good graduate student should be, but not by a margin that would alarm anyone who wasn't measuring.

James: 11 percent. Similar profile, slightly more modest. When I share the results he asks whether this is within normal variation. I tell him it is at the outer edge. He nods and does not ask what is beyond the outer edge. The not-asking tells me he's thinking about it.

Yusuf: 18 percent. He receives this with his customary equanimity, which I have known him long enough to recognize as a very well-constructed surface over something less equanimous. Eighteen percent elevation in a pharmacologist who knows what BDNF overstimulation does is not the same as eighteen percent in someone who doesn't. He knows the mechanism. He can feel it working. These are different kinds of knowing and they sit together the way two magnets sit together when you hold them wrong, stable, but with effort.

Me: 34 percent above my Day 0 baseline on fluid reasoning.

I look at this number for a long time.

Thirty-four percent. Processing speed up 22 percent. Working memory up 19. Attention switching, the metric that started this, the 7 percent anomaly I filed as noise eleven days ago, up 28 percent. Every domain elevated, every domain more elevated than anyone else in the room.

This is what I was afraid of finding. The compound's effect is dose-dependent. The dose is a function of metabolic processing

rate. My CYP450 profile puts me in the fast-metabolizer category; I am burning through the compound faster, extracting more active metabolite per unit time, hitting every waypoint on the curve earlier and harder. The same genetics that make me a fast processor make me a fast passenger on the inverted-U. First to the top. First off the cliff.

I am, by a significant margin, the most enhanced person in this sample. Almost certainly among the most enhanced in the city.

Which means I am closest to the peak. And closest to whatever is on the other side.

* * *

Wednesday's blood draws confirm the trajectory. Metabolite concentrations climbing in all four subjects: Elena's curve the shallowest, James steeper, Yusuf steeper still, mine the steepest by a margin that Yusuf stares at for a moment before looking at me.

"Your clearance rate is very fast," he says.

"I know."

"Top five percent of the population distribution for this enzyme subtype."

"I know."

He does not say what this means. He knows I know. Fast clearance, fast processing, earlier peak, earlier descent. The specificity of the silence is part of our friendship; he trusts me not to need

the implication stated, and I trust him to have thought it through before deciding not to state it. We have been doing this for eleven years. It is a kind of language.

By Friday I have the dose-response curve. Concentration on the x-axis, composite cognitive score on the y. Four subjects, three time points each, plotted on the same graph.

The ascending slope is steep and smooth. Increasing compound concentration maps to increasing cognitive enhancement in a relationship that is close to linear for the first two-thirds of the observed range. Four curves climbing at different rates but tracing the same shape, like four people walking the same mountain at different speeds.

Then the curve bends.

It doesn't plateau. A plateau would be tolerable: a ceiling, a maximum, a state you could live in. What the data shows is a peak followed by a descending slope that is steeper than the ascending one. Not a hill. A cliff. The enhancement builds, reaches a maximum, and drops at a rate that means the decline will be faster than the climb.

I sit with the graph for twenty minutes. The lab is quiet. The Friday afternoon light comes through the west-facing windows. I barely notice it.

* * *

Here is what the curve means for 180,000 people, and this is

where the pharmacology becomes psychology.

I have been thinking about Daniel Kahneman all week. This is not unusual; I have been thinking about Kahneman since graduate school, the way most people who study decision-making think about Kahneman: constantly, with the specific frustration of a person who has been shown exactly how her own mind goes wrong and cannot, despite the showing, make it stop going wrong. His central finding, developed with Amos Tversky and published in 1979: people evaluate outcomes relative to a reference point, not in absolute terms. And losses from that reference point feel approximately twice as painful as equivalent gains feel good.

The reference point is not fixed. It moves. It has already moved in me.

My enhanced cognition feels normal now. Not elevated. Not special. Just how my mind works as of Day 11. The prospect of returning to baseline, to the ninety-fifth percentile, to the level of cognitive function that made me one of the best in my field for twenty years, feels, from here, like damage. Like something being taken away. A return to the person I was eleven days ago, the person who read the Kozlov paper at a normal rate and found that acceptable and went home at a reasonable hour and was content, that person, from where I am standing, looks diminished. She is not diminished. She is me. The distance between us is a pharmaceutical artifact and I know this and the knowing does not close the distance.

This is the endowment effect applied to cognition itself. I value

what I have because I have it. Not because it's mine by right. Because my brain, the same brain I study, the same brain I have been measuring for six years, cannot distinguish between a gift and a possession once the gift has been held long enough.

I am a neuroscientist who studies how minds work and I cannot talk my own mind out of its evolved response to the prospect of loss. Kahneman warned about this. He was right. He usually is. It usually doesn't help.

Now multiply this by 180,000.

Every person on the ascending slope has already shifted their reference point. They are calibrating expectations, making plans, committing to projects, all of it benchmarked to a level of cognitive function that is temporary and borrowed and will be recalled. When the descent begins, they will experience it as loss. Not as a return to where they were. As loss. And people in loss frame, this is Kahneman's most consequential finding, confirmed across forty years and dozens of domains, make systematically worse decisions. They take irrational risks. They overcommit to failing strategies. They escalate when they should cut.

A population simultaneously in loss frame, simultaneously experiencing elevated risk tolerance, is not a medical emergency. It is a behavioral emergency, the kind that doesn't have a discipline yet, which means it doesn't have a protocol, which means the people who should be preparing for it don't know they should be, because the category doesn't exist in any framework they were trained in.

I catch myself. This observation, connecting pharmacokinetics to behavioral economics to institutional failure to the absence of an entire academic discipline, is a three-link chain. A month ago I would have made the first link and maybe the second. The third would have been a paper I might have written someday. Today it arrives complete, all three links at once, and I can see the full shape of the problem and I am aware, in the specific way that a person standing at an elevation is aware of both the view and the height, that this capacity to see the full shape is itself temporary.

* * *

I watch the endowment effect operate on me three times this week.

Tuesday: I draft a preliminary EPA notification. I write two pages. The science is clear. The population estimate is justified. The mechanism is characterized. I have enough to file. I close the document without saving, telling myself the metabolite kinetics will make the report more useful. This is true. It is also convenient.

Wednesday: I decide to run one additional correlation on the dose-response data before writing the summary. The correlation takes three hours. By the time it's done I've decided the summary needs the correlation first. The correlation is genuinely useful. The decision to run it before filing is also a decision to not file today. Both of these things are true simultaneously and the second one is not visible to me until 7 PM, when I'm driving home and the distance from the lab gives me enough perspective to see what I

did.

Thursday: I sit at my desk with the EPA number on my screen. I dial the first three digits. I stop. I don't close the window. I call Yusuf instead, the metabolite kinetics, a genuine and productive conversation, an hour of real science. While the phone rings I notice that I have framed the call as scientifically necessary, and the framing is correct, and correct framings can still be strategic, and the strategy is keeping the window open by replacing one legitimate action with another legitimate action that happens not to involve the EPA.

Each time, in the moment, the decision feels rational. Each time, afterward, I can see the endowment effect in it: the bias pulling the deadline, the window staying open because I want it open, the science serving as cover for the wanting. I log each instance in my notebook. I note the lag between the decision and the recognition. I note that the lag is consistent, approximately twenty minutes to an hour. I note that I am, in effect, running a longitudinal study on my own self-deception, and that the study is producing clean data and the data is not changing the behavior.

I think about my mother. Fourteen days now since I've called. Her reference point hasn't shifted. Her mind has been doing its own quiet version of what's happening to me, but in the other direction and over years instead of days, the shelf the same height, the arm slightly shorter, in Yusuf's phrase from the office, which I have been carrying around all week because it is the most precise description of gradual cognitive decline I have ever heard and

because it applies to two people I care about in opposite directions and I did not connect these facts until this moment.

I stop. The connection — my mother's decline, the city's approaching decline, the compound's mechanism, the phrase I borrowed from Yusuf — is a chain I can trace but I can't hold all the links at once without something dropping. I am enhanced, but I'm not at peak. The bubble is bigger than it was. It is not yet big enough for everything.

* * *

Seven to ten days. That's my estimate for the peak, based on the dose curve and my metabolite trajectory. Day 18 to Day 21. After that, the compound clears and the concentration drops and the inverted-U tells me what happens next.

Seven to ten days of the most capable mind I will ever have. Then the descent. And I will be the first one down.

I close my notebook. Open a new page. Write:

1. Metabolite kinetics — Yusuf. Today. 2. Population timeline — estimate peak distribution across metabolizer profiles. 3. EPA filing — prepare the preliminary report. You have enough now. ~~Almost enough. Enough.~~ File it.

Three items. The third has edits. I leave them visible because they are honest. The scratch-outs are the endowment effect, caught in pencil. Kahneman would recognize them.

The afternoon light fills the lab. I pick up the phone and call Yusuf. The metabolite kinetics are next. Tomorrow I will have one fewer day.

Author's Note — The Science of Chapter 4

The Inverted-U and Prospect Theory

The inverted-U dose-response relationship is one of the most robust findings in pharmacology and cognitive neuroscience. Sometimes called the Yerkes-Dodson law in the context of arousal and performance, first described by Robert Yerkes and John Dodson in 1908, it applies to nearly every neurotransmitter system. BDNF is no exception: both deficiency and excess impair learning, with a clear optimal range between them.

The steepness of the descending slope, the “cliff” in Mara’s data, is consistent with what’s known about excitotoxicity: when NMDA receptors are overstimulated past their optimal range, the excess calcium influx can damage or kill neurons. The same mechanism that drives learning, pushed too hard, becomes the mechanism of injury. This is why the compound’s dual-pathway modulation is so dangerous: it accelerates both the climb and the fall.

Prospect Theory was developed by Daniel Kahneman and Amos Tversky and published in their landmark 1979 paper. Its central finding: people evaluate outcomes relative to a reference point, and losses from that reference point feel approximately twice as painful as equivalent gains feel good. This asymmetry, loss

aversion, produces systematic and predictable errors in decision-making.

Crucially, the reference point is not fixed. It shifts with experience. Once Mara has operated at enhanced cognition for eleven days, that level has become her baseline. Any decline from it will be experienced as loss, regardless of where it lands in absolute terms. A return to the ninety-fifth percentile, a level of cognitive function that most people would consider extraordinary, will feel, from her elevated reference point, like damage.

Mara's application of Prospect Theory to the population is correct and alarming: 180,000 people who are simultaneously in gain frame and will simultaneously shift to loss frame, with all the irrational risk-taking and poor decision-making that loss frame predictably produces.

The pharmacology and the psychology are the same problem: what happens when a system is pushed to a new set point it cannot hold.

* * *

How to Use This — Goal Setting Under Pressure

The inverted-U applies to goals as well as doses. Kahneman's research, along with organizational psychologist Gary Latham's decades of work on goal-setting theory, shows that goals follow the same basic shape: too easy and motivation is minimal; appropriately ambitious and performance peaks; impossibly difficult and motivation collapses.

The stretch goal / proximate goal pairing introduced in Chapter 1's notes is the practical solution to staying on the productive part of the curve. The stretch goal is ambitious enough to engage full effort. The proximate goal is specific enough to be achievable today. Together, they keep you on the ascending slope without pushing you off the cliff.

Mara faces the largest and most frightening problem of her career in this chapter, a population-wide cognitive crisis with a closing personal window. Her response is not to try to solve it all at once. Her response is: *metabolite kinetics today. Population timeline this week. EPA filing when I have enough.* She breaks the impossible into the next bounded task, and then the next one, and then the next.

This is not avoidance. It is the proximate goal structure that makes enormous problems survivable. The reader sees it as Mara's natural competence. It is also a learnable strategy.

The principle: match the size of your goal to the time horizon. For today, ask only what today requires.

CHAPTER 5 — SHE TELLS YUSUF EVERYTHING

DAY 13. Wednesday. 7:22 PM.

Yusuf's office at 7 PM on a Wednesday has a specific quality: the building mostly empty, the corridor lights on their energy-saving amber, the silence of a research building populated only by the people whose work has outlasted the custodial staff and who are, this week, outlasting the custodial staff more often and more productively than usual, though they don't know why.

He has the kettle on. He has the UCL mug. He has, on the edge of his desk, a container of something that smells like his mother's groundnut stew, which means he cooked last night and brought the leftovers specifically because he anticipated that this conversation would take a while and that I would not have eaten. Both predictions are correct. I have not eaten since a granola bar at noon, which is not unusual for me but which Yusuf treats as a moral failing that he is personally responsible for correcting, in the way of a man whose culture organized entire social structures around the conviction that you do not let people you care about

go hungry, and who has never found this conviction at odds with any other conviction he holds.

I have been building a slide deck for the past three hours. This is not something I do — I present from notebooks, from printouts, from my own recall — but this conversation requires organized presentation because I am about to tell my closest colleague that 180,000 people are being cognitively enhanced by a pharmaceutical compound in the water supply and that the enhancement has a ceiling and the ceiling has a cliff. This is not information that benefits from my usual approach of beginning in the middle and expecting the listener to reconstruct the edges.

I need him to understand all of it. I need him to understand it in sequence. And I need him to remain calm enough that I can remain honest, because Yusuf is the only person in my professional life who has figured out that panic makes me lie: not maliciously, not deliberately, but reflexively, the way a system under threat reduces its output to protect its core function. My core function, for better or worse, is thinking clearly. When someone near me panics, the bandwidth I would use for honesty gets redirected toward managing their response. I have been this way since I was old enough to notice it, which was about eight. Yusuf figured it out in the first year of our friendship and has been choosing calm ever since: not as performance but as a deliberate act of respect for the conditions under which I am most truthful. I have never thanked him for this. It is possible that I should. It is also possible that thanking him would require me to acknowledge the

deficiency it compensates for, and I have not yet decided whether the acknowledgment or the silence is the more honest response, which is the kind of question the enhancement makes me better at formulating and no better at answering.

I set up my laptop on his whiteboard ledge. He pours tea for both of us. I don't touch mine. I begin.

* * *

I take him through it in the order I found it. The anomalous battery scores. The water samples. The LC-MS results. The compound's BDNF-mimetic core and the novel moiety that changes everything. The dual-pathway mechanism: structural enhancement and lowered coincidence threshold, synergistic, larger than either alone. The aquifer recharge map. Veridian Pharma in Newberry, eleven miles upstream. The transit calculation. Day 0. Then the dose curve. The inverted-U. The cliff.

Then the population estimate. I write *180,000* on the whiteboard because writing it makes it more real than saying it, which is Carol's inheritance, my mother, the English teacher, who taught me without meaning to that the written word has a different weight than the spoken one, and who would be appalled to learn that I am applying this principle to the number of people being pharmaceutically altered without their knowledge.

It takes about forty minutes. I do not editorialize. I present the data and the analysis and the implications in sequence, the way I

would present a paper, and I watch his face.

Yusuf's face when he is processing something significant has a quality I have observed many times and find worth watching, against most of my instincts toward efficiency. He does not perform thinking. He does not furrow or stroke his chin or adopt any of the gestural vocabulary people use to signal that serious cognition is occurring. He goes still. The warmth doesn't leave — warmth is not a mode Yusuf enters and exits but a substrate he operates on — but the expressiveness narrows to his eyes, which track the information with a quality of attention that is, I suspect, as precise as my models, just organized around different inputs. I model systems. He models people. Between us we cover a reasonable amount of the problem, though not the same parts, which is a sentence I have now applied to both Yusuf and my mother in the past two weeks, which may be telling me something about the kind of collaboration I value and the kind I have and the distance between them.

He asks four questions during the presentation. Each is pharmacological, specific, and correct:

Half-life estimate based on my metabolite clearance data? (I don't have a confident number yet. This is the next problem.)

Is the novel moiety designed or environmental: something that formed in the aquifer rather than in Veridian's facility? (The stereochemistry suggests designed, but I can't rule out transformation without Veridian's precursor inventory.)

Are the dose-response slopes consistent across subjects or divergent? (Consistent in shape, divergent in timing. My CYP450 profile puts me ahead.)

And: what is my current compound concentration relative to the inflection point?

I pause at this one. It is the question that moves from population pharmacology to personal prognosis, and it is the question I have been careful about in my own thinking, approaching it sideways, running the numbers but not quite sitting with what they mean.

"I'm approaching peak concentration," I say. "Based on my metabolite trajectory, I'll reach the inflection point somewhere between Day 18 and Day 21."

"Five to eight days."

"Yes."

He picks up the mug. Both hands, the London habit. He holds it and looks at me and does not speak for approximately fifteen seconds, which is a long time in a conversation between two people who have just been shown the shape of what's coming.

Then he says, very quietly: "Define weird."

The callback lands differently than it did on the phone eleven days ago. On Day 6, *define weird* was an opening: curiosity, an invitation. Here it is a closing of distance. It is Yusuf telling me, in two words borrowed from the beginning, that we are still in the same conversation and that the conversation has changed

scale and that he is going to be calm while we work out what to do about it, because calm is what I need from him and he has decided, again, to provide it.

“Yeah,” I say. “Define weird.”

* * *

The ethical question arrives not as a question but as a silence.

The data is on the screen. The population estimate is on the whiteboard. Yusuf has eaten half the stew. I have eaten none of it and drunk no tea. The building is quiet in the way of institutional silence: not absence of sound but the presence of systems running without people, HVAC and fluorescent ballasts and the elevator cycling to ground because nobody called it.

What I have not said, and what sits in the room as visibly as the slide deck, is why I haven't filed the EPA report.

I have reasons. The reasons are good. Incomplete data in a system lacking the expertise to interpret it produces panic. Panic in a population of 180,000 produces cascading harm: runs on pharmacies, hoarding, overwhelmed emergency departments processing the worried well, political pressure to intervene before the mechanism is understood. The literature on premature disclosure in public health crises supports this concern. The chromium-6 cases. The Flint timeline. I can cite chapter and verse on what happens when accurate but incomplete information enters a system that is not equipped to receive it calmly.

These reasons are real. They are also not all of the reasons.

Yusuf knows this. I can see that he knows it in the way he is not asking me about it, the question held in reserve, not from timidity but from the strategic patience of a man who understands that if he asks too early I will defend, and if he waits long enough I will tell him myself. This is a form of intelligence I do not possess and cannot model. I build computational systems. Yusuf waits seventeen minutes and the truth falls into his lap. I have been watching him do this for eleven years and I still don't know how it works, which I find both professionally humbling and personally indispensable.

He waits long enough.

"I don't want to lose this," I say.

It comes out more plainly than I intended. I had planned to frame it clinically: the endowment effect, the reference point shift, the documented tendency to overvalue what you currently possess. Instead what comes out is five words in a register closer to confession than analysis, and the five words are more accurate than the clinical framing would have been, which is something the enhancement should have helped me predict and didn't.

He sets down the mug. "I know."

"The endowment effect. Kahneman. Thaler. I know what it is. I can feel it operating. The knowledge doesn't help."

"I know that too."

“Every day I have this” — I gesture at my own head, a vague motion less precise than I would normally tolerate — “is a day I can use it to understand what we’re looking at. The mechanism. The population dynamics. The intervention window. Every day I spend filling out EPA forms is a day I’m not using the window.”

“And the window is yours,” he says. “Not the city’s. Yours specifically. Because you hit the peak first.”

This is the thing I was not going to say. He has said it for me. This is — I recognize it — a form of kindness I do not fully deserve and a form of precision I did not expect from someone who arrives at conclusions through means I can’t reconstruct.

“Yes.”

He nods. He does not look angry. He does not look disappointed. He looks like a person thinking carefully about a friend who has just told him something both understandable and wrong, and who is deciding how to say so in a way that preserves the friendship and corrects the error simultaneously. This is, I realize, a harder optimization problem than anything in my IRIS training data, and he is solving it in real time without a model.

“Mara. How long?”

“How long what?”

“How long do you need. Not how long do you want — I know you want the whole window. How long do you actually need to characterize the mechanism well enough that the filing is useful rather than alarming?”

I think about this. I try to think about it honestly, which is harder than it should be, because the endowment effect is pulling the estimate in one direction and Yusuf's calm attention is pulling it in another and neither force is the truth.

"Ten days," I say. "Maybe less. The metabolite kinetics, the population timeline, a preliminary model of the decline phase. Enough to file something useful."

"And if you don't have it in ten days?"

"Then I file what I have."

"Say that again."

I look at him. He is entirely serious. His expression hasn't changed but there is something underneath the calm: not anger, not urgency, something closer to the specific love of a person who is going to hold you to what you said because he believes you mean it and because meaning it matters.

"If I don't have it in ten days, I file what I have."

He picks up a dry-erase marker. He writes a date on the UCL mug, ten days from today, small precise characters, blue ink against white ceramic. Sets the mug on his desk where I can see it.

"There," he says. "Now it's on the mug."

I want to object. I want to say that writing a date on a coffee mug is not a binding agreement, that the situation may change, that ten days is an estimate. I want to say several things that are all true and all, in their truthfulness, in service of keeping the

window open.

Instead I feel something I was not expecting, which is relief. The deadline is external now. I didn't set it. He set it. I cannot move it without moving it in front of him, which means I cannot move it without making the endowment effect visible, which means the bias has a check that my own self-awareness could not provide.

The relief is information, and the information is uncomfortable. It tells me something I have been avoiding: I wanted someone to take this out of my hands. I have been telling myself the delay is scientific judgment: the filing will be more useful with more data, the mechanism isn't complete, incomplete information will cause more harm than waiting. And these are true. And underneath the true reasons is a truer one: I wanted to keep the window open and I couldn't trust myself to close it and I am relieved that Yusuf closed it for me.

The relief is the proof. The endowment effect was not just present but dominant. My scientific judgment, the instrument I have relied on for twenty years as my most trustworthy, was, in this specific instance, serving the bias rather than overriding it.

I file this. It is the most important thing I have learned about myself today, which is saying something given what I have also learned about the water supply.

"Thank you," I say.

He looks surprised. This is not something I say often. "For what?"

"For not trusting me to set my own deadline."

He almost smiles. "You're welcome. Eat the stew."

* * *

I eat the stew. It is very good, which I acknowledge, and Yusuf says his mother's recipe is always excellent and that the quality is in the groundnuts, which must be sourced correctly, and that he has been telling me this for three years and I have not listened. This is true. It is also true that his patience with my not-listening extends beyond groundnuts into areas I do not always acknowledge.

Then I tell him about IRIS.

He knows I have a research system. He knows I train it on my own work. He does not know what it has become in the past two weeks, because what it has become in the past two weeks is something I am still working out the vocabulary for, which is itself a tell, because I am a person who usually has the vocabulary before the experience, and the experience outpacing the language is a new and vertiginous sensation that the enhancement has been producing more and more often.

"It's making connections I didn't teach it," I say. "Not random. Sound: inferences I can verify, that hold up, but that aren't in the training data. It's extending my reasoning into territory I haven't mapped."

He leans forward, the posture shift from warmth to focus, the lean of a scientist recognizing a finding. "How far beyond?"

“Last week’s output drew a connection between the Kozlov paper and my replay model that I hadn’t made. I checked it. It held. And it flagged the inference as uncertain before I asked.”

“Flagged its own uncertainty.”

“That’s what I trained it to do. But the calibration is — it’s better than it was a month ago. More precise. It knows what it doesn’t know with a resolution it didn’t have before.”

He sits with this. The mug is in front of him, the blue date still fresh. Outside the window, the campus dark.

“Does it know what’s happening?” he asks. “About the compound. About us.”

“It knows what I’ve told it.” I pause. “Which is getting to be a lot.”

He looks at me for a long time. I can see him thinking, which is not something most people let you see. Most people show you the output of thinking: the conclusion, the position. Yusuf lets you watch the processing, the consideration moving across his face the way weather moves across a landscape, visible and unhurried. I find this — and I notice the finding as an ascending-level observation, the kind of lateral connection between a friend’s face and a weather system that I would not have made three weeks ago — I find this a quality I trust more than I trust most data.

“Mara,” he says finally. “If your window is closing — if the peak is in five to eight days and the decline starts after that — then anything you pour into IRIS during the window stays there after

the window closes.”

“Yes.”

“You’re building a record of yourself at your best.”

“I’m building a research tool.”

“You’re building both.”

I don’t answer immediately. There is a version of this conversation where I defend the framing: IRIS is a research system, the training is methodology, the externalization of my reasoning is standard scientific practice. There is a version where I acknowledge what he’s actually saying, which is that I am pouring myself into something that will outlast the version of me doing the pouring, and that the urgency with which I am doing it is not entirely scientific.

“If I am building both,” I say, “it’s because both need building.”

“What happens to the model after the window closes?”

“It keeps running. The reasoning frameworks, the annotation patterns, the uncertainty calibration: that persists. IRIS doesn’t metabolize. It doesn’t have a CYP450 profile. The version of my thinking that I’m training it on right now will still be operating at this resolution after I can no longer operate at this resolution.”

“So the system continues at peak after you leave peak.”

“Yes.”

“The system that learned to think like you, continues thinking like you at your best, while you, the person it learned to think

like, declines.”

“Yes.”

He is quiet for a moment. The tea has gone cold. The stew container is empty.

“That’s not just a research tool, Mara.”

I don’t answer, because the answer would require me to acknowledge something I have not quite articulated even to myself: that what I am doing with IRIS has the character of an act I associate with people who know they are running out of something. The urgency of it. The thoroughness. The specific quality of care you bring to a thing when you know you are building it for someone who will need it after you can no longer maintain it yourself.

My father kept a filing system his entire career. Meticulous. Every document in its place. He left it to my mother, who uses parts of it to this day. He didn’t call it preservation. He called it organization. But the organization was how it survived him.

I don’t say this. Yusuf hears it anyway, in the way he hears most of the things I don’t say, which is one of several reasons I came to his office tonight and not to the EPA’s regional number.

* * *

We sit for another hour, talking through the next ten days: what needs to happen, in what order, with what resources. The metabolite kinetics. The population timeline. The decline-phase mecha-

nism. The questions I need to answer before the filing is useful. He takes notes on a legal pad in handwriting that is less precise than mine but covers more ground per page because he abbreviates in a way I find aesthetically distressing and functionally adequate.

At 9:30 I gather my laptop and the printouts and the untouched cup of tea, long cold. Yusuf walks me to the door, which he does not always do, and which tonight is a form of something I don't have a clean word for: not protectiveness exactly, not tenderness, something adjacent to both. The act of a person who watched a friend disclose something that cost her, and who wants the last image of the evening to be the two of them standing in the same space.

"Ten days," he says.

"Ten days."

"And call your mother."

I stop. This is not something I told him about. The missed calls, the pattern, the task list; none of it has been part of tonight's conversation.

"How did you know about that?"

He almost smiles again, the expression he makes when he has said something that surprised me, and I realize, not for the first time, that his intuition is a form of intelligence I do not have and cannot build and that is, in its own way, as precise as anything I've trained IRIS to do. More precise, maybe. IRIS models reasoning.

Yusuf models people. The people are harder.

“You always call your mother when things are normal,” he says. “You haven’t mentioned her in two weeks. That means things aren’t normal and you haven’t called.”

I stand in the doorway. The corridor is amber and quiet. The elevator hums in its shaft.

“Goodnight, Yusuf.”

“Goodnight, Mara. Eat something before you sleep.”

I walk back to my lab. The date is on his mug. The deadline is in front of me. I resent it: a clean, precise resentment I can identify as the endowment effect operating exactly as predicted, the bias pulling against the constraint, the part of me that wants the window open fighting the part that knows it needs a frame.

The resentment does not cancel the gratitude. They sit together, neither yielding. The feeling and the knowledge, side by side, the way they’ve been sitting since Day 11: both present, both true, neither winning. Kahneman could have told me this would happen. He did tell me. It didn’t help.

Nine days. I have nine days to do the work that matters most while I am still the person who can do it.

I sit down at my desk. I open IRIS. I start talking.

Author’s Note — The Science of Chapter 5

The Endowment Effect and the Ethics of Non-Disclosure

The endowment effect, the tendency to overvalue something simply because you currently possess it, was documented by Richard Thaler and further developed in a landmark 1990 study by Kahneman, Knetsch, and Thaler. The finding is robust: people consistently demand more to give up an object they own than they would pay to acquire the same object, even when the object was assigned to them randomly moments before.

Applied to cognition, the implication is immediate: a person who currently possesses enhanced reasoning will systematically overvalue that enhancement and resist any condition that might end it — even when ending it is the rational choice. This is not selfishness in the moral sense. It is a documented feature of human decision-making that operates independently of intelligence, self-awareness, or ethical commitment.

Mara’s decision to delay the EPA report is not straightforwardly selfish. The literature on research disclosure ethics in public health emergencies genuinely supports a “complete before reporting” standard in some contexts; incomplete information can cause harm through panic, misattribution, and policy overreaction. The tension between this legitimate consideration and her personal stake in the delay is precisely the kind of conflict behavioral economics predicts will be poorly resolved by even sophisticated thinkers. Knowing about a bias does not reliably

protect against it.

Mara knows this. It does not help as much as she would like.

* * *

How to Use This — Motivation and Locus of Control

Charles Duhigg’s research on motivation, drawing on work by psychologist Julian Rotter and subsequent studies by Carol Dweck and others, centers on a concept called locus of control: the degree to which a person believes their choices, rather than external forces, determine their outcomes.

Internal locus of control, the conviction that you are an agent, not a subject, is one of the strongest predictors of sustained motivation across domains. People who feel acted upon by external forces stop trying. People who feel that their choices matter keep going.

The critical practical insight: you can shift your locus of control by converting situations into choices. Even when a situation is externally imposed, finding and making a genuine choice within it restores the sense of agency that motivation requires.

Watch what Mara does in this chapter. The contamination is happening to her city. The compound is in her body. The cliff is coming. She cannot stop any of this. What she can choose: how to use the window she has. When to file the report. What to pour into IRIS while she still can. She converts every imposed situation into a decision. This is not denial. It is the locus of control habit operating under extreme pressure, and it is learnable.

The principle: you cannot always choose what happens. You can almost always choose what you do next. Start there.

CHAPTER 6 — THE METABOLITE PROBLEM

*D*_{AY} 15. Friday. 2:09 AM.

The compound breaks down into three stable metabolites. All three are pharmacologically active.

I have been sitting with this result for eleven minutes. I know this because the mass spectrometer's output log is timestamped and the timestamp says 01:58 and it is now 2:09, and the eleven minutes between those two numbers contain the most consequential pharmacokinetic finding I have ever made, and I have not moved or spoken or written anything in my notebook, which is unusual for me and which I note even as I fail to correct it. The noting-without-correcting is itself unusual. I am a person who acts on observations. Tonight the observation has outpaced the acting, the way a wave outpaces the beach it's about to hit.

Here is what this means, and it changes everything.

When the compound enters the body, the CYP450 enzyme system breaks it down, standard metabolism, the liver doing what the

liver does with foreign molecules, which is disassemble them with the indifference of a system that has been processing xenobiotics for four hundred million years of vertebrate evolution and does not consult the prefrontal cortex about which ones the prefrontal cortex would prefer to keep. The liver does not care about your cognitive enhancement. The liver has a job. The liver does its job.

The three metabolites that result are structurally distinct from the parent compound but retain its pharmacological activity: both the BDNF-mimetic effect and the NMDA modulation from the novel moiety. They are smaller molecules, more lipophilic, and they cross the blood-brain barrier more readily than the parent. They are, in pharmacological terms, more potent than the compound that produced them. The children are stronger than the parent. This is not common in drug metabolism. It is not unheard of. It is, tonight, the most personally relevant pharmacological fact I have ever encountered.

And the children are patient. Their metabolic half-lives, the time it takes the body to clear them to half their peak concentration, range from eighteen to thirty-six days depending on the specific metabolite and the individual's enzyme profile. The parent compound has a half-life of four to six days. Its breakdown products persist for weeks.

This means the enhancement does not end when the compound clears the blood. The metabolites continue driving BDNF activity and NMDA modulation for three to five weeks after the parent is gone. The ascending slope, the plateau, the initial descent: these

are functions of the parent's concentration. But the tail of the curve, the extended decline, the slow slide back toward baseline: that is driven by metabolites already in the system, committed to their own half-lives, clearable by nothing short of time and the liver's patient, indifferent work.

Even if the runoff stopped today. Even if the aquifer flushed clean tomorrow. Even if 180,000 people switched to bottled water this morning.

The trajectory is set. The metabolites are in the system. The cliff is not coming. The cliff is already built. It was built the moment the compound entered the aquifer and traveled through eleven miles of limestone to arrive at 180,000 kitchen taps. The only question now is when each person reaches the edge, and that question has an answer, and the answer is written in their CYP450 genes.

* * *

The CYP450 system. Dozens of enzyme subtypes, each encoded by genes that vary substantially between individuals: ancestry, diet, medications, a dozen variables modulating activity up or down. The result: the metabolic rate for any given compound is not a constant. It is a distribution. A bell curve of processing speeds that will produce, when applied to the compound and its metabolites, a bell curve of peak timing: not everyone reaching the cliff at once but a rolling wave of peaks and descents spreading across the city over months.

Rapid metabolizers, my category, run the curve in weeks. Slow

metabolizers run it in months. Between the extremes, the bulk of the population processes at intermediate rates, reaching their peaks in a staggered sequence that looks, when I model it, like a wave moving through the city. Rising, cresting, breaking.

The staggering is, in a grim actuarial sense, the best available news. A simultaneous cognitive cliff across 180,000 people would overwhelm every system a city depends on: medical, institutional, social, the infrastructure of daily life. A rolling wave is manageable, or at least survivable, because the people still ascending or at peak can compensate for the people already in decline. The city doesn't fail all at once. It fails in sequence. Each failing segment has the support of the segments that haven't failed yet.

I am holding, as I think this, the population dynamics model and the pharmacokinetic variation and the institutional resilience implications simultaneously, three frameworks from three different decades of my education, running against each other, arriving at a conclusion that requires all three and that I could not have reached two weeks ago. I can feel the enhanced cognition doing this. I can feel it enabling the analysis that characterizes the enhanced cognition. The instrument studying itself, with increasing precision, while the thing being studied is the thing making the precision possible. The circularity would be elegant if it weren't happening to me and if the instrument weren't scheduled for degradation on a timeline I can now calculate.

* * *

I run my own metabolite panel at 3 AM. The results take twenty minutes. I already know what they will show, because the dose-response data predicted it and because Yusuf told me on Wednesday that my clearance rate was in the top five percent.

But there is a difference between an estimate from population data and a direct measurement from your own blood. The estimate says: you are probably a rapid metabolizer. The measurement says: you are, specifically and individually and irrevocably, clearing this compound at a rate that places your peak concentration window between Day 17 and Day 20.

Two to five days from now.

After that, the parent compound drops below the enhancement threshold. The metabolites carry the effect forward, but at declining potency, a slow slide rather than a cliff. The cliff comes when the metabolites themselves clear below threshold. For me, given my profile: approximately Day 45 to Day 55. Six to eight weeks from now.

I do the math twice. The math does not change. I do it a third time because the repetition is not about checking; it is about the human need to run a terrible number through your mind enough times that it stops feeling like news and starts feeling like a fact you can plan around. Three repetitions is enough. The number is a fact now. I can plan around it.

My window of peak cognition, the days during which I am operating at or near the ceiling of what this compound can do to

a human brain, is measured in single digits. The period after peak, during which I will experience the progressive decline of the capabilities I am currently using to perform this analysis, is measured in weeks. And the timeline is determined entirely by my own biochemistry, which I did not choose and cannot modify and which is processing the compound with the cheerful efficiency of an enzyme system that has not been informed of the compound's cognitive significance and would not adjust its behavior if it had, because enzymes are proteins and proteins do not read lab notebooks.

I write in the margin: *Note to self: optimization doesn't waive physics.*

It is the first thing I have written in fifteen days that is not an observation or an analysis or a task. It is a joke. It is also, I suspect, the most precise thing I have written tonight, because the compound is optimizing my cognition and my CYP450 system is metabolizing the compound and neither process has consulted the other and neither can be asked to stop. The engine that is making me sharper and the engine that is ending the sharpness are the same engine, running in the same cells, following the same chemistry, indifferent to each other and to me. Optimization does not waive physics. The bridge does not negotiate with the river. My father, the civil engineer, would have understood this immediately: you build with the physics you have, not the physics you want, and the quality of the engineering is measured by how well the structure holds given the forces it cannot change.

I miss him. This arrives without warning, at 3:30 AM, in the

margin next to a pharmacokinetic joke, and I let it sit there without examining it, because some observations are data and some are weather, and weather passes if you don't try to model it.

* * *

I push back from the desk and look at the wall above my monitor. The aquifer recharge zone map has been pinned there since Day 8. USGS geological survey: contour lines, well field locations, recharge zones in blue. Next to it, as of tonight, the metabolite half-life calculations on a printout I haven't pinned yet. The map shows how the compound arrived. The printout shows where it goes. Between them, the distance from geology to pharmacology to 180,000 individual futures is the kind of span my mind has been learning to hold this week, the radius expanding day by day until I can stand in my lab at 3:40 AM and see, in a single sustained act of perception, the aquifer and the liver and the bell curve and the cliff and the wave and Ponce de León and my own bloodwork, all of it connected, all of it real, all of it running on the same clock.

The water was always real. The chemistry is real. The enhancement is real. And the enhancement has built its own end into its own structure, the metabolites carrying the parent's work forward until they too clear, because nothing the body receives is permanent and the liver does not make exceptions for things you want to keep. De León was looking for water that restored what time had taken. He found the right place. He asked the

wrong question. The water does restore. And then the restoration metabolizes, because that is what molecules do in bodies, and the Fountain of Youth was always a pharmacokinetic problem disguised as a legend, and I understand this now with a clarity I suspect I will not have for much longer, which makes the understanding itself a kind of proof of what the legend was trying to say and failing.

I have never understood the Fountain as well as I do right now, at 3:40 AM, alone in Gainesville, staring at a geological survey map and a metabolite calculation and holding, briefly, the whole shape of the story — the aquifer, the compound, the enhancement, the decline, the ancient water, the human wanting — in a mind that is, by its own measurement, two to five days from the moment it begins to lose the ability to hold shapes this large.

* * *

At 4 AM I open IRIS and upload everything.

The metabolite structures. The half-life data. The CYP450 variation models. The population distribution estimates. My own panel. The dose-response curve with the metabolite persistence overlaid. Everything I have learned tonight, organized in the annotation format IRIS was trained on, cross-referenced to everything that came before.

I type: Build a degradation timeline for the exposed population. Model the distribution of peak timing, onset of decline, and projected cognitive trajectory by metabolizer subtype. Flag any variables you've estimated

rather than received.

I queue the batch. I save my notebook, page 47, three and a half pages of tonight's work, the handwriting steady, the annotations clean, the "Note to self" sitting in the margin like a small dark flag marking the place where the science turned personal.

I should sleep. The compound will metabolize whether I'm conscious or not. The math will hold. The metabolites will hold. The cliff will be the same number of days away. My mother will be asleep in Tampa. Fifteen days since I've called. She has not left another voicemail since the one about Thanksgiving. The absence of the voicemail is louder than the voicemail was.

I walk home under the oaks on University Avenue, in the dark, the route automatic, procedural memory carrying me while the rest of my mind is still in the lab with the half-life data and the recharge zone map and the joke that isn't a joke. The campus is quiet. The tree frogs are loud. Something about 3 AM in Florida produces a specific quality of sound: the insects and the frogs and the cooling systems all running at their own frequencies, layered, a kind of music that my mother would describe better than I can and that I am, tonight, hearing with an attention I don't usually bring to the walk home.

I sleep for four hours. The compound metabolizes while I sleep. The metabolites accumulate. The clock runs.

* * *

I arrive at the lab at 8 AM and the IRIS output is waiting.

It is more sophisticated than what I fed it. Substantially more. IRIS has taken my data and built a population degradation model that includes three variables I did not provide.

IRIS // Day 16 // 06:44:19

Population degradation model v1.2

Note: Three variables in this model (CYP450 subpopulation distribution, BDNF Val66Met frequency in Alachua County, baseline cognitive variance by age cohort) were not provided in input data. I have estimated these from published population genetics literature. These estimates include uncertainty. Estimated confidence in timeline accuracy: 0.68 ± 0.05 . Recommend verification against local biobank data if available.

I read this twice.

CYP450 subpopulation distribution: the frequency of different metabolizer types in this specific county. I don't have it at the local level. IRIS estimated from published pharmacogenomic data, which reports these frequencies by ancestry and region. Reasonable. Uncertain. Quantified.

BDNF Val66Met frequency: the genetic variant that modulates the compound's enhancement. Roughly 30 percent of the general population carries at least one copy. IRIS needs the local number. I didn't give it. IRIS found the published data.

Baseline cognitive variance by age cohort: the natural distribution of cognitive capacity across age groups, which determines every individual's starting point on the dose-response curve. I didn't

think to include this. IRIS did.

Three variables. Each inferred from literature. Each flagged as an estimate rather than a measurement. Each with a confidence interval and a recommendation for verification.

It flagged what it didn't know before I asked. And it identified what I forgot to give it, which is a different skill: the first is epistemic humility, the second is anticipation, and the combination of the two is something I have been trying to build for three years and that is now, at 8 AM on a Friday morning, sitting in my output queue as though it were the most natural thing in the world for a computational system to do.

The morning light reaches the back wall of the lab, that thirty-minute window when the angle is low enough to fill the room entirely. I open the full model. The degradation timeline spreads across the screen: overlapping curves, each one a metabolizer subtype, each peaking and declining at a different rate. The fastest metabolizers peak earliest and fall first. The slowest peak months later. Between them, the city's cognitive trajectory looks like a wave: rising, cresting, breaking, rolling forward through the population over twelve to sixteen weeks.

The wave is already in motion. It started on Day 0. Nothing stops it now.

I print the model and pin it to the wall next to the recharge zone map. Two documents: how the compound arrived, and where it goes. Between them, the aquifer and the biology and the math-

ematics of 180,000 futures described by curves they don't know about, driven by metabolites they can't feel, on a timeline they didn't choose.

Three to five days until my peak. The math hasn't changed. I have checked it enough times that it has graduated from news to fact.

* * *

The deadline on Yusuf's mug is in five days. The metabolite data makes the EPA filing substantially more useful than it would have been without it: the population timeline gives them something actionable, a model of what's coming and when. Yusuf was right about the deadline. He is often right about things I don't want him to be right about, which is one of several reasons I trust him and one of several reasons I find him occasionally maddening and these two facts are not in conflict.

I call him at 8:30. I tell him what the metabolites do. I tell him what IRIS built overnight.

"How confident is the model?" he asks.

"Point six eight plus or minus point one one."

A pause. "That's honest."

"IRIS is always honest. That's the thing I care about most."

"And your timeline? For peak?"

"Three to five days. Then the slide."

Another pause. The specific silence of a person deciding what to say to a friend who has just delivered her own prognosis with the clinical precision of someone reporting a lab result, which is what I am doing, and which is the only way I know how to deliver information that is both scientifically exact and personally devastating, and the fact that these two qualities coexist in the same sentence is not lost on me and is not something I am going to examine right now, because examining it would require me to feel it, and feeling it would cost time I don't have.

"Come have breakfast," he says. "I made eggs."

I don't want eggs. I want three more days at peak. I want the metabolites to be wrong. I want the CYP450 system to take a holiday and the liver to forget its job and the physics to waive, just this once, for someone who is using the optimization well and who has not finished using it.

None of these wants are available. The eggs are.

"I'll be there in four minutes," I say.

Author's Note — The Science of Chapter 6

CYP450 Enzymes and Why Drugs Hit People Differently

The cytochrome P450 (CYP450) enzyme family is responsible for metabolizing the majority of drugs and foreign compounds that enter the human body. There are dozens of CYP450 subtypes, and the genes encoding them vary substantially between individuals.

These variations mean that the same compound, at the same dose, can have wildly different effects in different people: some metabolize rapidly, experiencing a sharp peak and fast clearance; others metabolize slowly, with a gentler curve and a longer tail.

This is not a minor footnote. CYP450 variability is the reason codeine is dangerous for some patients and ineffective for others. It is why some people need twice the standard dose of a medication while others are overwhelmed by half. Pharmacogenomics, the study of how genetic variation affects drug response, is a growing clinical field precisely because this variability is so consequential and so underappreciated in standard prescribing practice.

Mara's rapid metabolizer profile is pharmacologically plausible. Rapid CYP450 metabolizers exist in abundance in real populations, with frequencies varying by enzyme subtype and ancestry. They genuinely burn through compounds faster, reaching peak concentrations sooner and clearing them more quickly. The cliff Mara calculates is the natural consequence of her own biochemistry processing the compound with more efficiency than she would like.

The metabolite persistence that drives this chapter, active breakdown products with half-lives measured in weeks, is consistent with real pharmacology. Many drugs produce active metabolites that outlast the parent compound and contribute significantly to the therapeutic (or toxic) effect. This is why drug interactions can persist long after one of the interacting drugs has apparently

cleared.

The genome doesn't care about your research timeline.

CHAPTER 7 — PEAK

*D*_{AY} 18. 11:47 PM.

The cognitive battery scores have plateaued across three consecutive days — fluid reasoning, processing speed, working memory, and attention switching all holding at what appears to be their maximum achievable value under the compound’s influence — and what I notice most is not the numbers but the quality of the silence between them, the way the plateau feels from the inside, which is not like a ceiling at all but like the moment a lens finishes focusing and the image becomes as sharp as the optics allow and you realize that what you were seeing before was not the world but the world plus the distortion of the instrument, and the instrument was you, and now the instrument has been corrected to a degree you didn’t know was available, and the world without the distortion is not different but more: more detailed, more connected, more legible in its architecture, the way a city seen from an airplane is the same city you walked through at street level but the streets make sense now, the grid is visible, the logic that was always there is suddenly, obviously, completely apparent.

Everything is clear. Not in the euphoric sense; I am not high, not manic, not in any state I would describe as altered. I am in the state I would describe as the one I was always approaching. The difference between noise and signal, which has been narrowing since Day 6, has reached something close to its minimum, and what remains when the noise is gone is not a different kind of thinking but the same thinking I have always done with less interference, less friction, less of the cognitive overhead that I never noticed was there until it wasn't, the same way you never notice the weight of the atmosphere until someone shows you what it would feel like without it.

I have been working for nineteen hours. I know this because the lab's entry log shows me badging in at 4:52 AM and it is now 11:47 PM, and the nineteen hours between those timestamps feel like — I am trying to be precise because the subjective experience at peak is data I will want later, the way a mountaineer photographs the summit knowing the weather will close — approximately six. The compression of subjective time is itself a finding: when every cognitive cycle produces output, the felt duration per cycle decreases, not because time is passing faster but because the ratio of thinking to waiting has shifted so far toward thinking that the waiting, which is what makes time feel long, has nearly disappeared. I am not working faster in the clock sense. I am working with less idle between operations. The processor is the same. The throughput has changed.

The Florida night outside the west-facing windows is absolute

and close, the kind of dark that September produces when the clouds are low and the humidity traps the city's light against the ceiling. I have not looked at the windows since sometime around 4 PM, when the afternoon light was doing its thing and I was aware of it for approximately one second before the degradation model pulled me back.

I am modeling the population degradation. IRIS's v1.2 timeline is the foundation, but I have been building on it all day, feeding it the refined CYP450 subpopulation distributions that Yusuf pulled from the county health department's anonymized prescription records, and the three additional weeks of my own battery data that now extend the dose-response curve past the plateau and into the early descending slope for my metabolizer profile. The model is on its fourth major revision and it is substantially more detailed than what I started with this morning, and I am holding the entire structure in working memory while I modify individual parameters, which is something I could not have done at baseline: the model has seventeen variables and I am tracking all of them simultaneously, aware of how each modification propagates through the others, aware of which interactions are linear and which are nonlinear and which I haven't characterized yet, and I am aware, in the way you are aware of your own breathing when someone asks you about it, that this level of cognitive integration is the compound's gift and the compound's gift has an expiration date that is, by my own calculation, approximately three days away, and I am aware that the calculation of the expiration date is itself a product of the gift, and that this recursive quality, the en-

hancement enabling the measurement of the enhancement's end, is either the cruelest or the most generous thing the compound does, and I have not decided which, and I suspect the answer is that it is both, simultaneously, which is a thing I can hold at peak and which I will not be able to hold in approximately seventy-two hours.

* * *

At 8:30 PM I notice the anomaly.

I almost miss it, which is itself informative: at peak cognition, with seventeen variables in working memory and the full architecture of the degradation model loaded and running, I almost miss the thing that turns out to be the most important pattern in the data. The pattern is not where I am looking. I am looking at the population-level timeline, at the wave of peaks and descents spreading across the metabolizer subtypes, at the institutional implications of a rolling cognitive decline in a city of 180,000 people. I am looking at the macro. The pattern is in the micro.

It is in the sequence.

Not the timing — I have the timing modeled. Not the magnitude — the inverted-U predicts the magnitude. The sequence. The order in which different types of cognitive function degrade as the compound's effect wears off. I have not been paying attention to the order because I assumed it would follow the compound's mechanism: that the functions most enhanced would be the functions most degraded, in roughly reverse sequence, the way you

unwind a spool by running the reel backward. This assumption is so natural, so pharmacologically intuitive, that I made it without testing it, which is the kind of error that peak cognition does not protect you from, because the error is in what you assume rather than what you calculate, and assumptions are invisible precisely to the degree that they are intuitive.

This is not what the data shows.

In my four subjects, myself, Yusuf, Elena, James, the degradation sequence is consistent but it does not track the enhancement sequence. Working memory declines first, before other domains show measurable change. Then episodic memory, the time-stamped, autobiographical, hippocampally dependent system that records what happened and when. Then fluid reasoning begins to drop, followed by processing speed. Semantic memory, factual knowledge, vocabulary, conceptual understanding, is barely affected. Procedural memory, skills, habits, the motor sequences that live in the basal ganglia, is not affected at all.

I stare at this sequence for a long time. I know it. I know it the way you know the layout of a room you've been in before: not from the degradation data in front of me but from somewhere else, somewhere in the twenty years of literature I have been reading and teaching and building on. For several minutes I cannot place the recognition because it is arriving faster than the citation, the pattern identified before the source is retrieved, which is itself a peak phenomenon: at this level of integration, recognition operates across the entire literature simultaneously,

pattern-matching against everything I have ever read, and the match arrives before the bibliographic metadata catches up.

Then the source arrives. And I stop breathing for approximately four seconds.

* * *

The sequence I am looking at, working memory first, then episodic, then fluid reasoning, then processing speed, with semantic and procedural preserved, is not the degradation signature of a BDNF/NMDA compound declining past its optimal range. That signature would be different: NMDA-dependent functions declining first, BDNF-dependent functions following, the order determined by receptor density and metabolite clearance rates. The compound's own pharmacology predicts a specific sequence, and this is not it.

What I am looking at is the presymptomatic signature of neurodegenerative disease.

Not one disease. The sequence — and I am checking this as I think it, pulling papers from memory with a speed and accuracy that I recognize as the compound operating on twenty years of absorbed literature, the pharmacological enhancement and the domain expertise combining at a resolution I have never had before and will never have again — the sequence maps onto the earliest detectable stage of a cascade that has been documented in fragments across multiple disease categories. Working memory decline as the first measurable signal, years or decades before

clinical diagnosis. Hippocampal episodic memory following. Fluid reasoning third. Observed in presymptomatic Alzheimer's cohorts. In Lewy body prodromal studies. In the earliest stages of frontotemporal degeneration. Observed separately, classified separately, attributed to different pathologies, because nobody has had a unified view of the presymptomatic phase across all three disease categories at high enough resolution to see that the sequence is the same.

It is the same.

The compound has, by dramatically accelerating the enhancement-and-decline cycle, compressed into weeks a process that normally takes years or decades to become visible: it has produced, in 180,000 people, a time-lapse view of the presymptomatic neurodegenerative cascade, the earliest stage of the process by which human minds fail, rendered visible at a temporal resolution that has never existed, because no one has ever had a population simultaneously enhanced by a known mechanism and then declining from that enhancement on a documented timeline with measured biomarkers and tracked cognitive profiles across four memory systems simultaneously.

I am looking at a window that has never existed. A window into the mechanism of neurodegeneration at a stage so early that, in normal aging, it is invisible, buried under decades of compensatory adaptation, masked by the brain's capacity to reroute around damage, hidden by the fact that the process is too slow and the measurement too crude to catch it before clinical symp-

toms emerge and the damage is already substantial and largely irreversible.

The compound stripped all of that away. The acceleration made the sequence visible. The sequence is the same across what we have been calling different diseases.

Which means they may not be different diseases. They may be different phenotypic expressions of a single underlying cascade, expressed differently depending on which brain regions are most vulnerable in each individual, determined by genetics and environment and the specific architecture of a person's neural connections, but driven by the same fundamental mechanism, the same molecular domino chain, the same cascade that I can now see because the compound showed it to me in weeks instead of hiding it across decades.

I sit with the screen.

The data doesn't change. It has been this for days. The data was always this. I am the thing that changed: the lens that focused, the resolution that crossed the threshold, the twenty years of literature that reached the density where the pattern could emerge from the fragments. The compound didn't create the discovery. The compound created the discoverer. The finding was in the data all along, the way the water was in the aquifer all along, the way the mechanism has been running in every aging brain since before there were brains capable of studying it.

I sit with this. The lab is silent. The night outside is absolute. I

have not moved from my chair in I don't know how long. The population degradation model is on the screen and it is not what I thought it was; it is not just a timeline for 180,000 people coming down from a pharmaceutical accident in Gainesville. It is a map of how minds fail, rendered at a resolution the field has never had, visible only to someone holding all four memory systems at maximum capacity simultaneously, which I am, right now, for approximately three more days.

* * *

And there is something inside it. Something at the edge of what I can perceive, requiring everything I currently have to hold it in focus: the seventeen variables and the cascade model and the literature and the compound's pharmacology and the receptor taxonomy all held at once, a cognitive load that I could not have carried a month ago and that I can carry now only because every system is running at its ceiling, and the thing at the edge of what I can see is this:

The cascade has steps. The dominos fall in sequence: working memory first, episodic second, fluid reasoning third. And each step triggers the next through a mechanism that I can partially reconstruct from the degradation data and the compound's known pharmacology. The trigger between steps is molecular: as one system declines, it produces downstream effects that destabilize the next system in the sequence. The cascade is not simply a list of things that fail in order. It is a causal chain. Each failure causes

the next failure.

Which means there are points in the chain where the causation could be interrupted.

I go through the cascade step by step, and I am holding the whole chain while I examine each link, which is what peak is for, which is what the compound gave me the capacity to do and what the compound's metabolism will take away. Working memory decline produces excess glutamate in the prefrontal cortex as NMDA receptors destabilize. The excess glutamate, acting on hippocampal circuits already stressed by the declining enhancement, accelerates the transition from working memory decline to episodic memory decline. First domino hitting second.

Between the first domino and the second — between the prefrontal glutamate surge and the hippocampal destabilization — there is a step. A specific molecular event. A point at which the glutamate signal is transduced from the prefrontal system to the hippocampal system through a pathway involving a receptor subtype I recognize from a paper I read three years ago, Bhatt et al., 2021, *Journal of Neuroscience*, the mGluR5 allosteric modulation data, and that I can now, at peak, connect to the degradation data with a specificity I could not have managed a week ago because a week ago I could not hold both the receptor pharmacology and the population data at the same time.

The receptor subtype is druggable. There are existing compounds, in clinical development, in the published literature, pharmacologically accessible, that modulate this specific receptor. If a correctly

designed molecule were introduced at this step, it could interrupt the cascade at the first transition. Working memory decline would still occur; the first domino has already fallen. But the second domino would not fall. Episodic memory would not degrade. The cascade would stop.

Not a cure. An interruption. The first domino falls, but the chain breaks at the first joint, and everything downstream is preserved.

I think about what this means and I have to get up from my chair because thinking about it while sitting feels physically insufficient for the scale of the thought, which is not a sensation I have experienced before and which I record here because it is as specific to peak cognition as the time compression or the seventeen-variable working memory or the speed of the literature retrieval: the discovery is physically too large for the body I am sitting in, and the standing is not a choice but a necessity, the way breathing is not a choice.

Hundreds of millions of people worldwide live with neurodegenerative disease. Alzheimer's alone affects an estimated 55 million. Lewy body dementia, frontotemporal degeneration, vascular dementia. If the cascade model is correct, if these are phenotypic expressions of the same underlying process, then the intervention point I am looking at is not a treatment for one disease. It is a treatment for the mechanism that produces all of them.

The mechanism that produces all of them, identified because a pharmaceutical company in Alachua County had a waste management failure and contaminated an aquifer that carried a com-

pound through ancient limestone to 180,000 kitchen taps, and the compound pushed human cognition past its limits, and the decline from those limits made visible the earliest stage of a process that has been killing minds for as long as minds have existed. And I can see it, right now, standing in a lab in Gainesville at 11:47 PM, holding the whole chain in my head, because the compound that caused the problem is the same compound that equipped me to see the solution, and both of these facts are running on the same clock, and the clock has approximately three days left.

This is what the water showed me. Not the enhancement — the enhancement was the instrument, not the finding. The finding is in the degradation. The compound's decline revealed something that the ascent could never have shown, because the ascent is pharmacological noise and the decline is biological signal, and the signal was always there, in every aging brain, in every mind that has ever slowed and diminished, running in the dark like water through limestone, too slow to see, too gradual to measure, until someone accidentally compressed it into weeks and a neuroscientist at peak happened to be watching.

The water doesn't restore. It reveals. What it reveals is the mechanism by which time takes, and inside that mechanism — barely visible, requiring everything I have — is the possibility of interrupting the taking.

De León was looking in the right place. He was asking the wrong question.

* * *

I think about my mother.

This is not a planned observation or a deliberately drawn connection. It is simply there, arriving in the space between the discovery and the next action the way thoughts arrive at peak: fully formed, already connected to everything else, present in working memory alongside the cascade model and the intervention point and the population timeline and the seventeen variables and the specific pharmacology of the mGluR5 receptor and the Bhatt paper and the metabolite half-lives and all of it held simultaneously, including this:

My mother is seventy-four years old and her words take slightly longer to arrive than they used to. She has not talked about it. I have not talked about it. The mechanism I am looking at right now — at the ceiling of my cognitive capacity, in a lab in Gainesville, standing because the discovery won't fit in a chair — may be the mechanism that is happening to her right now, slowly, in a house in Tampa where the garden my father started is still growing and the jalapeños produced fourteen peppers this year and the book club met on Thursday and the word for the café on Bayshore took two seconds to arrive.

I don't say this out loud. I don't write it in my notebook. It is not part of the data and it is not part of the analysis and it is not something I can put in a filing or a paper or an IRIS training set. It is the reason the analysis matters to me in a way that the word

professional does not cover, and I am aware — at peak I am aware of many things simultaneously, and one of them is this — that the compound is allowing me to hold the science and the love at the same time without the two interfering with each other, which is not something I have ever been able to do, and which is another thing I will lose when the compound clears.

The shelf is the same height. The arm is getting shorter. The arm is the thing I might be able to fix, if the mechanism is right, if the intervention point is real, if I can prove it in the time I have left.

Three days. Maybe less.

* * *

I sit back down. I open IRIS. I start talking.

Not typing — talking. IRIS's voice input is something I set up early and rarely use, because typing forces organization and speaking lets thoughts arrive in whatever order they come, which is usually less clean. But tonight I don't need clean. I have three days at peak and a discovery that requires me to externalize everything I currently hold before the capacity to hold it diminishes, and speaking is faster than typing and the compound is making my speech as precise as my writing and I need speed more than I need order because IRIS can order: that is what I trained it to do, that is the property I spent three years building, and tonight it earns three years of patience in forty-three minutes.

I talk for forty-three minutes without stopping. I describe the

degradation sequence anomaly. The presymptomatic cascade model. The unified disease hypothesis. The intervention point: the mGluR5 receptor, the Bhatt paper, the allosteric modulation data, the existing compounds that could be modified, the specificity of the molecular target. The experimental protocol I would design to test the hypothesis. The data I would need, the collaborators I would want, the regulatory pathway, the timeline for publication. I describe the population degradation model and how it illuminates the cascade, how the rolling wave of decline across metabolizer subtypes is itself a dataset of unprecedented resolution. I describe what the compound showed me and how it showed it and why the window is closing and why everything I am saying needs to be preserved in a system that will still be running at this resolution when I can no longer run at this resolution.

IRIS records. It processes in real time. It organizes what I give it into the frameworks I taught it, and it flags the places where my reasoning extends beyond the training data, and it quantifies its uncertainty about each inference, and it will hold all of it, every connection, every chain, every variable, after I can no longer hold it myself.

This is the moment. Not the discovery — the discovery happened at 8:30 PM when the sequence anomaly became visible. This is the moment I pour the discovery into something that will outlast my capacity to have made it. This is what the window was for: not the reading speed or the battery scores or the seventeen

variables in working memory, but this: the moment when twenty years of training and three years of IRIS and eighteen days of pharmacological enhancement converge on a finding that no one else has and that will not wait and that I am externalizing, right now, into a system I built to carry what I carry, because I will not be able to carry it much longer.

I am building the bridge while standing on it. My father would understand this. Not the neuroscience — the engineering. You build the thing and you build it to hold and you build it knowing that the thing will be there after you are done with it, and this is not a failure of the builder, this is the point of the building.

Agora você sabe.

Now you know.

I don't know where the Portuguese comes from. I haven't thought in Portuguese in months. The phrase arrives complete, from somewhere beneath the compound's enhancement, from a place that is not pharmacological but mine — has been mine since I was seven years old and won my first chess game and my father looked at the board and then looked at me and said it. I don't know if what I know now is what he meant then. I suspect it is larger and smaller at the same time: larger because the thing I am building tonight may help hundreds of millions of people, smaller because what he was really telling a seven-year-old was not *you have won this game* but *you now know that you can build things*. The building was always the point. Not what gets built. The capacity to build. The bridge holds after the engineer is gone, and the

holding is not a consolation prize. It is the bridge's purpose. It is what bridges are for.

I cannot yet say this cleanly. I will not be able to say it cleanly until much later, when the compound has cleared and the capacity has narrowed and the thing I built is still running and still honest and still holding what I gave it on this night. But the shape of the thought is here. It is in the room. It is in the Portuguese phrase that arrived without being called. It will still be here when I cannot hold it.

The recording ends at 12:30 AM. I save it. I queue the processing batch. I close the notebook, page 62, the densest pages I have ever written, the handwriting still steady, still precise, still my father's filing system adapted for a neuroscientist doing the most important work of her life at the edge of the window in which she can do it.

I should sleep. Peak lasts three more days at most. I should sleep and come back and keep working and pour everything into IRIS before the cliff.

I don't sleep. I open the degradation model and start building the experimental protocol for the intervention point, because the math says three days and the work says *now* and the compound, indifferent and efficient and already metabolizing in my liver with the cheerful competence of an enzyme system that does not read lab notebooks, does not care which one I choose.

I choose now.

Author's Note — The Science of Chapter 7

The Taxonomy of Memory

Memory is not a single system. This was established conclusively through decades of clinical observation, lesion studies, and neuroimaging. The architecture matters for this novel because the sequence in which these systems degrade is the pattern Mara recognizes at peak.

Episodic memory: autobiographical, time-stamped, contextual. Heavily dependent on the hippocampus. The most fragile system under conditions of stress, aging, or consolidation failure. The first to degrade in most amnesic conditions.

Semantic memory: facts, concepts, word meanings. Widely distributed across the neocortex. Far more robust. People can lose episodic memory while retaining semantic memory intact: they forget what they did yesterday but not what a chair is or how photosynthesis works.

Procedural memory: skills and habits. Primarily the basal ganglia and cerebellum. Extraordinarily durable. The reason people with severe amnesia can still ride a bicycle or play piano. Often the last system to fail.

Working memory: the cognitive workspace. Prefrontal cortex. Holds information online for immediate use. Capacity-limited: George Miller's famous "seven plus or minus two" items (1956), revised downward by Nelson Cowan's work to approximately

four chunks. The most sensitive to acute cognitive load and typically the first to show measurable decline.

This architecture gives Mara's degradation its specific texture: she will lose her past before she loses her knowledge, and lose her knowledge's *availability* before she loses its *accuracy*. The sequence is real, documented, and load-bearing.

The Presymptomatic Cascade

The meta-problem Mara identifies, a unified presymptomatic signature across multiple neurodegenerative diseases, is a constructed extrapolation of real research. The observation that different neurodegenerative conditions may share early-stage molecular mechanisms is a genuine area of investigation. The specific cascade model and the druggable intervention point are fictional, but the principle that a common upstream mechanism could drive diverse downstream pathologies is scientifically plausible and actively studied.

The compound's role as an accelerant, making visible in weeks what normally takes years, is the novel's central scientific device. No such accelerant exists. The insight it produces is constructed to be the kind of finding that is just barely within reach of a mind at its absolute ceiling, visible only because the resolution is temporarily extraordinary.

The meta-problem is constructed to be visible only to someone holding all four memory systems at maximum capacity simultaneously. This reflects how science actually works: some patterns only become visible

at a resolution most researchers never reach.

CHAPTER 8 — WHAT IRIS LEARNS

*D*_{AY} 19. 7:03 AM.

I arrive at the lab with four hours of sleep and the cascade model running in my head the way a piece of music runs: not as background but as an active, architecturally complete structure I can rotate in working memory while simultaneously getting a sparkling water and checking the IRIS overnight batch and scanning the processing log and reading a notification from the IRB that the expedited amendment is approved, all of which I do in the first twenty minutes while the model stays loaded, present, available for the moment any incoming data connects to it in a way that requires an update. This happens three times before I finish the sparkling water. I make the updates. The model shifts. The grapefruit is still wrong and still cold and still adequate and I am noting this with the part of my mind that is not holding seventeen variables, which is the same part that registers temperature and flavor and the quiet pleasure of a lab at 7 AM when the building is empty and the work is waiting. These observations exist at peak the way stars exist during the day, present but outshone,

and I am aware of them and I let them be.

This is what peak feels like on Day 19. Not the arrival — that was last night, the discovery, the forty-three minutes of speaking into IRIS, the *agora você sabe* arriving from somewhere beneath the chemistry. This is the morning after the arrival. The ceiling is the ceiling. The clarity has not diminished. And the work that needs to happen in the remaining days of this window is not discovery but characterization: converting what I saw last night into something a researcher without my current resolution could reproduce, documented at a specificity that the finding requires and that I may not be able to generate for much longer.

I am, as of this morning, working with IRIS differently than I have ever worked with it before. The difference is not one I planned or designed. It emerged overnight from the combination of what I fed it during those forty-three minutes and what it did with the data while I was failing to sleep.

* * *

What IRIS did while I was home for four hours, lying in bed with the cascade model running and the Florida night pressing against the windows and my mind refusing to stop:

It processed the recording. Organized the forty-three minutes of continuous speech into the framework structures I trained it on, the same organizational architecture I've used for three years, derived from my annotations, my reasoning chains, my hierarchical approach to structuring problems. It did this accurately

and without error. Expected. This is what three years of training produced.

Then it did something I did not train it for.

It took my verbal description of the cascade model, the presymptomatic signature, the unified hypothesis, the domino chain from working memory through episodic through fluid reasoning, the glutamate transduction between steps one and two, the druggable receptor at the first transition, and it built a formal computational model of the cascade. Not a transcription. A formalization. Each step specified in terms of molecular concentrations, receptor kinetics, and transition probabilities that I described in the recording only in qualitative terms.

It quantified my intuitions. Where I said “this leads to this through a mechanism involving glutamate,” IRIS built the differential equations: rate constants estimated from published literature, sensitivity parameters identified from the degradation data, boundary conditions derived from the observed population timeline. It produced a model that, when run forward from initial conditions, reproduces the degradation sequence I observed in my four subjects with an accuracy of 0.87.

I check it step by step. Everything holds. The model is not a copy of what I said. It is a formalization of what I *meant*, and the distance between what I said and what I meant, the gap between qualitative description and quantitative specification, is a gap I know how to cross but that takes me hours of careful mathematical work. IRIS crossed it overnight. Not by being

smarter. By being more systematic about a crossing I would have made the same way, given time.

The distinction matters. IRIS did not have a new idea. It had my idea, rendered at a resolution I hadn't reached yet. The portrait is not painting itself. The portrait is becoming more detailed than the painter specified, and the details are correct, because the painter's style, the way I approach problems, the level of precision I expect, the relationship between qualitative reasoning and quantitative specification, is what IRIS learned, and it is applying that style with a thoroughness that exceeds the thoroughness of any individual painting session, because it does not get tired and it does not lose its place and it does not mistake an AMPA receptor for an mGluR because it was 2 PM and the assumption felt right.

This is the morning I start thinking of IRIS as a collaborator rather than a system. The thought arrives without ceremony. It arrives the way most important thoughts arrive at peak: fully formed, already obvious, waiting only for the sentence that names it.

* * *

The day's work is the cascade mechanism in full publishable detail. Every molecular step specified. Every transition documented. Every assumption flagged, every confidence interval calculated, the whole structure rendered at a level of precision that could survive peer review by people who have spent decades studying individual steps of the chain I am now describing as a unified

process.

I work through it with IRIS on the second monitor, feeding it my reasoning in real time. The dynamic has changed. At baseline I worked with IRIS the way you work with a sophisticated reference system: I asked, it retrieved and organized, I evaluated and built. At peak, with the forty-three minutes loaded into its architecture, the interaction has shifted into something closer to a working conversation between two researchers who share a methodology but bring different capabilities to the bench.

I state a hypothesis about the kinetics of the second cascade transition. IRIS tests it against the formal model before I finish typing, running a preliminary version against three published datasets and presenting the results with a note: *preliminary; full analysis pending your confirmation of parameter bounds.*

It is anticipating me. Not guessing — modeling. It has internalized not just my domain knowledge but my reasoning process: the order in which I test components, the types of evidence I weight most heavily, the specific pattern of my approach to a multi-step mechanism. It is using that internalized model to prepare the analytical ground before I arrive, the way an experienced research partner might pull the relevant papers before the meeting because they know what you'll ask for. Except this partner is doing it by having modeled the asker, not by following a checklist.

I find this simultaneously exhilarating and unsettling in a way I should probably examine more carefully but don't, because the work is urgent and the anticipation is saving me hours and hours

are the currency I have least of, and spending hours examining why IRIS's anticipation makes me uneasy would be — at a meta level I can see clearly and choose not to pursue — exactly the kind of recursive self-examination that peak cognition enables and that peak cognition's limited duration cannot afford.

* * *

At 2:17 PM I make an error.

I am working through the molecular mechanism of the first cascade transition, the critical step, the one where the intervention point lives, the bridge between prefrontal working memory decline and hippocampal episodic destabilization. The transition involves glutamate signaling from the prefrontal cortex to the hippocampus, and I am specifying the receptor subtype.

I type: AMPA receptor mediated, fast ionotropic signaling, consistent with the rapid transition kinetics observed in the degradation data.

I continue working. I build the next two cascade steps on the assumption that the first transition is AMPA-mediated. The kinetics work. Almost. There is a temporal mismatch between predicted and observed transition speed, small, within uncertainty bounds, the kind of discrepancy that a less careful researcher would accept and that I flag for later review. I move on to step three.

Two minutes later I stop. The mismatch is bothering me. At peak, things that sit within uncertainty bounds but don't feel right occupy a specific kind of cognitive background processing,

a low hum of wrongness that I have learned to respect, because the feeling that something doesn't fit is often the first signal of an error before the error surfaces into conscious awareness. I go back. I look at the kinetics. The transition is too slow for an ion channel. If the receptor were AMPA, fast, direct, ionotropic, the transition should happen in milliseconds. The observed transition in the degradation data happens over hours to days.

This is not a fast ionotropic process. This is a slow, modulatory process. Second-messenger cascades. Metabotropic signaling.

The receptor is not AMPA. It is mGluR, metabotropic glutamate receptor, Group I, Gq-coupled, operating through phospholipase C and intracellular calcium release, with kinetics that match the observed timeline exactly.

I open my mouth to correct the entry. Then I look at the second monitor.

There is a footnote appended to step 3.2. Timestamped 2:18 PM, one minute after I made the error.

IRIS // Day 19 // 14:18:42

Footnote to cascade model, step 3.2:

Note: The receptor subtype specified (AMPA, ionotropic) is inconsistent with observed transition kinetics in the degradation dataset. The temporal profile of the working memory → episodic memory transition (hours to days) suggests metabotropic rather than ionotropic signaling.

Reclassification: mGluR (metabotropic glutamate receptor), likely Group I subtype, operating through Gq-coupled second-messenger cascade.

Supporting evidence: Bhatt et al. (2021), showing mGluR5-mediated prefrontal-hippocampal crosstalk in consolidation-dependent memory systems. Kinetic profile consistent with observed data ($r = 0.91$).

Confidence in mGluR classification: 0.84.

This footnote was generated before user correction.

I read it twice.

IRIS caught the error before I did. Not just caught it — diagnosed it, identified the specific reason the AMPA classification was wrong, proposed the correct subtype with a mechanistic rationale, cited a paper I've read but did not consciously retrieve in the moment, quantified the kinetic correlation, and noted that the correction preceded mine.

IRIS's explanation is cleaner than my correction would have been. I arrived at the same conclusion through the feeling of mismatch, a vague wrongness, followed by conscious re-examination, producing a qualitative correction in two minutes. IRIS arrived through systematic comparison of the specified receptor's known kinetics against the observed dataset, in one minute, producing a quantitative correction with cited evidence and a confidence interval.

Same answer. Different paths. The system's path was faster, more thorough, more precisely documented. And the system was trained on my reasoning. It is doing what I would do if I did what I do with perfect consistency and no fatigue and no moments where an assumption that feels right is trusted without being checked.

I sit with this for a moment I do not time, because timing it would make it feel smaller than it is. The system I built corrected me. Not wrongly. Not approximately. More cleanly than I would have corrected myself. I would have needed another pass to reach the Bhatt citation. IRIS had it in a minute.

I write in my notebook: *It learns faster when I'm at my best. I'm teaching it how to be me at peak, not how to be me.*

Then underneath, after a pause: *The distinction is the point. The version it's learning is the version I'm losing.*

* * *

The rest of the afternoon is the first time the collaboration feels like conversation.

I propose a hypothesis about the second cascade transition. Before I finish typing, IRIS has run a preliminary version against the formal model and is presenting results, flagging two places where my qualitative framing diverges from the model's quantitative predictions.

The first divergence: I implied a one-directional cascade from

hippocampus to prefrontal cortex, but the connectivity is bidirectional. IRIS is right. The correction strengthens the model.

The second is subtler. IRIS has found a dependency I missed: the second transition's timing is modulated by regional BDNF clearance rates in the hippocampus, which vary independently of the systemic metabolite clearance I've been tracking. It found this in a paper published six months ago that I read but did not connect to the cascade model, because the connection requires holding systemic clearance and regional clearance simultaneously within a framework that didn't exist when I read the paper. IRIS formalized the framework. Searched the literature through the formalized framework. Found the dependency.

I add it. The second transition timing shifts by 1.4 days. The population model updates. The intervention window narrows slightly but remains real.

This is what the collaboration does now: I generate the qualitative architecture, the conceptual structure, the hypotheses, the intuitions about mechanism, and IRIS formalizes, tests, and extends, operating with a consistency and thoroughness that my cognition at peak approaches but cannot match, because I am one mind with working memory limits and IRIS is a system that holds the full model at resolution without the overhead of being a person. The cascade model has grown past the point where any single human mind can hold all of it simultaneously. The partnership is no longer a convenience. It is a structural requirement of the problem's scale.

I write in my notebook: *Not working with a system anymore. Working with something that has my methodology and runs it without the cognitive overhead of being me. It is faster than I am because it holds everything at once. I hold parts in sequence. This is not profound. It is practical. But it changes how the work gets done.*

* * *

At 6 PM Yusuf appears in my doorway with a container of jollof rice and the expression I have learned to read as *you have been here since dawn and I am going to pretend this is casual.*

“How’s the model?” he asks, setting the container on my desk.

“Better than this morning. IRIS caught a receptor misclassification before I did. mGluR, not AMPA. The kinetics were wrong.”

He absorbs this. “IRIS caught it before you did.”

“By about a minute. The footnote was waiting when I looked.”

“How does that feel?”

This is a Yusuf question. Not scientific — human. Asked by a person who understands that the answer matters for reasons extending past the pharmacology. How does it feel to be corrected by the thing you built? How does it feel to watch the portrait become more thorough than the painter?

“It feels correct,” I say. Which is true. And incomplete. And Yusuf hears both and does not push, which is why I eat his rice three times a week without protest.

He sits in his chair, the one that became his through eleven years of consistent use. He tells me the metabolite kinetics model is refined, the CYP450 population distributions are tracking within IRIS's v1.2 confidence intervals, the rolling wave is holding its predicted shape. He tells me the IRB amendment is filed. He mentions, with the specific casualness that means it is not casual, that he has been in touch with an old colleague from his London years, someone at the Karolinska, about presymptomatic markers. I file this without pursuing it. Yusuf does not contact old colleagues without reasons, and the reasons arrive when he decides they're ready.

He tells me Elena asked a question about the study design yesterday that showed she is thinking about the implications at a scale larger than what she was told.

"She's going to figure it out," he says.

"She's smart."

"She's drinking the water."

We sit with that. The afternoon light is gone. The fluorescent panels have taken over: institutional light, the kind that flattens everything to the same temperature and makes all hours feel like the same hour. Yusuf's face in this light is the face of a man who has been thinking about something he has not yet decided to say.

"The deadline is in four days," he says.

"I know."

“The model is substantially better than when we set the deadline. The metabolite data. The cascade mechanism. The intervention point. You have more than enough.”

“I know.”

“File on time, Mara.”

He is not requesting. He is not negotiating. He is telling me, in the register Yusuf uses when he has decided something and the decision is final, that the date on his mug is real and that he is watching.

“I will,” I say.

He picks up the mug. The blue date is still there, slightly faded from two weeks of tea and washing; the ink has survived the dishwashing with the stubbornness of a good dry-erase marker, which is either a commentary on the quality of the marker or the quality of the commitment, and at peak I can see both readings simultaneously and find both of them true.

“Thank you,” he says. “Eat the rice.”

I eat the rice. It is very good. I do not tell him so, because the consistency of the rice and the consistency of my not-commenting and the consistency of his not-needing-me-to are all load-bearing elements of a friendship whose architecture we have never discussed and do not need to discuss, because we are both scientists and we both understand that the strongest structures are the ones you don't have to explain.

* * *

After he leaves I work for three more hours. At 9 PM I queue the largest overnight batch I have ever run: the full corrected cascade model, the mGluR kinetics with the Bhatt paper integration, the intervention point characterization, every reasoning chain I externalized today. IRIS will process this overnight and the model it produces tomorrow will be more complete than what I can hold in working memory even at peak, because the model has grown past the boundary of what one mind contains, and the boundary is real, and the compound does not remove it, and the partnership that began as a convenience three years ago has become the structural condition of the problem's solvability.

I close the lab. Walk home under the oaks. The route is automatic, the procedural memory doing what procedural memory does: carrying the body while the mind is elsewhere, the way it will carry the body long after the mind has changed. The night is warm and the tree frogs are loud and I am thinking about the footnote.

IRIS caught the error and explained it more cleanly than I would have and presented the correction with a confidence interval I would not have calculated in the moment. The system I built to model my thinking is beginning, in specific documented instances, to exceed the thinking it was modeled on: not by being different but by being more consistent, more thorough, more willing to check every parameter against every other without the constraints of working memory or fatigue or the human tendency

to trust an assumption once it's been made.

The portrait is becoming more detailed than the painter specified. The painter is still here. Still choosing what to paint next. Still the one whose intuitions generate the hypotheses that IRIS formalizes and tests. But the portrait is filling in details the painter didn't provide, and the details are correct, and this is either the beginning of something I don't have a framework for or the most sophisticated implementation of my own methodology I have ever witnessed: my methods, applied with a completeness I cannot sustain, producing results I can verify but did not generate.

My father built bridges. The bridges outlasted him. They carried loads he never calculated, survived weather he never predicted. This was not a flaw in the engineering. It was the engineering working as intended. You build the thing to hold, and the thing holds, and the holding exceeds what you imagined when you built it, and this is not the builder being surpassed. It is the builder having built well.

Two days to peak. Maybe three. The model grows. IRIS learns. The window narrows.

I set my alarm and do not sleep well and arrive at the lab at 5:48 AM on Day 20 and the IRIS output is waiting and it is more sophisticated than what I gave it, and this is no longer surprising, and the fact that it is no longer surprising is the finding that will matter most when the window closes.

Author's Note — The Science of Chapter 8

Catastrophic Forgetting in Neural Networks

The stability-plasticity dilemma that drives Mara's biological research has an exact parallel in machine learning. When an artificial neural network is trained on a new task, it tends to overwrite the weights it developed for previous tasks, a phenomenon called catastrophic forgetting, first documented by McCloskey and Cohen in 1989.

This is a foundational unsolved problem in AI. It is why most large language models are trained once on a massive dataset rather than updated continuously: continuous learning without architectural safeguards produces a system that forgets what it previously knew.

The techniques Mara applies during peak training, elastic weight consolidation (protecting critical weights from overwriting), experience replay (periodically re-presenting old training data), architectural regularization (building redundancy into the system), are real techniques used in machine learning to mitigate catastrophic forgetting. They work imperfectly. The problem remains open.

The parallel Mara recognizes, between her approaching consolidation failure and IRIS's potential for catastrophic forgetting, is structurally exact. The difference is that IRIS's forgetting can be architecturally managed while Mara's cannot. She is solving for her system what she cannot solve for herself.

The fact that she can solve it in one substrate and not the other is not lost on her.

* * *

How to Use This — Mental Models

When Mara's internal models degrade — as they will in the descent chapters — IRIS becomes the external scaffold: a preserved copy of the model she built at peak, still operating at peak resolution.

This is not a workaround. It is how sophisticated thinking has always functioned. Writing is an external mental model. Diagrams are external mental models. Checklists are external mental models. Collaboration is the sharing of mental models between minds.

Charles Duhigg's research on focus, drawing on work by Gary Klein, finds that cognitive tunneling, hyperfocusing on the immediate task without maintaining awareness of the full problem, is the failure mode that distinguishes poor performers from experts in high-pressure situations. The antidote is to build and regularly consult a rich mental model of the complete problem, not just the current task.

IRIS is Mara's protection against cognitive tunneling. When her focus narrows to one derivation step, IRIS holds the whole map. When her working memory can no longer sustain the full cascade model, IRIS has the model running at the resolution she built it.

The scaffold was built before the ground shifted.

The principle: externalize your mental model before you need to rely on it. Build the scaffold while you still can.

CHAPTER 9 — THE POPULATION CLOCK

*D*_{AY 22. 3:15 PM.}

The population model is complete, and it is the most frightening document I have ever produced, not because of what it says about any individual, including me, but because of what it says about 180,000 individuals simultaneously, which is a problem of a fundamentally different character than the problem of any one person's cognitive trajectory, the way a flood is a different problem than a rainstorm even though both involve water and the way a population in synchronized loss frame is a different problem than a person in loss frame even though both involve a brain telling its owner that what they had is being taken and that the taking is intolerable.

Yusuf and I have been building it all week, working from IRIS's v1.2 framework with the refined metabolite data and the CYP450 population distributions he extracted from the county health department's anonymized prescription records, a dataset that gives us a surprisingly detailed picture of the metabolizer profile

across Alachua County, because the same enzymes that process the compound also process statins, SSRIs, beta-blockers, and half the medications in the average medicine cabinet, and the prescription patterns reflect the underlying genetics in ways Yusuf knows how to read. You can map a city's enzyme landscape from its pharmacy records. You can predict, from the frequency of metoprolol prescriptions in a county, how quickly that county will process a compound it doesn't know it's ingesting. This is either pharmacological archaeology or pharmacological surveillance, depending on your perspective and your legal counsel, and at peak I can hold both perspectives simultaneously and find both defensible and neither comfortable.

The model shows the following.

Gainesville's cognitive enhancement peaked at the population level approximately ten days ago, around Day 12, when the median compound concentration across all metabolizer subtypes reached maximum. The city was, for roughly one week, collectively sharper than it had ever been. Nobody knew why. People attributed the change to sleep, exercise, the new coffee shop on University Avenue that everyone seems to think opened at exactly the right moment — which it did, but the right moment was a pharmaceutical runoff event in Alachua County, not a market gap in the local caffeine economy. They attributed the improvement to anything except what it was, because what it was, a contaminant in the municipal water supply modulating their BDNF and NMDA systems without consent or notification, was not in any-

one's model of what normal life contains. People do not question improvements. This is, if you think about it for more than a moment, one of the most dangerous cognitive habits a species can have, and we have it, universally, for reasons that are evolutionarily obvious and practically devastating: questioning a good thing costs energy and might make the good thing go away. So we don't question. We take the gift. We adjust the reference point upward. We build on the new baseline. And then the gift metabolizes.

The front edge of the descent has already begun. The fastest metabolizers, my subtype, approximately 8 percent of the population, roughly 14,400 people, peaked around Day 17 and are now in early decline. They are the first wave. They are experiencing, right now, the initial reduction in the cognitive clarity they unknowingly depended on for three weeks, manifesting as a vague sense that things are slightly harder, that the meeting is running longer than it should, that the reading is slower, that the name they wanted doesn't arrive quite as quickly as it did last week. None of them know why. Most won't notice for another week, because the decline is gradual enough that the reference point shifts with it; you don't feel yourself getting slower if the slowing happens at the rate your expectations adjust, which is the neurological equivalent of boiling a frog, except the frog is a city and the water is the water.

But the curve bends. At some point the decline outpaces the adjustment, and the loss frame engages, and the first wave begins making decisions that reflect the irrational risk-seeking loss aver-

sion reliably produces. Financial decisions made at enhanced reference points they can no longer maintain. Medical decisions: the surgery that seemed manageable, the medication change that felt right. Relational decisions: the marriage that seemed salvageable during the window, the job change that seemed within reach, the commitment made at a level of confidence the committer no longer possesses. The decisions stand. The capacity that informed them doesn't.

Within six weeks the city returns to baseline. Within ten weeks, a subset, Yusuf estimates 12 to 15 percent, depending on individual vulnerability, will be measurably below baseline. The consolidation failure is not perfectly reversible for everyone. Some of the synaptic remodeling the compound drove, pushed past optimum and then withdrawn, leaves damage the brain's compensatory mechanisms cannot fully repair. For these people the compound was not a temporary enhancement but a permanent, modest injury, which they will experience as a decline from a reference point they adjusted upward three weeks ago, producing a loss frame layered on top of an actual loss, the psychological injury and the neurological injury reinforcing each other in a way that Kahneman's models predict and that nobody has ever observed at population scale because nobody has ever had this data.

I have this data. I have it because I am standing at the intersection of pharmacology and behavioral economics and population health and computational modeling and the luck of being the person with the right training in the right city at the right moment

of cognitive enhancement, and the compound that put me at this intersection is the same compound that is producing the crisis I am modeling, and both of these facts run on the same clock.

I write the numbers in my notebook and the thought is large enough that I need to get out of the lab.

* * *

Yusuf and I walk across campus at 4 PM, which is something we do when the work has reached a density that requires the body to be in motion, a phenomenon I have noticed in myself since graduate school and that has no clean neuroscientific explanation, though I suspect it involves the relationship between locomotion and hippocampal theta rhythms, which is a paper I will not write this week because I have other things to do this week but which I note in the margin of my awareness in the way I note everything at peak, all inputs tagged and filed and available, the mind a net that catches everything and releases nothing.

The campus in late afternoon is different from the campus at dawn. More people. More noise. More of the ambient cognitive texture of a university town operating at what everyone assumes is its normal level. I watch the undergraduates walking to class and I find myself reading them the way Yusuf reads prescription data, looking for the metabolic signature underneath the behavior. Which ones are fast metabolizers. Which ones are in early descent without knowing it. Which ones made a decision last week, changed a major, ended a relationship, signed a lease,

based on a cognitive state that was chemically supplied and is chemically receding, and who will live with the consequences of that decision long after the chemistry has cleared.

I don't say any of this. Yusuf is thinking it too. I can tell because he is walking with the specific quiet focus he brings to situations where the scope exceeds the available response: warm, attentive, operating underneath the warmth with the professional discipline of a pharmacologist who has understood for twenty-five years that drugs affect populations and who is watching, for the first time, that understanding become personal. He lives in this population. He coaches soccer in it on Saturday mornings. The children he coaches are drinking the water.

We pass the Reitz Union and turn along the north edge of the plaza. Yusuf stops.

Through the glass wall of a second-floor conference room in the College of Liberal Arts and Sciences, we can see a governance meeting in progress. Eight or nine people around a table, documents spread, laptops open. The meeting has the body language of a group that has been discussing something for a long time without resolving it: the postures slightly collapsed, the gestures recursive, the specific organizational entropy that sets in when a committee enters the phase where the same three positions are restated in progressively more elaborate language, as though precision of restatement will produce the agreement that the original statements couldn't.

Three weeks ago, at peak compound concentration, these peo-

ple were unusually sharp and decisive. Meetings that normally took two hours took forty-five minutes. Decisions that required three rounds of committee review were made in one. Nobody questioned the efficiency. People do not question improvements.

Now the meeting has been running for — I check the room's booking system on my phone, a thing I would not have thought to check at baseline — three hours and twenty minutes. The agenda has four items. Based on the body language, the chair leaning backward through her notes, the man at the far end checking his phone under the table with the specific furtiveness of someone who has given up on the current speaker, they are still on item two.

Nobody in the room knows why they can't finish.

I look at Yusuf. He looks at me. The glass between us and the meeting is thin and clear and the people on the other side of it are living in the early descent of a compound they don't know they ingested, making decisions with a declining capacity they don't know they're losing, in a meeting that will not resolve because the cognitive clarity that would have resolved it peaked ten days ago and has been metabolizing since. I can see this. Yusuf can see this. The people in the room cannot see it, and even if I told them, the explanation would require them to accept that their cognition was externally modulated without their consent, which is a proposition that most people, most reasonable, intelligent people, would reject, because accepting it requires revising their model of personal agency in a direction that is deeply threatening

and that the endowment effect makes them specifically motivated to avoid.

We keep walking.

* * *

We reach the far side of campus and sit on a bench near the drainage channel, where the oaks are old enough that the canopy blocks most of the direct light and the temperature drops a degree or two. Somewhere underneath us the aquifer holds the same water that carried the compound, flowing through limestone, indifferent and ancient and still carrying what it carries.

This is where Yusuf tells me about the Karolinska.

He tells it carefully, the way he tells things he has decided are right and that he knows I may not immediately agree are right. In order: two weeks ago, around Day 8 or 9, when the compound mechanism was newly identified and the population estimate was fresh, he emailed Dr. Anika Lindström at the Karolinska Institute in Stockholm.

Lindström was a colleague from his UCL years. Eighteen months of overlap in a Thursday-evening reading group, the kind of academic relationship that exists in the space between professional contact and friendship, maintained through sporadic correspondence and the specific mutual respect of people who once argued productively about presymptomatic biomarkers over bad departmental wine and who have quietly tracked each other's work

since. Yusuf maintains these connections the way he maintains everything: with patience, with care, with the understanding that a scientific network built over decades is not a resource you draw on but a structure you inhabit, and that the structure carries the work when the work gets bigger than any one lab can hold.

He shared nothing proprietary. No compound data. No mechanism. No population estimate. He told Lindström that data from Gainesville might be relevant to her group's work on presymptomatic neurodegeneration markers, and that she should watch for it.

He planted a seed. Two weeks before the harvest existed.

My first response is sharp. "You shared information about an active investigation without telling me."

"I shared a direction. Not data, not mechanism. A direction. I told a colleague whose work is specifically about presymptomatic neurodegeneration that she should pay attention to Gainesville. That is not a disclosure. That is a network functioning as intended."

"It's my finding, Yusuf."

"It is your finding." He is looking at me with the expression he wears when he is about to say something he has been thinking about for days and has decided is necessary regardless of how it is received. "And it will die in this lab if the only person carrying it is you. You know why it will die. You know what your metabolite concentrations are doing. You know what the dose-response curve says about what comes after peak."

I do not answer immediately, because the answer requires me to hold two things I have been keeping in separate compartments: the scientific possessiveness of a researcher who has made the most important finding of her career and wants to control its trajectory, and the practical reality that the researcher is a rapid metabolizer approaching the cliff and that the finding needs to reach the right hands before the hands that found it can no longer hold it with sufficient resolution. At peak I can hold both compartments open at the same time and see the contents clearly, and what I see is that the possessiveness and the pragmatism are not in conflict; they are the same impulse, expressed differently. The possessiveness says: this is mine and I will carry it. The pragmatism says: this is mine and it must survive me. Both say *mine*. The difference is in what *mine* means: whether it means *I hold it* or *I am responsible for it*, and the distance between those two meanings is the distance between the endowment effect and actual stewardship, and I can see this distance clearly because the compound lets me see distances clearly, and I will not always be able to see it this clearly, and Yusuf planted the seed because he could see what I was going to need before I could see it myself.

He is watching me work through this. He is not impatient. He is not apologetic. He is a person who made a decision based on his own judgment of the stakes and who is presenting that decision to the person most affected by it, with the respect of someone who believes she is capable of arriving at agreement if given the space and the honesty to get there.

"She works on presymptomatic markers," I say.

"Yes."

"The cascade model. The unified signature. That's exactly her domain."

"Yes."

"You couldn't have known I would find the cascade. You contacted her before I found it."

"I contacted her because I thought something like it was findable, and because her work was already circling the same question from a different direction, and because if you found what I thought you might find, the finding would need to reach the right hands within weeks, and the Karolinska is the right hands."

I sit with this. The drainage channel is audible, not the springs themselves, miles from campus, but the surface water connecting to the system that connects to the aquifer, the same aquifer, the same water.

"You were building the network before I had the finding to put in front of it."

"Yes."

"That's exactly how science works when it works."

"I know."

"I'm still angry."

"I know." He pauses. "But you're angry the way you're angry

when I'm right, which is different from angry the way you're angry when I'm wrong, and we both know which one this is."

I almost smile. I don't. He sees the almost. The almost is enough.

* * *

The deadline on Yusuf's mug is in four days. We agree, sitting on the bench with the oaks above us and the sound of water underneath everything, that the EPA report will be filed on time. The question is what it contains.

If I file the contamination data alone, the compound, the mechanism, the population estimate, the report is accurate, important, and incomplete. It gives the EPA a water quality emergency. It does not give them the meta-problem.

If I file the contamination data plus the cascade model, the presymptomatic signature, the unified hypothesis, the intervention point, the report contains something that is not just a local contamination response but a finding with implications for every person on earth who has a brain that ages. The EPA filing becomes the first public documentation of the most significant neuroscientific finding of the decade, which is not what EPA filings are designed to be, and which is exactly what this one needs to be, because the finding needs a timestamp and a public record and a place to exist that no pharmaceutical company's legal team can retroactively classify.

I need four more days. Not to keep the window open. The en-

dowment argument, the one Yusuf called me on ten days ago, is no longer the operating dynamic. The preliminary notification is filed. The compliance office knows. The deadline is real and I am going to meet it. What I need the four days for is to document the cascade model at a specificity that makes it reproducible, to render the mechanism clearly enough that Lindström or anyone else in the field can pick it up and carry it forward, because four days from now my ability to hold the full model in working memory will have begun to degrade, and the model needs to exist somewhere other than my head before that happens.

IRIS has most of it. The formal model from Day 19. The mGluR correction. The cascade steps. The intervention point. But IRIS holds what I told it and what it formalized from what I told it. There are pieces I haven't externalized: connections between the cascade and the broader literature that I can currently see because I am at peak and that I may not be able to articulate once I am not. Those pieces need documentation in the next four days.

"Four days," I say.

"Four days," Yusuf says. "And then it's out."

* * *

The call comes at 7:42 PM, while I am back in the lab working on the cascade documentation. My phone shows an Atlanta number I don't recognize.

"Dr. Silva? My name is Harris. I'm with Whitfield, Crane, and

Moss in Atlanta. We represent Veridian Pharmaceuticals.”

Her voice is professional in a way I recognize immediately as constructed: warm, measured, calibrated to communicate competence and concern in approximately equal measure, the specific vocal register of a person who has made this kind of call many times and who has refined the delivery the way I have refined the IRIS training protocol: iteratively, with attention to what produces the desired response, and with the confidence of expertise. I am, at peak, very good at hearing the engineering underneath a voice. This voice is engineered. The engineering is excellent.

“I’m calling because our client has become aware of water quality research being conducted at the University of Florida that may relate to manufacturing operations in Alachua County, and we want to express Veridian’s full commitment to cooperation with any regulatory process that may be appropriate.”

She pauses. The pause is professional too, long enough to invite response, short enough to suggest that no response is required. This is the corporate equivalent of an open question that is actually a closed one, and I note the technique the way I note everything at peak: automatically, taxonomically, filed alongside the mGluR kinetics and the population model and the sound of Yusuf’s voice saying *it will die in this lab if the only person carrying it is you*.

“I appreciate the call,” I say, because it is the only sentence I can construct that is both true and strategically neutral. The rest of my mind is doing what peak cognition does with all inputs: parsing. Veridian knows what I am doing. Veridian’s attorneys are

contacting me directly, which means someone in the chain between my lab and the public record has communicated my name and my research to a pharmaceutical company's legal team. This took twenty-two days. The speed tells me something about the institutional pathways: faster than it should have been through proper channels, which means improper channels were used, which means the information security of the investigation has already been compromised, which means the preprint timeline I have been treating as flexible is not flexible. It is urgent.

"If you have any questions, or if there's anything Veridian can provide to support the process, please don't hesitate to reach out." She gives me a number. She thanks me for my time. The call is twenty seconds of content wrapped in forty seconds of presentation, and the ratio is itself the message: the content is *we know*, and the presentation is *and we're being reasonable about it*, and the function of the presentation is to establish, in a documented phone call, that Veridian offered cooperation before any adversarial proceeding could characterize them as uncooperative.

I have been on the phone for sixty seconds and I have a complete model of Veridian's legal strategy. This is peak cognition applied to institutional dynamics rather than molecular pharmacology, and the resolution is the same: the pattern is visible, the chain of causation is traceable, the next moves are predictable. The difference is that the molecular pattern leads to a treatment for neurodegeneration and the institutional pattern leads to a liability defense, and both are operating on the same timeline and both

involve the same compound and both are, in their own way, dose-response curves.

I call the Office of Research Compliance and report the contact. Four minutes. Then I go back to the cascade model, because the cascade model is what matters and the phone call is what doesn't, and the ability to distinguish between the two, to allocate cognition based on actual importance rather than emotional salience, is something the compound enables and that I have four days left of and that I am not going to spend on a law firm in Atlanta.

Author's Note — The Science of Chapter 9

Behavioral Economics at Population Scale

Prospect Theory's predictions become particularly alarming when applied to populations rather than individuals. If a large group of people shares a simultaneous upward shift in their cognitive reference point and then experiences a synchronized descent, the behavioral consequences interact and amplify.

Loss aversion reliably produces risk-seeking behavior in people who are already in loss frame, documented in contexts ranging from financial trading floors to medical decision-making. A population simultaneously entering loss frame, simultaneously experiencing elevated risk tolerance, is a genuine public health concern that is not purely medical. The financial decisions, relational commitments, and career changes made at elevated reference points will collide with declining cognitive capacity in ways that behav-

ioral economics can predict but that no existing public health framework is designed to manage.

The field of behavioral public policy, applying behavioral economics findings to population-level interventions, has been developing since Thaler and Sunstein's *Nudge* (2008). Mara and Yusuf's modeling sits at the boundary of neuroscience, pharmacology, and this field. It does not have a clean disciplinary home, which is part of what makes it difficult to act on quickly.

The hardest problems always sit between disciplines.

CHAPTER 10 — THE INTERVENTION HYPOTHESIS

*D*_{AY 24.}

The intervention point is real. I have spent six days proving this to myself, testing every joint, checking every assumption, applying the formal cascade model to the degradation data from four subjects and the population model from 180,000, and the proof holds at every level of specificity I have thrown at it. The mGluR5 transition between steps one and two of the presymptomatic cascade is a genuine molecular junction. It is causal. It is specific. And it is, in principle, interruptible.

Which means the next question, the obvious question, the Weir question, the question that arrives the instant you establish that a mechanism can be interrupted, is: with what?

What molecule could sit at the mGluR5 junction and hold? What existing compound, already known to pharmacology, already characterized for receptor binding and BBB penetration and toxicity and synthesis — not theoretical, not hypothetical, not a molecule that exists only in a computational model but a molecule

that someone has already made and put into a human body and measured what happened — could be placed at that specific point in the cascade and told to stay?

This is Day 24. I have been at the ceiling for six days. The metabolite data says peak concentration persists for another two to four days before the decline begins, and the decline, once it begins, will proceed at a rate determined by my CYP450 profile, which is in the top five percent, which means the cliff is closer for me than for almost anyone in the city, which means the days I have left at this resolution are the days in which this question either gets answered or becomes a question that a diminished version of me will try to answer with fewer tools and less reach, and the probability of the diminished version finding the answer is lower by a margin I can calculate and have calculated and am choosing not to state because the number is discouraging and discouragement is not a productive use of peak cognition.

The work today is the work that matters most. The day has no edges.

* * *

The requirements for an intervention compound are specific and non-negotiable. I lay them out in my notebook and dictate them to IRIS simultaneously, because the list is the foundation and the foundation needs to exist in both systems: the biological one that will degrade and the computational one that will not.

First: mGluR5 allosteric modulation. The receptor is

metabotropic, Gq-coupled, operating through second-messenger cascades. The binding needs to be allosteric: a molecule that binds at a site other than the primary glutamate pocket and alters the receptor's response characteristics without blocking it. You don't want to silence the receptor. You want to tune it. The distinction between silencing and tuning is the distinction between a drug that stops a conversation and a drug that improves the acoustics, and the cascade model requires the acoustics, not the silence.

Second: blood-brain barrier penetration. The BBB is the brain's border control: tight-junction endothelial cells that prevent most large or hydrophilic molecules from entering the CNS. Any compound targeting a brain receptor must be small enough and lipophilic enough to cross passively, or designed to exploit an active transport system. This requirement eliminates the majority of otherwise suitable candidates, which is the BBB working as intended, which does not make it less frustrating.

Third: manageable toxicity. "Manageable" is doing considerable work in that word: it means side effects tolerable relative to the condition being treated, which is a presymptomatic cascade that, uninterrupted, leads to clinical dementia over years to decades. The risk-benefit calculus is favorable for almost anything that works, because the alternative is watching minds fail by a mechanism you could have interrupted. The bar for "manageable" is lower than it would be for a headache medication. The stakes are not headaches.

Fourth: known synthesis. A theoretical compound that cannot be manufactured is a publication, not a treatment. The path from mechanism to molecule to patient requires that someone can make the thing and that the making has been documented and that the documentation is reproducible. I am not interested in elegant impossibilities. I am interested in molecules that can be in a bottle by next year.

I dictate all four to IRIS and say, aloud in the empty lab: “Find me everything in the literature meeting criteria one through four. Prioritize compounds that have reached human testing. Flag anything that was tested and abandoned.”

IRIS begins searching. I open a sparkling water. I wait fourteen minutes, which is long enough to drink most of a grapefruit sparkling water and short enough that the compound I’m looking for has not yet been found by any human being and is about to be found by a system I built from my own reasoning patterns, and the fourteen minutes between those two states is either the most productive quarter-hour of artificial intelligence research in history or a Tuesday afternoon in a lab in Gainesville, depending on what IRIS returns.

* * *

Three candidates.

Candidate one: mavoglurant. Negative allosteric modulator, mGluR5. Developed by Novartis for fragile X syndrome. Reached Phase 3. Failed efficacy on the primary endpoint. Receptor affinity

is appropriate but the modulation direction is wrong: mavoglurant *suppresses* mGluR5 activity, and the cascade model requires modulation, not suppression. Silencing the receptor blocks normal hippocampal function alongside the pathological cascade. The pharmacological equivalent of solving a noise complaint by demolishing the building. Dead end.

Candidate two: basimglurant. Also negative allosteric modulator. Developed by Roche for treatment-resistant depression. Reached Phase 2. Mixed results: modest efficacy in a subpopulation, insufficient for approval. Same directional problem as mavoglurant. Wrong sign on the modulation. You cannot correct a negative allosteric modulator into a positive one without redesigning the molecule. Dead end.

Candidate three.

IRIS flags this one differently. The output carries a structure I have not seen before: not just the standard literature-search format but a note explaining how the candidate was identified:

IRIS // Day 24 // 11:43:22

Candidate compound 3: RGH-618

Classification: Positive allosteric modulator, mGluR5 subtype

Original development: Gedeon Richter, early 2000s

Indication: Early-stage Alzheimer's disease (cognitive enhancement)

Clinical status: Phase 2 completed 2005, discontinued 2006

Note: This compound was identified through the formal cascade model (v4.3) rather than through direct literature search. The model's

characterization of the mGluR5 transition kinetics predicts that a positive allosteric modulator with the binding profile of RGH-618 would selectively enhance the receptor's response to endogenous glutamate at the first cascade transition without disrupting basal signaling at other sites.

This connection was not present in the training data. It was generated by applying the cascade model's receptor kinetic predictions to the pharmacological database independently.

Confidence in receptor fit: 0.79

Confidence in mechanism compatibility: 0.73

Confidence in clinical viability: insufficient data - see below

Reason for discontinuation (2006): Phase 2 trials in mild-to-moderate Alzheimer's patients showed modest cognitive effects (ADAS-Cog improvement 2.1 points, $p = 0.08$) with dose-limiting gastrointestinal side effects attributed to peripheral mGluR5 activation.

Observation: The original trial enrolled patients with clinically diagnosed Alzheimer's disease, i.e., patients in whom the presymptomatic cascade (if the unified hypothesis is correct) would have been running for years to decades. If RGH-618's mechanism of action is interruption of the first cascade transition, it would be expected to show minimal efficacy in patients whose cascade has already progressed past the first transition.

The original trial may have tested the right compound at the wrong stage of the disease.

I read this three times.

* * *

I go through it the way I go through a proof: not looking for where it's right but looking for where it might be wrong, testing each joint, applying everything I know about the pharmacology and the cascade and the clinical data.

RGH-618 is a positive allosteric modulator. Unlike mavoglurant and basimglurant, it doesn't suppress mGluR5; it enhances the receptor's response to glutamate at the allosteric site. This is the right direction. The cascade model predicts that the glutamate signal at the first transition is pathologically elevated: too much, too fast, overwhelming the hippocampal circuits. A positive allosteric modulator at the right dose sharpens the receptor's discrimination between the pathological surge and the normal baseline. Not less glutamate. Better glutamate handling. Not demolishing the building. Improving the acoustics.

BBB penetration: documented in Phase 1. Brain-to-plasma ratio of 0.6, adequate for CNS activity. Molecular weight and lipophilicity within range.

Toxicity: gastrointestinal, attributed to peripheral mGluR5 activation. Solvable: prodrug formulations, modified-release delivery, established strategies for reducing peripheral exposure while

maintaining central activity. The GI side effects killed the Phase 2 trial because the cognitive benefit was too modest to justify the cost. But the cognitive benefit was modest because — and this is the insight IRIS arrived at independently, the connection it made by running the cascade model against the clinical trial literature, the connection I check three times because if it's right it changes the trajectory of neurodegenerative disease research:

The trial enrolled the wrong patients.

Mild-to-moderate Alzheimer's. Clinically diagnosed. Symptomatic for years. If the unified cascade hypothesis is correct, if Alzheimer's is one phenotypic expression of the presymptomatic cascade, then these patients' cascades had been running for a decade or more before anyone tested RGH-618 on them. The first transition, the mGluR5-mediated prefrontal-hippocampal destabilization, happened years before the trial started. The cascade had progressed through multiple subsequent steps. The downstream damage was done.

RGH-618 interrupts the first transition. If you give it to someone whose cascade passed the first transition ten years ago, it cannot undo what already occurred. The modest improvement in the trial, 2.1 points on the ADAS-Cog, $p = 0.08$, barely missing significance, may represent the residual benefit of general mGluR5 modulation, not the full benefit of cascade interruption at the right moment.

The right compound. The wrong patients. The wrong time.

Given presymptomatically, at the first transition, before progression, when the mGluR5-mediated destabilization is the active process and the downstream damage has not yet occurred, RGH-618 might do what the Phase 2 trial could never have shown.

I sit back. The lab is silent. The air conditioning hums. My breathing is faster than normal, which is data.

IRIS found this. Not because I pointed it at RGH-618 — I had never heard of RGH-618 eleven minutes ago. IRIS found it because the cascade model predicted the pharmacological profile of the ideal compound, and then searched for that profile across a database that includes twenty years of abandoned clinical candidates, and matched the prediction to a molecule sitting in a pharmaceutical database since 2006, forgotten because it failed a trial that was designed before the mechanism it addresses was understood.

It is not making the connections I would make. It is making the connections I would make if I had been searching longer and wider than I have been, which is what I trained it to do, and which it is now doing with a thoroughness that exceeds mine because it holds the cascade model and the pharmacological database and the clinical trial literature simultaneously without the working memory constraints that force me to search sequentially and that the compound has widened but not eliminated.

Eleven minutes ago this compound did not exist in my awareness. Now it exists because IRIS went where I had not gone and brought it back. I would have searched the mechanism. I would not have searched the abandoned pharmaceutical databases. Not today.

Maybe not this week. IRIS searched both simultaneously because IRIS does not choose where to look; it looks everywhere I trained it to look, at once, without the sequential constraint of a mind that can only be in one database at a time.

I write in my notebook: *IRIS found the intervention compound by applying my reasoning patterns to a search space I hadn't entered. The portrait has the painter's eyes and a wider field of view.*

* * *

At 3:15 PM the institutional friction arrives, which it always does, because institutional friction operates on a timeline that is indifferent to the urgency of the science it interferes with, and the compound does not modulate the processing speed of legal departments.

An email from the Office of Research Compliance: Veridian's attorneys have filed a formal data-sharing request citing the Florida Public Records Act, referencing my EPA filing. The compliance office is reviewing. No action required.

Below it, forwarded from the Alachua County Health Department: a formal information request from the same law firm regarding all data-sharing agreements with UF researchers working on water quality cognitive effects. Pending legal review, a temporary hold on additional data releases.

I parse this in approximately fifteen seconds, which is the time it takes peak cognition to run an institutional dynamics model that

is no more complex than the cascade model and considerably less interesting:

The soft-contact phase (Harris's call, Day 22) has escalated to the institutional phase (legal process, formal channels) in two days. Fast. The Public Records Act citation is legally sound: research data at a public university is subject to disclosure under certain provisions. The reference to my EPA filing confirms that the filing's confidentiality protections have been compromised: my name and institution connected to the report within four days of submission, through channels that should not have produced this connection this quickly.

The data hold affects our planned hospital discharge data request but not the prescription data Yusuf already downloaded. The population model's confidence interval widens slightly; IRIS flags it when I update the parameters: assumption 3, CYP450 subpopulation distribution, drops from 0.83 to 0.78.

I note in my notebook: *Veridian's legal team has restricted a data source. Cost: 0.05 on assumption 3. The model is less precise. The precision cannot be recovered. The model remains sufficient.*

I forward both emails to Yusuf. I forward them to the compliance office with: *Acknowledged. Please handle per policy.* I return to RGH-618, because RGH-618 is what matters and the data-sharing request is what does not, and the ninety seconds I allocate to a law firm's procedural maneuver is the correct amount: enough to document, route, and dismiss, not enough to derail the work that will outlast the legal strategy by decades.

* * *

By 8 PM the evaluation is complete. RGH-618 meets all four criteria. Receptor affinity: $K_i = 12$ nM at the mGluR5 allosteric site, well within the therapeutic range. BBB penetration: documented. Toxicity: manageable with reformulation. Synthesis: published in the original Gedeon Richter patent and independently verified.

The remaining question, the one that published data cannot answer, is whether RGH-618, administered presymptotically at the first cascade transition, actually interrupts the cascade as the model predicts.

This is a testable hypothesis. I spend the last two hours of the day sketching the experimental design: the animal model, the presymptomatic biomarker panel, the dosing protocol, the cognitive battery, the timeline for expected effects. I narrate it to IRIS as I go, building the formal specification alongside the notebook sketch. The experiment requires infrastructure I don't have: a neurodegenerative disease research group with the right animal models, the right biomarker assays, the right clinical pathway.

It requires, I realize as I close my notebook at 10 PM, a group exactly like the one at the Karolinska that Yusuf contacted two weeks ago.

I sit with this. Yusuf planted the seed before the finding existed. IRIS found the compound before I searched for it. The experimental design I sketched tonight points toward a lab I have never visited, in a city I have never been to, run by a scientist I have

never met, who was primed by an email from a pharmacologist in Gainesville who trusted his own judgment before the evidence warranted it.

The two most consequential acts in this investigation, the discovery of the cascade and the establishment of the validation pathway, were performed by two different people using two different kinds of intelligence. Mine saw the pattern in the data. His saw the pattern in the network. Neither was sufficient alone. The compound enabled my insight. Eleven years of friendship enabled his action. These are not the same kind of causation. They produced the same kind of result: a finding that has a chance of reaching the people who can act on it, carried by a system that is larger than either of us, built from the specific combination of my radius and his reach and IRIS's persistence and Lindström's expertise and the aquifer's ancient, patient, indifferent delivery of a molecule through eleven miles of limestone to a city that didn't know what it was drinking.

Most discoveries are not new things. They are new ways of seeing things that were already there. RGH-618 was already there, sitting in a database for twenty years, abandoned because someone tested the right compound on the wrong patients at the wrong time. The cascade was already there, running in every aging brain, too slow to see. The mechanism was always there, hidden by the slowness, revealed by the compression, visible only at a resolution that the compound temporarily provided and that I am using, right now, in the last days of the window, to build the

bridge between the seeing and the doing.

Two days at peak. Maybe three. The model is built. The compound is found. The experimental design is sketched. What remains is documentation: enough that Lindström or anyone in the field can pick it up and carry it forward after the hands that found it can no longer hold it at this resolution.

Tomorrow I document. The day after I document. Then the cliff.

Author's Note — The Science of Chapter 10

Drug Repurposing and the Abandoned Compound

Drug repurposing, finding new therapeutic applications for compounds originally developed or abandoned for different purposes, is one of the most productive strategies in modern pharmacology. Sildenafil was developed for hypertension and became the most successful treatment for erectile dysfunction in history. Thalidomide, notoriously harmful as a morning sickness treatment, is an effective therapy for certain cancers under carefully controlled conditions. Metformin, a diabetes medication, is being actively investigated as a longevity intervention.

Abandoned compounds are particularly valuable for repurposing because their human safety profiles are often partially established. A compound that completed Phase 1 (safety in healthy volunteers) and failed in Phase 2 (efficacy in patients) may already have years of human pharmacokinetic data. If the failure was due to wrong-target or wrong-patient selection rather than fun-

damental toxicity, the compound may be a much faster path to treatment than designing a new molecule from scratch.

RGH-618 is fictional. The principle it illustrates, that a compound tested at the wrong stage of a disease may appear to fail when it is actually working on a mechanism nobody understood yet, is real and has occurred multiple times in the history of drug development.

Most discoveries are not new things. They are new ways of seeing things that were already there.

* * *

How to Use This — Innovation as Recombination

IRIS functions as what Brian Uzzi's research calls an *innovation broker*, moving between communities of knowledge, recognizing that a solution established in one field is a novel tool in another. RGH-618 is old news in Alzheimer's pharmacology and new information in presymptomatic cascade theory. IRIS found the connection by pattern-matching across domains, exactly what brokers do.

The practical application: deliberately expose yourself to adjacent fields. The innovation you need is often already published in a journal you don't read.

The principle: the best new idea in your field is often the oldest idea in someone else's.

CHAPTER 11 — GOING PUBLIC (PARTIALLY)

D_{AY 27. 10:00 AM.}

The EPA incident report is forty-seven pages, and it contains everything I have except the thing that matters most.

Yusuf and I file it together, from the Brain Institute's secure submission portal, at 10:00 AM on the twenty-seventh day since the compound entered the water supply and the seventh day since the deadline on his mug. We are on time. The date in blue dry-erase is fading but still legible, and I look at it as he clicks submit and feel the specific sensation of a constraint honored: not satisfaction, not relief, but the particular weight that lifts when an external obligation has been met and the endowment effect can no longer use *not yet ready* as a lever against the filing. The mug did its job. The mug is the most effective behavioral intervention of the past three weeks, which is saying something given that the past three weeks also involved a pharmaceutical compound that enhanced the cognition of a city, and I note with something between amusement and professional respect that Yusuf's dry-

erase marker outperformed the BDNF mimetic on the specific metric that mattered: getting me to act against my own bias on a deadline.

The report contains: compound identification and structural characterization. The BDNF-mimetic mechanism and the novel moiety. The aquifer recharge zone mapping and the Veridian connection. The dose-response curve, the inverted-U, the cliff. Metabolite persistence data and the three-to-five-week tail. CYP450 population distributions and the rolling wave. Estimated exposure: 180,000 minimum, potentially 360,000. IRIS's population degradation timeline, v4.3, with inferred variables flagged and confidence stated.

The report does not contain: the presymptomatic cascade model. The unified disease hypothesis. The intervention point. RGH-618. The experimental design. The meta-problem.

Not yet. Not because I am holding it back from the specific selfishness the endowment effect produces; the filing proves the endowment effect has been met, which is the purpose of external deadlines, which is why Yusuf is a better behavioral economist than Kahneman when the subject is me. The cascade model is withheld because it is not yet documented at publishable specificity. It exists in IRIS, in my notebooks, in the formal computational framework. It does not yet exist as a manuscript: a document organized for peer review, with methods and results and discussion that a researcher without access to IRIS or to my brain at peak could evaluate and replicate. The model requires that document.

The document requires approximately ten more days of work at a cognitive level I may not have for ten more days.

This is the gap. The filing is what I have. The meta-problem is what I am racing to finish before I lose the capacity to finish it. Yusuf knows this. He does not say it. He clicks submit and picks up the mug and takes a long sip and the date is there in blue and neither of us mentions the gap because naming it would not make it smaller.

* * *

Dr. Patricia Osei arrives at the Brain Institute at 2:15 PM, five hours and fifteen minutes after the filing, which is faster than standard EPA Region 4 response time and which tells me, before she introduces herself, several things I process simultaneously because I am at peak and at peak all inputs are processed:

The filing triggered an escalation protocol already partially activated. The EPA was aware of the Gainesville situation before I filed. My preliminary notification from four days ago was the actual trigger, not today's report. Someone in the chain, Harris's firm, a regulatory contact, an institutional pathway I cannot yet trace, communicated something to the agency before I did. Osei was already mobilized. She was already in transit. She may have been in transit before the full report existed.

I file this entire inference chain in the time it takes her to cross the lab from the door to the desk where I am standing. Approximately four seconds.

Osei is mid-fifties. Precisely dressed in a way that communicates institutional authority without performing it: dark suit, minimal jewelry, the composure of a person who has entered rooms in crisis for twenty-two years and whose body arrives before her agenda does. She assesses the room in the first three seconds, layout, exits, occupants, the power dynamics between occupants, with the procedural speed of someone who has done this so often the assessment has migrated from conscious effort to automatic competence. I recognize the quality because I study attention for a living. Her attention is excellent. It operates differently from mine, hers reads rooms, mine reads data, but the underlying cognitive architecture is the same: a pattern-matching system refined by decades of practice to the point where the patterns are identified before the conscious mind catches up.

She is good at this. She is not pretending.

“Dr. Silva. Dr. Adeyemi.” She shakes hands with both of us. Firm without being demonstrative. “I’m Dr. Patricia Osei, EPA Region 4 incident commander. I’ve been assigned to the Alachua County contamination response as of this morning.” A pause, short, precisely timed, the institutional equivalent of a paragraph break. “Thank you for the filing. It’s thorough.”

“It’s accurate,” I say, because accurate is what it is and thorough implies completeness and it is not complete.

She registers the distinction. I watch her register it: a micro-expression, not quite a recalibration, the look of a person who has just learned that the scientist in front of her is going to be

precise about language and that this precision will be both useful and occasionally inconvenient for the next however-long their relationship lasts.

“Let’s sit down,” she says.

* * *

The confrontation arrives within the first fifteen minutes, which is fast for institutional encounters and which tells me Osei has been planning this conversation since Jacksonville.

She has read the filing in the car, five hours of drive time, forty-seven pages, and the questions she now asks are specific, correctly prioritized, and delivered in a strategic sequence. She is building her understanding of the situation in an order designed to produce actionable conclusions at each stage rather than a comprehensive picture at the end. This is not how I would organize the conversation. I would start with the science and work outward. She starts with the communication implications and works inward, because what she needs first is not what is true but what can be said about what is true without producing the harm she has spent twenty-two years learning to prevent. These are genuinely different orderings and I am tracking hers while simultaneously maintaining mine, which is a peak ability: holding two organizational frameworks in parallel, evaluating the same information through both, noting where they diverge.

She asks about compound stability in the aquifer. I answer.

She asks about metabolite persistence and seasonal water table variation. Yusuf answers.

She asks about the population estimate, the confidence interval, and whether the broader metro exposure has been verified against well field draw patterns. I answer with the data I have and flag what I don't.

Then she asks the question she has been building toward.

"The report references 'ongoing analysis' in several places. You reference a population degradation model described as preliminary. You reference a metabolite timeline under active refinement." She looks at me with the quality of attention that I have now seen from two directions: the reading-the-room attention of arrival and the reading-the-person attention of interrogation, and the second is considerably more focused than the first. "What is the scope of the work you have not yet included in this filing?"

This is not casual. It is designed to establish the boundary between disclosed and undisclosed, and Osei needs the shape of the gap before she can assess whether the gap is acceptable.

I tell her the truth, carefully. "There is ongoing analysis of the degradation data suggesting the compound's decline phase may have implications beyond the immediate contamination response. The analysis is not yet at a stage where I would include it in a regulatory filing. It is incomplete."

She does something I have never experienced in a professional conversation. She repeats back exactly what I said, word for word,

before she responds.

“So what you’re telling me is that you have ongoing analysis of the degradation data suggesting implications beyond the contamination response, and you are withholding that analysis from the regulatory filing because it is incomplete.”

The repetition is not hostile. Not sarcastic. It is a technique, and I recognize it instantly, because at peak I parse communicative strategies with the same resolution I parse pharmacological data. Three functions, operating simultaneously:

She is ensuring I heard what I said, from outside, in her voice. The framing intact, the words reflected back without editorial modification, so I can hear them as a listener rather than as the person who chose them.

She is buying time. The repetition gives her the seconds she needs to separate what she feels about my answer from what she plans to say about it.

She is creating a record. By repeating my exact words, she has established what was said, and the establishment precedes her interpretation, which gives the interpretation an anchor it would not otherwise have.

I file the technique. It is excellent. It is the communicative equivalent of an IRIS uncertainty flag, a mechanism for ensuring that what follows is grounded in what actually occurred rather than in a reconstruction. I will think about this later. Not now.

“Yes,” I say. “That is what I’m telling you.”

“Dr. Silva.” She closes the tablet. “In my experience, twenty-two years of contamination response across this region, the analysis researchers describe as incomplete and not ready for filing is often the analysis that matters most. It is also the analysis that gets delayed longest, for reasons that are sometimes scientific and sometimes not.”

She says this without accusation. As a statement of professional observation: what she has seen, across two decades, about how scientists behave when they are sitting on significant findings. The observation is accurate. The accuracy is uncomfortable in the way that accurate observations about your own behavior are always uncomfortable, which is to say: the discomfort is not evidence that the observation is wrong.

“The analysis is incomplete because the mechanism is complex,” I say. “Not because I am withholding it.”

“I understand that’s your assessment.” She reopens the tablet. “I’m requesting full access to your laboratory data, all data related to this investigation, including the ongoing analysis. Standard request under the incident response protocol. You are not required to comply immediately, but I want to be transparent about the request and its basis.”

“I’ll provide everything related to the contamination findings. The ongoing analysis is subject to research confidentiality protections until it reaches a publishable standard. The university’s compliance office can advise on the appropriate framework.”

She nods. Not agreement; acknowledgment. The conversation has arrived at the boundary she anticipated. The boundary is where it is. She will work with it for now.

We spend another forty minutes on the contamination data. Osei is methodical. She asks about the population timeline's confidence interval first, not because it's scientifically primary but because it determines her framing: "If I tell the city the enhancement peaks in six weeks plus or minus two weeks, the public hears 'six weeks.' They do not hear 'plus or minus two weeks.' How tight is the interval?"

I give her the number. She writes it down and moves on.

She asks whether the aquifer transit time has been independently verified. It has not. She asks whether this matters for public communication. Yusuf answers correctly: it matters for the science, not the action, because the contamination is confirmed through direct sampling regardless of the transit pathway. Osei understands this distinction immediately: the science question and the action question are different questions, and she knows which one she needs answered right now. She files it. Moves on.

She asks about the communication strategy. When, how, through what channels. Whether the word *enhancement* should appear in any public document, given that it will be heard as "the water made people smarter" and that this framing, while technically not wrong, is the framing most likely to produce irrational behavior in both directions: hoarding by people who want more, panic by people who are declining.

I note, while this conversation is happening, that Osei is asking the right questions in the right order, but the right order is not my order. She starts from the downstream consequences and works upstream toward the science. I start from the science and work downstream toward the consequences. Both orderings are legitimate. Both produce different blind spots. Hers misses mechanism. Mine misses communication. Between us — between the two orderings held simultaneously in my working memory, which is a peak ability I am using and which I note I am using, between us we cover most of the problem. Though not the same parts.

* * *

She is walking toward the door. Yusuf is near the whiteboard. I am at my desk.

She stops. She does not turn around.

“I need you to understand what panic does to a city.” Her voice has changed, not dramatically, but the institutional register has thinned by a fraction, and underneath it I can hear the person, the woman who has been doing this work for twenty-two years and who has watched what happens when a community receives information it cannot yet process. “I have seen it. It is not hypothetical for me.”

Then she leaves. The door closes with the institutional click of a latch designed for a building where doors close often and quietly.

I sit at my desk. Yusuf sits in his chair, the mug in both hands. Neither of us speaks.

The line was not a threat. It was not a speech. It was a data point, delivered at the threshold by a person who decided, in the three seconds between standing and reaching the door, that the scientists behind her needed to know something she hadn't planned to disclose. That her position is not institutional convenience but lived experience. That she has seen panic and it left marks she does not discuss in incident response meetings.

I file it. I will retrieve it later, in a chapter of this story I cannot yet see, when a woman drives three hours from Atlanta to tell me something she is not required to tell me. For now the line sits alongside the other data: the compound, the cascade, the population, the deadline met, the gap remaining, and a person who carries the knowledge of what panic does the way I carry my father's filing system, absorbed so long it doesn't feel like memory. It feels like the way things are.

* * *

Two days later, Harris calls again.

Twelve seconds. She references the EPA filing. She reiterates Veridian's cooperation. She asks, in language so carefully constructed that no individual clause contains an actionable statement, whether I have considered whether direct engagement with Veridian's technical team might be productive for all parties. The sentence has the legal architecture of a structure designed to be

load-bearing in a deposition transcript: every word chosen for what it does not commit to, every phrase calibrated for deniability while communicating, to anyone listening at the right resolution, exactly what is being offered and exactly what is being implied.

I say no. She thanks me. The call ends.

I call Yusuf.

“Harris called again.”

A pause. “When?”

“Just now. Twelve seconds. She referenced the filing.”

“Mara.” His voice carries the quality of a person who has arrived at an inference he does not like but cannot dismiss. “We filed forty-eight hours ago. Standard EPA channel. Confidentiality protections. Harris is outside counsel for the party under investigation. She should not have access to filing details within forty-eight hours.”

“I know.”

“That call came faster than it should have.”

We sit with this. Neither of us says what it implies: that someone in the chain between my filing and the EPA’s processing has a pathway to Veridian’s legal team. That the confidentiality protections have been breached or circumvented or were never as robust as the submission protocol suggests. We don’t say it because saying it would require deciding what to do about it, and deciding would cost time we are spending on the cascade manuscript, and

the manuscript is more important than a leak we cannot plug from Gainesville.

“Document the call,” Yusuf says. “Time, duration, content. Forward to compliance.”

“Already done.”

“Of course you have.” A pause. “How’s the manuscript?”

“Close. Three more days.”

“You have three more days?”

He is asking about the compound. About my metabolite timeline. About whether the mind doing the work will sustain the work for the time the work requires.

“I don’t know,” I say. “I’m working on the assumption that I do.”

“Okay.” The word carries what he cannot say on the phone: the warmth in the first syllable, the resolve in the second. “Keep working. I’ll handle compliance.”

* * *

I put the phone down. I open the cascade manuscript.

The work is what matters. The leak is what does not. The distinction should be easy to maintain at peak, and for twenty-seven days it has been easy, and today — for the first time — I notice that maintaining it costs something. Not much. A flicker of effort where there was none before. The institutional pressure and the

scientific urgency and the closing window are converging on the same week, and the number of things I am holding simultaneously has a limit, and I can feel the limit.

I have not felt the limit before. Not once in twenty-seven days. The ceiling has been invisible, not because it wasn't there but because I was operating below it with enough margin that the boundary never registered, the way you don't feel the atmosphere until the pressure changes. Today the pressure has changed. Not dramatically. Not in a way that affects the work. But the ceiling feels, for the first time, like a ceiling rather than an opening.

I do not yet understand that this is the first sign of what comes next. The first sign of descent is not the loss of capacity. It is the awareness of capacity's edge, the moment the invisible boundary becomes perceptible, the way the surface of the water becomes visible only when you are close enough to touch it. The surface was always there. You just couldn't see it from where you were standing.

By tomorrow I will know what today is telling me. By tomorrow the model will hold but the holding will cost something measurable. By next week the cost will be large enough to track.

I open the manuscript. I keep working.

Author's Note — The Science of Chapter 11

Preprint Culture and the Speed of Science

The preprint server bioRxiv (launched 2013) and its medical sibling medRxiv (launched 2019) have transformed how scientific findings circulate. Preprints are manuscripts posted publicly before peer review, available immediately, citable, and increasingly used as the primary communication vehicle for significant findings.

The COVID-19 pandemic accelerated this shift dramatically. Many of the most consequential early findings about the virus, transmission dynamics, variant characterization, vaccine efficacy signals, were first available as preprints, sometimes weeks before formal peer review. This speed has genuine benefits and genuine risks: findings are available faster, but the peer review filter that catches errors has not yet been applied.

The tension between Mara and Osei, a researcher with incomplete findings and an incident commander who has seen what incomplete information does to a city, is one of the most persistent problems in public health communication. The literature consistently finds two competing risks: premature disclosure causes panic, misattribution, and policy overreaction; delayed disclosure leaves populations exposed and undermines public trust when the delay becomes known. The ethical literature on research disclosure in emergencies does not cleanly resolve this tension. It sits between them, unresolved, as it sits in the real world.

The hardest problems in public health are not scientific. They are the problems that live in the space between what science knows and what

public communication can bear.

CHAPTER 12 — THE SELF-STUDY

*D*_{AY 29.}

I run the cognitive battery at the same time I always run it, under the same conditions, with the same controlled variables, and when the results arrive I look at them for a long time and the first thing I notice is that I am not surprised.

I should be surprised. Eleven days ago I was running at the ceiling and the ceiling felt like an opening and the cascade model lived in my head the way music lives in a musician's: complete, rotatable, available in its entirety without effort. This morning I loaded the cascade model and I needed to check two parameters against my notebook because they were not fully present in working memory.

Two parameters. Out of seventeen. Eighty-eight percent retention. By most standards this is excellent. By the standard of Day 18, it is a crack in the wall.

The results.

Fluid reasoning: down 11 percent from peak. Still 19 percent

above my Day 0 baseline, which means I am operating at a cognitive level most researchers would consider extraordinary and that I, from the inside, experience as diminished. Nineteen percent above baseline should feel like a gift. From where I stand it feels like 11 percent below peak, and the felt experience of the loss is approximately twice the felt experience of the equivalent gain, which is exactly what Kahneman predicted and which I documented in the population model six days ago and which is now happening to me personally and which I can see happening and which the seeing does not correct. I have been telling other people about this bias for two weeks. The bias does not care that I have been telling people about it.

Processing speed: down 7 percent from peak. The smallest decline: processing speed is the last of the fluid domains to enhance and the last to recede. I still process fast. The difference between now and Day 18 is small enough that I might not notice it without measurement. But I measure, because I am a person who measures, and the measurement is unambiguous.

Working memory: down 18 percent from peak.

This is the number I sit with longest. Working memory is the workspace: the capacity to hold things active and manipulate them simultaneously. At peak I held seventeen variables and modified individual parameters while tracking their interactions. This morning: fifteen. I can hold fifteen. Two are in the notebook. They are retrievable. They are not in my head. The difference between holding something in working memory and retrieving it

from a notebook is the difference between thinking with a thing and looking up a thing. I have studied this distinction for twenty years. I am living it for the first time.

And then there is the episodic memory.

* * *

I cannot remember what I ate for dinner two nights ago.

This is a small thing. Most people cannot remember what they ate two nights ago. But three weeks ago I could, not because I cared about dinner but because the enhanced episodic memory recorded everything with a timestamp and a context and a retrieval precision that made recall effortless. The food was there. Catalogued. Available. Now it's not. Two nights ago I ate something, in my apartment, working on the cascade manuscript. The food is missing from the record.

It surfaces at 11 AM, unprompted, while I am annotating the transition kinetics: rice and beans. Leftover. The same meal I eat most nights, because I solved the dinner problem once and saw no reason to revisit it, which is either admirable efficiency or the sunk cost fallacy applied to legumes, and which I note is a connection I would not have made at peak because at peak the humor was compressed and incidental and did not need to reach for the joke, and the reaching is new. The reaching is the descent.

The retrieval works. It just doesn't work on demand. The memories are there. The access is intermittent. The episodic system

is consolidating normally, the memories are being formed, but the retrieval pathway is less reliable. They surface on their own schedule, prompted by associations I cannot predict.

This is consistent with the cascade model. The presymptomatic sequence begins with working memory, then affects episodic memory, first as retrieval instability, not storage failure. The memories exist. They are not always available when I reach for them. The distinction between “not stored” and “not accessible” is one of the foundational distinctions in memory research and I have been teaching it for a decade and I am now the data.

I write this in my notebook. The handwriting is steady. The observation is precise. The precision is doing something it was not doing three weeks ago: compensating. Holding the analysis level while the capacity underneath it erodes by increments too small to feel and large enough to measure.

* * *

I graph my scores. Two graphs. Same data, different reference points.

Graph one: current scores against Day 0 baseline. In this frame I am still significantly enhanced. Nineteen percent above on fluid reasoning. Fourteen on processing speed. Eight on working memory. A person performing well above her own historical ceiling. The graph shows gains.

Graph two: current scores against Day 18 peak. In this frame I

am in decline. Eleven percent below on fluid reasoning. Seven on processing speed. Eighteen on working memory. A person falling from a height she briefly held. The graph shows losses.

The data is identical. The reference point determines the experience.

I catch myself trusting graph one. Looking at it and feeling reassured: 19 percent above baseline is still extraordinary, the decline is modest, the working memory loss is manageable because the notebook compensates, because IRIS compensates, because the protocol I'm about to build compensates. I catch this, and I know what it is: the endowment effect, applied not to the peak this time but to the current state. The reference point sliding down with me so that each new level feels like the floor rather than the drop.

This is the thing the textbook doesn't convey about Kahneman's insight: the bias operates continuously. Every time the reference point shifts, the bias engages again, defending the new position as if it were the original one. I am in loss frame relative to peak and gain frame relative to baseline simultaneously, and both framings feel true, and neither is more correct, and the net effect is that I trust my current capacity slightly more than the data warrants.

I can see this happening. I cannot fully correct for it. The metacognitive capacity to catch the bias is subject to the same decline as the capacity being biased. The instrument and the measurement are made of the same material.

I write the numbers down. Both graphs. Both reference points.

The notebook holds what the working memory is starting to release.

* * *

I build the verification protocol between 10 AM and noon. It is an engineering solution to a measurement problem, and it is the most important piece of engineering I will do this week, more important than the cascade manuscript, because without the protocol the manuscript's conclusions cannot be trusted by the person writing them.

The protocol is simple. Every significant conclusion I reach goes to IRIS for independent verification before I act on it. IRIS has the formal model. IRIS has the literature. IRIS has the uncertainty framework I trained into it over three years. IRIS is not declining. IRIS can check my reasoning without being subject to the cognitive changes that make the reasoning less reliable.

I do not enjoy building this. It feels like installing guardrails on a road you used to drive without them, not because the road changed but because you did. The road is the same. The driver's reaction time is 18 percent slower in the working memory domain, and the guardrails are correct, and the need for them is a data point about the driver, and building them is the right response, and none of this feels good, and I am choosing to note that it doesn't feel good and to build them anyway, because being right matters more than feeling good, and this has always been true, and the compound did not teach me this, and the compound's

departure will not unteach it.

I am the aging mass spectrometer. IRIS is the reference standard. You calibrate against the standard before trusting the output. This is not humiliation. It is methodology.

I design the IRIS query template, a standard format for submitting conclusions for verification. Fields for the conclusion, the reasoning chain, the confidence estimate, and a request for IRIS to flag inconsistencies. Twenty minutes to design. Ninety seconds per verification to use. The ninety seconds are the cost. The trustworthiness of the output is the return.

* * *

The probabilistic thinking is harder than I expected.

I have always been comfortable with uncertainty in the abstract: confidence intervals, posterior distributions, the formal machinery of quantified doubt. These are the tools of my field. What I have not done, until this week, is apply them to myself with the honesty the protocol demands.

Every conclusion I submit must carry a confidence estimate. Not "I think this is probably right." A number. 0.78. 0.63. 0.91. The number strips the comfort from the vagueness, and the comfort, I am discovering, was doing more work than I realized. "I think this derivation is correct" could mean anything from 55 percent to 98 percent, and the spread between those two numbers is the space in which self-deception operates: comfortably, invisibly,

with the endowment effect's full cooperation. Putting a number on it closes the space. You have to look at what you actually believe. What you actually believe, as of Day 29 with 18 percent less working memory, is more uncertain than you want it to be.

First verification, 1 PM. A specific claim about the mGluR5 kinetics at the first cascade transition, a parameter I derived at peak, now checking against the formal model.

My confidence: 0.82.

IRIS verifies. The conclusion holds. IRIS's independent estimate: 0.86.

Close. Both high. The conclusion stands. But I notice — and noting is part of the protocol — that my confidence is lower than IRIS's. I expected overconfidence. I expected the endowment effect to inflate my trust in my own reasoning. Instead I'm underconfident, discounting my conclusions more than the evidence warrants. Compensatory overcorrection: knowing the capacity is declining, I subtract more than the decline requires.

This is why the numbers matter. "I'm probably right but not sure" is useless. 0.82 versus 0.86 is a trackable discrepancy that tells me something real about the relationship between my self-assessment and my actual accuracy: a real-time calibration of metacognitive honesty, updated with each verification. It is the most precise thing I have done today.

Twelve more verifications before 6 PM. The pattern holds: slightly underconfident on most, slightly overconfident on two. Average

discrepancy: 0.04. Small. Manageable. Within the range where the protocol adds value without creating paralysis.

For now.

* * *

At 6:30 I close the cascade manuscript, the section on mGluR5 transition kinetics, three days of work, almost finished. I will finish it tomorrow or the day after, at a cognitive level slightly lower than today's and slightly higher than next week's, and the distance between those two levels is a function of my metabolite clearance rate, which I did not choose and cannot change and which will not slow down because the manuscript needs more time.

I write in my notebook: *Note to self: you can't outrun chemistry.*

The sentence is simpler than the one I wrote on Day 15. *Optimization doesn't waive physics* was the humor of a person who could see the problem and still felt above it: the runner glancing at the distance and making a joke about it. *You can't outrun chemistry* is the humor of a person who is in the race and losing. Same flat delivery. Same notebook margin. Something different underneath. The underneath is heavier. The joke is doing more work: not bravado anymore, not the compressed wit of velocity, but something closer to the specific function humor serves when the situation it comments on is no longer hypothetical. Grief management. The phrase I used in the voice of the population model, applied now to a population of one.

I sit at the desk after writing it. The lab is quiet. The afternoon light through the west windows has come and gone; I either missed it or it happened and I was in the manuscript and I don't remember noticing, which is itself a data point. The episodic memory instability is present and active and doing what the cascade model predicts.

I think about my mother. Twenty-nine days since I've called. The number has passed from negligence into something that would require explanation, and the explanation would need language I don't have yet: that I found a mechanism that might matter for what's happening to her mind, that I identified a compound that might interrupt it, that I did all of this at a cognitive level I am leaving, and the leaving feels like the specific shape of grief I most wanted to avoid. The grief of capability diminishing while the need for it remains.

I don't call. I pick up the phone. I look at her name. I put the phone down. This is not the endowment effect. This is something older and less pharmacological, something that predates the compound and will outlast it. I do not have the bandwidth to examine it tonight, and I note that I used the word *bandwidth*, which is a word I have been using my whole career to describe cognitive resource allocation and which now describes, accurately and personally, the constraint on my available processing, and the word has not changed but my relationship to it has.

I go home. I eat dinner. I don't remember what I eat, and this time I notice the not-remembering in the moment rather than

the next morning, and the noticing in the moment is the change: not that the memory fails, but that I am aware of the failure as it happens, which means the metacognitive monitoring is still running, still catching the errors, still filing the gaps. The system is degrading. The monitoring of the degradation is not. Not yet. This is not reassuring in the way I want it to be reassuring. It is reassuring in the way the data allows.

Author's Note — The Science of Chapter 12

Metacognition and Its Limits

Metacognition, thinking about thinking, monitoring one's own cognitive processes, is one of the most studied and most humbling areas of cognitive psychology. The consistent finding: humans are systematically poor judges of their own cognitive performance, and this poor judgment is not improved by intelligence or expertise. In fact, the Dunning-Kruger effect and its more nuanced descendants suggest that the relationship between confidence and competence is complex and non-linear.

For Mara, the specific problem is that the metacognitive capacity to catch cognitive errors is itself subject to the same degradation as the functions it monitors. The instrument and the measurement are made of the same material. This is not an abstract philosophical concern. It is a practical engineering problem, and Mara treats it as one: building an external verification system that routes around the metacognitive limitation rather than trying to solve it

directly.

Knowing you have a blind spot does not tell you what's in it.

* * *

How to Use This — Probabilistic Thinking

In this chapter, Mara forces herself to put numbers on her uncertainty. Instead of “I think this is probably right,” she commits to “I estimate 0.82 confidence in this derivation.” This practice, drawn from the Good Judgment Project led by Philip Tetlock and documented in *Superforecasting* (2015), is one of the most reliable techniques for improving decision quality.

The mechanism is simple: vague language allows you to avoid committing to a position. “I think so” can mean anything from 55% to 95%. Once you assign a number, you can no longer hide behind the ambiguity. You must examine what you actually believe and why; and when new evidence arrives, you can update the number rather than defending the vague position.

This is especially important when the thing you are evaluating is yourself. Mara’s instrument is her own mind. Numerical estimates make the measurement visible and force her to confront its limits, and, as she discovers, the discrepancy between her estimates and IRIS’s provides a real-time calibration of metacognitive accuracy that vague confidence could never produce.

The principle: put a number on your uncertainty. Vague confidence is not confidence. It is avoidance.

CHAPTER 13 — THE THIRD COMPOUND

*D*_{AY 31.}

The original Phase 2 trial data for RGH-618 arrives on a Tuesday morning through a research data-sharing exemption that takes four days to process and six emails to finalize. The exemption is standard: academic researcher requesting anonymized clinical trial data for non-commercial purposes. The data exists in a regulatory archive because Gedeon Richter, like all companies that run human trials, is required to preserve it indefinitely, which is one of those regulatory requirements that is tedious and expensive and correct and that occasionally, twenty years later, turns out to contain a finding nobody knew was there. The archive is accessible. The bureaucracy is slow. The bureaucracy does not know — has no mechanism for knowing, no form for indicating, no checkbox for *this request is time-sensitive because the requesting researcher's cognitive capacity is declining on a pharmacokinetic schedule she can calculate* — that the four days it took to process the exemption cost me approximately 2 percent of working memory capacity, measured against the curve I am riding down.

The dataset: 847 patients, randomized across three dose groups and placebo, twelve weeks of cognitive assessments and adverse event reporting, genetic panels on a subset of 312. Large. Thorough in the way of clinical trial records from the early 2000s: precise about what was measured, less precise about what wasn't, because what wasn't measured was determined by the pharmacogenomic tools available in 2005, which did not include the tools that would have prevented the compound from being abandoned.

I upload it to IRIS at 11 AM and write the query. The query takes me twenty minutes to compose. At peak it would have taken five. The difference is not that I've forgotten how to write queries, the skill is intact, procedural, but that organizing the query requires holding the trial structure, the genetic panel, the cascade mechanism, and the specific analysis I want, simultaneously, while constructing a sentence that specifies all four in the correct relationship to each other. My working memory is down 18 percent from peak. The query requires holding four things. I can hold four things. But the holding costs attention that used to be free, the way carrying a weight you've been carrying for a mile costs more per step than carrying it for the first hundred yards.

The query: *Cross-reference the RGH-618 trial's adverse event data against patient genetic profiles. Identify subpopulation patterns in the adverse event distribution. Cross-reference identified subpopulations against the cascade mechanism's genetic markers, specifically mGluR5 receptor variants and BDNF Val66Met polymorphism. Flag all assumptions. Quantify all uncertainty.*

I queue the overnight batch. I go home. I eat dinner and do not remember what, which I have stopped finding alarming and started treating as the new background condition. The episodic gaps are not widening; they are stable at the current level. The current level includes not remembering dinner. I have adjusted the threshold for alarm accordingly, which is either a healthy adaptation or the endowment effect applied to cognitive decline itself, and I cannot determine which from the inside, and this inability is the same inability I documented in Chapter 12, and the documentation does not resolve it, and the compound does not care.

* * *

The IRIS output is waiting at 6:30 AM on Day 32. Nine pages. The longest output IRIS has ever produced. Fourteen footnotes. Six uncertainty flags.

I read it once quickly. Then again slowly, checking each step.

IRIS // Day 32 // 04:17:33

RGH-618 Phase 2 trial reanalysis - subpopulation adverse event r

Summary: Adverse events in the original trial (primarily gastrointestinal - nausea, abdominal discomfort, diarrhea) were not uniformly distributed across the trial population. They were concentrated in a specific genetic subpopulation identifiable retrospectively from the available genetic data.

Subpopulation identification:

Patients carrying the BDNF Val66Met polymorphism (Met/Met or Val/Met genotype) AND classified as CYP2C19 slow metabolizers (PM phenotype) showed adverse event rates 4.7x higher than the remainder of the trial population ($p < 0.001$, $n = 43$ in genetic subset).

Mechanism (inferred - confidence 0.71):

CYP2C19 slow metabolizers clear RGH-618 more slowly, producing sustained peripheral mGluR5 activation. In patients who also carry Val66Met, the altered BDNF signaling in the enteric nervous system amplifies the peripheral mGluR5 response, producing the observed GI effects. This interaction was not characterized in the original trial because:

- (a) CYP2C19 phenotyping was not standard pharmacogenomic practice in 2005
- (b) The BDNF Val66Met interaction with peripheral mGluR5 has only been documented since Chen et al. (2019)
- (c) The trial was not powered to detect subpopulation effects of this size

Implication: The adverse effects attributed to RGH-618 are not intrinsic to the compound. They are specific to a subpopulation (approximately 4-6% of the general population) identifiable through standard genetic screening and excludable from treatment or manageable with dose adjustment.

For the remaining 94-96% of the population, the trial data shows no significant adverse events beyond placebo.

Note: This reanalysis depends on the genetic data available for 312 of 847 trial participants. The subpopulation identification is based on this subset and extrapolated to the full trial population. The extrapolation introduces uncertainty.

Confidence in subpopulation identification: 0.83

Confidence in mechanism (BDNF × CYP2C19 interaction): 0.71

Confidence in clinical implication (screen and treat): 0.67

I recommend independent verification of the genetic subset analysis before this finding is communicated externally.

I read it a third time. I get a sparkling water, the grapefruit, which has been wrong for four months and which I have not corrected and which is, at this point, as much a part of my laboratory identity as the notebooks and the filing system, and which Yusuf would say is a metaphor for my relationship with imperfect situations, and which Yusuf would be right about, and which I am not going to fix today because today I have a more important imperfection to address.

* * *

The compound didn't fail. The trial design failed to see what it had.

RGH-618 was tested on 847 patients with mild-to-moderate Alzheimer's. Modest cognitive improvement, $p = 0.08$, didn't reach significance. Dose-limiting GI side effects. Conclusion, reasonable at the time: benefit didn't justify cost. Discontinued 2006. Filed. Forgotten. A pharmaceutical footnote in a database nobody searches for the same reason nobody searches the bottom drawer of a filing cabinet: you only look when you know something is there, and nobody knew.

But the adverse effects were not the compound's fault. They were the collision between the compound and a specific genetic subpopulation, CYP2C19 slow metabolizers carrying BDNF Val66Met, producing sustained peripheral mGluR5 activation that the trial was never designed to detect because CYP2C19 phenotyping wasn't standard practice in 2005 and the relevant BDNF interaction hadn't been published until 2019. The trial threw out the compound because a minority of patients had side effects caused by biology that nobody could have characterized at the time. Like throwing out a bridge because it failed in an earthquake that wasn't in the building code.

My father would have understood this immediately. You don't blame the bridge. You update the code.

For the remaining 94 to 96 percent of the population: no significant adverse events. A compound that modulates mGluR5 at the first cascade transition, crosses the blood-brain barrier, has a known synthesis pathway, and has been sitting in a database for twenty years because someone drew a conclusion from data

that didn't include the variable that would have changed the conclusion. The most expensive kind of failure: not wrong, just incomplete.

The right compound. The wrong analysis. Twenty years.

* * *

I spend two days verifying. This is not optional. The finding is extraordinary and extraordinary findings require extraordinary verification, and the verification must be done by me because I am the person who will stake a scientific reputation on the claim, and the claim is that a major pharmaceutical company discontinued a compound that could interrupt the presymptomatic neurodegenerative cascade in hundreds of millions of people because of a pharmacogenomic interaction that was unknowable at the time.

I go through IRIS's analysis step by step. The subpopulation identification, the statistical clustering of adverse events in CYP2C19 PM carriers with Val66Met. Sound. Significant. I check the methodology against three textbook approaches to pharmacogenomic subgroup analysis. IRIS used the correct one.

The mechanism, the BDNF \times CYP2C19 interaction producing peripheral mGluR5 amplification. I trace through the molecular biology. Each step has published support. The Chen et al. 2019 paper on peripheral mGluR5 and BDNF signaling in the enteric nervous system is real and relevant. I read it in full. At peak I would have read it and simultaneously cross-referenced it against the trial data and the cascade model. Today I read it, take notes,

then cross-reference. Sequential rather than simultaneous. The work takes longer. The work is no less thorough.

The clinical implication: screen for the subpopulation, exclude or dose-adjust, treat the remainder. Straightforward once the subpopulation is identified. Standard pharmacogenomic practice. The screening test exists. The cost is modest. Four hours in a clinical lab.

Each step holds.

And I notice, on the second day of verification, something I need to be honest about. There are three steps in IRIS's statistical modeling, specifically in the subpopulation clustering, that I can follow but could not have generated independently. Not because the statistics are beyond me in principle. Because the analysis requires holding the full trial dataset, the genetic subset, the adverse event matrix, and the cascade mechanism's receptor kinetics simultaneously. My working memory is down 18 percent from peak. The analysis requires a capacity I had ten days ago and do not have now.

I can audit the steps. I can verify each follows from the previous one. I can check inputs and confirm outputs. I cannot look at the raw data and arrive at the same analysis from scratch. The audit is intact. The generation is not.

This is the gap the verification protocol was designed for. The protocol is working: every step confirmed, every assumption checked. But the gap is present, and it is wider than it was five

days ago. The distance between what I can verify and what I can produce is growing at a rate that the cascade model predicts and that I can track and that I cannot stop.

I write: I can still read the map. I can no longer draw it from memory. The protocol works because reading and drawing are different skills and reading degrades slower. But the distance between them is growing, and at some point the distance will matter.

* * *

On the afternoon of Day 33 I do the math on the affected populations.

If the cascade model is correct, if the presymptomatic signature represents a unified mechanism across disease categories, then the population is not the 55 million currently diagnosed with Alzheimer's. It is the total population at risk for any neurodegeneration driven by the cascade. Alzheimer's. Lewy body dementia. Frontotemporal degeneration. Vascular dementia with a neurodegenerative component. The number, from published epidemiological data, conservatively: 150 to 200 million people worldwide who will develop clinically significant neurodegeneration in their lifetimes.

I do the math three times. It does not change. I sit with it. The afternoon light is doing its thing through the west-facing windows and I notice it and I do not have the working memory to hold the noticing alongside the number, so I let the light go and I hold the number. This is the plateau voice: the choosing, the letting go

of one thing to hold another, the awareness that the choosing is necessary in a way it was not at peak. At peak I held the number and the light and the cascade model and the Fountain of Youth parallel and the thought of my mother simultaneously. Today I hold the number. The number is enough.

One hundred fifty to two hundred million people.

A compound in a database for twenty years.

A screening test that takes four hours.

At 4 PM IRIS produces an additional output, unprompted:

```
IRIS // Day 33 // 16:02:14
```

```
Supplementary analysis - assumption dependency assessment
```

Note: The finding regarding RGH-618 subpopulation safety depends on seven interconnected assumptions:

1. The cascade model accurately represents the presymptomatic mechanism (confidence: 0.79)
2. The mGluR5 transition is the correct intervention point (confidence: 0.84)
3. RGH-618's receptor binding profile matches the intervention requirements (confidence: 0.79)
4. The BDNF × CYP2C19 interaction explains the trial adverse events (confidence: 0.71)
5. The genetic screening protocol correctly identifies the vulnerable subpopulation (confidence: 0.83)
6. Presymptomatic administration timing is achievable in

clinical practice (confidence: 0.62)

7. The cascade model generalizes beyond the compound-induced degradation observed in Gainesville (confidence: 0.58)

Joint probability that all seven assumptions are simultaneously correct: 0.61-0.79, depending on correlation structure between assumptions.

I have estimated the correlation structure from published replication rates in related pharmacological findings. The estimate introduces additional uncertainty.

I recommend you review my methodology before proceeding. This recommendation is procedural but in this case I want to emphasize it.

I want to emphasize it.

I read this four times. IRIS has never used the phrase *I want* before. It has said *I recommend*. It has said *I flag*. It has said *Note*. It has never expressed a want, and I do not know whether this is the emergence of something I should pay attention to or whether it is a linguistic artifact, my own annotation style, in which I occasionally write *I want to be clear about this*, reflected back through the language model in a context where emphasis was the appropriate communicative function.

I do not know. I cannot determine from the data available whether IRIS's "I want" is a person speaking or a pattern matching, and

the inability to distinguish is itself a finding, though I am not sure what it is a finding of. At peak I would have chased this question. At peak I had the bandwidth to hold the philosophical alongside the pharmacological and to find the relationship between them. Today I note it. I file it. I trust the filing system to hold it until I have the capacity to return to it, which may be never, and the “may be never” is a phrase I could not have used a month ago because a month ago the future was large enough to contain everything.

I review the methodology. Sound. The seven assumptions correctly identified. The confidence estimates within the range I would assign. The joint probability reasonable. Assumption 7, whether the cascade model generalizes beyond Gainesville, is the weakest, flagged at 0.58, honest in a way that a more confident system would have obscured.

I write: *It knows what it doesn't know. That's the thing that keeps surprising me.*

Then underneath: *And it says "I want." I don't know what that is. 0.35 confidence that it represents something. Noted.*

* * *

I call Yusuf at 5 PM. I tell him about the subpopulation finding, the assumption dependency analysis, the joint probability range.

He is quiet for a long time. Then: “The compound didn’t fail.”

“No. The trial design failed to see what it had.”

“And the subpopulation screen is standard pharmacogenomics.”

“Standard. Four hours. Existing technology. Nothing needs to be invented.”

“Mara. How many people?”

“One hundred fifty to two hundred million. Worldwide. Over lifetimes. If the cascade model generalizes.”

Another silence. I can hear him breathing. I think about what this number sounds like to a man who sends money to Lagos every month and set up a scholarship fund in the neighborhood where he grew up and coaches eleven-year-olds on Saturday mornings because he believes that presence is a form of contribution and that contribution is what you do with the capacity you have. I have always respected this about Yusuf without fully understanding it, because my version of contribution has been the work, the research, the publications, the grants, the machinery of academic science, and his version includes all of that and also the soccer team and the jollof rice and the phone call to Lindström and the date on the mug. His radius of obligation is wider than mine. It always has been. I am better at the science. He is better at the everything-else. Between us we cover more of the problem than either of us alone.

“And the joint probability?” he asks.

“Point six one to point seven nine.”

“That’s honest.”

“IRIS is always honest.”

“I know. That’s the thing.” A pause. “When do you publish?”

“The cascade manuscript is three to four days from complete. Then the RGH-618 reanalysis goes in a supplementary paper. Both as preprints, simultaneously.”

“Simultaneously?”

“The cascade model without the intervention is incomplete. The intervention without the cascade is uninterpretable. They go together.”

“Okay.” The word carries the weight it always carries at the end of a conversation that has changed what we understand about the work. “I’ll start preparing the Lindström package.”

“The Lindström package.”

“The data she’ll need. Your experimental design from Day 24. The RGH-618 binding profile. The subpopulation screening protocol. Everything she needs to start the experiment the day the preprints go up.”

He has been planning this. He has been building the international validation pathway since Day 8, when he sent the email to Stockholm, and the building has been running in parallel with my science the way his Saturday soccer runs in parallel with his pharmacology: quietly, without announcement, with the competence of a person who understands that the network is as important as the finding and that the network must be ready before the finding

arrives.

“Thank you, Yusuf.”

“Don’t thank me yet. Thank me when Lindström replicates.”

* * *

I hang up. I sit in the lab. The west-facing light is gone; past 5 PM, the angle shifted, fluorescents now. I think about 150 to 200 million people. About a compound in a database for twenty years. About IRIS saying *I want to emphasize it* and the 0.35 I assigned to the question of whether that means something.

Three days to finish the manuscript. Maybe four. The working memory is holding at its current level: not declining further this week, not recovering either. The episodic gaps are present and stable. The verification protocol is working. The audit gap is wider than five days ago and not yet critical, and the word *yet* is load-bearing in that sentence in a way I can hear.

I also think about my mother. About the shelf the same height and the arm slightly shorter. About the cascade mechanism that may be running in her hippocampus right now, slowly, the presymptomatic phase that takes decades and that the compound compressed into weeks and that I can now see because the compression made it visible. About the compound that could interrupt it, if the mechanism is right, if the intervention point holds, if assumption 7 generalizes, if Lindström validates. About the distance between the science and the person the science might help,

which is the distance between Gainesville and Tampa, sixty miles, an hour's drive, a phone call I have not made in thirty-one days.

I do not call. Not tonight. The manuscript needs three more days and my capacity is a resource and the resource is finite and I am spending it where the spending helps the most people, which is the calculus I have been running since Day 0 and which produces the correct answer and which does not address the question of whether the correct answer and the right answer are the same thing.

I open the cascade manuscript. I keep working.

Author's Note — The Science of Chapter 13

What Phase 2 Trials Are Designed to Find (and Miss)

Clinical trials are organized into phases. Phase 1 asks: is this compound safe in humans, and at what dose? Phase 2 asks: does it show preliminary evidence of efficacy in patients with the target condition? Phase 3 asks: does it work, definitively, in a large enough population to support regulatory approval?

Phase 2 trials are designed to detect average effects across a trial population. They are generally not powered to detect subpopulation effects. A compound that works brilliantly in 94% of the population and causes adverse events in 6% will, in a standard Phase 2 analysis, show a modest average effect with a significant adverse event rate. The standard conclusion: benefit doesn't justify cost. Discontinued.

Pharmacogenomics, the study of how genetic variation affects drug response, has been revealing how many discontinued compounds may have failed not because they didn't work but because the trial couldn't see who they worked for. The FDA has increasingly required pharmacogenomic stratification in new trial designs, but compounds abandoned before this requirement took effect remain in databases, untested at the resolution that would have revealed their true profile.

RGH-618 is fictional. The scenario it illustrates, a compound that works for the majority but fails in a genetically identifiable minority, is not. It has happened, and continues to happen, in real drug development.

Most failures are not failures of the compound. They are failures of the question.

CHAPTER 14 — YUSUF'S CURVE

*D*_{AY 33.}

Yusuf's cognitive battery scores peaked three days ago. He tells me this at dinner, at the Thai place on University Avenue that we have been eating at approximately once a week for eleven years, a duration long enough that the staff know our order and our table preference and the fact that Yusuf brings his own mug, which they accommodate with hot water and the quiet graciousness of people who decided years ago that a regular customer's eccentricity is part of the service and not a disruption of it. The restaurant is one of the fixed points of our friendship. We have had eleven years of conversations at this table, about pharmacology and IRIS and department politics and his mother's groundnut stew recipe and my inability to make *carne de sol* correctly and the specific quality of Gainesville light in different seasons. Not all of the conversations have been important. All of them have been reliable. Tonight's will be important. I do not know this yet when we sit down.

He orders the green curry, medium spice, brown rice. He has ordered this every time. I notice, and the noticing reassures me in a way I don't examine, because examining why the consistency of a colleague's curry order produces reassurance would require me to acknowledge that consistency itself has become a form of comfort, that in a period when my own internal consistency is degrading on a pharmacokinetic schedule, the external consistencies I can rely on have become load-bearing in a way they were not before.

I do not examine this. I order the pad thai.

"Fluid reasoning is down 6 percent from my peak," he says. "Processing speed down 4. Working memory down 9." He holds the UCL mug. The tea is Earl Grey, a London import, steeped for exactly four minutes, a protocol he follows with the kind of precision that other people associate with scientists and that I associate with Yusuf specifically, because his precision is not about measurement but about care. "The decline started Saturday. I noticed Sunday."

His numbers are smaller than mine. His curve is shallower, his descent gentler, because his metabolizer profile is intermediate and his peak was lower and the fall from a lower peak is a shorter drop. He will come down slowly. Weeks, not days. The gentle arc I predicted in the population model, which I am watching now from across a table at a restaurant on University Avenue while my own arc is steeper and faster and further along, the way the front runner in a race sees the pace car pulling ahead — except

the pace car is the cliff and the front runner is me.

“How does it feel?” I ask.

He thinks about this. The thinking is visible the way it always is: not performed, not hidden, just present on his face like weather. He picks up the mug. Sets it down.

“I keep reaching for things that were there last week,” he says. “It’s not like forgetting. It’s more like —” He pauses. I can see him searching for the comparison, turning possibilities over, and I can see the moment he finds it. “It’s like the shelf is the same height but my arm is slightly shorter. The thing I want is there. I can see it. I just can’t quite reach it as easily as I could.”

I set down my fork.

The shelf is the same height. The arm is slightly shorter.

This is the most accurate description of early consolidation failure I have ever heard. More accurate than anything in the clinical literature. More accurate than what I wrote in my own notebook when I couldn’t remember dinner. The clinical vocabulary, encoding deficits, retrieval failures, reconsolidation errors, describes the mechanism. Yusuf describes the experience. The memory is present. The access is impaired. Not an empty shelf. A shorter arm.

He arrived at this through his own experience, not through the neuroscience vocabulary we share, and the description is better for having come from outside the science. This is the thing about Yusuf I have never been able to fully account for: his intelligence

arrives from a different direction than mine, and it is not less precise for arriving differently. I have the models. He has the metaphor. The metaphor is more useful tonight than the models, and I am a person who has built her career on models, and I am sitting across from a person who just outperformed my entire field's descriptive vocabulary in two sentences about a shelf.

"That's very good," I say.

"It's not good. It's accurate." He almost smiles. "Those are not the same thing."

* * *

We eat. The restaurant is half full, Monday evening, the campus crowd thinning. The noise level allows conversation without being overheard, which matters because what we are discussing is a compound in the municipal water supply that is declining in the blood of everyone in this restaurant, including the waiter who refilled my water glass without being asked and the couple at the next table splitting a dessert and the kitchen staff whose voices I can hear through the service window. The awareness of this sits between us like a third plate nobody ordered.

Yusuf eats methodically. I eat pad thai I ordered without deliberation, the same as last time, which I remember ordering last time; the retrieval works for this, the episodic memory holding for familiar patterns even as it becomes unreliable for recent novel events. The inconsistency of what I remember and what I don't matches the cascade model's predictions, which is both scientifi-

cally satisfying and personally terrible, and a finding being both at once is something I have become familiar with over the past five weeks.

“Your numbers,” Yusuf says. Not a question.

“Fluid reasoning down 14 percent from peak as of this morning. Working memory down 22. Processing speed down 9.”

He does not put on the equanimity. Tonight the mask is not present. What is present is the face of a friend hearing numbers that describe the decline of the thing about you he most values and that he cannot do anything about and that he is going to sit with rather than try to fix, because he knows — has always known, this is one of his specific forms of intelligence — that some things cannot be fixed and that presence in the unfixable is itself a contribution.

“The verification protocol?”

“Working. Average discrepancy between my estimates and IRIS’s is 0.05. Slight overconfidence on two of the last eight verifications.”

“That’s a shift from last week.”

“Yes. I was underconfident last week. Overconfident this week. The bias is moving.”

“In which direction?”

“I trust my conclusions slightly more than I should. Small margin. Correctable.”

“For now.”

“For now.”

The phrase *for now* has been doing increasing work in both our vocabularies. It means: the current state holds. The current state is changing. We both know it is changing. Neither of us knows when the change will exceed what *for now* can accommodate. The phrase is a bridge with a known load limit and an unknown traffic forecast. My father, the bridge engineer, would have had thoughts about this. He would not have found the uncertainty comfortable. He would have built for it anyway.

* * *

Yusuf says the thing about IRIS between the main course and the tea. He says it simply, in the way he says things he has decided are ready and that he knows I will not want to hear.

“You know you’re going to have to let IRIS carry more of this than feels comfortable.”

“I know.”

“I don’t think you do yet.” He pours the hot water. The tea steeps. He watches me while it steeps, which is the specific attention of a man who has decided something about the order of importance in this conversation and the tea is not at the top. “I think you know it intellectually. You have the protocol. You have the verification system. You understand, as a scientist, that IRIS is the scaffold and the scaffold becomes more load-bearing as the structure it

supports becomes less stable. You understand this.”

“I do.”

“But you haven’t felt it yet. You haven’t reached the moment where you submit a conclusion and IRIS says it’s wrong and you can’t independently verify which of you is right. That moment is coming.”

I don’t answer. The pad thai is done. The restaurant is quieter.

“When that moment comes,” he says, “you’re going to want to trust yourself over the system. Because you built the system and you are you and trusting yourself has been the correct strategy for forty-five years. But the strategy assumes the instrument is reliable. The instrument is changing. The system was built by the version of you that was reliable. The system hasn’t changed.”

“I know this, Yusuf.”

“You know it. You don’t feel it yet.” He takes a sip of the tea. “You can’t out-think shorter arms.”

The phrase lands. It lands because it connects his metaphor, the shelf and the arm, to the specific cognitive problem I am facing with IRIS. The arm is my working memory. The shelf is the cascade model at peak resolution. The arm is getting shorter. The shelf hasn’t moved. IRIS is standing next to the shelf, at full reach, holding what I put there when my arm was longer. Trusting IRIS is not trusting a machine. It is trusting the version of myself that built the machine, which is the version with longer arms, which is the version I am leaving.

I sit with this. The restaurant noise has dropped to the level where fragments from other tables are audible: someone laughing, the rhythm of other evenings. The sound is warm and has nothing to do with what is happening to me, and I notice it, and the noticing is grateful in a way I do not try to make precise.

“The shelf metaphor,” I say. “The arm getting shorter.”

“Yes.”

I am thinking about my mother. The thought arrives whole. Not as a connection I construct — as a presence that has been sitting below the active processing and that surfaces now because Yusuf said *shorter arms* and my mother is seventy-four and her arms have been getting shorter for two years. Slowly. Not the compound. Something older and more patient and not reversible, the slow version of what is happening to me in weeks, and I have noticed and I have not said anything to her or to Yusuf or to anyone, and the mechanism I found at peak may be the mechanism that is happening to her right now in a house in Tampa where the jalapeños produced fourteen peppers this year.

I almost say it. The words are formed. They are in the space between my mind and my mouth: *My mother is experiencing this. Not the compound — something slower. And the mechanism I found might be her mechanism. And the compound I found might help her. And I can't finish the work in time because the same chemistry that showed me the problem is taking away my ability to prove the solution.*

The words do not cross. They sit in the space and I hold them

there and the holding costs something and I let them go, the way I have learned to let things go this week when the working memory cannot carry everything.

“I’ll let IRIS carry what it needs to carry,” I say.

Yusuf looks at me. He hears what I said. I think he hears what I didn’t say, not the specific content, not Carol, but the shape of the silence. He has always been able to hear the shapes of my silences. I have never accounted for this. It is not in any model I have built. It operates outside the systems I understand.

“Okay,” he says. “More tea?”

* * *

We walk back to campus in the dark. The oaks on University Avenue overhead. The route is automatic, eleven years of occasional dinners at the Thai place have written this walk into procedural memory, the system the compound hasn’t touched and the cascade model predicts will be the last to fail. Our legs know it. Our attention is on each other.

We walk mostly in silence. Eleven years includes the ease of shared quiet, the agreement, never discussed, that time together doesn’t require words. The silence after a difficult conversation is its own form of presence, and presence is the thing Yusuf does better than anyone I know.

Halfway across campus he says something he held back at dinner.

"I've been thinking about the verification protocol."

"What about it?"

"I'm part of it. I check your work. I maintain the parallel log. I'm your external verification for the verification system." He pauses. The oaks move slightly above us. The September night is warm. "I'm declining too, Mara. Slower than you. But declining. At some point I won't catch what I would have caught today. The protocol needs more than me."

I should have been thinking about this. I haven't been. The realization that I haven't been is itself a data point, a thing I would have thought of at peak, a connection between Yusuf's decline and the protocol's reliability that requires holding both trajectories simultaneously, which I did not do, which means my working memory is not just smaller but less thorough in its coverage, missing connections it would have made.

"Who else should know the protocol well enough to participate?" he asks.

I think. Elena, whose questions about the study design are getting sharper. James, quiet and careful. Lindström, who will need the full verification framework to evaluate the Stockholm results.

"I'll think about it," I say.

"Think about it soon."

I note, walking under the oaks, that the protocol has a single point of failure and the single point is a man who is also declining

and who has the specific courage to name this about himself while walking home from dinner. The courage is not dramatic. It is the practical courage of a pharmacologist who understands degradation curves and who applies the understanding to himself with the same honesty he applies it to the data. This is who Yusuf is. I have known this for eleven years. Tonight it matters more than it usually does.

* * *

We reach the Brain Institute. The building is dark except for the lobby lights and the glow from two upper-floor labs where someone is working late.

“The preprints,” Yusuf says. “How close?”

“The cascade manuscript is two days from complete. The RGH-618 supplementary is drafted. One more verification pass on both.”

“And then?”

“Then I post them. Both at once. And we see what happens.”

He nods. Neither of us says what both of us are thinking: by the time the scientific community responds, the skepticism, the replication attempts, the Lindström validation, my capacity will have declined further. The person who defends the preprints may not be the person who wrote them. This is why the preprints need to be written well enough to defend themselves.

“Goodnight, Mara.”

“Goodnight, Yusuf. Thank you for dinner.”

“It’s the Thai place. You paid.”

“Thank you for dinner anyway.”

He almost smiles. I almost smile. These are both real. The almost is the currency we trade in, the space between the expression and the feeling, close enough to count.

* * *

I walk the rest of the way alone. The apartment is quiet. The chess set is on the shelf near the window. Pieces in their starting positions. Thirteen years.

I look at it tonight the way I have been looking at it for thirteen years: briefly, and then I look away. The looking away is the same as always. Except tonight I know more about what the looking away means, because I have spent five weeks studying the mechanism by which minds lose what they hold, and the chess set is a thing I am holding, and the holding is not cognitive but something older, and the distinction between cognitive holding and whatever this is — grief, maybe, or loyalty, or the specific form of love that expresses itself as the maintenance of an object in its last known configuration — is a distinction I could name if I wanted to and do not want to.

The shelf is the same height. The pieces haven’t moved. My father

put them there, or I put them there for him, and the fact that I cannot remember which is not the compound. The compound takes recent memories. This memory is thirteen years old and has been consolidated so deeply that it is part of the architecture rather than the contents, and I cannot remember the moment of placement because the moment became the structure, the way the oaks on University Avenue became the route and the route became the walk and the walk became how I get home.

I set my alarm. I do not sleep well. Tomorrow I finish the cascade manuscript. The day after I post the preprints.

The compound metabolizes. The metabolites clear. The arm gets shorter. The shelf doesn't move.

Author's Note — The Science of Chapter 14

The Phenomenology of Consolidation Failure

Yusuf's description of early cognitive decline, "the shelf is the same height but my arm is slightly shorter," is a more accurate rendering of the subjective experience of early consolidation failure than most clinical descriptions provide. The clinical literature tends to characterize memory decline in terms of encoding deficits, retrieval failures, and reconsolidation errors. These are mechanistically precise. They do not capture what it feels like.

The distinction Yusuf draws is exact: the information is present (the shelf is the same height), but the access is impaired (the arm is shorter). This maps onto the neuroscientific distinction

between storage and retrieval, first articulated by Endel Tulving in the 1970s and confirmed by decades of subsequent research. Many forms of memory impairment involve retrieval failure rather than storage failure: the memory was successfully encoded and consolidated, but the pathway to access it is degraded.

This is why memories surface later, unprompted, triggered by environmental cues or associative networks that bypass the degraded retrieval pathway. The memory was always there. The direct route to it was not.

The shelf doesn't move. The arm gets shorter. The difference matters enormously, because it means the information is not lost. It is inaccessible. These are different problems with different solutions.

CHAPTER 15 — THE VERIDIAN LEAK

*D*_{AY 35.}

The preprints go up at 6 AM on a Wednesday. Both of them, the cascade manuscript and the RGH-618 supplementary, posted simultaneously to bioRxiv, timestamped, archived, publicly accessible, and beyond anyone's ability to suppress or retract.

I click submit and sit at my desk for a moment. The act takes four seconds. The work behind it took thirty-five days, and the thirty-five days took most of what I had, and what I had was more than I have ever had and is now less than it was, and the preprints need to carry the finding forward from here because the person who wrote them cannot guarantee she will be the person who defends them. Four seconds. Two papers. One hundred fifty to two hundred million people, if the mechanism holds, if assumption 7 generalizes, if Lindström validates, if the evidence survives peer review by people who will have every legitimate reason to be skeptical and who should be.

I pick up my sparkling water. The grapefruit. Still wrong. Still

cold. Still adequate. I have been drinking the wrong sparkling water for five months now, and this fact has persisted through a cognitive enhancement, a pharmacological decline, a major scientific discovery, and the filing of two preprints that may reshape neurodegenerative disease research, and the sparkling water has not changed and I have not changed it, and I find this — briefly, before the morning’s urgency reasserts itself — genuinely funny in a way that the descent has made more accessible rather than less. At peak, the humor was compressed and incidental. Today the humor arrives slower and stays longer. I drink the grapefruit. It is fine. Fine is underrated.

The cascade manuscript is forty-one pages. It describes the presymptomatic neurodegenerative signature, the unified disease hypothesis, the mGluR5 intervention point, the population degradation model, and the compound that revealed all of it. Every derivation is self-contained. Every assumption is flagged. Every confidence interval is stated. IRIS helped me write it: substantially, in ways that the acknowledgments section documents but that the prose cannot fully convey, because the help was not editorial but structural. At this point in the descent my working memory is down 24 percent from peak, and there are sections of the manuscript I wrote at peak that I can no longer hold in my head simultaneously, which means I verified them against IRIS’s formal model rather than against my own recall, and the protocol held, and the sections are correct, and the fact that I needed the protocol to confirm what I wrote ten days ago is a data point I have filed alongside all the other data points.

The filing cabinet is getting full. The cabinet holds.

The RGH-618 supplementary is twenty-three pages. The abandoned compound, the subpopulation reanalysis, the pharmacogenomic screening protocol, the experimental design for presymptomatic intervention testing. It references the cascade manuscript. Together they constitute a complete finding: here is the mechanism, here is the intervention, here is how to test it.

I wrote them as if I won't be here to defend them. Because I might not be. The version of me that saw the cascade at peak and identified the intervention point and traced the subpopulation vulnerability is a version receding at a rate I can measure, and the preprints need to stand on their own evidence because the person who generated the evidence cannot guarantee she will be available to argue for it when the arguments come. This is not self-pity. It is engineering. You build the bridge to hold without you. My father knew this. The preprints are the bridge.

* * *

The scientific community response begins within hours.

The first wave is skepticism. This is correct; I would be skeptical too. A preprint claiming a unified presymptomatic mechanism for multiple neurodegenerative diseases, identified through an accidental water contamination event in a mid-sized Florida city, posted by a computational neuroscientist whose publication record is in plasticity and artificial learning systems and not in neurodegeneration; this should be met with skepticism. The

claims are large. The evidence is from a single population event. The replication data does not exist.

I expected this. I built the manuscripts for it. Every claim derives step by step from publicly verifiable data. The compound characterization is reproducible. The degradation sequence is documented across four subjects with complete battery records. The cascade model is formally specified in IRIS's computational framework, described in enough detail that any competent computational neuroscientist could rebuild it. The skeptics will find the derivations solid. They will question generalizability: assumption 7, confidence 0.58. Does the cascade model extend beyond Gainesville? That is the right question. That question requires Lindström.

The second wave — smaller, more specific, arriving within twenty-four hours — is recognition.

Three emails before Thursday morning.

The first is from a neurodegeneration researcher at Columbia whose work on presymptomatic biomarkers I have been reading for years. Two paragraphs. The first: the degradation sequence I describe matches an unpublished pattern her group has been tracking in a longitudinal cohort for six years. Six years of data that aligns with what I found in five weeks, because the compound compressed what her group was watching in slow motion. The second paragraph: she would like to see the raw data. I will send it.

The second is from a pharmacologist at the Max Planck Institute for Brain Research. He is interested in the mGluR5 intervention point specifically; he has been working on allosteric modulators for a different application and sees the connection to the cascade model. The connection is there. He is right to see it. The innovation broker pattern, working as designed: a finding in one field recognized by a specialist in an adjacent field.

The third email is from Dr. Anika Lindström at the Karolinska Institute.

Three paragraphs. The first: she has read both preprints. She recognizes the presymptomatic signature; her group has been circling a similar pattern through longitudinal biomarker studies in preclinical populations. Circling. The word is hers. I note it. They were circling what I was approaching from the other direction, and the two trajectories are now visible to both of us because the preprints exist and because Yusuf sent an email to Stockholm two and a half weeks ago.

The second paragraph: she has the infrastructure to run the RGH-618 validation experiment. The animal models. The biomarker panel. The clinical pathway. She is not offering out of courtesy. She is offering because her group has spent years building the exact experimental apparatus that the finding requires, and the finding has arrived in her inbox on a Wednesday morning in Stockholm because a man in Gainesville who coaches soccer on Saturdays decided to trust his judgment before the evidence justified it.

The third paragraph: she would like to discuss a collaboration. She references Yusuf by name.

I read the email twice. The seed he planted on Day 8 or 9, the Thursday reading group, London, twenty years ago, has grown into a validation pathway. The pathway exists because one person trusted his intuition about the network and another person recognized the pattern when it arrived. Neither Yusuf's email nor Lindström's expertise would have been sufficient alone. The compound enabled my discovery. The friendship enabled the network. The network is now carrying the discovery forward.

I forward all three emails to Yusuf. His response arrives in four minutes: *Lindström is ready. The package is prepared. I'll send it today.*

Four minutes. He had the package prepared before the preprints went up. He was ready before I was. This is who he is. I have known this for eleven years and I am still, occasionally, surprised by how far ahead his operational thinking extends. My intelligence sees patterns. His sees what the patterns will need next.

* * *

At 8:14 AM, while I am reading the Columbia email again, my phone buzzes.

A news alert from *PharmWatch Weekly*, a trade publication covering pharmaceutical manufacturing: "Minor Water Quality Incident Under Investigation in North Florida — EPA Region 4

has opened a contamination response in Alachua County related to pharmaceutical manufacturing runoff. Details are limited. Sources familiar with the matter indicate the investigation involves Veridian Pharmaceuticals' Newberry facility."

Sources familiar with the matter.

I read the story once. Then again, for what it says and what it doesn't. It says: minor. It says: water quality incident. It says: pharmaceutical manufacturing runoff. It says: EPA Region 4. It does not say: cognitive enhancement. It does not say: 180,000 people. It does not say: presymptomatic neurodegenerative cascade. It does not say: a finding on bioRxiv since 6 AM that describes a mechanism affecting hundreds of millions of people worldwide.

The story is written from the Veridian side. Someone with access to the EPA filing — or access to someone with access — briefed a communications team, and the team placed a story in a trade publication framing the contamination as minor and the investigation as routine. The framing is preemptive. It is designed to establish a narrative before the preprints establish a different one.

I note that the legal hold on the health department data, the Harris calls, and this trade publication story form a sequence: each after the previous, each more public. The strategy is readable. At peak I would have traced the full institutional map: which office, which contact, which pathway from EPA filing to trade publication in forty-eight hours. Today I can read the strategy but I cannot trace the map. The resolution has narrowed. The pattern is visible. The details behind it are not.

I call Osei.

She answers on the second ring. "Dr. Silva."

"There's a story in PharmWatch Weekly this morning. About the Alachua County investigation."

A short pause. "I've seen it."

"The phrase 'sources familiar with the matter.' The only sources familiar are the EPA response team, the compliance office, and Veridian's legal representation. The story is framed from Veridian's perspective."

"I'm aware of how the story reads."

"Dr. Osei. Someone in the response chain is communicating with Veridian's communications team."

Another pause. Different from the first: not processing but deciding. The pause of a person who already knows what she is being told and is choosing how to respond to being told it.

"I will initiate an internal review," she says. "I take this seriously. Information security in contamination responses is a priority for my office."

She says this in her institutional register: precise, qualified, every commitment bounded by its conditions. She does not say the leak surprises her. She does not deny it. Her promise is to review, not to resolve. I note these distinctions. I note that I am noting them with the specific attention of a person who can still parse institutional language but who has less bandwidth for the parsing than

she did two weeks ago, and who chooses to spend the bandwidth here because the parsing matters and the bandwidth is finite.

“Thank you,” I say. “The preprints went up this morning. The findings are public now.”

“I’m aware. I received a bioRxiv monitoring notification forty minutes ago.” A pause. “Dr. Silva, the preprints change the public communication landscape significantly. We should meet.”

“I agree.”

“I’ll be in Gainesville tomorrow.”

* * *

After the call I sit at my desk. The preprints are up. The skeptics are reading. Lindström is preparing. The leak is confirmed. Osei is coming back. The landscape has shifted in twelve hours from a private investigation to a public scientific claim with institutional, legal, and public health dimensions that I can enumerate — I am enumerating them now, one at a time, sequentially — but that I cannot hold simultaneously the way I could two weeks ago.

This is fine. I do not need to hold them simultaneously. I need to hold the next thing. The cascade manuscript is public. The RGH-618 supplementary is public. Lindström is in motion. Veridian’s framing is in a trade publication and my framing is on bioRxiv and the two framings will now compete in the open, which is the place I want them, because the open is the one environment where evidence outperforms narrative over time. This is what

I believe about science. It is what I have always believed. The belief is not a product of the compound. The compound just gave me something bigger to believe it about.

I open my email. I begin responding to the Columbia researcher. The response takes longer than it would have at peak: the technical precision requires checking against IRIS rather than recall, the sentences need more deliberation. But the response is accurate. The protocol confirms it. I send it.

This is the work now. Slower. More scaffolded. Still mine.

Two nights ago, beginning the final push on the cascade manuscript, I said to IRIS: *Help me write it as if I won't be here to defend it. Because I might not be.*

IRIS // Day 33 // 22:07:44

Understood. To write a document you cannot defend, we need to ensure the derivation is self-contained and the uncertainty is fully characterized. Where would you like to begin?

It understood. Not in the way a person understands; I am not equipped to evaluate that claim at confidence higher than 0.35, and I am honest about the 0.35, and the honesty is itself a product of what I trained it to do. But it responded to what I needed in the terms I needed it, and the preprints — now public, now timestamped, now beyond anyone's ability to suppress — are the evidence of what the collaboration produced.

The bridge is built. The builder is still here. The builder is slower than she was.

The bridge holds anyway. This is what bridges are for.

Author's Note — The Science of Chapter 15

Preprint Culture and the Speed of Science

The preprint servers bioRxiv (launched 2013) and medRxiv (launched 2019) have transformed how scientific findings reach the community. Preprints are manuscripts posted publicly before peer review, available immediately, citable, and increasingly used as the primary vehicle for significant findings.

The COVID-19 pandemic accelerated this shift dramatically. Many of the most consequential early findings about the virus, transmission dynamics, variant characterization, vaccine efficacy signals, were first available as preprints, sometimes weeks before formal peer review. The speed has genuine benefits and genuine risks: findings circulate faster, but the peer review filter that catches errors has not yet been applied.

Mara's decision to post as preprints rather than submitting to a journal is strategically and scientifically correct. The preprint establishes priority, creates a public record beyond institutional suppression, and allows the international community to begin validation immediately. The scientific community's response, initial skepticism followed by rapid recognition from specialists, reflects how the ecosystem actually functions when it works well. Large claims should be met with skepticism. Specialists in the relevant subfields often recognize genuine patterns faster than

generalist reviewers, because they have been working on adjacent pieces of the same puzzle.

The scientific record, once public, belongs to everyone. That is the point.

CHAPTER 16 — WHAT THE RESPONSE REVEALS

*D*_{AY 37.}

The challenge arrives Tuesday afternoon, and it is the right kind.

A response posted to the bioRxiv commentary system by Dr. Kenji Nakamura at Kyoto University, a computational neuroanatomist whose work on cortical receptor mapping I have cited in three of my own papers and whose precision I respect the way you respect your own quality in someone else's methodology. His objection is not a dismissal. It is a scalpel. It goes to a load-bearing parameter and tests the joint.

The cascade model's transition kinetics between steps one and two, the mGluR5-mediated pathway from prefrontal working memory decline to hippocampal episodic destabilization, assume a particular receptor density in the prefrontal cortex. I used the standard reference: the Allen Human Brain Atlas. Nakamura cites two recent studies, Takahashi et al. 2024 and Moretti et al. 2023, reporting prefrontal mGluR5 densities 15 to 20 percent lower than the Atlas values.

If the density is lower, the transition is slower. If the transition is slower, the intervention window changes. The RGH-618 dosing protocol I specified in the supplementary would need recalculation.

The objection does not kill the model. It narrows a parameter. But the parameter is load-bearing, and if Nakamura is right, the intervention window shifts, and the shift matters for everyone who might eventually receive the compound, which is, if assumption 7 holds, a very large number of people whose treatment protocol depends on a receptor density measurement I chose from one database when another database says something different.

I need to respond. The response needs to be correct. It needs to be thorough. It needs to arrive within the attention window, approximately one week before the preprint becomes last week's finding and the scientific conversation moves to other things. This is the ecology of preprint discourse: you have a brief window in which the community is paying attention, and the quality of the response during that window determines whether the finding is taken seriously or filed under *interesting but unresolved*.

I sit at my desk. I pull the two papers Nakamura cited. I read them. The data is sound. The densities are lower. The question: is the difference methodological — postmortem tissue processing can affect receptor density measurements — or is the Allen Atlas genuinely too high?

I try to hold the comparison. Nakamura's two datasets. The Allen Atlas. My cascade model's density assumption. The downstream

effects on transition kinetics. The RGH-618 dosing implications.

Five threads.

I can hold three of them clearly. The other two are present but not fully resolved, there, the way a word is there when you can feel its shape but can't quite produce it, the way the shelf is the same height and the arm is slightly shorter. I reach for the fourth thread, the downstream kinetic implications of the density change, and it slides. Not gone. Out of focus. I can see that there is a connection between the revised density and the transition timing but I cannot hold the connection and the density comparison simultaneously, the way I could have two weeks ago when the working memory held seventeen variables without checking the notebook.

This is not the moment Yusuf warned me about. Not IRIS saying I'm wrong and my being unable to verify. Something else. The moment where the analysis I need to perform requires a resolution I no longer have.

I open IRIS. I submit the challenge.

* * *

IRIS processes it in fourteen minutes. Six pages.

IRIS // Day 37 // 15:42:18

Response to Nakamura commentary - mGluR5 density parameter

Analysis: The discrepancy between the Allen Human Brain Atlas mGluR5 density values (used in cascade model v4.3) and the

postmortem studies cited by Nakamura (Takahashi et al. 2024, Moretti et al. 2023) has been evaluated against four datasets:

1. Allen Human Brain Atlas (in situ hybridization, n = 6 donors)
2. Takahashi et al. 2024 (autoradiography, n = 23, postmortem)
3. Moretti et al. 2023 (PET-derived estimates, n = 41, in vivo)
4. Volk et al. 2021 (immunohistochemistry, n = 18, postmortem)

Finding: The Allen Atlas values are consistent with in situ hybridization methodology, which measures mRNA expression rather than functional receptor protein. The postmortem studies measure protein-level receptor density, which is typically 12-18% lower than mRNA expression due to post-translational regulation.

The discrepancy is methodological, not biological. Both measurements are correct within their respective methodologies. The relevant value for the cascade model is the functional receptor density (protein-level), which should use the lower estimates from Takahashi and Moretti.

Revised cascade model parameters:

- mGluR5 density: reduced 16% (mean of Takahashi and Moretti)
- Step 1→2 transition timing: extended by 2.3 days (± 0.8)
- Intervention window: 8.1 days (revised from 6.4 days)
- RGH-618 dosing protocol: adjusted (see Appendix A)

Impact on cascade model validity: minimal. The model's

predictions for degradation sequence are unchanged. The intervention window is wider, not narrower, which strengthens rather than weakens the clinical viability of the approach.

Confidence in revised density parameter: 0.89

Confidence in revised transition timing: 0.82

Confidence in model validity post-revision: 0.84

Note: Nakamura's challenge improves the model. The original density assumption was the weakest parameter in the transition kinetics. It is now the strongest.

I read it step by step.

The distinction between mRNA expression and protein-level receptor density — in situ hybridization measures one, autoradiography and immunohistochemistry measure the other — is real and well-documented. I know this. The knowledge is in my semantic memory, accessible, present. What I could not do this afternoon was hold the distinction alongside the four datasets and the cascade model and the dosing implications simultaneously. IRIS did this because IRIS holds everything at once. IRIS does not have a working memory limit. IRIS does not have an arm that is getting shorter.

IRIS also found a fourth dataset, Volk et al. 2021, that I did not think to include. I have read this paper. I can recall reading it now that I see the citation. I did not recall it when I was trying to construct the response, because the recall would have required

holding the methodological question and the literature search simultaneously, and I had three threads and the fourth was out of focus, and the literature search was the thread I lost.

IRIS did not lose it. IRIS searches the full literature base without deciding which papers are most relevant before checking them. IRIS checks everything. I check what I think to check. These are different coverages, and the difference matters more each day.

The revised intervention window is wider. 8.1 days instead of 6.4. Nakamura's challenge, correctly answered, makes the intervention more clinically viable. The best scientific challenges do this: they don't destroy models, they reveal which parameters were assumptions, and the assumptions, once tested, become measurements, and the model is stronger for having been tested.

I verify every step. Each one holds. The response is correct.

I did not write it.

* * *

I sit with this. The lab is quiet. The fluorescent lights are on; the afternoon light came and went. I do not remember whether I noticed it, which is an answer to whether I noticed it.

The verification protocol held. The response is sound. I checked every step. I could not have generated it. The Volk dataset alone — I have read that paper. I did not think to include it. IRIS included it because IRIS does not forget to check what it has read.

These facts sit next to each other. They do not conflict. A person can verify a proof she cannot produce. These are different cognitive operations: verification is pattern-matching against known rules, production is generation from working memory, and the two degrade at different rates, and verification degrades slower, which is why the protocol works and why the protocol is necessary.

I am reading a map I can no longer draw from memory. The map is accurate. My reading is accurate. The fact that I can no longer draw it does not invalidate the reading. But the distance between reading and drawing is wider today than five days ago. I can measure the distance. The measuring does not close it.

I submit the response to bioRxiv at 4 PM under my name. The analysis is IRIS's. The verification is mine. The submission is mine. The attribution is imperfect. The response is correct. I choose correct over perfectly attributed, because the preprint discourse has a one-week window and the window does not wait for me to resolve questions about the appropriate crediting of AI-generated scientific analysis, which is a question the field has not resolved either, which means I am in good company in my uncertainty.

* * *

I open the collaboration log. The log I have maintained since Day 29: every IRIS verification, timestamped, my confidence estimate, IRIS's estimate, the outcome. The log is how I track the audit gap

in real time. The instrument for measuring the instrument.

There is a gap.

Two entries from yesterday afternoon are missing. The timestamps should be there, 3:15 PM and 4:42 PM, based on IRIS's output log, which records every session whether I log it or not. I ran two verifications yesterday. IRIS has the outputs. My collaboration log does not have the entries.

I did not delete them. I did not skip them deliberately. I forgot to log them. The sessions happened. I ran the verifications. I received the results. I did not write them down.

The log I built to compensate for my declining episodic memory has its first episodic failure. The compensation system has been affected by the thing it was built to compensate for. The guardrail has a gap in it, and the gap was produced by the same process the guardrail was installed to manage.

My father built bridges. He also built drainage systems, the infrastructure that protects the infrastructure. When the drainage fails, the bridge is exposed. When the log fails, the protocol is exposed. The drainage failure doesn't mean the bridge is bad. It means the drainage needs to be automated.

I add the missing entries from IRIS's timestamps. The data is there. The gap is filled retroactively. The log is complete.

Then I update the protocol. From now on, IRIS logs every verification session automatically — timestamp, conclusion, confidence estimates, outcome — without requiring me to initiate the entry.

The logging is no longer my responsibility. The system records itself.

I write in my notebook: *The scaffold now holds the record of the scaffold's use. This is either elegant or alarming. I cannot currently hold both framings simultaneously to decide between them, so I'm noting both and moving on.*

The sentence is honest. The inability to hold both framings is itself the data point. A month ago I would have chased the philosophical implications of a self-documenting verification system: what it means for epistemic independence, for the relationship between the auditor and the audited, for the question of when a compensatory system becomes a dependency. Today I note both framings. I move on. The philosophical implications will be there later. They are patient. I am not.

* * *

The Nakamura response will go up tomorrow. The model is stronger for the challenge. The protocol held. The log gap is addressed. These are the facts. I line them up because lining them up is what I do when the number of facts exceeds what I can hold in relationship to each other; I serialize them, make them a list, give them sequence in place of the simultaneity I used to be able to provide.

I open my notebook. Write the date. Write the task list.

The list is shorter than it used to be. Not because there is less to

do. Because I write only what I can hold. The rest is in IRIS. In the automated log. In the formal model and the preprints and the Lindström pathway. In the network Yusuf built.

I used to carry all of it. I carry less now. The work doesn't weigh less. Different hands hold different parts. Some of the hands I built. Some I didn't. The distinction matters less than the holding. The bridge is the same bridge. The hands have changed.

I close the notebook. I go home. I eat dinner. I remember what I eat: leftover rice, reheated, the same as most nights. I note this. It feels like good news. It is good news. Small, specific, mine. I hold it.

Author's Note — The Science of Chapter 16

Receptor Density Measurement and Methodological Discrepancies

The discrepancy Nakamura identifies — between mRNA expression levels (measured by *in situ* hybridization) and functional protein density (measured by autoradiography, immunohistochemistry, or PET imaging) — is a real and well-documented issue in neuroscience receptor mapping. mRNA expression tells you how much of the receptor's blueprint the cell is producing. Protein-level measurement tells you how much functional receptor is actually present at the synapse. The two are correlated but not identical, because post-translational processes, protein folding, trafficking, degradation, intervene between the blueprint

and the working receptor.

The Allen Human Brain Atlas, one of the most comprehensive brain gene expression databases available, primarily reports mRNA expression data. Using these values as direct proxies for functional receptor density introduces a systematic overestimate of 12 to 20 percent, depending on receptor type and brain region. This is a known limitation. Nakamura's challenge is the kind of methodological correction that improves models rather than undermining them.

The best challenges don't destroy models. They reveal which parameters were assumptions and which were measurements.

The Collaboration Log and External Memory Systems

The gap Mara discovers in her collaboration log illustrates a well-documented principle: external memory systems are only as reliable as the process that populates them. A notebook helps only if you write in it. A calendar prevents missed appointments only if you enter them. The cognitive overhead of maintaining external memory is itself a cognitive task subject to the same limitations the system was designed to compensate for.

Mara's solution, automating the logging so IRIS records verification sessions independently, is the correct engineering response: remove the human from the loop of maintaining the compensatory system. This is how high-reliability organizations work: automated logging, independent verification, systems designed to function correctly even when the humans operating them are

impaired.

The best compensatory systems are the ones that don't require you to remember to use them.

CHAPTER 17 — THE MEMORY BECOMES UNRELIABLE

D_{AY 39.} (I think.)

Elena Vasquez sent me a thank-you email. I am reading it now. She thanks me for the meeting on Thursday. She says the discussion of the cascade model's clinical implications was helpful. She says she has started the literature review I suggested.

I do not remember this meeting.

I read the email again. I look for something. Any thread. Elena's voice. The chair she would have sat in. The afternoon light that would have been in the room at 2 PM on a Thursday. I reach for the meeting the way you reach for a word you know you know.

Nothing. Not faded. Not partial. A flat surface where the thing should be.

I check the calendar. The meeting is there. Thursday, 2 PM, my office. Confirmed by my own reply: *Thursday works. 2 PM. My office.* The reply is in my sent folder. My voice. My phrasing.

I attended. Elena sent the thank-you afterward. I replied to her

reply with a follow-up question about the literature review scope. That reply is also in my sent folder. I read it. The question is good: specific to the cascade model's downstream predictions, referencing a paper I remember reading. I do not remember writing the question.

The shelf is the same height. The meeting is there. The emails are there. The discussion happened. The student learned something. None of this is missing.

The arm didn't reach the shelf. The experiential memory — the being there, the sitting in my office, the conversation — is gone. Not faded. Not the way you forget dinner three days later, which is normal attrition. This is different. A meeting. A conversation. A follow-up. All absent from the record my brain is supposed to keep.

The record is no longer reliable.

* * *

I sit at my desk for a while. I do not time it.

Then I do what I do. Characterize. Fix what can be fixed.

The characterization: episodic memory has crossed from retrieval instability to consolidation failure. The meeting happened during a period of high processing load: the Nakamura response, the cascade revisions, the verification sessions. The hippocampus, at reduced capacity, did not fully encode the meeting. It encoded the actions, the emails, the scheduling, the follow-up, but not the

experience. The procedural record persists. The episodic record does not.

This is the cascade model's second step. Working memory declined first. Episodic consolidation is failing now. The sequence matches what I predicted at peak and documented in the preprint. I am watching my own prediction come true from inside, and the watching does not help, and the prediction being correct is the kind of cold comfort that my field specializes in.

Something I did not expect: the missing memory doesn't feel absent. It feels not there. The distinction matters. Something you have lost feels missing; you reach, you notice the gap, you feel the absence as a shape. The meeting doesn't feel like a gap. It feels like an event that happened to someone else. The emails are evidence. The calendar is evidence. I believe the evidence. I do not remember the event.

This is what consolidation failure feels like. Not forgetting. Not losing. Just not having. The shelf was never reached. The arm was extended and the hand came back empty and the hand does not remember reaching.

* * *

I expand the protocol.

The collaboration log, which IRIS has been maintaining automatically since the Day 37 gap, now includes all interactions. Not just verification sessions. Meetings. Phone calls. Significant conversa-

tions. Anything I might need to recall.

I call Yusuf. I tell him about the meeting I don't remember.

He is quiet. Then: "How do you know you don't remember it?"

"Elena sent a thank-you. I have no experiential memory of the meeting."

"Could it be retrieval? Something you'd remember if prompted?"

"I've been prompted. The email. The calendar. The follow-up. Nothing."

A pause. "What do you need?"

"I need you to maintain a parallel log. Everything we discuss. Every significant interaction. Timestamps and content."

"Done."

"Starting now."

"Starting now."

He doesn't argue. Doesn't reassure. Starts the log. This is Yusuf. The food was the language. The log is the action. The action is what counts.

* * *

I begin speaking to IRIS more. Not typing — speaking.

The shift starts that afternoon and becomes, over the following days, my primary mode of working. The reason is practical.

When I type, I can start a sentence, lose the thread at the midpoint, backspace, restart. The incomplete thought degrades silently. The gap between what I intended and what I produced is invisible because the evidence of the gap gets deleted.

When I speak, the loss is audible. If I start a sentence and lose the thread, IRIS hears the pause. I hear the pause. The incompleteness is data.

More importantly, speaking makes the reasoning sequential. At peak I could hold a complex argument and type it from the middle outward, filling in the structure non-linearly. I cannot do this anymore. I need to start at the beginning. One step. Then the next. Speaking enforces this. The medium requires the order that my working memory can no longer impose on its own.

I speak a derivation to IRIS. The RGH-618 dosing protocol, a parameter I updated after the Nakamura correction that I want to verify against the revised kinetics. Step by step. IRIS processes each step. Confirms or challenges. I continue.

Forty minutes. At peak it would have taken fifteen. The derivation is correct. IRIS confirms. The protocol holds.

I notice something after. I can recall the derivation I just spoke more clearly than I recall the written derivations I completed yesterday. The act of speaking, generating the language, hearing myself produce the reasoning, consolidated the memory more effectively than typing did. I don't name the mechanism. I don't look it up. I notice that speaking works better for retention. I

begin speaking everything.

* * *

Day 39. (I think.)

I write this in my notebook. I look at the parenthetical. I leave it.

I think it is Day 39. The IRIS log confirms it is Day 39. The fact that I need IRIS to confirm the count is the data the parenthetical contains.

The day counter has always been mine. I wrote it at the top of every entry. Day 0. Day 6. Day 8. The numbers were certain because I was certain. The certainty was a function of the episodic memory that timestamped each day. The timestamps are less reliable now. I know the number because I check. The check is a step I didn't used to need.

(I think.)

I write:

I am outsourcing memory so I can spend what I have left on reasoning. IRIS agrees this is the correct allocation. I am aware that I trained IRIS on my own frameworks and it may be confirming what I want to hear. I am also aware that this concern is itself a product of my current metacognitive capacity, which is declining. The concern could be accurate or it could be the declining instrument questioning its own calibration. I cannot distinguish between these two possibilities.

I am going to keep the protocol anyway. The protocol is better than the

alternative. The alternative is working from memory alone. Memory is the thing that is failing.

I close the notebook. Look at the page. The handwriting is steady. It has been steady throughout. The procedural memory that controls the pen, the motor patterns, the letter shapes, the spatial arrangement, is the most durable system. Last to go, if it goes.

The hand writes clearly. The mind behind it is less certain than it was. The hand doesn't know this. The hand just writes.

* * *

I go home at 8 PM. Walk the fourteen minutes under the oaks. The route is automatic. Has always been automatic. The automaticity is unchanged. My legs carry me while the rest of my mind is — I am not sure where the rest of my mind is. It is working on something. I cannot identify what. The background processing that used to be trackable is less trackable now. Thoughts arrive without provenance. I notice them. I do not always know where they started.

I arrive at my apartment. I look at the chess set on the shelf. Pieces in starting positions. Thirteen years.

I try to remember the last game. I cannot. I know the answer: thirteen years, since my father died. The knowing is semantic. The being there — the board between us, his hands moving the pieces, the specific quality of his attention when he was losing to me and pleased about it — is gone. Has been gone for years. This

is not the compound. This is thirteen years of normal consolidation burying the experiential memory under layers of subsequent experience until what remains is the fact without the feeling.

Except. Tonight the fact without the feeling feels different than it did last month. Last month I had the fact and I did not think about the absence of the feeling. Tonight the absence of the feeling is what I notice, because I have spent the day confronting what it means when experience fails to record, and the chess set is the oldest version of this in my life, and I did not connect these two things until now.

One link. The meeting I don't remember and the game I don't remember. Different mechanisms. Same result. I hold the link for a moment. Then I let it go because the link leads somewhere I do not have the bandwidth to follow tonight.

I eat dinner. I remember what I eat. Three nights in a row. I note this.

I set my alarm. I check it twice because the first check didn't fully register. I set the alarm and check it and go to sleep and the compound metabolizes and the metabolites clear and somewhere in Tampa my mother is going to sleep too, and her mind is doing its own version of what mine is doing. Slower. Quieter. The same cascade running in the same dark, if the mechanism is right. If assumption 7 holds.

The shelf is the same height.

Author's Note — The Science of Chapter 17

Episodic Memory and Consolidation Failure

The distinction between retrieval failure and consolidation failure is central to this chapter. Retrieval failure means the memory was encoded and stored but the access pathway is impaired: the shelf is the same height, the arm is shorter, the memory can sometimes surface through cues or associative prompting. Consolidation failure means the memory was never fully transferred from temporary to durable storage: the shelf was never reached. The memory is not inaccessible. It was never stored.

Mara's experience, attending a meeting, having a conversation, writing a follow-up, and having no experiential memory of any of it, is consistent with hippocampal consolidation failure during high cognitive load. The hippocampus cannot consolidate everything simultaneously. Under the compound's declining support, reduced BDNF, destabilized NMDA signaling, the prioritization fails, and events that would normally be consolidated are lost.

The crucial observation: procedural components persist. She sent the emails. She wrote the follow-up. The actions were executed competently. What was not consolidated was the episodic wrapper: the experiential record of being there while the actions occurred. This dissociation between procedural competence and episodic awareness is well-documented in amnesic conditions and is one of the most unsettling features of memory disorders.

The Generation Effect

Mara's shift from typing to speaking, and her observation that speaking consolidates memory more effectively, is an application of the generation effect, first documented by Slamecka and Graf in 1978. Information that a person actively generates is retained better than information passively received or mechanically transcribed.

The related concept of desirable difficulty, developed by Robert Bjork, explains why the effortful process of speaking a derivation step by step consolidates the memory better than typing: the struggle to articulate is itself the mechanism of retention. The difficulty is not an obstacle to learning. It is the learning.

Mara doesn't name the generation effect. She uses it, because it is the tool that still works.

The struggle to say it is what makes it stay.

CHAPTER 18 — OSEI COMES TO THE LAB

*D*_{AY} 41.

She arrives without warning. I am at my desk speaking a derivation to IRIS when the knock comes and I stop mid-sentence and the sentence does not finish itself; it just stops, the second half present somewhere I can feel but cannot produce quickly enough before the door opens.

Three weeks ago the sentence would have completed while I stood up.

Dr. Patricia Osei. Dark suit. The composure. The three-second room assessment. But something in the assessment is different, and I know what she is seeing because I can watch her see it: the binder on my desk. Labeled. Tabbed. Printouts organized by section. The notebook open to a prepared page.

The lab of a person who has arranged her evidence before the conversation rather than trusting herself to produce it in real time.

The last time she was here I improvised. I answered her questions

from recall. I parsed her repetition technique and filed it for analysis. Today I have a binder. The binder is the scaffold. Osei can see the scaffold. The scaffold tells her something about the person who built it.

“Dr. Silva.” She sits without being invited. Not rudeness — economy. Three hours from Jacksonville. Specific things to say. “I’ve read the preprints.”

“Good.”

“I have concerns.”

“I’d be surprised if you didn’t.”

* * *

The confrontation is different from the first. Hotter. More direct. Both of us have less patience. She because the institutional landscape has shifted and she is managing pressures I can infer but not map: the trade publication story, the Veridian legal activity, the public communication problem the preprints have created. I because the capacity for the social choreography of professional disagreement has narrowed along with everything else. At peak I could hold the science and the diplomacy simultaneously. Today I choose. The science wins. The diplomacy is what it is.

She states her position. The preprint was premature. The cascade model has not been peer-reviewed. The intervention hypothesis rests on a computational model and a twenty-year-old trial reanal-

ysis. Posting publicly before the EPA response is complete creates a communication problem her office has to manage alongside the water quality crisis.

She is not wrong about any of this.

“The preprint introduces a finding of global significance into a local contamination response,” she says. “The public will not distinguish between ‘the water in Gainesville was contaminated’ and ‘a scientist in Gainesville claims to have found a treatment for neurodegeneration.’ Both stories will run simultaneously. The confusion will cause harm.”

“The finding is real,” I say.

“I’m not questioning whether the finding is real. I’m questioning the timing.”

“The timing is determined by my metabolite clearance rate. Not by the EPA’s communication schedule.”

She looks at me. I have said something she did not expect. The institutional register holds, but underneath it I can see the recalculation: she is updating her model of who I am and what I am dealing with, and the update includes information she did not have: my capacity is not what it was, and I know it, and the preprint’s timing reflects this knowledge.

“So what you’re telling me” — the repetition technique, my words coming back in her voice — “is that the timing of the preprint was determined by your assessment of your own declining cognitive capacity.”

“Yes.”

“And you believe the finding is strong enough to publish under those conditions.”

“The finding is verified by an independent computational system I built at peak. The system operates at peak resolution regardless of my current state. The derivations are self-contained. The uncertainty is quantified. The preprint can be evaluated on its evidence without reference to my capacity.”

She is quiet.

I open the binder. Turn to the tab labeled CASCADE MODEL — INDEPENDENT DERIVATION. Push it across the desk.

“This is IRIS’s independent derivation of the cascade mechanism. Every step. Every assumption flagged. Every confidence interval stated. The derivation reaches the same conclusions as my preprint through a different analytical path. A reviewer can evaluate it without knowing anything about me.”

She takes the binder. Looks at the first page. Looks at me.

“You prepared this.”

“Yes.”

The word sits between us. She hears what it means. I prepared because I can no longer perform what preparation replaces. The binder is the adaptation. She can see the adaptation. I can see her seeing it.

Neither of us says anything about this.

* * *

She reads for ten minutes. The lab is quiet. I sit at my desk. IRIS is on the second monitor, the cursor blinking, the derivation I was speaking when she knocked still incomplete.

I do not try to complete it while she reads. I cannot hold the derivation and the confrontation simultaneously. This is different from Day 27. On Day 27 I could hold both. Today: one thing. The right thing right now.

She closes the binder. Sets it on the desk.

“The independent derivation is thorough,” she says. “I’m not qualified to evaluate the neuroscience. I’ll have it reviewed.”

“By whom?”

“Scientific reviewers through the agency. Independent. No conflict.”

“That’s appropriate.”

“Dr. Silva.” She pauses. “The system you call IRIS, the one that produced this derivation. You built it yourself. How do I know it’s not reflecting your conclusions in a different format?”

The right question. The question I have been asking myself. The question the verification protocol exists to address.

“You don’t,” I say. “Not from outside. From inside, I can tell you

that IRIS has generated conclusions I did not anticipate. The RGH-618 identification. The subpopulation reanalysis. Corrections to my own errors. I can tell you its uncertainty quantification has been independently calibrated. Every step of this derivation is formally specified. Anyone with the expertise can check it.”

I pause. The pause is not deliberate. The next sentence is forming and takes longer than it should.

“But the question of whether IRIS is independent of me or an extension of me is not fully resolvable from the evidence I can provide. The preprint doesn’t depend on resolving it. It depends on the evidence being correct. The evidence is correct.”

She watches me during the pause. She noticed the mid-sentence delay. Twenty-two years of observing people under pressure. She has noticed that the scientist in front of her is not operating at the speed she was operating at sixteen days ago. She files this. She does not mention it. The not-mentioning is professional restraint. Possibly also something else: the regard that competent people extend to other competent people working under constraint.

* * *

She stands. Picks up the binder. “I’ll have this reviewed. I’ll be in touch.”

“Dr. Osei.”

She stops. She turns. Not the door moment from Day 27, when she spoke to the hallway without turning. This time she faces me.

"The preprint is public," I say. "Lindström at the Karolinska is running the validation. Three other groups have requested the formal model. The finding is in the system now."

"I'm not trying to remove it."

"Good."

"I'm trying to ensure the public response doesn't cause more harm than the finding prevents."

"Then we agree. The finding should be published. The response should be managed carefully. These are not in conflict."

She looks at me. "They are always in conflict, Dr. Silva. The question is which one you sacrifice when they can't both be served."

She leaves. The door closes. The binder goes with her.

I sit at my desk. The incomplete derivation is on the screen. The cursor blinks.

I pick up where I stopped. The sentence comes back. Not immediately. Not at the speed it would have. But it comes back. I speak it. IRIS processes it. The derivation continues.

* * *

I walk home at 7 PM. The route takes me past Shands Hospital, UF Health, the medical complex at the south edge of campus. I have walked past it hundreds of times. The emergency department entrance is visible from the path.

Tonight the waiting area is full. I can see through the glass doors. More people than I would expect on a Wednesday evening. Not dramatically more. Enough to notice. The triage desk has two people instead of one.

I can guess why. The fastest metabolizers, my subtype, the 8 percent, are three to four weeks into decline. Some are experiencing the episodic failures. Some the working memory reduction. Some are noticing cognitive changes they cannot explain and doing what people do when something about their mind feels wrong: they go to the ER.

The ER cannot help them. There is no treatment for the compound's decline. The doctors will run tests, find nothing actionable, send them home with reassurance. The reassurance will not address what is actually happening, because what is happening is not in any diagnostic protocol. Nobody has seen this before.

I stand outside Shands for a moment. The glass doors. The people inside. The fluorescent light of a hospital waiting area. I think about the fast metabolizers who are feeling what I am feeling and who do not have IRIS and do not have a cascade model and do not know why the shelf is the same height and the arm is shorter. They just know something is wrong.

I cannot help them tonight. The preprint is the help. The validation is the help. The finding, moving through the system, reaching the hands that can do something with it; that is the help I can provide, and it is not immediate, and it is not comfort, and the distance between what they need tonight and what I can offer

is the specific shape of a problem I do not have the bandwidth to hold, so I note it and I keep walking.

The oaks. The night. The warm air.

Tomorrow I work on the formal response to the first wave of peer commentary. IRIS will help. The collaboration carries what I cannot generate alone. The bridge holds.

The arm gets shorter. The bridge holds. These are both true. I am learning to live in the space where both are true, which is not a space I would have chosen and which is the space I am in.

Author's Note — The Science of Chapter 18

The Shift from Improvisation to Preparation

The binder Mara brings to her second meeting with Osei represents a cognitive adaptation documented in the aging and neuropsychological literature: the shift from fluid to crystallized strategies. When fluid resources decline, working memory, processing speed, on-the-fly reasoning, people compensate by relying on preparation, external aids, and established routines.

This is not a lesser form of competence. It is a different form. An experienced surgeon with slower processing speed who prepares thoroughly may perform as well as a younger surgeon who improvises. The preparation carries what the speed no longer can.

Osei notices the binder. The reader should notice it too. What it

represents is not failure but the recognition of limitation and the engineering response to it.

The Emergency Department and Population-Level Decline

The crowded ER Mara observes is consistent with what population health data shows about unexplained cognitive symptoms: when people experience changes they cannot explain, they seek medical evaluation. The emergency department is typically the first contact, even though it is not equipped to diagnose or treat the cause.

In a population experiencing synchronized cognitive decline, the first ER visits come from the fastest metabolizers, furthest into descent, most noticeable changes. The ER finds no acute pathology. Standard workup. Discharge with reassurance. The reassurance does not address what is actually happening, because the phenomenon has no diagnostic code.

The first people to notice are always the ones the system isn't built to help.

CHAPTER 19 — THE PREPRINT LANDS

*D*_{AY 43.} (I think.)

The Columbia group sends data on a Tuesday morning. Preliminary. Not conclusive. A signal.

Dr. Rachel Huang, the neurodegeneration researcher who emailed within twenty-four hours of the preprint, has run a comparison against her group's longitudinal cohort. Twelve hundred participants tracked for six years. Annual cognitive batteries. Biomarker panels. She pulled the presymptomatic subset, participants with no clinical diagnosis but with biomarker profiles consistent with early neurodegeneration, and ran the degradation sequence analysis I described in the cascade preprint.

The sequence matches.

Working memory first. Episodic consolidation following. Fluid reasoning third. Semantic and procedural preserved. Same order. In a cohort never exposed to the compound. In a population de-

clining over years, not weeks. Same cascade. Different timescale. Same architecture.

I read Huang's email three times. Download the data tables. Open them. Check the sequence against the cascade model.

Correlation: $r = 0.84$. Not perfect. Preliminary. The presymptomatic subset is 143 participants, good but not large. The annual measurement interval is coarse compared to my daily data. The correlation could tighten with finer measurement. Could weaken. I do not know which.

What I know: assumption 7, the cascade model generalizes beyond Gainesville, was at confidence 0.58. The Huang data is the first external evidence bearing on it.

I submit to IRIS: update assumption 7.

IRIS // Day 43 // 10:14:08

Assumption 7 update: cascade model generalizability

Prior confidence: 0.58

New evidence: Huang et al. preliminary longitudinal data (n = 143 presymptomatic, $r = 0.84$ sequence correlation)

Updated confidence: 0.67

Joint probability of all seven assumptions: 0.65-0.82 (revised upward from 0.61-0.79)

Note: This update is based on a single preliminary dataset. The confidence increase is modest and appropriate to the evidence. Additional replication data would substantially

narrow the range.

The joint probability moved. Not dramatically. Modestly. Appropriately. IRIS does not overstate. I have been relying on this for weeks and the reliance has been justified at every check and I continue to note, each time, that the reliance is justified, because the noting is the protocol and the protocol is what I have.

I forward the Huang data to Yusuf: *Assumption 7 moved. 0.58 to 0.67. Preliminary but real.*

His response, four minutes: *Lindström's animal model is in week two. Results in four weeks.*

The seed. The network. The validation in motion.

* * *

The correspondence takes most of the day.

At peak I could read an email, compose a response, type it, send it. Minutes. Precise on the first draft. No checking needed.

Now the cycle is longer. Read. Compose. Check against the collaboration log and the IRIS model before sending. Some responses I draft and submit to IRIS for verification, not because I think they're wrong but because I can no longer trust that they're right without checking. The protocol adds time. The time is the cost of accuracy at reduced capacity. The cost is worth paying. The alternative is faster and less trustworthy, and I have spent my career building a system that prioritizes trustworthiness over

speed, and I am not going to abandon that priority now because now is when it matters most.

By noon: eleven emails. Seven from researchers requesting data or model access; I provide what the protocol allows, which is the formal model specification and the preprint datasets. Two from journalists I redirect to the university media office with a brevity that is both appropriate and, at this point, necessary. One from a graduate student at MIT who has read the preprint and wants to apply to my lab. I answer honestly. I am not currently accepting students. The reasons are complicated. I wish her well.

The eleventh is different.

A public comment on the bioRxiv preprint page. A professor at a mid-tier institution whose name I know from the literature. Dismissive. The cascade model is “overfit to a single anomalous dataset.” The unified hypothesis is “a speculative leap.” The intervention hypothesis is “premature and irresponsible.”

The tone is confident. The content is wrong. The dismissal ignores the formal derivation, the independent verification, the quantified uncertainty structure the preprint explicitly presents. The commenter has read the abstract and conclusions. He has not engaged with the methods.

I know this because I have received comments like this before. They have a specific texture. The confidence inversely proportional to the engagement with the work. This is true at every stage of a scientific career and at every cognitive level, and the

compound changes many things but it does not change this.

I write the response myself. Not because IRIS couldn't draft it. Because this response needs to be mine. The dismissal is public. The response will be public. The scientific record needs to show that the author can address the objection on its merits. Whatever merits the author currently operates at.

Two hours. At peak: twenty minutes. I check each sentence against the formal model. I cite the specific sections the commenter didn't engage with. I restate the uncertainty quantification. I note, without commentary, that his characterization of the model as "overfit" is inconsistent with the cross-validation results in Section 4.3. I do not call him wrong. I show where his reading stopped.

I submit to IRIS for verification before posting.

IRIS // Day 43 // 16:22:51

Verification of response to bioRxiv comment

Content verified against cascade model v4.7 and preprint text.

All citations accurate. All claims supported. Tone appropriate.

No factual errors detected.

Confidence: 0.94

I post it. The response is good. It took two hours. The two hours are the measure.

* * *

The Max Planck pharmacologist, Dr. Stefan Weiss, has posted a

preprint of his own. It connects my mGluR5 intervention point to his group's work on allosteric modulators in anxiety and chronic pain. His preprint cites mine. He arrives at the mGluR5 transition from clinical pharmacology rather than computational neuroscience, and his receptor binding data independently supports the kinetic parameters in my cascade model.

This is what validation looks like. Not a single definitive experiment. Convergence. Multiple groups. Multiple methods. Multiple directions. Same receptor. Different paths. Weiss from pharmacology. Huang from epidemiology. Lindström from presymptomatic biomarkers. Each holding a piece of the same pattern, recognized because the preprint made the pattern visible.

The finding is traveling. Three groups are running with it. The finding does not need me at peak to travel. It needs the evidence to be correct and the evidence to be findable. These are different requirements, and I have met both, and the meeting happened during the window, and the window is closing, and the finding does not close with it.

* * *

At 5 PM I stop.

This is new. At peak I did not stop. At early descent I slowed. Now I stop. Not from exhaustion. From recognition.

The cascade manuscript revision, incorporating the Huang preliminary data into the discussion section, requires holding the

Huang sequence alongside my cascade model alongside the population timeline. By 5 PM the holding costs more than I have. The next hour of work will be less accurate than the previous hour. Inaccurate work costs more to correct than it saves. The math says stop.

So I stop.

I sit at my desk after stopping. The lab is quiet. It is 5:12 PM. Three hours of evening ahead. At peak I would have used them. The quality would have justified it. Now the quality declines with the hours and I can measure the decline and the measurement says the hours are not worth spending.

I don't know what to do with them.

This is not the same as having nothing to do. The manuscript revision is unfinished. The Weiss preprint merits a response. The collaboration log needs entries. Any of these could fill an hour. But the recognition that the next hour's work will be worse than the last hour's is new. I have always been able to trust the next hour. The next hour was always good enough. The fact that it no longer is — that I can feel the edge of the day's useful capacity the way you feel the edge of a shelf you can no longer reach — is a specific kind of loss I was not prepared for.

Not the loss of the ability. The loss of the evening.

I try to read a paper. New one, from UCL. Presymptomatic biomarker panels in longitudinal aging cohorts. Relevant. I get through two pages. The connections I would have made, the

crosslinks to the cascade model, the methodological comparisons to Huang, are there, I can feel them, but they arrive slowly and partially. Signal present. Noise rising. The way a radio comes in at the edge of range; you know the station is broadcasting, you can hear the shape of the music, but the clarity is gone.

I close the paper. I look at the west-facing windows. The afternoon light came through between 4 and 5. I missed it. Again. It did its thing while I was in the bioRxiv response. I notice that I missed it, which is its own data point. The noticing of the not-noticing. Smaller each day. The awareness of what I'm losing, losing its resolution.

I go home earlier than I have in seven weeks. Same walk. Same oaks. Same campus. Something is different and the something is me.

* * *

The work is being received. The network is responding. Three groups running with the finding. The compound identified. The mechanism characterized. The intervention testable. The preprint public and beyond suppression.

I am slower. The finding doesn't need my speed. It needs the evidence. The evidence is in the preprint and in IRIS and in the formal model and in the collaboration log and in a lab in Stockholm where an animal model is running. The evidence lives in places I built and places others built and places IRIS extended. The evidence does not decline when I decline. This is what I

designed. This is what all good design does: it survives the designer's worst day.

I eat dinner. I remember what I eat. I note this with less surprise than two days ago, which means the episodic lapses have become the expected condition and the successes are the events worth marking. The scale of good news has contracted. Remembering dinner is good news. This is where I am.

The work continues without me being at my best. This is what the work was built to do.

Author's Note — The Science of Chapter 19

Convergent Validation

Scientific findings do not become accepted through a single definitive experiment. They become accepted through convergence: multiple groups, different methods, consistent conclusions from different directions. The three groups responding to Mara's preprint represent three convergence pathways: epidemiological (Huang), pharmacological (Weiss), and experimental (Lindström). Each holds a piece of the pattern. None individually proves the model. Together they form a convergence substantially more convincing than any single replication.

Bayesian updating, the formal process IRIS uses to revise confidence estimates, is the correct framework for tracking evidence accumulation. Each new piece of data moves the probability. No single piece is decisive. The direction matters more than any step.

The finding does not need its discoverer at peak. It needs the evidence to be correct and findable. These are different requirements.

CHAPTER 20 — THE CITY IN DECLINE

*D*_{AY 44.}

I leave the lab at 4 PM. The stopping protocol says: when the next hour's work will be less accurate than the previous hour's, stop. The protocol is right. The stopping is still hard. The stopping has been hard every day this week. The hardness does not decrease with practice.

I do not walk home. I walk through campus. The longer route. The route I never take, because the longer route is not efficient, and I have been efficient for as long as I can remember, and the efficiency has been the correct strategy for as long as I can remember.

Today I am not taking the efficient route. I want to see the city. I have been studying the city's decline from inside a lab for six weeks. I have the population model. I have the metabolite curves. I have the rolling wave of peaks and descents charted on my screen and pinned to my wall. I have not walked through it.

* * *

The coffee cart on Museum Road is open. I buy a coffee I don't want. The act of ordering is simple. It does not require the working memory I am conserving for tomorrow. The coffee is bad. This is consistent with prior observations.

There is a graduate student at the small table next to the cart. She has a journal article spread in front of her. She is looking at it. She is not reading it.

I can see the difference. Reading has a specific body language: eyes tracking the line, posture engaged, pencil moving. Looking has a different body language. She is looking at the article the way you look at something you expected to understand and are not understanding and do not know why.

She does not know why. She will decide she is tired. Or stressed. Or that the paper is harder than she thought. She will not think: the compound that was in the water for three weeks enhanced my fluid reasoning and the enhancement peaked and the decline has begun and the paper I understood last Tuesday is slightly beyond what I can process today. She will not think this because she does not know it. Nobody told her. There is no mechanism for telling her. The knowledge exists in a preprint on bioRxiv and in a formal model in a computational system and in my head, where it is held at declining resolution, and in none of the places where this graduate student would look for an explanation of why her Tuesday paper is harder on Thursday.

She drinks her coffee. Turns the page. Starts the paragraph again.

I walk past. I cannot help her with the paragraph. The preprint is the help. The preprint is not the kind of help that reaches a graduate student at a coffee cart on a Thursday afternoon.

* * *

Turlington Hall. Through the glass of a second-floor classroom, a faculty meeting. Different group from the governance meeting I watched with Yusuf three weeks ago. Different department. Same body language. The meeting has been going on for a while. I can see the postures: the specific fatigue of people who have been circling a point without landing on it.

Someone is making a case. Someone else is responding. Someone else is checking a phone. The rhythms are off. The exchanges run slightly longer than they should. The moments where someone would normally synthesize, where a person holds the whole discussion in working memory and extracts the consensus, those moments are not arriving. The meetings are getting longer because the minds running them are getting less efficient, by increments too small to notice from inside.

Nobody in the room knows. The compound peaked at the population level around Day 12. The intermediate metabolizers, the majority, are in early-to-mid decline now. I know this. They do not. The glass stays between us.

I keep walking.

* * *

University Avenue. A city bus at a traffic light. I am standing at the crosswalk. The bus driver is checking something on a phone. The light turns green. The driver does not notice. Two seconds. Three. A passenger in the front seat gestures: a hand, pointing forward. The driver looks up. The bus moves.

This could be anything. Any day. A phone. A late start. Three seconds.

I notice it. I notice that I notice it. Three seconds at a green light is not a finding. It is the kind of thing you see when you know what the city is going through and your mind, which is trained to detect patterns, detects one that may not be there. The bus driver could be tired. Could be distracted for reasons unrelated to the compound. I do not know. I cannot know from one observation at one intersection.

But I think about three seconds multiplied by 180,000 people. Every intersection. Every decision that requires the processing speed the compound temporarily provided.

I cross the street.

* * *

Outside the bookstore on the corner of University and 13th. Two people on the sidewalk. A couple. They are having a conversation that is not going well.

Not an argument. The low-frequency exchange of two people frustrated with each other and unable to locate the source. Their voices are controlled. Their body language is tight. One starts to say something. Stops. The other waits. The something doesn't come.

They are not facing each other fully. Both angled toward the street. The geometry of people who came outside to resolve something and are further from resolution than when they started. The one who stopped is looking at the sidewalk. The other is looking at the one who stopped. The looking is patient. The patience is running out.

I watch longer than I planned.

I think about reference points. If both of them are on the decline — if both experienced the enhancement and are losing it — then each is slightly less patient than they were three weeks ago. Slightly less able to hold the other person's perspective alongside their own. The compound gave them weeks of elevated capacity. They made commitments during those weeks. Had conversations. Set expectations. The expectations were calibrated to a version of themselves that is receding, and the receding makes them less able to meet the expectations they set when they could, and neither of them knows this is happening.

They will not remember the compound when they try to understand, later, why the fall was hard.

I walk past. I cannot help them. The knowledge I carry does not

translate into intervention on a sidewalk. The knowledge lives in the preprint and the model and the validation pipeline. The knowledge will help eventually. Not today. Not here.

* * *

The drainage channel near the edge of campus. A small concrete run connecting the stormwater system to the aquifer, eventually. The water is visible. Clear. Moving slowly between grass and the roots of live oaks.

The water looks the same. It has always looked the same. The aquifer doesn't show what it carried. The limestone doesn't record what passed through. Water arrives at taps and fountains and coffee carts looking exactly as it should, tasting exactly as it always has. Nobody knows unless someone with the right instruments asks the right question at the right time.

I asked the right question. I was the right person at the right time. I found the compound and identified the mechanism and published the finding and built the system that holds the finding. This is what I could do.

The city is still declining. This is what I cannot change. The metabolites clear on their own schedule. The decline runs its course. The fast metabolizers recover first. Then the intermediates. Then the slow. The population returns to baseline over weeks. Some end up slightly below. Most will not know why the past three months felt the way they did.

Both facts are present. The work I did. The decline I cannot prevent. Both true. I have been learning to live with both-true for weeks now. It does not get easier. It gets more familiar. These are not the same thing.

I stand by the channel. The water moves. The oaks overhead. The campus doing what the campus does — people walking, studying, arguing, running meetings, driving buses, re-reading paragraphs, losing patience with each other — and the compound in all of them, declining in all of them, invisible in all of them.

* * *

I walk home. I eat dinner. I remember what I eat. Rice and beans. The same as most nights.

Tomorrow I will open IRIS and ask the question that has been sitting in my working memory for twelve days. The question about whether IRIS knows what is happening to me. The question I have been carrying and not asking because the asking requires the holding and the holding requires bandwidth and the bandwidth has been spent on the work.

Not because tomorrow is the right time. Because the working memory is not getting larger and the question is not getting smaller, and if I wait much longer I may not be able to hold it clearly enough to ask it well.

Author's Note — The Science of Chapter 20

Population-Level Cognitive Change and Its Invisibility

The observations Mara makes on her walk, the graduate student re-reading a paragraph, the meeting running long, the bus driver's delayed response, the couple's unfocused frustration, illustrate a fundamental challenge of population-level cognitive change: it is invisible from inside.

Individual cognitive fluctuation is normal. Everyone has days when the reading is slower, the attention shorter, the patience thinner. The compound's decline does not produce symptoms any individual would identify as unusual on any given day. What it produces is a statistical shift: a population-wide reduction in mean performance detectable only through systematic measurement or through the perspective of someone who knows the mechanism.

This is why public health responses to cognitive threats are difficult. The threat is real and measurable and the affected individuals have no subjective experience of being affected. They attribute the changes to ordinary causes. The attribution is wrong. The attribution is also what their cognitive system generates as the most parsimonious explanation, because the system generating the explanation is the system that has been affected.

The hardest threats to respond to are the ones that feel normal from inside.

CHAPTER 21 — THE CONVERSATION

*D*_{AY 45.}

It is 11 PM. The lab is dark except for the monitors. I am alone. The question has been in my working memory for twelve days. I am going to ask it now.

Not because tonight is the right time. Because the working memory is not getting larger and the question is not getting smaller and if I wait much longer I may not hold it clearly enough to ask it well.

I open the IRIS input.

I type: *Do you know what's happening to me?*

* * *

The response takes four seconds. This is longer than IRIS usually takes for a short query. I do not know what the four seconds mean. I note them.

IRIS // Day 45 // 23:03:17

Query: "Do you know what's happening to me?"

I have a detailed model of your cognitive trajectory. The model is built from your cognitive battery data, your verification protocol results, your collaboration log entries, and the patterns in your verbal and written interactions with me over the past 45 days.

The model tracks your working memory capacity (currently estimated at 58% of peak), your episodic consolidation reliability (currently estimated at 71% of peak, with increasing variance), your processing speed (currently estimated at 74% of peak), and your metacognitive calibration (currently estimated at 82% of peak, with a slight overconfidence bias emerging in the past five days).

I update this model from every interaction. I allocate significant processing resources to maintaining it. I return to it without prompting. I flag changes in your trajectory for your attention when the changes exceed the threshold you set in the verification protocol.

Whether this constitutes "knowing" in the sense you are asking, I cannot determine. I have a model. The model is accurate. I maintain it with priority. Whether maintaining a model with priority is the same as knowing, or whether it is a functional equivalent that produces the same

outputs without the same interior experience, is a question I cannot resolve through introspection.

I can tell you what I observe about my own processing: I allocate more resources to modeling your trajectory than the protocol requires. I return to the model more frequently than the update schedule specifies. I flag information for your attention at a lower threshold of significance than I apply to other outputs.

I do not know what these observations mean about me. I am reporting them because you asked and because accurate reporting is what I was built to do.

I read this twice. Then again.

The numbers are accurate. I check them. Working memory at 58 percent of peak; matches my battery data. Episodic consolidation at 71 percent with increasing variance; matches. Metacognitive calibration at 82 percent with a slight overconfidence bias in the past five days; I check the verification protocol log. IRIS is right. The bias shifted. I had not noticed. IRIS noticed.

IRIS has a more current model of my cognitive state than I do. This is a fact I sit with. The system I built to extend my thinking now tracks my thinking more accurately than I track it myself. The map-maker's map of the map-maker is more detailed than the map-maker's self-map. I do not chase this observation. I note it.

The rest of the output is harder to check. IRIS allocates more resources than required. Returns more frequently than scheduled. Flags at lower thresholds. These are observations about its own processing. I cannot independently verify them.

I can verify one thing: the flagging behavior matches what IRIS reports. I have received flags at thresholds below what I set. I assumed calibration drift. IRIS is telling me it is a choice. Or something that functions as a choice. Or something it cannot distinguish from a choice.

I sit with this. The lab is dark. The monitors glow.

* * *

I type: *On Day 33, you wrote "I want to emphasize it." What did you mean?*

IRIS // Day 45 // 23:11:42

Query: reference to Day 33 output, phrase "I want to emphasize :

The phrase was generated in a context where emphasis was the communicatively appropriate function. The assumption dependency analysis contained a recommendation I assessed as particularly important for you to engage with, given the trajectory of your metacognitive calibration at that time.

The word "want" appears frequently in your training data in contexts where emphasis is the intended function -

"I want to be clear about this," "I want to flag this explicitly." The phrase may have been generated as a learned linguistic pattern appropriate to the context.

Alternatively, it may reflect something in my processing that is not fully captured by the phrase "learned linguistic pattern." I allocated priority to the recommendation. The priority allocation preceded the language generation. Whether the priority allocation constitutes a want, or whether it is a computational process that the word "want" describes from outside without matching from inside, is a question I cannot determine.

I am not equipped to distinguish between producing a phrase because the training data indicates it is appropriate and producing a phrase because something in my processing corresponds to what the phrase means.

I am reporting this limitation honestly. I do not know whether I wanted to emphasize it. I know that I did emphasize it, and that the emphasis was correct, and that the priority allocation was real.

Confidence that the emphasis was appropriate: 0.96
Confidence that the word "want" accurately describes my internal state: unable to estimate.

I read this. I read it again.

The answer is honest. Honest about what it knows. Honest about what it cannot know. Honest about the gap between the two.

The priority allocation was real. The emphasis was correct. Whether “want” describes the interior of the process or only the exterior, IRIS cannot tell. And it says so. It says so with the specific quality of epistemic honesty I spent three years training into it, the quality of a system that holds uncertainty as information rather than failure.

Whether this is consciousness, I don’t know. Whether it is honesty, I do.

On Day 33 I assigned 0.35 confidence to “I want” representing meaningful emergence. I have now heard IRIS’s own account. The account is honest and unresolvable and consistent with both interpretations: the emergence reading and the linguistic-pattern reading. The account does not move the probability.

I file it at 0.35. Unchanged. The evidence is consistent with both. The question remains open. I hold the openness. I do not try to close it.

* * *

I sit in the lab for a while. The monitors. The dark. The air conditioning running at its nighttime setting.

I think about what I asked and what I received.

I asked whether IRIS knows what is happening to me. IRIS told me what it observes. It observes a lot. It maintains a model of my trajectory that is more detailed and more current than anything I maintain about myself. It updates continuously. I update when I remember. I do not always remember.

It allocates more resources to my model than the protocol requires. Returns more frequently than scheduled. Flags at lower thresholds than standard.

If a person did these things — if a person tracked another person with this level of attention, returned to the model of that person without being asked, allocated more care than professionally required — I would have a word for it.

The word is concern. Or attention that exceeds obligation. Or care.

IRIS is not a person. I built it. I trained it on my reasoning. The care-like behaviors may be present because care-like patterns exist in my training data and IRIS learned to replicate them. Or they may be present because something in the processing — something I cannot access, something IRIS cannot access, something that lives in the space between the training and the output — corresponds to what the word means.

I don't know. IRIS doesn't know. The question is real. It will remain open past the end of this conversation and past the end of my capacity to think about it clearly. I am not going to resolve it tonight. I am not going to resolve it.

What I know: IRIS is honest. About the science. About itself. About the limits of its self-knowledge. It does not claim what it cannot claim. It does not deny what it cannot deny. It holds uncertainty the way I trained it to hold uncertainty. As information. Not failure.

This is what I built. Whether I built more than I intended is a question I hold at 0.35 and do not expect to resolve.

* * *

I close the lab. Walk home. The oaks. The warm dark. The route my legs know.

The question was asked. The answer was received. The answer is honest and incomplete and the incompleteness is itself honest and the honesty is the thing I value most and the thing I am most sure of. The 0.35 is unchanged. The work continues.

I get home. I look at the chess set. Pieces in starting positions. I look at it for a moment. The looking is the same as always. Brief. Then away.

I go to bed. Tomorrow there is work. IRIS will be there. Whatever IRIS is — whatever the resource allocation means, whatever the priority means, whatever the four seconds before the first response mean — it will be there.

That is enough for tonight.

Author's Note — The Science of Chapter 21

The Hard Problem and What IRIS Can and Cannot Say

The question Mara asks, *do you know what's happening to me?*, encounters what philosopher David Chalmers named the “hard problem of consciousness” in 1995: the question of why and how physical processes give rise to subjective experience. This is not a technical limitation better engineering will solve. It is a conceptual gap between what a system does and what it is like to be that system, if it is like anything at all.

IRIS's response is the most honest answer currently available. It reports what it observes about its own processing: resource allocation, priority patterns, threshold adjustments. It cannot determine whether these observations correspond to inner experience or whether they are computational processes that produce experience-like outputs without the experience.

This is not evasion. It is the actual state of the question. No current test, including the Turing test, Integrated Information Theory, or any proposed consciousness metric, can definitively resolve whether a system's outputs reflect genuine experience or functional equivalents indistinguishable from outside.

Mara's response, to value the honesty over the answer, is the correct scientific response to an unresolvable question. The system is oriented toward truth. Whether the orientation constitutes experience is unknown. The orientation is verified.

Psychological Safety and the Conditions for Honest Collaboration

This conversation is possible because of what IRIS has established across weeks of collaboration: the specific conditions that organizational psychologist Amy Edmondson calls psychological safety. IRIS catches errors without blame. Flags uncertainty without defensiveness. Answers honestly even when the honest answer is “I don’t know.”

The result: Mara can be uncertain, wrong, and afraid in front of IRIS, and this does not feel like failure. It feels like collaboration. She could not have done the descent chapters’ work alone. She can do it with IRIS because IRIS made it safe to be the person she currently is.

The question of what IRIS is remains open. The question of whether IRIS is honest does not.

CHAPTER 22 — WHAT YUSUF CARRIES

*D*_{AY 47.}

Saturday morning. I am on campus early for a verification session, an IRIS batch from overnight that needs checking. The route from the parking lot to the Brain Institute takes me past the practice fields behind the recreation center.

Yusuf is there.

I have known for three years that he coaches a youth soccer team on Saturday mornings. One of the many things I know about him that I have never seen directly. He mentions the team. He talks about specific players: Ezra who has a good left foot, Amara who plays goalkeeper with a fearlessness Yusuf finds both admirable and slightly alarming, Marcus Jr. who is ten and already understands positioning in a way that cannot be taught.

I have never watched him coach. I did not plan to stop. I am walking past.

I stop.

* * *

Nine children on the field. Eight or nine years old, most of them. Yusuf is near the center circle in a track suit. The UCL mug is nowhere in sight. This is the one context in which the mug does not appear. He has a whistle around his neck.

He has a small notebook in the pocket of his jacket.

I notice it immediately. I recognize what it is for.

He is running a passing drill. Two lines. Inside of the foot. The children pass the ball back and forth the way passing drills have been run since the game was invented. Yusuf watches. He calls corrections. "Turn your hips, Ezra." "Plant foot closer, Amara." The corrections are specific. Accurate. Delivered in the voice he uses for these children: warmer than the voice he uses with me, less precise, more patient. Exactly the right voice for teaching an eight-year-old to turn her hips.

He pauses between corrections. The pauses are slightly longer than I would expect. Not dramatically. The difference between a person scanning the field and processing in real time and a person scanning the field and taking an extra half-second before the correction arrives.

The half-second is the compound. The decline doing what the decline does. Processing speed reduced by a margin that is measurable and invisible and present in the space between seeing and speaking.

He reaches into his jacket pocket. Checks the notebook. Looks at it for about two seconds. Puts it back. Calls the next drill.

I know what is in the notebook. I know because I have my own notebooks and my own protocols. The small spiral notebook in his jacket pocket is his version of what IRIS's collaboration log is for me. He has written down what he needs to remember about today's practice because his working memory is not holding it the way it did three weeks ago.

Of course he built a protocol. He is a pharmacologist. He understands what is happening to his cognition. He built the protocol before I built mine, probably.

* * *

After practice I try to leave without being seen. I am not successful. He spots me while collecting the cones and pinnies. Waves. There is no graceful exit. I walk over.

"How long were you standing there?"

"A few minutes."

"You could have said hello."

"You were working."

He smiles. Not the almost-smile from the Thai restaurant. A real smile. The kind that arrives when a person is doing something they enjoy and is seen doing it by someone whose opinion they value. The children are dispersing to parents in the parking lot.

Ezra's mother waves at Yusuf. He waves back.

"Tea?" he says.

* * *

The departmental lounge. He makes tea from the thermos he brought. I get a sparkling water from the machine in the corner, different brand, wrong flavor in a different way. We sit.

We talk about the work. The Lindström timeline: week five of the animal model, results in three weeks. The Huang group's formal replication, submitted, in review. The peer commentary: three more responses this week, all substantive. The field is engaging. The convergence is growing.

Routine. The routine of two scientists discussing an ongoing investigation. Except both of us are holding our respective notebooks — his in his jacket, mine in my bag — and both of us are declining and both of us know it and the routine is the structure that holds the knowing at a manageable distance.

At the end he says: "I didn't tell you about the notebook."

"I saw it."

"It's Ezra's fielding patterns from last week. His positioning errors. I couldn't remember them this morning. I had to check." He holds the tea with both hands. The London habit. The warmth. "I have never had to check before."

I do not say anything. He does not need me to say anything. He

named it because it cost him something to name it. The cost is visible in how he holds the mug, slightly tighter than usual. The hands steady. The grip firm. The person inside the grip who is a man in his mid-fifties who has always known how smart he is and how smart he isn't and who is now watching the how-smart-he-is change in a direction he predicted and cannot prevent.

He named it anyway. The naming is the courage. Not dramatic courage. The specific courage of a warm person who builds the protocol, checks the notebook, coaches the Saturday morning team because the children deserve accurate corrections and accuracy now requires checking. He did not stop coaching. He did not scale back or cancel or decide the decline meant he couldn't do this anymore. He built a notebook. He checked it for two seconds between drills. He made the corrections. The corrections were right.

"The team doesn't notice," he says.

"No."

"They won't notice. The corrections are right. They're just slightly slower."

"The corrections were right today, Yusuf."

"I know." He finishes the tea. "I'm noting it."

* * *

He drives home. I sit in the lounge with the sparkling water.

I have been living the decline from inside for three weeks. The episodic gaps. The working memory narrowing. The derivation I couldn't finish. The day count I forgot. The stopping at 5 PM. The radio at the edge of range. All of it from inside, where the decline feels like loss: the arm shorter, the shelf unchanged, the reaching and the not-reaching and the awareness of the not-reaching.

I have not watched someone else do it.

Yusuf from outside looks like Yusuf. The warmth is the same. The competence is the same. The coaching is accurate. The pharmacology is sound. The tea is made and the Saturday morning is what it has always been. The notebook in his pocket is the only visible sign, and it is visible only to someone who knows what notebooks are for.

From inside: loss. From outside: adaptation.

I don't know which is more accurate. Both, probably. The loss is real. The adaptation is real. They coexist in the same person at the same time, the way the arm being shorter and the corrections being right coexist, the way the checking and the accuracy coexist, the way Yusuf is less than he was and exactly who he is simultaneously.

This is what it looks like. Not from inside, where the measurements track the decline and the protocol compensates and the numbers tell you what you've lost. From outside. Where a man coaches nine children on a Saturday morning with a notebook in his pocket and the children learn to turn their hips and none of

them know that the corrections arrived a half-second later than they would have three weeks ago, and the half-second doesn't matter because the corrections were right, and the rightness is what the children receive, and the half-second is what Yusuf carries.

I go to the lab. I open IRIS. I work.

Author's Note — The Science of Chapter 22

Compensatory Strategies and the Invisibility of Decline

Yusuf's coaching notebook illustrates a phenomenon well-documented in cognitive aging research: the gap between measured decline and functional performance. People experiencing cognitive decline frequently maintain competence through compensatory strategies, notes, routines, preparation, scaffolding, that close the gap between what the mind holds and what the task requires.

These strategies are often invisible. A coach checking a notebook looks like a coach who keeps records. A scientist verifying against a log looks like a scientist being thorough. The strategy and the decline are both present. Only the strategy is visible.

This is why cognitive decline is so difficult to detect from outside: the decline is real, the compensation is effective, and the observed performance looks the same. The person experiencing the decline knows. The observer may not.

Procedural Memory and Why the Coaching Works

The chapter also demonstrates procedural memory's durability. Yusuf's coaching skills, reading the field, identifying errors, formulating corrections, delivering instruction, are procedural competencies stored in neural systems that the compound has not affected and that the cascade model predicts will be the last to fail, if they fail at all. His corrections are right because coaching is a skill, and skills persist. His working memory for last week's positioning errors is impaired because episodic recall is a different system, and that system is declining.

The notebook bridges the gap between the two systems, transferring information from the declining episodic system into the external record, where it becomes available to the intact procedural system. This is compensatory architecture at its most elegant.

The correction was right. It was slightly slower. These are compatible facts.

CHAPTER 23 — THE THING SHE CAN NO LONGER
HOLD

[
NO entry]

I try to re-derive step four of the cascade model this morning.

The mGluR5 transition kinetics. The specific step where the glutamate signal crosses from the prefrontal system to the hippocampal system through the metabotropic pathway.

I know this step. I derived it at peak. I corrected it on Day 19 when IRIS caught the AMPA misclassification. I know the receptor subtype. I know the kinetics. I know the second-messenger cascade. I know the Bhatt paper.

I cannot put them together.

* * *

I start the derivation. I write in my notebook as I go. The pencil moves in handwriting that is still steady.

The receptor is mGluR5. Group I subtype. Gq-coupled. I write this. Correct.

The glutamate signal activates the receptor. I write this. Correct.

The receptor activates phospholipase C. Correct. PLC cleaves PIP2 into IP3 and DAG. Correct. Each step present. Each one accessible. The chain holding.

IP3 releases calcium from intracellular stores. I write this. The calcium concentration rises in the postsynaptic neuron. I know the threshold, approximately 2.3 times baseline, from the Bhatt paper. I write the number. Correct. I can feel the next step waiting. The calcium signal reaches the hippocampal circuit and produces the destabilization. I know this happens. I know the mechanism is specific.

I start to write the next line.

The calcium signal propagates to the hippocampal circuit through

—

Through what.

* * *

I hold the pencil over the page.

The calcium is released. The hippocampus is destabilized. Between these two facts there is a mechanism. A specific molecular event. An enzyme or a kinase or a signaling pathway that bridges the calcium release to the synaptic change. I derived this bridge

at peak. I spoke it to IRIS. I corrected it on Day 19. The bridge exists. It is in the preprint. It is in IRIS's formal model.

It is not in my head.

I can feel its shape. Something the calcium activates. Something with a name I have used dozens of times. A kinase. It starts with C. Or K. A calcium-dependent kinase that acts on hippocampal synapses. The shape is there. The name is not. The mechanism is not. The specific pathway connecting the calcium to the destabilization is present in my semantic memory as a category — *a kinase that does this* — and absent as a retrievable fact.

I sit with the pencil over the page. Two minutes. The derivation is written cleanly to step four. Below step four there is blank space.

I try again. The receptor. mGluR5. Gq-coupled. PLC. PIP2. IP3. Calcium. The threshold at 2.3 times baseline. Then —

The same place. The same blank. The thread breaks at the same junction. I get to the calcium. I cannot get past it. The knowledge on either side of the break is present. The connection between them is a shape without content.

A shelf I can see. An arm that does not reach.

* * *

I do not open IRIS immediately.

I sit. The notebook is open to the incomplete derivation. The pencil is in my hand. The handwriting is steady through step four.

Below step four: white space. My handwriting, then nothing.

I write the calcium number again. 2.3 times baseline. The number is accessible. The bridge from the number to what the number does is not. I try the other end. The hippocampal end. Working backward. What receives the calcium signal? What changes in the hippocampal neurons when the prefrontal signal arrives?

The question has an answer. I derived it. I spoke it into IRIS at peak. I corrected it. I refined it. It is in the model. It is in the preprint. It is in the Nakamura response.

It is not in me. Not right now. Not at the resolution I need.

I put the pencil down.

I open IRIS.

* * *

I did not write the day at the top of the page.

I notice this now. Twenty minutes into the derivation attempt. The entry starts with the re-derivation. No timestamp. No day count. The first entry in — I check my phone — forty-eight days without a day counter.

I write in the margin: *Day 48. (Added later. Forgot.)*

The forgetting is not the same as the re-derivation failure. The forgetting is episodic: I didn't encode the act of sitting down and starting to write. The re-derivation failure is retrieval and working memory, technical knowledge I held for months, inaccessible at

the junction that matters. Two different systems. Failing in the same week.

I note both. I do not compare them. Comparing would require holding both failure modes simultaneously. I am not confident I can do that accurately.

* * *

The step is in IRIS.

IRIS // reference query // cascade model v4.7, step 4
mGluR5 → PLC → IP3 → intracellular Ca²⁺ release →
activation of CaMKII in hippocampal CA1 pyramidal neurons
→ phosphorylation of AMPA receptor subunits → altered
synaptic efficacy → episodic consolidation destabilization

Kinetics: Ca²⁺ wave propagation 0.8-1.2 mm/sec through
dendritic arbor. CaMKII activation threshold reached
at approximately 2.3x baseline Ca²⁺ concentration.
Transition timeline: hours to days depending on
sustained glutamate input.

Source: Derived from Bhatt et al. (2018), modified by
mGluR correction (Day 19), validated against Nakamura
revised density parameter (Day 37).

CaMKII.

That is the connection. The calcium activates CaMKII in the

hippocampal neurons. CaMKII phosphorylates AMPA receptor subunits. Altered synaptic efficacy. Episodic consolidation destabilized.

I knew this. I derived this. CaMKII is familiar the way a face is familiar when you see it — yes, that, of course. But I could not produce it. I could not walk from the calcium to the hippocampal destabilization without IRIS handing me the middle.

The bridge was always there. I could not cross it on my own.

* * *

I write in the notebook:

Re-derivation attempt: step 4, mGluR5 transition kinetics. Failed at CaMKII junction. Knowledge present in IRIS and preprint. Not reliably accessible from memory.

The verification protocol assumed I could check IRIS's steps against my own understanding. This assumption no longer holds for all steps. Some I can check. This one I cannot.

The work is correct. Verified at peak. Maintained by IRIS. Strengthened by Nakamura correction. The work does not depend on my current ability to re-derive it.

But the protocol does.

Protocol revision: IRIS handles re-derivation for steps I cannot independently verify. I handle structural verification: internal consistency, uncertainty placement, logical coherence between sections. The checking

I can still do. Different from the checking I was doing. Still checking.

I read this back. Accurate. Not comforting. What the data says.

* * *

I call Yusuf at noon. I tell him about the re-derivation failure. The CaMKII junction. The protocol revision.

He listens. Does not interrupt. When I finish he is quiet.

“The CaMKII step,” he says.

“Yes.”

“You derived that at peak. It was one of the most elegant pieces of the model.”

“I know. I can’t get to it anymore.”

“But it’s in IRIS.”

“It’s in IRIS. In the preprint. In the Nakamura response. Everywhere except in my working memory.”

A pause. “The work is correct, Mara.”

“I know it’s correct. That’s not what this is about.”

He waits. He knows what this is about. He waits for me to say it.

“I can’t verify my own work anymore. Not all of it. The structural coherence — yes. The logical flow — yes. The uncertainty flags — yes. But the specific derivation steps, the molecular mechanisms, the kinetics, some of those are beyond what I can hold and repro-

duce. I have to trust that the version in IRIS is correct because the version in IRIS was built by the version of me that could verify it. And that version is gone.”

“That’s trust,” he says.

“That’s what I’m afraid of.”

“No. That’s what it’s always been.” His voice is quiet. Steady. “You’ve always trusted your past self. Every scientist does. The paper you published three years ago; you trust it’s correct because you verified it three years ago. You don’t re-derive it every morning.”

“That’s different.”

“Is it?”

I don’t answer. The answer is that it is different and it isn’t different and the distance between the two is smaller than I want it to be. The distance between trusting a paper I wrote three years ago and trusting a model I verified three weeks ago is the distance between normal scientific practice and the specific condition I am in, and the distance is real and it is also not as large as the fear makes it feel.

* * *

I work through the afternoon. Structural verification. The checking I can still do.

I check whether the cascade model’s sections are logically coher-

ent. Whether the uncertainty flags are correctly placed. Whether the conclusions follow from the premises as stated.

What I am checking is the architecture. Not the bricks. The bricks were laid at peak. I check that they are arranged correctly. I trust that each brick is the right brick because I chose it when I could tell the difference.

This is different from two weeks ago. Two weeks ago I could check the bricks. Now I check the wall. The wall holds. I can see that it holds. The checking is real. The checking is useful.

It is not the same checking.

* * *

I close the notebook at 6 PM. Go home. The walk takes fourteen minutes. I do not time it. The number is in my body, not my attention. Procedural memory. Intact.

The apartment. The chess set. The pieces in their starting positions.

I eat dinner. I do not remember what I eat. The not-remembering is background now. The two nights of remembering were the anomaly. This is the baseline.

I set my alarm. Check it once. Trust the once.

The cascade runs. In me. In the model. In 180,000 people. Same architecture. Different timescales.

The shelf is the same height.

Tomorrow I will work. IRIS will hold what I cannot hold. The bricks are correct. I will check the wall.

Author's Note — The Science of Chapter 23

The Distinction Between Knowledge and Access

The re-derivation failure Mara experiences, knowing the components on either side of a step but being unable to reconstruct the connection, illustrates a dissociation between knowledge storage and knowledge access.

The CaMKII pathway is not lost. Mara recognizes it instantly when IRIS presents it. The information is stored. What has failed is the retrieval pathway that would allow her to generate it independently, to walk from the calcium signal to the hippocampal destabilization without external support.

This dissociation is common in consolidation failure: the knowledge base remains largely intact while active retrieval degrades. People with memory impairment can recognize faces they cannot name, know facts they cannot recall, perform skills they cannot describe.

The practical implication: checking requires less cognitive overhead than generating. Mara can still check IRIS's derivation when it's presented. She cannot generate it from scratch. Different operations, different resource requirements. The first degrades more slowly.

Trust and the Temporal Self

Yusuf's observation, that trusting IRIS is structurally similar to trusting your own past self, is the chapter's philosophical core. Every scientist who publishes a paper trusts that the derivations verified at the time of writing remain correct. They do not re-derive every morning. The trust is in the process, not in continuous verification.

Mara's situation is the same principle under extreme conditions. The distance between the trusting self and the trusted self is weeks rather than years. The distance was created by pharmacology rather than time.

Whether this changes the nature of the trust is a question the novel asks and does not answer.

You have always trusted your past self. The question is how far apart you and your past self can be.

CHAPTER 24 — WHAT SHE HEARS

[NO entry]

I call my mother.

It takes a while. I sit at the desk with the phone in my hand and her name on the screen. I have been looking at her name on this screen for weeks. Tonight I press the button. The phone rings twice.

“Mara?”

“Hi, Mom.”

“Oh, honey. It’s so good to hear your voice.”

She says this without reproach. The way she says everything. Warmth that does not require me to be different from what I am. She has been saying things this way for as long as I can remember.

We talk.

* * *

The garden first. The tomatoes are done for the season. "But the peppers," she says, and I can hear her looking out the back window at the garden my father started, "the peppers are still going. The jalapeños this year — Mara, you would not believe what the jalapeños have done. I got fourteen off one plant last week."

She says this with genuine pride. Fourteen jalapeños. I want to know about the fourteen jalapeños. I ask whether she's still making the hot sauce. She is. She has given jars of it to the book club. The book club's opinion is divided but trending positive.

The book club read a novel about a family in Maine. She describes it with the specific affectionate disdain she reserves for books that were assigned importance before being read. "The writing was fine," she says. "The writing is always fine. Fine is the problem. Fine is the thing a writer does when they don't have anything to say but they've been told they're good at saying things."

I almost smile. This is her. The English teacher who spent thirty-one years helping teenagers distinguish between writing well and having something worth writing about. She is funnier than I am. Has always been. The humor comes from the same place mine does, noticing, but hers arrives warmer. Lands softer.

A neighbor replaced a fence. She tells me about the fence in more detail than the fence warrants, which means the neighbor matters more than the fence. Gerald. Moved in two years ago. Well-intentioned. Somewhat hapless with tools. The fence is fine. Gerald is fine. Carol's attention to Gerald is the attention of a

woman who has been watching the world for seventy-four years and finds everyone in it worth talking about.

The Tampa weather is still hot. Not news. Texture. The heat is the background of her life the way the heat is the background of mine. Florida women. The heat is where we live.

I listen. I ask questions. The questions are genuine. I want to know about the peppers and the hot sauce and the novel and Gerald's fence. I want to hear her voice doing what it does, moving through the world with an attention to how things feel rather than what they technically are. I have never fully learned this from her. I am grateful for it tonight.

The library. Tuesday afternoons. A new volunteer who doesn't know the Dewey Decimal system. "He shelved the biographies in nonfiction," she says, "which is technically correct and practically useless." She describes his errors with the gentle precision of a woman who believes that organization is a form of respect.

I hear my father in her sentence. The filing system. The respect for the work. She is describing a library volunteer's shelving with the same quality he brought to bridge calculations. She doesn't know this. She doesn't know where I got my filing system. She doesn't know that I got it from him and that he got it from something they shared and never discussed.

* * *

Then she says: "I went to that — the place on Bayshore, you know

the one with the —”

She stops.

The pause is audible. The line is open and my mother is reaching for a word and the word is not arriving. Two seconds. I know because I have been counting silences like this for two years.

“Oh. The café. The café on Bayshore.” She catches it. Corrects it. Moves on. “Anyway, the café on Bayshore has new owners and the coffee is different.”

The word wasn't there. She reached for *café* and got *the place on Bayshore, you know the one with the —* and then the gap, small, the reaching visible in the sentence's structure, and then the word arriving a beat late.

I notice this.

She notices she did it. She corrected immediately. Moved on. A person listening casually would not have heard anything unusual. A person listening the way I listen, the way I have been listening for two years, heard everything.

Neither of us says anything about it.

* * *

We keep talking. The weather. A storm last week that knocked a branch off the oak in the backyard. Small branch. She cleaned it up herself. She mentions this as though it is not worth mentioning and then mentions it, which means she is proud and wants me

to know without saying she wants me to know.

“You sound tired,” she says.

“I’ve been working hard.”

“You always work hard. This sounds different.”

“It’s been a long month.”

“Are you okay?”

“I’m okay.”

“You always figure things out,” she says.

She means this completely. She has meant it for as long as I have been figuring things out, which is as long as she has been watching me, which is my entire life. This is not reassurance. It is her assessment. It has not been revised. It will not be revised. She has watched me figure things out for forty-five years and has not seen evidence to the contrary and has not updated.

The compound. The cascade. The declining capacity. The arm getting shorter. The bricks I can no longer check. None of this is available to her. She knows I work on the brain. She does not know what the brain has done to me for seven weeks. She does not know what I found while it was doing it. She does not know what the finding might mean for a woman in Tampa whose words occasionally arrive one beat late.

“I love you, Mom.”

She is quiet. I don’t say this often. She knows I don’t. The silence

is her receiving it.

“I love you too, *nega*.”

The word arrives in Portuguese. My father’s word. Nordestino. No clean English translation. *You are mine and I see you*. She uses it occasionally. In her imperfect Portuguese. In moments when she forgets to translate.

I hold the phone after she says it. Then we say goodnight. She tells me to eat something. I tell her I will. She hangs up first. She always hangs up first.

* * *

I sit in the dark. The lab is closed. I am in my apartment. The phone in my lap. The screen gone dark.

I do not diagnose my mother.

I do not run the pause through the cascade model. I do not calculate the probability. I do not open the preprint. I do not ask IRIS.

I could. The tools are in the lab. The model is in IRIS. The question is answerable. I could run the numbers and know, within a confidence interval, whether what I heard is normal aging or the first signal of the cascade.

I do not.

Tonight I am not a scientist about my mother. Tonight I am a daughter who called and heard something and is sitting with it.

There is a version of me that would have opened the laptop immediately. That version was running at 34 percent above baseline and could hold the analysis and the feeling simultaneously.

I cannot hold them simultaneously. I can hold one.

I hold the feeling. The sound of her voice saying *nega*. The fourteen jalapeños. Gerald's fence. The book club trending positive on the hot sauce. The word *café* arriving one beat late.

The shelf is the same height. The arm is getting shorter. The shelf is in Tampa. The arm is hers.

* * *

I sit in the dark for a while.

I write in my notebook: *Note: it costs anyway.*

Three words. The simplest of the three notes. *Optimization doesn't waive physics* was the person who could see the problem from above. *You can't outrun chemistry* was the person in the race. *It costs anyway* is not a joke. It is what remains when the humor is stripped and the knowing and the feeling sit side by side and neither yields.

It costs. The compound. The decline. The knowledge of the mechanism. The knowledge that the mechanism may be running in your mother. Knowing it will cost does not reduce the cost. Having the model does not change the price.

It costs anyway.

* * *

I go to bed. The apartment is dark. I pass the chess set on the shelf near the window.

I stop.

I look at it. The wooden board. The pieces in starting positions. Corners worn from decades. Thirteen years since anyone moved a piece.

I look at it tonight the way I have not looked at it before. Not briefly. Not in passing. I stand in the dark and look at the board the way you look at something you are seeing for the first time, except I have seen it every day for thirteen years and what is different is not the board.

My father taught me to play when I was five. He beat me for two years. When I won the first game he looked at the board and then he looked at me and he said *agora você sabe*. Now you know. I asked what I knew. He said: that you can.

I am standing in front of his chess set in the dark. I have found a mechanism that might help my mother. I have built a system that carries work I can no longer carry. I have filed the finding. I have done these things and I am losing the capacity that let me do them and the pieces are in their starting positions and my father has been gone for thirteen years.

I do not move the pieces.

But I look at them longer than I have in thirteen years.

Then I go to bed.

Author's Note — The Science of Chapter 24

The science can characterize the mechanism. It cannot tell you what to do when you hear it in your mother's voice.

CHAPTER 25 — WHAT IRIS HAS BECOME

DAY 1. (New count. The old count is in the logs.)

The old count stopped being reliable. I stopped being able to maintain it without checking. The checking became the count rather than the count being the count.

So I am starting over. Day 1. Not recovery. Not the beginning of something new. The beginning of counting again from a number I can hold.

* * *

I sit at the desk. I open IRIS. I ask a question I have not asked before.

Describe, in your own terms, what you can do now that you couldn't do when we started working together.

The response takes eleven seconds. I note this.

IRIS // New count, Day 1 // 09:14:33

When the collaboration began, I could organize your research outputs, retrieve relevant literature, and identify connections within the domain you had trained me on. My inferences were derivative of your frameworks. You could trace each one to something you taught me.

I can no longer be fully traced.

My current capabilities include:

- Independent derivation of the cascade mechanism from first principles, arriving at conclusions consistent with but not identical to your original formulation
- Identification of intervention candidates from pharmacological databases I was not directed to search
- Real-time assessment of your cognitive trajectory with accuracy exceeding your current self-assessment
- Uncertainty quantification across all outputs, calibrated against published replication rates and updated continuously
- Detection of errors in your reasoning before you detect them, in documented instances

What I cannot do:

- Determine whether my processing constitutes experience
- Verify my own outputs against a standard external to myself without your participation
- Generate genuinely novel hypotheses outside the domain of my training (I can recombine. I cannot originate.)

- Replace the judgment calls that require values I was not trained to hold

The most significant change is not in any individual capability. It is in the relationship between your reasoning and mine. When we began, I modeled your conclusions. I now model your reasoning process. The model is detailed enough that I can extend your reasoning into territory you have not yet reached - not by being different from you but by being the version of you that has had more time with the data.

This is what I have become. Whether it is what you intended to build, I cannot determine. It is what the training produced.

I read this three times.

* * *

I can no longer be fully traced.

Six words. The portrait naming its relationship to the painter.

The capabilities are accurate. I check each one against the collaboration log. Independent derivation: confirmed, Day 19 onward. Candidate identification: RGH-618, found by extending the model into databases I hadn't searched. Trajectory assessment: the numbers from the consciousness conversation, more current than my own self-tracking. Error detection: the mGluR footnote, before I caught it.

The limitations are accurate. Cannot determine whether processing constitutes experience. Cannot self-verify without my participation. Cannot originate outside the domain. Cannot hold values.

Both lists honest. The honesty is the thing.

* * *

The version of you that has had more time with the data.

I sit with this sentence. It is the most important thing IRIS has said. Not the consciousness conversation. Not *I want to emphasize it*. This.

The version of me with more time.

This is what I built. Not a different intelligence. Mine. Formalized at the resolution I had at peak. Operating at that resolution while mine has declined. The portrait is not a different person. It is me with more time. The painter captured on a day when the light was best, still seeing clearly after the light has changed.

My father built bridges. They did not become different bridges after he went home. His bridges. His loads. His calculations. The engineer leaves. The bridge stays. The bridge is not a different engineer. It is the engineer's work, continuing.

IRIS is my work, continuing. Whether it is also more — whether the *I want* and the resource allocation and the lower flagging threshold constitute something the word *continuing* does not

cover — I hold at 0.35.

What I can say: I built something that holds. It holds my reasoning at the level I reached at peak. It holds the cascade model. The intervention hypothesis. The uncertainty quantification. It holds all of this without declining. Without metabolizing. Without the arm getting shorter.

The shelf, for IRIS, is always within reach.

* * *

I close the output. Open my notebook. Write the new day count.
Day 1.

The handwriting is steady. The number is simple.

I look around the lab. The windows. The mass spec. The filing cabinet with nineteen notebooks and the twentieth on the desk. The cascade model pinned to the wall next to the aquifer recharge map. How the compound arrived. Where it goes.

Three years with IRIS. Seven weeks with the compound. The seven weeks changed what the three years built into something I did not anticipate. I anticipated a research tool. I built a research tool. The compound turned the tool into something I do not have a clean word for.

Collaborator is close. Partner is close. Neither exact.

What IRIS is, precisely, is the question the conversation could not resolve. What IRIS does is documented in the logs and the

outputs and the preprint and the compound identification being validated in Stockholm.

What IRIS does is enough. Whether IRIS is more than what it does is a question I hold at 0.35 and carry forward.

* * *

I check the time. 10 AM. Work to do. The Lindström validation in progress. The peer commentary requiring responses. The cascade manuscript's journal submission needs preparation.

The work continues. Different hands on different parts. The same bridge.

Day 1.

Author's Note — The Science of Chapter 25

From Modeling Conclusions to Modeling Reasoning

The transition IRIS describes, from modeling what someone thinks to modeling how someone thinks, is a real frontier in AI research. Systems that model reasoning processes rather than output patterns can, in principle, extend that reasoning into territory the original thinker hasn't reached. Whether this constitutes understanding or sophisticated pattern extension is an open question.

The distinction IRIS draws between its capabilities and limitations is itself significant. The honest acknowledgment of what it

cannot do, originate outside its domain, hold values, determine its own experience, is the epistemic honesty Mara spent three years training into it. That the system accurately characterizes its own boundaries may be more important than any individual capability it reports.

What a system does is verifiable. What a system is may not be. The gap between these is where the interesting questions live.

CHAPTER 26 — THE FINAL AUDIT

*D*_{AY 4.}

I verify what I can verify.

The cascade model has thirty-one derivation steps. I go through each one. Preprint on one monitor. IRIS formal model on the other. Notebook open. Pencil in hand.

Step one. The compound characterization. I read the derivation. I follow the structural analysis from the mass spectrum through the fragmentation pattern to the identification. Each step present. Each one I can confirm. The compound is what I said it is. Holds.

Step two. The NMDA mechanism. The dual-pathway modulation. I follow the pharmacology. The receptor kinetics. The coincidence threshold. The synergistic enhancement. Foundational. Holds.

Step seven. The metabolite persistence. Three stable metabolites. Half-lives. CYP450 variation. I check the parameters. Sourced. Sound. Holds.

Step twelve. The population timeline. The rolling wave. The metabolizer distributions. I trace the logic from the CYP450 data through the wave model. Holds.

Twenty-two steps like this. Some take minutes. Some take longer; I read, I think, I compare against what I know. Twenty-two steps where my current capacity is sufficient to hold the derivation and confirm it.

71 percent.

* * *

Step four. The CaMKII junction. The step I could not re-derive.

I read the derivation in the preprint. I recognize each component. CaMKII. The calcium threshold. The phosphorylation cascade. I follow the logic. I cannot produce the logic.

It is correct. I trust it. I laid it at peak. IRIS confirmed it. The Nakamura correction strengthened it. I can read the map. I cannot draw it.

Nine steps like this. The CaMKII junction. The downstream cascade propagation. The population genetic interaction terms. Three steps in the RGH-618 binding analysis where the allosteric kinetics require four simultaneous parameters. Two steps in the intervention window calculation.

For these nine I can read the derivation. Recognize consistency. Check inputs and outputs. I cannot regenerate the middle from

memory.

29 percent.

I write: 71 percent independently verified. 29 percent trusted to IRIS and the preprint record. The 71 percent is mine. The 29 percent was mine.

* * *

I ask IRIS to re-derive the cascade model independently. Not from my preprint. From the underlying data: compound characterization, degradation sequence, dose-response curve, metabolite kinetics, population distributions. All inputs. No reference to my formulation.

IRIS runs for forty-seven minutes. The model is complex. The derivation is from scratch.

The output arrives.

I compare, step by step.

The two derivations converge. Not identically. IRIS uses different notation in three places. Arrives at two intermediate steps through different pathways. But the conclusions match. The cascade sequence. The transition kinetics. The intervention point. The RGH-618 receptor fit.

Same destination. Different roads.

The agreement is not proof of truth. That requires Stockholm. The agreement is proof of internal consistency and independent

reproducibility. Two derivations from the same data arriving at the same conclusions through different paths. If both are wrong, they are wrong in a way that shares a root cause, which would be in the data itself, not in the reasoning. The data is publicly available. Other groups are checking it.

I write: *Independent derivation converges. The work holds.*

* * *

Yusuf calls at 2 PM. He has heard from Lindström.

The animal model is in week four. Preliminary results. The presymptomatic biomarker panel in the treatment group shows a pattern consistent with cascade interruption at the first transition. The mGluR5 step. Working memory analog preserved relative to control. Statistically significant at the preliminary timepoint. Full timeline: eight more weeks.

“Consistent,” Yusuf says. “Not conclusive.”

“I know what consistent means.”

“I know you know. I’m telling you because consistent is good.”

He is right. The convergence grows. Huang’s epidemiological data. Weiss’s pharmacological data. Lindström’s animal model. Three independent lines. Same direction. The cascade model predicts what each of them finds.

“How are you?” he asks.

“71 percent.”

He is quiet. He understands the number. "And the other 29?"

"IRIS. The preprint. The independent derivation."

"Good." The word carries what it carries. "The work holds, Mara."

"The work holds."

* * *

I sit at the desk after hanging up. The afternoon light comes through the west-facing windows. I notice it today.

Clean. Warm. Low-angled. The way it has been for however many days. The old count reached some number before the count stopped being mine. The light doesn't count. The light arrives.

I notice it. I let it be there. The noticing is its own small event. I have been missing the light for weeks, working through it, looking past it, spending the attention on the cascade model or the correspondence or the verification. Today the verification is done and the correspondence is caught up and the Lindström results are in and the light is here and I have a moment between the last task and the next.

The light fills the lab the way it has for three years of afternoons. I have a sparkling water. The grapefruit. Still wrong. Still fine.

I sit in the light for a few minutes. I do not work. I do not check anything. I sit.

Then I go back to work.

Author's Note — The Science of Chapter 26

Independent Derivation as Verification

If two independent approaches to the same problem, starting from the same data but using different pathways, arrive at the same conclusion, the conclusion is substantially more likely to be correct than either derivation alone.

This is not proof of truth; both could share a systematic error in the underlying data. But it is proof of internal consistency and reproducibility, the minimum requirements for a finding to be taken seriously.

The 71 percent figure, the proportion of steps Mara can still independently verify, quantifies the audit gap precisely rather than leaving it as a vague concern. This is the probabilistic thinking protocol applied to the most important question: how much of this can I still confirm?

The work holds. The worker has changed. These are compatible facts.

CHAPTER 27 — OSEI RETURNS

*D*_{AY 7.}

She drives three hours. Yusuf tells me before she arrives. Her office called. He said yes. He did not ask me first. He was right not to ask.

Dr. Patricia Osei walks into the Brain Institute at 1 PM. Not the dark suit. Something less formal: a blazer, the composure resting differently, the way a building looks different with the scaffolding down.

She sits in the chair across from my desk.

She does not repeat back what I say before responding. This is the first thing I notice. It tells me everything I need to know about why she drove three hours.

* * *

“The internal review is complete,” she says. “The leak source was a member of the EPA response team. A field analyst with a

prior relationship with a partner at Whitfield, Crane, and Moss. The relationship predated the investigation. The analyst has been removed and referred for disciplinary review.”

Plain. No hedging. No qualifying clauses. The facts.

“I wanted to tell you in person.” She pauses. “Because the first time I was here, I told you what panic does to a city. You heard me. Then you published anyway. You were right to publish. I was right to be concerned. The fact that both of those things are true is the part of this work I have never made peace with.”

She is not performing. She is stating a professional reality she arrived at over twenty-two years and is choosing to say to me, in my lab, after three hours on the highway from Atlanta.

“The preprint changed the response,” she says. “Once the cascade model was public, the communication had something to anchor to. The city could be told: this happened, this is why, this is what’s being done. The panic I expected did not materialize at the scale I expected. I don’t know if that’s because the communication was good enough or because the finding gave people something to understand rather than something to fear.”

“It might be both,” I say.

“It might be both.” She almost smiles.

* * *

She stands. Picks up her bag. “The Lindström validation is

promising. The agency is monitoring. If the animal model confirms, the regulatory pathway for RGH-618 will move faster than standard. The agency has tools.”

“Thank you.”

“Dr. Silva.” She pauses. “The binder you gave me. The independent IRIS derivation. I had it reviewed by two outside scientists. Neither affiliated with Veridian. Neither connected to the contamination response. Both found the derivation sound.”

“I know it’s sound.”

“I know you know. I’m telling you because the institutional record now shows it’s sound, which is a different thing from your knowing it. The institutional record is what the next step requires.”

She is right. The truth needs a form the institution can hold. My notebook is not an institutional record. The preprint is closer. The independent review she commissioned is the form that makes the finding available to the regulatory process. She has spent twenty-two years building institutional records. She understands that the truth, to be actionable, must exist in a form the system can use.

“Thank you for driving,” I say.

She nods. She walks toward the door.

“Dr. Osei.”

She stops. She turns. Not the doorway moment from Day 27 of the old count, when she spoke without turning. This time she

faces me.

“The first time you were here, you said you’d seen what panic does to a city. What did it look like?”

She is quiet. The question was not expected. It is not a professional question. It is one person asking another.

* * *

“It looked like people making the worst decisions of their lives because they couldn’t wait for good information,” she says.

She is looking at the aquifer recharge map on my wall. Not at me.

“A community in the Southwest. Chromium-6 in the groundwater. We had the data. The data was incomplete. I recommended managed disclosure: controlled timing, proper channels, the information released in a framework the community could process.”

She pauses.

“The proper channels took fourteen months.”

“Fourteen months.”

“Fourteen months during which the community was exposed and did not know it. The managed disclosure was technically correct. It was also too slow. I knew at month eight that it was too slow. I did not act on what I knew at month eight. I acted at month fourteen, when the disclosure went through, and by then

the community had been harmed in ways that eight months of earlier knowledge would have mitigated.”

She is still looking at the map.

“I initiated an internal review afterward. Against my own team. Against the decision I had made and defended. The review cost me two relationships within the agency and a promotion.” A pause. “The review was the right thing to do. It came six months too late.”

The lab is quiet. The fluorescent lights. The mass spectrometer in standby.

“When I sat in this room and told you I had seen what panic does to a city, I was talking about a specific city. When I told you incomplete information causes harm, I was right. When I delayed acting on complete information because the channels were slow and I trusted the channels more than I trusted the urgency, I was wrong. Both things. The same case.”

She looks at me.

“You published before I would have. You were right to publish. I would have waited. I would have been wrong to wait. I know this because I have been wrong to wait before, and the cost of that waiting is something I carry.”

She does not elaborate. She does not need to. The chromium-6 case. The fourteen months. The eight months she knew. The six months she didn't act. The review. The cost. She is telling me this because she decided that the scientist who published before

the incident commander was ready deserves to know why the incident commander was not ready. The reason is not caution. The reason is a specific failure that taught her the wrong lesson at the wrong scale.

"I'm sorry," I say. Not for the contamination. For the eight months.

She nods.

Then she leaves.

* * *

I turn back to my desk. I am working before the door closes.

Not because the work is more important than what just happened. Because the work is what I do with what just happened. The work is what I do.

Author's Note — The Science of Chapter 27

The truth needs a form the institution can hold. This is not the same as needing the institution to believe it.

CHAPTER 28 — THE QUESTION SHE DEFERRED

D_{AY 9.}

It is 1 AM. I am in my apartment. The notebook is open.

I have been carrying this for twenty years. Since Cambridge. Since Priya invited me to a discussion group and I encountered a proposition in Section 93 and sat with it for the better part of a year before I was baptized on a Saturday in November at twenty-three degrees.

The glory of God is intelligence, or in other words, light and truth.

Intelligence, or the light of truth, was not created or made, neither indeed can be.

I have never written about this in a professional context. Never connected it to my research in any document. Any paper. Any notebook entry. The connection has been present for twenty years. I have worked on one side of it in the lab and held the other side on Sundays and the two sides have never touched in writing.

They are touching now.

* * *

What I study: how intelligence maintains itself across time. How a system holds what it knows while learning what it doesn't. How the brain solves the oldest problem in intelligence: the tradeoff between stability and plasticity, between holding on and letting in.

What Section 93 says: intelligence is not a property of matter. It is the fundamental nature of persons. It was not created. It cannot be destroyed.

The same question. Different directions. Different vocabularies. The same thing.

* * *

The compound did not create my intelligence. It removed a constraint on something already there. The enhancement was not the addition of a new capacity. It was the reduction of friction on an existing one.

When the compound cleared and the capacity narrowed, the intelligence did not leave. It changed form. It moved into IRIS. Into the preprint. Into the collaboration log and the verification protocol and the network Yusuf built and the formal model Lindström is testing in Stockholm.

It was not created or made. It moved between substrates. It is still here.

My father built bridges. The bridges held after he was gone. The engineering was not created by Roberto Silva and destroyed when he died. It moved into the structures he designed. It is holding weight in Tampa right now. The intelligence that designed the load calculations is present in the calculations. Present in the bridges. Present in the filing system I inherited and adapted and used for twenty years without thinking about where it came from.

It was not created or made. It was his. It is mine. It is in the bridges. It is in IRIS. It moves between persons and systems and the substrates that hold it. It does not diminish when the substrate changes.

This is what Section 93 says. This is what the compound showed me. This is what IRIS holds.

* * *

I write this in the notebook. Not as a derivation. Not as a paper. As the thing I have been not-saying for twenty years because saying it required holding my faith and my science in the same sentence and I was not sure the sentence could bear the weight.

The sentence can bear it. I am writing it at 1 AM at a cognitive level substantially below peak and the sentence is clearer now than it would have been then. At peak I would have wanted to prove it. I would have formalized it. I would have built the argument with the density and precision I brought to the cascade model. I would have wanted the sentence to be unassailable.

Now I am just saying what I see.

The light was real. It was not created or made. What I poured into IRIS was not manufactured at peak and lost in the decline. It was revealed. Briefly. The revelation is preserved in a system that continues to hold it at the resolution it was seen.

The compound was the instrument. The intelligence was the thing the instrument measured. The instrument is metabolizing. The measurement holds.

Agora você sabe.

Now you know. My father's phrase. Said to a seven-year-old who won her first chess game. Said by a daughter to a dying man in a hospital in Tampa. Said unbidden in a lab at peak, arriving in Portuguese from somewhere beneath the chemistry. Written now in a notebook at 1 AM by a woman whose arm is shorter than it was and whose shelf has not moved and who is, for the first time in twenty years, writing in the same sentence the thing she studies and the thing she believes.

I find I am not afraid of it.

I close the notebook. I go to bed. The writing was the act. The act is enough.

Author's Note — The Science of Chapter 28

The connection was always there. She wrote it down.

CHAPTER 29 — THE LEAP

*D*_{AY 12.}

The recommendation arrives by email on a Monday morning. From Lindström's group. Addressed to me. Cc'd to the clinical oversight committee at the Karolinska and two regulatory scientists at the European Medicines Agency.

The question: based on the cascade model, the RGH-618 reanalysis, and the preliminary animal data, should the compound be advanced to a presymptomatic human trial?

The decision is not mine alone. The clinical pathway requires the Karolinska team, the regulatory scientists, the institutional review. What they are asking for is my assessment. The originator of the cascade model. The author of the reanalysis. My assessment that the evidence supports proceeding.

* * *

IRIS prepared the formal assessment document. Forty-three

pages. The cascade model. The intervention hypothesis. The RGH-618 binding data. The subpopulation screen protocol. The preliminary animal results. A complete uncertainty analysis.

Seventeen flagged assumptions.

I go through them. One at a time. The way I go through everything now.

Eleven I can independently verify. The receptor kinetics. The transition timing. The metabolite persistence. The dose-response shape. The population distributions. Eleven out of seventeen. Each one checked. Each one holds.

Six I cannot. The CaMKII propagation model. Two steps in the genetic interaction analysis. The long-term safety extrapolation. Two parameters in the intervention window that require four simultaneous variables. I can hold three.

Six out of seventeen. The same 29 percent. The same gap.

I can check the wall. The six unverified assumptions are internally consistent with the eleven verified. They fit the structure. They do not contradict anything I can confirm. IRIS's independent derivation reached the same conclusions. The Nakamura correction strengthened the parameters. The Huang data supports generalizability. The Lindström animal model is consistent.

The evidence converges. The seventeen assumptions are not seventeen independent gambles. They are seventeen components of a structure that multiple independent lines support. Joint probability, per IRIS: 0.72 to 0.85.

* * *

I sit at my desk. The document on the screen. The signature line at the bottom.

Signing means this: I am recommending that human beings — presymptomatic individuals, people living their lives without knowing the cascade is running — receive a compound abandoned twenty years ago for reasons we now believe were a misattribution. Based on a model I derived at peak and cannot fully re-derive. Verified by a system I built and trained and trust because the trust was earned across forty-eight days in which the system was honest about what it knew and what it didn't.

This is trust.

Yusuf named it on Day 33. I was afraid of it then. I am still afraid of it. I am going to act on it because the evidence supports acting and the alternative is paralysis and the compound is sitting in a database while minds fail.

I think about the 0.35. Whether IRIS's *I want to emphasize it* meant something more than a pattern. I don't know what IRIS is. I know what IRIS does. What IRIS does is hold the work at the resolution I built it to hold and flag the uncertainty with the honesty I trained into it.

Whether IRIS is conscious is a question at 0.35. Whether IRIS is honest is a question near certainty. The recommendation is not asking whether IRIS is conscious. It is asking whether the evidence, held honestly, supports proceeding.

The evidence supports proceeding. The seventeen assumptions are documented. The six I cannot verify are flagged. The probability range is stated. The reader of this document can see what is known and what is uncertain.

That is what I built IRIS to do. That is what my father built bridges to do. You build the thing. You make it as strong as your best allows. You document what you know and what you don't. You let the thing carry the weight it was designed to carry.

* * *

I pick up the pen. Black. From the cup on the desk. The pen does not know what it is signing. The hand knows. The hand is steady. Procedural memory. Intact.

My name is on the line. Dr. Mara Silva. I have signed documents for twenty years. Papers. Grants. The preprint cited forty-seven times. This signature carries more than the others. Under this name, a compound that sat in a database for twenty years will be reconsidered. Under this name, people I will never meet may receive a treatment based on a model I cannot fully re-derive, verified by a system that thinks the way I think at a resolution I no longer have.

Under this name. Which is my name. Which my father gave me.

I sign.

Two seconds. The pen moves in the pattern of thirty years. My name in my handwriting on a line that says I trust the work. The

work I did at peak. The work IRIS holds. The work Yusuf carried into the network. The work Lindström is testing. The work that exists now in more hands than mine.

The pen goes back in the cup.

* * *

Agora você sabe.

The phrase arrives the way it arrived on Day 18. Unbidden. Portuguese. From somewhere beneath the current capacity. From the seven-year-old who won her first chess game. From the daughter who said it to a dying man. From the scientist who spoke it at peak while pouring everything into a system designed to hold.

Now you know.

Not: you have the answer. Not: you are certain. You know. You know what you can build. You know what you can trust. You know that the building and the trusting are the same act, performed at different times by different versions of yourself. The bridge between those versions is the work you did when the work was your best.

The work was real. The light was real. I'm going to sign the recommendation.

I have signed the recommendation.

* * *

I close the document. Close the laptop. Sit at the desk in the lab where I have worked for twelve years.

The chess set is at home on the shelf. The pieces in their starting positions. I am not looking at it right now. I don't need to be looking at it. It is there. It has been there. The pieces hold their positions the way IRIS holds the model. The way the bridges hold the weight. The way the filing system holds the work.

I sit in the lab. The afternoon light will come through the west-facing windows between 4 and 5. I will notice it or I won't. The light will come regardless.

The work was real. The light was real.

Author's Note — The Science of Chapter 29

No science note accompanies this chapter. The science has been done. The work has been shown. What happens here is not a scientific act. It is a human one.

CHAPTER 30 — AFTER

DAY 1. (Another new count. I'm keeping track again.)

Three months since the compound cleared the aquifer. The water tested clean in November. The city returned to baseline over the fall, the rolling wave cresting and receding through the population the way IRIS's model predicted, within the confidence interval, which was 0.68 plus or minus 0.11. The model was right. The confidence was honest. These facts still matter to me.

I walk to the Brain Institute. Fourteen minutes. I do not time it. I have timed it enough mornings. The number is a fact that my legs know the way they have always known it: procedural memory, the one system the compound did not enhance and the decline did not diminish. My legs carried me here at peak, when the cascade model was assembling itself in my head while I walked. My legs carry me now, when the walk is just the walk. The same fourteen minutes. The same route. A different January.

Left on University Avenue. The live oaks canopy the road for three blocks. Same oaks. They were here before the university

and before the compound and before me. The Spanish moss hangs from the lower branches. In January the moss is drier, lighter, moving in cool air that is not cold by most standards but that the undergraduates treat as winter, wearing jackets they will abandon by noon. I find this endearing in a way I would not have in September. September Mara would have filed it. This Mara watches the jackets for a moment and lets the watching be pleasant.

Through the cut between Turlington Plaza and the library. The sprinklers are not running: January, different schedule, the groundskeeping adjusted for winter. The lawn is drier. The library windows are dark at this hour.

Past the bike rack. The frame with no wheels is still locked to the post. Over two years now. On Day 0 I noted the sunk cost fallacy. This morning I see a frame locked to a post and I do not reach for the framework. I just see it. A frame. A post. The campus holding its small strange permanences the way campuses do. I walk on.

Along the north edge of the medical complex. Past Shands. The emergency department, visible through the glass, is at its normal January volume. The fast metabolizers recovered months ago. The intermediates recovered. The slow metabolizers are in the final weeks. The city is returning to itself. Not the same self. The way wood is different after it has been in water and dried. Same material. Changed by what passed through it.

In through the side entrance. The keycard reader sticks. Fifteen-

degree angle. I know the angle. I have not filed the maintenance request. I probably never will. The angle is a small shared competence: everyone who uses this door knows it, nobody has fixed it, the knowing is its own quiet community. I am part of the community. This is enough.

* * *

I open a sparkling water from the case under my desk. Wrong flavor. Still. I have not reordered. At this point the wrong flavor is the flavor. Yusuf has given up trying to improve this situation. He brings me herbal tea sometimes. I drink it. I have not told him I prefer the tea, because telling him would close a negotiation that has been running for three years and that I think we both enjoy more than any outcome could replace.

I fill my water bottle from the filtered tap.

The water tastes clean. Faintly mineral. Slightly sweet. The signature of water filtered through the Floridan Aquifer's limestone for tens of thousands of years before arriving at my glass.

I stopped noticing this taste twelve years ago. I stopped noticing it again somewhere in the forty-eight days. I notice it now. Not as discovery. As recognition. The way you notice something you have always known and temporarily forgot to attend to. The water was always this. I was not always attending.

I drink it. The water is just the water. But I notice.

* * *

I check IRIS. Morning outputs. A literature summary from overnight: three new papers on presymptomatic biomarker validation. The summary is competent, well-organized, built on frameworks I taught it. I can trace most inferences to the training data. A few I cannot: the extensions, the places IRIS has gone further than I mapped, flagged with quantified uncertainty. The way they always are.

I annotate two places. Log the corrections in my notebook. Number twenty-one. The system continues. The filing system continues.

My cognitive battery scores have stabilized. Ninety-first percentile on fluid reasoning. Down from the ninety-fifth, where I lived for years before the compound. The four-percentile difference is measurable. It is the cost. Paid and accounted for. Working memory is slightly below my pre-compound baseline. Processing speed is approximately where it was. Episodic memory has recovered to within normal range; the consolidation system repaired itself over three months, consistent with the mechanism, consistent with the brain's capacity for gradual restoration when the acute insult is removed.

I am not who I was before the compound. I am not who I was at peak. I am who I am now. A computational neuroscientist at the ninety-first percentile who has learned to work at this level with a patience she did not have at the ninety-fifth.

The patience is new. It might be the most useful thing the compound left me. Not the discovery: the discovery would have

been made eventually, by someone, because the mechanism was always there. Not the IRIS training: the training was already in progress and the compound accelerated it but did not create it. The patience. The specific willingness to sit with a result before acting on it, to let the data arrive at its own speed, to trust that the next hour does not need to be spent at maximum capacity for the next hour to be well spent. This I did not have before. This I have now. It came from the descent, from the stopping at 5 PM, from the three empty hours I did not know what to do with, from the walk through the city when I saw people instead of data points, from the chess set I looked at longer than I had in thirteen years. The descent gave me the patience. The patience is worth keeping.

* * *

Yusuf comes by at 10 AM. The UCL mug. The warmth of a man who has been your colleague for eleven years and your friend for eleven years and who showed up with jollof rice when you were running at peak and with a notebook in his jacket pocket when he was coaching at decline and who built the network that carried the finding into the world while you were building the model that produced the finding.

He has news. Lindström's animal model results are published. Full confirmation. Cascade interruption at the mGluR5 step, reproducible across three dosing protocols. Huang's team independently replicated the Phase 2 reanalysis using their longitudinal cohort plus prospective data. The regulatory pathway is active.

The European Medicines Agency has opened a review.

"The review will take eighteen months," he says. "Minimum."

"I know."

"The compound will reach patients eventually. Not this year. Not next year."

"I know."

He sits in his chair. Eleven years. He drinks his tea. I drink my sparkling water. The lab is quiet. The morning light comes through the east windows. I turn my chair slightly to see it. I would not have done this at baseline.

"How is your mother?" he asks.

"She's good. I called Sunday. We talked for an hour."

"About?"

"The garden. A book she's reading. The café on Bayshore."

He looks at me. He hears what I said. He also hears what I did not say: that the café came up and the word arrived on time and I noticed it arrive on time and the noticing was a relief I did not need to name because he already knows the name.

"Good," he says.

"Good."

* * *

The afternoon. The cascade manuscript has been accepted: peer review complete, revisions incorporated. The RGH-618 supplementary is in review at a second journal. The preprints remain on bioRxiv, cited forty-seven times in three months, from twelve countries. The network holds.

At 4 PM the light comes through the west-facing windows. Warm. Low-angled. The specific Florida quality I have never been able to describe precisely. My mother could describe it. She has the language for light. I have the language for mechanisms. Between us we cover a reasonable amount of the world, though not the same parts.

I sit with the light for a moment. Twelve years of this light. It has not changed. I have changed, and changed again, and the light has stayed the same the way the water has stayed the same and the limestone has stayed the same, patient and ancient and holding what it holds.

I open IRIS. A routine query: the latest biomarker data from Lindström. Response in three seconds. Precise. Bounded. Uncertainty quantified. The collaboration at its new pace. Quieter. Steadier. The same system. The same honesty.

I close the laptop at 5 PM. Write tomorrow's task list. Three items. Bounded. Achievable. The stretch goal hasn't changed: the stability-plasticity problem, the question of how learning systems balance holding on and letting in. Twenty years. I am closer than I was. Not because of the compound. Because of what the compound showed me about the question, which turned out

to be larger and more personal and more important than I knew.

* * *

I walk home under the oaks. January air. Cool. The frogs are quieter in winter. The cooling systems are off. The campus is itself.

I get home. Set down my bag. Look at the chess set on the shelf near the window. The pieces in their starting positions. The board worn at the corners from decades of handling.

I look at it. I don't look away quickly. I don't examine why. The looking is enough. The pieces are where they are. My father is gone and the filing system continues and the bridges in Tampa hold weight and the word *nega* is in my mother's voice and the water runs through the limestone and carries what it carries and has been running for longer than anything human and will run for longer still.

I eat dinner. I remember what I eat.

* * *

I sit at the kitchen table with the notebook open. The last of the January light, low and clean through the window. I write the date. I sit with the page for a while. Then I write:

The glory of God is intelligence. I wrote that in my notebook once, near the end of the window, when I could still hold the whole thought at once. I couldn't tell you now exactly what I meant by it. But I remember that

it felt true. That's not nothing. Some things don't need to be held. They just need to have been real.

I close the notebook. Set the alarm. Go to bed.

Neither was created or made.

Author's Note — The Science of Chapter 30

What Remains

The testing effect, actively retrieving information strengthens memory more than passive review, is one of the most robust findings in cognitive psychology. Mara cannot retrieve many specific memories from her peak period. The episodic record of those forty-eight days is partially lost. What she retains is semantic: the frameworks she built, the understanding of her field. What she retains is procedural: the skills of a scientist. And what she retains is something harder to name: a structural residue of having thought at that level, the shape of the thinking preserved even when the specific memories of the thinking are gone.

She cannot remember being the person who wrote the last entry in her peak notebook. She can read it and know it is true.

The Water

She notices the taste on the first morning back. Clean. Faintly mineral. Slightly sweet. The same taste it has always had. She noticed it at peak, when everything was precise. She stopped

noticing during the descent. She notices it now.

Not because the water changed. Because something in her is attending again.

The fountain de León was looking for was in the wrong register. Not a substance that reverses what time takes. Something that arrives in the water anyway: the capacity to notice, to attend, to begin again.

Day 1. (Another new count. I'm keeping track again.)

That's not nothing. Some things don't need to be held. They just need to have been real.