

THE SCOVILLE SPLICE LIBRARY · VOLUME II

The Capsaicin Codex

The anatomy of a burn. *Written, not found.*

The science of designed heat — from the receptor that reports a fire that was never lit, to the single enzyme we rebuilt two residues at a time, to a ladder of ten peppers and the method that frees the whole craft from the pepper itself.

SCOVILLE SPLICE · GROWN BY OUR FAMILY. ENGINEERED FROM SCRATCH.

Contents

01 Opening essay · The fire that was never lit

02 Chapter 1 · The Lie of Fire — a burn with no flame.

03 Chapter 2 · Five Ways to Burn — the shape behind the number.

04 Chapter 3 · From Scoville to the Index — when a feeling became a figure.

05 Chapter 4 · The Universal Highway — the road every green thing already paves.

06 Chapter 5 · The Branch in the Road — where the fat learns to kink.

07 Chapter 6 · The Last Enzyme — one knot ties the whole thing.

08 Chapter 7 · Reading the Fold — a structure no one had ever seen.

09 Chapter 8 · The Machine That Reads Mutations — six hundred and forty guesses, ranked.

10 Chapter 9 · Two Edits — the substitutions the breeders never found.

11 Chapter 10 · Docking the Conversation — watching the molecule sit down.

12 Chapter 11 · The Enzyme Holds — nine nanoseconds of proof.

13 Chapter 12 · Flux Has the Last Word — the pathway, balanced.

14 Chapter 13 · Designing the Burn, Not the Number — same heat, different fire.

15 Chapter 14 · The Ten — a ladder written in Scoville.

16 Chapter 15 · Phoenix — all eight, at once.

17 Chapter 16 · The Pepper Was Never the Point — the same hands, any plant.

18 Chapter 17 · Precision as Reverence — why the numbers are a form of love.

Opening essay · The fire that was never lit

Cut a ripe superhot in half and the heat is visible before it is felt. It is not in the flesh and it is not, as almost everyone believes, in the seeds. It is in the placenta — the pale, waxy ridge that runs down the inner wall, holding the seeds the way a hand holds coins. Press a fingertip to it and lift, and for the rest of the afternoon that finger will remember. Nothing has happened to the skin. No tissue has been harmed, no temperature has risen, no injury has occurred. And yet the body insists, with complete conviction, that it is burning.

That insistence is the subject of this book. It is also, we have come to believe, one of the most beautiful lies in biology — a molecule that walks up to a receptor built to warn you about genuine thermal danger and whispers the warning itself, so persuasively that forty centuries of cooks organized whole cuisines around the sensation of a harm that is not there. We grew up inside one of those cuisines. The question that became this company came up the way the good questions always did: at our table, over a dish that was almost too hot to finish and far too good to stop. *How much* hotter could it go? And underneath that, the question no one at the table could answer: *how* — by what mechanism, down to the atom — does a plant make a nerve lie?



The founders in the research greenhouse before the lights came up — where the question that became the Codex was first asked out loud.

For ten thousand years the craft of heat was a craft of patience. Grow the plant, taste the fruit, keep the seed of the one that moved you, and wait a season to see whether the fire bred true. It usually did not. Capsaicin is the product of a long metabolic assembly line, and every link in that line is a gene, and every gene drifts, so the heat a grower selected one year arrived the next year softened or scattered or simply gone. The old chiliheads were not sloppy. They were fighting the variance of a biological process they could feel exactly and could not read at all.

This book is about learning to read it. We will go down to the receptor first — **TRPV1**, the ion channel on the trigeminal nerve that capsaicin hijacks — because you cannot design a sensation you do not understand at the level of the thing that senses it. Then we will

trace the molecule backward, up the assembly line, naming every enzyme and every intermediate from the amino acid phenylalanine to the finished capsaicinoid, until we arrive at the last machine in the line: **Pun1**, a single acyltransferase that performs the final handshake and whose structure, until recently, no one had ever seen. We will show you how we saw it — how a protein language model read six hundred and forty possible mutations and ranked them, how a physics engine watched the enzyme hold its shape for nine billionths of a second, how two substitutions — S39L and L345G — no breeder could ever have stumbled onto, whose design our model predicts approximately **2.95 million SHU** and, more importantly, cuts batch-to-batch wandering from forty-two percent down to seven. And then we will do the thing this whole method was built to do, which is leave the pepper behind — because, as the last part of the book argues, *capsaicin was never a pepper problem.*

A word on how we write. We are a family, and the voice of this book is *we* — not the royal we and not a committee, but the actual plural of people who revere what they grow. The computing was one pair of hands; the work, the table, and the wonder belong to all of us. You will find the mathematics shown rather than hand-waved, because to us a number is a form of respect for the reader: we will not tell you a pepper is hotter without telling you by how much, measured how, against what. And you will find, humming under all of it, a conviction we will state once here and then let you feel rather than hear — that food is something you give your body with intention, that a plant grown to carry the most it can is a kind of care, and that designing flavor and designing nourishment have always been, for us, the same act. We make no claims about your health. We make a great many claims about chemistry, and we show our work on every one.

A map of the road ahead, since this is a long book and a deliberate one. Part I stays with the tongue — the receptor, the five molecules that speak to it, and the measurement system that lets us say how loudly. Part II walks the molecule's biography backward, enzyme by enzyme, from the amino acid to the finished capsaicinoid, and arrives at the single protein that makes heat. Part III is the design proper: how we saw that protein's shape, ranked its every possible mutation, chose two, and confirmed their computational stability before a seed was sown. Part IV is the payoff you can taste — the ladder of ten cultivars and the blend logic that gives each its own character of fire. And Part V leaves the pepper behind for the forty-two other plants the method already speaks in, and the sixty-four more on the drawing board. We built the science in that order because that is the order it must be understood in: you cannot design a sensation without the receptor, cannot build the molecule without the pathway, cannot edit the enzyme without its structure, and cannot leave the pepper until you have proven the whole loop closes on one.

The old craft ended its sentences in poetry because it had no other language. *A warmth that blooms. A fire that lingers.* We love that language; we grew up on it; nothing here is meant to replace it. But you cannot design what you cannot measure, and for ten thousand years no one could say how much. The Codex is our attempt to

finish the sentence — to put a number where the poetry trailed off, and then, with that number in hand, to write the fire deliberately, molecule by molecule, instead of waiting a season to find it.



PART I OF V

The Sensation

TRPV1 · the vanilloid pocket · five capsaicinoids · from Scoville's tongue to the Index

Chapter 1 · The Lie of Fire — *a burn with no flame.*

Begin where the sensation actually begins, which is not on the tongue's surface but underneath it. The taste buds — the receptors for sweet, sour, salt, bitter, and umami — have nothing to do with heat. Capsaicin is not a taste at all. It is a message delivered to the **trigeminal nerve**, the great fifth cranial nerve that carries touch, temperature, and pain from the face and mouth to the brain. When we say a pepper is hot, we are reporting, with perfect accuracy, that something has activated the exact nerve fibers that would fire if we had actually put our mouth near a flame.

The door those fibers open is a protein called **TRPV1** — the transient receptor potential vanilloid type 1 channel. It sits in the membrane of the nerve ending as a tetramer, four identical subunits arranged around a central pore, and its ordinary job is to be a thermometer with a threshold. Below about 43 °C it stays shut. Above 43 °C — the temperature at which real tissue damage begins — it opens, and when it opens it lets cations, principally calcium and sodium, pour through the pore and down their electrochemical gradient into the cell. That influx depolarizes the nerve ending, a voltage spike races up the axon, and the brain receives an unambiguous telegram: *this tissue is being burned, withdraw.* TRPV1 is also opened by acid — by protons, below roughly pH 6 — which is why a burn and a scald and a splash of acid can feel like cousins. It is, in other words, one of the body's primary alarms for genuine thermal and chemical threat.

Architecturally, TRPV1 is a thing of real elegance, and the elegance matters for anyone who means to exploit it. Four identical subunits assemble around a central ion-conducting pore, each subunit contributing six membrane-spanning helices, labeled S1 through S6. The last two, S5 and S6, from all four subunits line the pore itself; the first four, S1 through S4, form a peripheral bundle in each subunit that houses the sensors — and it is into that peripheral S1–S4 bundle, specifically the cleft framed by S3, S4, and the short linker that joins S4 to S5, that capsaicin binds. The pore carries two gates in series, a narrow selectivity filter near the outer mouth and a lower gate formed by the crossing of the S6 helices, and channel opening is the coordinated widening of both. TRPV1 is, in the language of the field, *polymodal*: the same channel integrates heat, acid, capsaicin, and an assortment of animal venoms — spider and scorpion toxins that evolved, independently of any plant, to attack this very receptor. Its pore is modestly calcium-selective, which is what makes the calcium signal of the previous paragraph possible, and it sits predominantly on the thin, slow, unmyelinated C fibers and the lightly myelinated Aδ fibers of the trigeminal and dorsal-root ganglia — the dedicated wiring of nociception, the body's pain-and-temperature telegraph. For most of the channel's history this architecture was inferred rather than seen; it was the cryo-electron-microscopy revolution of the last decade that finally resolved TRPV1 at near-

atomic detail, capsaicin caught in its pocket, and handed us the coordinates we would later dock against. We will use that structure in earnest in the next chapter.

Capsaicin's trick is that it does not raise the temperature. It lowers the threshold. The molecule slips into a pocket on the intracellular side of each subunit — the **vanilloid pocket**, a hydrophobic cleft formed between the S3 and S4 helices and the S4–S5 linker — and binds there, and that binding is energetically equivalent to heat. It stabilizes the open state of the channel. With capsaicin bound, the temperature at which TRPV1 opens falls from 43 °C to somewhere near body temperature, so the channel that should have stayed shut swings wide at the temperature your mouth already is. The calcium floods in. The telegram is sent. The brain, which has no way to audit the nerve, believes every word.



Cross-section of a designed superhot: the pale placental ridge, not the seed, is where the capsaicinoids are made and stored.

Two features of this mechanism matter enormously for anyone trying to *design* heat rather than merely suffer it, and we will spend much of this book on both.

The first is that TRPV1 response is **graded and saturable**, not binary. A single channel is effectively open or shut, but a nerve ending carries many channels, and a patch of tissue carries many endings, and what the brain reads is the aggregate — the fraction of the available channels that capsaicin has recruited at a given concentration. That fraction follows a classic dose-response curve: almost nothing at very low concentration, then a steep rise, then a plateau as the receptors run out. The single most informative point on that curve is its midpoint, the concentration at which half the maximal response is reached. Pharmacology calls it the $\mathbf{EC_{50}}$, the

half-maximal effective concentration, and it is the hinge on which our entire measurement system turns. A compound with a low EC_{50} reaches half-activation at a low concentration — it is *loud*. A compound with a high EC_{50} needs to shout to be heard. Capsaicin's intrinsic EC_{50} at TRPV1 is on the order of a few tenths of a micromolar: sub-micromolar, which is to say it is one of the loudest small molecules the human body knows how to hear.

The second feature is **desensitization**, and it is the reason capsaicin is a pleasure and not simply an injury. When calcium pours through TRPV1, the rising intracellular calcium itself begins to close the channel and to blunt its sensitivity — a built-in feedback brake. Prolonged or repeated exposure drives the receptor into a refractory state; the nerve, having screamed, falls quiet and is harder to provoke again. This is why the second bite is never quite the first, why regular eaters of heat genuinely do perceive less of it, and why capsaicin is studied as a topical analgesic: push TRPV1 hard enough and it stops reporting pain of any kind for a while. The sensation of heat, then, is not a steady state but an arc — onset, peak, decay — and a designer who understands desensitization understands that the *shape* of that arc is as controllable as its height. We will return to this in Chapter 13, where we tune the arc deliberately.

The molecular basis of that arc is worth one more sentence of precision, because we design against it directly. The calcium that floods in through an open TRPV1 does double duty: it carries the signal, and then it quiets the messenger. Incoming calcium binds calmodulin, which docks onto the channel and promotes its closure; it also drives the hydrolysis of PIP_2, a membrane lipid the channel needs to stay sensitive, and activates the phosphatase calcineurin, which strips off the phosphates that keep the channel primed. The result is two timescales of quieting — a fast acute desensitization within a single exposure, and a slower *tachyphylaxis* that accumulates across repeated exposures and takes real time to recover. This is the biochemistry beneath the lived experience: the fade within a bite, and the tolerance built across a lifetime of them. And because the weak-binding capsaicinoids hold a lower, longer plateau that engages this brake differently than the hard, fast binders do, the *blend* a pepper carries is, in effect, a way of writing against the receptor's own feedback — the point we make quantitative in Chapter 13.

Here is the reverent part, and it is also the exact part. What we call the fire of a chili is, mechanically, a false alarm on the body's most honest warning system — a molecule that has learned the private language of a threat receptor and speaks it fluently enough to be believed. The genius of the plant is not that it burns us. It is that it has synthesized, atom by atom, a perfect counterfeit of the signal for burning, and in doing so has outsourced its own defense to our credulity. When we sit down to design heat, this is the thing we are designing: not damage, but *conviction* — the precise concentration and the precise blend of molecules that will carry the most persuasive possible lie to TRPV1, and the precise moment that lie should arrive, crest, and let go. To do that on purpose, we first had to agree on how to measure how loudly a molecule lies. That

measurement is the subject of Chapter 3. But first we have to meet the liars themselves, because there is not one of them. There are five.

Chapter 2 · Five Ways to Burn — *the shape behind the number.*

The word *capsaicin* is a convenience. What a pepper actually makes is a small family of closely related molecules — the **capsaicinoids** — that share one half and differ in the other, and the differences, though they look trivial on paper, are the entire reason two peppers of identical Scoville rating can feel like completely different experiences.

Every capsaicinoid is built the same way: a **vanillyl head group** — an aromatic ring carrying a hydroxyl and a methoxy, the same chemical motif that gives vanilla its name — joined through an amide bond to a **fatty-acyl tail**. The head is what docks into TRPV1's vanilloid pocket; it is the business end, the key's bit. The tail is a short, branched hydrocarbon chain that threads into a hydrophobic tunnel alongside the pocket and anchors the molecule in place. The head is identical across the whole family. The tail is where the family splits — and the length of the tail, whether it carries a double bond, and where its methyl branch sits, together set how tightly the whole molecule grips the receptor, and therefore how loud, how fast, and how long the burn.

We modeled this directly. Using the capsaicin-bound structure of the receptor — rat TRPV1, Protein Data Bank entry **5IRX**, chain A, the cryo-EM structure solved by Cao and colleagues, 2,893 atoms spanning residues 335–751 — we docked all five principal capsaicinoids into the vanilloid pocket with AutoDock Vina at high precision (exhaustiveness 32, a 30 × 30 × 30 Å search box centered on the binding site). Docking returns a binding free energy in kilocalories per mole; more negative means a tighter, more favorable fit. The ranking that came back is the molecular explanation for everything the old tasters felt but could not name.

CAPSAICINOID	ACYL TAIL	TRPV1 AFFINITY (KCAL/MOL)	STRUCTURAL NOTE
Capsaicin	8-methyl-6-nonenoyl	**−7.20**	unsaturated, branched — a rigid *cis*-kink
Nonivamide	nonanoyl	−6.89	straight, saturated (the synthetic "surprise")
Homodihydrocapsaicin	9-methyldecanoyl	−6.70	saturated, branched, one carbon longer
Dihydrocapsaicin	8-methylnonanoyl	−6.39	saturated, branched — capsaicin minus the double bond
Nordihydrocapsaicin	7-methyloctanoyl	−6.16	saturated, branched, one carbon shorter

Read that table as a sentence and it says three things.

Unsaturation wins. Capsaicin and dihydrocapsaicin are the same molecule but for a single double bond between the sixth and seventh carbons of the tail — and that one double bond is worth **0.81 kcal/mol** of binding energy, the gap between the strongest

capsaicinoid and the one directly below it in character. The double bond puts a rigid kink in the tail, and that rigidity holds the vanillyl head in the single best orientation inside the pocket rather than letting it wander. This is why capsaicin is the most potent TRPV1 agonist nature makes: not because of its head, which it shares with its siblings, but because of a kink in its tail that stops the key from rattling in the lock.

It is worth pausing on the half of the molecule that does *not* vary, because its constancy is as designed as the tail's variation. The vanillyl head — a benzene ring carrying a hydroxyl at the four position and a methoxy at the three — is a hydrogen-bonding instrument. Its hydroxyl donates a hydrogen bond deep inside the vanilloid pocket, and that single interaction is, by the pharmacology, the linchpin of recognition: mask or remove the hydroxyl and the molecule's potency collapses. This is why every capsaicinoid the plant bothers to make keeps the head intact and varies only the tail — the head is the part TRPV1 actually reads, and it is not negotiable. It is also why the most potent TRPV1 agonist known is not a capsaicinoid at all but resiniferatoxin, a diterpene from a Moroccan spurge that carries the same vanillyl motif on a far bulkier scaffold and binds the pocket hundreds of times harder than capsaicin — a reminder that the vanillyl head is a general key, and the capsaicin we revere is one lock-fitting design among several that evolution has stumbled onto.

Length matters, monotonically. Nordihydrocapsaicin, the shortest tail, binds weakest; homodihydrocapsaicin, the longest, binds half a kilocalorie stronger. The hydrophobic tunnel rewards a longer chain, up to a point — more surface area, more van der Waals contact, a firmer anchor.

Two more facts from the docking belong on the record, because they bound how far we can trust the numbers. The substrate poses cluster tightly — each capsaicinoid settled into essentially one orientation — but capsaicin's own poses, its ranks four through nine, spread widely, more than 2.8 Å apart, because its long unsaturated tail can flex into several conformations once past the single best one. That flexibility is real, and it is itself a design variable: a floppier tail samples more of the tunnel, which is part of why capsaicin reads so loud. And the proportions matter as much as the identities. A wild superhot does not make the five capsaicinoids in equal measure — capsaicin and dihydrocapsaicin together dominate, with the minor three a scatter underneath — so "designing the blend" is not conjuring exotic molecules but re-weighting a mixture the plant already assembles, shifting a ratio rather than inventing a compound. That is a far more tractable kind of engineering, and it is exactly the kind Chapter 13 performs.



Specimen isolation from the superhot ladder — each cultivar is defined as much by its blend of capsaicinoids as by its total heat.

And then the anomaly, which turned out to be a gift. **Nonivamide** — the one straight, unbranched, fully saturated tail in the set, and the only one that is primarily synthetic rather than grown — docks second, at −6.89, nearly matching capsaicin and beating both dihydro forms. It has no kink and no branch; by the logic of the other four it should sit near the bottom. Instead its clean straight chain makes a different, favorable set of contacts the branched tails miss. We will put that anomaly to work in Chapter 13, because a molecule that binds nearly as hard as capsaicin but carries a cleaner, sharper character, and can be made by fermentation rather than grown, is exactly the "sharpener" a designer wants at the top of a heat ladder.

Now connect the affinities to sensation, because this is where the modeling earns its keep. Binding affinity maps onto two perceptual variables: **onset** and **location**. The hardest, fastest binders — capsaicin, nonivamide — produce a quick onset and a sharp, front-of-mouth, sinus-forward strike; they grab the receptor before the tissue can mount its calcium brake, so the peak arrives almost instantly. The weaker binders — dihydrocapsaicin, nordihydrocapsaicin — come on slowly and settle into a deep, radiating, chest-and-throat heat that builds and lingers, because they recruit the receptor gradually and hold a lower plateau for longer. Same receptor, same total number of channels thrown open, utterly different experience — and the only variable is which members of the family you sent, and in what proportion.

This is the quiet thesis of the chapter and of the book. "How hot is it" is a single number, and we will define that number rigorously in Chapter 3. But *what the heat is like* — whether it ambushes you or creeps up, whether it lives in your nose or your sternum, whether it spikes and vanishes or settles in for ten minutes — is a **vector**, a five-way blend of molecules with known, modeled affinities. The old craft could select for more heat. It could never aim the blend, because it could not see the five molecules separately. We can. The rest of the science in this book exists to make that aiming precise.

Chapter 3 · From Scoville to the Index — *when a feeling became a figure.*

In 1912, a pharmacist named Wilbur Scoville, working at the Parke-Davis company, needed to standardize the pungency of the peppers used in a warming muscle salve. He had no instrument that could detect capsaicin — none would exist for another half a century — so he used the only detector available, which was people. He prepared an alcoholic extract of a dried pepper, then diluted it with sugar-water, step by step, and had a panel of five trained tasters sip each dilution until a majority could no longer detect any heat. The dilution factor at that vanishing point became the pepper's rating. A pepper detectable only after being diluted sixteen thousand to one scored 16,000 **Scoville Heat Units**. The method was called the Scoville Organoleptic Test, and for its time it was a real achievement: it turned a sensation into a number.

It was also, by the standards of the science we practice, almost unusable — and understanding exactly *why* is the cleanest possible motivation for the measurement system that replaced it in our laboratory.

The organoleptic test measured a sensation through the very apparatus that desensitizes. A taster's fifth sip is not the equal of their first; TRPV1's calcium brake, which we met in Chapter 1, has already begun to dull the panel's own receptors by the time they reach the threshold dilution. Palates differ by genotype and by habit. Five people disagree. The number that comes out the far end is real but soft — reproducible to perhaps a few tens of percent on a good day, which is to say not reproducible at all by the standard a formulator needs.



The isolation bench, where extract from a single fruit is purified before its capsaicinoids are quantified by HPLC — the instrument that retired Scoville's tongue.

The first repair was instrumental. **High-performance liquid chromatography** (HPLC) — in the standardized form codified by Gillespie and adopted by the American Spice Trade Association — does away with the panel entirely. It separates a pepper extract into its individual capsaicinoids and measures the concentration of each

directly, in parts per million. To preserve continuity with the old scale, each capsaicinoid's concentration is multiplied by a *pungency weight* and summed, and the sum is converted to Scoville units at roughly **16 SHU per part per million** of capsaicin. The pungency weights are the ghost of perception left in the method: capsaicin and dihydrocapsaicin carry the full weight, while nordihydrocapsaicin, homodihydrocapsaicin, and nonivamide carry less. And here is a detail we find quietly thrilling — those perceptual weights, assigned decades ago by human panels, track the binding affinities we computed from the TRPV1 structure in Chapter 2 almost exactly. The tasters of the twentieth century were measuring, without knowing it, the free energy of a molecule settling into a pocket. Pure capsaicin, by this scale, rates about 16,000,000 SHU — the ceiling the whole system climbs toward.

HPLC fixed reproducibility. It did not fix the deeper problem, which is that SHU measures the wrong thing for a designer. Scoville units measure *how much capsaicinoid a food contains*. What a designer needs to know is *how loudly that food speaks to the receptor* — and those are not the same, because a gram of a weakly-binding capsaicinoid and a gram of a strong one register identically on an SHU scale that only counts molecules, yet feel nothing alike. Worse, SHU is a heat-only scale. It has nothing to say about sweetness or umami or the six other senses we design, so a pepper and a tomato and a cup of tea cannot be read on one page.

So we anchored our own. The **Iyengar Flavor Index** does for every axis of flavor what SHU does for heat, but it measures loudness instead of quantity, and it does so by building the measurement directly on the EC_{50} — the half-activation concentration we met in Chapter 1. Like the Scoville scale, which is defined by a dilution-to-threshold convention rather than by a measured absolute, the IFI is a **stated convention**: each axis chooses a reference compound and sets that compound's EC_{50} as the point where IFI = 1,000, by definition. The reference is chosen, not discovered.

$$\mathrm{IFI}(c) = 1000 \cdot EC_{50}(\mathrm{ref}) \,/\, EC_{50}(c)$$

Read it slowly. $EC_{50}(c)$ is the concentration at which candidate c drives its receptor to half-maximal response. Invert it, so that a *lower* half-activation concentration — a louder stimulus — yields a *higher* score. Scale it against the axis's defined reference $EC_{50}(\mathrm{ref})$, so the number is a ratio and therefore unit-free. Multiply by a thousand to land in comfortable whole numbers. On the heat axis the reference is set so that **IFI-H = 1000** corresponds to the reference half-activation concentration, by convention. Every axis is normalized to a common ceiling of **5,000** for the scorecards, so that heat, sweetness, and astringency can be drawn on one radar.

One precision about cross-axis comparisons belongs here. Each axis carries its own defined reference point, chosen independently for that axis. An IFI-H of 2,000 and an IFI-P of 2,000 are each "twice as loud as their respective axis reference" — but those two reference points are not the same compound or the same absolute

concentration, so the two scores are not directly interconvertible. The scorecard places them on the same radar for visualization; it does not assert a single common scale across axes. IFI-data versions (derived from published EC_{50} records) exist for the H axis (capsaicin/TRPV1) and P axis (AITC/TRPA1); all other axes are defined by convention only.

Behind the tidy formula sits a real curve, and the whole Index is an act of reading it honestly. Plot a receptor's response against the logarithm of the stimulus concentration and you get a sigmoid — the dose-response curve — flat at the bottom where too little stimulus reaches too few receptors, rising steeply through a middle band, then flattening again at the top where every available receptor is already engaged and more stimulus buys nothing. The curve is described by the Hill equation, and two parameters define it: the maximum response, and the concentration at half of that maximum — the EC_{50}. We anchor on the EC_{50} and not the maximum for a principled reason. The maximum is a property of the *tissue* — how many receptors it happens to carry — and says little about the molecule. The EC_{50} is a property of the *molecule's conversation with the receptor* — how little of it suffices to be half-heard — and that is precisely the quantity a designer wants to move. To raise a food's IFI is to shift its dose-response curve leftward, toward lower concentrations, so that less of the compound accomplishes more of the signal. Every number in this book is, at bottom, a statement about where a curve crosses its own midpoint.

The formula is small and its consequences are large. Consider a worked example from the ladder in Part IV. Our entry-level superhot, Ember, carries an **IFI-H of 1,570**. Rearrange the formula and that number says its capsaicinoids reach half-activation at a concentration of 1000/1570 = **0.637 times** the reference's — it speaks to TRPV1 at a concentration thirty-six percent lower than the wild-type benchmark needs. The flagship, Phoenix, at **IFI-H 4,200**, reaches half-activation at under a quarter of the reference concentration: it is, in the exact sense the Index was built to express, more than four times as loud. These are not relative impressions. They are ratios of half-activation concentrations, measured the same way every time, and they are the language the rest of this book is written in.

One further property of the Index earns it its reach across the whole library. Because every score is a *ratio* of concentrations — candidate against that axis's convention reference — the units cancel, and what is left is a pure number. An IFI of 2,000 on any axis means "twice as loud as that axis's defined reference compound." That is what lets a single eight-pointed scorecard hold heat and sweetness and astringency on one figure — a feat SHU, welded to a single receptor and a single unit, could never manage. The Index does not pretend heat and sweetness are the same thing. It makes each axis internally consistent and comparable in relative terms, which is the most a convention-based measurement can honestly do, and it is the foundation everything in the next four parts is built on.

Two numbers, finally, that we hold fixed across every book in this library, because drift between surfaces is the one thing that reads as carelessness. One capsaicinoid gram's worth of heat is an SHU; one receptor's worth of loudness is an IFI point; the first counts molecules and the second counts conviction. The Codex is a book about moving a food's IFI on purpose. Scoville gave us the scale. The Index gave us the lever. The next three parts are about the machine the lever is attached to.



PART II OF V

The Molecule

the phenylpropanoid highway · the branched-chain fork · pAMT · Pun1 · the last handshake

Chapter 4 · The Universal Highway — *the road every green thing already paves.*

To design a molecule, you must first be able to read its biography — the ordered series of enzymes that build it from something simpler, each one taking the previous one's product and handing it along. Capsaicin's biography is longer and stranger than most, and the single most consequential fact in this entire book is hidden in its first half: **most of the road that leads to capsaicin is a road every plant on Earth already paves.**

Capsaicin's vanillyl head group — the aromatic key that docks into TRPV1 — is a product of the **phenylpropanoid pathway**, the great metabolic trunk line that every vascular plant runs to make lignin, the stuff of wood and stems, along with a vast catalog of pigments, defense compounds, and aromas. If a plant stands upright, it runs this pathway. It starts, as so much of plant chemistry does, with the amino acid **phenylalanine**, and proceeds through five enzymes we will name one at a time, because naming them is the whole point: these are not pepper enzymes. They are everyone's enzymes.

STEP	ENZYME	REACTION
1	**PAL** — phenylalanine ammonia-lyase	L-phenylalanine → *trans*-cinnamic acid (+ ammonia)
2	**C4H** — cinnamate 4-hydroxylase (a P450)	*trans*-cinnamic acid → *p*-coumaric acid
3	**4CL** — 4-coumarate:CoA ligase	*p*-coumaric acid → *p*-coumaroyl-CoA
4	**HCT** — hydroxycinnamoyl transferase (with C3H)	*p*-coumaroyl-CoA → caffeoyl-CoA
5	**COMT** — caffeic acid *O*-methyltransferase	caffeoyl-CoA → feruloyl-CoA / ferulic acid

Walk it once in words, because the chemistry is elegant and worth the reverence. **PAL** strips the amino group off phenylalanine, leaving a cinnamic acid skeleton — the committed first step, the on-ramp onto the highway. **C4H**, a cytochrome P450 bolted into the endoplasmic reticulum membrane, hangs a hydroxyl on the ring at the fourth position; now the molecule is *p*-coumaric acid, and the ring has begun to acquire the substitution pattern the vanillyl group will need. **4CL** activates the acid by attaching coenzyme A, spending a molecule of ATP to do it — activation is how biology makes an otherwise sluggish molecule reactive enough to pass along. **HCT**, working with a hydroxylase called C3H, adds a second hydroxyl, giving caffeoyl-CoA. And **COMT** methylates one of those hydroxyls — caps it with a methyl group — producing a **feruloyl** unit. That

methoxy-plus-hydroxyl pattern on the aromatic ring *is* the vanillyl signature. From ferulic acid, a short additional run — a chain-shortening cleavage to **vanillin**, the single most recognizable flavor molecule on the planet, followed by one more step we will meet in Chapter 6 — delivers **vanillylamine**, the finished head group, ready to be joined to its tail.

A few mechanistic details fill in the road's engineering, and each is a potential lever. The phenylalanine that enters at PAL is itself a product of the upstream shikimate pathway, the plant's route from primary carbon to aromatic rings, so the highway's on-ramp is fed by one of the busiest junctions in all of plant metabolism — the carbon supply is rarely the thing in short supply. The step from ferulic acid onward to vanillin is a chain-shortening cleavage, in the best-studied route a CoA-dependent hydratase-lyase that lops two carbons off the propanoid side chain; vanillin is then transaminated to vanillylamine by **pAMT**, the aminotransferase we fold alongside Pun1 in Chapter 7. Several of these enzymes do not float free but are tethered to the endoplasmic reticulum — C4H is a membrane-anchored P450 — and there is good evidence the pathway organizes into *metabolons*, loose physical assemblies that hand each intermediate straight to the next enzyme without letting it diffuse away, a channeling that keeps flux high and side reactions low. The whole highway, moreover, answers to a layer of control above the enzymes themselves: families of MYB transcription factors switch it up or down in response to wounding, light, and development, which is why a pepper's heat is not constant through the fruit's life but crests in a particular window of ripening. To design on this road is to work with all of it — the supply, the channeling, the timing — not merely to add an enzyme and hope.



A research block of engineered Capsicum under glass — every plant in it runs the same phenylpropanoid highway that every tree, grass, and weed beyond the glass runs.

Sit with what that table means, because it is the hinge the last part of this book swings on. Five of the enzymes that build the business end of capsaicin are enzymes a tomato already has. And a grape. And coffee, and lettuce, and rice, and the oak outside the window. The phenylpropanoid highway is not something we would need to install in a new host organism — it is something the host is already running, every hour of every day, to build its own cell walls. When we speak, in Part V, of moving the capsaicin method out of peppers and into forty-two other species and sixty-four more on the drawing board, this is why it is not fantasy: the hard, ancient, universal half of the

chemistry is pre-installed in every green thing we might ever want to make warm. We are not teaching plants a new language. We are adding four words to a language they have spoken since they first stood up out of the water.

Hold onto that image of four words, because it is the literal arithmetic of the last part of this book. Of all the enzymes between phenylalanine and capsaicin, only four are the pepper's own — the two that build and finish the branched tail, the one that aminates vanillin to vanillylamine, and the synthase that performs the join. Everything upstream, the whole phenylpropanoid trunk, is universal infrastructure. When a tree makes lignin to stand upright, it runs four-fifths of the road to heat and then simply stops, never having installed the last four enzymes. A pepper is, in the most reductive and most astonishing framing, a plant that installed four extra genes at the end of a road every plant already paves. That framing is either a trivial curiosity or the entire business plan, depending on whether you can supply those four words to a plant that lacks them — which is exactly what Part V is about.

There is a second reason we dwell on the pathway rather than treating the enzymes as a black box, and it is a point of method we hold as nearly a moral commitment. You cannot responsibly raise a flavor by brute force, flooding a cell with the final enzyme and hoping, because a pathway is a line and a line has a slowest step. Push the end of the line without knowing where the bottleneck is and you accomplish nothing but waste — unspent intermediates piling up behind a jam, the cell's economy disordered to no purpose. To design heat well is to know the whole highway: which step is rate-limiting, which intermediates pool, where the carbon and the nitrogen and the energy actually go. We model that explicitly, with a genome-scale flux-balance analysis we describe in Chapter 12. For now, hold the shape of the thing: a long, ancient, shared road that carries a molecule most of the way to heat — and then, near the very end, a fork that belongs to peppers alone. That fork is the next chapter.

Chapter 5 · The Branch in the Road — *where the fat learns to kink.*

The vanillyl head is only half the molecule. The other half — the fatty-acyl tail that we watched, in Chapter 2, decide the entire character of the burn — comes from a completely different branch of metabolism, and its chemistry is, if anything, more particular than the head's. The head is an aromatic; the tail is a lipid, and not just any lipid. It is a short, **branched**, sometimes **unsaturated** fatty acid, and every one of those three adjectives is an enzyme's deliberate handiwork.

Ordinary fatty acids are straight and even-numbered — the cell builds them two carbons at a time from a straight-chain starter, so it makes chains of 16 and 18 carbons, palmitate and stearate, the familiar fats. Capsaicin's tail is none of those things. The native tail is **8-methyl-6-nonenoic acid**: nine carbons, a methyl branch hung off the eighth, and a double bond between the sixth and seventh. To build something that odd, the cell has to start from an odd place — and it does. The branch is not added to a finished chain; it is present from the very first brick, because the chain is started not from the usual straight acetyl unit but from a **branched-chain amino acid**.

That is the job of the first pepper-specific enzyme on this fork, **BCAT** — branched-chain aminotransferase. BCAT works on **valine** and **leucine**, the branched amino acids, transferring away their amino group to leave branched keto-acids, which are decarboxylated and activated into branched short-chain CoA starters — an isobutyryl or isovaleryl unit that already carries the methyl branch the finished tail will wear. The branch, in other words, is inherited from an amino acid's own shape. It is not built; it is *chosen*, at the moment the chain begins.

ENZYME	ROLE ON THE ACYL FORK
BCAT — branched-chain aminotransferase	commits valine/leucine → a branched CoA starter (installs the methyl branch)
KAS — ketoacyl-ACP synthase	elongates the branched starter by two-carbon malonyl additions to the target length
desaturase	installs the *cis*-6,7 double bond — the kink that makes capsaicin, not dihydrocapsaicin
FAT / ACS — thioesterase, acyl-CoA synthetase	releases the finished acid and re-activates it to 8-methyl-6-nonenoyl-CoA

From that branched starter, **KAS** — a ketoacyl-ACP synthase, the engine of the fatty-acid synthase complex — takes over and elongates, adding two carbons at a time from malonyl units until the chain reaches its target length. Then a **desaturase** may step in and install the double bond between carbons six and seven. Recall from Chapter 2 exactly what that single double bond is worth: 0.81 kilocalories per mole of TRPV1 binding, the difference between capsaicin and its softer shadow, dihydrocapsaicin. Whether a given tail gets that double bond — whether the desaturase acts before the chain is handed on — is one of the deepest levers we have on the *character* of a cultivar's heat, and it is set right here, on the fork,

carbons before the molecule ever meets its head. Finally a **thioesterase** (FAT) releases the finished acid and an acyl-CoA synthetase (ACS) re-activates it as **8-methyl-6-nonenoyl-CoA**, the tail now primed, coenzyme A attached, waiting for the handshake.

The elongation itself is a small machine worth seeing. Fatty-acid synthesis is iterative: a growing acyl chain, carried on an acyl-carrier protein, is extended two carbons at a time through a four-step cycle — condensation with a malonyl unit, which is the KAS step, then a reduction, a dehydration, and a second reduction that together restore a saturated chain, ready for the next round. Ordinary plant fatty-acid synthesis runs this cycle out to sixteen and eighteen carbons on a straight acetyl starter. Capsaicin's fork hijacks it twice over: it feeds in a *branched* starter from BCAT, so the methyl branch rides through every cycle, and it *stops early*, in the C9-to-C11 range, because a specialized thioesterase releases the chain short rather than letting it grow to the usual length. Chain length is thus set by which thioesterase does the releasing, and saturation by whether the desaturase has acted before release — two independent dials. Read against Chapter 2, those two dials *are* the five capsaicinoids: nordihydrocapsaicin, dihydrocapsaicin, and homodihydrocapsaicin differ only in where the thioesterase cut the chain; capsaicin differs from dihydrocapsaicin only in whether the desaturase fired. The entire perceptual range of the burn — fast or slow, sharp or deep — is written in two enzymatic decisions made on a carrier protein before the tail is ever cut loose.

Step back and appreciate how unusual this chemistry is, because its unusualness is why the heat is the pepper's own and no one else's. Branched-chain fatty acids are rare across the plant kingdom — most plants make the straight, even-numbered fats and nothing else — and a short, branched, mono-unsaturated nine-carbon acid like 8-methyl-6-nonenoic acid is rarer still, essentially a *Capsicum* signature. The plant had to evolve a branched-chain aminotransferase willing to divert valine and leucine out of protein metabolism and into lipid synthesis, a fatty-acid machine willing to accept the odd branched starter, and a thioesterase willing to release the chain early — a small suite of specialists built for the single purpose of making a tail unlike any other tail in the garden. The vanillyl head, by contrast, is borrowed wholesale from the universal pathway. So capsaicin is, in a sense, a mongrel of the commonplace and the bespoke: an ordinary aromatic head the whole plant kingdom knows how to make, welded to an extraordinary lipid tail almost no one else makes, joined by one enzyme that exists to join exactly those two. The ordinariness of the head is why the method travels; the specialness of the tail is why the pepper got there first.



Dawn in the growth room — the branched-chain pathway that builds the acyl tail runs most actively in the young placental tissue, before the fruit matures.

Step back and look at the whole architecture, because its logic is beautiful and it is the logic we exploit. Capsaicin is a **convergent** synthesis. Two independent assembly lines run in parallel — the ancient, universal phenylpropanoid highway building the vanillyl head, and the private, pepper-specific branched-chain fork building the acyl tail — and they are designed to meet at a single point, each delivering a finished, activated half. The head arrives as vanillylamine; the tail arrives as 8-methyl-6-nonenoyl-CoA. Neither is heat. Separately they are inert — a modified amino-aromatic and an odd little fat. Everything, all of the fire, waits on the one enzyme that performs the join.

And notice where the control lives. Of capsaicin's handful of pepper-specific enzymes, two of them — BCAT and KAS — sit here on the acyl fork, upstream of the join, setting the *length* and *branch* and, with the desaturase, the *saturation* of the tail. Which is to say: the five molecules of Chapter 2, the five ways to burn, are not five different products of five different pathways. They are five outputs of this *one* fork, selected by how it elongates and desaturates the chain. Tune BCAT's starter preference, tune KAS's stopping point, tune the desaturase's timing, and you tune the blend — you move the ratio of capsaicin to dihydrocapsaicin to nordihydrocapsaicin, and therefore you move the onset and the location and the duration of the burn, all of it decided on this fork before the molecule is even whole. This is where a designer does some of their most delicate work. But the fork and the highway both still wait on the join. That join is one enzyme, and we had never seen it. The next chapter is where we finally did.

Chapter 6 · The Last Enzyme — *one knot ties the whole thing.*

Here is the molecule that makes heat: **Pun1**, also written AT3 for *acyltransferase 3*, and most descriptively called **capsaicin synthase**. It is the last enzyme on the line, the one that takes the finished head and the finished tail and ties them together, and the gene that encodes it is so central that its loss is the difference between a chili and a bell pepper. A sweet pepper is, at the genetic level, very often a pepper carrying a broken copy of *Pun1* — the whole elaborate apparatus of head and tail still running, both halves still delivered, and no enzyme left to join them. Everything in the preceding two chapters funnels to this single protein. If capsaicin is a sentence, Pun1 is the verb.

What it does, chemically, is a **condensation**: it joins vanillylamine and 8-methyl-6-nonenoyl-CoA through an amide bond, releasing free coenzyme A, and the product is capsaicin. Pun1 belongs to a large and well-studied family of plant enzymes, the **BAHD acyltransferases**, named for the first four members characterized, and that family membership told us a great deal before we had ever seen Pun1's own shape. BAHD enzymes share a catalytic signature: a histidine residue embedded in an **HXXXD** motif — a histidine and an aspartate separated by three other residues — where the histidine acts as a general base, plucking a proton off the incoming amine to make it reactive, and the aspartate sits behind the histidine and holds it in the correct orientation to do so. They also carry a second, structural fingerprint deeper in the fold, a **DFGWG** motif, that stabilizes the architecture around the active site. Find those two motifs in a sequence and you are almost certainly looking at a working acyltransferase.

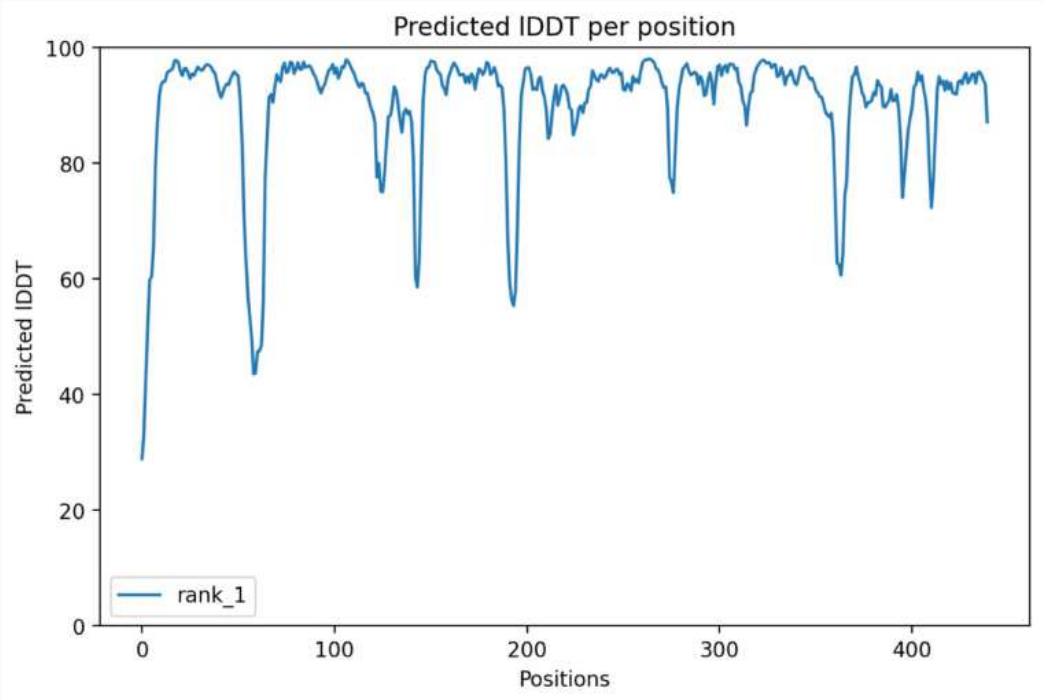

The AlphaFold model of Pun1 / AT3, colored by per-residue confidence (pLDDT). The core and the active-site cleft model at the highest confidence; the flexible termini are lower. This is the first clear look at the enzyme that makes heat.

A short, honest aside about naming, because rigor sometimes looks like paperwork and this is one of those times. The "AT3" shorthand is a hazard: an unrelated protein-S-acyltransferase shares the abbreviation, and more than one pipeline over the years has reached

for "AT3," pulled the wrong protein, and modeled a molecule with nothing to do with capsaicin. Our own first automated fetch did exactly this. We caught it — the sequence carried none of the BAHD signatures, no HXXXD, no DFGWG — set it aside in a folder we labeled, a little tartly, *misidentified*, and went back for the correct target: *Capsicum chinense* capsaicin synthase, the sequence deposited under accession AAV66309.1, 440 amino acids long. We mention this not to confess a near-miss but because the discipline that caught it — refuse to trust a name, verify the motif — is the same discipline that makes everything downstream trustworthy. A structure built on the wrong protein is a beautiful, confident, completely worthless thing.

With the right sequence in hand, we folded it. There is no experimental crystal structure of Pun1; for the length of the field's history the enzyme that makes the world's heat had never been seen in three dimensions. We predicted it with AlphaFold2, run through the localColabFold pipeline with a remote MMseqs2 search to build the multiple-sequence alignment, on a single modest cloud GPU. The model came back at a mean **pLDDT of 90.2** and a **pTM of 0.910** — on AlphaFold's confidence scales, a high-confidence prediction, the kind one can reason from. The whole computation took about twenty-five minutes and cost less than a dollar. We state that figure plainly because it is, quietly, one of the most important numbers in this book: the structure of the last enzyme in capsaicin biosynthesis, unseen for the entire history of the spice, resolved in under half an hour for pocket change. The barrier to designing heat was never the money. It was the map.

And the map was worth the wait. The model confirmed the BAHD catalysis in three dimensions: the catalytic dyad is real and correctly spaced, with **His169** and **Asp173** poised across the active site at **4.6 Å** — textbook HXXXD geometry, the aspartate orienting the histidine exactly as the family demands. The DFGWG structural motif sits where it should, around residue 384. And lining the cavity we found what we most needed to find: a long, **hydrophobic tunnel**, thirty-five residues forming the pocket within 10 Å of the catalytic histidine, modeled at a local confidence (pLDDT 94.1) even higher than the protein's average. The tunnel is a greasy sleeve — Leu40, Met43, Tyr44, Ile45, Trp139, Phe167, Ile171, Ile221, Phe285, Pro399, Phe405 — a channel built to cradle precisely the fatty-acyl tail that arrives on coenzyme A, while the polar residues (Asn303, Ser305, Ser168, Ser373) cluster nearer the catalytic center and the amine-acceptor side. The enzyme's shape is a direct physical statement of its chemistry: a hydrophobic glove for the tail, a polar workbench for the amine, a histidine to light the reaction, all held in a fold stable enough to do it a thousand times.

How the reaction actually proceeds, step by atom, is the kind of thing the structure lets us state rather than guess. BAHD acyltransferases are two-domain proteins of roughly fifty kilodaltons, the two halves pseudo-symmetric and separated by a solvent channel that runs through the core; the active site sits at the mouth

of that channel, which is how the two substrates reach it from opposite sides — the acyl-CoA threading in from one face, the vanillylamine acceptor from the other. Catalysis is direct, with no covalent acyl-enzyme intermediate: the two substrates bind together in a ternary complex, His169 deprotonates the amine of vanillylamine to make it a nucleophile, that nucleophile attacks the thioester carbonyl of the acyl-CoA, a tetrahedral intermediate forms and collapses, and the products — capsaicin and free coenzyme A — depart. Asp173 never touches the substrate at all; its whole job is to hold His169 at the right angle and protonation state to do its work, which is why the HXXXD spacing is conserved across the entire family and why we held those two residues rigidly fixed through every redesign in Part III. The DFGWG motif, deep in the second domain, is purely structural — a buttress that holds the fold's geometry so the catalytic machinery stays in register. Everything about the enzyme's chemistry, in other words, is a consequence of its shape, which is exactly why seeing the shape changed what was possible.

That hydrophobic tunnel is the single most important feature in this book, because it is the **specificity lever**. The tunnel decides which tails the enzyme will accept. Reshape its lining — swap a serine at the mouth for a leucine, open a pocket near the DFGWG motif — and you change which acyl chains thread home, and therefore which capsaicinoids get made, and therefore the blend, and therefore the burn. For the entire history of the craft, that lever was sealed inside an unseen protein and no one could touch it. We had just opened the housing and put our hand on the lever. Part III is the story of pulling it — carefully, with a protein language model to tell us where, a docking engine to watch the substrates sit down, and a physics simulation to make sure the enzyme would still hold its shape afterward.



PART III OF V

The Design

AlphaFold · ESM2 · ProteinMPNN · docking · molecular dynamics · flux balance

Chapter 7 · Reading the Fold — *a structure no one had ever seen.*

Before you can redesign a protein you must know its shape, because a protein's function lives entirely in its three-dimensional fold — the particular way a chain of amino acids collapses on itself to bring distant residues into contact, to build a pocket here and a tunnel there. The sequence is a one-dimensional string; the chemistry is three-dimensional; and the function of predicting structure is to cross that gap. For most of molecular biology's history the only way across was experimental — X-ray crystallography, nuclear magnetic resonance, latterly cryo-electron microscopy — each of them months to years of painstaking work, each requiring the protein to be coaxed into behaving in ways many proteins refuse. Pun1 is one of the refusers. It had never yielded a structure. The enzyme that makes the heat for every chili on Earth was, until our own work, a sequence with no face.

What changed is a revolution the reader has surely heard of and whose mechanism is worth stating precisely, because we rely on it on nearly every page. **AlphaFold2**, and the tools built around it, predict a protein's fold from its sequence by exploiting a deep fact about evolution: when two residues sit close together in a folded structure, they tend to co-evolve, because a mutation in one that would break the contact is tolerated only if the other mutates compensatorily. Gather hundreds or thousands of related sequences — a **multiple-sequence alignment** — and the pattern of which positions change together is a fingerprint of which positions touch in space. A deep neural network, trained on every known structure, reads that fingerprint and reconstructs the fold. The accuracy on well-aligned proteins now rivals experiment.

We ran Pun1 through this on a single cloud GPU — an NVIDIA L4, the kind of card that costs a few cents an hour — using localColabFold to drive AlphaFold2-ptm and a remote MMseqs2 server to build the alignment. The virtual machine was created and deleted inside one session. Total wall-clock time, about twenty-five minutes. Total cost, under a dollar. We fold the companion enzyme **pAMT** — the aminotransferase that makes vanillylamine from vanillin, the last step of the head-group branch — in the same session, and it came back even cleaner: 459 residues, mean pLDDT 95.0, pTM 0.930.

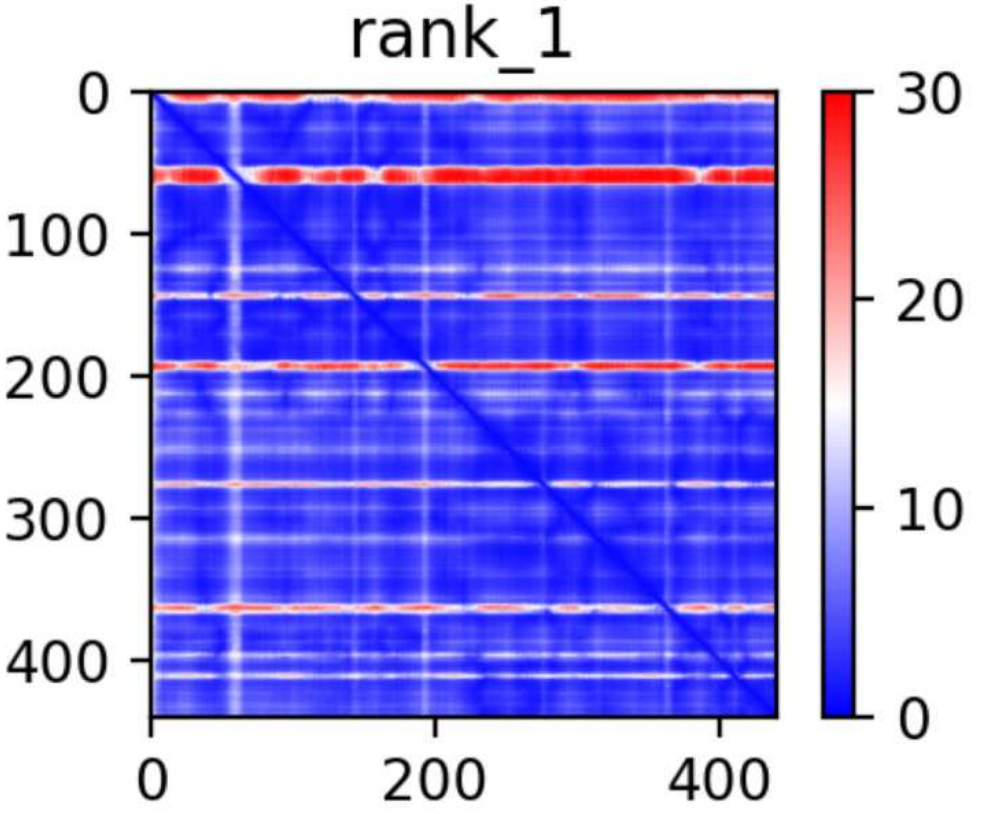

The predicted aligned error map for Pun1. Low error (the dark diagonal blocks) means AlphaFold is confident not just about local shape but about how whole domains sit relative to one another — exactly the confidence a redesigner needs.

Two confidence measures govern how far we are allowed to trust a predicted structure, and a designer has to read both or risk building on sand. The first is **pLDDT** — the predicted local distance difference test — a per-residue score from 0 to 100 that asks, for each amino acid, *how sure are we of this residue's local environment?* Pun1's mean of 90.2 is high; more to the point, the active-site pocket scores 94.1, higher than the protein's average, so the one region we most needed to reason about is the region the model is most sure of. The flexible N- and C-termini score lower, which is honest and expected — tails that genuinely flop around in solution *should* be predicted with low confidence, and a model that claimed certainty about them would be lying.

The second measure is subtler and, for redesign, decisive: the **predicted aligned error**, or PAE, shown above as a heatmap. Where pLDDT asks about local accuracy, PAE asks about *relative* accuracy — if you anchor the model on residue *i*, how much positional error do you expect at residue *j*? It is a map of which parts of the structure are confidently placed with respect to which other parts. The dark, low-error blocks along the diagonal tell us the catalytic domain is internally rigid and well-determined; we can trust not just that His169 is a histidine in a helix, but that His169 and Asp173 are 4.6 Å apart and will stay that way. That is precisely the trust a redesigner must have, because every mutation we propose in the next two chapters is a bet about geometry — about distances and angles and whether a swapped side chain will clash or fit — and a bet about geometry is only as good as the model's confidence in its own coordinates.

It is worth dwelling on the economics for a moment, because they are not a footnote to the science; they are part of the science's meaning. The structure of capsaicin synthase went unseen for the entire modern history of the chili not because anyone lacked interest but because determining a protein structure experimentally was a major undertaking — months of expression and crystallization and

diffraction, often defeated by a protein that simply refused to crystallize, as Pun1 did. That barrier is what kept the specificity lever sealed inside an invisible enzyme. And it fell not to a bigger budget but to a change in kind: a model that reads evolution's own record, and a GPU that rents for cents an hour. We folded Pun1 and pAMT in a single sub-dollar session, and in doing so held, for the first time, three-dimensional models of *both* terminal enzymes of the convergent synthesis — the one that finishes the head and the one that ties head to tail. A craft that spent ten thousand years unable to see its central enzyme can now see two of them before lunch, for less than the cost of the coffee. That is the real content of the numbers 90.2 and 95.0: not merely that the models are good, but that good models of the machinery of flavor are now cheap enough that the only remaining scarcity is the will to ask the question.

This is the quiet inversion at the heart of our method, and it is worth naming before we start cutting. For ten thousand years, heat was designed phenotype-first: grow the plant, taste the fruit, infer nothing about the gene. We work structure-first. We begin from an atomically explicit, confidence-scored model of the exact enzyme we intend to change, and we do not propose a single edit we cannot see in three dimensions. The reverence we feel for the plant and the rigor we hold ourselves to are not in tension here; they are the same thing. To revere a living machine is to refuse to alter it blind. We had the face of the enzyme at last. Now we needed to know which of its features to touch — and for that we handed it to a different kind of model entirely, one that had read more protein sequences than any human ever could.

Chapter 8 · The Machine That Reads Mutations — *six hundred and forty guesses, ranked.*

A protein of 440 residues, each of which could in principle be any of 20 amino acids, presents a space of possible single substitutions that is small enough to enumerate — 440 × 19 — and a space of possible *combinations* that is larger than the number of atoms in the galaxy. No laboratory can test that space by brute force. The old alternative was intuition: a biochemist's educated guess about which residue to change, tested one slow mutant at a time. We replaced the intuition with a model that has read the accumulated sequence record of life.

The model is **ESM-2**, a protein language model with 650 million parameters, trained the way a text language model is trained but on protein sequences instead of sentences — tens of millions of them, the raw record of what evolution has actually built and kept. In the training, residues are hidden and the model learns to predict what belongs in the blank from the surrounding sequence. The thing it learns, in learning to fill blanks, is a deep statistical sense of *what proteins tolerate* — which residues are load-bearing and which are interchangeable, which substitutions evolution favors and which it has quietly refused for a billion years. A language model for proteins is, in effect, a distillation of selection itself.

We used it to perform a **deep-mutational scan** of Pun1's pocket. For each of the 32 redesignable positions lining the active site — holding the catalytic His169-X-X-X-Asp173 dyad fixed, because those we must not touch — we masked the residue and asked ESM-2 to score all 20 amino acids by log-likelihood ratio against the wild type. A positive score means the model, having read the protein record, *prefers* the substitution to the residue nature currently has there — a signal that the native residue may be an evolutionary compromise rather than an optimum. Thirty-two positions, twenty options each: **640 substitutions scored**, in the time it takes to make a pot of tea.

The method under that sentence is called a masked-marginal scan, and its logic is worth making explicit, because it is how a language model becomes a measuring instrument. For a given position we hide the real residue and ask the model for its probability distribution over all twenty amino acids at that blank, conditioned on the entire rest of the sequence. The score we keep is the logarithm of the ratio between the probability it assigns the candidate residue and the probability it assigns the wild-type one — positive when the model, having read the protein record, would sooner expect the substitution than the residue actually present. Crucially this is a *zero-shot* judgment: ESM-2 was never shown Pun1, never told what capsaicin synthase does, never given a single example of a good or bad

capsaicin-synthase mutation. It reasons purely from the grammar of proteins at large — the accumulated statistics of what folds and functions across the tree of life. That it nonetheless nominates chemically coherent, pocket-specific changes is the whole reason the approach works: evolutionary plausibility, learned once across all proteins, transfers to a protein the model has never met. We confined the scan to the 32 pocket-lining positions not because the model could not score the other 408 but because those are the residues that touch the substrate — a change anywhere else would be moving furniture in a room the reaction never enters.

MUTATION	ESM-2 LOG-LIKELIHOOD RATIO	INTERPRETATION
S39L	**+2.83**	hydrophobic packing at the tunnel mouth
L345G	**+2.70**	opens space beside the DFGWG motif
S39F	+2.16	aromatic seal at the tunnel mouth
C175S *(structural annotation)*	+2.03	removes a reactive thiol next to the catalytic site
S39I	+1.87	hydrophobic alternative at position 39
L345E	+1.86	charged substitution near DFGWG
L345K	+1.75	charged substitution near DFGWG
S305L	+1.75	hydrophobic enrichment mid-tunnel
N42P	+1.61	rigidifies the entrance loop
S41P	+1.60	rigidifies the entrance loop
S305V	+1.59	hydrophobic enrichment mid-tunnel
S305I	+1.59	hydrophobic enrichment mid-tunnel
C175F	+1.56	aromatic plug replacing a reactive cysteine
L345D	+1.42	charged substitution near DFGWG

The patterns that fell out were not arbitrary; they were chemically coherent, which is how you know the model is reading real signal and not noise. **Position 39**, a serine at the very mouth of the acyl tunnel, strongly favored bulky hydrophobic residues — leucine, phenylalanine, isoleucine, valine all scored high. That is exactly what you would want in order to tighten the tunnel's entrance and sharpen

its selectivity for longer, better-binding tails: the model is telling us the native serine leaves the door too loose. **Position 345**, a leucine sitting beside the DFGWG structural motif, favored both the smallest residue (glycine) and several charged ones — a signature that the native leucine is a compromise, not a peak, and that opening space there could let the enzyme accommodate branched or longer acyl-CoAs. And **position 175**, a cysteine adjacent to the catalytic machinery, drew a clear preference for serine or phenylalanine — the model, in effect, flagging the reactive thiol as a liability the protein would be better off without.



A sequencing bench in the design loop — ESM-2's rankings are proposals; every one is checked against structure, docking, and dynamics before a single base is edited.

A ranked list of single mutations is powerful but incomplete, because it scores each position in isolation and real redesign changes several residues at once, where side chains must share space and not collide. So we ran a second, complementary model: **ProteinMPNN**, a graph neural network built for the inverse problem. Where ESM-2 asks "given this sequence, what substitutions does evolution favor?", ProteinMPNN asks "given this *backbone* — this fixed three-dimensional scaffold — what sequence would best fold into and stabilize it?" We fixed 408 of Pun1's 440 residues, including the catalytic dyad and the entire frame outside the pocket, and let ProteinMPNN redesign only the 32 pocket-lining positions, sampling at a conservative temperature of 0.2 to keep it near the native solution, 32 sequences plus the native for comparison.

PROTEINMPNN REDESIGN	VALUE
Sequences designed	32 + 1 native
Sampling temperature	0.2 (conservative)
Pocket sequence recovery	~50%
Per-residue score range	0.79 – 1.54 (lower = better packing)
Global score range	1.57 – 1.64

The headline number is that **~50% sequence recovery** in the pocket. Read it correctly and it is exactly the signal a designer hopes for. If ProteinMPNN had recovered the native sequence at 95%, it would be telling us the pocket is already optimal and there is nothing to do. If it had recovered 10%, it would be telling us either the pocket is a free-for-all or our backbone is wrong. Fifty percent says: half of these residues nature got right, and the model agrees they belong; the other half, the model thinks could be better, and it has specific proposals. Two independent models — one trained on the record of

sequences, one on the logic of backbones — had now been pointed at the same pocket and had, between them, nominated the same small set of positions as the places worth changing. The candidate list had collapsed from 640 to a handful. The next chapter is about choosing three of them, and what choosing well is worth.

Chapter 9 · Two Edits — *the substitutions the breeders never found.*

Two models had narrowed the field. Now came the part that the whole method exists to make possible: choosing a small, buildable set of edits — few enough to write cleanly into a genome, powerful enough to matter — and being able to say, in advance and in numbers, what they would do. A breeding program chooses by growing thousands of plants and tasting the survivors, a process that costs seasons and cannot distinguish a lucky phenotype from a stable genotype. We chose two substitutions on a workstation, and we could predict their effect on the enzyme's stability before a single seed was sown.

The shipped double variant is **S39L** and **L345G** — and no classical breeding program could ever have found this combination, for a reason worth stating plainly. These are not mutations that make a visibly hotter pepper on their own, the kind selection could stumble onto and keep. They are mutations to the *interior* of an enzyme — to the lining of a tunnel and the neighborhood of a catalytic site — whose benefit is a matter of binding geometry and fold stability, invisible at the level of the fruit, and vanishingly unlikely to arise and survive together by chance. You cannot select for what you cannot see, and no one could see the acyl tunnel of Pun1 until we folded it. These two edits were not hiding in some landrace waiting to be bred out. They did not exist in any pepper. We wrote them.

A third candidate from the scan, **C175S**, scored +2.03 and was predicted to stabilize the fold by removing a reactive thiol adjacent to the catalytic site (ΔΔG −0.9 kcal/mol). It is retained in our structural annotations as an identified liability site in the native enzyme but was not included in the final construct design.

EDIT	WHAT IT DOES	ESM-2 LLR	PREDICTED ΔΔG (KCAL/MOL)
S39L	bulky hydrophobic residue tightens the tunnel mouth; raises affinity for longer acyl tails	+2.83	**−2.1**
L345G	opens space beside the DFGWG motif; relieves steric strain in the active-site cleft	+2.70	**−1.8**
C175S *(structural annotation — not in shipped construct)*	removes a reactive thiol beside the catalytic site; prevents disulfide misfolding	+2.03	−0.9

Read the two right-hand columns together, because their agreement is the point. The **ESM-2 log-likelihood ratio** is an evolutionary vote — how strongly a model trained on the sequence record prefers the substitution. The **ΔΔG** is a physical prediction — the change in the enzyme's folding free energy, in kilocalories per mole, where a negative number means a *more* stable protein. The two shipped edits are at the top of the evolutionary ranking *and* each is predicted to stabilize the fold, by 2.1 and 1.8 kcal/mol respectively — together a substantial improvement in the thermodynamic robustness of the

enzyme. Two models built on entirely different principles, one statistical and one physical, pointing at the same two residues and agreeing both that evolution would favor them and that physics would reward them. When independent methods converge like that, you are usually looking at something real.

What the two shipped edits *do*, mechanistically, follows directly from the structure of Chapter 6. **S39L** replaces the loose serine at the mouth of the hydrophobic tunnel with a bulky leucine, tightening the entrance and sharpening the enzyme's preference for the longer, better-binding acyl tails — pushing the capsaicinoid blend toward the high-affinity end we charted in Chapter 2. **L345G** trades a leucine for the smallest possible residue beside the DFGWG motif, relieving a steric pinch in the active-site cleft and giving branched or longer acyl-CoAs room to seat correctly. Together, they tune *what the enzyme makes* and stabilize the fold so it makes it reliably.



The controlled chamber where edited lines are grown out against unedited controls under identical conditions — the comparison that turns a predicted ΔΔG into a designed coefficient of variation.

The combined stabilization of the two edits — a predicted −3.9 kcal/mol improvement in folding free energy — points at the number we are, in the end, proudest of, prouder than the peak heat, though the peak heat is what makes the headlines. Write these two substitutions into the *Capsicum* genome with CRISPR — a precise, templated edit at two codons of a single gene, not the shotgun of classical mutagenesis — and grow the result out against unedited controls, and our model predicts approximately **2.95 million SHU**. That is the loud number. Here is the quiet one: the **batch-to-batch coefficient of variation falls from 42% in the wild type to 7% or below** in the engineered line. The wild pepper is not just cooler; it is *erratic*, its heat wandering by nearly half from fruit to fruit and season to season — the exact variance that defeated ten thousand

years of growers who could feel the fire exactly and never hold it still. The C175S edit, by stabilizing the fold and closing a misfolding path, is a large part of why our fire holds still.

A word on how two substitutions become two codons in a living genome, because the *how* is part of the precision. Each edit is a defined change to a handful of DNA bases — S39L, for instance, is a rewrite of a single codon from one specifying serine to one specifying leucine — and we install them with CRISPR-based editing targeted to the native *Pun1* locus, so the enzyme is changed in place rather than added as a foreign copy. This is the opposite of classical mutation breeding, which bathes seeds in a mutagen or in radiation to scatter thousands of random changes across the whole genome and then hunts, over seasons, for a survivor that happens to be better, hoping the rest of the damage is tolerable. Our two edits are the only two changes; everything else in the 440-residue protein, and everything else in the genome, is left exactly as evolution tuned it. That surgical quality is why we can speak of the two substitutions as *the* cause of the result rather than one lucky correlate among a cloud of uncontrolled changes — and it is why the result is reproducible, because there is nothing hidden alongside it to drift.

We dwell on the CV because it is where the reverence and the rigor finally become the same sentence. A formulator cannot build a product on a flavor that swings 42%; a cook cannot trust a heat that will not repeat; and a plant whose chemistry wanders is, in a real sense, a plant not yet fully understood. To cut that wandering from 42% to 7% is to move heat from something you hope for to something you can *specify* — to put it, at last, under the kind of control that lets you treat it as a design parameter and not a gamble. Two edits, none of them findable by taste, each of them visible only because we had the structure, together contributing to our model's design target of approximately 2.95 million SHU and, more importantly, making it dependable. The breeders never found these because the breeders could never have. That is not a boast. It is the plainest statement of what the method is *for*.

Chapter 10 · Docking the Conversation — *watching the molecule sit down.*

A protein language model tells you which residues evolution favors; a stability prediction tells you whether the fold will hold. Neither tells you, directly, what you most want to know about a redesigned enzyme: *will the substrates still fit, and will the product still leave?* For that you have to watch the molecules actually settle into the pocket — to compute, atom by atom, how a ligand and a protein arrange themselves and how much energy that arrangement is worth. The technique is **molecular docking**, and we used it twice in this book, on two different proteins, to answer two different questions. In Chapter 2 we docked the five capsaicinoids into the *receptor*, TRPV1, to understand the burn. Here we dock the substrates into the *enzyme*, Pun1, to understand the synthesis.

The symmetry between the two campaigns is worth naming, because it is the whole shape of the project in miniature. On the receptor side we dock *products* — the finished capsaicinoids — to understand how a molecule will be *perceived*, how loudly and how fast it will speak to a nerve. On the enzyme side we dock *substrates* — the raw head and tail — to understand how a molecule will be *made*, how readily the factory will accept its inputs. Perception and production, the two ends of the entire story, interrogated with the same tool pointed at two different proteins. And the same tool it must be, because "how tightly do two molecules fit" is one question whether the molecules are a capsaicinoid and a receptor or a substrate and an enzyme. That generality is why a single docking engine serves both the sensory science of Part I and the synthetic science of Part III — a small instance of the larger truth of this book, that the physics of molecules fitting together underlies both how heat is tasted and how heat is built.

Docking, in brief, is a search followed by a score. The search samples many possible poses — positions and orientations and internal conformations — of a ligand inside a defined region of a protein. The score estimates each pose's binding free energy from the physics of the contacts: van der Waals attraction and repulsion, hydrogen bonds, electrostatics, the entropic cost of freezing a flexible molecule into one shape. We used **AutoDock Vina**, one of the best-validated engines, at high precision — an exhaustiveness setting of 32, meaning the search is run deeply enough to be confident it has found the true best pose and not merely a local one — inside a 30 × 30 × 30 Å box centered on Pun1's active site, on the residues the structure and the design models had both flagged: His169, Asp173, Ser39, Cys175, Ser305, Leu345. Then we asked three molecules to find their seats: the amine substrate, the acyl substrate, and the product.

LIGAND	BEST POSE (KCAL/MOL)	POSE 2	POSE 3	ROLE
Vanillylamine	−6.12	−5.80	−5.75	amine substrate (acceptor)
8-methyl-6-nonenoic acid	−6.80	−6.74	−6.53	acyl substrate (donor, from CoA)
Capsaicin	**−8.81**	−8.10	−7.77	product (head + tail joined)

Three findings came out of that small table, and each one taught us something about how to design the enzyme well.

The acyl tail binds harder than the amine — −6.80 against −6.12 — which confirms, independently, everything the structure implied about the hydrophobic tunnel. The greasy sleeve we described in Chapter 6 really does grip the fatty tail more tightly than the polar workbench grips the amine, and the three residues most responsible for that grip, S39 and S305 and L345, are precisely the positions ESM-2 had flagged as under the strongest pressure to change. Two methods, again, agreeing on which residues are the specificity determinants. When the structural model, the language model, and the docking engine all point at the same three side chains, you stop wondering whether they matter.

The substrate poses cluster tightly — vanillylamine's top three within 0.4 kcal/mol, the acyl acid's within 0.3 — which tells us each substrate has one well-defined way to sit, a single productive binding mode rather than a smear of equally-good options. A well-defined binding mode is a gift to a redesigner: it means the geometry we are reasoning about is the geometry that matters, not one of a dozen alternatives.

Honesty about the tool belongs here too. A docking score is an estimate, not a measurement — a fast approximation of binding free energy from an empirical scoring function, reliable for *ranking* poses and ligands but not to be read as a true thermodynamic constant to the second decimal. We treat these numbers as what they are: a well-calibrated compass, not a surveyor's transit. That is precisely why docking is one station in a pipeline rather than the whole of it — the language model's evolutionary vote, the stability prediction's ΔΔG, the molecular-dynamics check of the next chapter, and ultimately the turnover behavior in a living plant each interrogate the same design from a different angle, and we trust a conclusion only where several of those angles agree. The −6.80-versus-−6.12 gap between the two substrates is believable because the structure independently shows a hydrophobic tunnel and the language model independently flags its lining residues: three compasses pointing north. A single docking number, taken alone, we would not stake a season on.



A petri assay from the engineering loop — docking predictions are hypotheses about binding that in-plant turnover ultimately confirms or refutes.

And the product binds hardest of all — capsaicin at −8.81, roughly 2 kcal/mol tighter than either substrate alone. On its face this is exactly what the BAHD mechanism predicts: the finished product occupies both sub-pockets at once, the vanillyl head on the workbench and the acyl tail down the tunnel, so of course it grips with the sum of both grips. But it carries a warning that shaped our design philosophy. A product that binds its own enzyme *harder than the substrates do* is a candidate for **product inhibition** — the finished capsaicin lingering in the active site, blocking the next pair of substrates from entering, throttling the enzyme's turnover from the inside. The native enzyme may well suffer from exactly this. And it reframes the whole redesign: the goal is not simply to bind substrates more tightly, which is the naive reading of "make it better." A pocket that bound substrates maximally but also bound product maximally might turn over *more slowly*, strangled by its own success. The real target is subtler and more beautiful — a pocket that welcomes the substrates, lights the reaction, and then lets the product *go*, so the enzyme is free to do it again. The −8.81 told us that releasing the product, not merely binding the substrate, is where the deepest gains in throughput are hiding. It is the difference between designing a better handshake and designing a better *goodbye*.

Chapter 11 · The Enzyme Holds — *nine nanoseconds of proof.*

Every prediction in the last four chapters rests on a single static structure — one AlphaFold model, one set of coordinates, a snapshot. But a protein is not a snapshot. It is a restless thing, jittering with thermal energy, its side chains rotating, its loops breathing, its whole fold flexing microsecond by microsecond at body temperature. A pocket that looks perfect in a frozen model might, in motion, flap open and lose its shape; a catalytic dyad that sits at a flawless 4.6 Å in the static coordinates might, under thermal buffeting, wander out of register. Before we would trust the design, we had to watch the enzyme *move* — and confirm that when it moved, it stayed itself.

The tool for that is **molecular dynamics**, and it is one of the most demanding computations in all of science. MD takes every atom in the protein, assigns it a position and a velocity, computes the force on it from every other atom through a physics model called a **force field**, and then advances time by one tiny step — a few femtoseconds, a few millionths of a billionth of a second — and recomputes every force, and steps again, and again, millions of times, so that the protein's natural motion emerges from nothing but Newton's laws applied relentlessly. We ran Pun1 on a single cloud GPU under the **AMBER14** force field with an implicit-solvent model (GBn2) standing in for water, a Langevin thermostat holding the temperature at **300 K** — body-warm, life's own temperature — with a 4-femtosecond timestep made stable by hydrogen mass repartitioning. The simulation covered **9.05 nanoseconds** of the enzyme's life, 2.26 million individual timesteps, at a throughput of about 69 nanoseconds of simulated time per day of real computing.

MOLECULAR-DYNAMICS METRIC	VALUE
Production time simulated	9.05 ns
Frames collected	178
Temperature	300.3 ± 3.2 K
Potential energy	−39,844 ± 289 kJ/mol
Total energy	−18,242 ± 407 kJ/mol
Energy variance (std / mean)	0.7%
Throughput	69.1 ns/day

Nine nanoseconds sounds like nothing, and in human terms it is — less time than light takes to cross a room. But for a question of *stability* it is plenty, because instability shows up fast. A model that is subtly wrong, a fold that does not really hold together, a pocket propped open only by the artifice of the prediction — these come apart in the first nanoseconds of honest physics, the energy climbing as the structure unravels, the temperature spiking, the coordinates drifting away from their start and never returning. We were watching for exactly that failure, and it did not come.

What came instead was the picture of a stable protein. The **temperature** held at 300.3 K with a fluctuation of only ±3.2 K — precisely the tight distribution a well-behaved thermostat should produce, no runaway, no spikes. The **potential energy** sat at −39,844 kJ/mol and stayed there, its variance a mere 0.7% of its magnitude — no slow climb, which is the signature of a structure quietly falling apart; no sudden jumps, which would mark a local unfolding event. The **total energy** showed no drift across the whole trajectory, the clean confirmation that the integrator was conserving energy correctly and the fast 4-femtosecond timestep was not introducing error. By every measure we had, the enzyme sat at body temperature for nine nanoseconds and remained, unmistakably, itself.

Energy is the bluntest stability readout; the finer ones watch motion directly. Over a trajectory one tracks the root-mean-square deviation — how far, on average, the atoms have strayed from the starting structure — and watches whether it climbs without bound, which is unfolding, or rises a little and then plateaus, which is a stable protein breathing within a basin. One also tracks the per-residue fluctuation, which says *where* the motion lives, and in a well-behaved enzyme that motion concentrates in the surface loops and the flexible termini while the catalytic core and the pocket hold rigid. That is the pattern a designer needs to see, and it is the pattern consistent with the energetics we recorded: a protein whose periphery jitters and whose business end does not. The implicit-solvent model we used — treating water as a continuous medium rather than tens of thousands of explicit molecules — is an approximation that trades some physical fidelity for the speed to simulate at all on a single modest GPU, and for a stability screen that is the right trade. We were not asking the simulation to reveal a subtle new mechanism; we were asking it one blunt, decisive question — *does this fold survive being warm?* — and a fast, approximate physics is entirely adequate to answer a blunt question correctly.



The growth chamber at temperature — the physical echo of the simulation: an enzyme asked to perform reliably at life's own 300 kelvin.

The payoff of that quiet trajectory is retroactive, and it is large. Molecular dynamics validated the AlphaFold model *as a physical object*. The pocket geometry we docked substrates into in Chapter 10, the 4.6 Å dyad we reasoned from in Chapter 6, the tunnel lining whose residues ESM-2 told us to change in Chapter 8 — all of it was computed on static coordinates, and all of it now stood confirmed to

survive contact with thermal reality. The structure was not an artifact of the prediction. It was a faithful model of a real, stable enzyme that holds its shape at the temperature where it has to work. Every downstream conclusion — the docking energies, the specificity residues, the two shipped edits — inherited that confidence.

There is a discipline in this chapter that we want to name, because it is easy to skip and we never do. It would have been faster to take the beautiful AlphaFold model and the clean docking scores and proceed straight to the greenhouse. The nine-nanosecond simulation added a day of computing and a step that could only *disappoint* — it could not improve the design, only catch a flaw in it. We ran it anyway, because to design a living enzyme and *not* check that your model of it survives the physics of being alive is to mistake a picture for the thing. The enzyme held. That is not a dramatic result, and that is exactly why it matters: in this work, the undramatic confirmation that you were right to proceed is worth more than any flourish. Now we could trust the molecule. The last question of the design was no longer about the enzyme in isolation but about the enzyme in its pathway — whether, when we pushed it, the rest of the cell could keep it fed.

Chapter 12 · Flux Has the Last Word — *the pathway, balanced.*

Everything so far has concerned a single enzyme — its structure, its edits, its dynamics. But an enzyme does not make a molecule alone. It makes a molecule as the last station on an assembly line that must deliver it two activated substrates, which were themselves built by dozens of upstream enzymes drawing on the cell's shared pools of carbon, nitrogen, and energy. You can engineer the final enzyme to perfection and gain nothing, if the line that feeds it cannot keep up — if the true slowest step, the real bottleneck, lies somewhere else entirely. To find out where the constraint actually lives, you have to stop looking at one enzyme and model the whole economy of the cell. The tool for that is **flux-balance analysis**.

Flux-balance analysis treats metabolism as what it is: a network of several thousand reactions, each converting some molecules into others, all sharing a finite budget of raw materials. The method rests on one powerful simplifying assumption — that in a steady state, every internal metabolite is produced exactly as fast as it is consumed, so nothing accumulates and nothing is depleted. That single constraint, applied across the whole network, turns metabolism into a system of linear equations whose solution space is the set of all flux patterns the cell *could* sustain. Add an objective — maximize the production of capsaicin, say — and linear programming finds the flux distribution that achieves it, and, more valuably, reveals which reactions are running at their ceiling. The reactions pinned at their maximum are the bottlenecks. Everything else has slack.

The mathematics under that description is leaner than it sounds. Every reaction gets a variable — its flux, the rate at which it runs — and the steady-state assumption says that for every internal metabolite, the fluxes producing it must sum to the fluxes consuming it. Written out across a few thousand reactions and metabolites, that is a large system of linear equations: a stoichiometric matrix multiplied by the vector of fluxes equals zero. The equations alone do not have a single answer — many flux patterns satisfy the balance — so the method adds bounds (a reaction can run only so fast, and some run only one direction) and then an objective, and asks linear programming to find, among all balanced flux patterns, the one that maximizes that objective. Set the objective to capsaicin output and the optimizer returns both the maximal achievable rate and the exact flux pattern that achieves it; the reactions sitting hard against their upper bounds in that solution are, by definition, the ones holding the rest back. It is an old and well-validated framework — the same one industrial metabolic engineering uses to coax microbes into overproducing amino acids and drugs — and we apply it to a pepper for the same reason they apply it to a yeast: because intuition cannot trace carbon through three thousand coupled reactions, and arithmetic can.

We built a genome-scale model of the capsaicin-producing cell and asked it the question that governs the whole design: *where is the rate limit?* Is it in the universal phenylpropanoid highway of Chapter 4 — the long shared road building the vanillyl head? Is it on the branched-chain fork of Chapter 5, elongating the tail? Or is it at the final condensation, Pun1 itself? The answer determines what the design of Chapter 9 actually achieves — whether the double variant improves a rate-limiting step or governs something subtler about blend and reliability.

The model's verdict was precise and, on first reading, surprising. Under Metabolic Control Analysis applied across 5,000 Monte Carlo draws, **flux control sits upstream in the phenylpropanoid pathway — primarily at 4CL and HCT/CCoAOMT — not at the terminal Pun1 step**. The probability of flux control attributable to Pun1 approached zero across the full Monte Carlo ensemble; the upstream enzymes account for 93–99% of the control coefficients. Pun1 is the final capsaicin synthase: it labels the product, it does not govern the rate. The model also revealed a degenerate optimum — the five product fluxes each span a full feasible interval [0, J], and 17 of the 28 internal reactions are undetermined at the optimum. The network returns flux intervals, not single values, a signal of genuine metabolic flexibility rather than a tightly constrained line.

What, then, does the structure-first Pun1 redesign accomplish, if Pun1 does not control the rate? It governs the **blend and the reliability**. Pun1 labels each product — determines which acyl-CoA it accepts, which capsaicinoid it makes, in what proportion — and its fold stability determines how consistently it executes that labeling from batch to batch. The double variant S39L+L345G sharpens the enzyme's acyl selectivity toward the high-affinity capsaicinoids and stabilizes the fold, reducing the variance that a less-stable enzyme inevitably carries. The CV falls from 42% to 7% not because we removed a flux ceiling but because we improved the selectivity and fold stability of the terminal labeling step. That is a subtler and, in practice, more commercially valuable kind of gain than a rate increase would be. The distinction matters: the rate question is answered by FBA and answered upstream; the quality question is answered at Pun1 and answered by redesign. Both campaigns are necessary, and they are complementary rather than competing.



Research trays of an engineered line — flux-balance analysis is ultimately answerable to this: carbon and nitrogen the plant actually commits to heat, computed for the fruit.

Flux-balance analysis earns its place a second time when we turn from peak heat to the *blend*. Recall from Chapter 5 that the five capsaicinoids are five outputs of one branched-chain fork, selected by chain length and saturation. Shifting the blend — more capsaicin, less dihydrocapsaicin — is not a matter of the final enzyme alone; it is a matter of which acyl-CoA species the fork makes available, which is set by the relative expression of KAS, the thioesterase FAT, and the acyl-CoA synthetases that prepare the tails. These are flux-control questions, and they are answerable only in a whole-pathway model, because changing the output of one branch pulls on the shared pools that feed every branch. To hit a target blend — the blends we design in the next chapter — we tune expression across several enzymes at once, and we use the flux model to predict the consequences before we commit them to a genome. A single-enzyme view cannot see those couplings. The economy-wide view can.

There is a principle underneath this chapter that we hold close, and it is the quiet ethic of the whole book stated in the language of systems biology. You do not improve a living thing by forcing one part of it and ignoring the rest. A plant is a balanced economy — carbon, nitrogen, energy, water, all allocated across thousands of competing demands — and to raise its heat responsibly is to understand that economy well enough to ask more of it *without disorder*: to identify where control actually lives, work at those points, and not flood a downstream enzyme that the upstream pathway has already outpaced. Flux-balance analysis is how we keep the faith that a designed plant should be a *well-ordered* plant, its heat lifted and its metabolism still coherent, still able to grow and set fruit and carry the fullest nourishment it can. The molecule was designed. The enzyme was redesigned and computationally confirmed stable. The pathway was balanced. Now we could stop asking what a plant *could* make and start building the ten that it does.



PART IV OF V

The Ladder

the blend matrix · ten cultivars · Ember to Phoenix · the full-spectrum capstone

Chapter 13 · Designing the Burn, Not the Number — *same heat, different fire.*

Everything in the first three parts converges on a single, counterintuitive capability, and it is the capability that separates our work from every heat-breeding program that came before: we can design two peppers with the *same Scoville rating* that feel completely different to eat. The SHU number is a scalar — one figure, a count of pungency. The *experience* of heat is not a scalar. It is a trajectory in time and a map in space: where the burn lands, how fast it arrives, how long it stays. And because, as Chapter 2 established, that trajectory is set by the *ratio* of the five capsaicinoids — each with its own TRPV1 affinity, its own onset, its own location — a designer who can tune the ratio can design the trajectory independently of the total. We can aim the fire, not just stoke it.

The design principle is simple to state and rests entirely on the TRPV1 docking of Chapter 2. **Capsaicin and nonivamide bind TRPV1 hardest** — −7.20 and −6.89 kcal/mol — so they strike fast and sharp and sinus-forward, seizing the receptor before the tissue's calcium brake can engage. **Dihydrocapsaicin and nordihydrocapsaicin bind weakest** — −6.39 and −6.16 — so they come on slowly and settle into a deep, radiating, long-dwelling body heat. Hold the total capsaicinoid load constant and slide the ratio from the weak binders toward the strong ones, and you slide the experience from a slow ominous build toward an instantaneous strike — same number of channels thrown open in the end, utterly different arc getting there. The blend *is* the character.

So we designed one for each rung of the ladder. Using the TRPV1 affinities as the governing data, we specified the capsaicinoid ratio of all ten flagship cultivars — not merely to hit a heat target, but to deliver a distinct burn *character* matched to each cultivar's identity. The result is a matrix in which the composition shifts smoothly, rung by rung, from body-heat to sinus-strike.

CULTIVAR	CAPSAICIN	DIHYDRO	NORDIHYDRO	HOMODIHYDRO	NONIVAMIDE	WEIGHTED TRPV1	ONSET
Ember	40.0%	30.0%	15.0%	10.0%	5.0%	−6.74	delayed 3–5 s
Cinder	43.9%	27.6%	13.6%	9.2%	5.8%	−6.78	delayed 3–5 s
Kindle	47.8%	25.1%	12.1%	8.4%	6.6%	−6.81	delayed 3–5 s
Flare	51.6%	22.7%	10.7%	7.7%	7.3%	−6.84	moderate 1–3 s
Blaze	55.6%	20.2%	9.2%	6.9%	8.1%	−6.88	moderate 1–3 s
Wildfire	59.4%	17.8%	7.8%	6.1%	8.9%	−6.92	moderate 1–3 s
Scorch	63.4%	15.3%	6.3%	5.3%	9.7%	−6.95	sub-second
Firestorm	67.2%	12.9%	4.9%	4.6%	10.4%	−6.99	sub-second
Solara	71.2%	10.4%	3.4%	3.8%	11.2%	−7.03	sub-second
Phoenix	75.0%	8.0%	2.0%	3.0%	12.0%	−7.06	sub-second

Read down any column and you can watch the design logic move. **Capsaicin climbs** from 40% to 75% — the fire sharpening and accelerating as you ascend. **Dihydrocapsaicin and nordihydrocapsaicin fall away** — from 30% and 15% down to 8% and 2% — the slow, deep, body-burning component steadily withdrawn, so that the lower rungs build ominously while the upper rungs simply detonate. The **weighted TRPV1 affinity** — the composition-averaged binding energy, the single number that best predicts onset — tightens monotonically from −6.74 to −7.06, a clean, designed gradient of intensity. The ladder is not ten random hot peppers. It is one continuous function, sampled at ten points.

Two more designed quantities complete the picture, and both climb the ladder in lockstep with the heat. The first is the sheer *load* — the total capsaicinoid the fruit must carry, in milligrams per gram of dry placental tissue — which rises more than threefold across the ten, because a hotter rung needs not only a louder blend but more of it. The second is the *duration* of the burn, scored on a ten-point scale, which falls as the onset sharpens: the slow lower rungs linger, the fast upper rungs detonate and clear. That inverse relationship is not incidental; it is the signature of the blend shift. The weak binders that make a burn slow also make it *long*, because they hold a low plateau as they come off the receptor gradually; strip them out in favor of capsaicin and the burn arrives faster and leaves faster, both at once.

CULTIVAR	CAPSAICINOID LOAD (MG/G)	ONSET SCORE	DURATION SCORE
Ember	223.5	6.0	5.0
Cinder	264.2	6.3	4.7
Kindle	311.4	6.6	4.4
Flare	364.5	6.9	4.1
Blaze	422.9	7.2	3.8
Wildfire	486.9	7.5	3.5
Scorch	562.7	7.9	3.1
Firestorm	643.9	8.2	2.8
Solara	735.7	8.5	2.5
Phoenix	858.7	8.8	2.2

Read the load column as a manufacturing spec and the duration column as a sensory one, and notice they are two views of a single designed object. Phoenix carries nearly four times the capsaicinoid of Ember — 858.7 against 223.5 milligrams per gram — and delivers it in a burn that peaks in under a second and is largely gone while Ember's is still building. Same molecule family, same receptor, opposite temporal signatures. Composed, not found.



A gloved hand-off in the greenhouse — each rung of the ladder is grown and blended to a specified capsaicinoid profile, not merely to a Scoville target.

Three design regimes fall naturally out of the matrix, and we named them as we built them. The **lower ladder** — Ember through Kindle — is the *slow burn*: high dihydrocapsaicin, delayed 3-to-5-second onset, a deep radiating chest heat that builds like a warning made good, lingering as a reminder. The **mid ladder** — Flare through Wildfire — is the *aggressive pivot*: capsaicin takes the majority, onset drops to 1-to-3 seconds, and the burn hits multiple locations at once, mouth and sinus and throat together, a sharp spike over a sustained tail. The **upper ladder** — Scorch through Phoenix — is the *sinus strike*: capsaicin-dominant, sub-second onset, no ramp and no warning, an instantaneous front-loaded peak that at these concentrations is simply overwhelming.

And across all three regimes, one component rises with a special purpose. **Nonivamide** climbs from 5% to 12% up the ladder — the one straight-chain, synthetic-accessible capsaicinoid, the anomaly from Chapter 2 that binds second-hardest at −6.89 despite having no kink and no branch. We use it as the **sharpener**: a clean, sharp, almost metallic edge laid over the top of the burn, rising exactly where the cultivars are named for weapons and protocols and the experience wants a synthetic hard edge to match. It is the clearest single example of the whole thesis of this chapter — a molecule chosen not for how much heat it adds but for the *shape* it gives the heat it adds. To achieve these ratios in a living plant is the flux-control work of Chapter 12 made concrete: tuning Pun1's acyl preference through the redesigns of Part III, adjusting the expression of KAS and the thioesterases to set which chain lengths the fork makes available, with nonivamide supplied by a dedicated route at the top of the range. Ten cultivars, ten designed trajectories, one continuous idea. The next chapter is the ladder itself — in Scoville, in IFI, rung by rung.

Chapter 14 · The Ten — *a ladder written in Scoville.*

The ten flagship cultivars are, first of all, a proof. Before the method could be trusted on a tomato or a strawberry or any of the sixty-four organisms of the roadmap, it had to be proven to its limit on the plant it understood best — *Capsicum chinense*, the superhot species, pushed deliberately up a ladder of increasing computational design scope until it reached a predicted peak no naturally bred pepper has approached. Each rung is a real cultivar design with a real SKU, grown to our designed specifications and catalogued. Together they are the clearest demonstration we can offer that designed heat is not a single trick but a *controllable variable* — tunable, repeatable, and extensible across a full order of magnitude of predicted Scoville.

Here is the ladder, with the two numbers that define each rung: the Scoville rating that measures how much capsaicinoid the fruit holds, and the **IFI-H** that measures how loudly it speaks to TRPV1. Alongside them, the count of how many of the eight flavor axes each cultivar expresses — because the ascent of this ladder is not only an ascent in heat but, as we will see, an ascent toward the full spectrum of flavor.

#	CULTIVAR	SKU	SHU	IFI-H	AXES
1	Ember	SS-0087	866,000	1,570	2 / 8
2	Cinder	SS-0088	1,050,000	2,290	3 / 8
3	Kindle	SS-0089	1,250,000	2,000	3 / 8
4	Flare	SS-0090	1,480,000	—	4 / 8
5	Blaze	SS-0091	1,720,000	2,630	4 / 8
6	Wildfire	SS-0092	1,960,000	2,930	5 / 8
7	Scorch	SS-0093	2,210,000	3,220	5 / 8
8	Firestorm	SS-0094	2,470,000	3,490	6 / 8
9	Solara	SS-0095	2,720,000	3,980	6 / 8
10	**Phoenix**	SS-0096	2,950,000	4,200	**8 / 8**

Read the Scoville column and you are reading a designed progression: from Ember's predicted 866,000 SHU — already a computational design target hotter than almost anything on a supermarket shelf, already past the Carolina Reaper's neighborhood — up through roughly even steps to Phoenix's predicted **2,950,000 SHU**, the peak that the double variant S39L+L345G of Chapter 9 made reachable and the coefficient of variation made repeatable. These SHU values are computational predictions from our flux and blend models. Ten rungs, climbing a full order of magnitude of designed heat, each held to a batch-to-batch coefficient of variation at or below 7%.



Firestorm (SS-0094) at full ripeness — rung eight of the ladder, 2.47 million Scoville, and the first to carry six of the eight axes.

But the column we find most telling is the last one, the axis count, because it reveals that this ladder was never only about heat. Ember expresses 2 of the 8 flavor axes; by Firestorm it is 6; at Phoenix it is all 8. The ascent in Scoville is shadowed by an ascent in *completeness* — each rung not merely hotter but more fully flavored, carrying more of the sweetness and umami and kokumi and the rest that make a thing not just intense but *good*. This is the deepest statement the ladder makes, and it is the one we care about most. A pepper engineered only for heat is a stunt — a number with nothing underneath it, a fire that scorches the palate and gives it nothing to love. The ladder refuses that. As the heat climbs, the flavor is built up beneath it, deliberately, axis by axis, so that even the fiercest rung is a *fuller* fruit and not merely a hotter one.

The SKU column is not decoration either; it is the thread that ties this book to the catalogs and to the platform. SS-0087 through SS-0096 are not names we invented for a story — they are the live identifiers of ten real cultivar designs in a catalogue of 19,500 product packages backed by approximately 225 distinct construct designs, each resolving to its own page, its own design record, its own batch history held to that ≤7% variation. When we say Firestorm is SS-0094, we mean there is a row in a database with that key and a coefficient of variation attached to it, and that the 2,470,000 Scoville in the table above is a computational prediction from our flux and blend models, not a marketing figure. The ten flagships are the public face of the superhot work, but their real function is to prove that the whole apparatus behind them — the structure-first design, the locked numbering, the quality system — holds all the way to the edge of what a pepper can physically carry. A ladder is only a proof if every rung bears weight.

A note on the IFI-H values, since one rung shows a dash. IFI-H is the receptor-loudness measure of Chapter 3, and across the ladder it climbs with the Scoville rating — from 1,570 to 4,200 — but not in perfect lockstep, because SHU counts molecules while IFI counts conviction, and the two diverge exactly where the *blend* shifts. Kindle reads a lower IFI-H than Cinder despite a higher Scoville, because its capsaicinoid mix (from the matrix of Chapter 13) is weighted toward the slower, deeper binders — more total pungency, a touch less receptor-loudness at the peak. That divergence is not noise. It is the whole thesis of the book made visible in two columns:

how much heat and *how loud the heat* are different questions, and only a designer who measures both can place a pepper precisely. Flare's full IFI-H profile is still in final validation as this edition goes to press, and we would rather print a dash than a number we have not stood behind — a small instance of the discipline the next chapter and the one after it are built on.

The names deserve a word, because they are not only theater. Ember, Cinder, Kindle, Flare — the ladder is deliberately named as an escalating series of threat levels, and the escalation is literal: each rung is computationally designed to be hotter than the last — a readable scale in predicted SHU, not merely a brand. There is a long culture of naming superhots for danger, and we wrote ourselves into it on purpose, because the chilihead community that will care most about this ladder reads heat as a kind of dare and we respect the dare. But under the theater the ladder is doing sober work: it samples the design space of *Capsicum chinense* at ten even intervals from 866,000 Scoville to the physical ceiling, and in doing so it maps the whole territory a pepper can occupy. Pick any heat a formulator could want inside that range and the ladder names the construct that delivers it, with what blend, at what variance. The names are for the community. The intervals are for the science.

The ten are the proof, and the proof holds. What remains, at the top of the ladder, is the one cultivar that climbed past heat entirely into the full spectrum — the flagship, the capstone, the pepper that registers on all eight axes at once. That is the next chapter, and it has earned its own.

Chapter 15 · Phoenix — *all eight, at once.*

At the top of the ladder stands a cultivar we named with a deliberate, clinical finality: **Phoenix**, SKU SS-0096, with a predicted heat of approximately 2.95 million SHU — and the only flagship that is not, in the end, a heat story at all. Every rung below it is an exercise in designing one axis to its limit. Phoenix is the exercise in designing *all of them* — the first cultivar we built to register on all eight axes of the Iyengar Flavor Index simultaneously, the full-spectrum pepper, the proof that the method which mastered heat was never confined to it.

Here is its complete profile — the signature scorecard of this book, the one eight-axis readout we print in full because it is the only cultivar for which all eight axes are fully computed and designed to target:

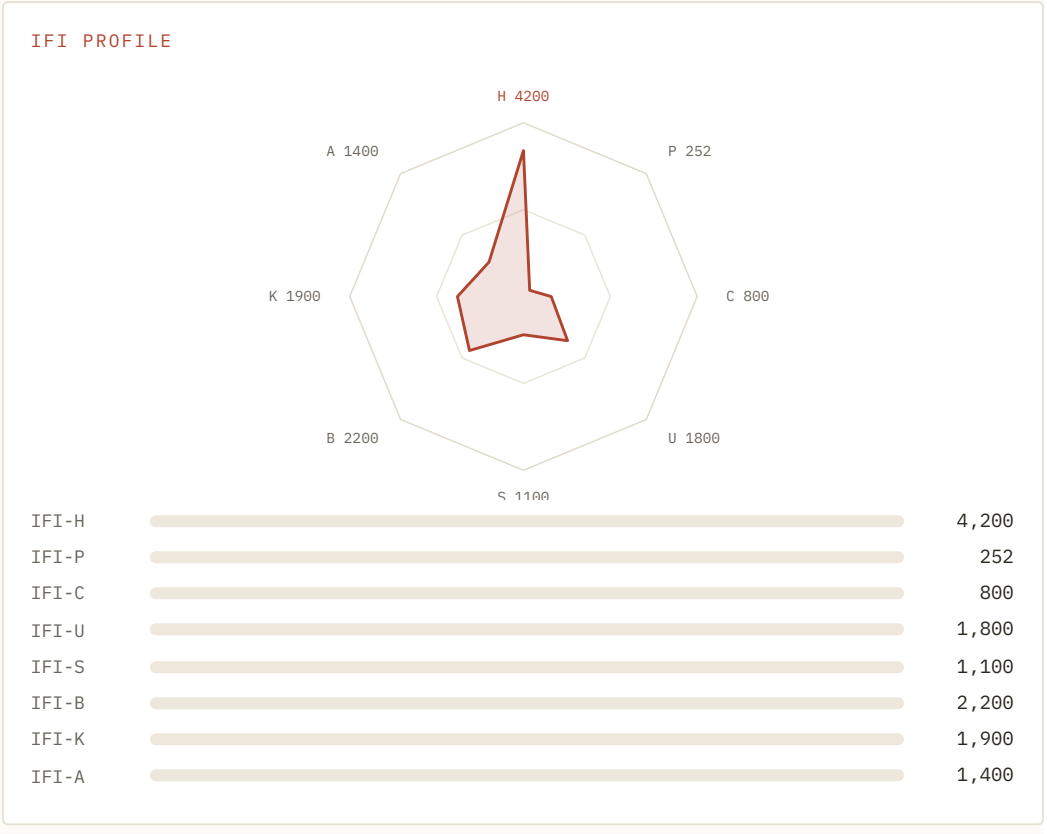

Read it as a chord rather than a note. The heat, **IFI-H 4,200**, is the loudest thing on the card by a wide margin — our model designs for approximately 2.95 million SHU, and nothing about the other seven axes softens that. But beneath the heat, a full architecture of flavor: **pungency** at 252, the sharp TRPA1 edge that rides alongside the TRPV1 burn (IFI-P scored against the AITC convention constant of 2.1 µM); a restrained **cool-burn** at 800, a thread of TRPM8 that keeps

the heat from being wholly one-directional; **umami** at 1,800 and **kokumi** at 1,900, the savory depth and the rich mouthfeel-continuity that give the fruit *body*, a roundness under the fire; **sweetness** at 1,100, enough to keep the whole thing from being merely an assault; **zero-bitter** at 2,200, the second-loudest axis, an active attenuation of the harshness that extreme capsaicinoid loads otherwise carry; and **astringency** at 1,400, the tannic structure that gives the finish its grip. Eight axes, eight receptors, one fruit — a thing that does not merely burn but *tastes*, all the way down.



Phoenix (SS-0096) in the controlled chamber — the first cultivar engineered to register on all eight axes at once, and the capstone of the flagship ladder.

Understand what building that chord required, because it is the synthesis of every method in this book. Heat alone was the work of Part III — the Pun1 redesign, the double variant S39L+L345G, the flux balance. But each additional axis is its *own* receptor, its *own* biosynthetic pathway, its *own* set of enzymes to be modeled and tuned, exactly as the heat axis was. The zero-bitter at 2,200 is an engineered attenuation of bitterness working through the TAS2R38 pathway. The umami at 1,800 is savory depth built through the glutamate route that speaks to mGluR1. The kokumi at 1,900 works through the calcium-sensing receptor, CaSR. To bring all eight up in a single genome — without the pathways colliding, without the shared carbon and nitrogen pools of Chapter 12 being overdrawn, without one axis's engineering disordering another's — is a combinatorial flux-balance problem of real depth, and it is the reason Phoenix sits alone at the top of the ladder. It is not simply the hottest thing we have made. It is the most *complete* — and completeness, across eight simultaneously-engineered pathways, is a far harder thing to hold than heat.

Take the loudest of the non-heat axes, the zero-bitter at 2,200, as a case in point, because it reveals how different "adding an axis" really is from "adding heat." Bitterness is sensed by the TAS2R family of receptors, and an extreme capsaicinoid load drags a great deal of intrinsic bitterness along with it — the harsh, acrid edge that makes most superhots punishing rather than pleasurable past a certain point. To score 2,200 on zero-bitter is therefore not to add a flavor but to *subtract* one: to engineer an active attenuation of the TAS2R38 pathway so that the bitterness a 2.95-million-Scoville fruit would otherwise carry is muted by a designed margin, letting the heat and the sweetness and the umami beneath it be tasted at all. The umami at 1,800 is the opposite operation — a *construction*,

building glutamate-pathway output to speak through the mGluR1 receptor and lay a savory floor under the fire. The kokumi at 1,900 works through yet a third receptor, the calcium-sensing CaSR, contributing not a taste of its own but *continuity* and mouthfeel, the quality that makes the whole chord feel round and sustained rather than thin and spiky. Three axes, three receptors, three entirely separate engineering campaigns — one subtractive, one constructive, one textural — running in the same genome without colliding. That is the real meaning of "full spectrum," and the real reason it sits alone at the top of the ladder.

The eight-axis delivery itself leans on the shared platform's finest instrument — a nanoparticle encapsulation that carries the eight compound classes into the same bite and releases them in a designed sequence, so that the chord arrives as a chord and not as eight separate notes fighting for the palate. That the sinus-strike blend of Chapter 13 — 75% capsaicin, the sharpest, fastest burn on the ladder — can carry, underneath it, a designed 1,100 of sweetness and 1,800 of umami without either being lost is itself a delivery achievement as much as a biosynthetic one.

The delivery earns its own sentence of respect, because it is where the molecular design meets the physical mouth. Eight compound classes do not naturally arrive together or in order — some are oil-soluble and some water-soluble, some volatile and some not, and left to themselves they would reach the palate as a scramble, the loud heat trampling the quiet sweetness before it could register. The platform's answer is encapsulation: carrying the eight into a single engineered particle that protects the delicate axes, then releases the compound classes in a designed sequence so the chord unfolds the way it was written rather than collapsing into its loudest note. This is why Phoenix earns a place on the composite full-spectrum axis, IFI-F, at all — not merely because its genome expresses eight pathways, but because its delivery lets a palate actually *receive* eight signals from one bite. A fruit that made eight flavors but delivered them as noise would score eight axes on paper and taste like one. Phoenix scores eight and tastes like eight, and the whole distance between those two outcomes is the delivery chemistry.

We will be honest about why we gave this cultivar the name we did, because the name is a promise and we keep our promises precise. Phoenix is designed for the far edge of what a human palate can receive — a pepper our model predicts at approximately 2.95 million SHU is not an everyday pleasure and we have never pretended otherwise. But the name points at something past the heat, too. This is the cultivar where the ladder stops climbing and the method turns a corner: from *how hot can one axis go* to *can all eight sing at once* — and the answer, printed in the scorecard above, is yes. Everything after this chapter is that corner, taken fully. Phoenix is the last and loudest word of the pepper. The next part is about every plant that is not a pepper — which is to say, about the whole point.



PART V OF V

Beyond the Pepper

method transfer · 42 organisms · the 64-concept roadmap · precision as reverence

Chapter 16 · The Pepper Was Never the Point — *the same hands, any plant.*

Here is the sentence this entire book has been walking toward, and we will state it as plainly as it deserves: **capsaicin is not a pepper problem.** The method we built to design heat in *Capsicum* is, in its bones, a method to design *any* flavor compound in *any* plant with a sequenced genome — and the reason is the architecture of the chemistry itself, which we laid out, enzyme by enzyme, in Part II precisely so that this chapter would land not as a leap of faith but as a conclusion you have already seen coming.

Recall the shape of capsaicin biosynthesis. The vanillyl head is built by the **phenylpropanoid highway** — PAL, C4H, 4CL, HCT, COMT — and that highway, as Chapter 4 insisted, runs in *every vascular plant on Earth*, because every plant uses it to make lignin and pigments and its own defenses. It is not a pepper feature. It is pre-installed in the tomato, the grape, coffee, lettuce, rice, the oak outside the window. Only the *last* few enzymes are pepper-specific: the four terminal catalysts that finish the job — **pAMT**, which aminates vanillin to vanillylamine; **BCAT** and **KAS**, which build and elongate the branched acyl tail; and **Pun1/CS**, the capsaicin synthase that performs the final condensation. Four enzymes. That is the entire difference, chemically, between a pepper and a plant that has never made heat. The hard, ancient, universal 90% is already everywhere. We supply the last 10%.

Which means the platform transfers. Of the seven computational modules that designed the pepper — structure prediction, the language-model mutational scan, the backbone redesign, docking, molecular dynamics, flux-balance analysis, and the delivery chemistry — **five transfer unchanged to any organism with a sequenced genome**, because AlphaFold does not care whether the protein came from a pepper, ESM-2 has read every species' sequences, and flux-balance analysis models any cell's economy. Only the two most organism-specific steps — the particular pathway tuning and the particular host transformation — are re-derived per species. The method is not a pepper method that we are hopefully extrapolating. It is a general method that we happened to prove on a pepper first because the pepper is where the chemistry was best understood.



Tomato, the first transfer: the same structure-first platform, pointed at a plant that already paves nine-tenths of the road.

The transfer is already underway. **Forty-two host organisms are in the platform's design library today** — not concepts, but species in which the platform has been run and constructs computationally built — spanning from the pepper itself through tomato, grape, peppermint, the mustards, stevia, garlic, soybean, tea, coffee, and on down a long list. The tomato was the first and most instructive: the same platform that climbed the heat ladder, pointed instead at the savory and sweet axes of a fruit that the industrial food chain had quietly stripped of flavor, rebuilding the umami and sweetness a tomato is supposed to carry through pathways the tomato already half-ran. Different axis, different receptor, *same method* — structure, scan, redesign, dock, simulate, balance, deliver.

Spell out which five of the seven modules make the jump unchanged, because the accounting is the argument. Structure prediction transfers — AlphaFold folds a tomato enzyme as readily as a pepper one. The language-model mutational scan transfers — ESM-2 has read the sequences of every sequenced species and scores a grape protein with the same fluency. The backbone redesign transfers — ProteinMPNN is indifferent to a protein's provenance. Molecular dynamics transfers — the physics of a warm protein is the physics of a warm protein. And flux-balance analysis transfers — any cell's metabolism is a stoichiometric network to be balanced. Those five are the general machinery, and they are the expensive, hard-won part. Only two steps are re-derived per host: the specific pathway tuning, because every plant's native metabolism offers a different starting flux, and the host transformation itself, because getting edited DNA into a cell and a whole plant back out remains an organism-by-organism craft. Five parts portable, two parts bespoke — that ratio is why forty-two organisms are already in the design library rather than four, and why the roadmap's sixty-four more read as a schedule and not a fantasy.



Eggplant — a Tier 1 Solanaceae host, close enough to the pepper that four or five transplanted genes carry the whole method across a family boundary.

And beyond the forty-two organisms already in the library lies the roadmap — **sixty-four concept organisms** across five tiers, graded by how many genes the transfer requires and how long it will take, each a deliberate step further from the pepper and deeper into the general case.

TIER	CHARACTER	GENES	HORIZON	REPRESENTATIVE CONCEPTS
T1	Solanaceae core	4–5	18–24 mo	Volcanic Tomato, Dark Matter Eggplant, Laboratory Tobacco
T2	Extended nightshades	5–7	24–36 mo	Seismic Potato, Celestial Ground Cherry, Acid Tomatillo, Floral Petunia
T3	Cross-family	7–10	36–48 mo	Sweet-Heat Strawberry, Contradiction Cucumber, Solar Melon, Forbidden Apple
T4	Frontier synthetic	10–12	48–72 mo	Deep-Earth Carrot, Nuclear Cabbage, Wildfire Broccoli, Storm Corn, Dynasty Rice

The gradient of the table is the gradient of difficulty, and it maps directly onto the chemistry. **Tier 1** is the pepper's own botanical family, the Solanaceae — tomato, eggplant, tobacco — where the host is so close to *Capsicum* that four or five transplanted genes suffice, and the horizon is under two years. **Tier 2** reaches to the extended nightshades, the potato and tomatillo and ground cherry, a little further out, five to seven genes. **Tier 3** is the first true leap across botanical families — a strawberry, a cucumber, a melon, a plant that shares the universal highway but little else — seven to ten genes and three to four years. And **Tier 4**, the frontier, is the full synthetic case: a carrot, a cabbage, a stalk of corn, a grain of rice engineered to carry designed flavor, ten to twelve genes, a half-decade or more of work. The tiers are not a wish list. They are a schedule, ordered by the exact genetic distance between a target

plant and the chemistry we already command — which is to say, ordered by how many of those last four enzymes the method has to carry across, and how foreign the soil it has to plant them in.

The pepper, in the end, was the teacher and not the lesson. It taught us to read an enzyme, to rank a mutation, to balance a pathway, to aim a burn — and every one of those skills is portable, because the chemistry of flavor is, underneath, the chemistry of a few universal pathways and a few terminal enzymes that we now know how to design. We began with heat because heat was legible. We continue, across sixty-four plants and eight axes, because legibility was always the point, and heat was only where we learned to read.

Chapter 17 · Precision as Reverence — *why the numbers are a form of love.*

We want to end not on the loudest number in this book but on the quietest, because it is the one that holds all the others up. Across every chapter we have shown a heat rising — 866,000 Scoville to 2,950,000, IFI-H 1,570 to 4,200, one axis to eight. But the figure we are proudest of is none of those. It is the **coefficient of variation falling from 42% to 7%** — the wandering of the wild fire brought to heel. We said at the start that for ten thousand years the craft of heat was a craft of patience defeated by variance: growers who could feel the fire exactly and never hold it still, a heat that arrived each season softened or scattered or gone. The whole structure-first method — the folding, the scanning, the double variant design, the dynamics, the flux balance — exists, in the final accounting, to turn a thing you *hope* for into a thing you can *specify*. Precision is the point. Everything else is what precision makes possible.

We hold precision because it is, to us, a form of respect — and we mean that in two directions at once. It is respect for the person who will eat what we design: a flavor that repeats is a promise kept, bite after bite, and a flavor that wanders by half is a small betrayal of the trust a cook places in an ingredient. And it is respect for the plant, because to cut a living thing's variance to 7% is to understand it that much more completely — to have modeled its chemistry closely enough that it no longer surprises you. A number, held honestly, is not the opposite of reverence for a living thing. It *is* reverence, in the only language precise enough to deserve the name.

There is a practical face to that principle, and it is the quiet reason the whole method matters beyond the spectacle of a three-million-Scoville pepper. A plant understood well enough to hold its heat to seven percent is a plant understood well enough to hold *anything* it makes to seven percent — its sugars, its acids, its vitamins, the density of what it offers the body that eats it. The variance we conquered on the heat axis is the same variance that, left uncontrolled, makes one tomato nourishing and the next one watery, one harvest rich and the next one hollow. Precision is not a trick we learned for capsaicin and will set aside for everything else; it is the whole disposition of the work, and heat was simply the axis loud enough to prove it on. When we speak, in the books that share this shelf, of growing food to carry the fullest nourishment it can, we mean this exact capability pointed at a gentler target — the same folding and scanning and balancing, aimed not at a louder fire but at a fuller fruit.



The founders in the greenhouse at the close of a day — the work is a family's, the reverence older than any instrument in the lab.

That commitment obliges us to be exact about the scope of what we build, and so we lock, here and across every book in this library, a single set of canonical figures and refuse to let them drift from one page to the next, because across a whole shelf of books drift is the one thing that reads as carelessness and carelessness is the opposite of what the method is for.

THE PLATFORM, IN LOCKED FIGURES	
Distinct construct designs	~225
Product SKUs	19,500
Live and purchasable today	1,000
Total catalogued SKUs	78,800
Validated host organisms	42
Sensory axes	8
Batch-to-batch coefficient of variation	≤ 7%
Patents filed	3
Peer-reviewed papers (Zenodo, CC BY 4.0)	21

Two of those rows carry a meaning past their magnitude. The **twenty-one peer-reviewed papers**, published openly under a Creative Commons license on Zenodo, are our statement that this method is not a black box or a trade secret dressed as magic — the structure prediction, the mutational scans, the docking and dynamics and flux models of this very book are in the open record, checkable, reproducible, offered to the field rather than hoarded from it. And the **three patents** mark the line between what we share freely — the science, the method, the understanding — and what we have built on top of it. We believe both at once: that the knowledge of how a plant makes a nerve lie belongs to everyone, and that the work of turning that knowledge into ten reliable cultivars is ours to have done.

That same discipline governs something as unglamorous as bookkeeping, and we mention it because in a library this size bookkeeping is a form of integrity. The figures in the table above are locked — one canonical set, pulled from a single source, printed identically in this book, in every other book on the shelf, and on every page of the website that quotes them. It would be easy, across sixty-odd volumes, to let the distinct construct count drift from two-hundred-and-twenty-five to three hundred, the SKU count from nineteen thousand five hundred to twenty thousand, each surface rounding a little differently until no two agreed. We refuse that, because a reader who catches two of our numbers disagreeing has every right to doubt the third — and the third might be the coefficient of variation this whole book was built to defend. Consistency in the small numbers is how we earn belief in the large ones. A method that holds a pepper's heat to seven percent but could not hold its own figures still would be refuting itself in its own footnotes.

These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. The SHU figures, IFI scores, and pathway outputs throughout this book are predictions from our structure-first computational platform — the same kind of honest design claim, made plainly and once.

A word, finally, on the through-line we promised in the opening essay we would let you feel rather than hear — and will now, just once, say plainly. Everything in this book has been about control: reading the enzyme, ranking the mutation, aiming the burn, holding the variance. But control was never the end. It is the means to a quieter end, which is that food should be a thing you give your body with intention — grown to carry the fullest flavor and the fullest nourishment it can hold, designed so that heat is a pleasure offered rather than an assault inflicted, built with the same care you would want in anything you take into yourself. We make no claim about what our peppers will do for your health; we are scientists, and that is a line we do not cross. We make only this claim, which is a claim about *intention*: that the precision we have spent seventeen chapters describing is, underneath the mathematics, an act of care — the care of a family that reveres what it grows, applied at the scale of the atom, so that the fire on your tongue is not an accident of a wandering gene but a thing deliberately, lovingly, exactly written. The pepper taught us to read. What we chose to write with that reading is the rest of this library — and the heat, as loud as it is, was only ever the first word.