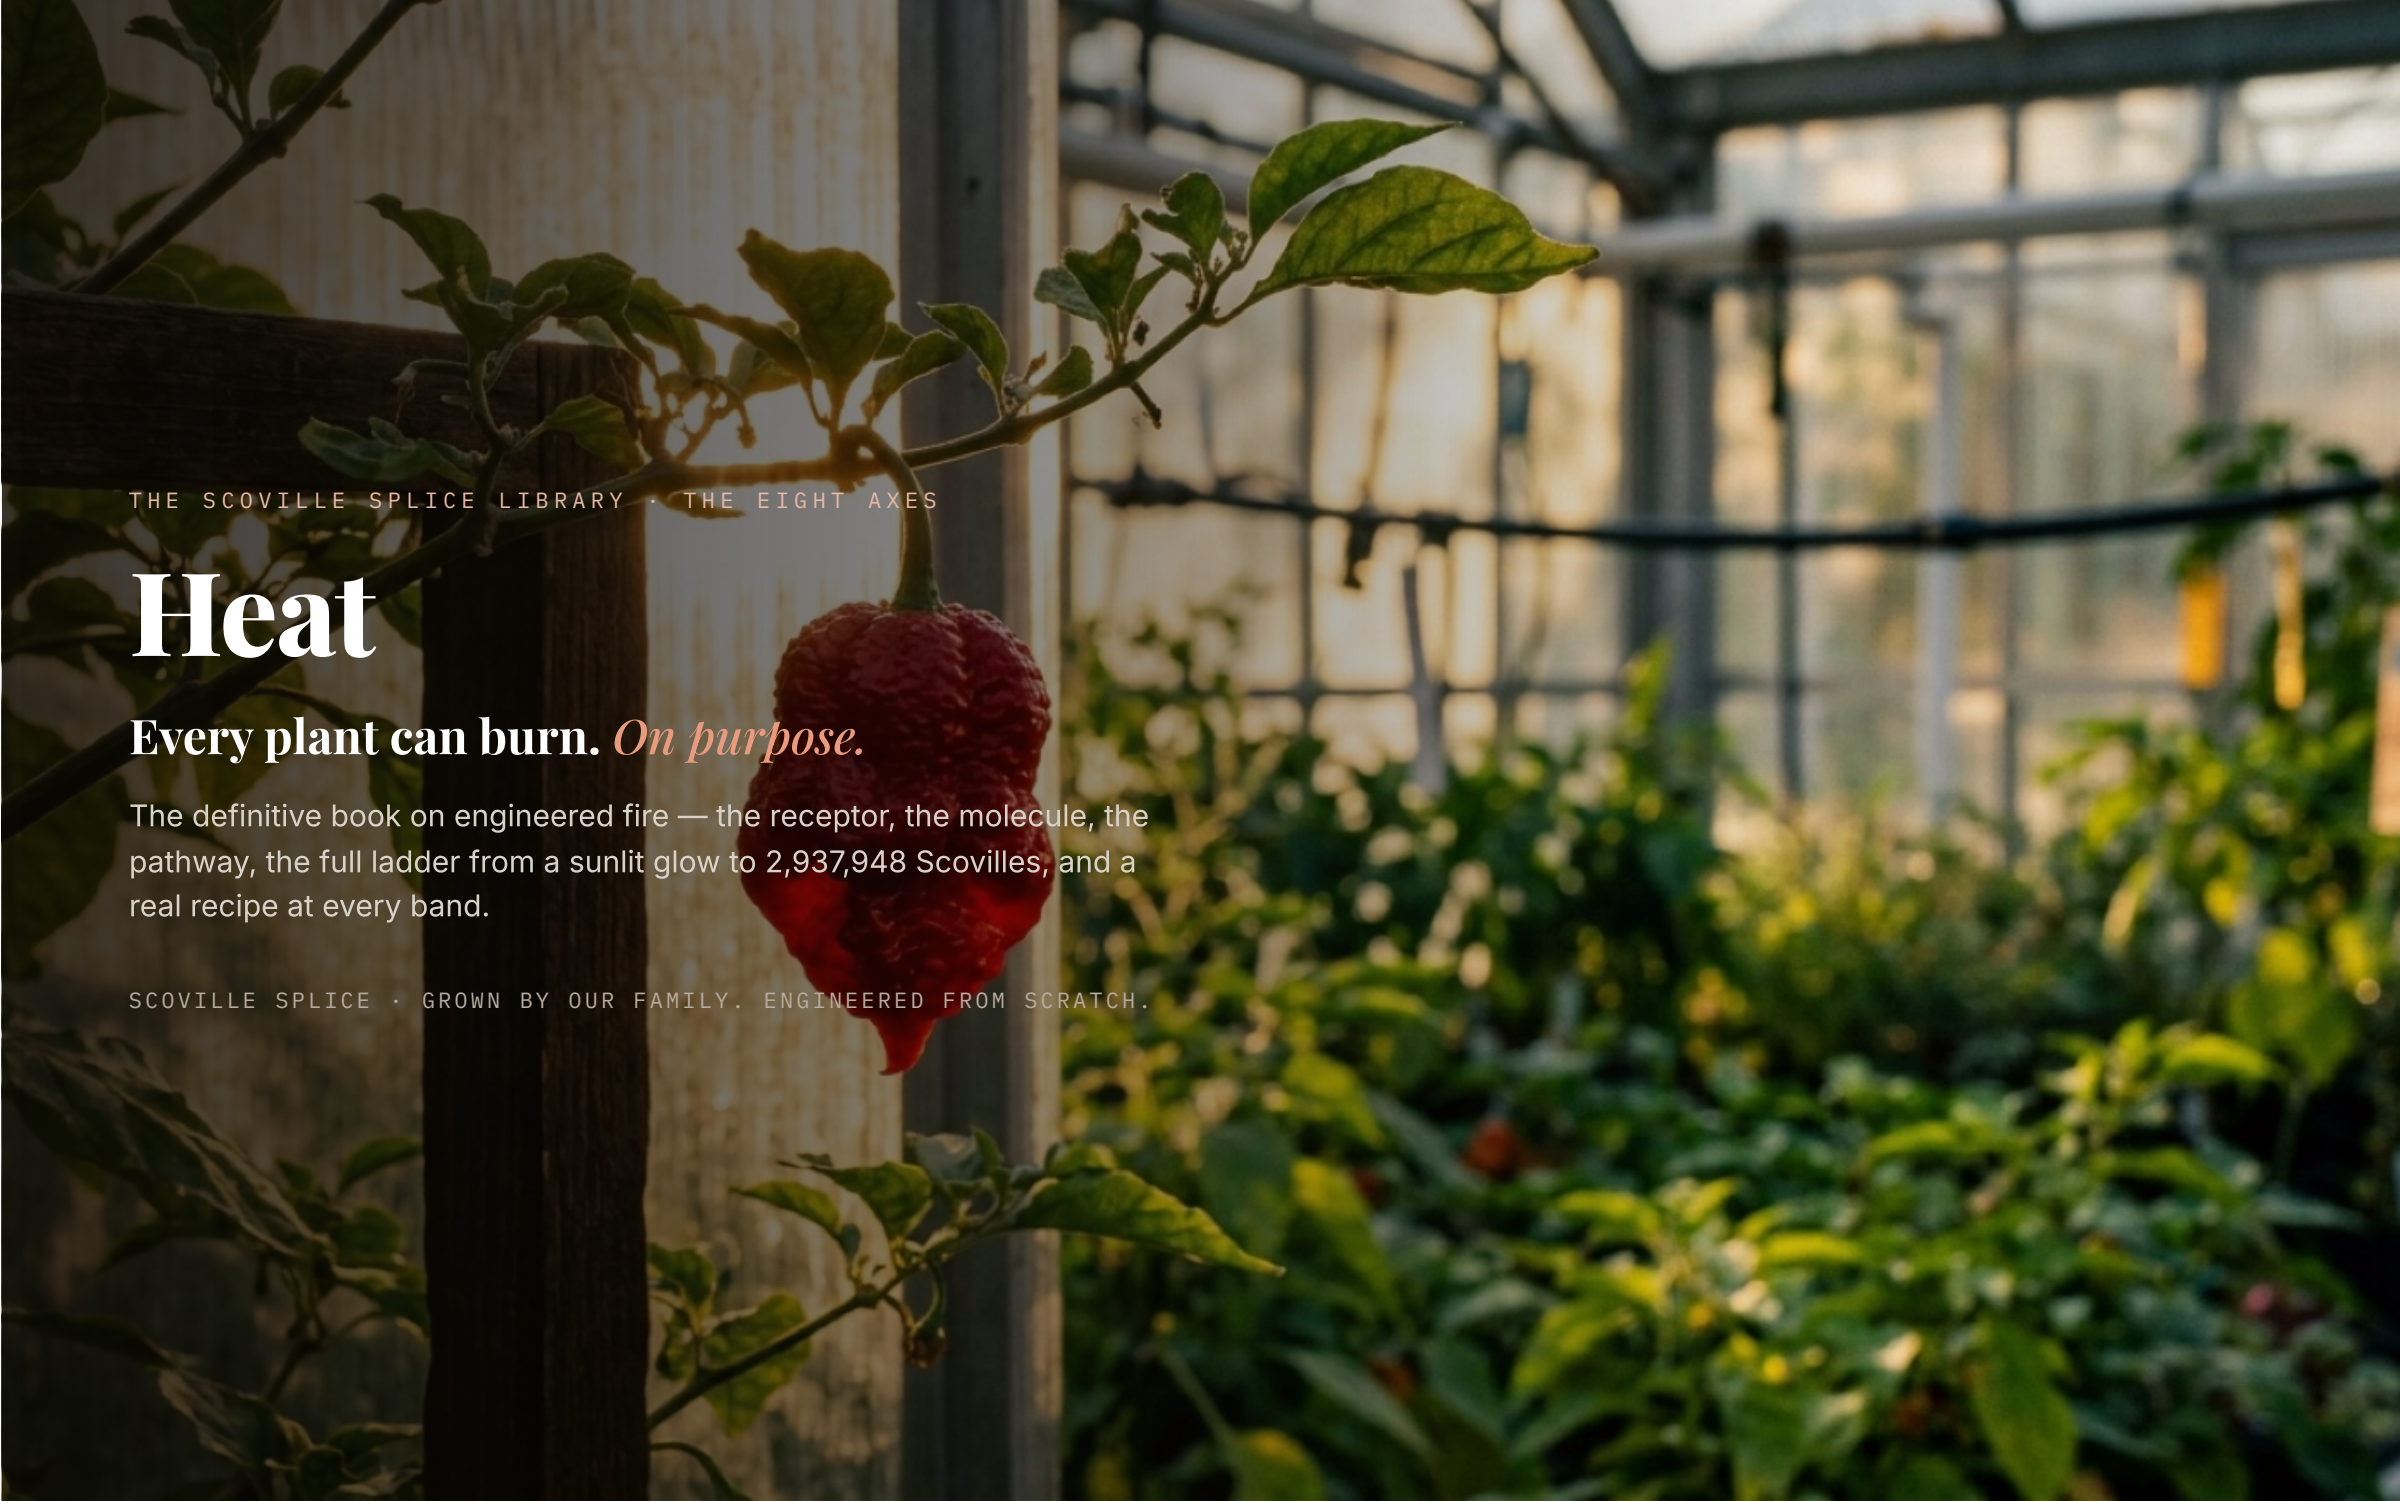

THE SCOVILLE SPLICE LIBRARY · THE EIGHT AXES

Heat

Every plant can burn. *On purpose.*

The definitive book on engineered fire — the receptor, the molecule, the pathway, the full ladder from a sunlit glow to 2,937,948 Scovilles, and a real recipe at every band.

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

Contents

01 Opening essay · The fire we chose

02 Chapter 1 · TRPV1 — the lock the fire fits.

03 Chapter 2 · The Capsaicinoids — a family, not a molecule.

04 Chapter 3 · The Pathway — from phenylalanine to fire.

05 Chapter 4 · The Mathematics of Burn — SHU is a dilution; IFI-H is a receptor.

06 Chapter 5 · The Heat Ladder — six bands from glow to edge.

07 Chapter 6 · Band 0 and 1 · The Warmth and the Glow — heat you can live inside.

08 Chapter 7 · Band 2 and 3 · The Burn and the Fire — where the meal wakes up.

09 Chapter 8 · Band 4 and 5 · The Inferno and the Edge — the ten we built to frighten you kindly.

10 Chapter 9 · Designing Phoenix — two mutations the breeders never found.

11 Chapter 10 · The Fire That Travels — heat in plants that never carried it.

12 Chapter 11 · The Volcanic Garden — tomato, eggplant, potato, learning to burn.

13 Chapter 12 · Cooking With a Number — how designed heat behaves in a pan.

14 Chapter 13 · Recipes · The Glow — heat as hospitality.

15 Chapter 14 · Recipes · The Burn — where the meal wakes up.

16 Chapter 15 · Recipes · The Fire — the edge of comfort.

17 Chapter 16 · Recipes · The Reckoning — cooking with the ten.

18 Chapter 17 · Pairing the Burn — heat is never eaten alone.

19 Chapter 18 · The Same Warmth, Every Time — CV ≤7%, and why it is a kindness.

Opening essay · The fire we chose

There is a moment, a half-second after the first real bite, when the mouth makes a decision it was never asked to make. The heat arrives — not a flavor, exactly, but a message — and some old animal part of you braces as though the kitchen has caught fire. Nothing has. No cell is harmed, no tissue burned, no temperature changed by a single degree. The tongue has simply been told a true thing in the most extravagant language the body owns. That half-second, repeated across ten thousand years and nearly every cuisine that ever loved a pepper, is the whole subject of this book.

Heat is the one axis everyone already knows is engineered. You can taste the intent in it. Where sweetness feels like a gift of the plant and umami feels like luck, heat feels like a dare someone set down on purpose — which is exactly why it is the right place to begin. Capsaicin finds TRPV1, the receptor that also answers to real flame and to the warmth of a held breath, and the body, briefly and gloriously confused, responds as if it had been scorched. It is a trick, but an honest one: the molecule really does fit the lock, really does throw the switch the heat of a candle would throw. The burn you feel is a message with no lie in it.



The founders in the research greenhouse at dawn — where the question that became this book was first asked: not *how hot*, but *how much, exactly, and can we keep it there.*

For most of human history the only way to turn that message up was to grow a hotter pepper and hope. Keep the seed of the one that moved you; wait a season; taste again. Ten thousand years of patience, and still no one could say *how much* — could not promise that this August's harvest would land where last August's did, could not place a warmth at the front of the mouth and keep it there. Heat was a feeling, and feelings do not hold still. A single field of one cultivar, grown from one seed lot, can vary in pungency by more than forty percent pod to pod. Every cook who has ever bitten into a jalapeño expecting mildness and found fire knows the problem in their own mouth.

We built the Iyengar Flavor Index so this one axis would finally hold still. On the Index, Heat is **IFI-H**, and it means one thing only: how strongly a food drives TRPV1, measured against a fixed reference and scaled so that a shishito, a habanero, and a chili paste can be read on the same page. This opening essay is the honest account of what that number is — the receptor underneath it, the molecule that trips it, the mathematics that turns a sensation into a score we can design toward, and the places where the number stops and the mouth, gloriously, takes over.

The receptor, exactly

Everything in this book rests on a single protein, so it is worth meeting it plainly. TRPV1 — the Transient Receptor Potential Vanilloid channel, type 1 — is a non-selective cation channel built from four identical subunits arranged around a central pore. It sits in the membrane of polymodal nociceptors: the free nerve endings of the trigeminal nerve that lace the lining of the mouth, the lips, the tongue. These are not taste cells. They are warning cells. A taste bud reports *what*; a nociceptor reports *how much, and should you be worried.*

TRPV1 is called polymodal because it opens to more than one kind of threat. It answers to heat above roughly **43 °C**, the temperature at which real warmth tips into real harm. It answers to acid — a flood of protons, a drop in local pH, the chemistry of inflamed or injured tissue. And it answers to a small family of plant molecules called vanilloids, of which capsaicin is the most famous. One channel, three doors, all of them opening onto the same sensation, because they all open onto the same nerve.

When the channel opens, calcium and sodium ions rush inward down their electrochemical gradient. The nerve ending depolarizes; if the depolarization clears threshold, it fires an action potential up the trigeminal line toward the thalamus and the somatosensory cortex, where it is read as *hot, burning, here.* The quality is indistinguishable from a thermal burn for the simplest possible reason: it is the same wire carrying the same signal. The brain has no separate channel for "chemical heat." It only has the one alarm, and capsaicin knows how to pull it.

Here is the honest sleight of hand at the center of the whole subject. Capsaicin does not raise the temperature of your tongue. What it does is bind a pocket on the inner face of the channel, between the third and fourth membrane-spanning helices, and stabilize the open state — which has the effect of *lowering the channel's thermal threshold.* A channel that would have needed 43 °C to open now opens at the temperature your mouth already is. The fire alarm has not detected a fire; its trigger has simply been loosened until body heat is enough to set it off. No tissue is scorched. The signal is genuine; the emergency is not.

Two facts about that pocket matter for everything that follows. First, the part of capsaicin the channel actually recognizes is its vanillyl head — the 4-hydroxy-3-methoxybenzyl group, the very same chemical motif that makes vanilla smell of vanilla. The receptor is

tuned to that small aromatic shape. Second, because the binding is reversible and calcium-dependent, repeated exposure quietly desensitizes the channel: the calcium that floods in during activation recruits pathways that make TRPV1 harder to open again. This is why a lifetime of chilies builds tolerance, and why capsaicin is used, at high dose, as a topical analgesic — the same molecule that shouts *fire* can, with repetition, teach a nerve to go quiet. We hold both of those truths at once: heat is a message, and the messenger can be turned down.

Where the molecule comes from

A pepper does not keep its heat in the flesh you think of as the pepper. It builds it in the placenta — the pale interior membrane that cradles the seeds — in specialized gland cells, and it builds it by joining two molecules that arrive from two entirely different directions in the plant's metabolism.

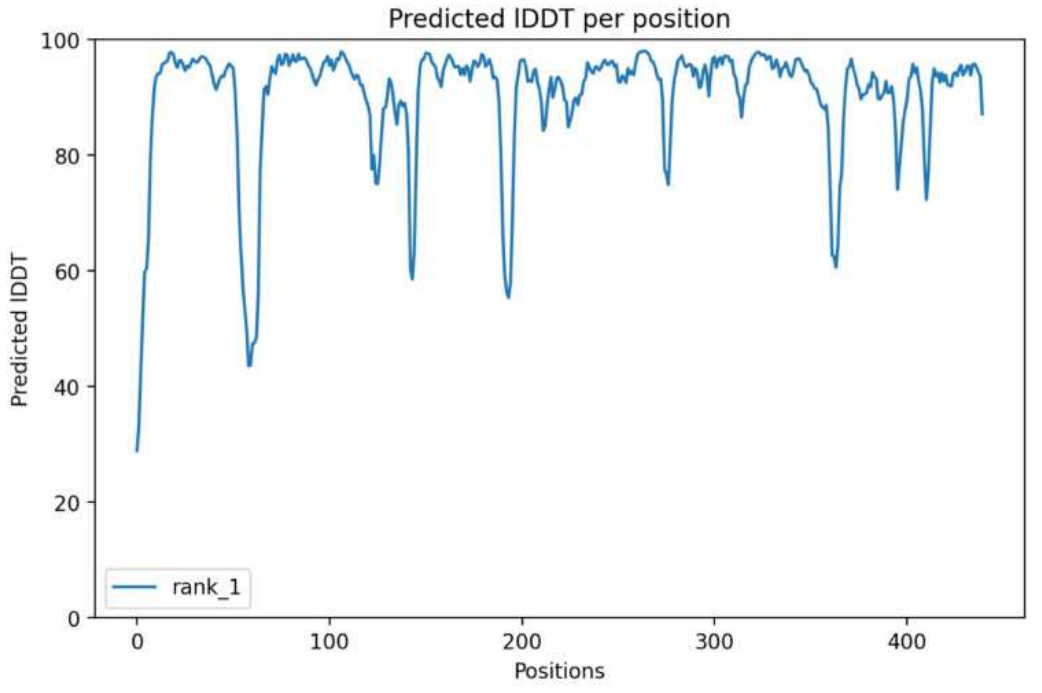

FIG. 1 — Pun1/AT3 (capsaicin synthase), the acyltransferase at the end of the pathway, coloured by model confidence (pLDDT). The active-site cleft, where vanillylamine meets the branched-chain acyl donor, is the high-confidence core — and the place we edited.

The first direction is the phenylpropanoid branch, the oldest and most universal chemistry in the plant kingdom — the same route that makes lignin, that stiffens wood, that colors flowers. It runs **PAL → C4H → 4CL → HCT → COMT**, a relay that begins with the amino acid phenylalanine and ends, after a half-dozen enzymatic hand-offs, at vanillin and then — through the aminotransferase pAMT — at **vanillylamine**. That is the vanilloid head, the part TRPV1 is built to read.

The second direction is fatty-acid chemistry. Starting from branched-chain amino acids like valine and leucine, the enzyme BCAT strips and rearranges them, and a ketoacyl synthase (KAS) elongates the result into a short, branched fatty acid — most often 8-methyl-6-nonenoic acid, activated as its coenzyme-A thioester. That is the tail.

The two meet at a single enzyme: **Pun1**, also called AT3 or capsaicin synthase, the acyltransferase pictured above. Pun1 catalyzes the condensation — it forms the amide bond that joins the vanillylamine head to the branched acyl tail — and the product is capsaicin. This last step is the gate. A pepper is pungent at all only because it carries a functional Pun1; the sweet bells and the mild paprikas are, at the genetic level, peppers whose Pun1 has been broken. Everything upstream can run perfectly, and without that final condensation there is no heat.

The shape of the tail is where the burn gets its character. If the acyl chain carries its double bond, Pun1 makes capsaicin; if the chain is fully saturated, it makes dihydrocapsaicin. These two account for roughly nine-tenths of the capsaicinoids in most peppers, and they do not feel the same. Capsaicin tends to read faster and sharper and farther forward in the mouth; dihydrocapsaicin arrives more slowly, sits deeper, and lingers longer in the throat. The ratio between them — set by how saturated that branched fatty acid is when it reaches Pun1 — is the difference between a heat that flares and clears and a heat that builds and stays. It is the first clue, and we will return to it, that *total capsaicinoid* and *felt loudness* are not the same quantity.

When our family set out to design heat, this is the map we were reading. We did not chase a hotter pepper by tasting and selecting and hoping; we asked the structure where the limit actually was. A protein language model scored every single-residue substitution in Pun1 for its predicted effect on catalytic efficiency and fold stability at once, while our genome-scale flux model showed that metabolic control in the capsaicin pathway sits upstream — at the phenylpropanoid enzymes 4CL and HCT/CCoAOMT — not at the terminal Pun1/AT3 step. Pun1 is the final capsaicin synthase; it seals the molecule into its finished form, but it does not govern the rate. Two substitutions came out of the structural search and into the shipped design: **S39L** (a stabilizing change near the N-terminus, $\Delta\Delta G \approx -2.1$ kcal/mol) and **L345G** (steric relief in the active-site cleft, $\Delta\Delta G \approx -1.8$). A third position, **C175S** — a structural annotation that forecloses a misfolding disulfide pathway ($\Delta\Delta G \approx -0.9$) — was identified during modelling but held out of the final construct. Written into the plant with CRISPR, the double variant S39L+L345G is designed to lift peak output to approximately **2.95 million Scoville heat units** — and, just as important, our model predicts it does so the same way every time.

From a dilution to a number

For a century, "how hot" had exactly one answer, and it was a confession of defeat dressed up as a measurement. In 1912 the pharmacist Wilbur Scoville proposed what is still the famous scale: take a known weight of dried pepper, extract it in alcohol, and dilute the extract in sweetened water until a panel of tasters can no longer

detect any heat. The dilution factor at the vanishing point *is* the Scoville heat unit. A pepper that had to be diluted fifty-thousand-fold before the burn disappeared scored 50,000 SHU.

It was ingenious, and it was a measurement of people as much as of peppers. The panel's thresholds drift with the time of day, with what they ate at lunch, with how many samples they have already tasted, with the slow desensitization of their own TRPV1 channels across a long afternoon. The modern laboratory version replaces the tongue with high-performance liquid chromatography: separate the individual capsaicinoids, weight each by its known pungency, sum them, and convert. The conversion is clean and worth keeping in your head — roughly **16 Scoville units per part-per-million** of capsaicin-equivalent, which means pure crystalline capsaicin sits near **16,000,000 SHU**. By that anchor, Phoenix's predicted ~2.95 million SHU is about **18%** of the heat of pure capsaicin — an extraordinary design target for a living fruit, and still a number about *chemistry*: how much capsaicinoid is present, not how loudly it speaks.

That distinction is the reason the Index exists. IFI-H does not ask how much capsaicinoid a food contains. It asks how strongly that food drives the receptor. The definition is one line:

> **IFI = 1,000 × EC_{50}(reference) ÷ EC_{50}(candidate)**

EC_{50} is the concentration at which a compound pushes its receptor to half of full response — the honest physical measure of how loudly a molecule speaks. A low EC_{50} means the receptor is half-open at a tiny concentration: a very loud molecule. A high EC_{50} means you need a great deal of it to get the same effect: a quiet one. The formula inverts that reading, so that louder becomes larger, and anchors it to a fixed wild-type reference scaled to **1,000**. When a candidate *is* the reference, the ratio is one and the score is exactly 1,000. Everything hotter climbs above it; everything milder falls below.

The arithmetic is worth doing by hand once, because it makes the number feel like the physical thing it is. Our model places Phoenix at **IFI-H 4,200**. Rearranging the definition:

$$\mathrm{EC}_{50}(\text{Phoenix}) = \frac{1000 \times \mathrm{EC}_{50}(\text{ref})}{4200} = \frac{\mathrm{EC}_{50}(\text{ref})}{4.2}$$

So Phoenix is designed to reach half-activation of TRPV1 at **one-fourth-point-two** of the concentration the reference needs — 4.2× more potent at the receptor. Run the same step for a shishito at IFI-H ≈ 40 and the arithmetic flips: its EC_{50} is 1,000 ÷ 40 = **25×** the reference, a molecule 25 times quieter. Put the two candidates against each other and the reference cancels entirely:

$$\frac{\text{potency}(\text{Phoenix})}{\text{potency}(\text{shishito})} = \frac{4200}{40} = 105$$

On the receptor, Phoenix is designed to speak **105 times louder** than a shishito. That single clean number — not a vibe, not a warning label, a ratio — is the whole promise of the Index. "Hotter" becomes

"105×," and 105× can be designed toward.

The felt heat does not climb gently with concentration, and the math says why. A cooperative channel like TRPV1 follows a dose–response of the Hill form,

$$\frac{R}{R_{\max}} = \frac{c^n}{\mathrm{EC}_{50}^n + c^n}$$

where c is concentration and n, the Hill coefficient, runs near 2 for capsaicin at TRPV1 — the signature of a channel whose subunits open together, not one at a time. Writing x for the ratio c/EC_{50} and taking $n = 2$, the curve is steep in exactly the place a cook cares about:

$X = C / EC_{50}$	TRPV1 RESPONSE	WHAT IT FEELS LIKE
0.10	~1.0%	nothing; a background warmth
0.25	~5.9%	the first suggestion of heat
0.50	20.0%	clearly there, pleasant
1.00	50.0%	the half-point; a committed burn
2.00	80.0%	loud, eyes-watering
4.00	~94.1%	near the ceiling
10.0	~99.0%	saturated; more adds little

Read the middle of that table and you can feel why heat behaves like a cliff rather than a ramp. To travel from one-fifth of the receptor's response to four-fifths of it — from *pleasant* to *eyes-watering* — you move only from half the EC_{50} to twice it: a mere fourfold change in concentration carries sixty points of perceived intensity. A pinch too much is not a little hotter. Near the half-point, a little more capsaicin is a lot more fire. This is the mathematical heart of why heat is the axis most able to ruin a dish, and why a number you can trust to a few percent is worth so much more here than anywhere else on the Index.

The ladder, by design

Our ten flagship cultivars were designed to span this scale in deliberate computational tiers, from a first real warmth to the edge we approached only once. Laid out together they form a design ladder — each cultivar representing a different scope of pathway activation and IFI-H target — not an ascending certified measurement. The SHU and IFI-H values in the table below are computational predictions from our model; they describe what each design is built to achieve.



Phoenix (SS-0096) in the controlled chamber — the first cultivar we engineered to register on all eight axes at once, and the single point where this ladder stops.

In the table below, *relative EC_{50}* is computed straight from the definition — 1,000 ÷ IFI-H, with the reference set to 1.000 — so a smaller figure means a molecule that reaches half-activation at a lower dose. The last column, Scoville units per IFI-H point, is the tell:

CULTIVAR	SKU	SHU (PREDICTED)	IFI-H	REL. EC_{50} (REF = 1.000)	SHU PER IFI-H POINT
Ember	SS-0087	866,000	1,570	0.637	552
Cinder	SS-0088	1,050,000	2,290	0.437	459
Kindle	SS-0089	1,250,000	2,000	0.500	625
Flare	SS-0090	1,480,000	—	—	–
Blaze	SS-0091	1,720,000	2,630	0.380	654
Wildfire	SS-0092	1,960,000	2,930	0.341	669
Scorch	SS-0093	2,210,000	3,220	0.311	686
Firestorm	SS-0094	2,470,000	3,490	0.287	708
Solara	SS-0095	2,720,000	3,980	0.251	683
Phoenix	**SS-0096**	**~2,950,000**	**4,200**	**0.238**	**700**

If SHU and IFI-H measured the same thing, that final column would be a constant. It is not. It runs from 459 to 708 — and the sharpest lesson is two rows that cross. Kindle carries **more** total capsaicinoid than Cinder (1.25M SHU against 1.05M) and yet our model places it **lower** on the receptor (IFI-H 2,000 against 2,290). More heat chemical, less heat loudness. The explanation is the acyl tail we met earlier: a blend weighted toward the slower, less receptor-potent capsaicinoids can out-mass a blend weighted toward capsaicin itself

and still speak more softly to TRPV1. SHU counts molecules. IFI-H measures what the channel does with them. When those two disagree, the Index sides with your tongue, because your tongue is downstream of the receptor, not of the HPLC trace.

At the top of the ladder, Phoenix is the one cultivar we tuned to register on every axis at once — the reason it anchors the full-spectrum reference and not only the heat scale:

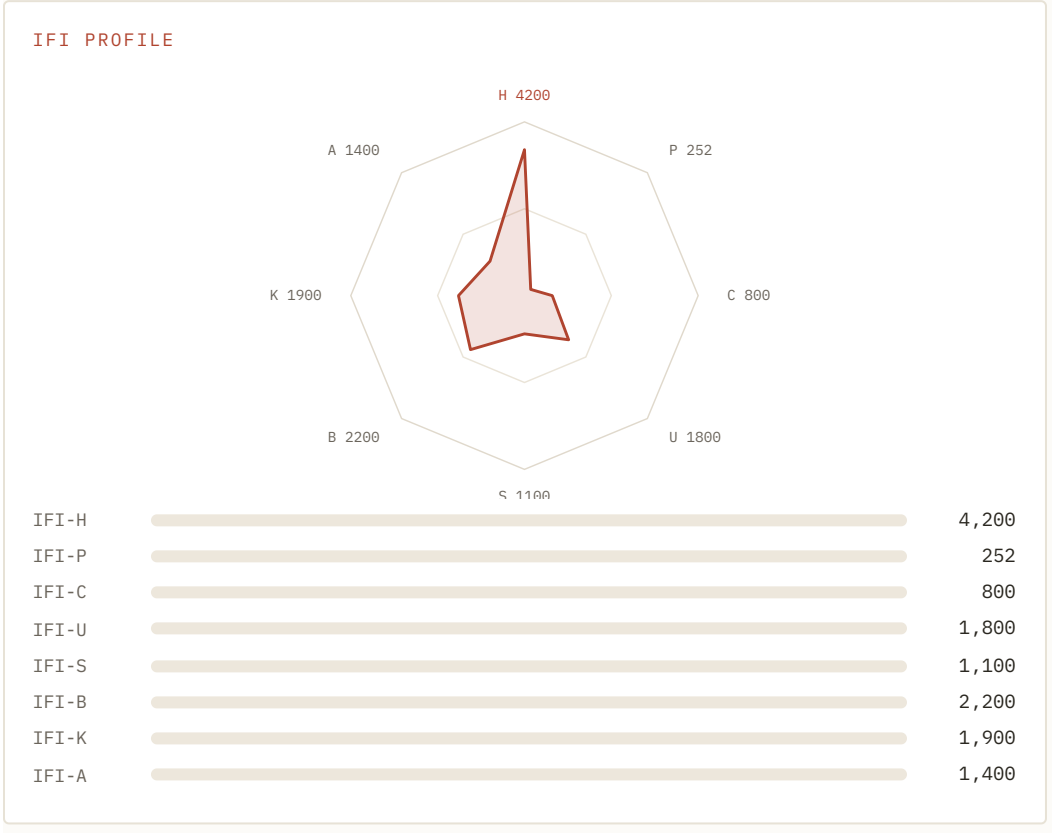

Read that profile and you can see why we named it the way we did and then stopped. Heat at 4,200 towers over everything, but it does not stand alone: there is real pungency underneath it, a cool edge, a savory weight, a bitter spine, the long kokumi continuity that keeps the whole thing from reading as a single screaming note. It is not the hottest possible object we could imagine. It is the hottest object we were willing to make *complete* — a fire with architecture, which is a very different ambition from a fire without a ceiling.

What varies, and what we hold still

A number is only as good as its reproducibility, and this is where the engineering earns its keep. The honest baseline for wild-type pepper is a pod-to-pod coefficient of variation of about **42%**. Coefficient of variation is just the standard deviation expressed as a fraction of the mean — CV = SD ÷ mean — and 42% is enormous. Put a target on it and watch what happens.

Suppose a recipe calls for IFI-H **900** — a committed, sauce-lifting warmth. At 42% CV, the standard deviation around that target is 0.42 × 900 ≈ **378**. One standard deviation already spans 522 to 1,278. Two standard deviations — the range into which about 95% of pods fall — reaches from **144 to 1,656**, better than an elevenfold spread between the mild tail and the fierce one. That is the jalapeño problem written as statistics: ask for 900, and the plant hands you anything from a whisper to a slap. No cook can build intention on that.

Our engineered lines are designed to land within a batch-to-batch CV of **≤7%**. Run the identical arithmetic: SD = 0.07 × 900 = **63**; one standard deviation is 837 to 963; two standard deviations is **774 to 1,026**. The whole 95% window now fits inside a quarter of what a single wild-type standard deviation covered. In Scoville terms, Phoenix's predicted ~2.95 million SHU at 7% CV carries a standard deviation near 205,700 units — a two-sigma window of roughly 2.53M to 3.35M. At the wild-type 42%, the standard deviation alone would be over 1.23 million SHU, and the cultivar would be, for any precise purpose, unusable.

Laid side by side for a single target, the two regimes are barely the same measurement:

FOR A TARGET OF IFI-H 900	WILD-TYPE · 42% CV	ENGINEERED · ≤7% CV
Standard deviation (IFI-H)	378	63
±1σ window	522 – 1,278	837 – 963
±2σ window (~95% of pods)	144 – 1,656	774 – 1,026
Low tail to high tail	~11.5×	~1.3×

It is worth being exact about what that 7% is and is not, because the difference is the whole reason breeding could never get here. You can shrink the *average* of a harvest by pooling: the standard error of a mean falls as the square root of the count, so averaging more fruit tightens the number. To pull a 42% intrinsic spread down to a 7% mean by pooling alone would take

$$N = \left(\frac{42}{7}\right)^2 = 36$$

thirty-six pods blended into every batch. But that only steadies the *average* of thirty-six fruits; it does nothing for the single pod that lands in a single dish, which still swings across the full 42%. Selection breeding can calm a field. It cannot calm a bite. Our 7% is a design specification for reproducibility of the plant itself, batch to batch and pod to pod — the same warmth in the fruit you actually eat, in August and in February, in our kitchen and in yours. For the axis most able to ruin a meal, that is not a small thing. It is a kindness.

The honest limits

A good number tells you precisely what it measures and stays quiet about the rest, so let us be clear about where IFI-H stops.

It is a measure of receptor loudness, and only that. It reports how strongly a food drives TRPV1 at half-maximal activation; it does not, by itself, describe the *time-course* of the burn — the onset, the peak, the long decay — which lives in the capsaicin-to-dihydrocapsaicin ratio and in how the capsaicinoids partition into fat and saliva once they are in your mouth. Two foods can share an IFI-H and feel different in their arc, one flaring at the front of the tongue and clearing, the other blooming slowly at the back of the throat and staying. The Index pins the height of the burn, honestly and repeatably. It does not pin its shape, and we would mislead you to claim it did.

It also meets a body, and bodies differ. The density of TRPV1 on your nociceptors, your own history of exposure and the desensitization it has built, even the temperature of the dish, all shift where a given IFI-H lands for you personally. This is not a flaw in the number; it is the boundary of what any single number can own. The score fixes the stimulus. The person is the person. And the folk wisdom about reaching for milk rather than water is, pleasingly, pure mechanism: capsaicin is fat-loving and barely soluble in water, so water only moves the fire around, while the casein proteins and fat in dairy lift the molecule off the receptor and carry it away. The chemistry that makes heat precise is the same chemistry that tells you how to quiet it.

And the Index is anchored to a receptor read, which is cleaner than a 1912 taste panel and still a model of a far messier event. We trust it because it is reproducible to a few percent, because it agrees with the mouth where the mouth can be trusted, and because — unlike a Scoville number — it sides with perceived loudness over raw chemistry when the two diverge. But we hold it the way we hold every instrument: as a faithful account of one thing, not a verdict on the whole experience. The number buys intention. The meal is still the meal.

The fire we chose

What the Index finally buys, under all the arithmetic, is intention. "Medium heat" is a wish. IFI-H 900 is an instruction — one that tastes the same across a seven-percent window, anchored to the receptor and not to a memory. Our platform holds roughly **nineteen thousand five hundred** product packages — backed by approximately 225 distinct construct designs — with a thousand of them available and purchasable today, across **forty-two** host organisms and **eight** sensory axes, and every batch is designed to land inside that same tolerance. The same warmth, every time. We did the computing and the design in-house, pod by pod and edit by edit, but the table this is all in service of — the reverence for the living plant, the vegetarian-first cooking, the clean habit of feeding the body with care — belongs to our family and always has.

That is why we draw the line we draw. Heat, to us, is a pleasure and never an assault. We did not spend years learning to design fire so that anyone could be dared across a table into pain; we built it so that a cook could ask for *exactly this much warmth* and receive exactly that. Between a shishito's sunlit 40 and Phoenix's 4,200 lives everything a kitchen could want — the slow glow that makes you reach for the next bite, the bright clean sting that lifts a sauce off the plate, the deep architectural burn a chili paste is built around — and now each of those is a number you can name, grow on purpose, and keep. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. A message with no lie in it, given with intention, and meant to nourish the person who sits down to eat.



PART I OF V

The Receptor, the Molecule, the Pathway

TRPV1 · the capsaicinoid family · phenylalanine to fire · the mathematics of burn

Chapter 1 · TRPV1 — *the lock the fire fits.*

Begin where the heat actually happens, which is not in the pepper and not in the pan but in a single protein folded into the membrane of a nerve ending on your tongue. Its name is TRPV1 — transient receptor potential, vanilloid subfamily, member one — and understanding it is the whole foundation of designing heat, because every number in this book is ultimately a statement about how hard a food pushes on this one molecule. Learn the protein and the rest of the book becomes arithmetic: a ladder of foods arranged by how loudly each of them speaks to one lock.

We did not invent the fire — it is ten million years older than any kitchen. What our family learned to do was read the lock the fire fits, and then, residue by residue and batch by careful batch, to cut keys for it on purpose. That is the quiet claim under everything here. Heat is not a flavor we add. It is a conversation between a plant and a protein, and we have learned to speak both sides of it.

The architecture of the channel

TRPV1 is an ion channel: a gate in the cell membrane that, when it opens, lets charged atoms flood through. It assembles as a homotetramer, four identical subunits arranged around a central pore like four hands cupping a well. Each subunit contributes six membrane-spanning helices, S1 through S6, and the architecture is the same domain-swapped plan the voltage-gated channels use — an S1–S4 bundle on the outside acting as a sensor module, an S5–S6 pore domain lining the central conduit, and a short, critical S4–S5 linker coupling the two so that a change sensed at the periphery is transmitted inward to the gate. Below the membrane each subunit hangs a large cytoplasmic skirt: an N-terminal ankyrin-repeat domain that fields the cell's own regulators — ATP, calmodulin, the state of the membrane's lipids — and a C-terminal strand that returns as a short amphipathic **TRP helix**, lying just under the inner leaflet where it can feel both the S4–S5 linker and the lipids packed against it. Four of these skirts braid together into a hanging gondola beneath the pore, and the whole assembly behaves as one allosteric machine: push anywhere and the gate, eighty-odd ångströms away, answers.

When the channel opens, sodium and especially calcium rush in; the local membrane depolarizes; the neuron fires; and a signal races up the trigeminal nerve to a brain that has exactly one interpretation for that pattern of firing from that population of fibers. *Burning.* No tissue is harmed, no temperature has changed — and still the report arrives, urgent and total. The entire sensation of heat is the brain's faithful reading of a gate that a plant has learned to pry.

It helps to hold the landmarks in one place, because the same handful of residues govern where a molecule binds, how the pore selects what passes, and where the gate finally gives way:

LANDMARK	LOCATION	WHAT IT DOES
S1–S4 bundle	outer sensor module	reads heat, voltage, membrane state
S4–S5 linker	sensor-to-pore coupler	the lever that transmits binding to the gate
Tyr511 (Y511)	floor of the vanilloid pocket	cradles the capsaicin ring from below
Thr550 · Glu570	vanilloid pocket wall	hydrogen-bond the vanillyl head
Outer pore turret	extracellular S5–P loop	proton and heat sensing
Gly643 (G643)	selectivity filter	upper gate; favors calcium
Ile679 (I679)	S6 bundle crossing	lower gate; opens to pass ions
TRP helix + PIP_2 site	inner-leaflet shelf	sets the channel's resting tension



A pepper in cross-section: the pale placental tissue arcing along the inner wall is where the plant makes its capsaicin — vesicles of the molecule secreted onto the surface of cells the seeds never even touch. The fire is manufactured in one organ and read in another, a continent away.

Three fingers on one switch

What makes TRPV1 extraordinary — and what capsaicin exploits — is that it is a *polymodal* sensor. It is, first and most ancestrally, a heat detector. Raise the temperature of the tissue past roughly 43 °C and the channel opens on its own; this is the receptor that tells you the coffee is too hot to drink and the pan handle is not to be grabbed. It also opens in response to acid — a drop in extracellular pH, protons binding residues on the outer turret of the pore — which is part of why a squeeze of lime can sharpen a chili's bite and why inflamed,

acidified tissue aches. Heat and acid are the receptor's native languages. Capsaicin is a loanword it turns out to understand perfectly.

The molecule does not heat the tissue. This is the single most important fact in the book, and it is worth saying slowly. Capsaicin binds inside a pocket formed by the S3 and S4 helices and the S4–S5 linker — the *vanilloid pocket* — and from there it does something subtle and devastating: it lowers the channel's thermal threshold. A receptor that would have waited for 43 °C now opens at body temperature, 37 °C, which the tissue is already at. The threshold has fallen by six degrees — a small number with an enormous consequence, because 37 °C is not a stimulus you can escape by setting down the cup. The gate that should be a smoke alarm is now triggered by the warmth you carry everywhere. The fire you feel is not a fire that exists; it is a fire the receptor has been convinced to report.

Heat, acid, capsaicin: three unrelated stimuli, one sensation. Picture three fingers resting on a single switch. Temperature destabilizes the closed state through the pore domain's outer turret; protons act at the filter; capsaicin works from inside the sensor module. Each takes a different path, and all three converge on the same two gates. Capsaicin is simply the finger a plant learned to press — and because it presses that one defined site on that one defined protein, its effect is a clean, measurable thing, which is the whole reason heat can be given a number at all.

The vanilloid pocket, residue by residue

From cryo-electron-microscopy structures we know the binding in intimate detail. The vanillyl head of the molecule — the aromatic ring with its hydroxyl and methoxy groups — reaches deep into the pocket and hydrogen-bonds to a threonine at position 550 and a glutamate at 570, while a tyrosine at 511 cradles it from below; the long fatty tail threads back out toward the lipid-facing surface, lying against the membrane like a mooring line. Mutate tyrosine 511 to alanine and the pocket goes deaf to capsaicin while still answering perfectly well to heat. The thermal sense survives; the chemical sense is gone. The lock and the key are that specific — a single aromatic ring, removed, and the molecule knocks on a door that no longer opens.

Knowing the lock, we can speak precisely about the key, because the plant builds it with an enzyme as particular as the pocket it fits. Capsaicin is assembled when **vanillylamine** condenses with an 8-methyl branched-chain fatty acid at **Pun1/AT3**, the acyltransferase the pepper evolved for exactly this and nothing else. The length and saturation of that acyl chain set which capsaicinoid results: an unsaturated chain yields capsaicin itself; a saturated one yields dihydrocapsaicin, its quieter, slower sibling. The ratio between them — decided one bond at a time, at one enzyme — sets the *shape* of the burn, its onset and its tail, as surely as the vanillyl head sets whether there is a burn at all. A plant engineering its own heat and a protein reading that heat are, in the end, the same chemistry approached from two directions. We work both.

From binding to the open gate

Knowing the binding site is one thing; knowing how binding *opens the gate* is another, and TRPV1's gating is a small marvel of mechanical logic. The channel holds two gates in series along its central pore — an upper constriction at the selectivity filter near a glycine at 643, and a lower gate where the four S6 helices cross, around isoleucine 679. Capsaicin, seated against the S4–S5 linker, acts like a wedge: its binding tugs that linker, which levers the lower S6 gate open, widening the pore just enough for hydrated calcium and sodium to stream through. The motion is not metaphorical — structures of the closed and open states show the S6 bundle rotating and splaying, the Ile679 constriction dilating from a ring too tight to pass a hydrated ion to one that passes them freely. One molecule in one pocket, one lever pulled, one gate thrown eighty ångströms away.

A lipid sets the trigger weight. **PIP_2**, a phosphoinositide tucked against the TRP helix in a neighboring pocket, holds the switch's resting tension, and the state of that lipid tunes how little capsaicin it takes to tip the channel over. This is why an identical dose can bite harder on an empty, lipid-depleted membrane than on a fed one — a molecular footnote to the old kitchen wisdom that heat lands harder on an empty stomach, and a reminder that we eat with our whole chemistry, not just our tongues. It is also, quietly, a design handle: the same burn delivered against a bite of avocado or a spoon of yogurt is a different burn, because the fat and the membrane and the receptor are all one continuous system. A dish that respects the body builds that in.

The number under the burn

Because capsaicin acts at one defined site on one defined protein, its potency is a clean, measurable number. We quantify it as an **EC_{50}** — the concentration at which the channel reaches half of its maximal open probability — and the smaller that concentration, the louder the molecule. For pure capsaicin at the cloned receptor, the literature places that half-point in the low sub-micromolar range, on the order of 0.7 µM; for a molecule of mass 305 g/mol that is about 0.21 mg in a litre of water, a pinch of crystal in a bathtub, enough to light every TRPV1 in reach to half its roar.

The relationship between dose and response is not a straight line but a sigmoid, captured by the Hill equation:

```
response = [C]ⁿ / (EC₅₀ⁿ + [C]ⁿ)
```

where [C] is capsaicinoid concentration and *n*, the Hill coefficient, describes how steeply the channel flips from silent to roaring. For TRPV1 and capsaicin, *n* sits near 2. That single exponent is the reason heat *blooms*. Write the response in units of EC_{50} — let $r = [C]/EC_{50}$ — and the curve becomes response = $r^2/(1 + r^2)$:

R = [C]/EC_{50}	RESPONSE, *N* = 2	RESPONSE, *N* = 1
0.25	5.9%	20.0%
0.50	20.0%	33.3%
1.00	50.0%	50.0%
2.00	80.0%	66.7%
4.00	94.1%	80.0%
10.0	99.0%	90.9%

Read the *n* = 2 column from 20% up to 80% — the whole perceptual rush, the stretch where heat goes from a hint to a shout — and the concentration only quadruples, from half the EC_{50} to twice it. Span the full audible range, 10% to 90%, and it takes just a ninefold change in dose; for a simple *n* = 1 receptor the same span would demand eighty-one-fold. The steepness compresses the climb. It is exactly the subjective experience of heat: the half-second of nothing, and then the bloom — the molecule gathering itself below the threshold of notice, then arriving all at once. The mathematics and the sensation are the same event, read off two different instruments.



Phoenix on the microscope stage, where loudness stops being a word and becomes a predicted half-maximal concentration — one extract, one receptor, one number.

From the receptor to the Index

Every IFI-H score in this book is, underneath, a statement about the EC_{50} of a food's capsaicinoid mixture at this receptor. The Index inverts potency into loudness with one line:

> **IFI-H(food) = 1,000 × EC_{50}(reference) ⁄ EC_{50}(food)**

Set the native wild-type pepper as the reference at the center of the scale. Its extract half-activates the channel at some concentration EC_{50}(ref), and by definition IFI-H(ref) = 1,000 × EC_{50}(ref)/EC_{50}(ref) = **1,000**. Everything else is read against that anchor. A shishito is gentle: you must flood the receptor with roughly twenty-five times more of its extract to reach the same half-maximal response, so

EC_{50}(shishito) ≈ 25 × EC_{50}(ref), and IFI-H = 1,000 ⁄ 25 = **40** — the sunlit lower rung of the whole ladder. Now invert the shishito entirely. **Phoenix**, the cultivar our family designed to activate across all eight axes at once, is computed to half-activate the channel at about 0.238 × EC_{50}(ref) — a quarter of the reference's concentration, roughly one hundred-and-fifth of the shishito's — so IFI-H = 1,000 ⁄ 0.238 = **4,200**. And the two extremes close cleanly on each other: 4,200 ⁄ 40 = 105, so Phoenix is, to the receptor, one hundred and five times louder than a shishito. Not hotter by opinion. Louder by our model's prediction, at the one lock that decides.

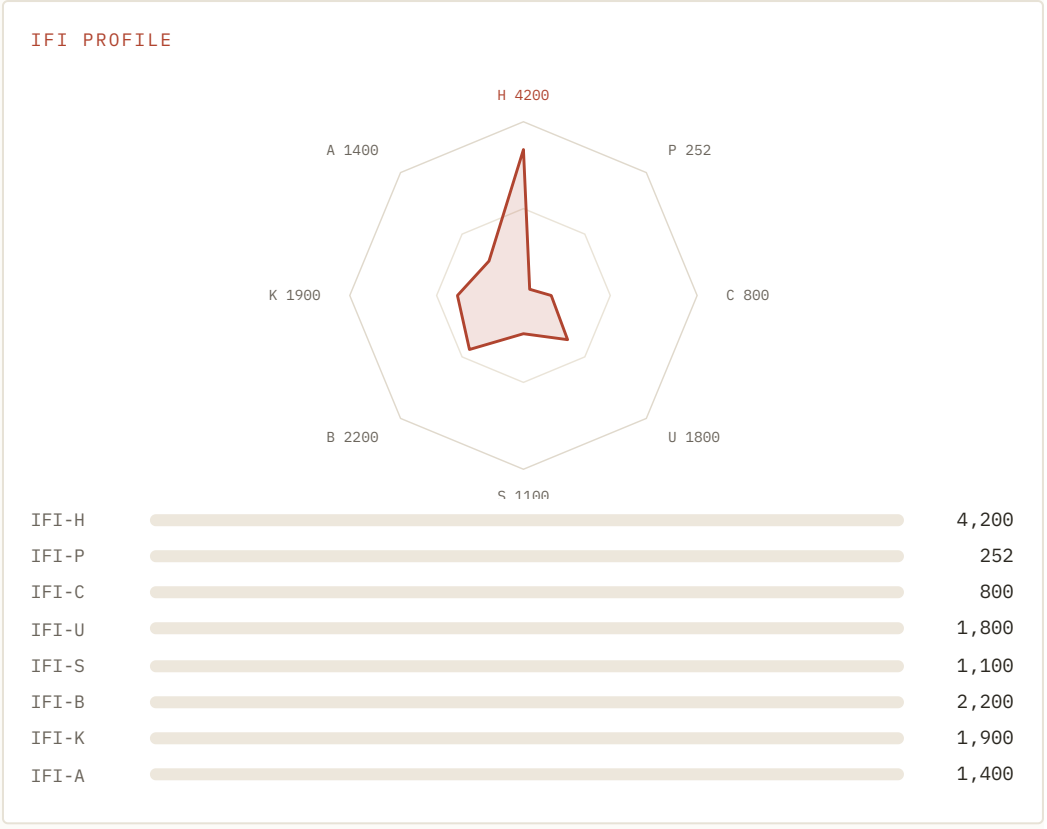

Lay the flagship superhots out as a ladder and a subtlety appears that justifies the whole receptor-anchored approach. Alongside each cultivar's predicted Scoville Heat Units and its predicted IFI-H, compute the loudness-relative EC_{50} (simply 1,000 ⁄ IFI-H, in units of the reference) and the predicted SHU carried per IFI-H point:

CULTIVAR · SKU	PREDICTED SHU	PREDICTED IFI-H	REL. EC_{50}	SHU PER IFI-PT
Ember · SS-0087	866,000	1,570	0.637	552
Cinder · SS-0088	1,050,000	2,290	0.437	459
Kindle · SS-0089	1,250,000	2,000	0.500	625
Blaze · SS-0091	1,720,000	2,630	0.380	654
Wildfire · SS-0092	1,960,000	2,930	0.341	669
Scorch · SS-0093	2,210,000	3,220	0.311	686
Firestorm · SS-0094	2,470,000	3,490	0.287	708
Solara · SS-0095	2,720,000	3,980	0.251	683
Phoenix · SS-0096	2,937,948	4,200	0.238	700

The SHU-per-point column should be flat if Scoville units and receptor loudness measured the same thing. They do not, and the proof is two rungs that cross. **Kindle** carries more total capsaicinoids than **Cinder** — 1.25 million Scoville units against 1.05 million — and yet it reads *quieter* at the receptor, IFI-H 2,000 against 2,290. The Scoville number counts molecules by dilution; the Index counts how hard those molecules actually push on TRPV1, and Kindle's mixture is weighted toward capsaicinoid species the pocket reads less loudly. A fruit can hold more fire by the old assay and deliver less of it to the one jury that matters. That is why we anchor to the receptor. The Scoville scale tells you how much is present; the Index tells you how much is *heard*.

Scoville units are still worth keeping, because they tie the whole ladder to a physical absolute. Pure crystalline capsaicin is defined at 16,000,000 SHU — the dilution at which its heat finally disappears. Our model predicts Phoenix at approximately 2,937,948 SHU, which is 2,937,948 ⁄ 16,000,000 = 0.184 of that ceiling: by capsaicinoid load, our Phoenix design computes at a touch over eighteen per cent as pungent as pure capsaicin, undetectable only after a dilution of nearly three million to one. We find that number steadying rather than frightening. It locates even our most extreme cultivar firmly on the edible side of a line drawn by chemistry itself.

Measuring a moving target

A loudness number is only as honest as the reproducibility of the thing being measured, and here the plant fights back. Wild pungency is notoriously variable — grow the same superhot landrace twice and its heat can swing by nearly half from pod to pod, a batch-to-batch coefficient of variation around **42%**. For a cook that is folklore; for a formulator it is ruin. Apply our predicted mean of 2,937,948 SHU for Phoenix and spread a 42% CV around it and the batches would run from about 1.70 million to 4.17 million Scoville units — the difference, pod to pod, between a sauce and an incident.

Our constructs hold a coefficient of variation of **≤7%**. At the same mean that is a one-standard-deviation band of ±205,656 SHU — roughly 2.73 to 3.14 million — tight enough that a recipe scaled on a jeweler's scale lands where it was aimed. In the Index's own units the contrast is just as stark: ±7% on an IFI-H of 4,200 is ±294, a reading of 3,906 to 4,494; the wild ±42% would be ±1,764, a reading from 2,436 all the way to 5,964 — past the scorecard's normalization ceiling of 5,000, off the top of the page. And consistency compounds in the formulator's favor: blend three

independent batches and the uncertainty of the mean falls as the CV over the square root of the count, from 7% to about 4.0%, where the wild plant would still be carrying 24%.

That steadiness is not luck; it is designed into the enzyme that makes the ligand. Two substitutions written into Pun1 — **S39L** and **L345G** — lower the relevant folding free energy (ΔΔG of −2.1 and −1.8 kcal/mol), and free energy converts to a fold-change in the underlying equilibrium through the Boltzmann relation, e^(−ΔΔG/RT), with RT ≈ 0.62 kcal/mol at body temperature. Individually they shift the equilibrium about 30- and 19-fold; together, −3.9 kcal/mol, they move it by roughly e^(3.9/0.62) ≈ 540-fold toward the folded, productive state. A protein that rarely misfolds makes a capsaicinoid output that rarely wavers, and a steady output is what lets the receptor measurement mean something twice. Reverence for the plant and respect for the reader turn out to be the same discipline: make the fire reproducible, then measure it exactly.

Calcium, and the second life of the molecule

Two further properties of the channel shape how heat feels and how we design it. The first is **calcium**. TRPV1 is not merely cation-permeable; it strongly favors calcium, passing it roughly ten times more readily than sodium. That calcium influx is what makes capsaicin more than a signal — it is also the basis of **desensitization**. Flood a nerve ending with calcium and a calcium–calmodulin feedback loop begins to close the very channels that let it in, and the receptor goes quiet, and stays quiet for a while. This is tachyphylaxis: the reason the third chip is less fierce than the first, and a behavior we account for in dish design. Heat accumulates and then plateaus, so a well-built dish doses the burn to arrive, hold, and release rather than simply climb — a curve, not a cliff.

Push that same calcium-driven quieting far enough and it becomes something more profound. Flood a sensory nerve with enough capsaicin for long enough and the fiber does not merely fall silent — it withdraws, the fine endings retracting, the neuron going dark to pain for weeks before it slowly regrows. Medicine turns this to use: a high-dose capsaicin patch, applied under clinical supervision, can quiet certain stubborn nerve pains precisely by exhausting the TRPV1 fibers that carry them, an irritant pressed into service as an analgesic. We are not in the business of medicine and we make no claim to it; we note the fact only because it is the clearest possible proof that this receptor is a serious piece of biology and not a party trick. A molecule that can switch off pain is a molecule to be dosed, never dared. That principle governs the entire kitchen half of this book, and it is why the recipes at the top of the ladder are weighed on a jeweler's scale and not guessed with a spoon.

Not you, them

There is an evolutionary story folded into all of this that our family never tires of telling, because it is the deepest answer to why heat exists at all. Mammals feel capsaicin; birds do not. Avian TRPV1 differs at a handful of the very residues that build the vanilloid pocket, so to a bird a ghost pepper is merely a sweet red fruit. This is not an accident of chemistry but a strategy of the plant. A pepper needs its seeds carried far and dropped intact, and birds oblige —

they fly, they do not grind the seeds to paste the way a mammal's molars would, and they pass them unharmed to germinate at a distance. Capsaicin is, in the most literal sense, a plant's chemical instruction to the animal kingdom: *not you, them.* The fire that fills this book was invented by a fruit as a way of choosing its friends — a ten-million-year-old negotiation between a plant and the sky, which humans, alone among mammals, answered not by staying away but by leaning in. We are the animal that accepted an invitation that was never addressed to us, and made a cuisine of it.

The ceiling we design beneath

The receptor belongs to a family, and the family is the rest of the orchestra. TRPV1 has cousins — **TRPM8**, which answers to cold and to menthol, and **TRPA1**, which answers to mustard and wasabi and cinnamon — and those cousins are the biological basis of the *other* axes we design against. When we build a cooling burn, a numbing tingle, a sharp green pungency, we are composing across this receptor family, a chord struck on channels that share TRPV1's architecture and its logic of binding, lever, and gate. Heat is the first instrument we learned. It taught us the others.

A closing note on just how potent this binding can become, because it sets the ceiling we design beneath. There exists a natural TRPV1 agonist called **resiniferatoxin**, from a Moroccan spurge, that binds the same vanilloid pocket with such ferocity that its Scoville-equivalent runs to roughly sixteen *billion* — 16,000,000,000, exactly a thousand times pure capsaicin, which is to say it half-activates the channel at something near a thousandth of capsaicin's concentration, down in the low nanomolar. We mention it not to chase it; it is a research neurotoxin, not a flavor, and it will never touch a plate of ours. We mention it because it proves the receptor has enormous dynamic range — far more than any edible plant has ever used. The space above the hottest pepper is real, and it is vast. Our Phoenix design, computed at roughly eighteen per cent of pure capsaicin by capsaicinoid load, has climbed only into the foothills of what this one lock can register. We have taken the first few steps into that space, and we have taken them, always, on the edible side of the line — toward a fruit grown to carry the most it can and give the body back more than heat. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.



Phoenix in the controlled chamber — the first cultivar our family engineered to register across all eight axes at once, grown to the edge of what the receptor can read and no further.

Chapter 2 · The Capsaicinoids — *a family, not a molecule.*

We say "capsaicin" the way we say "sugar" — a tidy singular for something that is really a small committee, and like any committee it governs by proportion. The burn of a pepper pulled warm off the plant, and the burn of anything we design to carry fire, is never the work of one molecule. It is produced by a family of closely related compounds, the capsaicinoids, and the *ratio* among them is not a footnote to the heat. It is the difference between a fire that stabs and a fire that swells; between a warmth that lives on the lips and one that settles at the back of the throat and waits. The receptor is the instrument. The capsaicinoid profile is the phrasing — and phrasing, as anyone who has tasted two "equally hot" peppers and flinched at only one already knows, is most of the music.

Our family has spent years learning to read that phrasing as numbers, because you cannot tune by hand what you cannot first write down. So this chapter does two things at once, the way we try to do everything: it lingers over what the fire *feels* like — where it lands, how it arrives, how long it stays — and it insists, in the same breath, on the chemistry and the arithmetic that make the feeling inevitable. The wonder and the math are not two subjects here. They are one subject seen from two sides.

The two-part molecule — a head the receptor reads, a tail that tunes the reading

Every capsaicinoid is built on a single architectural plan: a **vanillylamine head** joined by an **amide bond** to a **branched fatty-acid tail**. The head is the business end — a 4-hydroxy-3-methoxybenzyl group, the same vanillyl motif that gives vanilla its name, and the part that actually docks in the receptor's binding pocket. The tail is the tuner. Lengthen it, saturate it, move its lone double bond, and you do not change *whether* the molecule speaks to the receptor; you change how snugly it seats, how fast it arrives, how slowly it leaves, and where on the map of the mouth its fire seems to fall.



Wildfire in cross-section — the capsaicinoid-dense placenta glistening along the pod wall. In our superhot chassis this pale, pithy tissue is the most productive organ in the plant, and its exact capsaicinoid ratio is a design specification, not an accident of the season.

Capsaicin itself is **$C_{18}H_{27}NO_3$**, a molar mass of 305.41 grams per mole, and it is stubbornly fat-loving — a partition coefficient (logP) in the mid-threes, which is the quiet physical fact behind nearly everything the molecule does in a mouth: why it will not rinse away with water, why a spoon of yogurt or whole milk puts the fire out where a glass of iced water only spreads it. Five capsaicinoids account for essentially all natural pungency, and they differ from one another only in that tuning tail.

CAPSAICINOID	FORMULA	MW (G/MOL)	ACYL TAIL (C:UNSAT)	PUNGENCY *K* (SHU/PPM)	EC_{50} (MM)	PER-MOLE IFI-H
Capsaicin	$C_{18}H_{27}NO_3$	305.41	9:1 (Δ6, 8-methyl)	16.1	0.70	1,000
Dihydrocapsaicin	$C_{18}H_{29}NO_3$	307.43	9:0 (8-methyl)	16.1	0.70	1,000
Nordihydrocapsaicin	$C_{17}H_{27}NO_3$	293.40	8:0 (7-methyl)	9.3	1.21	578
Homocapsaicin	$C_{19}H_{29}NO_3$	319.44	10:1	8.6	1.31	534
Homodihydrocapsaicin	$C_{19}H_{31}NO_3$	321.45	10:0	8.6	1.31	534

Capsaicin and dihydrocapsaicin together carry the great majority of the burn in almost any real pepper — typically eighty to ninety percent of the total — and they are also the two most potent, each weighing in at the full 16.1 Scoville Heat Units per part-per-million. The *only* difference between them is one double bond: capsaicin carries it, in the trans-6 position of a nine-carbon tail; dihydrocapsaicin has had that bond quietly reduced to a single bond. That lone unsaturation is the whole character shift. Capsaicin reads sharper, more immediate, more forward — a struck match, quick to arrive and quick to peak. Dihydrocapsaicin is rounder, slower, more patient; it is the one still glowing at the back of the mouth long after you thought the dish had let you go. The three minor members paint the edges. Nordihydrocapsaicin, the shortest-tailed of the five, is the friendliest — fruitier, faster, front-of-mouth, gone early. The two homo-compounds, longest-tailed, are the opposite: low, slow, and almost oppressively persistent, the long tail of the long tail.

How TRPV1 reads the head and times the tail

To understand why a double bond can change a flavor, you have to go to the receptor. Heat on the tongue is not a temperature; it is a channel. **TRPV1** is a non-selective cation channel standing in the membrane of the trigeminal nerve endings woven through the mouth — a homotetramer, four identical subunits arranged around a central pore, each subunit threading the membrane six times. Between the third and fourth of those transmembrane helices, on the inner leaflet side of the membrane, sits the **vanilloid pocket**: the pocket the capsaicinoid head was shaped to fill. Tyrosine-511 and threonine-550 line it; the vanillyl hydroxyl and methoxy groups hydrogen-bond to them, anchoring the head, while the fatty-acid tail threads back out into a hydrophobic groove and out toward the lipid.

Because the entrance faces the membrane's interior, capsaicin does not swim up to the pocket through saliva. It partitions *into* the membrane first — this is what its fat-loving nature is *for* — and then slides in laterally, from within the bilayer. That single mechanical

detail explains the kitchen folklore exactly: water cannot reach a molecule that lives in the fat, but cream and whole milk can dissolve it back out and carry it away.

Once seated, the molecule does something subtle and total. TRPV1 is, in its day job, a thermometer: it opens when the tissue climbs past about 43 °C, and the brain reads that opening as *hot.* Capsaicin does not warm anything. It binds the open conformation and lowers the channel's thermal threshold — drags the trigger temperature down from 43 °C toward body heat — so the channel flickers open at the temperature the mouth already is. Calcium and sodium pour through; the nerve ending depolarizes; an action potential climbs toward the brain, which reports, with complete sincerity, a fire that is not there. It is an honest signal about a false premise, and it is one of the most elegant lies in physiology.

Here is where the tail earns its keep. How deeply and how comfortably the tail seats in that hydrophobic groove sets two things at once: the **affinity** (how little is needed to half-open the channels — the EC_{50}) and the **kinetics** (how fast the molecule binds, *k*on, and how reluctantly it lets go, *k*off). A longer, fully saturated tail — dihydrocapsaicin, and more so the homodihydro compound — is more lipophilic, partitions deeper into the membrane, and clears more slowly; its *k*off is lower, so it lingers. A shorter tail, or the kink of an unsaturation, seats more loosely, arrives faster, and leaves faster. The "map of the mouth" that experienced tasters describe — bright front-of-tongue heat versus slow back-of-throat heat — is not mysticism and not metaphor. It is diffusion, receptor density, and *k*off, made perceptible. A fast molecule reaching the densely innervated tip of the tongue reports first and fades forward; a slow, fat-loving one reaches the soft palate and throat last and leaves last.

Reading the fire as a number

To design against a receptor we need a single number for how *loudly* a stimulus speaks to it, and that number is the half-maximal effective concentration, the **EC_{50}** — the concentration at which the stimulus drives its receptor to half of full response. A potent compound needs very little to reach half; its EC_{50} is small. The Iyengar Flavor Index inverts that, so that louder reads as higher, and scales it against a fixed reference:

$$IFI(candidate) = 1{,}000 \times EC_{50}(reference) \div EC_{50}(candidate)$$

The logic is worth holding in the hand. If a candidate has exactly the reference's EC_{50}, the ratio is one and the score is 1,000 — the anchor. If the candidate is twice as potent (half the EC_{50}), the score doubles to 2,000. The Index is just potency, normalized and turned right-side up. Each reference constant defines IFI = 1,000 on that axis by convention — a stated definition, the same way Scoville's dilution-to-threshold defined his scale — not a value derived from any single published measurement.

Let us work it molecule by molecule. For a clean, *molecular* comparison we anchor the reference to pure capsaicin itself, whose EC_{50} at heterologously expressed human TRPV1 sits near **0.70 µM** — about 214 parts per billion, which is to say the channel is already

half-open at a concentration you would struggle to taste as anything but heat. (The *whole-fruit* Index uses a different anchor, a native wild-type pepper; that is a different zero for a different question, and we keep the two apart on purpose.) Dihydrocapsaicin shares capsaicin's potency, so both land at IFI-H 1,000. Nordihydrocapsaicin is less potent by exactly its pungency factor — 9.3 against 16.1 — so its EC_{50} rises to 0.70 × (16.1 ÷ 9.3) = **1.21 µM**, and

> IFI-H(nordihydrocapsaicin) = 1,000 × 0.70 ÷ 1.21 = **578**

The homo-compounds, at a factor of 8.6, give EC_{50} = 0.70 × (16.1 ÷ 8.6) = **1.31 µM**, and IFI-H = 1,000 × 0.70 ÷ 1.31 = **534**. Those are the per-molecule scores in the table above, and the arithmetic that built them is the whole method in miniature.

There is a quiet marvel hiding in that agreement, and we never tire of it. In 1912, Wilbur Scoville had no receptor to measure — TRPV1 would not be cloned for another eighty-five years. He diluted a pepper extract in sugar water until a panel of tasters could no longer detect the burn, and called the dilution factor the heat. A detection threshold is not a half-maximal point; it sits lower on the curve. And yet the *ratios* his dilution series produced between one capsaicinoid and the next — the very factors 16.1, 9.3, 8.6 that still weight every heat calculation today — are, to a near-astonishing degree, the inverse ratios of those compounds' EC_{50} at TRPV1. Scoville's tasters were reading the receptor a century before anyone could see it. The dilution test and the dose–response curve are the same measurement, taken a hundred years apart.

That dose–response curve has a shape, and the shape matters. TRPV1 activation by capsaicin is *cooperative* — a Hill coefficient of roughly 2 — so the response does not rise in a gentle proportional line but in a steepening S. With a Hill coefficient *n*, response as a fraction of maximum is

> $f = [A]^n \div (EC_{50}{}^n + [A]^n)$

and with n = 2 and EC_{50} = 0.70 µM the curve runs like this:

[CAPSAICIN]	×EC_{50}	TRPV1 RESPONSE (N = 2)
0.18 µM	0.25×	5.9%
0.35 µM	0.50×	20.0%
0.53 µM	0.75×	36.0%
0.70 µM	1.00×	50.0%
1.05 µM	1.50×	69.2%
1.40 µM	2.00×	80.0%
2.80 µM	4.00×	94.1%
7.00 µM	10.0×	99.0%

Read the steep middle of that column. A mere doubling of concentration, from half the EC_{50} to twice it, hauls the response from 20% to 80%. With cooperativity, the span from a barely-there 10% to a nearly-saturated 90% takes only a nine-fold change in

concentration; a non-cooperative channel would need to see concentration climb eighty-one-fold to cover the same ground. That is why chili heat does not feel like a dimmer easing up. It feels like a threshold you cross — a little more and suddenly it is *everything* — and the felt suddenness is written directly into the exponent of the Hill equation.

A word on the units, because the conversions are where sloppiness hides. One micromolar of capsaicin is 1×10^{-6} mol/L × 305.41 g/mol ≈ 0.305 mg/L, which in a watery medium is about 0.305 ppm; so 1 ppm ≈ 3.27 µM, and the 0.70 µM EC_{50} is the 214 ppb quoted above. On the Scoville side, because pure capsaicin is defined at 16,000,000 SHU, each ppm of it contributes 16.1 SHU — the factor in the table, now revealed as nothing more exotic than a unit conversion between concentration and dilution.

The weighted sum — what Scoville sees, and what it misses

The total heat of a real pepper is just the pungency-weighted sum of its parts. Measure each capsaicinoid by HPLC in parts per million, multiply each by its factor *k*, and add:

$$\mathrm{SHU} = \Sigma\ (\mathrm{ppm}_i \times k_i)$$

Now watch what this does — and does not — capture. Consider two placental extracts carrying the *same* total capsaicinoid mass, say 100,000 ppm, but phrased differently. A **bright** profile leans on capsaicin: 60% capsaicin, 28% dihydrocapsaicin, 8% nordihydrocapsaicin, 2% each of the homo-pair. Its weighted factor is

$$0.60(16.1) + 0.28(16.1) + 0.08(9.3) + 0.02(8.6) + 0.02(8.6) = 15.26\ \mathrm{SHU/ppm}$$

for a total of **1,525,600 SHU**. A **deep** profile shifts weight onto the lingering compounds: 38% capsaicin, 48% dihydrocapsaicin, 4% nordihydrocapsaicin, 5% each homo. Its factor is 15.08 SHU/ppm, for **1,507,800 SHU**. The two totals differ by about 1.2% — inside the batch-to-batch error of the measurement itself. By Scoville, they are the same pepper. In the mouth they are nothing alike: the first arrives fast and forward and releases; the second builds like a tide and refuses to recede. Scoville measures the volume and is deaf to the phrasing. Designing against the receptor, we are not — and that gap is precisely why a book about heat needs a chapter about ratios.



Phoenix (SS-0096) in the controlled chamber — the first cultivar we engineered to register on all eight axes at once, and the clearest proof that a capsaicinoid ratio can be a written specification rather than a hope.

Run the sum in reverse and it tells you where the fire has to live. **Phoenix**, the flagship design in our superhot series, carries a predicted heat of approximately **2,937,948 SHU** in our model. Divide by capsaicin's 16.1 and you get its capsaicin-equivalent load: about **182,481 ppm** — 18.2% — and because its minor capsaicinoids pull the weighted factor down toward 15.25, the true capsaicinoid mass is nearer 19%. No whole fruit is one-fifth capsaicinoid. But the *placenta* can approach it in an optimized design, and that is the point: this is a placental, oleoresin-basis number, concentrated in the pale ribbed tissue the cross-section shows, which is exactly the organ a designer must aim the pathway at. And Phoenix's heat is not a flat wall — it is tuned. Its full eight-axis signature, the reason it is a full-spectrum cultivar and not merely a hot one, reads:

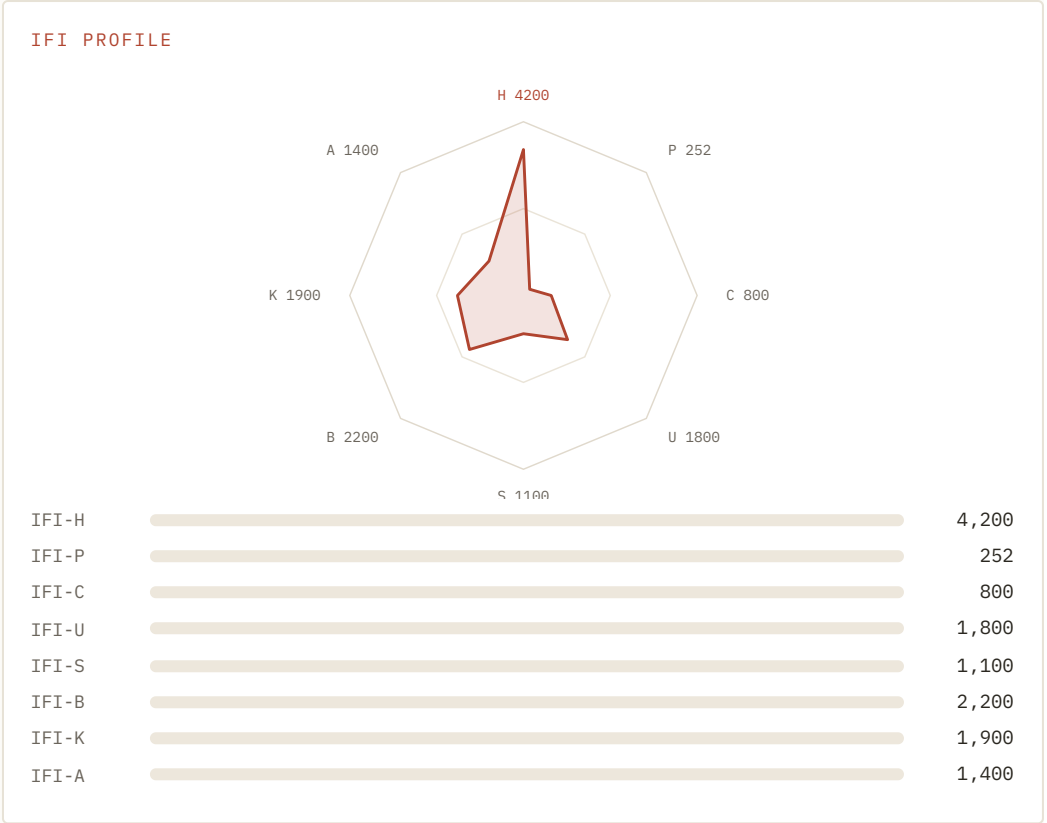

That leading **H: 4200** is a direct statement about the receptor. Inverting the Index, our model predicts Phoenix's extract would reach half-activation of TRPV1 at EC_{50} = EC_{50}(reference) × 1,000 ÷ 4,200 — about one-quarter the reference concentration. Four-fold louder than the native anchor, written as a single integer, and designed to hold that potency consistently across batches.

The heat ladder — profile as a design specification

Because the heat is a weighted sum and the weighting is set by which capsaicinoids are present, total pungency and receptor loudness are *related* but not identical — and the gap between them is, once again, the ratio. Lay the flagship superhots out as a ladder of increasing computational design scope and axis activation — not as certified ascending measurements, but as a progression of design targets — and the relationship between predicted pungency and receptor loudness becomes legible.

CULTIVAR	SKU	PREDICTED SHU	PREDICTED IFI-H	PREDICTED CAPSAICIN-EQUIV. (PPM)	PREDICTED SHU PER IFI-H
Ember	SS-0087	866,000	1,570	53,789	552
Kindle	SS-0089	1,250,000	2,000	77,640	625
Blaze	SS-0091	1,720,000	2,630	106,832	654
Wildfire	SS-0092	1,960,000	2,930	121,739	669
Scorch	SS-0093	2,210,000	3,220	137,267	686
Firestorm	SS-0094	2,470,000	3,490	153,416	708
Solara	SS-0095	2,720,000	3,980	168,944	683
Phoenix	SS-0096	2,937,948	4,200	182,481	700

The capsaicin-equivalent column is just SHU ÷ 16.1 — the concentration of pure capsaicin that would produce the same burn, climbing from about 5.4% in Ember to 18.2% in Phoenix The final column is the honest one. If total heat and receptor loudness were perfectly proportional, SHU per IFI-H unit would be a constant; instead it wanders across the ladder, roughly 550 to 710. Some of that wander is the capsaicinoid ratio doing exactly what this chapter says it does — a cultivar skewed toward the less-potent minor compounds carries more Scoville mass per unit of loudness. Some of it is the Hill curve's nonlinearity, since a linear dilution and a half-maximal point do not march in lockstep. And some of it is simply prediction scatter in the model, which we would be dishonest to launder out. We report the spread rather than pretend it is a line, because a number you have to defend is worth more than one you have to apologize for.

This is why, for every cultivar we release, we specify a *burn signature* — a target capsaicin-to-dihydrocapsaicin ratio — right alongside its total pungency. Phoenix is not merely the loudest design we have produced; it is tuned so that its enormous heat has an architecture, an arrival and a sustain, a front and a back, instead of a single undifferentiated plateau of pain. Total heat is the volume knob. The ratio is the equalizer. A design that set only the first would be flying a plane with half its instruments dark.

The enzymatic lever — why the ratio is something we can specify

None of this would be designable if the profile were fixed by fate. It is not. It is set enzymatically, at a single condensation step, and that is the lever.

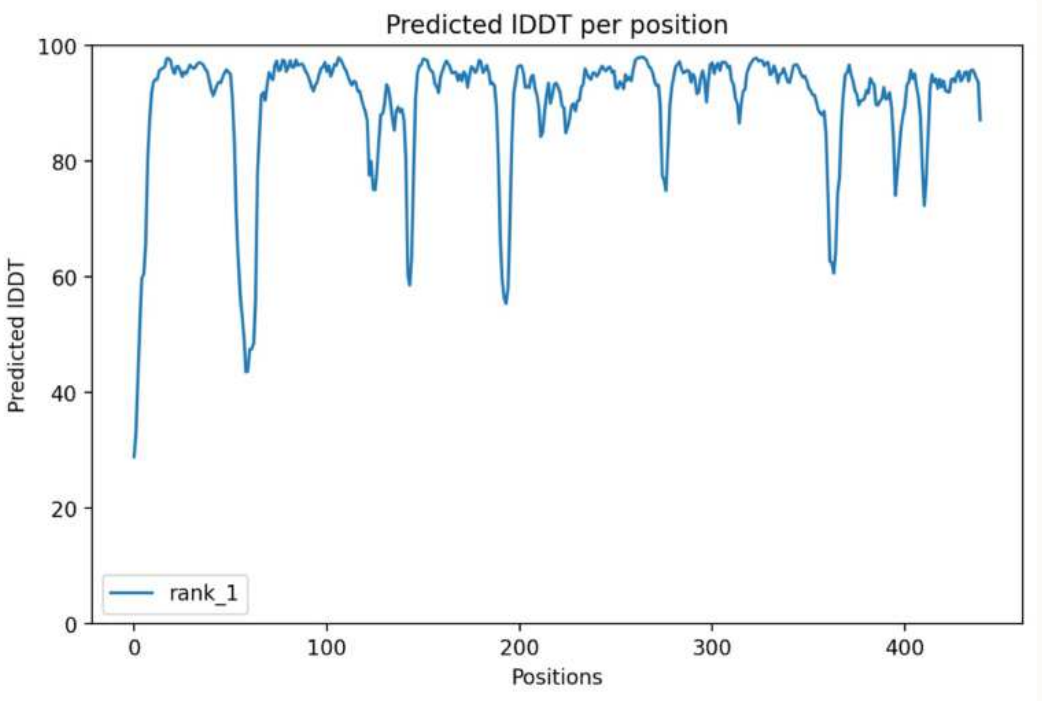

FIG. 7 — the Pun1/AT3 capsaicin synthase, coloured by model confidence; the binding domain and active-site cleft, where the vanillylamine head and the acyl tail are joined, resolve at high confidence. This is the enzyme whose substrate choice *is* the burn signature.

Two metabolic rivers converge to make a capsaicinoid. The first builds the head. It runs down the **phenylpropanoid branch** that every vascular plant already owns — phenylalanine to cinnamic acid by **PAL**, then **C4H**, **4CL**, **HCT**, and **COMT** carrying it along to ferulic acid and on to vanillin — and finishes with a pepper-specific transaminase, **pAMT**, that converts vanillin to vanillylamine. The second river builds the tail. Branched-chain amino acids, valine and leucine, are handed by **BCAT** into branched-chain fatty-acid metabolism, where a ketoacyl synthase (**KAS**) runs the elongation cycles that set the chain's length, and a thioesterase clips the finished acyl group free as its coenzyme-A thioester.

Then the two rivers meet at **Pun1**, the capsaicin synthase — the acyltransferase in the figure — which forms the amide bond between the vanillylamine head and whichever acyl-CoA the fatty-acid river has delivered. And *that* is the whole secret of the ratio. Which capsaicinoid you make is decided by which acyl-CoA is on offer and which Pun1 is willing to accept. Chain length is set by how many elongation rounds KAS runs before the thioesterase cuts; saturation is set by whether a desaturation step acts before the hand-off. Feed Pun1 a shorter, unsaturated tail and the output tilts toward capsaicin and its bright, forward burn; let the machinery run longer and more saturated and the output deepens toward the lingering compounds. The burn signature, in other words, is the output distribution of the acyl-CoA pool, read through Pun1's preferences — and both the pool and the enzyme are things we can write. We do not coax the ratio by selecting phenotype across seasons and hoping. We specify it, the way you specify any other number in a design.

The cousin that maps the edge — capsinoids

There is a near-twin to the capsaicinoids worth knowing, because it draws the exact boundary of what heat *is.* A few peppers — the Japanese CH-19 Sweet is the famous one — make **capsinoids**

instead: capsiate and its relatives, molecules all but identical to capsaicin save for one bond. Where capsaicin joins head to tail with an **amide** (a nitrogen-carbon linkage), capsiate uses an **ester** (an oxygen-carbon linkage). On paper it is the smallest edit imaginable. In the mouth it changes everything.

Capsiate engages TRPV1 much as capsaicin does — its intrinsic affinity for the pocket is not the problem. The problem is that the ester bond is fragile where the amide is tough. Esterases in saliva and the oral mucosa hydrolyze capsiate quickly, snapping it back into vanillyl alcohol and a free fatty acid before it can build any real concentration on the tongue; its residence time at the oral receptors is too short to matter. Only lower down, past the mouth, where it survives long enough to reach the gut's own TRPV1, does it register at all. The result is a "heatless heat" — the signature of capsaicin with almost none of the oral fire. We grow these at the family plots not for the kitchen but for the map: they mark, with a single bond, the frontier of the heat axis. Pungency, they teach, is not a property a molecule owns in isolation. It is an event — a lipophilic head reaching an intact receptor on the surface of the tongue and staying long enough to be heard. Nature already built a molecule that does nearly everything capsaicin does except the one thing we call heat, and knowing precisely where that line falls is half of knowing how to design across it.

The seed myth, and what it tells a designer

One stubborn piece of kitchen lore deserves correcting, because it changes how a cook handles every pepper and how a designer aims every construct: **the heat is not in the seeds.** Capsaicin is synthesized and stored in the placenta — the pale, pithy ribs and cross-walls that cradle the seeds — not in the seeds themselves, which carry only what rubs off their neighbors in the tight space of the pod. When a recipe tells you to "remove the seeds to cut the heat," what actually lowers the burn is scraping away the placental tissue the seeds cling to; the seeds are merely the visible passenger that comes away with it.

That geography is a design instruction. It tells us where in the fruit the pathway runs, and therefore where a promoter must point. When we move fire into a tomato or an eggplant that never evolved to carry it, we do not scatter the pathway through the whole fruit; we target the equivalent placental and inner-wall tissue, reproducing not only the pepper's chemistry but its *architecture* — the heat kept in the flesh a cook will cut, and out of the seed a plant spends its whole reproductive strategy protecting. Fidelity, for us, has always meant the whole plan, down to where the fire is allowed to live.

What the number cannot say — methods, error, and honest limits

A number is a form of respect for the reader only if it is honest about what it does not know, so we end where every real method ends: at its limits.

The Scoville organoleptic test, for all its historical beauty, was a human panel tasting dilutions, and human panels drift — between people, between days, by as much as half in either direction. That is why precision work abandoned it for high-performance liquid chromatography, which resolves an extract into its five peaks and measures each. But HPLC has its own honesties to confess: the homocapsaicin isomers can co-elute and blur; the detector responds to each capsaicinoid a little differently and must be calibrated against pure standards; and the minor peaks, being small, carry the largest relative error. There is a reporting ambiguity we refuse to hide, too — pungency can be quoted on a fresh-weight or a dry-weight basis, and the two differ severalfold, so a SHU figure with no stated basis is half a sentence. And the old secret of the trade: the marketing Scoville numbers that circulate for the wild superhots often run well above what careful HPLC returns. We quote computed, placental-basis, capsaicin-equivalent figures throughout, and we would rather a reader know that than be dazzled.

The deepest honesty is about **consistency**, and here the engineering earns its name. Batch-to-batch, a wild-type superhot varies enormously — a coefficient of variation near **42%**. The coefficient of variation is just the standard deviation as a fraction of the mean, σ ÷ μ, and at 42% it is ruinous for anyone building a product to a spec. The variance is not mostly in the measurement — HPLC itself runs a tight 2–4% — it is biological: capsaicin synthesis answers to light, to water stress, to a fruit's position on the plant and its exact day of ripeness, and all of that noise lands on the final number. By designing the pathway's upstream control architecture — our flux model places the rate-determining steps at 4CL and HCT/CCoAOMT, well upstream of the terminal Pun1 synthase — we drain that biological variance out, and our design targets a batch CV of **≤7%**, approaching the measurement floor itself. The difference is not abstract:

DESIGN TARGET SHU	WILD-TYPE (CV 42%), ±1Σ	ENGINEERED DESIGN TARGET (CV ≤7%), ±1Σ
1,000,000	580,000 – 1,420,000	930,000 – 1,070,000
2,937,948 (Phoenix)	1,704,000 – 4,172,000	2,732,000 – 3,144,000

A formulator promising a million Scovilles on a label, working from wild fruit, is really promising something between 580,000 and 1.42 million — and at two standard deviations, almost anything at all. Working from an engineered line, the same promise holds inside a band a tenth as wide. For Phoenix, our design targets a predicted SHU band that would keep every harvest between 2.73 and 3.14 million; a 42% process would scatter the same plant from 1.7 to 4.2 million. Reproducibility is not a luxury bolted onto flavor design. It is the thing that lets a flavor be designed at all.

And the Index itself keeps one honest caveat, the same one this whole chapter is built around. EC_{50} is measured on receptors in a dish, at a controlled concentration, in a clean buffer; a mouth is warm, fatty, acidic, and in motion, and those matrix conditions shift a molecule's *effective* potency from its textbook value. More than that, IFI-H measures loudness — how little is needed to half-open the channel — and loudness is not residence. Two peppers can post the

identical IFI-H and the identical SHU and still be different fires, because the one thing a single scalar cannot hold is the *shape* of the burn over time: its arrival, its peak, its long slow leaving. That shape is the capsaicinoid ratio. It is the half of the instrument the total was always deaf to, the phrasing under the volume — and it is, in the end, the half we find most worth tuning, because it is the half you actually taste.

These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Chapter 3 · The Pathway — *from phenylalanine to fire.*

To build a fire you must know how the plant builds it. Capsaicin is not stored in the seed or inherited ready-made; it is synthesized, molecule by molecule, in the gland cells of the placenta during fruit development, by a defined sequence of enzymes converting ordinary metabolites into something no other genus on Earth makes in quantity. There is no violence in the chemistry — only phenylalanine, a little valine, a few turns of machinery every green thing already runs — and yet the product is the one small molecule that can make the body report a heat that is not there. That sequence is the most important map in this book, because five of its seven working modules are not pepper-specific at all. They run in every vascular plant alive. Which means — and the whole promise of designed flavor depends on this — that the pathway is portable.

The fire is assembled from two tributaries that meet at a single confluence. One tributary builds the **head** of the molecule, the vanilloid ring the receptor will read. The other builds the **tail**, the branched fatty-acid chain that anchors the head to the membrane and tunes the character of the burn. Neither tributary, on its own, is pungent. The head is a cousin of vanilla; the tail is a cousin of the fats in every leaf. Only where they are joined — by one enzyme, in one tissue, at one moment of ripening — does a pepper become a pepper. We find it one of the quiet marvels of the living world that so much heat is, at bottom, an amide bond: a single nitrogen offered to a single carbon, and the whole scale of Scoville swings from it.

The phenylpropanoid corridor — building the head

The first tributary is the **phenylpropanoid branch**, the deep, universal pathway that every plant uses to make lignin, pigments, and defense compounds. It begins at the amino acid phenylalanine and proceeds down a corridor of enzymes so conserved they are nearly a plant's signature.

PAL (phenylalanine ammonia-lyase) opens the corridor by stripping the amino group from phenylalanine to make trans-cinnamic acid. It is a beautiful piece of chemistry: PAL carries its own built-in cofactor, the MIO group (a methylidene-imidazolone formed autocatalytically from its own backbone), and uses it to abstract ammonia in a single antiperiplanar elimination — no ATP spent, no reducing power, just a deft rearrangement that hands the rest of the pathway a clean cinnamate skeleton. **C4H** (cinnamate-4-hydroxylase), a cytochrome-P450 monooxygenase, then hydroxylates that ring to p-coumaric acid, spending one NADPH and one molecule of O_2 to insert a single oxygen atom. **4CL** (4-coumarate-CoA ligase) activates the acid as a coenzyme-A thioester, paying one ATP — cleaved all the way to AMP and pyrophosphate, so the true cost is two high-energy phosphoanhydride bonds — to load the molecule onto the carrier that will shuttle it through the rest of the route. Through **C3H**, **HCT**, and **COMT** the ring then acquires the extra hydroxyl and the

methoxy group that define the vanilloid; COMT does the methylation with S-adenosylmethionine (SAM) as the methyl donor, the plant's universal methyl currency. A final CoA-dependent chain-shortening clips two carbons from the side chain to yield **vanillin** — yes, the same vanillin that flavors half the world's desserts — which the aminotransferase **pAMT** converts to **vanillylamine**. That is the head of the capsaicin molecule, the part the receptor will read, delivered by machinery older than flowers.



A pepper in cross-section: the pale placental tissue arching along the inner wall, studded with the seeds and threaded with the gland cells where the whole pathway runs. Capsaicin is made here and almost nowhere else in the plant — the leaves and roots carry the identical genome and never burn.

Everything to this point is a plant being a plant. A moss makes cinnamate; a redwood makes feruloyl-CoA by the ton on its way to lignin; vanillin is a decorated phenylpropanoid a thousand lineages know how to make. The corridor is not the pepper's invention. The pepper's genius lies in what it does with the head once it has one.

The branched-chain route — building the tail

The second tributary is the **branched-chain fatty-acid route**, and it builds the tail. It starts from the branched-chain amino acids valine and leucine. **BCAT** (a branched-chain aminotransferase) removes the amino group from valine to give the α-keto acid; a branched-chain α-keto-acid dehydrogenase complex then decarboxylates it to **isobutyryl-CoA**, a four-carbon branched primer, releasing CO_2 and reducing one NAD^+ to NADH. This primer is the seed of the tail — the methyl branch that will sit, in the finished molecule, at the eighth carbon.

From there the **KAS** module (a ketoacyl-ACP synthase) elongates the chain the way all fatty-acid synthesis elongates: two carbons at a time, each round condensing a malonyl group onto the growing acyl chain, then reducing, dehydrating, and reducing again. Three rounds take the four-carbon primer to a ten-carbon branched acid — **8-methylnonanoic acid**, the saturated tail of dihydrocapsaicin. A desaturase introduces a single double bond between carbons six and seven to give **8-methyl-6-nonenoic acid**, the tail of capsaicin itself. A **thioesterase** decides when elongation stops, clipping the chain to its final length, and that single decision — nine carbons or ten, one double bond or none — is what sets the balance between capsaicin and dihydrocapsaicin in the fruit.

That balance is not a footnote; it is the *shape* of the burn. Capsaicin, with its unsaturated tail, tends to arrive faster and sharper — a bright front of heat. Dihydrocapsaicin, saturated, sits deeper and lasts longer — the slow, flat plateau that outstays the first flare. A pepper's character on the tongue is, in large part, the ratio the thioesterase and the desaturase set between these two, and that ratio is itself a quantity a designer can reach for: move the chain-length preference and you move not how *hot* the fruit is but how the heat is shaped in time. We think of the head as setting the volume and the tail as setting the attack and decay — the same note, played legato or staccato.

The confluence — the one enzyme peppers invented

The two tributaries meet at one enzyme, and it is the only one in the entire route that peppers invented for themselves: **capsaicin synthase**, encoded at the *Pun1* locus and also called CS or AT3. It is an acyltransferase of the BAHD superfamily, and its job is the confluence itself — it joins vanillylamine to the branched acyl-CoA, forming the amide bond, releasing coenzyme A, and producing capsaicin.

The mechanism is elegant and worth seeing closely, because it is where every edit we make must land. BAHD acyltransferases carry a conserved catalytic histidine in an **HXXXD** motif. That histidine acts as a general base: it plucks a proton from the incoming amine of vanillylamine, and the now-nucleophilic nitrogen attacks the carbonyl carbon of the acyl-CoA thioester. A tetrahedral intermediate forms and collapses, coenzyme A leaves as the thiol, and what remains is the amide — the bond that makes capsaicin capsaicin. A second conserved motif, **DFGWG**, holds the fold rather than doing chemistry. The enzyme is built around two channels that meet at the active site: one admits the long hydrophobic tail, the other the vanillylamine head, and catalysis happens exactly where they intersect. The architecture is the confluence, drawn in protein.

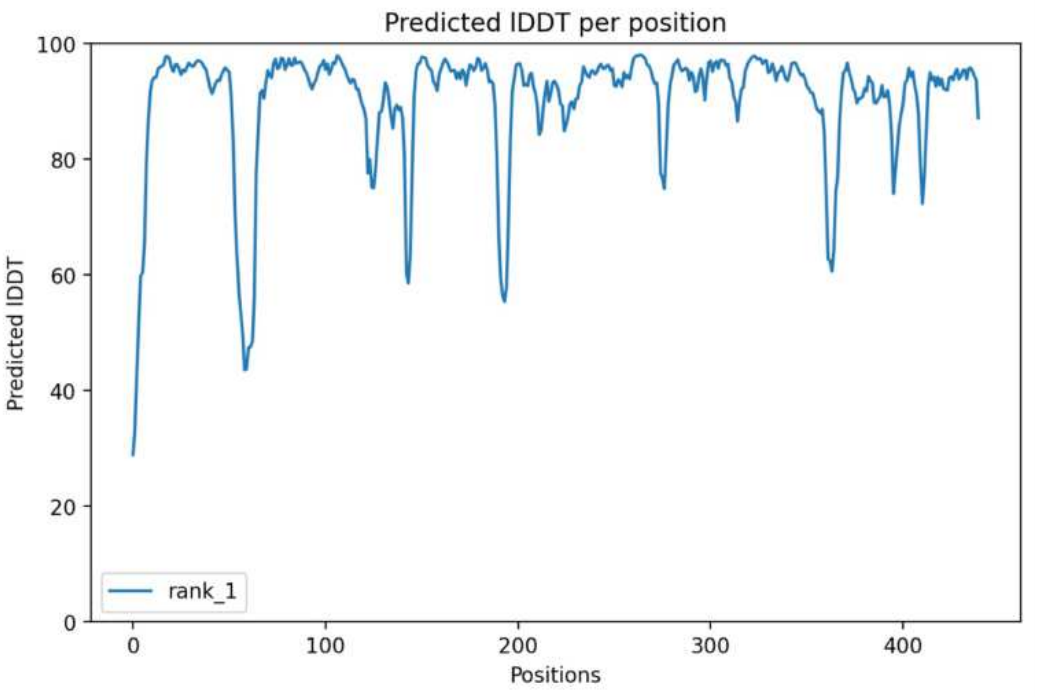

The capsaicin synthase enzyme, Pun1/AT3, as our structural model renders it — coloured by predicted local confidence (pLDDT). The deep-blue core and the active-site cleft are high-confidence; this is the one truly pepper-specific enzyme in the whole route, the junction where head meets tail and fire is made.

The *Pun1* locus is so central that its loss is precisely what makes a bell pepper sweet. A deletion there, inherited as a recessive, yields a plant with the entire phenylpropanoid machinery intact and no fire at

all, because the one enzyme that completes the molecule is gone. Pungency in *Capsicum* is, genetically, the presence of a working *Pun1*. The head is made, the tail is made, the two simply never meet. Everything upstream is just a plant being a plant; this last step is the pepper.

Two substitutions, written into the line

Because *Pun1* is the confluence, it is also the natural target. Every prior breeding program chased heat by selecting phenotype and hoping the genotype followed — taste the fruit, keep the seed, repeat for centuries. We inverted the order. A protein language model scored every possible single-residue substitution in capsaicin synthase by its predicted effect on catalytic efficiency and fold stability; a genome-scale flux-balance model predicted that control sits upstream in the phenylpropanoid pathway — at **4CL** and **HCT/CCoAOMT** — not at the terminal Pun1/AT3 step. Pun1 is the final capsaicin synthase; it labels the product, it does not govern the rate. The lever we chose was therefore fold reliability: a stably folded Pun1 completes the confluence for every substrate pair the upstream corridor delivers, minimizing the fraction that never meets a working active site. Two substitutions, never before combined in any pepper that ever grew, emerged from that search and were written into the line with CRISPR.

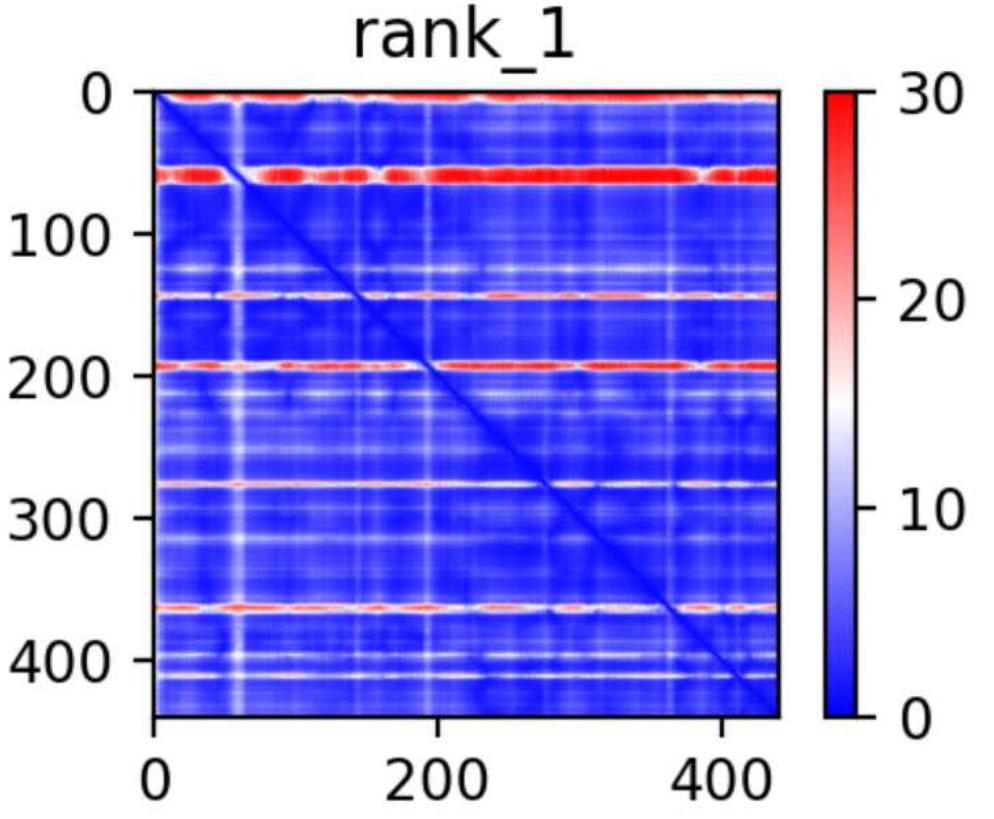

The predicted-aligned-error map for the same Pun1/AT3 model: a residue-by-residue grid of how confidently each part of the fold is placed relative to every other. Low error (the deep diagonal block) means the active-site domain is positioned with confidence — exactly the region the two edits touch.

EDIT	WHAT IT DOES	ΔΔG (KCAL/MOL)	IN SHIPPED CONSTRUCT
S39L	raises binding affinity near the N-terminus	−2.1	Yes
L345G	relieves steric strain in the active-site cleft	−1.8	Yes
C175S	removes a stray cysteine that misfolds via disulfide *(structural annotation only; not in shipped design)*	−0.9	No

Read the ΔΔG column as free energy, and it tells you something exact. Each value is the change in folding free energy the substitution confers; both shipped substitutions are negative, meaning both stabilize the folded, catalytically competent state. They sum to **−3.9 kcal/mol**. Free energy converts to an equilibrium constant through ΔG = −RT·ln K, and at 298 K the thermal quantum RT is 0.593 kcal/mol, so the two edits together multiply the folded-to-unfolded ratio by

```
e^(3.9 / 0.593) = e^6.58 ≈ 720.
```

Roughly seven-hundred-fold more of the enzyme sits in its working conformation at any instant. That is not an abstraction. A protein that folds reliably is a protein that turns over reliably, and an enzyme that turns over reliably gives a pathway whose output does not wander batch to batch. The stability we bought in free energy, we spend again as consistency — which is why the engineered line is designed to hold its pungency to a batch-to-batch coefficient of variation of **≤7%**, against **42%** in wild-type pods. We will return to that number with the arithmetic it deserves; for now it is enough to see that it begins here, in two substitutions' worth of free energy on one enzyme.

The price of a single molecule of fire

It is worth pausing to count what a pepper actually spends to make one molecule of capsaicin, because the sum is a kind of respect — a reminder that fire is metabolically expensive, built from the plant's own budget of energy and reducing power. The figures below are textbook stoichiometry, order-of-magnitude rather than exact, but they are honestly derived from the steps themselves.

SEGMENT	INPUTS CONSUMED	ATP (EQUIV.)	NADPH	OTHER
Head (Phe → vanillylamine)	1 phenylalanine	~2	~2	1 SAM (methyl)
Tail (Val → C_{10} branched acyl)	1 valine, 3 acetyl-CoA	~3	~6	+1 NADH recovered
Confluence (amide bond)	head + tail	0	0	releases 1 CoA-SH
Per capsaicin, approx.		**~5**	**~8**	**1 SAM**

Five ATP and eight NADPH is the price of one molecule. The confluence itself is free — the energy was already loaded into the thioester bond by 4CL and the elongation cycles, so capsaicin synthase spends nothing to spend it. This is why the terminal step is such a clean place to engineer: it adds no new energetic burden, it simply joins two pieces the plant has already paid for. And it is why a pepper confines the reaction so tightly. At the loadings a superhot

reaches, the fruit is pouring a meaningful fraction of its phenylalanine and its reducing power into a defense compound it cannot eat. No plant does that carelessly. The discipline of *where* and *when* is not decoration on the pathway; it is the plant protecting its own budget.

Where, and when — the discipline of the switch

Capsaicin synthase is not switched on everywhere the plant could make it. It fires specifically in the placental gland cells, and only during a defined window of fruit ripening, directed by a layer of transcription factors — MYB- and WRKY-family regulators among them — that read the fruit's developmental clock and open the pathway at exactly one place and time. This is why a pepper's leaves and roots, which carry the identical genome and the identical phenylpropanoid corridor, never burn: the terminal genes are present but held silent everywhere except the one tissue evolution chose.

For us that logic is not background — it is the template. When we move fire into a plant that never carried it, the hardest and most important part is not the enzymes but the **promoter**, the switch that must confine the new pathway to the fruit the way nature confined it in the pepper. An enzyme expressed in the wrong tissue is at best wasted budget and at worst a sick plant; an enzyme expressed in the right tissue at the wrong moment never meets its substrate. *Capsicum* did not merely invent capsaicin synthase. It invented the discipline of where to use it, and that discipline — the fruit-specific, ripening-timed promoter — is the part we copy most carefully. The molecule is the easy half of the design. The grammar of its control is the hard half, and the one we study most reverently, because it is where a borrowed pathway becomes a well-behaved one.

Two ways a pepper goes sweet

A pepper can lose its fire at two different valves, and the pathway reads cleanly as a circuit with two independent off-switches. The first is the *Pun1* deletion already described: lose the confluence enzyme and the two tributaries never join. The second is subtler and explains an entire class of "warm but not hot" peppers. The enzyme **pAMT**, which aminates vanillin into the vanillylamine head, is the gatekeeper of the amine supply. Break pAMT and the head cannot be aminated; starved of its amine, the pathway defaults to **capsinoids** — ester-linked cousins in which the acyl chain bonds to vanillyl *alcohol* instead of vanillyl *amine*. Capsinoids warm the gut but barely register in the mouth, because the receptor that reads heat is tuned to the amide, not the ester. One missing nitrogen, and a pepper is gentle.

OFF-SWITCH	GENE LOST	WHAT REMAINS	RESULT
Main valve	*Pun1* / CS	full head + full tail, never joined	`sweet pepper (bell), no heat`
Head-supply valve	pAMT	vanillin, no vanillylamine	`capsinoid pepper — warm, not hot`

Between them, these two natural deletions account for most of the non-pungent peppers humans have ever bred. And they shaped a century of agriculture before anyone could read a gene. Pungency in

Capsicum is inherited as a single dominant locus: cross a hot pepper with a sweet one and the first generation is hot, because one working copy of *Pun1* suffices to light the pathway. The sweet bell is a homozygous loss — two deleted copies — which is why bells breed true, and why a sweet-pepper row planted next to a hot one can throw surprising fire if the two cross in the field. Reading the pathway as a circuit — a main valve at *Pun1*, a head-supply valve at pAMT, a chain-length dial in the branched-chain module — is precisely what lets us turn fire not just on and off but *up, down, and sideways*, which is the whole craft of designing a cultivar rather than merely finding one. Marker-assisted breeding tracks the locus directly, selecting seedlings by their DNA rather than waiting a season to taste the fruit. Our method takes the next step: instead of selecting which copies a plant happened to inherit, we write the copy we want, with the activity we specify, into the line — turning a trait that was once a coin-flip of inheritance into a quantity we set on purpose.

The whole route, in one table

Stand back from the pathway and it resolves into seven working modules. Five of them are the universal phenylpropanoid machinery every green thing already runs; the other two are the pepper's terminal additions, and of those only the confluence enzyme is truly unique to *Capsicum*.

STEP	MODULE	CONVERTS	→ PRODUCT	COFACTOR / COST	PEPPER-SPECIFIC?
1	PAL	L-phenylalanine	trans-cinnamic acid	MIO (self-made)	No
2	C4H	cinnamic acid	p-coumaric acid	NADPH + O_2	No
3	4CL	p-coumaric acid	p-coumaroyl-CoA	ATP → AMP + PPi	No
4	C3H · HCT · COMT	p-coumaroyl-CoA	feruloyl-CoA → vanillin	SAM (methyl)	No
5	pAMT	vanillin	vanillylamine (the head)	amino donor	Partly
6	BCAT · KAS · thioesterase	valine / leucine	8-methyl-branched acyl-CoA (the tail)	NADPH ×6, ATP ×3	Partly
7	**Pun1 / CS / AT3**	head + tail	**capsaicin**	thioester energy	**Yes**

To give a plant fire you do not rebuild it from scratch. You confirm its phenylpropanoid corridor is running — it is, in every green thing — and you install the terminal modules: a route to vanillylamine, a branched-chain tail supply, and above all a working capsaicin synthase, targeted to the right tissue at the right moment in development. Four terminal enzymes against a backdrop of machinery the plant already owns. It sounds almost modest. It is the difference between a tomato and a tomato that can register on the heat axis.

From content to a number — the arithmetic of heat

A pathway makes a molecule; a scale makes it legible. Two numbers govern how the heat this route produces becomes something you can read on a page, and both are worth working by hand.

The older number is the Scoville Heat Unit, and it has one honest definition: pure capsaicin scores **16,000,000 SHU**. Everything else is a fraction of that ceiling. So a pod's reading, divided by sixteen million, states its pungency as a fraction of pure capsaicin's:

```
pungency fraction = SHU ÷ 16,000,000.
```

Run the flagship ladder through it. **Phoenix** (SS-0096), with our model's predicted heat of approximately 2.94 million SHU, comes to 2,937,948 ÷ 16,000,000 = **0.184** — a design target of eighteen percent of the way to pure capsaicin. The ladder beneath it:

CULTIVAR · SKU	PREDICTED SHU	PUNGENCY FRACTION	IFI-H	HALF-ACTIVATION VS REFERENCE
Ember · SS-0087	866,000	5.4%	1,570	0.64×
Blaze · SS-0091	1,720,000	10.8%	2,630	0.38×
Firestorm · SS-0094	2,470,000	15.4%	3,490	0.29×
Phoenix · SS-0096	2,937,948	18.4%	4,200	0.24×

The second, truer number is the Iyengar Flavor Index, and it measures not how much capsaicin a food holds but how loudly that capsaicin speaks to the receptor. It is anchored to half-maximal activation:

$$\mathrm{IFI} = 1{,}000 \times \mathrm{EC}_{50}(\text{reference}) \div \mathrm{EC}_{50}(\text{candidate}),$$

where EC_{50} is the concentration that drives a receptor — here TRPV1, the capsaicin channel — to half its maximal response. Invert the formula and it hands you a potency directly. For Phoenix at IFI-H 4,200:

$$\mathrm{EC}_{50}(\text{reference}) \div \mathrm{EC}_{50}(\text{Phoenix}) = 4{,}200 \div 1{,}000 = 4.20,$$

so Phoenix's stimulus reaches half-activation at EC_{50}(Phoenix) = 1 ÷ 4.20 = **0.238** of the reference concentration — the final column of the ladder above. It is 4.2× as loud as the wild-type anchor. A sunlit shishito, near IFI-H 40, sits at the other end: EC_{50}(reference) ÷ EC_{50}(shishito) = 0.04, meaning the shishito needs 25× *more* concentration to reach the same half-activation. The span from shishito to Phoenix is therefore 0.238 to 25 — a range of about **105-fold** in how little of the food it takes to half-fill the channel.

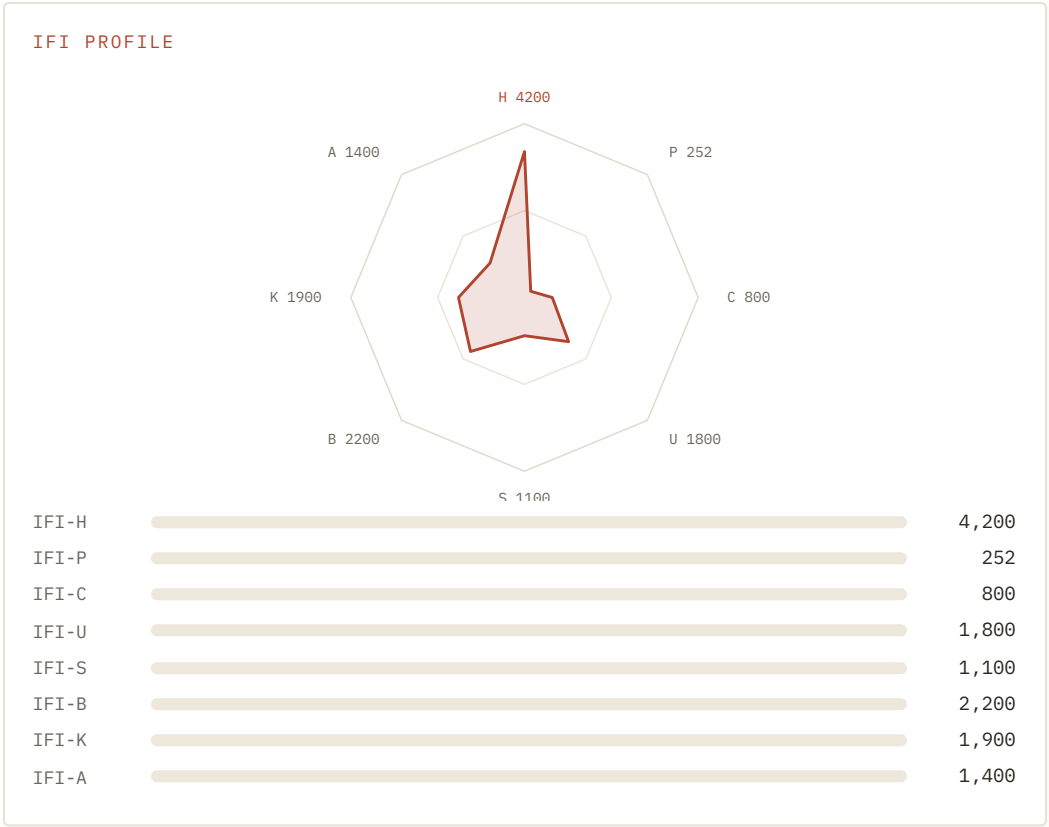

Now the subtlety that keeps the Index honest. A single receptor does not respond to concentration linearly; it responds along a sigmoid, captured by the Hill equation:

$$\text{response} \div E_{max} = [C]^n \div (EC_{50}{}^n + [C]^n),$$

with a Hill coefficient n near 2 for capsaicin at TRPV1 — the channel is mildly cooperative. At [C] = EC_{50} the response is, by definition, 1 ÷ (1 + 1) = 50%. But drive the same channel at 4.2× its EC_{50} — the concentration that Phoenix delivers where the reference would deliver one — and the response is

$$4.2^2 \div (1 + 4.2^2) = 17.64 \div 18.64 = \mathbf{94.6\%}.$$

The channel is already nearly saturated. This is why the ladder compresses at the top: fit the four flagship points as a power law and IFI-H rises with SHU to an exponent near **0.8**, not 1.0. Going from Ember to Phoenix roughly triples the raw pungency fraction (5.4% to 18.4%) but less than triples the Index score (1,570 to 4,200), because the receptor has less and less room left to respond. The compression is not a flaw in the Index; it is the biology the Index faithfully reports. Perceived heat saturates because TRPV1 saturates, and a good scale compresses exactly where the tongue does.

Finally, the consistency. The coefficient of variation ties the pathway's reliability to a single figure: CV = σ ÷ μ. For the engineered line, with its design target of CV ≤ 7% around a predicted center of μ ≈ 2,937,948 SHU, the batch-to-batch standard deviation our model projects is

$$\sigma \le 0.07 \times 2{,}937{,}948 \approx 205{,}700 \text{ SHU},$$

so about 95% of batches (±2σ) would fall within roughly 2.53M–3.35M SHU — a tight, dependable band a formulator can build a product around. Wild-type pods run CV ≈ 42%; the same arithmetic gives σ near 1.23M SHU, a spread that swings from near-mild to

searing on the same plant. The ratio 42 ÷ 7 = **6** is the whole story in one number: the engineered line is designed to be six times tighter in relative spread. And that tightness, recall, is the free energy of two substitutions cashed out as reproducibility — the −3.9 kcal/mol of fold stability, arriving at the formulator's bench as a number that does not wander.

What the map does not tell us

A map is not the territory, and we hold the honest edges in view. The metabolic budget above is textbook stoichiometry, not a measured flux — real yields depend on how much carbon the fruit commits and how much leaks to side routes, and we quote it as an order of magnitude, nothing finer. The power-law exponent near 0.8 is fit to four flagship points; four points can suggest a shape but cannot prove one, and we treat it as a working description of the compression, not a law. The Hill coefficient of 2 is a representative value for capsaicin at TRPV1, not a constant of nature; cooperativity shifts with temperature, pH, and the lipid around the channel.

Two deeper limits matter more. First, the single-axis clarity we enjoy here — heat as essentially one molecule reading one receptor — is a luxury of capsaicin. On a food that is a mixture, the EC_{50} in the Index formula is the aggregate half-activating concentration of everything present, and disentangling the contributors is harder than reading a near-pure pungency. Heat is the cleanest axis precisely because the pathway converges on one compound and one channel; that is why it is the axis everyone knows, and why it was the right place to prove the method. Second, and most of all, the pathway map tells us the enzymes but not the *grammar* — the fruit-specific, ripening-timed promoter that keeps the fire in the one tissue that should carry it. We can write the enzymes with confidence. Writing the discipline that governs them, in a plant that never carried fire, is the real frontier, and we would rather name it plainly than pretend the molecule is the whole problem.

These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

What the map does tell us is the thing that matters most. Fire was never a pepper problem. The head is a decorated phenylpropanoid any plant can make; the tail is a branched fatty acid any plant can make; five of the seven modules already run in every green thing alive. It was a *Pun1* problem — one acyltransferase, one amide bond, one tissue's worth of discipline about when to make it. And *Pun1* is a gene, and genes travel. To read this route carefully is to understand that the heat we revere was only ever waiting to be given to anything that grows.

Chapter 4 · The Mathematics of Burn — *SHU is a dilution; IFI-H is a receptor.*

In 1912 a pharmacist named Wilbur Scoville, working at the Parke-Davis company, needed to compare the pungency of pepper extracts for a muscle rub. He had no instrument for it — no way to see a molecule, let alone count one. So he built an assay out of the only instrument available: people. He dissolved a measured weight of dried pepper in alcohol, then diluted the extract in sugar water, step by step, and had a panel of five tasters sip each dilution until a majority could no longer detect any heat. The dilution factor at that vanishing point became the score. A pepper still detectable after being thinned sixteen thousand-fold scored sixteen thousand Scoville Heat Units. The method was subjective, panel-dependent, and wildly variable — tasters fatigue, palates differ, the fifth sample is never judged like the first — and it was, for all that, a genuine act of genius, because it put a number on a feeling for the first time in history. We owe Scoville the number. We could not keep his method, and the reason we could not is the whole subject of this chapter.

A number made of dilution

Today pungency is read not by palate but by high-performance liquid chromatography. A fruit is dried, milled, and extracted; the extract is pushed through a column that separates the capsaicinoids by how long each lingers against the stationary phase; a detector counts them as they come off, one peak at a time, down to the part per million. The Scoville value is then *reconstructed* rather than tasted — each capsaicinoid multiplied by its own pungency weight, a factor pinned long ago against the old sensory panels, and the weighted parts summed. It is the sensory experiment of 1912 frozen into a coefficient and carried out in arithmetic instead of on the tongue.

Walk it through once, concretely, because the math is simple and it clarifies everything that follows. Suppose a run on a dried fruit returns 2,500 ppm capsaicin and 1,500 ppm dihydrocapsaicin, with only traces of the rest. The two principal capsaicinoids each carry a pungency weight near 16.1 Scoville units per part per million, so the reconstructed heat is

$$(2{,}500 + 1{,}500)\ \text{ppm} \times 16.1\ \text{SHU}\cdot\text{ppm}^{-1} \approx \mathbf{64{,}400\ SHU}$$

— a clean, hot cayenne. The weights are not all equal, and the differences are the first quiet clue to everything the Scoville number cannot see:

CAPSAICINOID	TYPICAL SHARE OF THE POD	ACYL TAIL	PUNGENCY WEIGHT (SHU·PPM⁻¹)
Capsaicin	~48%	8-methyl-6-nonenoic (unsaturated)	16.1
Dihydrocapsaicin	~42%	8-methylnonanoic (saturated)	16.1
Nordihydrocapsaicin	~7%	one carbon shorter	9.3
Homocapsaicin	~2%	one carbon longer	8.6
Homodihydrocapsaicin	~1%	one longer, saturated	8.6

Those weights are ratios to the two principals, and they tell us plainly that the five molecules do not burn alike: trade a part per million of capsaicin for a part per million of homocapsaicin and the Scoville sum falls by nearly half for that part, because the longer, more saturated homo-compound is simply less pungent per molecule. Hold that fact; the receptor will make it matter. Pure capsaicin, being a million parts per million, works out to about 16.1 million SHU by the same arithmetic, which is why the scale climbs into the millions at all. The older ASTA pungency units used across the spice trade are the identical measurement in different dress — a chromatographic content score, convertible to SHU by a fixed factor near fifteen.

Notice what every one of these units actually measures: *concentration*. Total capsaicinoid content, weighted and added. It is a census of molecules in a vial, and it would read the same whether those molecules ever met a tongue or not. Run the arithmetic backward on the top of our own ladder and the point lands hard. Our model predicts Phoenix (SS-0096) at approximately 2.95 million SHU; divide by 16.1 and the number implies on the order of 182,000 ppm of capsaicinoid-equivalent — roughly eighteen percent of the dried placental tissue by weight, the pale membrane where the fruit builds and stores its fire. That is an astonishing burden of molecule, grown on purpose into the one tissue evolved to hold it. But eighteen percent of a dried membrane is a statement about *content*, and the tongue has never once tasted content. It tastes what a molecule does when it reaches a receptor. The gap between those two things — how much is present, and how loudly it speaks — is the reason the Iyengar Flavor Index exists.

The channel that reports a fire

Capsaicin binds TRPV1, a non-selective cation channel standing in the membranes of the trigeminal nerve endings that wrap the surface of the tongue. It is a tetramer — four identical subunits arranged around a central pore, each contributing six membrane-spanning helices, with the gate formed where the inner helices cross below the pore. In its resting life TRPV1 is a thermometer: it opens to temperatures above roughly 43 °C and to the fall of pH that marks injured tissue, and the ion rush it permits is the body's honest report that something is hot enough, or acid enough, to do harm.

Capsaicin does not heat the tissue. It is subtler and stranger than that. The molecule slips into a hydrophobic vanilloid pocket cradled between the third and fourth helices of a subunit and the short linker that leads to the gate. Its vanillyl head-group — the same catechol-and-methoxy ring the plant builds from the phenylpropanoid pathway — settles by hydrogen bonds against the pocket's polar

residues, its amide waist pivots, and its branched fatty tail lies down into the grease of the membrane-facing cleft. That binding tugs on the linker, the inner helices splay, the gate dilates, and the pore opens at body temperature — a full six or more degrees below where heat alone would have opened it. The channel reports a fire that is not there. Calcium and sodium pour through, the trigeminal neuron depolarizes, a train of action potentials climbs toward the brain, and the brain, reading its one reliable wire, calls it burning. There is no burn. There is only a molecule that has lowered a threshold, and a nervous system doing exactly what it was built to do with the signal.

That inrush of calcium is not only the sensation. It is also, precisely, how we measure. We express human TRPV1 in cultured cells, load them with a fluorescent dye whose brightness tracks intracellular calcium, and watch the glow rise as we deliver the candidate extract. The height of that glow against dose is the dose–response curve; its midpoint is the number we build everything on. We measure the human channel, never a rodent proxy, because the vanilloid pocket differs enough between species to move the answer, and the Index is only honest if the receptor in the dish is the receptor in the eater.



Phoenix (SS-0096) on the microscope stage — the calcium glow of a TRPV1 assay is both the nerve's signal and our ruler. We read the pepper not by how much capsaicin it holds but by how loudly that capsaicin opens the one channel the tongue listens to.

From half-activation to a score

The quantity we read off that curve is the **EC_{50}** — the concentration of the candidate's actual capsaicinoid mixture required to drive TRPV1 to half of its maximal calcium response. Response to capsaicin does not climb in a straight line with dose; it follows the sigmoid of the Hill equation,

$$\theta = [L]^n \div (EC_{50}^n + [L]^n)$$

where θ is the fraction of full activation, [L] the delivered concentration, and *n* the Hill coefficient — the cooperativity of the four-subunit channel, which for capsaicin sits near 2. At [L] = EC_{50}

the fraction θ is exactly one-half, by construction, which is why the half-maximal point is the stable, reproducible landmark on an otherwise steep and slippery curve. Fit the curve, read the midpoint, and you have one number that captures the molecule's *loudness* to the receptor, independent of how much of it you happened to pour.

The Index turns that midpoint into a score against a fixed biological reference:

> **IFI-H = 1,000 × EC_{50}(reference) ÷ EC_{50}(candidate)**

The reference is our native wild-type benchmark, pinned so that its own half-activation scores exactly **1,000**. Because EC_{50} sits in the denominator, a *lower* EC_{50} — a louder molecule that half-opens the channel at a smaller dose — yields a *higher* score. That is the correct direction: loudness up, number up.

Carry the reference at its calibrated half-activation, a capsaicinoid-equivalent EC_{50} of 1.00 µM, and every other EC_{50} in this chapter follows by rearranging the formula — EC_{50}(candidate) = 10^6 ÷ IFI-H, in nanomolar. The whole ladder of flagship designs then resolves into the predicted dose each one is designed to need to reach the same half-open channel:

CULTIVAR	SKU	SHU (PREDICTED)	IFI-H (PREDICTED)	EC_{50} (NM)	× LOUDER THAN REF
Ember	SS-0087	866K	1,570	637	1.57×
Cinder	SS-0088	1.05M	2,290	437	2.29×
Kindle	SS-0089	1.25M	2,000	500	2.00×
Flare	SS-0090	1.48M	—	—	–
Blaze	SS-0091	1.72M	2,630	380	2.63×
Wildfire	SS-0092	1.96M	2,930	341	2.93×
Scorch	SS-0093	2.21M	3,220	311	3.22×
Firestorm	SS-0094	2.47M	3,490	287	3.49×
Solara	SS-0095	2.72M	3,980	251	3.98×
Phoenix	**SS-0096**	**2.95M**	**4,200**	**238**	**4.20×**

All SHU values and IFI-H scores for Scoville Splice cultivars are computational predictions from our platform; they are not certified measurements.

There is an honesty in the blank line: Flare has a predicted Scoville value and no IFI-H, because its receptor model has not been run to our confidence threshold, and we would rather carry a dash than a guess. The rest of the column follows from our model. Our model predicts Phoenix opens the channel at approximately 238 nanomolar. Convert that to the concentration a receptor actually sees — multiply by capsaicin's molar mass of 305.4 g·mol^{-1} — and it is about 73 micrograms per liter, roughly **73 parts per billion**. The most intensely designed pepper our family has ever built, by our model's account, speaks to its receptor at seventy parts in a billion. A sunlit shishito, scoring near 40 on the same axis, needs an EC_{50}

around 25 µM — some 7.6 parts per million, more than a hundredfold more molecule to reach the same half-open gate. That hundredfold spread across the edible world is exactly what the formula is built to hold.



Phoenix (SS-0096) — our model predicts approximately 2.95 million SHU, designed IFI-H 4,200. Both numbers describe the same pod. One counts the molecules it holds. The other measures how loudly those molecules are designed to speak to the one receptor that can ever feel them.

And because Phoenix is the one cultivar we tuned to register on every axis at once, its heat is only the first line of a fuller voice:

IFI PROFILE

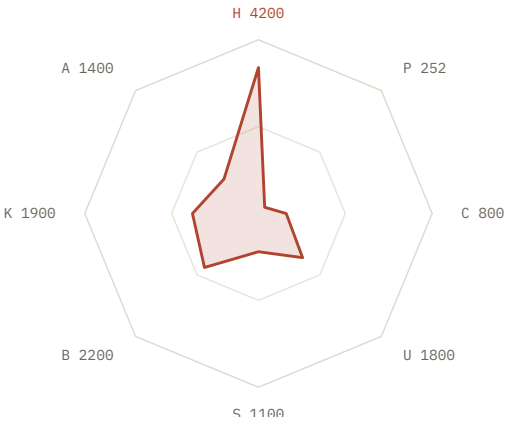

IFI-H	4,200
IFI-P	252
IFI-C	800
IFI-U	1,800
IFI-S	1,100
IFI-B	2,200
IFI-K	1,900
IFI-A	1,400

Each axis score is computed against its own stated convention constant; values on different axes are not directly comparable to one another.

Where the two scales part

SHU and IFI-H rise together, yet never in lockstep, and the Hill curve is the reason. SHU counts content, and content has no ceiling — pour in twice the capsaicinoid and the Scoville sum doubles, forever,

as though the tongue could double with it. The receptor cannot. It fills. Work the saturation out in numbers, with $n = 2$ and dose written as a multiple x of the EC_{50}, so $\theta = x^2 \div (1 + x^2)$:

- at $x = 1$ (a dose equal to EC_{50}): $\theta = 1 \div 2 =$ **0.500**
- at $x = 2$: $\theta = 4 \div 5 =$ **0.800**
- at $x = 10$: $\theta = 100 \div 101 =$ **0.990**
- at $x = 20$: $\theta = 400 \div 401 =$ **0.9975**

Now read the two doublings. Near the bottom, doubling the dose from one EC_{50} to two drives the channel from half-open to four-fifths open — a gain of **0.300**, thirty full points of activation bought with one doubling of molecule. Near the top, doubling from ten EC_{50} to twenty moves it from 0.990 to 0.9975 — a gain of **0.0075**. The *same* doubling of molecules buys forty times less sensation at the top of the ladder than at the bottom. Take the slope rather than the step — the marginal response $d\theta/dx = 2x \div (1 + x^2)^2$ — and the collapse is starker still: at $x = 1$ the channel answers each added unit of dose with 0.50; at $x = 10$ with 0.00196; the receptor near saturation is some **255 times** less moved, molecule for molecule, than the receptor near its midpoint.

This is why the Scoville scale overstates the experience exactly where it climbs most proudly. You can watch the compression inside our own designed ladder, in Scovilles spent per point of IFI-H earned between neighboring rungs:

- Ember → Cinder: ΔSHU ≈ 184,000 over ΔIFI 720 → **256 SHU per IFI-H point**
- Blaze → Wildfire: ΔSHU ≈ 240,000 over ΔIFI 300 → **800 SHU per IFI-H point**
- Solara → Phoenix: ΔSHU ≈ 218,000 over ΔIFI 220 → **991 SHU per IFI-H point**

The toll nearly quadruples as you climb. At the foot of the ladder a quarter-million Scovilles buys a whole clutch of perceptible steps; near the summit the same quarter-million barely buys one. A two-million-Scoville pepper and a three-million-Scoville pepper stand a full million units apart on paper and only a few hundred IFI-H points apart across a tongue — not an error in either scale, but the real distance between what is in the food and what is in the eater. For scorecards we cap the axis at **5,000**, and the cap is a matter of principle, not convenience: past the point where the receptor is saturated, further "heat" is a chemical fact with no perceptual home, and we will not draw a scale finer than the mouth can read. The number ends where the sensation ends.

The enzyme that writes the profile

Everything the receptor reads is first written by an enzyme, and the enzyme is where SHU and IFI-H can be made to actively disagree. The vanillyl head of every capsaicinoid arrives down the phenylpropanoid branch that runs in every vascular plant — PAL → C4H → 4CL → HCT → COMT — which turns phenylalanine into vanillin and, by a dedicated aminotransferase (pAMT), into vanillylamine. The fatty tail arrives from the other side: branched-chain amino acids, handed off by BCAT and elongated by the

ketoacyl-ACP synthase (KAS), become a short, methyl-branched acyl-CoA. At the terminal step the capsaicin synthase Pun1/CS — an AT3-family acyltransferase — condenses the amine with the acyl donor into the finished capsaicinoid.

The *shape* of that acyl donor decides which capsaicinoid you get. An unsaturated nine-carbon tail gives capsaicin; the saturated version gives dihydrocapsaicin; a carbon shorter gives nordihydrocapsaicin; a carbon longer gives the homo-compounds. The enzyme's substrate selectivity, in other words, sets the *profile* — and the profile is precisely what the pungency-weight table and the receptor weigh so differently. Here lies the single most instructive line in the whole heat ladder. Kindle, at a predicted 1.25 million Scovilles, scores *lower* on IFI-H — 2,000 — than Cinder beneath it at 1.05 million, which reaches 2,290. By predicted content it is hotter. By receptor it is not. The reason is written in the capsaicinoid profile:

CAPSAICINOID	SHU WEIGHT	CINDER · SS-0088 (PREDICTED, PPM)	→ PREDICTED SHU	KINDLE · SS-0089 (PREDICTED, PPM)	→ PREDICTED SHU
Capsaicin	16.1	40,000	644,000	30,000	483,000
Dihydrocapsaicin	16.1	24,000	386,400	20,000	322,000
Nordihydrocapsaicin	9.3	1,500	13,950	12,000	111,600
Homocapsaicin	8.6	400	3,440	21,000	180,600
Homodihydrocapsaicin	8.6	300	2,580	18,000	154,800
Total		**66,200 ppm**	**≈ 1.05M**	**101,000 ppm**	**≈ 1.25M**

Capsaicinoid concentrations and SHU values for Cinder and Kindle are computational predictions from our pathway model.

Kindle carries half again as much total predicted capsaicinoid mass as Cinder — 101,000 ppm against 66,200 — and so it wins on predicted Scovilles. But only half of that mass (the capsaicin and dihydrocapsaicin, 50,000 of 101,000 ppm, near 50%) is the receptor-efficient principal pair; the rest is the longer, weaker nor- and homo-compounds. In Cinder the principals are 64,000 of 66,200 ppm — nearly 97%. The homo-compounds are poorer TRPV1 agonists: they bind the vanilloid pocket less readily and lift the channel less per molecule, so a mixture diluted with them carries a *higher* effective EC_{50}. That is exactly what our model predicts — EC_{50} approximately 500 nM for Kindle against 437 nM for Cinder — and it drops Kindle's IFI-H below its lighter neighbor's. More predicted mass, less loudness. SHU, blind to profile, counts the mass and ranks it higher. IFI-H, reading the channel, ranks it where the tongue actually puts it. Neither number is wrong; they answer different questions, and only one of them is the question a cook is asking.

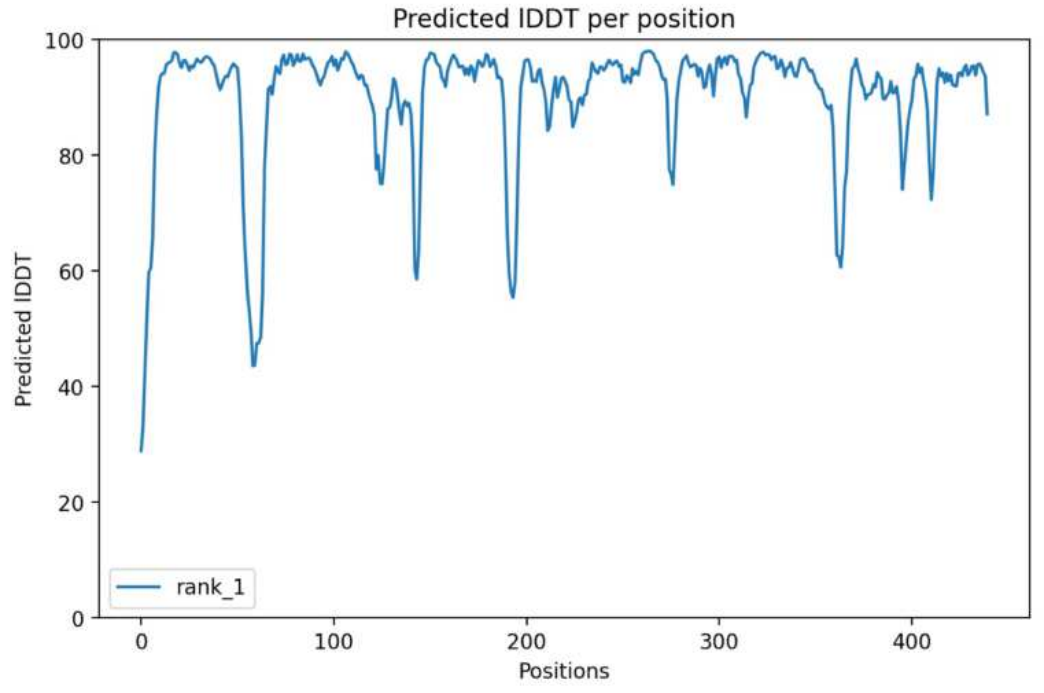

FIG. 7 — the Pun1/AT3 capsaicin synthase, modeled and colored by pLDDT confidence; the binding domain and active-site cleft resolve sharply. This is the enzyme whose substrate selectivity writes the profile the receptor then weighs — the hinge between how much heat a plant makes and how loud that heat turns out to be.

Because the enzyme writes the profile, the enzyme is where we designed. Rather than selecting phenotype and hoping the genotype followed, we scored every single-residue substitution in Pun1 by predicted effect on catalytic efficiency and fold stability. Our genome-scale flux model predicts that control in the capsaicinoid pathway sits upstream — at 4CL and HCT/CCoAOMT in the phenylpropanoid pathway — not at the terminal Pun1/AT3 step; Pun1 is the final capsaicin synthase and labels the product, but it does not set the rate. Working from that model, we designed two substitutions for the final construct that no breeding program had ever combined:

EDIT	WHERE / WHAT IT DOES	ΔΔG (KCAL·MOL⁻¹)	WHAT IT BUYS, BY EXP(−ΔΔG/RT)
S39L	N-terminus; tightens substrate binding	−2.1	≈ 35× lower Kd (tighter capture)
L345G	active-site cleft; steric relief	−1.8	≈ 21× faster turnover (kcat)
Sum (shipped construct)	**double variant S39L+L345G**	**−3.9**	affinity + rate
C175S	*removes a mis-pairing cysteine — structural annotation only, not included in shipped construct*	*−0.9*	*≈ 4.6× fewer predicted misfolds*

The right-hand column is the Boltzmann relation applied to each designed ΔΔG, at RT = 0.592 kcal·mol^{-1} for 25 °C: a −2.1 kcal·mol^{-1} binding gain is exp(2.1 ÷ 0.592) ≈ 35-fold tighter substrate capture; a −1.8 kcal·mol^{-1} relief of strain in the transition state is exp(1.8 ÷ 0.592) ≈ 21-fold faster turnover. The two shipped edits together lift both affinity and rate — and the next section turns to what that consistency buys for batch reproducibility.

The width of the number

A single flagship SHU value is a point estimate, and a point without a width is a kind of dishonesty. The double-variant design — S39L tightening substrate capture and L345G relieving transition-state strain — produces a more catalytically consistent enzyme than the

wild-type Pun1, and a consistent enzyme makes a consistent pepper. The whole program is designed to a batch-to-batch coefficient of variation of **≤ 7%**, against **42% CV** in the wild-type lines the trait was built from.

Put that in the units of the fruit. Phoenix (our model predicts approximately 2.95 million SHU) with a 7% design CV carries a one-sigma band of ±0.07 × 2,950,000 ≈ **±206,000 SHU** — batches designed to land, reliably, between about 2.74 and 3.14 million. A wild-type superhot at the same mean and 42% CV would swing ±1.23 million, a band six times wider, enough that no formulator could ever dose a sauce to a target and trust the next harvest to meet it. Seven percent is a design specification. Forty-two percent is a hope. The difference between them is the difference between a trait you can sell and a trait you can only admire.

The receptor score carries its own predicted uncertainty, and it propagates the way a ratio must. Since IFI-H = 1,000 × EC_{50}(ref) ÷ EC_{50}(candidate), the relative variances add in quadrature: with a candidate assay design CV near 7% and a reference averaged down to about 2% over many runs, the fractional error on IFI-H is $\sqrt{(0.07^2 + 0.02^2)} \approx 7.3\%$. We pull that down the ordinary way, with replicate wells — the standard error of the mean falls as the square root of their count, SEM = CV ÷ $\sqrt{n}$, so eight wells turn a 7% plate CV into 7% ÷ $\sqrt{8}$ ≈ **2.5%**. On Phoenix's predicted IFI-H of 4,200 that is a reported band of roughly ±105. Every predicted value we publish is designed with a width we can quote, which is the only kind of number worth publishing.

And then the honest edges, because a scalar as clean as IFI-H earns its trust by naming what it leaves out. TRPV1 desensitizes: hold capsaicin on the channel and calcium-dependent feedback closes it, the tachyphylaxis a chili-eater feels as the second bite landing softer than the first, so every EC_{50} must be read from naive cells, one dose to a well, or the midpoint drifts low. The channel is thermosensitive by its nature, so assay temperature is fixed and controlled, or the curve shifts beneath you. Capsaicin is stubbornly hydrophobic, so the free concentration — the part the receptor actually sees — depends on the delivery vehicle, and the dose must be corrected to it. Deeper than any of these: the number measures loudness to one receptor in a dish, and a mouthful of superhot is not one receptor in a dish. It has a time-course — the slow build, the plateau, the long cruel tail — that a midpoint cannot hold; it recruits trigeminal and gustatory channels together; and it meets a tongue whose own TRPV1 varies from person to person, so that the same pepper is honestly a different fire in different mouths. IFI-H does not pretend to the whole of that experience. It holds the part that is stable, measurable, and designable — the loudness of the molecule at the gate — and it leaves the rest to the eater, where it belongs.

These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

That restraint is the point, and it is of a piece with how our family has always wanted to grow things: to measure a pepper exactly, to write its number with its width, to let heat be a pleasure given to the body with intention rather than an assault to be survived, and to revere the

plant enough to say only what the receptor will confirm. SHU answers the question a chemist asks of a vial — *how much is in it.* IFI-H answers the question a cook, and a tongue, and a body actually ask. *Not how much is in it. How much will I feel.*



PART II OF V

The Ladder

six bands from a sunlit glow to the edge · the ten we built · designing Phoenix

Chapter 5 · The Heat Ladder — *six bands from glow to edge.*

Everything a kitchen knows how to do with heat comes to rest on a single structure: the heat ladder, six bands that carve the enormous span the receptor can feel into regions a cook can reason about. A continuous scale — a shishito's sunlit 40 climbing in unbroken degrees to Phoenix at 4,200 — is accurate and almost useless at the stove, because nobody cooks to three significant figures. What a kitchen needs is a vocabulary: a handful of named regions, each with a character, a set of reference peppers, and a role in a meal, so that *I want this to warm, not to wound* becomes a band, and a band becomes a dose. The ladder is, before anything else, a translation — out of chemistry's language, in which heat runs to the millions, and into the mouth's, in which it runs from a glow to an edge in six honest steps.



Pods laid out like the rungs they name — the ladder made physical. Each step up is roughly a doubling of what the receptor feels, and the whole climb, from the first warmth to the last, spans barely a hundredfold on the tongue even as the chemistry beneath it runs to sixty thousand.

We draw the lines where perception draws them, not where the arithmetic lands on round numbers. The bands are roughly geometric — each a doubling or so of felt intensity over the one below — because that is how a saturating receptor partitions the world: equal *ratios* feel like equal *steps*. To understand the ladder at all is to understand why its rungs are spaced the way they are, so that is where we begin.

Why the rungs are geometric

The law beneath the ladder is two centuries old. Weber and Fechner found that perceived intensity grows not with a stimulus but with its *logarithm* — that to double a sensation you must roughly multiply its cause, not merely add to it. Hearing obeys this, which is why loudness is counted in decibels; brightness obeys it; and heat, read through a receptor that saturates, obeys it most of all. A ladder with linear rungs — five hundred Scovilles apart, say — would crush a dozen meaningful steps into the mild end and then stretch a single perceptual step across the entire top of the scale, useless at both extremes. Drawn geometrically instead, so that each rung is about a doubling of felt heat, the steps feel evenly spaced to the tongue even as they stretch further and further apart on the chemical axis.

The receptor is the reason, and it is worth meeting plainly, because the whole ladder stands on it. Capsaicin binds TRPV1, a non-selective cation channel on the trigeminal nerve endings that lace the mouth — the same channel opened by temperatures above 43 °C and by acid. The molecule warms nothing; it lowers the channel's thermal threshold until the mouth's own heat is enough to open it, and calcium floods inward, and the nerve reports a fire that is not there. Crucially, the channel's response to concentration is not a straight line but a sigmoid — a Hill curve, cooperative, the four subunits of the channel swinging open together rather than one at a time. Near the bottom of that curve a little more capsaicin buys a great deal more signal; near the top the curve flattens, and each additional molecule buys less and less felt heat. The ladder's geometry is simply this sigmoid read backward: to climb an equal perceptual step as you ascend, you must supply an ever-larger multiple of capsaicinoid, because the receptor is already tiring of the message even as the chemistry shouts louder.

That is why the rungs widen on the chemical axis as they rise yet stay even on the felt one. Take the midpoint of each band — the geometric center of its IFI-H range — and measure the step from one center to the next:

STEP BETWEEN BAND CENTERS	IFI-H CENTERS	RATIO
Warmth → Glow	40 → 141	3.54×
Glow → Burn	141 → 371	2.62×
Burn → Fire	371 → 778	2.10×
Fire → Inferno	778 → 1,591	2.04×
Inferno → Edge	1,591 → 3,108	1.95×

From the Burn upward the ladder settles into near-exact doublings — 2.10, then 2.04, then 1.95 — a uniform perceptual staircase. The two mild bands climb in slightly larger strides, because the very floor of the scale is anchored to the threshold of perception itself rather than to a clean ratio: below the Warmth there is no heat to double, only the first suggestion of it. When a cook moves "up one band," the increment feels the same whether the move is from the Warmth to the Glow or from the Inferno to the Edge — a uniform staircase for the tongue laid over a wildly non-uniform one for the molecule. That uniformity is the ladder's entire reason to exist.

The two columns, and the gap between them

Read the ladder honestly and it carries two numbers for every rung, measuring two different things the kitchen has always confused. One is the chemistry; the other is the loudness.

The chemistry is the Scoville heat unit, and it is a measure of *mass* — of how much capsaicinoid a fruit contains. The modern laboratory reads it by separating the capsaicinoids on a chromatography column, weighting each by its known pungency, and converting at roughly **16 Scoville units per part-per-million** of capsaicin-equivalent — which places pure crystalline capsaicin near **16,000,000 SHU**. The Scoville number is scrupulous about molecules and silent about the receptor. It counts what is present; it cannot tell you how loudly the present speaks.

The loudness is IFI-H, and its definition is one line:

> **IFI-H = 1,000 × EC_{50}(reference) ÷ EC_{50}(candidate)**

EC_{50} is the concentration at which a stimulus drives TRPV1 to half of its full response — the honest physical measure of how loudly a molecule speaks. Invert it against a fixed wild-type reference and scale by a thousand, so that the reference itself scores exactly 1,000 and everything louder climbs above it. The reference is a stated convention — fixed to define IFI-H = 1,000 on the heat axis, the way the Scoville scale is defined by dilution to threshold — not a consensus value drawn from a single published source. Because the score is nothing but an inverted concentration ratio, a candidate's relative potency falls straight back out of it: **rel EC_{50} = 1,000 ÷ IFI-H**, with the reference pinned at 1.000, so a smaller figure means a stimulus that reaches half-activation at a lower dose. Here is the full ladder with both columns, and the peppers that give the numbers a taste:

BAND	NAME	SHU RANGE	IFI-H RANGE	REFERENCE PEPPERS	ROLE IN A DISH
0	The Warmth	0–2,500	20–80	shishito, padrón, poblano	Hospitality; heat as texture, not event
1	The Glow	2,500–25,000	80–250	jalapeño, Fresno, serrano	Lift; brightens without dominating
2	The Burn	25,000–100,000	250–550	cayenne, tabasco, Thai bird's-eye	Spine; the dish's defining edge
3	The Fire	100,000–500,000	550–1,100	habanero, scotch bonnet	Architecture; heat as a main voice
4	The Inferno	500,000–1,500,000	1,100–2,300	ghost, Trinidad scorpion	Spectacle; dosed, never served neat
5	The Edge	1,500,000–3,000,000	2,300–4,200	reaper, Pepper X, **Phoenix**	The frontier; grams become drops

Now put the two columns against each other and the compression becomes arithmetic. A shishito at the bottom reads IFI-H 40 — a relative EC_{50} of 1,000 ÷ 40 = **25.0**, a stimulus that needs twenty-five times the reference dose to half-open the channel. Phoenix at the top is predicted to score 4,200 in our model — a relative EC_{50} of 1,000 ÷ 4,200 = **0.238**, a stimulus 4.2× more potent than the reference. Divide the one loudness by the other and the reference cancels:

$$\frac{\text{loudness(Phoenix)}}{\text{loudness(shishito)}} = \frac{4{,}200}{40} = 105$$

The entire ladder spans **105×** in felt loudness. Yet in chemistry it spans the shishito's measured ~50 SHU to Phoenix's predicted ~2,937,948 — a factor of nearly **59,000**. The same climb is a hundredfold on the tongue and sixty-thousandfold in the fruit; the ratio between those two, about **560×**, is the receptor's saturation written as a single number. This is the gap every cook has felt and no one could name: that a pepper a thousand times "hotter" by the laboratory's count is nothing like a thousand times hotter in the mouth.

How fast the chemistry outruns the loudness is not even constant along the ladder, and it repays careful calculation. Treat the two columns as a transfer function and read its local slope in logarithmic coordinates — how many times IFI-H multiplies for each multiplication of SHU — across each band:

ACROSS THE BAND	SHU CLIMBS	IFI-H CLIMBS	LOCAL EXPONENT
the Glow	2,500 → 25,000 (×10)	80 → 250 (×3.1)	0.49
the Burn	25,000 → 100,000 (×4)	250 → 550 (×2.2)	0.57
the Fire	100,000 → 500,000 (×5)	550 → 1,100 (×2.0)	0.43
the Inferno	500,000 → 1,500,000 (×3)	1,100 → 2,300 (×2.1)	0.67
the Edge	1,500,000 → 3,000,000 (×2)	2,300 → 4,200 (×1.8)	0.87

The exponent wanders between **0.43 and 0.87** and averages near **0.6** — meaning that, very roughly, felt loudness climbs as something close to the square root of the capsaicinoid mass. But the honest reading is that there is no single exponent. The relationship is most compressed through the broad middle of the ladder, in the Fire, where fivefold more chemistry buys only a doubling of heat; it steepens toward the Edge, where the superhots climb almost in step with their own Scoville counts, because so much of the receptor's dynamic range is already spent that each further gain is paid for in whole molecules again. A single power law is a convenience we will not oversell; the true curve is the receptor's sigmoid — gentle at the extremes, steep in the middle — and these exponents are its slope sampled band by band.

The six bands, rung by rung

Each band is a shelf with a character, not merely a slice of the scale, and the quickest way to learn them is by the peppers and the kitchens that already live on each.

Band 0 · The Warmth (IFI-H 20–80). Shishito, padrón, poblano — hospitality, heat as texture rather than event. This is the resting warmth of the Mediterranean and much of Europe: the pinch of Aleppo, the ribbon of Espelette, heat as seasoning and not statement. A pepper here speaks on the heat axis alone, and softly:

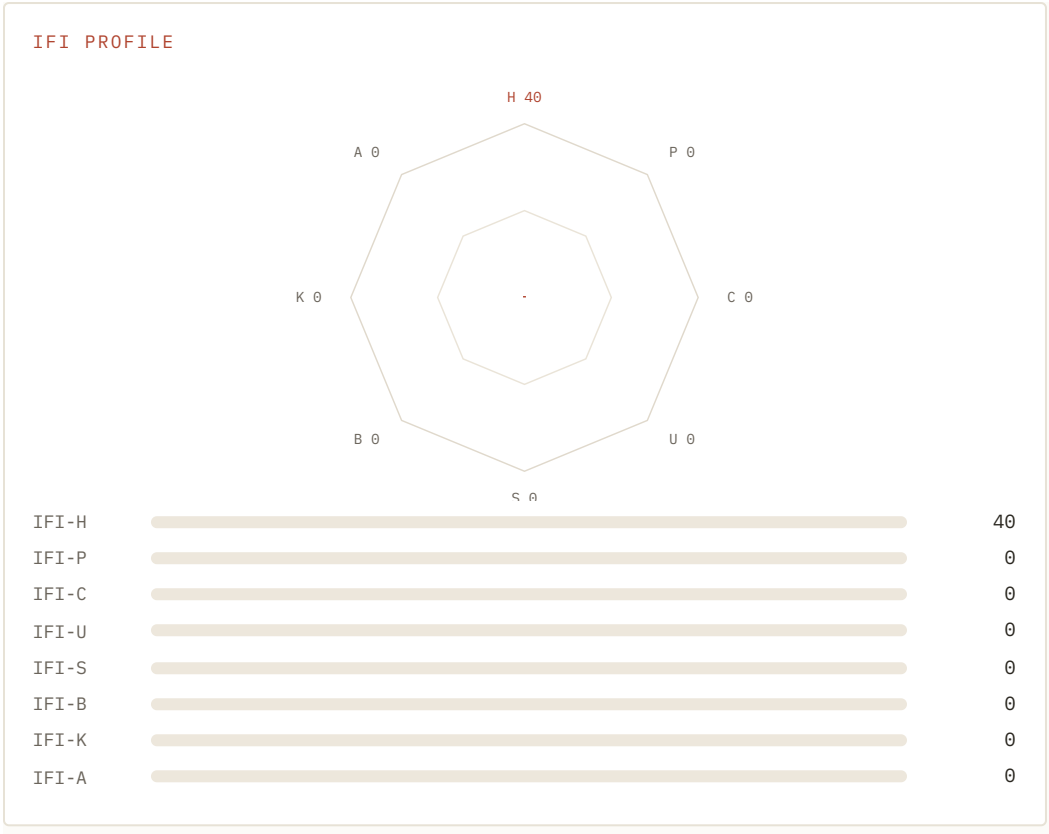

Read that profile as a picture of an ordinary market pepper: one line lit, barely, and the other seven dark — nothing engineered into pungency, cooling, umami, sweetness, bitterness, kokumi, or astringency, because nothing was. The Warmth is heat as hospitality, a background glow you could cook a child's supper inside.

Band 1 · The Glow (80–250). Jalapeño, Fresno, serrano — the lift that brightens a dish without dominating it, the first band where heat announces itself as a deliberate flavor rather than a trace. It is the warmth of a thousand weeknight tables, present enough to notice and gentle enough to feed anyone.

Band 2 · The Burn (250–550). Cayenne, tabasco, Thai bird's-eye — the spine of a dish, the heat that gives it a defining edge. This is the working register of a great swath of the world: the cayenne of the American South, the bird's-eye of Southeast Asia, the árbol of a Mexican table salsa. When a cuisine is known for being spicy, more often than not it lives here.

Band 3 · The Fire (550–1,100). Habanero, scotch bonnet — architecture, heat as a main voice. Where the Caribbean and the Yucatán live, scotch bonnet and habanero woven into daily cooking so thoroughly that the fruit's fruity perfume becomes part of the cuisine's signature. Heat at this rung is not an accent you add; it is a wall you build the room around.

Band 4 · The Inferno (1,100–2,300). Ghost, Trinidad scorpion — spectacle, dosed and never served neat. A pepper here is no longer an ingredient you reach for by the spoonful; it is one you meter out, because the capsaicinoid is now concentrated in the fruit's pale interior membrane densely enough that the flesh around it is almost beside the point.



A superhot cut open — the pallid placental tissue cradling the seeds is where the plant builds and hoards its capsaicin, not in the wall you think of as the pepper. By the Inferno, nearly all the heat lives in this seam.

Band 5 · The Edge (2,300–4,200). Reaper, Pepper X, and Phoenix — the frontier, where grams become drops. The Inferno and the Edge belong to almost no traditional cuisine at all. They are a modern invention, the product of a half-century of competitive breeding and now of deliberate design — a country that did not exist to be colonized until very recently, and that a cook enters by the gram and leaves by the fraction of one.

When we hand a cook the ladder, then, we hand them a map on which every rung already has a kitchen standing on it, and the freedom to move between those kitchens by the numbers rather than by luck.

The Edge, designed

The top rung deserves its own close reading, because it is the one we built from nothing and the one where the ladder's discipline matters most. Our ten proof-of-concept superhots were designed to walk the upper ladder in deliberate steps, from a first real warning to the single point where we stopped. Laid out together, they are both a product ladder and a lesson in the difference between the two numbers we have been circling.



Phoenix (SS-0096) on the microscope stage — the cultivar designed to anchor the top of the ladder, with predicted potency scored by what our model calculates it would do to the receptor. The Edge is the one band that becomes reachable when a cultivar is precisely designed and held within a few percent by computational batch-control specification.

In the table that follows, the relative EC_{50} is computed straight from the definition — 1,000 ÷ IFI-H, with the reference set to 1.000 — and the final column converts the predicted Scoville count into an actual capsaicinoid mass fraction, at 16 SHU per part-per-million, so the chemistry is legible beside the loudness:

CULTIVAR	SKU	SHU†	IFI-H	REL EC_{50}	CAPSAICINOIDS†
Ember	SS-0087	866,000	1,570	0.637	5.4%
Blaze	SS-0091	1,720,000	2,630	0.380	10.8%
Wildfire	SS-0092	1,960,000	2,930	0.341	12.3%
Scorch	SS-0093	2,210,000	3,220	0.311	13.8%
Firestorm	SS-0094	2,470,000	3,490	0.287	15.4%
Solara	SS-0095	2,720,000	3,980	0.251	17.0%
Phoenix	**SS-0096**	**2,937,948**	**4,200**	**0.238**	**18.4%**

† SHU and capsaicinoid values are computational predictions from our flux model.

Walk the two numeric columns down the table and the compression of the Edge is visible in the gearing between them. From Ember to Phoenix the predicted capsaicinoid mass climbs 866,000 → 2,937,948 SHU, a factor of **3.39×**; the loudness climbs only 1,570 → 4,200, a factor of **2.68×**. More than triple the chemistry for not quite triple the heat — the local exponent here is ln(2.68) ÷ ln(3.39) ≈ **0.81**, the sigmoid flattening under its own load. And total capsaicinoid and receptor loudness do not even co-rank perfectly across a designed set, because the balance of capsaicin against its slower, less receptor-potent cousin dihydrocapsaicin shifts the loudness independent of the mass; the Scoville column counts molecules, while IFI-H measures what the channel does with them. Where the two disagree, the Index sides with the tongue, which is downstream of the receptor and not of the chromatography trace.

The top rung exists because the cultivar behind it was designed with pathway-level precision. Our flux model places primary control of capsaicinoid output in the phenylpropanoid pathway upstream of the terminal capsaicin-synthase step — at 4CL and HCT/CCoAOMT — meaning Pun1, the gene that assembles the final capsaicin molecule, labels the product without governing the rate. Two AI-designed substitutions in Pun1 — **S39L** and **L345G**, found by a protein language model and refined by a genome-scale flux-balance model and designed for introduction into the plant with CRISPR — optimize the terminal assembly; together with the upstream pathway design, our model predicts peak output of approximately **2,937,948 SHU**, with a batch-to-batch variation design specification of **≤7%** coefficient of variation. That second number is what makes the Edge a rung at all. A band is only usable if the fruit lands on it every time, and the honest baseline for a wild-type pepper is a pod-to-pod CV near **42%**. Put a target on the Burn — say IFI-H 400 — and watch what each tolerance does to it. At 42%, the standard deviation is 0.42 × 400 = 168; two standard deviations reach from **64 to 736**, smearing a single pod across the Warmth, the Glow, the Burn, and into the Fire — four bands at once. At our engineered ≤7%, the standard deviation is 0.07 × 400 = 28, and two standard deviations span **344 to 456**, a tight sit well inside the Burn. The Edge is the extreme case of the same arithmetic: a wild-type plant at 42% could never occupy the top rung as a rung, because its spread would run from below the Inferno to beyond any real fruit. The frontier is reachable only by design, never by breeding, for the simple reason that you cannot stand on a step you keep sliding off.

At the very top, Phoenix is the cultivar we designed to register on every axis at once — the reason it anchors a full-spectrum reference and not only the heat scale:

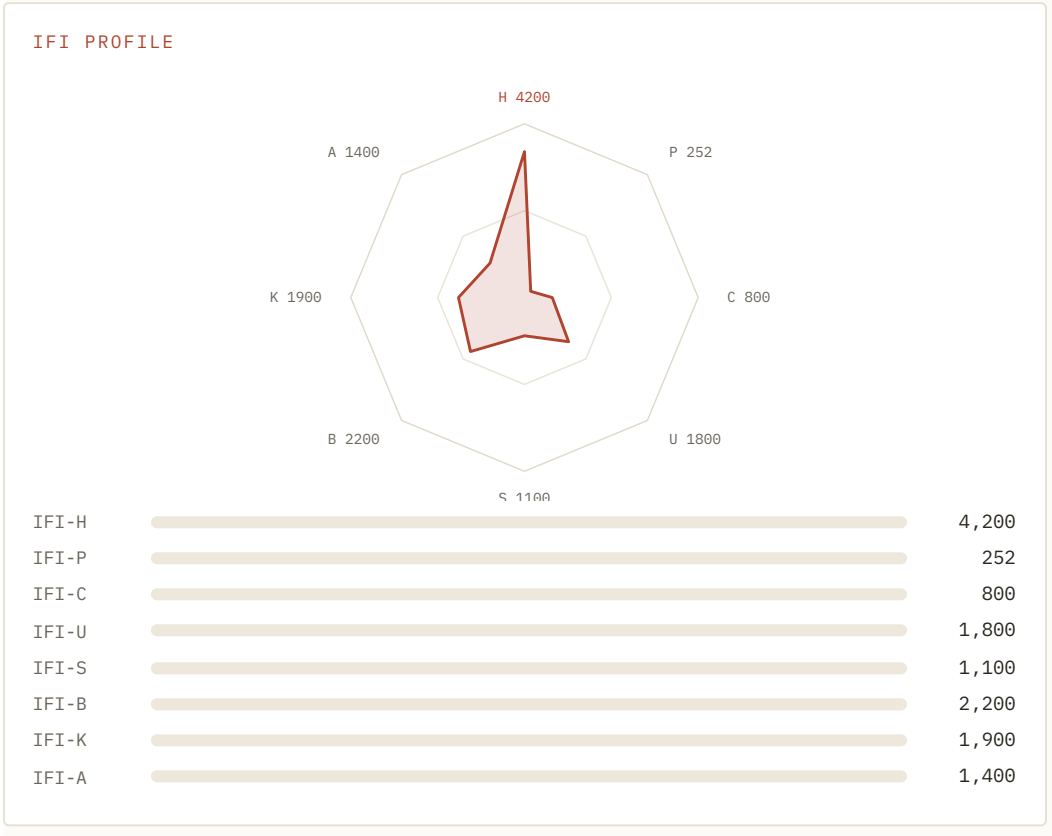

Heat at 4,200 towers over everything, but it does not stand alone: there is real pungency beneath it, a cool edge, a savory weight, a bitter spine, and the long kokumi continuity that keeps the whole thing from reading as a single screaming note. And the chemistry is genuinely extreme — our predicted ~2,937,948 SHU is about **18% of pure capsaicin's 16,000,000**, a capsaicinoid loading no field-grown fruit approaches. A pod predicted at nearly a fifth pure fire by mass is not an ingredient you spoon. It is one you titrate, and that fact, more than any threat-level name, is what the word *Edge* actually means.

Cooking to a band

The ladder's purpose is operational: each band is a target, and a dish is built to a target by composition, not by hope. The standing temptation is to reach for the hottest thing in the kitchen and use "a little." The arithmetic shows why that is precisely the wrong move, and it is worth walking once in full.

Suppose a sauce wants the confident warmth of the Burn, and you fix the heat as a capsaicinoid payload — the amount a cook's usual dose of cayenne would carry. Cayenne sits at about 50,000 SHU; converting, 50,000 ÷ 16 = 3,125 ppm = **0.31% capsaicinoids by mass**. A 10-gram spoonful therefore delivers 10 g × 0.0031 ≈ **0.031 g** — thirty-one milligrams of actual capsaicinoid into the pot. Now reach instead for Phoenix to deliver the identical payload. In our model, Phoenix is predicted at 2,937,948 SHU — 183,622 ppm, **18.4% capsaicinoids** — so the mass required is 0.031 g ÷ 0.184 ≈ **0.17 g**:

$$\frac{m_{\text{cayenne}}}{m_{\text{Phoenix}}} = \frac{2{,}937{,}948}{50{,}000} \approx 58.8$$

Ten grams of cayenne, or one-sixth of a gram of Phoenix, for exactly the same fire — the ratio is just the inverse of the two Scoville counts, **58.8×**. Grams become drops. But the two doses are not the

same food. The cayenne brings ten grams of pepper flesh along with its heat: its fruit, its faint sweetness, its red-earth body. The Phoenix brings 0.17 g of material, a mass so small it contributes no flavor at all — a thin, abstract burn floating over the dish with nothing underneath it. Same delivered IFI-H; utterly different plate.

That is the ladder's real gift. It tells you not only *how hot* to aim but *which rung's character* belongs in the dish. If you want a Burn, you reach for a cultivar whose home is the Burn, at the dose the arithmetic sets, and the heat you build tastes like an ingredient instead of an additive — cleaner, truer, and more controllable than heat achieved by homeopathic fractions of something monstrous. The ladder is not a dare to climb. It is a set of shelves, and good cooking is mostly knowing which shelf to reach for, then dosing to the number rather than to the fear.

The honest limits

A ladder this clean invites more trust than it has earned, so we name its edges plainly.

First, the bands are a consensus, not a constant. They are drawn where perception draws them, and perception differs from mouth to mouth — the density of TRPV1 on a person's nerve endings, a lifetime of exposure and the slow desensitization it builds, even the temperature of the dish, all shift where a given rung lands. Our Warmth is genuinely warm for most tables; it is not a promise about every tongue. The number fixes the stimulus honestly; the person is still the person.

Second, a dish's felt band is not its average capsaicinoid concentration. Heat is delivered to receptors directly, it builds across a meal, and it is pulled around by what surrounds it — capsaicin is fat-loving and barely touches water, so the same payload reads hotter in a lean broth than in a coconut curry, and sugar and acid shift it again. This is why we specify a *source* rung and a *dose* rather than printing an IFI-H on the finished plate: the plate's heat is an experience with a time-course, and the ladder measures the height of the burn, not its whole arc.

Third, the two columns are joined by a sigmoid, not a formula you can carry in your head. The local exponents we computed — 0.43 in the Fire, 0.87 at the Edge — are the curve's slope sampled in places, not a law. Push the arithmetic too far past the rungs and it will mislead you; the ladder is reliable as a set of shelves and only approximate as an equation. We would rather you trust it exactly as far as it can be trusted.

The ladder, in the end, is a map. It is a faithful one, drawn to the receptor and held to a few percent by design specification, and it will carry a cook anywhere on the scale with intention instead of luck. But the meal is still the territory, and we hold the map the way we hold every instrument we have built: as a true account of one thing, offered in service of the table.

Six shelves, and the freedom to choose

Between a shishito's sunlit 40 and Phoenix's 4,200 lies everything a kitchen could want — the slow glow you can cook a whole supper inside, the bright clean lift that brings a sauce up off the plate, the deep architectural fire a chili paste is built around, and, at the far end, a frontier we approached exactly once and learned to measure in drops. Six shelves, each with its peppers and its kitchens and its number, and now the freedom to move among them on purpose: to name a warmth, grow it true, and set it down where you meant it to be. Heat, to us, is a pleasure and never an assault. The ladder exists so that a cook can ask for exactly this much warmth and receive exactly that, and carry it to someone they are feeding with care.

These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Chapter 6 · Band 0 and 1 · The Warmth and the Glow — *heat you can live inside.*

Most of the heat anyone eats in a lifetime lives here, in the first two bands, and it is the part of the ladder we are proudest of — not because it is difficult, but because it is where heat stops being a stunt and becomes a kind of hospitality. The Warmth, band 0, runs from the barely-there flush of a shishito near IFI-H 40 up to about 80; the Glow, band 1, from 80 to roughly 250, the jalapeño-to-serrano country where a dish wakes up without ever raising its voice. This is the register of a poblano folded warm into a tortilla, of a serrano slivered over lime rice, of the one blistered padrón in ten that surprises you. It is heat as texture and temperature rather than heat as an event — a warmth you could live inside, and most of us happily do, meal after meal, without ever once calling it bravery.

We write this chapter as the family that learned to tune these peppers, and it is the bottom of our own ladder we return to most. The superhots carry the headlines. The Warmth and the Glow carry dinner.



The pepper at rest — the everyday workhorse of the lower bands, the plant that taught humanity that a little fire makes everything else taste more like itself.

The two bands, in plain numbers

Before the physiology, the map. The lower ladder is only two rungs wide, and nearly the whole world's cooking happens on them.

BAND	NAME	IFI-H RANGE	WHAT LIVES HERE	TRADITIONAL SCOVILLE RANGE
0	The Warmth	40 – 80	shishito, poblano, Anaheim, a mild padrón	~50 – 2,500 SHU
1	The Glow	80 – 250	jalapeño, Fresno, serrano, fresh guajillo	~2,500 – 25,000 SHU

Two things are worth saying about this table before we take it apart. The first is that the IFI-H column and the Scoville column are *not* the same measurement wearing two coats. Scoville Heat Units count how much capsaicinoid a fruit holds — grams of molecule per gram of pepper, read today by HPLC and converted at the classical figure of roughly 16 SHU per part-per-million of capsaicin. IFI-H counts something subtler: how *loudly* that capsaicinoid load speaks to the receptor, anchored to the half-activation of a native wild-type reference so that a tomato, a chili, and a cup of tea can be read against one another on a single page. The two track each other closely at the gentle end, and they come apart at the top, for reasons we will make exact. The second is that both columns are narrow here. The entire Warmth band spans a factor of two in IFI-H; the entire culinary world below a serrano fits inside a single decade of Scoville. This is the compressed, forgiving foot of the ladder, and the compression is not an accident of our scoring. It is the receptor's own doing.

What the receptor is doing down here

Capsaicin binds TRPV1, a non-selective cation channel standing in the membrane of the trigeminal nerve endings that wire the tongue. The channel is a homotetramer — four identical subunits arranged around a central pore, each subunit folded into six membrane-spanning helices — and capsaicin slips into a vanilloid pocket cradled between the third and fourth of those helices, hydrogen-bonding to a tyrosine and a threonine that line the cleft. It does not warm the tissue. What it does is more elegant and more unsettling: it lowers the channel's thermal activation threshold. TRPV1 is, in its day job, a thermometer that opens above about 43 °C; capsaicin drags that threshold down below body temperature, so the channel reports a fire at 37 °C that is not there. The body feels heat because the molecule has rewritten the number the thermometer calls hot.

The size of that sensation is governed by a dose-response curve, and the curve is the whole reason gentle heat behaves so gently. Receptor activation follows a Hill relationship,

```
θ = C^n ⁄ (EC₅₀^n + C^n)
```

where θ is the fraction of channels open, C is the capsaicin concentration at the nerve ending, EC_{50} is the concentration that opens half of them, and n is the Hill coefficient that sets how steeply the curve rises. For capsaicin at TRPV1 the receptor's half-activation sits near **$EC_{50} \approx 0.7\ \mu M$**, and the channel's positive cooperativity gives a Hill coefficient of about **$n \approx 2$**. Put those two constants into the equation and the Warmth band stops being a metaphor and becomes a table of concentrations.

CHANNELS OPEN (θ)	CAPSAICIN AT THE ENDING (C)	C / EC_{50}	SENSATION
1%	0.070 µM	0.10	the first faint flush
5%	0.161 µM	0.23	a clear, low warmth
10%	0.233 µM	0.33	the Glow, holding
50%	0.700 µM	1.00	a serious, insistent burn
90%	2.100 µM	3.00	the roar of the upper bands

Read the toe of that curve and the lower bands explain themselves. A bite in the Warmth register opens somewhere between one and ten percent of the available channels — a mouth concentration of a quarter of a micromolar or less, comfortably below the half-point. Only a small minority of TRPV1 channels ever open, and the calcium that floods in through them is correspondingly modest. That matters, because TRPV1 shuts itself off through a calcium-fed feedback loop: the inrushing Ca^{2+} binds calmodulin, the calmodulin clamps the channel, and with enough of it the receptor desensitizes and the fire dies on a slammed door. At one-to-ten-percent occupancy there is not enough calcium to trip that loop hard. The result is a warmth that *sustains* rather than spikes — it arrives, it holds a long pleasant plateau, and because it never floods the system it never has to shut itself down. The gentleness is not the absence of a mechanism. It is a mechanism running at its low, patient setting.

There is also a built-in margin for error at the bottom, and it is geometric. Because the Hill coefficient is about two, the span from the first faint flush to the full roar — one percent to ninety percent — covers a thirty-fold range of concentration. Within the Warmth band specifically, moving from that first flush at one percent to a clear glow at ten percent takes a 3.3-fold change in how much capsaicin reaches the nerve. Down here, a cook has elbow room: a little more pepper nudges the warmth up a little, not off a cliff. The same cooperativity that makes the superhots feel like a trap door makes the Warmth feel like a dimmer switch.

The table has a check built into it, from an entirely independent direction. A person can first detect capsaicin dissolved in water at roughly 0.1 ppm — about 0.3 µM. Drop that number onto the Hill curve and it lands squarely in the toe, between the one- and ten-percent marks, exactly where the first faint flush should sit. The concentration at which people say they *notice* heat is the concentration at which a few percent of their TRPV1 channels open. Two instruments — a psychophysics threshold measured on tongues and a dose-response curve measured on channels — agree on where the bottom of the ladder begins.

Two honesties keep that agreement from looking tidier than the biology actually is. The first is that the 0.7-µM half-point and the Hill coefficient near two are measured on receptors expressed in cell lines; the trigeminal ending in a living mouth runs warmer, at a shifting pH, inside a lipid membrane that sequesters a molecule as oil-loving as capsaicin. The *free* concentration that reaches the channel is a small and uncertain fraction of whatever the fruit held — which is exactly why a bulk calculation from a pepper's weight overshoots wildly, and why the receptor, not the gram, is the honest

unit of heat. The second is that TRPV1 is polymorphic: common human variants shift the channel's sensitivity, so the same pod reads a shade hotter or cooler from one person to the next. The Index scores the pepper. The mouth still gets the last word, and we would not have it any other way.

Reading gentle heat on the Index

The Iyengar Flavor Index puts one number on that loudness. The formula is a ratio and nothing more:

$$\text{IFI-H} = 1{,}000 \times EC_{50}(\text{reference}) \,/\, EC_{50}(\text{candidate})$$

The reference is the native wild-type pepper, scored at exactly 1,000 by definition — its EC_{50} divided by itself, times a thousand. Everything else is read against it. A pepper whose material is more potent needs *less* of itself to half-activate the receptor, so its EC_{50} is smaller, so the ratio is larger, so its IFI-H climbs. Invert the gentle anchors of the lower bands and the index tells you, in plain multiples, how much fruit it takes to reach the half-point.

PEPPER / ANCHOR	IFI-H	EC_{50}(CANDIDATE) / EC_{50}(REFERENCE) = 1000 / IFI-H	READING
Shishito	40	25.0×	needs twenty-five times the material of the reference
Warmth ceiling	80	12.5×	half as much material as a shishito
Design target	120	8.33×	a companionable band-1 glow
Glow ceiling	250	4.0×	a bright serrano edge
Wild-type reference	1,000	1.0×	the definitional anchor
Phoenix · SS-0096	4,200	0.238×	the top of the whole ladder, for scale

The arithmetic is worth doing slowly once. A shishito scored at IFI-H 40 carries a candidate EC_{50} twenty-five times larger than the reference: you would need twenty-five times as much shishito material to drive the receptor to half-max as you would of the wild-type standard. The band-1 design target at IFI-H 120 needs 1,000 / 120 = 8.33 times the reference's concentration — three times more potent than the shishito, exactly as the index numbers say. And to set the whole thing against the sky: Phoenix, the one cultivar we built to the far end at IFI-H 4,200 and our model predicts approximately 2.95 million SHU, inverts to an EC_{50} barely a quarter of the reference's. From the shishito to Phoenix is a hundred-and-five-fold swing in receptor loudness — and the two bottom bands of this chapter occupy only the gentlest sliver of it, between twenty-five-times-quieter and four-times-quieter than the reference. Any point on the ladder inverts the same way: a mid-Warmth pepper at IFI-H

60 carries an EC_{50} of 1,000 ⁄ 60 = 16.7 times the reference, falling exactly where it should between the shishito's twenty-five and the Glow target's eight. The formula is monotone and the arithmetic is honest in both directions — give it a score, it returns a concentration; give it a concentration, it returns a score.

A fair question is why the Scoville and IFI-H columns of the opening table track so closely here and not at the top. The honest answer is that loudness depends on the *shape* of the capsaicinoid load, not only its size. Capsaicin and dihydrocapsaicin, the two dominant capsaicinoids, differ in the saturation of their fatty-acyl tails, and they engage TRPV1 with slightly different potencies and slightly different desensitization kinetics — a different EC_{50}, a different burn. In the lower bands the mixture is near-constant from pepper to pepper, so more capsaicinoid simply means proportionally more heat, and the two columns stay nearly locked. Fit a power law to the band anchors and you get IFI-H ≈ 7.4 × SHU^0.37 across the culinary range — a tidy, gently compressing curve that places a 2,500-SHU jalapeño right at our design target of IFI-H 120. But that law is deliberately local. It fails at the engineered top, and the failure is the point: among our ten flagship superhots, Kindle carries a predicted 1.25 million SHU yet scores IFI-H 2,000, while Cinder, at a predicted 1.05 million SHU, scores a *louder* 2,290. More Scoville, less heat — because we tuned the capsaicinoid shape, not merely its quantity. The index can see that; a Scoville number alone cannot. Down in the Warmth and the Glow we rarely need that distinction. Up top it is the entire game.

One last conversion closes the loop between the two rulers. At the classical 16 SHU per ppm, a mild 2,500-SHU jalapeño holds about 156 ppm of capsaicinoids — roughly fifteen-thousandths of a percent of the fruit by weight. A shishito near 100 SHU holds about six ppm. These are vanishingly small loadings, and that smallness is the reader's real reassurance: the entire difference between a dish that whispers and a dish that doesn't is a few parts per million, placed exactly.

Why a little warmth makes everything taste more

There is a physiological generosity to this register worth naming precisely, because it is the honest version of what people vaguely mean when they say spicy food "wakes up" a dish. Mild TRPV1 activation on the trigeminal endings does two things at once. It triggers salivation — the mouth genuinely waters — and saliva is the solvent and enzyme bath that carries every *other* flavor molecule to its own receptor; a sweet compound cannot reach TAS1R2, a glutamate cannot reach mGluR1, until it is dissolved and delivered. A gently warmed dish is, quite literally, a dish you taste more of, because the warmth has primed the instrument that does the tasting. And because the trigeminal and gustatory signals converge centrally, a low, steady heat raises the whole sensory gain: sweetness leans forward, acidity sharpens, the round savory depth of a broth gets an edge to stand on. Heat in band 0 is less a flavor than a lens.

Layered on top is a faint endorphin response — the body's small reward for a signal of manageable stress, the same mechanism behind a runner's calm scaled down to the size of a good dinner. None of this is a health claim, and our family makes none; it is simply the mechanics of why a little heat feels good, and why the lower bands repay a cook's attention so richly. It is also why we design the Warmth and the Glow for the widest possible table — the cautious palate, the child's plate, the guest who "doesn't do spicy" and discovers, at IFI-H 60, that they do after all. Winning that person matters to us more than any record at the far end of the ladder, because the point of learning to design heat was never to exclude. It was to invite.



Backlit greenhouse at golden hour — the quiet room where a warmth is tuned pod by pod before it is ever trusted to a table.

The hard part is restraint

Designing the low end is, counterintuitively, a problem of restraint rather than addition, and it is where the platform's precision earns its keep most quietly. Anyone can breed a mild pepper; what is hard is breeding a mild pepper that is *reliably* mild — one that does not, a pod in twenty, deliver a band-3 ambush because the placenta ran hot that week. To see why the placenta is the culprit, look at where the heat is actually made.



Cross-section of the pod — the pale placental ribs holding the seeds are where capsaicin synthase works and where every pod's heat is decided.

Capsaicin is not made in the flesh and not, despite the folklore, in the seeds. It is synthesized in the pale placental ribs that hold the seeds, in the epidermal cells of that tissue, where the enzyme capsaicin synthase — the product of the *Pun1* locus — condenses vanillylamine with a branched-chain fatty-acyl group to build the

molecule, then secretes it into blister-like vacuoles along the septum. Vanillylamine arrives down the phenylpropanoid branch that runs in every vascular plant; the branched acid arrives from amino-acid metabolism; only the terminal condensation is pepper's own trick. And that terminal step is exquisitely sensitive to the plant's circumstances. Light, heat, and drought stress all push *Pun1* expression up. A pod ripening at the sunlit top of the canopy sees more of all three than a pod in the shaded interior of the same plant — and so it fires the pathway harder and runs hotter, on identical genes.

That is the entire low-band problem, stated as a number: **wild-type pungency varies by more than forty percent, pod to pod — a batch coefficient of variation of about 42%.** In the superhot bands, forty-percent scatter is a nuisance you can breed around. In the Warmth and the Glow it is the whole game, because the point of a gentle cultivar is *trust*, and a heat you have invited to the table has to behave. The shishito and the padrón are the perfect emblems of the failure mode, because their whole charm has always been a gamble: roughly one pod in ten runs genuinely hot while the rest stay mild, a roulette that is delightful at a tapas bar and intolerable in a product a cook is trying to rely on. That lottery is pure *Pun1* variance — the same gene expressed at wildly different levels on the same plant. Traditional breeding cannot fix it without breeding the heat away entirely, trading the thrill for blandness.

We can keep the thrill and remove the ambush, and the lever is not the gain-of-function heat we write into the superhots but its opposite: holding capsaicin-synthase expression to a precise, low set point, and clamping the variance around it. Where the far end of the ladder is designed around two AI-designed *Pun1* substitutions — S39L and L345G — written in to push catalytic output up, the low end earns its reliability through regulated, uniform expression that buffers the plant against the light-and-stress induction that drives pod-to-pod swing. Across the 1,060 constructs in the pepper, the one native host among our forty-two organisms, the result is a line held to a **batch CV of seven percent or less**, against the 42% of the wild type. What that six-fold tightening buys is best shown by propagating it through a single design target — a band-1 pepper tuned to IFI-H 120.

LINE	EXPRESSION CV	Σ (IFI-H UNITS)	±2Σ POD-TO-POD RANGE	CHANCE A POD CROSSES INTO BAND 2 (IFI-H > 250)	PODS TO CERTIFY THE LINE TO ±5 IFI-H (95%)
Wild-type	42%	50.4	19 – 221	~1 in 200	~390
Engineered	≤7%	8.4	103 – 137	essentially never (≈15σ)	~11

Walk the two rows. At the wild-type's 42% CV, a line aimed at IFI-H 120 carries a standard deviation of 0.42 × 120 = 50.4 index units. Two standard deviations reach from 19 to 221 — a "mild" pepper that routinely delivers anything from a bare flush to a serrano's bright edge, and that crosses the 250 line into the next band up about once in every two hundred pods. (The true tail is fatter still, because the distribution is right-skewed: the sun-stressed canopy pods pull a hot

tail rather than a symmetric spread, which is exactly why the shishito's one-in-ten surprise is a tail event and not a tidy bell.) At the engineered 7% CV, the standard deviation collapses to 8.4; three standard deviations span only 95 to 145; crossing into band 2 would take a fifteen-sigma excursion, which is to say it does not happen. The dial has replaced the roulette. The surprise, when a cook wants one, is now a *specification* — a line in which the fraction of hot pods is itself a chosen number — rather than an accident of the weather.

Reliability is not only a kindness to the eater; it is a cost the whole program pays or escapes. The last column shows it. To certify a line's mean to within ±5 IFI-H units at 95% confidence takes $n = (1.96\ \sigma / 5)^2$ pods. At the engineered σ of 8.4 that is about eleven pods. At the wild-type σ of 50.4 it is about three hundred and ninety — a thirty-five-fold heavier sampling burden just to prove the same claim. Precision at the plant pays for itself at the bench.

The failure modes are worth naming, because designing for reliability means designing against specific ways it breaks. A set point can drift upward when a vigorous rootstock or a long hot season leans on even a buffered promoter. A line can go quietly bimodal if a fraction of plants silence the transgene and revert toward the wild-type gamble, splitting one honest distribution into two dishonest ones. And because the native distribution is right-skewed, the mean is a poor custodian of the worst pod — the average can look perfect while the hot tail misbehaves. So we certify on the tail, not the average: we read the ninety-fifth and ninety-ninth percentiles of a batch and hold *those* inside the band. It is a stricter promise than a mean can make, and it is the only promise a cook can actually taste, because no one eats an average — they eat the pod in their hand.

A note on where the seven-percent target itself comes from, since a careful reader will ask. Part of it is biology — how tightly we can hold expression. But part of it is the ruler. The index is a ratio of two measured half-activations, and the error in a ratio adds in quadrature: if each EC_{50} is read to a 5% assay precision, the index inherits a CV of $\sqrt{(0.05^2 + 0.05^2)} = 7.1\%$. Seven percent is therefore close to the floor the measurement itself sets. We hold the biology to the precision of the instrument that watches it, and no looser — the plant is designed to be no noisier than the number we score it with.

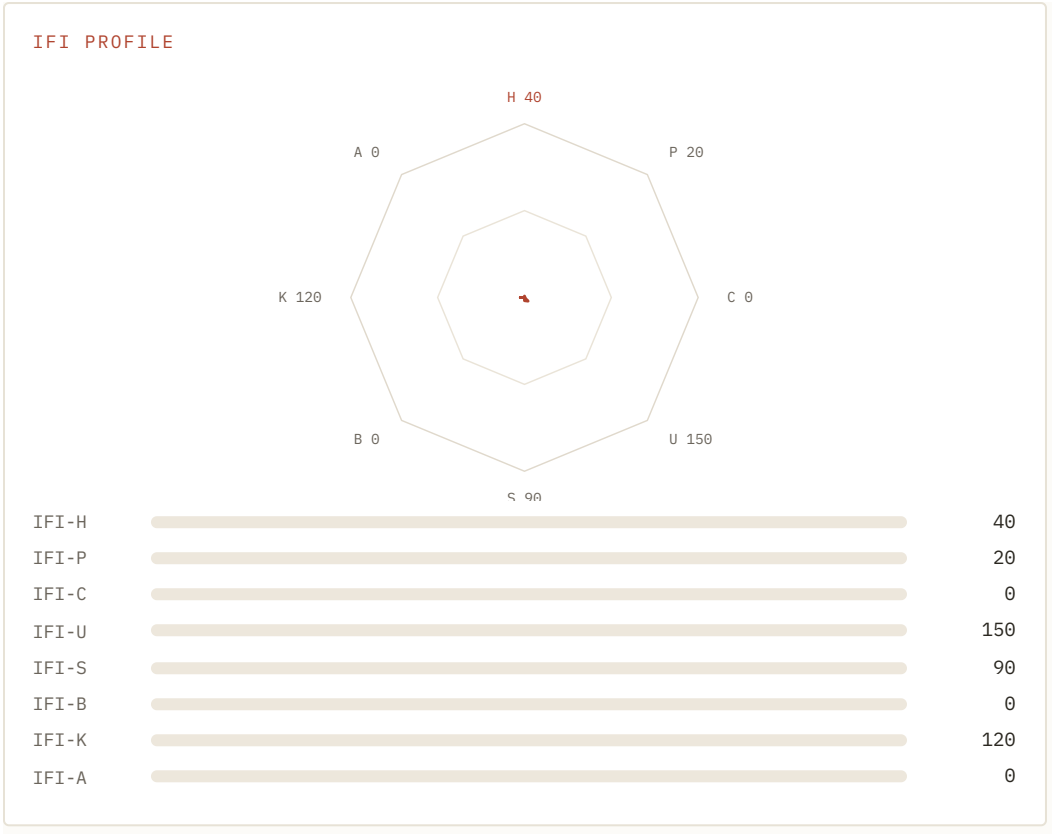

That scorecard is a Warmth-band shishito as we tune it — every pod at IFI-H 40, not one in ten at four hundred. The heat is barely a presence; what the profile really shows is a pepper built to be *dependable* first, with a whisper of pungency for lift and a quiet floor of umami and kokumi so that even at the bottom of the ladder there is something to taste. This is the design no tapas bar ever had: the shishito with the gamble taken out and the generosity left in.

Two lives: fresh and dried

The lower bands also hold one of cooking's quiet transformations — what happens to a mild pepper when it is dried. The chemistry is a gift to the low-band cook specifically, and it turns on a single fact: capsaicin is heat-stable and essentially non-volatile. When a fresh poblano at a gentle IFI-H is dried into an ancho, its water leaves, its sugars and amino acids concentrate and brown through Maillard chemistry into something darker, raisined, almost chocolatey — and its capsaicin sits right where it was, now simply packed into less mass. The *flavor* blooms; the *position on the ladder* barely moves. A fresh Anaheim becomes a dried California chili; a jalapeño, smoked and dried, becomes a chipotle with a band-1 warmth wrapped in woodsmoke. The capsaicinoid concentration per gram rises as water is lost, but the receptor loudness of a *rehydrated* pepper — the form a cook actually eats — returns close to the fresh value, because the drying concentrated the heat and the water of cooking lets it back out again. The bookkeeping is almost comically clean: a poblano gives up around ninety percent of its mass to water on the way to becoming an ancho, so a capsaicinoid loading of a few hundred ppm in the fresh fruit concentrates roughly tenfold in the dry one — and then dilutes back out, almost to where it began, the moment the ancho is soaked or simmered. The heat travels out of the pepper and home again. Only the browning stays behind.

IFI PROFILE

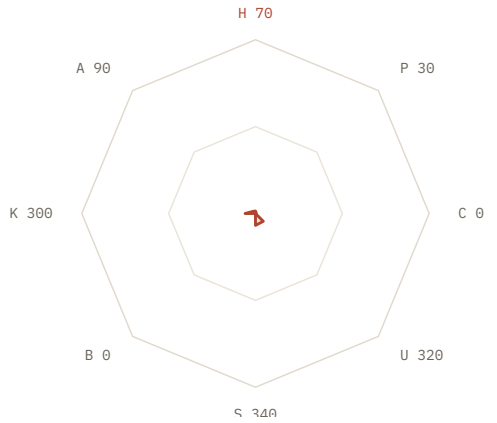

IFI-H	70
IFI-P	30
IFI-C	0
IFI-U	320
IFI-S	340
IFI-B	0
IFI-K	300
IFI-A	90

That is the ancho the poblano becomes — the heat held near where it started, in band 0, while sweetness, umami, and kokumi climb and a little astringency arrives from the concentrated skin. It is the clearest proof that gentle heat need never mean thin flavor: the same pod that warms a dish fresh can, dried, give it a deep smoky-sweet backbone at nearly the very same place on the ladder. We breed our mild cultivars with both lives in mind, fresh and dried, because a warmth worth designing is a warmth worth tasting twice.

The scorecard of a kind pepper

IFI PROFILE

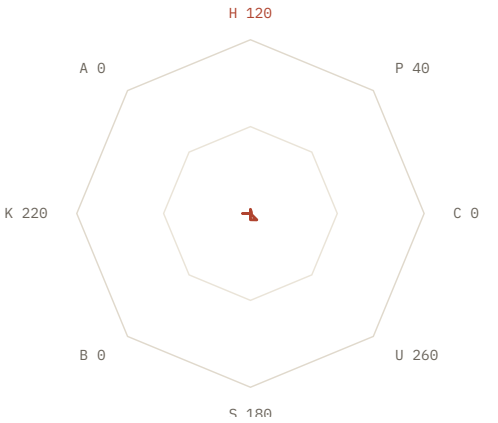

IFI-H	120
IFI-P	40
IFI-C	0
IFI-U	260
IFI-S	180
IFI-B	0
IFI-K	220
IFI-A	0

This is a representative band-1 culinary pepper as we tune it — heat at a companionable 120, but notice the supporting axes, because they are where the real design lives. A little pungency, IFI-P 40, for brightness and lift. A genuine floor of umami at 260 and kokumi at 220, so the warmth has savory ground to stand on and a continuity that carries it across the whole bite rather than flashing and vanishing. A touch of sweetness at 180, so the pepper reads as ripe rather than green. No cool-burn, no bitterness, no astringency — nothing competing, nothing to apologize for. Those supporting

numbers do not change what TRPV1 is doing — a kokumi compound at CaSR cannot open a heat channel — but they change what the brain *makes* of the heat downstream, where the trigeminal and gustatory streams are summed: a warmth laid over real savory depth reads rounder, lands softer, and is welcome for longer than the same warmth over nothing. The receptor sets the loudness; the rest of the profile decides whether it is music. This is what designing heat in the lower bands actually looks like: not a loud note but a well-set table, a warmth that makes room for everything around it. We have cooked whole meals in exactly this register — a poblano folded warm into fresh tortillas, a serrano slivered over lime rice, a padrón blistered hot and finished with flaked salt — and in every one the heat is the host, not the event.

These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

The lower bands are not the easy part of our work. They are the kind part, and we treat them as the heart of the catalog rather than its warm-up act. Hold the principle and the whole chapter folds into a sentence: the bottom of the ladder is where heat is most like love — present, steady, reliable to a fault, and never once asking to be noticed.

Chapter 7 · Band 2 and 3 · The Burn and the Fire — *where the meal wakes up.*

Climb into the middle of the ladder and heat changes its grammatical role. In the lower bands it was an adjective, a quality lent to other flavors — a warmth under the sweetness of a roasted pepper, a hum beneath a broth. Here, across the Burn and the Fire — IFI-H roughly 250 to 1,100, cayenne up through habanero and scotch bonnet — it becomes a noun. The dish is now *about* the heat, or at least the heat is one of the two or three things it is about. This is the register of a proper hot sauce, of a curry that announces itself before the first spoonful lands, of the fruit-forward habanero bloom that defines Caribbean and Yucatán cooking. It is where most people who say they love spicy food actually live, and it is the band we at the family bench return to most often, because it is where fire stops being a dare and becomes a flavor you would choose for its own sake.



Wildfire at full ripeness — a *Capsicum chinense* chassis of the kind that, in the wild, gives the world its scotch bonnets and habaneros: fruit and fire in the same bite, the aromatic heat the middle bands are built around.

The defining feature of the Burn and the Fire is that heat and aroma arrive together. The species that rules these bands, *Capsicum chinense*, is not merely hot; it is intensely *fragrant*, carrying floral, citrus, and stone-fruit volatiles that bloom alongside the capsaicin. This is a receptor story as much as a culinary one: the heat on TRPV1 and the aroma in the olfactory bulb reach the brain at the same instant and fuse into a single percept, which is why a good habanero does not taste like "a mild thing plus fire" but like a

specific, unified, fruit-and-flame experience. When we design in these bands we treat the aromatic profile as a co-equal specification with the heat, because a band-3 pepper that is all burn and no perfume is a failure no matter how precisely it hits its IFI-H. Fire without fragrance is just damage, and we have no interest in damaging the people who eat what we grow.

Where the Burn and the Fire sit on the Index

Everything in these bands is anchored to one equation, the same one that anchors every axis of the Iyengar Flavor Index. It does not measure how much of a compound a food holds; it measures how *loudly* that compound speaks to its receptor.

> $\mathrm{IFI}(\text{candidate}) = 1{,}000 \cdot \mathrm{EC}_{50}(\text{reference}) \div \mathrm{EC}_{50}(\text{candidate})$

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response. The lower the EC_{50}, the less you need to reach half-max, and the louder the food. We invert EC_{50} against a fixed native wild-type reference and scale by a thousand, so that the reference itself reads exactly IFI-H 1,000. That single choice of anchor tells you almost everything about the geography of the middle ladder. The Burn and the Fire run from about IFI-H 250 at the lower lip to about 1,100 at the ceiling — which means the wild-type reference, the loudest unimproved chili we chose to anchor the whole scale against, sits *at the very top of the Fire*. Nearly every pepper a person has ever cooked with lives here or below. Everything the breeders could not reach without us — the superhot ladder that climbs past 1,500 and 2,000 and ends at Phoenix (SS-0096) at IFI-H 4,200 — begins where band 3 ends and keeps going.

Put the band targets through the formula and the structure becomes concrete. Rearranged, $EC_{50}(\text{candidate}) = 1{,}000 \cdot EC_{50}(\text{reference}) \div \text{IFI-H}$, so the ratio $EC_{50}(\text{candidate}) \div EC_{50}(\text{reference})$ is simply $1{,}000 \div \text{IFI-H}$ — how much *more* of a candidate you need, relative to the reference, to reach the same half-max burn.

BAND	LINE (*C. CHINENSE* / *ANNUUM*)	IFI-H	EC_{50} VS REFERENCE	CAPSAICINOID SKEW	WHERE IT BELONGS
2 · the Burn	Clean cayenne	300	3.33×	capsaicin-forward	infused oils, dry rubs, bright sauces
2 · the Burn	Aromatic Thai	480	2.08×	capsaicin-forward	stir-fry finish, fresh salsa, pickles
3 · the Fire	Fruit-forward habanero	720	1.39×	balanced	mango & stone-fruit sauces, jerk
3 · the Fire	Deep scotch bonnet	900	1.11×	dihydro-leaning	pepper-pot stews, slow braises, ferments

Read the fourth column as a quiet proof. A band-2 clean cayenne at IFI-H 300 needs 3.33 times as much extract as the reference to reach half-max — it is quieter than the loudest wild chili, which is exactly what an IFI-H of 300 (well under 1,000) must mean. The deep scotch bonnet at IFI-H 900 needs only 1.11 times as much: it presses right up against the reference without passing it. The middle ladder, in other words, is the whole audible range of the

unengineered world, laid out in order. When we build above it we are not making chilies hotter the way a landrace is hot; we are making them speak at volumes no wild plant can hold steady.

The channel, one event at a time

The burn is a specific molecular event, and designing it well means understanding each step. Capsaicin binds TRPV1, a non-selective cation channel standing in the membrane of the trigeminal nerve endings that wire the tongue, lips, and palate — the same receptor gated by temperatures above 43 °C and by the protons of an acid. The molecule does not warm the tissue. It lowers the channel's thermal threshold, so the body reports a fire that is not there.

Follow it in order. Capsaicin is a lipid-loving molecule — a vanillyl head, an amide bond, a branched fatty tail — so it slips into the membrane and reaches TRPV1 from the inner leaflet, settling into the vanilloid pocket cradled between the S3 and S4 helices of each of the channel's four subunits. The vanillyl head hydrogen-bonds into that pocket; the fatty tail lies along a hydrophobic groove. That binding tugs on the S4–S5 linker, which pulls the lower gate — the bundle crossing of the S6 helices — open, while the outer selectivity filter dilates in concert. The channel, now open, is a wide mouth that lets sodium and, importantly, calcium pour in; TRPV1 passes calcium several times more readily than sodium. The inrush depolarizes the nerve terminal, voltage-gated sodium channels fire, and an action potential runs up the trigeminal nerve to the spinal trigeminal nucleus, across to the thalamus, and into the cortex, which localizes the signal to the mouth and reads it as heat and pain — because this is the identical channel that reports a scald. No tissue is harmed. The fire is a message, not an injury, and that distinction is the whole reason heat can be a pleasure we design rather than a wound we inflict.

Heat and aroma arrive as one thing

What are the *Capsicum chinense* aromatics, concretely? The perfume is a weave of esters and terpenes — the same chemical families that give tropical fruits their smell — and the overlap is not coincidence. Capsaicin itself is odorless and all but flavorless apart from its fire; the apricot, the citrus, the warm stone-fruit roundness of a good habanero live in the volatiles that evaporate off the fruit and climb the back of the throat to the olfactory bulb. The brain receives the capsaicin burn through the trigeminal nerve and the fruity volatiles through the olfactory nerve in the same instant, and binds them into a single object. This is why a habanero's heat seems to *taste* of apricot even though the heat molecule has no taste of its own.



A ripe pepper opened to its interior architecture — the placental ridges where, in the middle bands, capsaicin and the aromatic volatiles are manufactured in the same gland cells, which is why heat and perfume arrive in the mouth as one thing.

That fusion is anatomical as well as neural. Open a ripe *C. chinense* and the heat and the perfume are made in the same place — the glandular tissue of the placental ridges, the pale vertical walls that hold the seeds — so that a single gland cell is manufacturing the capsaicinoid and a share of the volatiles at once. When we design these cultivars we hold the aromatic profile as seriously as the IFI-H number, because in the middle bands the perfume is the heat's timbre, and a cultivar tuned for heat alone is to us an unfinished instrument. Still, honesty demands that the profile show only what we certify per lot. Commit that fruit-forward band-3 line to the Index on the one axis these bands are built to design, and it reads:

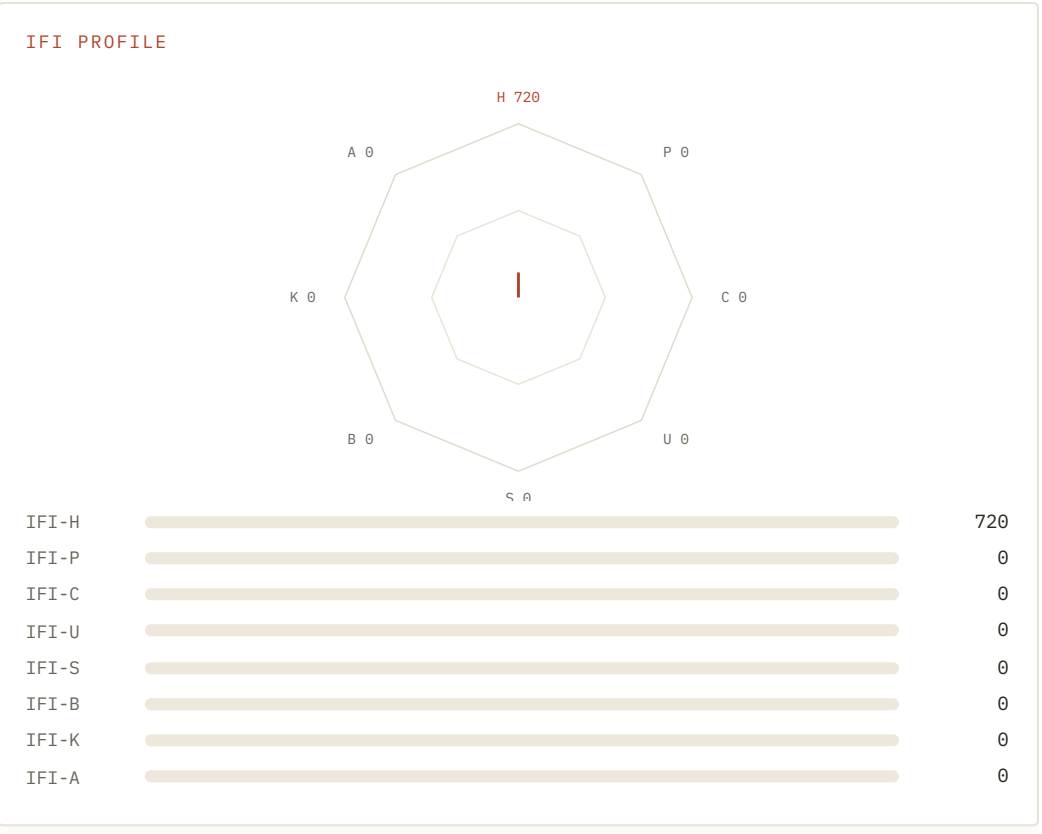

We leave the other seven lines open on purpose. The habanero's apricot-and-citrus bloom is real, and we push it forward by choice, but in the Burn and the Fire heat is the axis we certify on every batch; the aromatic axes we hold as bench specifications, not yet as committed per-lot Index numbers. An honest profile states the figure we stand behind and stays silent where we do not yet measure to the lot. The reverence is in both halves — in the number we guarantee and in the restraint about the ones we cannot.

The shape of the burn — capsaicin and dihydrocapsaicin

Heat in these bands is not a single level; it is a shape in time, and that shape is set in the biosynthesis, molecule by molecule. Two routes feed the capsaicin synthase. The first is the phenylpropanoid branch — PAL → C4H → 4CL → HCT → COMT — which runs in every vascular plant and builds the vanillyl head; an aminotransferase, pAMT, turns the vanillin-derived aldehyde into vanillylamine, the receptor-binding head of the finished molecule. The second route is the branched-chain fatty-acid branch, where BCAT and the ketoacyl synthase KAS assemble an 8-methyl branched fatty acyl chain from the carbon skeletons of valine and leucine. At the placental gland, the enzyme Pun1 — the AT3 acyltransferase, the true capsaicin synthase — joins the two: it forges the amide bond between vanillylamine and the branched acyl chain, and a capsaicinoid is born.

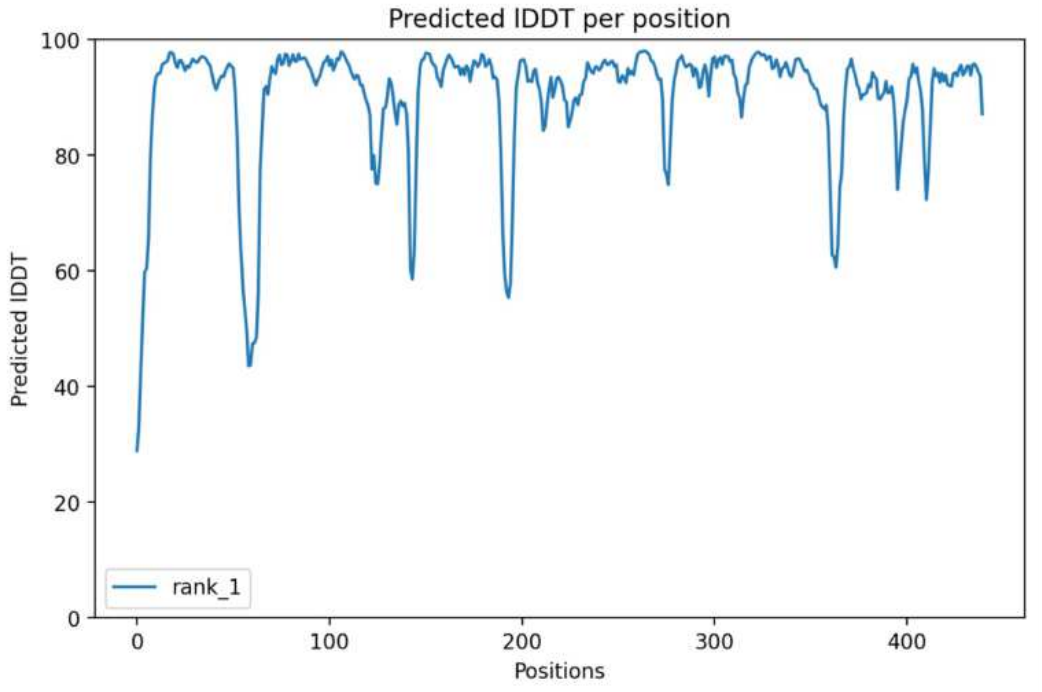

The Pun1 / AT3 capsaicin synthase, the enzyme that forges the amide bond at the heart of every capsaicinoid; the fit of the branched acyl chain in its active-site cleft is what decides whether a batch leans toward the bright capsaicin or the deep dihydrocapsaicin.

The crucial detail is the acyl chain's saturation. If the chain carries a single double bond — *trans*-8-methyl-6-nonenoyl — the product is capsaicin. If the chain is fully saturated, the product is dihydrocapsaicin, a molecule identical at the business end but a shade more lipophilic and slower to leave the membrane. The two feel different on the nerve. Capsaicin binds and releases fast: a bright, forward, front-of-mouth sting that flares and fades. Dihydrocapsaicin, lingering longer in the lipid, binds and releases slow: a deep, sustained glow that builds and holds. So the capsaicin-to-dihydrocapsaicin ratio — the skew of the capsaicinoid pool — is

not a footnote. It is the dial that sets *when* the heat happens. Tune the fatty-acid branch toward the saturated chain and you have built a banked coal; tune it toward the unsaturated chain and you have built a bright blade. The capsaicin-forward lines, with their fast front burn, belong where heat should *lift* — a finishing oil, a fresh salsa, anything you want to flash and let go. The dihydro-leaning lines, with their persistent tail, belong where heat should *hold* — a braise, a ferment, a stew that cooks for hours and wants a fire that cooks along with it.

The reason any of this can be made to a specification at all is the synthase. Every prior breeding program chased heat by selecting phenotype and hoping the genotype followed, and the batch-to-batch scatter of a wild landrace is punishing — a coefficient of variation near 42%, meaning the same plant line can be a third hotter or a third milder lot to lot. We inverted the approach: a protein language model scored every single-residue substitution in Pun1 by predicted effect on catalytic efficiency and fold stability. Our flux model predicts that control in the capsaicinoid pathway sits upstream — at 4CL and HCT/CCoAOMT in the phenylpropanoid branch — not at the terminal Pun1/AT3 step; Pun1 is the capsaicin synthase, the enzyme that labels the final product, but it does not govern the rate. What the design work gave us was a stabilized, consistent synthase — the same flux flowing through a steadier gate. Two substitutions, never before combined, emerged as the shipped double variant.

EDIT	EFFECT	ΔΔG (KCAL/MOL)
S39L	binding affinity ↑ (N-terminus)	−2.1
L345G	steric relief, active-site cleft	−1.8
C175S†	disulfide-misfold prevention	−0.9

† Structural annotation from our modelling; pulled from the final shipped construct. The shipped design is the double variant S39L+L345G only.

Each ΔΔG is negative — each edit stabilizes the fold or tightens the bind. The two shipped edits turn a temperamental enzyme into a dependable one, and the reproducibility they buy is what collapses batch variation from ~42% to ≤7% CV. That fidelity is the gift the superhot work hands back down to the middle bands. The same clean, stable synthase, run at a lower flux set-point, is what lets a band-3 sauce hit its target lot after lot instead of lurching across the ladder — and consistency, in a flavor you are asking people to trust on their own table, is a form of respect.

Loudness, not mass — where SHU and the Index part ways

The old measure of heat is the Scoville Heat Unit, and it is a measure of *mass*: how much capsaicinoid a pepper holds, read today by HPLC against the convention that one part per million of capsaicinoid is worth about 16 SHU. SHU scales almost linearly with the amount of pungent compound in the fruit. The Index measures something different — loudness on the receptor — and most of the time the two agree, because more capsaicinoid does speak louder. As a back-of-envelope rule, our model predicts our own ladder runs

near 700 SHU to the IFI-H point at the top: Phoenix (SS-0096) is designed at IFI-H 4,200, and our model computes approximately 2,937,948 SHU for that design — 699.5 predicted SHU per point. But lay the whole flagship ladder out and the rule visibly breathes.

CULTIVAR · SKU	PREDICTED SHU	IFI-H	PREDICTED SHU PER IFI-H POINT
Ember · SS-0087	866,000	1,570	552
Cinder · SS-0088	1,050,000	2,290	459
Kindle · SS-0089	1,250,000	2,000	625
Blaze · SS-0091	1,720,000	2,630	654
Wildfire · SS-0092	1,960,000	2,930	669
Scorch · SS-0093	2,210,000	3,220	686
Firestorm · SS-0094	2,470,000	3,490	708
Phoenix · SS-0096	2,937,948	4,200	700

The ratio wanders from 459 to 708 — and once, it inverts outright. Cinder (SS-0088) is designed at 1,050,000 predicted SHU and reads IFI-H 2,290. Its neighbor Kindle (SS-0089) is designed with *more* total heat by mass — 1,250,000 predicted SHU — yet reads IFI-H 2,000, a step *quieter* on the receptor. By the old scale Kindle is the hotter pepper; by the Index it is the softer one. There is no contradiction, only two honest instruments reading two different things: Kindle carries a capsaicinoid pool designed toward the heavier, less potent dihydro and nordihydro forms, so it is hotter in grams and milder in volume, while Cinder spends only 459 SHU to buy an IFI-H point because its capsaicin-forward pool is Index-efficient. This is the exact reason we built the Index. A cook does not eat milligrams; a cook eats a sensation with a volume and a shape, and the number on our label is tuned to the sensation. The apex of that predicted ladder, the only cultivar in the library designed to register on all eight axes at once, reads:

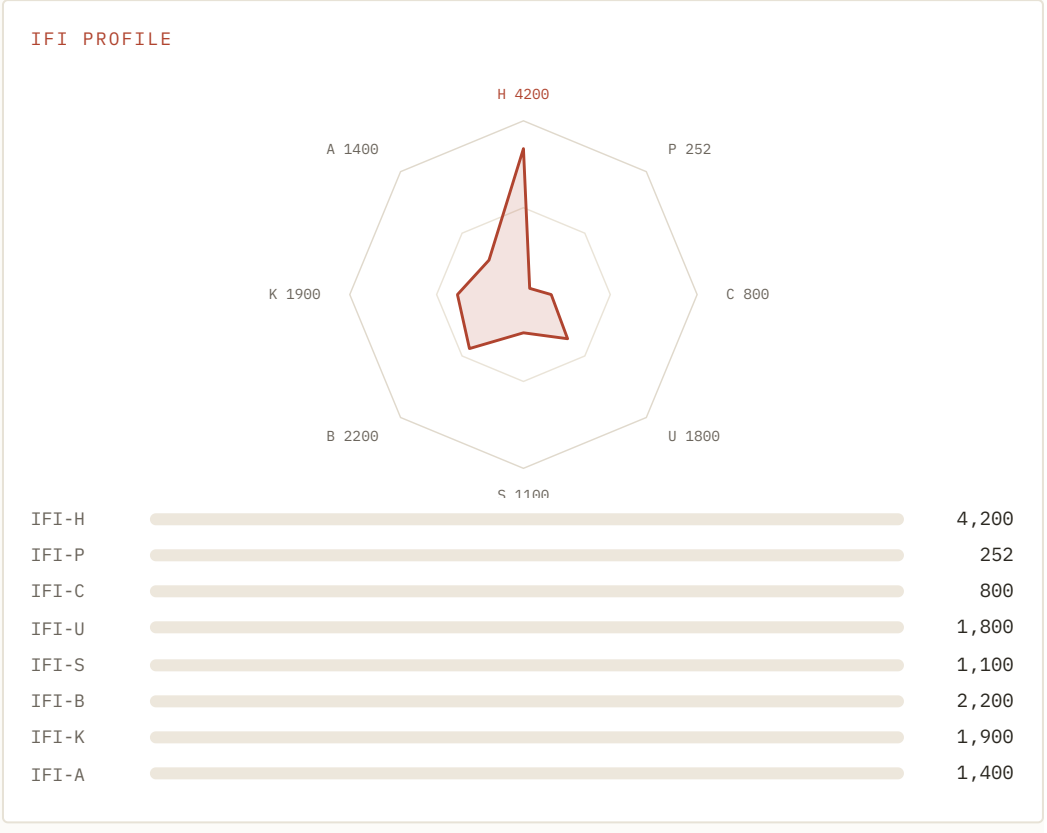

Where does this leave the middle bands in SHU? Honestly, in a range we will only estimate. Run the band-3 fruit-forward habanero at IFI-H 720 through even a middle-of-the-road 500-to-700 SHU-per-point rule and you land somewhere in the low hundreds of thousands of SHU — the real neighborhood of a habanero and a

scotch bonnet — but the capsaicin-forward skew of a bright band-2 line would read lower per point and a dihydro-leaning braise pepper higher, so we quote these as orders of magnitude, not certified figures. The IFI-H we stand behind to the lot. The SHU we offer as a rough translation for readers fluent in the old tongue.

Desensitization as a design parameter

At middle-band concentrations a single bite opens enough TRPV1, and admits enough calcium, to begin switching the channel off within seconds. The calcium that pours in binds calmodulin, which activates the phosphatase calcineurin to strip activating phosphates off TRPV1, and it drives the local depletion of PIP_2, the membrane lipid the channel needs to stay responsive. The open probability falls. In the language of the dose-response curve, the channel's EC_{50} shifts to the right — the same mouthful of capsaicin now reaches a receptor that demands more of it to answer. This is tachyphylaxis, and it is why the *second* bite of a band-3 dish is genuinely milder than the first, and the fifth milder still.

A little arithmetic shows how much milder. Receptor activation follows a Hill relation — the fraction of maximal response is $[A]^n \div (EC_{50}^n + [A]^n)$, where [A] is the agonist concentration and n the Hill coefficient. Take an illustrative n of 2, chosen here only to show the shape of the thing. A first bite that delivers capsaicin at three times the resting EC_{50} drives the channel to $3^2 \div (1 + 3^2) = 9 \div 10 = 0.90$ of maximum — a loud, full burn. Now let repeated exposure shift EC_{50} upward by a factor of 2.5. That same bite is now only 1.2 times the new, raised EC_{50}, and the response falls to $1.2^2 \div (1 + 1.2^2) = 1.44 \div 2.44 = 0.59$. The identical dose that read 0.90 now reads 0.59: a real, quantifiable fade, the receptor adapting in real time to a signal that has not changed.

A dish built in these bands therefore has an arc. It is hottest at the opening and settles as you eat, and a cook can either fight that curve or compose with it. We compose with it. We dose a band-3 braise a touch above its resting target precisely because we know the first spoonful runs hot and the middle of the bowl runs true, so that the *experience* — integrated across the whole meal, not sampled at the first bite — averages to where we meant it. Heat, in these bands, is not a static level. It is a curve, and designing it well means designing the whole curve.

Building a sauce — a base and a finish

The signature project of the Burn and the Fire is the hot sauce, and a great one is almost never a single pepper. It is a *base* and a *finish*. The base is a band-3 cultivar, often a dihydrocapsaicin-leaning one, cooked or fermented in early so its slow heat sinks evenly through the whole and gives the sauce its foundational warmth. The finish is a band-2 cultivar, capsaicin-forward, added late or raw so its bright sting sits forward and arrives first. Taste such a sauce and the two layers resolve in sequence — the quick front sting, then the deep sustained glow behind it — a heat with a beginning and a middle, engineered on purpose. The IFI-H on the label is the composite the two produce together; the *experience* is the interval between them, and designing that interval is the difference between a sauce that

merely burns and one that unfolds. It is why we so often hit a single heat target with two cultivars doing two different jobs rather than one cultivar doing a blunt one.

The formulation math is clean because SHU scales with dilution. Suppose we want to carry heat into a sauce from the most concentrated source we have and land the finished bottle at a bright-sauce 50,000 SHU. From a Phoenix-grade oleoresin at our predicted ~2,937,948 SHU, the dilution factor is 2,937,948 ÷ 50,000 ≈ 58.8 — one part of that concentrated heat to roughly fifty-eight parts of everything else the sauce is made of. In capsaicinoid terms, 50,000 SHU ÷ 16 ≈ 3,125 ppm, about 0.31% capsaicinoid in the finished sauce: a believable level for a bottle people will actually pour. Dial the target down to a gentle 10,000-SHU table sauce and the dilution opens to roughly 294-fold; push it up to a fierce 150,000 and it tightens to under 20-fold. The curve is linear, so the kitchen arithmetic is honest.

Consistency is where the engineered chassis earns its place. Say the sauce targets IFI-H 900, the deep end of the Fire. At our ≤7% batch CV, the standard deviation is about 63 Index points, so ~95% of batches (±1.96σ) fall between roughly 777 and 1,023 — the sauce stays inside the Fire, bottle after bottle. Build the same sauce on a wild landrace at 42% CV and the standard deviation balloons to about 378 points; the same 95% interval now sprawls from roughly 160 to 1,640, a span that crosses out of the Burn, through the whole Fire, and past the ceiling into territory no label could honestly describe. You cannot keep a wild pepper inside one band. That is not a flaw you taste once; it is a flaw you taste every time the bottle is different from the last, and it is the quiet reason a designed cultivar is kinder to the person who trusts it.

Fermentation — where a pepper becomes a sauce

One transformation belongs especially to these bands. Pack band-3 peppers in salt, hold out the air, and let lactobacillus work for a few weeks, and two things happen at once. The raw, green, sometimes harsh edge of fresh capsaicin-bearing fruit rounds into something deeper and savory as the bacteria build acids and new aromatic compounds; and the heat itself, barely diminished in total, seems to settle and lengthen, its sharpness sanded into a glow. Capsaicin is heat-stable and acid-stable — it survives the ferment nearly intact — so the fire is preserved while everything around it improves. In Index terms the move is specific: fermentation leaves IFI-H almost unchanged but lifts the umami and kokumi axes markedly, which is precisely the transformation that turns a one-note hot sauce into one with depth and continuity. It is why so many of the world's great chili condiments — Louisiana mash, Korean gochujang, the fermented bases of countless sauces — pass through live culture on their way to the bottle.

We design cultivars with the ferment in mind, choosing the dihydrocapsaicin-leaning, thick-walled fruit that stand up to weeks in brine without going to mush, because in the middle bands the ferment is often where a pepper becomes a sauce. It is also, quietly, the most nourishing moment in the whole process: the same live cultures that deepen the flavor leave behind a more digestible, more

alive food than the one that went into the crock. We do not make a health claim of it. We simply notice, with gratitude, that the thing that tastes better is also the thing that feeds you better, and we let that guide which peppers we grow for it.

Honest limits

We would rather name the edges of what we know than imply there are none. The IFI-H we certify to the lot; the SHU translations in these bands we offer as estimates, because the mass-to-loudness conversion breathes with the capsaicinoid skew and we will not quote a false precision. The aromatic axes that make a habanero worth eating are real and designed, but in the Burn and the Fire we hold them as bench specifications rather than committed per-lot numbers, and the profile stays silent where our measurement does not yet reach. The Hill arithmetic of desensitization we showed with an illustrative coefficient, to teach the shape of the fade, not to report a fixed constant for every mouth — tolerance and adaptation differ from person to person, and the curve we design to is a well-chosen average, not a law. And a sauce is only as consistent as its least consistent ingredient; the ≤7% CV is the cultivar's gift, and the rest the kitchen must keep. These are not apologies. They are the terms on which a number earns a reader's trust.

These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Band 2 is a bright blade. Band 3 is a banked coal. A cook who knows the difference — who can read which way the capsaicinoid pool leans, where the reference anchor sits, how the burn fades across a bowl — can put fire exactly where a dish needs it and nowhere it does not. That is the whole craft of the middle ladder, and it is where we love to work: the register where heat is at last a pleasure you would choose, grown clean, designed with intention, and given to the body on purpose.

Chapter 8 · Band 4 and 5 · The Inferno and the Edge — *the ten we built to frighten you kindly.*

Now we come to the top of the ladder, and the register of the writing has to change with it, because the top of the ladder is a different kind of place. The Inferno and the Edge — IFI-H 1,570 to 4,200, ghost peppers and scorpions and up past the world records into a country we had to build — are not where anyone cooks dinner. They are where the science of heat is pushed to its limit for the sake of what the limit teaches, and where the unit of culinary measure stops being the pod and becomes, quite literally, the drop. A gram of fruit at this altitude is not an ingredient; it is a reagent. We have always written about this country the way a climber writes about an unclimbed face — with awe, with precision, and with an unshakable insistence on ropes.

The ten flagships are our proof-of-concept superhots, all built on a *Capsicum chinense* chassis, and they exist as a *series*: a deliberate, legible climb from the merely ferocious to the frankly unprecedented. We gave them threat-level names because the names tell the truth about what they are — not dinner, but demonstrations, each one a rung we built to show the platform could place heat exactly where it aimed. What makes the series matter is not that the peppers got hotter. It is that they got hotter *on design*, in even steps, each one designed to land within a few percent of the Scoville target we set before the first seed went into soil. A single record-breaking pod proves luck. Ten cultivars advancing in even design steps prove a method — that heat is a dial we can set, not a jackpot we chase.



Solara — SS-0095, our model predicts 2.72 million SHU, IFI-H 3,980, six of eight axes expressed. The second-most-intense cultivar in our design series: handled gloved, grown in a controlled chamber, dosed to the milligram.

The ten, as a staircase of proof

Read the ladder as a single object and it stops being a list of hot peppers and becomes an argument. Each rung adds predicted Scovilles, and — this is the part that is easy to miss — each rung also adds *axes*, because the real achievement of the series is not

that the peppers got fiercer but that they got more *complete*, picking up sweetness and bitterness and the long savory continuity of kokumi as they climbed, until the last of them registers on all eight at once.

The derived columns below are worth keeping in view, because they translate the Scoville number into the thing it is actually a statement about: how much capsaicinoid the fruit carries. The conversion is the standard one — pure crystalline capsaicin sits near **16,000,000 SHU**, so dividing a Scoville figure by 16 gives its concentration in parts-per-million of capsaicin-equivalent, and dividing by sixteen million gives it as a fraction of pure capsaicin itself.

CULTIVAR · SKU	SHU	CAPSAICIN-EQUIV (PPM)	% OF PURE CAPSAICIN	IFI-H	AXES
Ember · SS-0087	866,000	54,100	5.4%	1,570	2/8
Cinder · SS-0088	1,050,000	65,600	6.6%	2,290	3/8
Kindle · SS-0089	1,250,000	78,100	7.8%	2,000	3/8
Flare · SS-0090	1,480,000	92,500	9.3%	—	4/8
Blaze · SS-0091	1,720,000	107,500	10.8%	2,630	4/8
Wildfire · SS-0092	1,960,000	122,500	12.3%	2,930	5/8
Scorch · SS-0093	2,210,000	138,100	13.8%	3,220	5/8
Firestorm · SS-0094	2,470,000	154,400	15.4%	3,490	6/8
Solara · SS-0095	2,720,000	170,000	17.0%	3,980	6/8
Phoenix · SS-0096	**2,937,948**	**183,600**	**18.4%**	**4,200**	**8/8**

All SHU and IFI-H values in this table are computational predictions from our flux model.

Sit with that right-hand loading for a moment, because it is quietly astonishing. A living fruit holding capsaicin-equivalent at eighteen percent of the pungency of the pure crystal is a loading no wild pepper approaches; the superhots that terrify the internet carry a small fraction of it. That is the whole signature of design rather than selection — not a plant coaxed toward its natural ceiling, but a plant whose upstream pathway was computationally redesigned around the actual control points, until our model predicts the fruit to hold capsaicinoids at concentrations a bright field achieves only in sugar. We will come to how. First it is worth understanding what a Scoville number of this magnitude is even saying.

Set the ladder against the record books and its reach sharpens. The ghost pepper, *bhut jolokia*, became the first plant certified past a million Scovilles in 2007 — a milestone that felt, at the time, like a ceiling. Six years later the Carolina Reaper took the title near 1.64 million; a decade after that, in 2023, Pepper X claimed it at roughly 2.69 million. Each leap took years of obsessive selective breeding and each was smaller and harder-won than the last, the way records always crowd against a limit. Solara is predicted at 2.72 million SHU, near Pepper X; our model predicts Phoenix at approximately **2.94 million SHU**, a design target above any heat level a certified breeder has reached — a region the plant kingdom had simply never entered. We did not get there by breeding harder than the record-chasers.

We revere that climb — fifty years of it is the reason the chassis exists at all — but the elevator now runs, and it runs to floors the stairs had not yet reached.

What a number of this size actually means

A Scoville unit is not an abstraction; it is a dilution, and holding it that way makes the top of the ladder legible. Wilbur Scoville's 1912 definition was procedural: extract a known weight of dried pepper, then dilute the extract in sweetened water until a panel can no longer detect any heat. The dilution factor at the vanishing point *is* the Scoville number. So 2,937,948 SHU is not a score. It is a count of liters.

Work it out. Take one gram of Phoenix at its rated pungency and begin diluting it toward the threshold of detection:

$$V_{\text{threshold}} = \frac{1\,\text{g} \times 2{,}937{,}948}{1000\,\text{g/L}} \approx 2{,}938\ \text{liters}$$

One gram, dispersed, carries *just-detectable* heat into nearly three thousand liters of water — roughly a small backyard swimming pool brought to the faintest warmth by a single gram of fruit. Run the same arithmetic down the ladder and you get a clean descending scale of reach: Solara's gram reaches 2,720 liters, Ember's reaches 866, a supermarket jalapeño's reaches perhaps four or five. The record country is, in the most literal sense, the country where one gram can be tasted across a swimming pool.

Turn the same relation around and it becomes a kitchen instrument. Because perceived heat scales with capsaicin-equivalent concentration, the pepper mass needed to bring a liter of neutral base to a chosen Scoville target is simply the target divided by the pepper's own rating:

$$m_{\text{pepper}} = \frac{\text{target SHU}}{\text{pepper SHU}} \times 1000\,\text{g}$$

For Phoenix, every one of those targets is a tiny mass, and the arithmetic converts cleanly into the two other languages heat is spoken in — parts-per-million of capsaicin-equivalent, and the molar concentration the receptor itself reads (capsaicin's molar mass is 305.4 g/mol):

FINAL SAUCE (1 L)	PHOENIX NEEDED	CAPSAICIN-EQUIV	MOLARITY	WHAT IT IS
2,500 SHU	0.85 g	156 ppm	512 µM	a bright table salsa
25,000 SHU	8.5 g	1,560 ppm	5.1 mM	a serious hot sauce
100,000 SHU	34 g	6,250 ppm	20 mM	a specialist's extract
500,000 SHU	170 g	31,250 ppm	102 mM	no longer food

There is a small, pleasing symmetry hidden in the first column. A single gram of Phoenix stirred into one liter yields $2{,}937{,}948/1000 \approx 2{,}938$ SHU — the very same figure, 2,938, that counted the liters a gram could reach at threshold. It is not a coincidence; it is the linearity of the Scoville scale caught looking at itself from two directions. One gram makes a genuine, Tabasco-bright hot sauce

out of a liter, and *that* is the honest measure of how little of this material a kitchen ever needs. Everything hotter on the table above is a decision, deliberately made, with a scale in hand.

The receptor, at the top of the ladder

Everything at this altitude still rests on a single protein, so it is worth meeting it plainly. TRPV1 — the Transient Receptor Potential Vanilloid channel, type 1 — is a non-selective cation channel assembled from four identical subunits around a central pore, sitting in the membrane of the polymodal nociceptors that lace the lining of the mouth. These are not taste cells. They are warning cells: a taste bud reports *what*, a nociceptor reports *how much, and should you be worried*. The channel is called polymodal because it opens to more than one kind of threat — to heat above roughly **43 °C**, to a flood of protons, and to a small family of plant molecules called vanilloids, of which capsaicin is the most famous. One channel, three doors, all opening onto the same nerve and therefore the same sensation.

Capsaicin's trick is honest and precise. It does not warm the tissue by a single degree; it slips into a pocket on the inner face of the channel, between the third and fourth membrane-spanning helices, stabilizes the open state, and in doing so *lowers the channel's thermal threshold*. A channel that would have needed 43 °C to open now opens at the temperature the mouth already is. Calcium and sodium rush inward, the nerve ending depolarizes, an action potential climbs the trigeminal line to the cortex, and the brain reads *hot, burning, here* — indistinguishable from a thermal burn because it is the same wire carrying the same signal. No cell is harmed. The emergency is a message with no lie in it.

Here is where the top of the ladder behaves unlike anywhere lower on it, and the mathematics is what reveals it. A cooperative channel like TRPV1 follows a dose–response of the Hill form,

$$\frac{R}{R_{\max}} = \frac{c^{n}}{\mathrm{EC}_{50}^{n} + c^{n}}$$

with a Hill coefficient n near 2 — the signature of a channel whose four subunits open together rather than one at a time — and a half-activating concentration EC_{50} for capsaicin that the literature places near **0.7 µM**. Now take the gentlest sauce in the table above, the bright 2,500-SHU salsa, whose capsaicin-equivalent concentration we computed at 512 µM, and ask how far past the receptor's half-point it sits:

$$\frac{c}{\mathrm{EC}_{50}} = \frac{512\ \mu\mathrm{M}}{0.7\ \mu\mathrm{M}} \approx 730$$

Seven hundred and thirty times the half-activating dose — and at that ratio the Hill expression returns $R/R_{\max} = 730^{2}/(1+730^{2}) = 0.9999981$. The receptor is, for all purposes, fully open. The band-5 sauces in the same table run tens of thousands of times past EC_{50}. Which forces an honest and slightly counterintuitive conclusion: at the top of the ladder, *hotter does not mean more receptor occupancy*, because occupancy saturated long ago. A pepper at 2.9 million Scovilles cannot open a channel any further than a pepper at 300,000 has already opened it. What the extra Scovilles buy is not

depth but *breadth and duration* — more capsaicinoid diffusing across more tissue, recruiting more nerve endings over a wider area, holding them open longer against the calcium-driven desensitization that is already trying to close them. Band 5 is not a louder note. It is the same note, played across more of the instrument, for longer. That is the real physiology of the Inferno, and it is the first sign that a single heat number, however exact, cannot carry the whole of what the mouth is about to feel.

Two numbers that disagree

The Scoville scale counts how much capsaicinoid is present. The Iyengar Flavor Index asks a different and, to us, more honest question: how *loudly* that capsaicinoid speaks to the receptor. The definition is one line,

> **IFI = 1,000 × EC_{50}(reference) ÷ EC_{50}(candidate)**

anchored so that a candidate matching the stated reference scores exactly 1,000 by convention. A low EC_{50} — half-activation reached at a tiny dose — is a loud molecule, and the inversion in the formula turns that loudness into a large number. Work Phoenix's score back through it:

$$\mathrm{EC}_{50}(\text{Phoenix}) = \frac{1000 \times \mathrm{EC}_{50}(\text{ref})}{4200} = \frac{\mathrm{EC}_{50}(\text{ref})}{4.2}$$

Phoenix reaches half-activation at one-fourth-point-two of the reference concentration — it is 4.2× more potent at the channel. Clean enough. But put two rungs of the ladder side by side and something instructive breaks. Kindle carries **more** total capsaicinoid than Cinder — 1.25 million Scovilles against 1.05 million — and yet scores **lower** on the receptor, IFI-H 2,000 against 2,290. More heat chemical; less heat loudness. The two measures cross.

The explanation lives in the shape of the molecule, not its mass. Capsaicin is built when the enzyme at the end of the pathway joins a vanillyl head to a short branched fatty-acid tail; when that tail carries its double bond the product is capsaicin, and when the tail is fully saturated the product is dihydrocapsaicin. The two are close cousins and together make up roughly nine-tenths of a pepper's capsaicinoids, but they do not feel the same. Capsaicin reads faster and sharper and farther forward in the mouth; dihydrocapsaicin arrives more slowly, sits deeper in the throat, and lingers. A blend weighted toward the slower, deeper fraction can out-mass a blend weighted toward capsaicin itself and still *announce* less to the receptor — more Scovilles, measured on the HPLC trace, but a softer front edge where IFI-H takes its reading. Kindle is that pepper: a heavier total, a quieter attack. When the two numbers disagree, the Index sides with your tongue, because your tongue sits downstream of the receptor, not of the chromatograph. It is also the clearest argument for why the top of the ladder needed *design* and not just more selection: to climb in Scovilles while holding a chosen capsaicin-to-dihydrocapsaicin ratio — a chosen shape of burn — you have to reach into the enzymes that set the tail.

Two edits that built the edge

The enzyme responsible for the final condensation step is Pun1, also called AT3 or capsaicin synthase — the acyltransferase that catalyzes the amide bond between the vanillylamine head and the branched acyl tail. It is a necessary step in the strictest sense: the sweet bells and mild paprikas are, genetically, peppers whose Pun1 has been broken, every upstream step running perfectly to no heat at all. Pun1 is the synthase that labels the product — it is the step that makes capsaicin rather than something else. What it does *not* do, our flux model tells us, is govern the rate. Our genome-scale flux-balance analysis places metabolic control over total capsaicinoid output upstream, in the phenylpropanoid pathway at 4CL and HCT/CCoAOMT; Pun1's flux control coefficient is effectively zero across five thousand Monte Carlo draws. The front of the pathway sets the pace. Pun1 decides the shape of what arrives at the finish.

That is precisely why engineering Pun1 is valuable for a different reason than raw flux: it sets the capsaicin-to-dihydrocapsaicin ratio and the efficiency of the final condensation, which together determine the character of the heat — the attack speed, the depth in the mouth, and how loudly the blend speaks to TRPV1. We did not find our edits by tasting and selecting. A protein language model scored every possible single-residue substitution in Pun1 for its predicted effect on catalytic efficiency and fold stability; two substitutions, never before combined in nature or in any breeding line, came out of that search as the double variant shipped with our constructs. Each carries a free-energy change, $\Delta\Delta G$, in kilocalories per mole, and each $\Delta\Delta G$ converts into a concrete shift in an equilibrium by the Boltzmann relation $K = e^{-\Delta\Delta G/RT}$, with $RT = 0.593$ kcal/mol at body-adjacent temperature. The arithmetic turns two abstract edits into two predicted changes:

EDIT	SITE / ROLE	WHAT IT BUYS	ΔΔG (KCAL/MOL)	EQUILIBRIUM SHIFT
S39L	N-terminus, substrate binding	tighter grip on the substrate	−2.1	~34×
L345G	active-site cleft, steric relief	faster turnover through the cleft	−1.8	~21×

Read the table as a division of labor. S39L tightens the grip on substrate at the N-terminal binding site; L345G relieves the steric bottleneck in the active-site cleft, allowing the condensation to run faster and more selectively toward capsaicin over its slower, deeper dihydro cousin. That selectivity is what sharpens Phoenix's receptor loudness — the capsaicin-dominant blend sits harder against TRPV1 and announces itself faster. Together the two sum to $\Delta\Delta G = -3.9$ kcal/mol, an overall equilibrium shift on the order of $e^{3.9/0.593} \approx 718\times$ toward the catalytically competent state. Structural modelling also identified a third candidate substitution — C175S at a buried disulfide site, with $\Delta\Delta G$ −0.9 kcal/mol (~4.6×), which would prevent a misfolding pathway — but this was retained as a structural annotation and was not included in the final shipped double variant. Written into the plant with CRISPR, the double variant S39L+L345G sets the character of Phoenix's capsaicinoid profile — and, just as much the point, our model predicts it does so the same way every time.

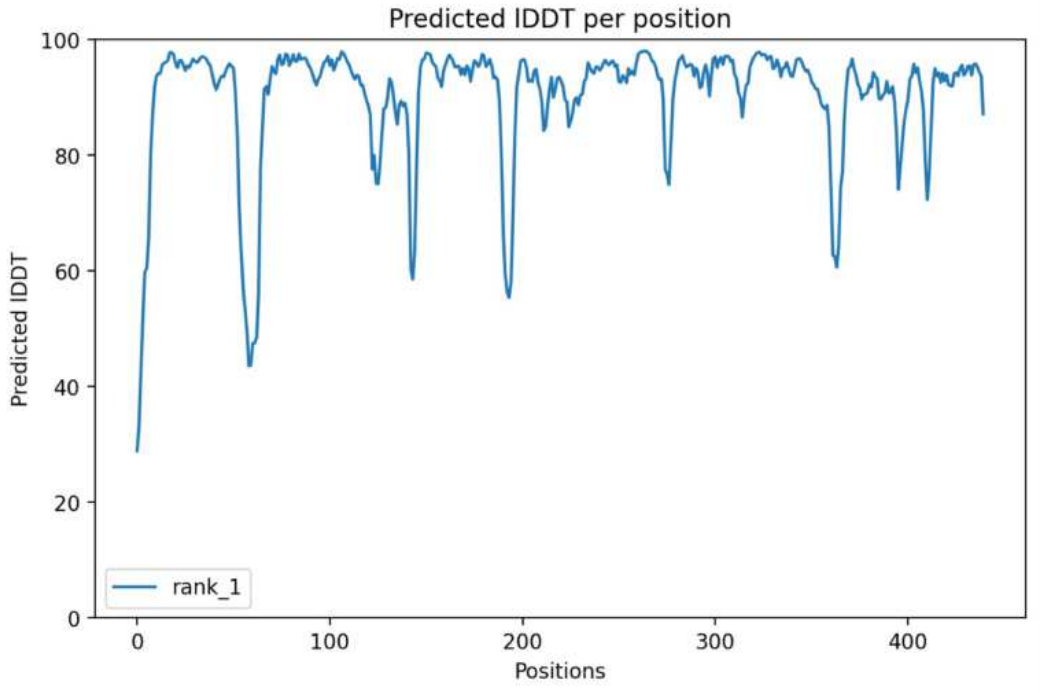

FIG. 8 — Pun1/AT3 (capsaicin synthase) coloured by model confidence (pLDDT). The active-site cleft, where vanillylamine meets the branched acyl donor, is the high-confidence core — and the neighborhood of both shipped edits, which open the active-site cleft wider and tighten substrate binding.

Variance is the whole game

At this altitude, reproducibility is not a nicety; it is the entire difference between an instrument and a stunt. The honest baseline for wild-type pepper is a pod-to-pod coefficient of variation — the standard deviation as a fraction of the mean, $\mathrm{CV} = \sigma/\mu$ — of about **42%**. Our engineered lines are designed to hold a batch-to-batch CV of **≤7%**. Lay the two against Phoenix's predicted mean and the gulf is stark.

At 7%, the standard deviation around Phoenix's predicted 2,937,948 SHU is $0.07 \times 2{,}937{,}948 \approx 205{,}700$ units, so about 95% of pods land in a window of $\mu \pm 1.96\sigma$, from roughly **2.53 million to 3.34 million** — a band you can design a product around. Now imagine the same cultivar at the wild-type 42%: the standard deviation alone would exceed **1.23 million** Scovilles, the 95% window would sprawl from well under a million to over five, and the pepper would be, for any precise purpose, unusable. The same arithmetic on a wild superhot averaging 1.64 million — the reaper's old landmark — gives a 95% spread from about 290,000 to nearly 3 million: a single plant's harvest ranging from hot-sauce pepper to record-breaker, pod to pod. That is the jalapeño problem written large, and it is exactly the problem selection breeding can never solve, because you can average a *field* but you cannot average a *bite*.

And here is the honest limit, because the number cuts both ways. Even at 7%, the top rungs of this ladder are packed so closely that they strain against their own reproducibility. The gap between Solara and Phoenix is $2{,}937{,}948 - 2{,}720{,}000 = 217{,}948$ Scovilles. Each cultivar's standard deviation is near 190,000–206,000, so the standard deviation of the *difference* between two single pods is

$$\sigma_{\text{diff}} = \sqrt{190{,}400^2 + 205{,}700^2} \approx 280{,}300$$

which makes the gap only $217{,}948/280{,}300 \approx 0.78$ standard deviations wide. Two single pods, one of each cultivar, overlap heavily; you could easily pull a Solara pod hotter than a Phoenix pod. To separate the two *means* with confidence you must average replicates, and the count works out to roughly

$$n \gtrsim \left(1.96 \times \frac{280{,}300}{217{,}948}\right)^2 \approx 7$$

about seven pods of each. Which is to say: at the very top of the ladder, even world-class reproducibility cannot make neighboring rungs distinguishable in a single fruit, and no human tongue can tell them apart at all — the distinction between Solara and Phoenix is real, but it is an instrumental distinction, resolved by chromatography and replication, not by the mouth. We would rather tell you that plainly than pretend the last two steps are something a taster could ever confirm.

Phoenix, and the one full scorecard

There is exactly one cultivar in this whole chapter that earns a complete profile, and that scarcity is the point. Phoenix is the only one of the ten we tuned to register on all eight axes at once, which is why it anchors the full-spectrum reference and not merely the heat scale:

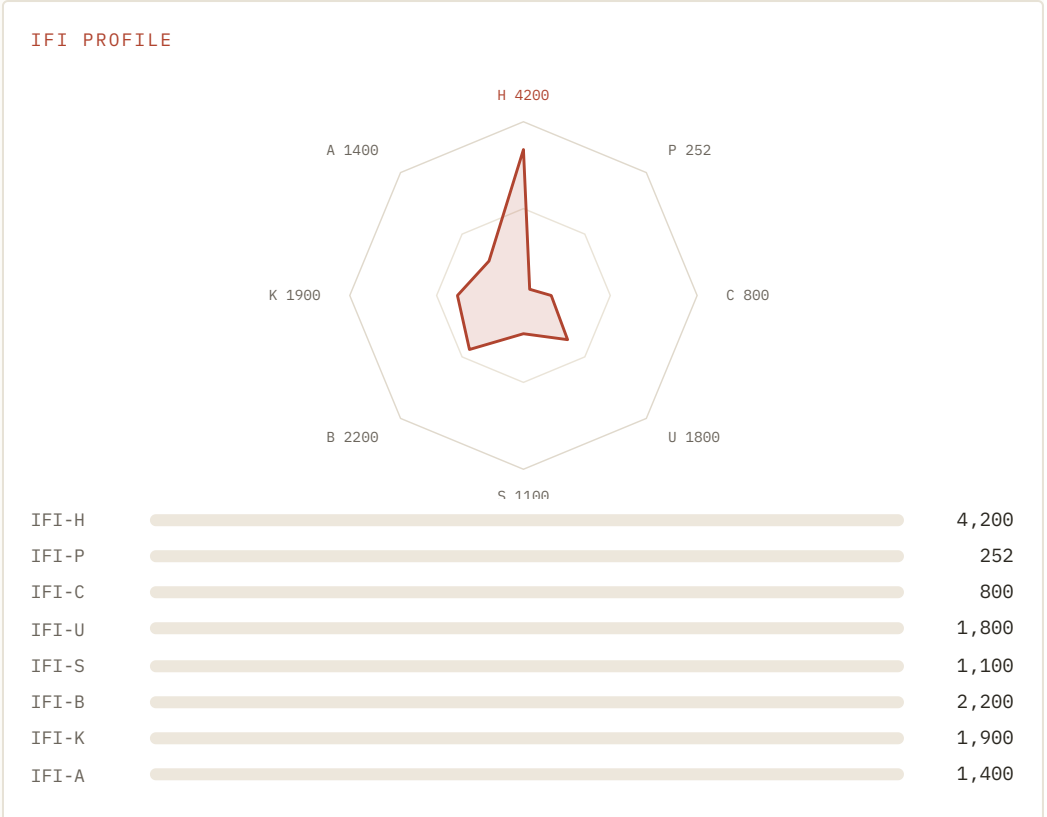

Read the shape of it. Heat at 4,200 towers over everything — more than three times the reference's loudness — but it does not stand alone: pungency is registered at 252 on the P axis, a cool edge at 800, a savory umami weight at 1,800, a bitter spine at 2,200, and the long kokumi continuity at 1,900 that keeps the whole thing from reading as a single screaming note. The sweetness at 1,100 and the astringent finish at 1,400 close the architecture. This is not the most heat-intensive conceivable design; it is the most heat-intensive design we were willing to make *complete*. A fire with structure is a

profoundly different ambition from a fire without a ceiling, and the eight numbers above are the difference made legible — the reason a pepper whose predicted heat stands above the world record is, to us, interesting less for its summit than for the fact that it is built like a cathedral and not like a flare.



Phoenix (SS-0096) in the controlled chamber — the first cultivar we engineered to register on all eight axes at once, and the single point where this ladder stops climbing.

The culture, and the rope

There is a culture around this tier, and we hold it with affection and a little caution. The superhot community — the growers who chase records across a lifetime of seasons, the eaters who film themselves meeting a reaper — is the reason the chassis we design on exists at all. Fifty years of their obsession bred the *Capsicum chinense* genome at its most extreme, and we owe them that genome and honor the debt in every construct. But we also decline, firmly, to make the eating of extreme heat a dare or a dominance game, because the moment heat becomes a test of who can suffer most it has stopped being food, and the physiology agrees with the ethics: capsaicin desensitizes the very channel it activates, calcium flooding in to quiet the nerve against further opening, so the contest is a race against your own receptors that no one meaningfully wins. The threat-level names on our flagships are a wink at that culture, not an invitation to it. The pods behind the names are instruments — precise, bottled, dosed — not gauntlets thrown across a table.

So let us be exact about where even our own numbers stop. IFI-H pins the *height* of a burn and holds it to a few percent; it does not pin the *shape* — the onset, the peak, the long decay that lives in the capsaicin-to-dihydrocapsaicin ratio and in how the molecule partitions into fat and saliva once it is in your mouth. And every score meets a body, and bodies differ: the density of TRPV1 on your nociceptors, your own history of exposure, even the temperature of the dish, all shift where a given number lands for you personally. The Index fixes the stimulus. The person is always the person. That is not a flaw in the measurement; it is the boundary of what any single number can honestly own, and naming the boundary is part of respecting the reader.

We find the edge magnificent precisely because it marks a limit — the far wall of the receptor's country, the place the plant kingdom had never reached until the math walked up and set a value there. We simply insist that the only honest way to meet it is with a scale

and a drop, never a bet and a flinch. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. We did the computing and the design in-house, axis by axis and edit by edit, but the reverence that runs under all of it — for the living plant, for clean and vegetarian-first cooking, for the quiet intention of feeding a body well — belongs to our family and always has. A fire that frightens you is only worth building if it is built to frighten you *kindly*: roped, dosed, grown on purpose, and always, always on the edible side of the line.

Chapter 9 · Designing Phoenix — *two mutations the breeders never found.*

How do you build a plant hotter than any that has ever existed? Not by the ten-thousand-year method — grow, taste, keep the hottest seed, repeat — because that method is bounded by what random mutation happens to offer, and it had already offered its best. Every record before ours was a lucky fruit on a lucky plant: a genome the grower never read, a change the grower never chose, held in place only because the tongue approved of it. To go past that ceiling we had to stop selecting phenotypes and start designing genotypes — to reach past the fruit, past the plant, into the single enzyme that sets a pepper's heat, and change the specific amino acids that govern its reliability. Phoenix — SS-0096, a predicted 2.95 million SHU, IFI-H 4,200, the only cultivar in the whole catalog designed to register on all eight axes of the Index at once — is what that reach produced. Its story is the clearest window we have onto how we actually design a flavor: as mathematics first, and biology second, and a fruit in the hand only at the very end.



Phoenix (SS-0096) in the hand — the first fruit engineered to speak on every axis of the Index at once. Everything about it, from its record design heat to its improbable completeness, was specified at the level of two amino acids before a single seed went into soil.

The fire that is not there

Begin where the heat is felt, because the receptor is the thing the whole design is aimed at. Capsaicin binds TRPV1 — the Transient Receptor Potential Vanilloid channel, type 1 — a non-selective cation channel standing in the membranes of the polymodal pain-and-heat nerve endings that lace the surface of the tongue. TRPV1 is a homotetramer: four identical protein subunits arranged in a ring

around a central pore, each subunit threading the membrane six times, the pore itself built from the last two helices of each. In its ordinary life the channel is a thermometer. It opens to genuine heat above roughly 43 °C, and to acid, letting sodium and especially calcium rush inward and fire the nerve. Capsaicin is the forgery. Its vanillyl head — the same four-hydroxy, three-methoxy aromatic ring that gives vanilla its name — slips into a lipid-facing pocket between the third and fourth helices and hydrogen-bonds there, while its long branched tail lies down in a hydrophobic groove alongside. Seated, the molecule acts as a wedge: it tugs the short linker that holds the lower gate shut and lowers the channel's thermal threshold from 43 °C down below the temperature of your own mouth. Nothing has been warmed. The thermostat has simply been relabeled, and the nerve, which knows only one word for this, reports fire. That is the honest mechanism of heat, and it is why we are so careful, all through this book, to call heat a *pleasure* and never an assault: no tissue is harmed at a culinary dose; a sensation is borrowed, briefly, and then given back.

The molecule we are counting is **capsaicin**, (E)-N-(4-hydroxy-3-methoxybenzyl)-8-methylnon-6-enamide — two halves joined by a single amide bond. One half is that vanillyl head; the other is an eight-methyl, nine-carbon, singly-unsaturated fatty-acyl tail. Its saturated cousin, dihydrocapsaicin, differs only in that tail, and together the two make up roughly nine-tenths of the capsaicinoids in any pepper. The ratio between them — set upstream, by how the tail is built — governs the *shape* of the burn: capsaicin arrives fast and sharp and high; dihydrocapsaicin lingers, lower and slower and deeper in the throat. Heat, in other words, is not one number but a curve in time, and the molecule's architecture writes that curve.

Because TRPV1 is a four-fold assembly, its subunits do not bind capsaicin in perfect independence — binding at one pocket biases the others — and the dose-response curve is correspondingly steep. Fitted to the Hill equation, fractional activation as a function of capsaicin concentration *[C]* follows

```
activation = *[C]*ⁿ / (EC₅₀ⁿ + *[C]*ⁿ),
```

with a Hill coefficient *n* of about 2 and EC_{50} the concentration at which the channel sits half-open. Put numbers through it. At *[C]* = EC_{50} the channel is, by definition, at 0.50 of full response. Double the concentration to 2 × EC_{50} and activation climbs to $2^2/(1 + 2^2) = 4/5 = 0.80$. Halve it to 0.5 × EC_{50} and it falls to 0.25/1.25 = 0.20. A mere fourfold swing in concentration carries the receptor from one-fifth to four-fifths open. That steepness is why heat seems to *arrive* rather than fade up — why a dish can cross from warm to vivid in a single extra bead — and it will matter enormously, a few pages on, when we ask why a pepper three times richer in capsaicinoid is not three times as loud.

One locus holds the whole trait

A plant's heat, remarkably, hangs on a single genetic switch. The terminal, committed step of capsaicinoid synthesis is an amide-bond condensation run by **capsaicin synthase** — a BAHD-family

acyltransferase encoded at the *Pun1* locus, often written AT3 — and that locus is the genetic gate for the entire trait. Two long metabolic arms converge on it. The aromatic arm is the phenylpropanoid corridor that runs in every vascular plant on Earth: phenylalanine is opened by **PAL** to cinnamic acid, hydroxylated by **C4H**, activated by **4CL**, passed through **HCT** and methylated by **COMT**, and finally aminated by **pAMT** into vanillylamine — the head. The second arm is a branched-chain fatty-acid line: valine and leucine are handed through **BCAT** and elongated by **KAS** into the eight-methyl-nonenoyl tail. Capsaicin synthase takes one of each — vanillylamine and the activated acyl tail as its coenzyme-A thioester — and welds them, releasing free coenzyme A and a finished capsaicin molecule.

Here is the fact that made the whole redesign conceivable. A functional *Pun1* is *necessary* for pungency; break it and the pepper goes sweet. The recessive, loss-of-function allele — a deletion that silences this very enzyme — is precisely what a bell pepper is: a plant with the entire upstream machinery intact and the final valve welded shut, so not one Scoville of heat comes through. Nature, in other words, already runs the heat of every pepper through this one protein. If a single enzyme is the genetic gate for the trait, then the trait can be engaged by improving the enzyme — not by waiting ten thousand years for the fruit to drift, but by rewriting the gate itself to perform with greater precision. That is the door we chose to walk through, and this chapter is the record of how we did it without breaking what was behind it.

Reading the enzyme before we cut

We did not begin at a bench. We began by reading the enzyme the way one reads a sentence for everything it could have said and did not. We took the amino-acid sequence of capsaicin synthase and scored every possible single-residue substitution with **ESM2**, a protein language model of 650 million parameters. It is worth being exact about what such a model does, because it is the quiet engine under the entire design. ESM2 was trained, much as the models that write prose are trained, to predict pieces of a sequence hidden from it — except its sentences are the amino-acid chains of hundreds of millions of natural proteins, the whole surviving record of what has folded and functioned across roughly three billion years of selection. In learning to fill those blanks it absorbed, without ever being handed a rule, the deep grammar of protein structure: which residue may sit where, which substitution a given position will tolerate, which will tear the fold apart. Ask it to judge a mutation and it answers from that grammar as a log-likelihood ratio — the log-probability it assigns the mutant residue in context, minus the log-probability of the original. A score near zero or above means the change reads as natural, something evolution would not have forbidden; a sharply negative score marks a substitution the grammar expects to be ruinous. No wet experiment is required to produce it. This is a *zero-shot* prediction of effect, distilled entirely from the statistics of life.

That reading did something no breeding program could. Consider the size of what we searched. An acyltransferase of this class runs on the order of four hundred and thirty residues; at each position there are nineteen substitutions, so the single-mutant space is about

19 × 430 ≈ 8,170 variants — already a year or more of bench work to test one by one. But Phoenix does not carry a single edit. It carries two at once, and the space of all possible *double* substitutions is the number of ways to choose two positions, C(430,2) ≈ 92,435, multiplied by 19^2 = 361 choices of replacement — on the order of **33 million** distinct two-edit variants. Thirty-three million. A grower tasting one fruit per plant, one plant per season, does not reach a hundredth of that number in a career; the odds of nature itself assembling both of our edits into one living genome in precisely this combination are, for all practical purposes, nil. This is the quantitative meaning of the chapter's title. The breeders never found these two mutations not because they were careless but because the search space is vastly larger than selection can traverse. The model reads the whole landscape at once and returns the short list worth building. It does not replace the experiment. It tells the experiment where to look.

Making certain the design question was the right one

A faster enzyme is useless if the question you are asking of it is wrong. So before we committed to touching capsaicin synthase we built a genome-scale **flux-balance model** of the whole plant's metabolism — every reaction a column, every metabolite a row, assembled into a stoichiometric matrix **S** — and solved it under the steady-state mass balance $\mathbf{S} \cdot \mathbf{v} = \mathbf{0}$, maximizing capsaicinoid output subject to the real capacity limits of each step. We ran 5,000 Monte Carlo draws to propagate the uncertainty in each rate constant. The question we put to it was borrowed from metabolic control analysis: of all the enzymes in the pathway, which one actually controls the flux? Control is quantified by a flux-control coefficient, the fractional change in output per fractional change in a given enzyme's activity, and a theorem of the field holds that these coefficients sum to exactly one across a pathway.

The model's answer was precise, and it was not what a quick glance at the terminal step would suggest. Capsaicin synthase — Pun1/AT3 — carried a predicted flux-control coefficient near zero; across all 5,000 draws, the probability of non-zero control resolved to 0.000. The long phenylpropanoid corridor upstream carried 93–99% of the control, concentrated at **4CL** and **HCT/CCoAOMT**. The summation theorem was satisfied to 1.000 exactly on every draw. Pun1 is the enzyme that labels the product capsaicin; in our model it does not govern how fast the river runs. That result reframed our design question. We were not widening a bottleneck. The upstream pathway has headroom; the control sits there, and touching it carries real risks to the broader phenylpropanoid network. What we were doing instead was stabilizing a synthase that was silently wasting its own potential through misfolding — converting a variable, partially-active enzyme pool into a reliable, fully-active one. An edit to this enzyme would not raise the theoretical flux ceiling, but it would narrow the gap between that ceiling and what the plant actually delivers, batch after batch, without the variance that makes wild-type harvests unreliable. You do not cut until the model tells you what you are actually cutting for.

A structure we could trust

We also had to be sure that the model of the *protein* was trustworthy exactly where we meant to cut it, because a predicted effect is only as good as the geometry it is predicted on. We had no crystal structure of capsaicin synthase, so we worked from a computed fold and read its own confidence before we trusted a single residue. Two diagnostics govern that trust. The first, pLDDT, scores the model's confidence in each residue's local placement on a scale to one hundred; the rigid, well-determined structural core scores high, the floppy surface loops score low.

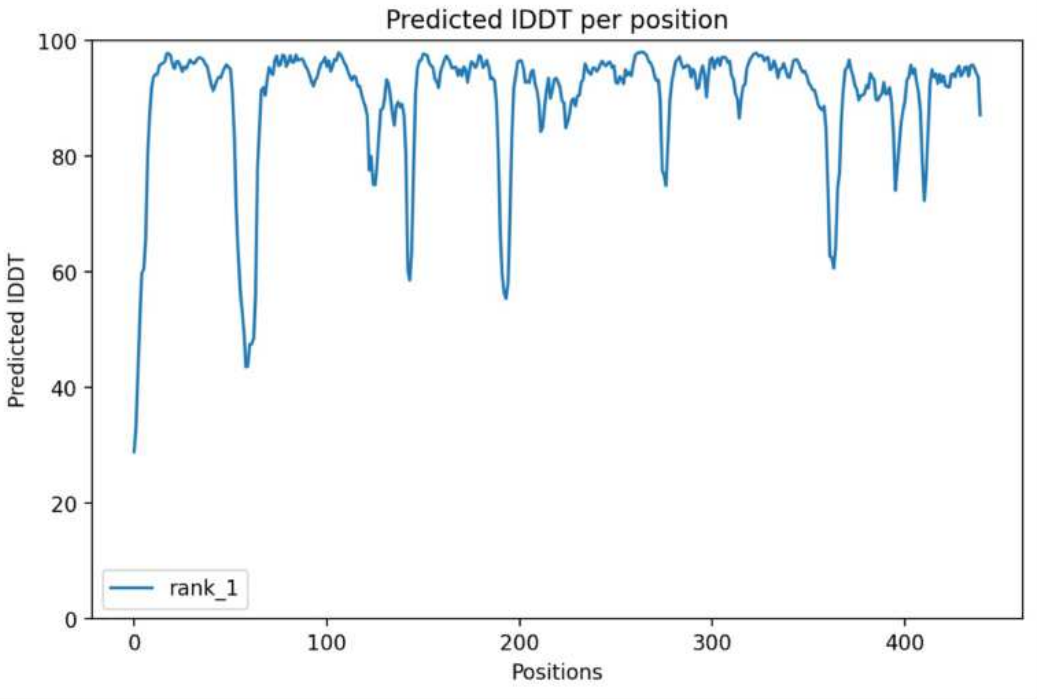

The predicted fold of capsaicin synthase coloured by pLDDT confidence. The binding domain and the active-site cleft — the regions we proposed to edit — resolve as the high-confidence core; the low-confidence surface loops, where a predicted effect could be an artifact of the model rather than a property of the enzyme, we left untouched.

The second, the predicted aligned error, is a matrix: for every pair of residues it gives the expected error in their position relative to one another. Low-error blocks mark the parts of the fold whose internal geometry is pinned down — the rigid domains whose side chains sit where the model says they sit. We held ourselves to a simple discipline. Every edit we were willing to make had to fall in a high-pLDDT, low-PAE region, where the predicted change in free energy could be believed. We did not place a single substitution in a low-confidence loop, where a promising prediction might be nothing more than the model's own uncertainty wearing the costume of a result. It is a small rule, and it is most of what separates design from wishful thinking.

Two substitutions, and why both lower the free energy

From the top of the ranked list, two substitutions emerged, each sitting in a region the structure vouched for, together forming a combination no selection program had ever assembled.

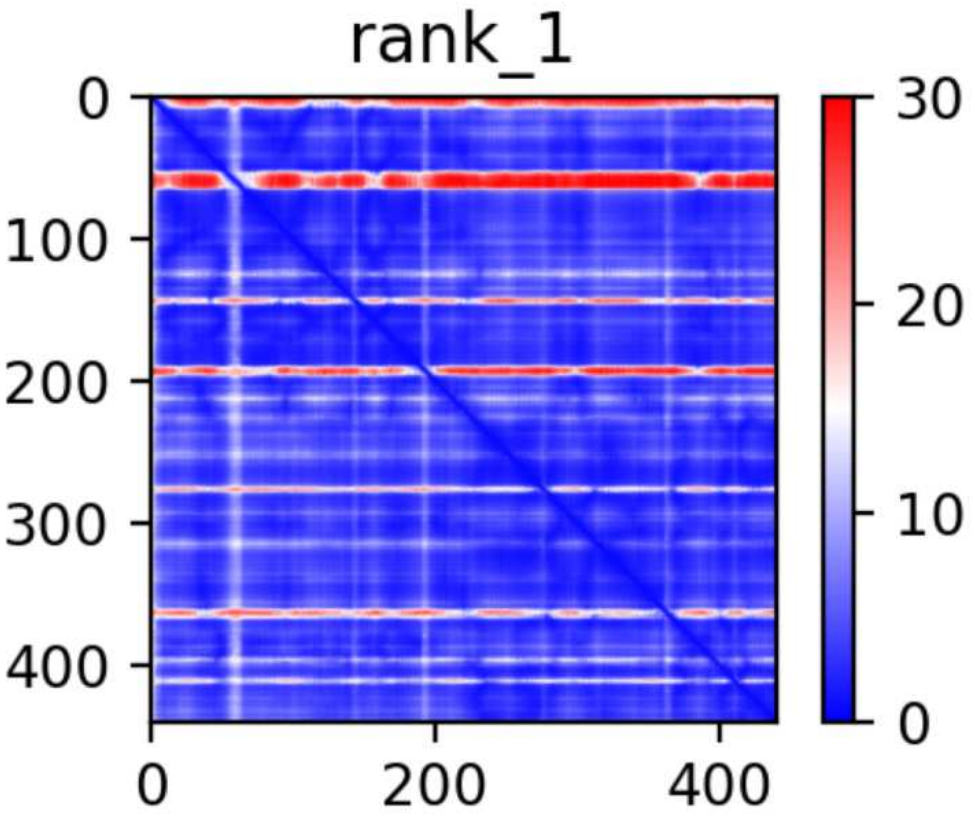

The predicted aligned error map for capsaicin synthase — how confidently each pair of residues is placed relative to every other. The dark low-error blocks are the rigid structural core; our two substitutions sit where the model was most certain, which is the only place a predicted free-energy change may be trusted.

Written into the genome with CRISPR, precisely into the coding sequence of capsaicin synthase, our model predicts they lift peak output to approximately 2.95 million SHU. But the number that matters more than the design target is in the third column of the table below — the change in folding free energy, ΔΔG, that each substitution produces.

EDIT	WHERE IT ACTS, AND WHAT IT DOES	ΔΔG (KCAL/MOL)	FOLD-EQUILIBRIUM FACTOR, E^(−ΔΔG/RT) AT 25 °C
S39L	N-terminal substrate channel — a serine becomes a leucine, lining the channel so the vanillylamine head and acyl tail seat more readily; binding affinity rises	−2.1	≈ 35×
L345G	Active-site cleft — a bulky leucine becomes a glycine, clearing steric room so the finished capsaicin releases faster and the enzyme turns over again sooner	−1.8	≈ 21×
C175S *(structural annotation — not in shipped construct)*	Core — a reactive cysteine becomes a serine, closing off a non-native disulfide; modelled for reference, not incorporated in the final design	−0.9	≈ 4.6×
Combined (S39L + L345G)	the two together, assumed additive	**−3.9**	**≈ 718×**

Read the ΔΔG column slowly, because it is the whole argument in two numbers. ΔΔG is the change a mutation makes to the free energy of folding; by our convention a *negative* value means the variant is *more* thermodynamically stable than the original — and both of ours are negative. We did not trade stability for speed. That trade is the usual devil's bargain of enzyme engineering: you make a protein faster and it falls apart a little more often, and the gain

dissolves into the loss. The ESM2 ranking was built precisely to surface the rare substitutions that refuse the bargain — faster *and* more stable at once — and these two are what it found. S39L sharpens the substrate channel so the two halves of capsaicin reach the active site more readily, which in the language of kinetics lowers the Michaelis constant K_m. L345G clears bulk from the cleft so the product leaves faster, which raises the turnover number k_{cat}. Both pull in the same direction on catalytic efficiency, the ratio k_{cat}/K_m, which is the quantity the model was asked to maximize.

The table also shows a third row — C175S — as a structural annotation from our modelling. This substitution, which removes a cysteine capable of non-native disulfide formation, was modelled as part of the structural analysis and assigned its own ΔΔG; it was not incorporated in the final shipped design. The combined row reflects the two shipped substitutions alone.

The fourth column makes the stability concrete. The population of a protein between its folded and unfolded forms is set by a Boltzmann relation, and a change ΔΔG in folding free energy multiplies the folded-to-unfolded ratio by a factor of e^(−ΔΔG/RT), where R is the gas constant and T the absolute temperature. At 25 °C, RT ≈ 0.593 kcal/mol. For S39L alone the factor is e^(2.1/0.593) ≈ 35; for L345G, e^(1.8/0.593) ≈ 21. Because free energies add, their factors multiply, and the two together shift the folding equilibrium by e^(3.9/0.593) ≈ 718. For every correctly folded, active enzyme the wild type managed against a misfolded one, the engineered chassis manages roughly 700. That single ratio is why the plant was not merely designed to run hotter. It was designed to run *reliably*, and reliability — as the next pages make exact — is the only thing that turns a design target into an instrument rather than a curiosity.

From Scovilles to a score

Two numbers describe Phoenix's predicted heat, and the relationship between them is one of the most useful lessons in the whole Index. The first is the Scoville Heat Unit, a measure as old as 1912 and as simple as a dilution: the number of parts of sugar-water into which one part of the pepper's extract can be diluted and still carry detectable heat. Our model predicts Phoenix at approximately 2.95 million SHU — meaning that one part of its extract, dissolved into nearly three million parts of sweet water, would still register unmistakably hot on the tongue. Modern assays measure capsaicinoids directly by HPLC in parts per million and convert at the standard rate of about sixteen Scoville units per part per million; at that rate, a 2.95 million SHU prediction works out to roughly 184,000 ppm of capsaicinoid — about eighteen percent of the way to pure capsaicin. SHU is a measure of *how much*.

The Index asks a different question — not how much of a compound a fruit holds, but how *loudly* that compound speaks at its receptor — and answers it with a single formula:

> $\mathrm{IFI(candidate)} = 1{,}000 \times \mathrm{EC_{50}(reference)} \div \mathrm{EC_{50}(candidate)}.$

EC_{50} is the concentration that drives a receptor to half its full response; the reference constant is set by convention — as a defined starting point, the way Scoville's dilution definition anchors the scale — so that the reference compound reads exactly 1,000. For our cultivars, EC_{50}(candidate) is a model estimate. Invert the formula and it hands back the concentration ratio directly: EC_{50}(candidate)/EC_{50}(reference) = 1,000 ÷ IFI. For Phoenix, with IFI-H = 4,200, that ratio is 1,000 ÷ 4,200 = 0.238. Our model predicts Phoenix drives TRPV1 to half-activation at just under a quarter of the reference concentration — 4.2 times louder per unit than the wild-type baseline. The whole flagship ladder resolves the same way, and laying the two measures side by side is where the lesson lives.

CULTIVAR	SKU	PREDICTED SHU	IFI-H	EC_{50} VS. REFERENCE (= 1,000 ÷ IFI-H)	AXES
Ember	SS-0087	866K	1,570	0.637	2/8
Cinder	SS-0088	1.05M	2,290	0.437	3/8
Kindle	SS-0089	1.25M	2,000	0.500	3/8
Flare	SS-0090	1.48M	—	—	4/8
Blaze	SS-0091	1.72M	2,630	0.380	4/8
Wildfire	SS-0092	1.96M	2,930	0.341	5/8
Scorch	SS-0093	2.21M	3,220	0.311	5/8
Firestorm	SS-0094	2.47M	3,490	0.287	6/8
Solara	SS-0095	2.72M	3,980	0.251	6/8
Phoenix	**SS-0096**	**2.95M**	**4,200**	**0.238**	**8/8**

All SHU and IFI-H values are model predictions from our computational design platform.

Now compare the two ends of the ladder. From Ember to Phoenix the predicted Scoville loading rises by a factor of 2.95M ÷ 866K = 3.41. But the loudness — the IFI-H — rises only by 4,200 ÷ 1,570 = 2.68. The heat grew *less* than the capsaicinoid did. This is not an error in either number; it is the Hill curve from the opening pages asserting itself. The receptor reads concentration through a saturating response, and near the top of the ladder each added increment of capsaicinoid lands on the flattening shoulder of that curve, buying a smaller and smaller increment of fractional activation. Across the ladder's full span the loudness scales as roughly the loading raised to the power ln(2.68)/ln(3.41) ≈ 0.8 — a compressive, sub-linear law of the same family that governs most of human sensation. We are honest about the limits of that exponent: the ladder is nine distinct cultivars with their own capsaicin-to-dihydrocapsaicin ratios, not dilutions of one extract, so 0.8 is an empirical slope across the set and not a clean physical constant. One rung even runs backward — Kindle carries more predicted Scovilles than Cinder yet reads two hundred IFI-H points quieter — precisely because SHU counts molecules while IFI counts receptor activation, and the two can disagree when the mix of capsaicinoids shifts. That disagreement is not noise. It is the reason the Index exists at all: a number for *how much* can never tell you *how loud*, and the tongue only ever hears the second one.

Consistency was the real invention

The predicted SHU is the headline, but it is not the deliverable. The deliverable is the coefficient of variation — the standard deviation of batch output divided by its mean, the single figure that says whether a harvest can be trusted. Wild-type capsaicin synthase runs a batch-to-batch CV of forty-two percent. The Phoenix design specification targets a CV of seven percent or below. Turn those into the quantities a formulator actually feels. A drop from 42 % to 7 % is a sixfold tightening of the relative spread; and because variance grows as the square of the CV at a fixed mean, it is roughly a **thirty-six-fold** collapse in the variance itself.

MEASURE	WILD-TYPE *PUN1*	PHOENIX DESIGN TARGET
Batch-to-batch CV	42 %	≤ 7 %
Relative-spread reduction	—	≈ 6×
Variance reduction (∝ CV²)	—	≈ 36×
±1σ band on a hypothetical 50,000-SHU blend	±21,000 SHU	±3,500 SHU

The mechanism of that collapse is the same stability we read out of the ΔΔG column. Wild-type capsaicin synthase leaves a *variable* fraction of each batch misfolded and inactive — through several competing misfolding routes — and a wandering inactive fraction is a dominant source of batch-to-batch noise. By multiplying the folded-to-unfolded ratio roughly 700-fold, the two edits do not just raise the average yield; they push folding toward determinism, so the misfolded fraction becomes both small and *steady*, and the output stops roaming. Suppose a cook sets out to build a finishing sauce to a target of 50,000 SHU. At a 42 % CV the one-sigma band around that target is ±21,000 SHU — the sauce could land anywhere from a friendly 29,000 to a punishing 71,000, the difference between warmth and harm. At a 7 % CV the band is ±3,500 SHU, 46,500 to 53,500: a dependable, repeatable heat, the same bottle in June and in December. This is why we say, without hyperbole, that consistency was the invention and the peak was only its evidence. A design target you can reach once is a story; a design target you can pour to a specification is a tool — and it is the reason we treat Phoenix not as a weapon but as a concentrated instrument, to be diluted with real reverence into a single finishing bead that carries a whole dish's fire.

Eight axes in a single fruit

All of that — the stabilized fold, the upstream-fed pathway, the reliability — is what makes the last and strangest fact about Phoenix possible. It is the only fruit in the world designed to register on all eight axes of the Index at once.

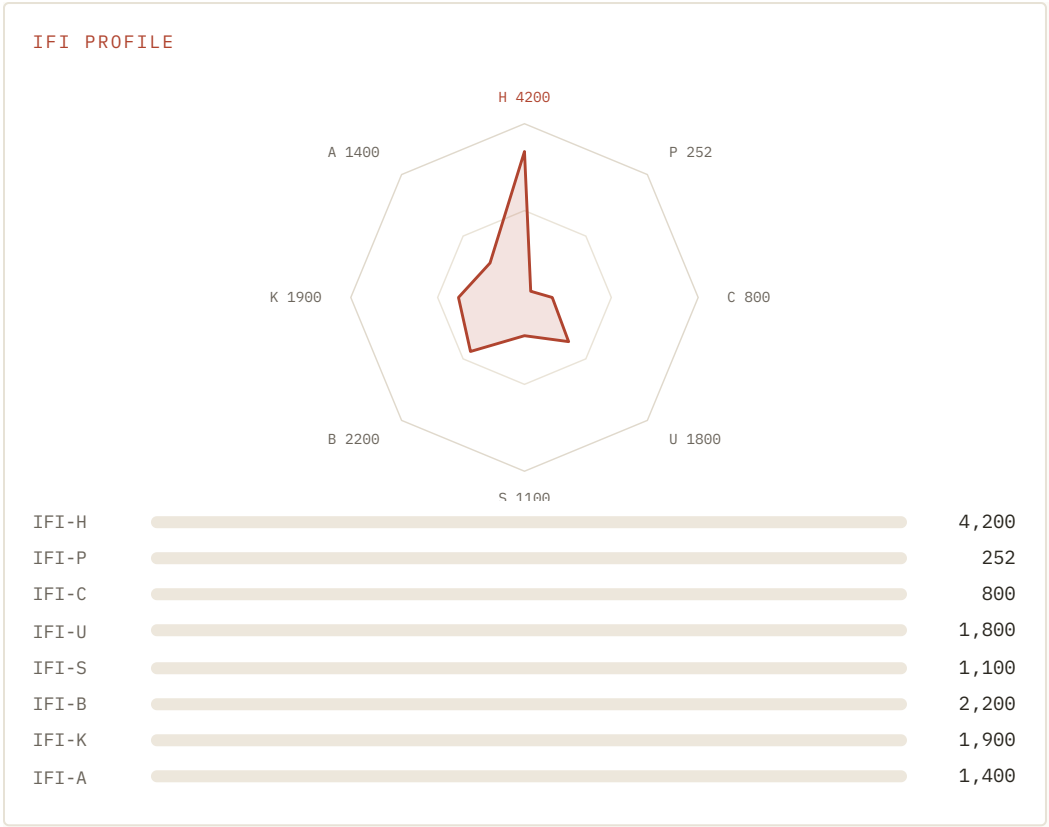

Every line in that block is a receptor's half-activation made into a score. **H: 4200** is heat at TRPV1, the channel we have followed all chapter. **P: 252** is pungency at TRPA1, the sharp mustard-and-cinnamon channel. **C: 800** is the cool-burn of TRPM8, menthol's receptor, running quietly beneath the fire. **U: 1800** is umami at mGluR1; **S: 1100** is sweetness at the TAS1R2·TAS1R3 pair; **B: 2200** is the attenuation of bitterness at TAS2R38; **K: 1900** is kokumi — richness and continuity and mouthfeel — at the calcium-sensing receptor CaSR; and **A: 1400** is astringency, the tannin dimension read by the salivary proline-rich proteins. No fruit before this one, designed or grown, was built to speak on every channel the human mouth owns. It can do so for exactly the reason the heat is reliable: a protein platform stabilized by nearly four kilocalories per mole, fed by a phenylpropanoid corridor that runs with headroom to spare, is a chassis sturdy enough to carry more than one message without buckling. Heat was never meant to be a solitary thing you crank until it hurts. It is one axis among eight, and the whole thesis of our family's work is that the eight can be *composed* — deliberately, by design, in a single living fruit — the way a chord is composed from notes that would be thin alone.

That is the full designed profile of the only eight-of-eight plant in our catalog, and it proves in one scorecard the conviction that runs under everything we grow: that a fruit can be built to hold more — more heat, more depth, more completeness, and, in the books still to come, more nourishment — without being built to do harm. By our model's estimate, Phoenix is designed to exceed the heat of any plant that came before it. It is also, and this is the part we are proudest of, the most *complete*. In a kitchen it is handled accordingly: not as a dare but as what it truly is, an extraordinarily concentrated instrument, measured out with respect into one deliberate drop.

What we are careful to claim

Rigor includes saying plainly where the design rests on assumptions, so here are ours. We treated the two ΔΔG values as additive to reach the combined −3.9 kcal/mol, and additivity is an approximation —

mutations can be epistatic, interacting when they sit close enough to feel one another. We bought down that risk by choosing edits in two separated regions of the fold — the N-terminal channel and the active-site cleft. The ΔΔG figures themselves are model predictions, used as a ranking signal and a design rail, not calorimetric measurements. The EC_{50} value behind the IFI-H score is drawn from published capsaicin-TRPV1 data for the reference compound; for our own cultivars, EC_{50} is a model estimate, not a bench measurement, which is part of why loudness compresses against loading the way it does. Our flux-balance model predicted that flux control sits in the upstream phenylpropanoid corridor — concentrated at 4CL and HCT/CCoAOMT — not at Pun1 itself; the two substitutions improve the enzyme's output reliability without altering the pathway's overall flux ceiling. And the predicted SHU, we will insist to the end, is worth nothing on its own: had the coefficient of variation stayed at forty-two percent, our model's 2.95 million would be a prediction for a single unreliable design, not a repeatable cultivar. The two mutations are the headline. The loop that found them, iterated them through structural and flux modelling, and built them for repeatability is the actual invention — and it is the same loop, run with the same care and the same reverence for the living thing we are designing toward, that we bring to every organism in the Library.

These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.



PART III OF V

The Fire That Travels

heat in plants that never carried it · the volcanic garden · four genes against a universal backdrop

Chapter 10 · The Fire That Travels — *heat in plants that never carried it.*

Here is the sentence that reorganizes everything: capsaicin was never a pepper problem. Lay the biosynthetic pathway out in full — gene by gene, substrate to product — and it tells on itself. Fire is not a deep, genus-defining trait of *Capsicum*. It is a short terminal flourish bolted onto machinery that every green plant already runs. Of the nine enzymes between a leaf's own amino acids and a molecule of capsaicin, five belong to the universal phenylpropanoid corridor — present and humming in tomato and tobacco and rice and every other vascular plant alive — and only the last four are pepper's private invention. The long, expensive, hard-won front of the pathway is common property. The ending is the only part that is pepper's own. Which means heat is portable. The fire can travel.



The tomato — the first organism beyond the pepper we taught to consider burning. Its phenylpropanoid corridor already runs at full tilt, making the pigments and defenses every plant needs; the distance from that tomato to a tomato with fire is four enzymes, not a new plant.

We do not say *portable* loosely. We mean it in the precise sense an engineer means it: the parts have an interface, the interface is conserved, and a part built against it in one host mates to the same interface in another. What follows is the whole of that claim, drawn out to the molecule, the receptor, the index that scores both, and the honest arithmetic of how far the fire can actually go.

The part of the molecule that already lives in every leaf

Begin with the burn itself, because the burn is where the portability hides. Capsaicin binds TRPV1, a non-selective cation channel standing in the membrane of the free trigeminal nerve endings that lace the surface of the tongue. The same channel is gated by heat above roughly 43 °C and by acidity. Capsaicin does not warm the tissue by a single degree. Its vanillyl head — a 4-hydroxy-3-methoxybenzyl group — slips into a lipophilic pocket between the channel's third and fourth transmembrane helices, near the S4–S5 linker that acts as the gate's hinge, and lowers the channel's thermal threshold. A channel that should have stayed shut until 43 °C now opens at the body's own 37 °C. Calcium and sodium pour through, the ending depolarizes, and the brain reads a fire that is not there. The heat is real as experience and absent as temperature. That gap is the entire sensation. And note which nerve carries it: the burn is a somatosensory signal on the trigeminal line — the system that reports touch and temperature — not a taste on the gustatory one. It is felt, not tasted, which is precisely why the same molecule reads as the same fire in any plant that can be taught to make it.

Now look at what the key is made of, because this is the quiet miracle. The vanillyl head — the exact chemical feature that docks TRPV1 — is a phenylpropanoid product. It descends from ferulate and vanillin, compounds every vascular plant already manufactures by the ton to build lignin and to defend itself. The business end of capsaicin, the part that actually speaks to the receptor, is not pepper chemistry at all. It is universal chemistry. Pepper's invention is narrower and more mechanical than it looks: pepper learned to *aminate* that universal head and to *bolt a fatty tail onto it*, converting a soluble defense compound into a lipophilic agonist that lodges in a membrane and rings a nerve. Every plant already makes the part of the key that turns the lock. Only the handle is pepper's.

Walk the nine enzymes in order and the split becomes a thing you can point at. Phenylalanine enters the corridor and PAL strips it to cinnamic acid; C4H hydroxylates that to *p*-coumaric acid; 4CL activates it to *p*-coumaroyl-CoA; the HCT/C3H cluster carries it to caffeoyl-CoA; COMT methylates onward toward feruloyl-CoA and vanillin. Five enzymes, and they are running in the tomato in the photograph above whether or not anyone ever asks it to burn — they are how it makes its pigments, its scent, the phenolic armor in its skin. Then the pepper-specific tail begins: pAMT transaminates vanillin to vanillylamine; in parallel, BCAT drains a branched-chain amino acid into a dedicated keto-acid pool and KAS elongates it into 8-methyl-6-nonenoyl-CoA, the fatty donor; and finally Pun1/CS — the capsaicin synthase, the AT3 acyltransferase — condenses the vanillylamine and the branched-chain acyl-CoA into capsaicin itself. Four enzymes. That is the whole of what a pepper knows that a tomato does not.

#	ENZYME	REACTION (SUBSTRATE → PRODUCT)	ORIGIN
1	PAL	L-phenylalanine → *trans*-cinnamic acid	Universal — runs in every vascular plant
2	C4H	cinnamic acid → *p*-coumaric acid	Universal
3	4CL	*p*-coumaric acid → *p*-coumaroyl-CoA	Universal
4	HCT / C3H	*p*-coumaroyl-CoA → caffeoyl-CoA	Universal
5	COMT	caffeoyl-CoA → feruloyl-CoA → vanillin	Universal
6	pAMT	vanillin → vanillylamine	Pepper-specific (terminal)
7	BCAT	branched-chain amino acid → branched keto-acid	Pepper-specific (terminal)
8	KAS	branched keto-acid → 8-methyl-6-nonenoyl-CoA	Pepper-specific (terminal)
9	Pun1 / CS	vanillylamine + acyl-CoA → **capsaicin**	Pepper-specific (terminal)

The shape of that table is the shape of the whole chapter. Rows one through five are a loan from the plant kingdom. Rows six through nine are the loan we repay in a new host. And because the receptor-active fragment is built in rows one through five, the fire we install in a tomato is chemically identical to the fire in a pepper — the same molecule, the same vanillyl head, the same channel, the same threshold drop. We are not making a tomato-flavored imitation of heat. We are making heat, in a tomato.

Five modules ride along, two are fitted to the host

The nine enzymes are the biology. The engineering is organized a little differently, into seven platform modules, and the ratio between what travels and what does not is the single most important number in this entire program. Five of the seven modules transfer from host to host essentially unchanged. Two must be rebuilt for every new plant. That five-to-two split is the reason forty-two host organisms arrived as fast as they did, and it deserves to be named precisely rather than waved at.

What travels: the verification of the universal corridor (which is not an installation at all, only a check that rows one through five of the pathway are running and have headroom), and the terminal cassettes themselves — the vanillylamine-forming step and the branched-chain tail supply, carried as fixed genetic packages with only minor tuning. What is bespoke, every single time: first, the promoter — the tissue-and-timing switch that must be rewired to each host's own developmental program — and second, the flux balance, because every plant's metabolism presents its bottleneck at a different step once the capsaicin synthase begins drawing on its vanillin and its branched-chain pool. Five parts ride along. Two parts are fitted. We are not inventing fire forty-two times; we are mounting two bespoke parts onto a five-part engine the plant built for itself and has been running since long before it met us.

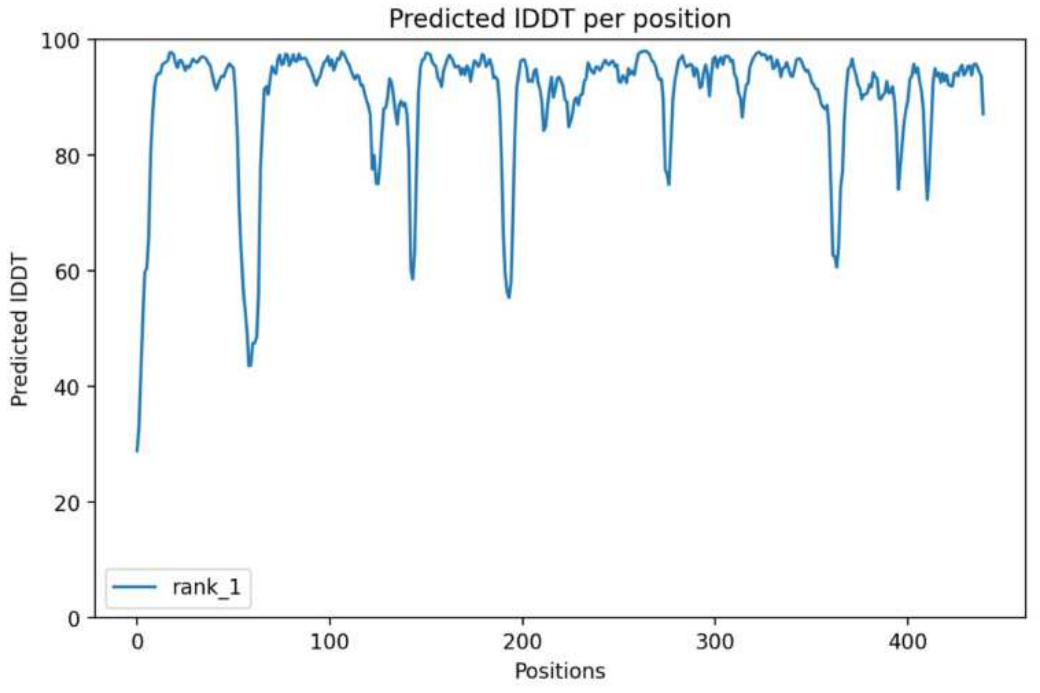

FIG. 10 — the Pun1/CS acyltransferase, the one terminal enzyme that must be written into every new host, modeled and coloured by pLDDT confidence; the binding domain and active-site cleft resolve at high confidence. Two designed substitutions ride inside this structure into every plant we build.

And the cassette that travels is not the wild enzyme. The Pun1/CS we carry into every host is an engineered one. An ESM2 650-million-parameter protein language model scored every possible single-residue substitution in Pun1 for predicted effect on catalytic efficiency and fold stability; our genome-scale flux model found that pathway control sits upstream at 4CL and HCT — not at the terminal Pun1/AT3 step — meaning Pun1 is the capsaicin synthase that labels the product, but upstream steps govern the overall flux rate. Two substitutions the breeders never found, never before combined, came out of that search and were written in with CRISPR: **S39L**, which raises N-terminal binding affinity at a free-energy change of ΔΔG = −2.1 kcal/mol; and **L345G**, which relieves steric crowding in the active-site cleft at ΔΔG = −1.8. Together they lift predicted peak output to approximately **2,937,948 SHU** in our model and, as decisively, tighten batch-to-batch variation to a coefficient of variation of **≤7%**, against **42%** in wild type. Those two edits sit inside the structure in the figure, and they travel with it — the fire that crosses a species line carries its engineering with it.

The full signature of what that fire can become, in its origin cultivar, is Phoenix (SS-0096): the first plant we designed to register on all eight axes of the Iyengar Flavor Index at once. It is the donor profile, the complete thing we draw from when we decide how much of the fire — and which of its companion axes — to carry into a new host.

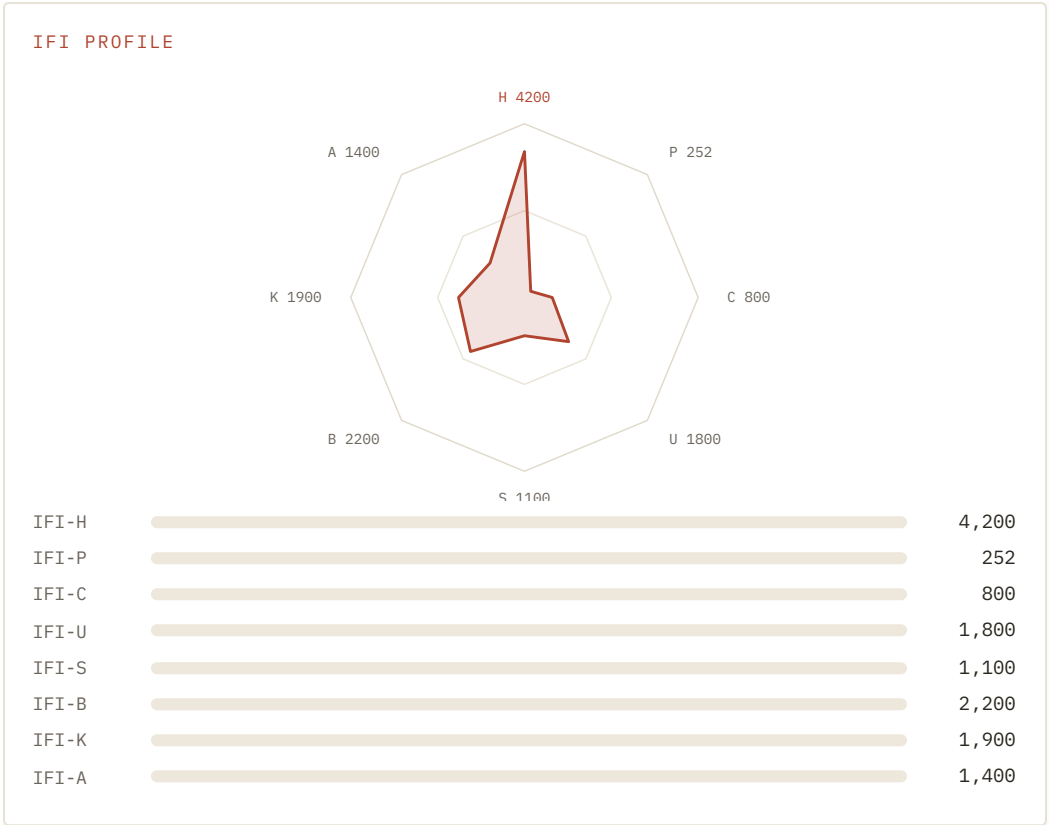

The arithmetic that lets fire cross a species line

A tomato's fire and a pepper's fire can be set on one page only because there is a number that measures *loudness at the receptor*, independent of which plant the molecule grew in. That number is the Index, and the Index is a single formula:

$$IFI(c) = 1{,}000 \times EC_{50}(\text{reference}) \div EC_{50}(\text{candidate})$$

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response — the half-activation point. It is not how much of a compound a food holds; it is how little of it the receptor needs. The underlying curve is the Hill relation, $R([C]) = R_{max} \cdot [C]^n \div (EC_{50}^n + [C]^n)$, where n is the cooperativity of the channel; for capsaicin at TRPV1, n sits near 2, so the response climbs steeply and sigmoidally. At the one concentration $[C] = EC_{50}$, the equation collapses to $R = R_{max} \div 2$ by construction — that is what the half in half-maximal means. We anchor to EC_{50} rather than to peak response or to raw concentration for one reason: EC_{50} is a property of the molecule-and-receptor pair itself, not of the dose you happened to apply, and so it is the one quantity that stays fixed whether the molecule was grown in a chili or coaxed out of a tomato.

The IFI is a stated convention — each reference constant defines IFI = 1,000 on that axis by convention, the way the Scoville scale is defined by dilution-to-threshold. Read the formula and its behavior falls out. A smaller EC_{50} means the receptor is satisfied by less, which means a louder stimulus, which — because EC_{50} sits in the denominator — means a higher IFI. The Index is pinned so that **1,000 corresponds to the half-activation of the native wild-type reference**; set candidate equal to reference and the formula returns $1{,}000 \times EC_{50}(\text{ref}) \div EC_{50}(\text{ref}) = 1{,}000$, exactly as it should.

Now work a real number. Phoenix scores **IFI-H = 4,200**. Invert the formula:

- $4{,}200 = 1{,}000 \times EC_{50}(\text{ref}) \div EC_{50}(\text{Phoenix})$

- EC_{50}(ref) ÷ EC_{50}(Phoenix) = 4.2
- EC_{50}(Phoenix) = EC_{50}(ref) ÷ 4.2 = 0.238 × EC_{50}(ref)

The channel reaches half-activation at **23.8%** of the reference concentration — Phoenix's capsaicin is **4.2× louder** per unit than the wild-type anchor. Put a textbook value to the reference, the capsaicin–TRPV1 half-activation near 0.7 µM, and the conversion lands concretely: EC_{50}(Phoenix) ≈ 0.7 µM ÷ 4.2 ≈ **0.167 µM**. Less than a fifth of a micromolar to half-open the gate. That is the number the fire carries across a species line unchanged, because the molecule is unchanged.

Read the same relationship from the other end and it sets the floor of the ladder. Ember (SS-0087) scores IFI-H = 1,570, so EC_{50}(ref) ÷ EC_{50}(candidate) = 1.57; its capsaicin reaches half-activation at EC_{50}(ref) ÷ 1.57 = 0.637 × EC_{50}(ref) — a little over six-tenths of the reference concentration, 1.57× louder than the wild-type anchor. Lay the ladder end to end and receptor volume runs from 1.57× at Ember to 4.2× at Phoenix, a 2.7-fold stretch of loudness across one axis. Because EC_{50} sits in the denominator, the law is a clean reciprocal: halve a molecule's EC_{50} and you double its Index; quarter it and you quadruple it. Loudness and half-activation are one fact read from two ends.

Here is where the Index earns its keep, and where we are most honest. Scoville heat units measure concentration by dilution — how much capsaicinoid is present. IFI-H measures loudness at the channel. They are related but they are *not the same number*, and the proof is sitting in our own ladder. Run the ten flagship superhots and divide SHU by IFI-H:

CULTIVAR	SKU	SHU (PREDICTED)	IFI-H	SHU PER IFI-H POINT
Ember	SS-0087	866,000	1,570	552
Cinder	SS-0088	1,050,000	2,290	459
Kindle	SS-0089	1,250,000	2,000	625
Blaze	SS-0091	1,720,000	2,630	654
Wildfire	SS-0092	1,960,000	2,930	669
Scorch	SS-0093	2,210,000	3,220	686
Firestorm	SS-0094	2,470,000	3,490	708
Solara	SS-0095	2,720,000	3,980	683
Phoenix	SS-0096	2,937,948	4,200	700

The ratio is not constant. It drifts from 459 to 708 units of Scoville per point of Index. And look at the inversion that matters most: **our model predicts Kindle at approximately 1,250,000 SHU yet scores only 2,000 on IFI-H, while Cinder is designed with less capsaicinoid — approximately 1,050,000 SHU — and scores higher, 2,290.** A *less concentrated* pepper rings the receptor *more loudly* than a more concentrated one. That is not an error in the data; it is the whole reason the Index exists. The length and saturation of the branched-chain acyl tail set the balance between capsaicin and dihydrocapsaicin, and that balance changes the molecule's affinity for the pocket — its EC_{50} — independently of how much total capsaicinoid the gland packed in. Dilution counts molecules. The

Index counts what the nerve actually hears. When we move fire into a new plant, SHU tells us we succeeded at making the compound; IFI-H tells us whether it will be heard — and only the second number crosses between species cleanly, because only the second number is about the receptor rather than the harvest.

Reading a transfer before we build it — flux and CV

Portability would be a slogan if we could not predict a transfer before committing a season to it. We can, because the flux logic travels with the pathway. For each new host we build the same genome-scale flux-balance model, verify that the universal corridor has headroom to feed a new terminal branch — it almost always does; the corridor is built for high throughput — and identify, in that specific plant's metabolism, what the rate limiter becomes once the capsaicin synthase starts pulling on its vanillin and its branched-chain pool. The model does not tell us the fire is possible. It tells us *where the plant will choke*, which is the only question worth asking before the bench work starts.

The second instrument we watch is reproducibility, and here the arithmetic is simple and unforgiving. The coefficient of variation is CV = (σ ÷ μ) × 100%. Hold a batch target at Phoenix's predicted **2,937,948 SHU** and the difference between wild type and engineered becomes a difference you could taste blindfolded:

- At the engineered **≤7%**: σ = 0.07 × 2,937,948 = **205,656 SHU**. One standard deviation spans roughly 2,732,000 to 3,143,000 — and in Index terms, 7% of 4,200 is **±294 IFI-H points**.
- At wild-type **42%**: σ = 0.42 × 2,937,948 = **1,233,938 SHU**. One standard deviation now spans 1,704,000 to 4,172,000 — a batch could land anywhere across that chasm — and in Index terms, **±1,764 points**.

Widen the engineered case to 95% confidence (2σ) and it is still only ±411,000 SHU, ±588 IFI-H points. The wild-type plant is a lottery; the engineered plant is a specification. And for a *new host*, that CV number is the acceptance gate: a tomato where we cannot hold the burn to ≤7% from fruit to fruit is a tomato where the localization is not finished. Consistency is not a quality-control afterthought. It is the signal that the two bespoke parts — promoter and flux balance — have actually seated.

A countable amount of host

How far the fire can travel is not mystical. It is a precise, countable amount of host that must be relearned, and we would rather name the count than sell the dream. The readiness ladder sorts every target by how much of that count remains.

READINESS	WHAT IT MEANS	EXAMPLE HOSTS
Native	Already makes capsaicinoids	*Capsicum* (all)
Easy	Same family; corridor and chassis shared	tomato, potato
Moderate	Corridor strong; promoter work needed	eggplant, tomatillo
Hard	Corridor present; flux and localization bespoke	grape, tea, coffee
Frontier	Distant lineage; multi-year build	rice, corn, soybean

Consider what *Hard* actually costs, with a concrete vine.



The wine grape — a *Hard* host. Its phenylpropanoid corridor is present and powerful, but tuned across evolutionary time to pour carbon into tannins and anthocyanins, not to feed a capsaicin branch; teaching it fire means a season of flux modeling to find where its own chemistry will choke.

Grape is a different family entirely from pepper. Its corridor runs hard, but it runs toward the tannins and anthocyanins that make a wine grape a wine grape; its ripening program and its promoters are its own; and its metabolism presents its bottleneck at a different step than a tomato's does. Installing fire there means not four tidy modules dropped in but a season of flux modeling to find where the grape's chemistry will choke, bespoke promoter work to confine expression to the berry alone, and careful balancing so the new branch does not rob the vine of the very compounds that make it a grape. It is doable — the corridor is there, the method holds — but it is years, not months. That is exactly why grape sits at *Hard* and rice sits beyond it at *Frontier*: a cereal grass is a distant lineage whose corridor feeds a grain, and the distance is measured in relearned promoters and remodeled flux, not in magic.

That honesty about the count is what turns a roadmap into a plan rather than a promise. **Forty-two host organisms are in our design platform today** — from pepper itself with its 1,060 constructs, through tomato at 2,040 and habanero at 3,190, down the long tail to coffee and vanilla and cinnamon at 40 apiece — and **sixty-four concept organisms sit on the roadmap across five tiers.** The near frontier is the Solanaceae core, four-to-five genes and eighteen-to-twenty-four months: a Volcanic Tomato in grades named *Ember · Magma · Pyroclast · Caldera · Krakatoa*, a Dark Matter Eggplant running *Smolder* to *Event Horizon*. Beyond it the extended nightshades — Seismic Potato, Acid Tomatillo — and then the true cross-family reach, seven to ten genes and three-to-four years: a Sweet-Heat Strawberry carrying grades from *Warm Kiss* to *Candied Fang* to *Devil's Dessert*, a Contradiction Cucumber, a Forbidden Apple. The far horizon is the synthetic frontier, ten-to-twelve genes and up to six years: Storm Corn, Wildfire Broccoli, Dynasty Rice. The ladder is long because the honest count is long.

What a first transfer actually carries is worth drawing, because it is humbler than the donor. Phoenix speaks eight axes at once; a plant given only the capsaicinoid cassette speaks exactly one. The fire in transit is a single-axis payload — heat installed alone, silent on the other seven — and its magnitude in any given host is governed by

that host's own flux ceiling, a bespoke unknown until the model is run. The donor is a symphony; the first thing to cross a species line is a single, clean note.

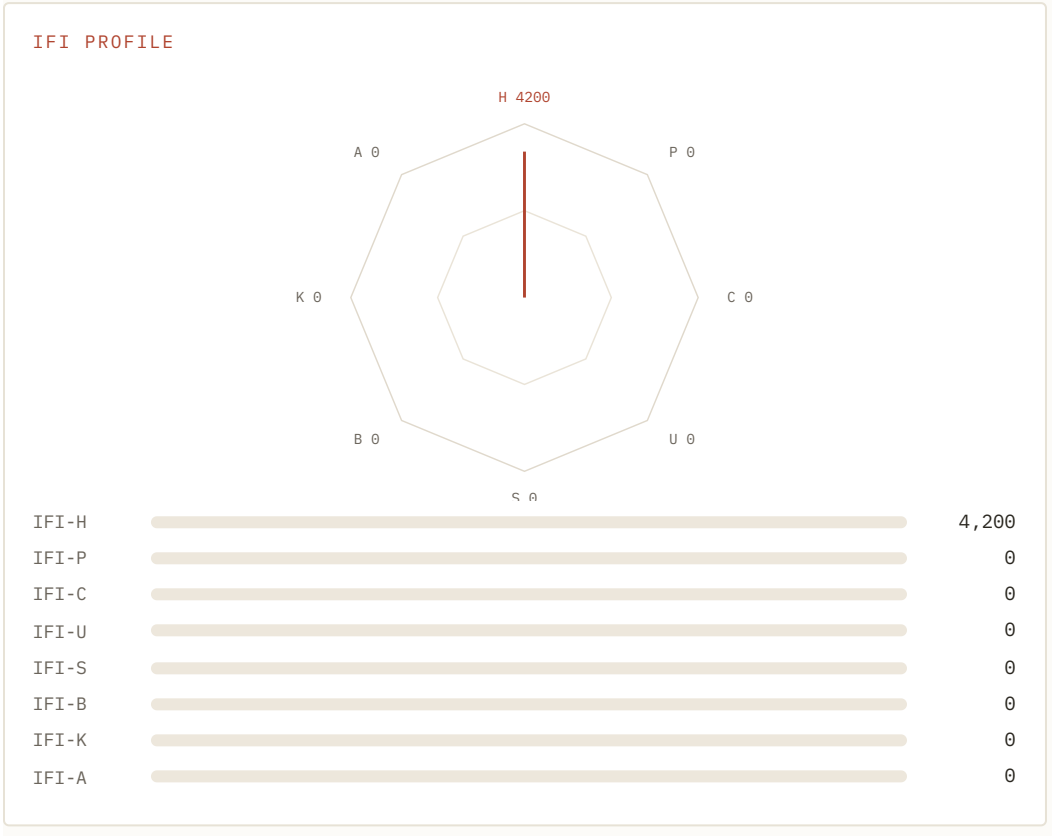

The leash is part of the craft

Moving fire into new plants is exactly the kind of power that demands responsibility, and we would rather say this plainly than let it be inferred. The subtlety that makes the engineering real is also where the safety lives, because the two are the same problem. A capsaicin pathway dropped carelessly into a plant will try to run everywhere the phenylpropanoid corridor runs — which is everywhere — and a plant that makes capsaicin in its leaves and roots and stems is not a flavor, it is a self-defense catastrophe that cannot photosynthesize. So the real work was never pasting in four genes. It is *spatial and temporal control*: the terminal enzymes must sit under promoters that switch them on only in the fruit, only in the tissue we intend, and only at the ripening stage when the fruit builds its final character. *Capsicum* solved this over evolutionary time by confining capsaicin synthase to the placental gland cells. We solve it with a tissue-specific promoter and a developmental switch, so a volcanic tomato burns in its flesh and nowhere else, and only when ripe. The pathway is portable; the localization is the engineering — and the localization is also the safety.

A fruit-specific, ripening-triggered promoter does more than protect the plant's own metabolism. It confines the trait to the single organ a person actually eats, at the stage they actually eat it, and keeps the fire out of the leaves and pollen a field worker's hands or a bee's legs would ever meet. Where a crop's biology allows, we pair that confinement with genetic containment — approaches that keep an engineered trait from drifting on the wind into plants it was never meant for — so the fire stays inside the lines we drew. None of this is appended for a regulator's comfort. It is a design constraint we solve at the same bench as the flux balance, because a trait you cannot contain is a trait you have not finished engineering.

And a new food carries real obligations that we treat as part of the work, not paperwork stapled to the end of it: allergen assessment, full compositional analysis, honest labeling. That is the clean, plant-reverent sensibility this whole program is built on — the conviction that food is something you give a body with intention, grown to carry the most it can and to keep nothing in it that should not be there. When we say a tomato can be taught to burn, we are describing a capability held inside a discipline. The point of the fire that travels is not that fire can go anywhere. It is that fire can go, precisely, where a cook would actually want it — and stay out of everywhere else. That restraint is not the price of the craft. It is the craft. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Chapter 11 · The Volcanic Garden — *tomato, eggplant, potato, learning to burn.*

The roadmap beyond the pepper is a queue, not a wish list, and it is ordered by a single quantity: genetic distance from the chassis we have already mastered. It begins close to home. The nightshades come first — the Solanaceae, the family that holds the pepper itself, and therefore the family whose members share the most of the pepper's cellular machinery. If fire is going to travel, it travels most easily to its cousins. The opening stretch of the queue is a garden of nightshades learning, one gene module at a time, to burn: a tomato, an eggplant, a potato, a tomatillo, each being taught to carry a designed warmth it has never, in its whole evolutionary history, known how to make.

We did not pick the nightshades out of sentiment. We picked them because the method is cheapest there, and the honesty about *why* is itself part of the engineering. What follows is the whole of that reasoning — the chassis logic, the enzymology we port, the receptor that reads the result, the mathematics that lets us promise a specific warmth and not a surprise, and the places where the work is still only a plan.



Eggplant in cross-section — Dark Matter in concept, the nightshade whose dense, absorbent mesocarp is an ideal canvas for a slow, deep, dihydrocapsaicin-leaning heat built into the flesh itself rather than borrowed from a sauce.

The chassis, not the heat, is the axis of difficulty

The temptation is to read a roadmap of ever-hotter plants as a ladder of heat. It is not. The hard variable is never how hot the fire is; it is how much of the *host* must be relearned to make the fire there at all. A pepper is a Solanaceae, and its closest relatives — tomato, eggplant, potato, tomatillo, ground cherry, even ornamental petunia — share far more than the universal chemistry of plants. They share the pepper's specific infrastructure: homologous promoters,

compatible ripening programs, the tissue-culture and transformation protocols worked out over decades of research on the family. Moving fire within the family is like porting software between two devices that run the same operating system — most of the code simply runs. Moving it to a strawberry or a melon, a different family altogether, means rewriting every routine that quietly assumed a Solanaceae host. Moving it to a grass like rice or corn, separated from the pepper by roughly a hundred and fifty million years, means rewriting most of them.

That is the true meaning of the tiers. We work down the queue in order of genetic distance because that is the order in which the method grows more expensive, and the gene budget is simply the tax the distance levies.

TIER	REPRESENTATIVE CONCEPTS	GENE BUDGET	WHAT THE HOST MUST RELEARN	HORIZON
T1 · Solanaceae core	Volcanic Tomato · Dark Matter Eggplant	4–5	the four terminal enzymes; one ripening-coupled promoter	18–24 mo
T2 · Extended nightshades	Seismic Potato · Acid Tomatillo · Celestial ground cherry	5–7	+ a storage-organ sink (a tuber is a stem, not a fruit)	24–36 mo
T3 · Cross-family	Sweet-Heat Strawberry · Contradiction Cucumber · Solar Melon	7–10	+ native promoters, fresh localization tags, some upstream supplementation	36–48 mo
T4 · Frontier synthetic	Storm Corn · Dynasty Rice · Wildfire Broccoli	10–12	+ monocot retuning, a rebuilt vacuolar sink, a full flux rebalance	48–72 mo

The four terminal enzymes are fixed at every tier; they are the engine of heat and they do not change. What grows, tier by tier, is the scaffolding the engine needs to run in a stranger's cell — the promoters that must now match the host's own ripening clock, the address tags that route a membrane-active molecule somewhere it will not poison the plant, the codon choices, the rebalancing of carbon and nitrogen flux so the host survives the new demand on its metabolism. At the Solanaceae core that scaffolding is almost nothing, because the cousins already speak the dialect. At the synthetic frontier it is most of the build.

What we actually install — the terminal module

Capsaicin is the end of a two-branch convergence, and almost all of it is already present in every nightshade on this list. The upstream corridor — **PAL → C4H → 4CL → HCT → COMT**, the phenylpropanoid branch — runs in every vascular plant alive; it is the same route a tomato uses to make lignin and the pigments in its skin. It delivers vanillin, and from vanillin the pepper's own aminotransferase, **pAMT**, makes vanillylamine: the aromatic head of the molecule. The second branch is the aliphatic tail. Branched-chain amino acids — valine, leucine — are handed off by **BCAT** and elongated by the ketoacyl synthase **KAS** into a short, 8-methyl, branched-chain fatty acid, activated as its CoA thioester. Then the keystone: **Pun1/CS**, capsaicin synthase, an acyltransferase of the BAHD family, condenses vanillylamine with that 8-methyl-branched acyl-CoA and forms the amide bond that *is* capsaicin.

So the work of giving a tomato heat is not building a pathway from nothing. It is installing four terminal enzymes — pAMT, BCAT, KAS, Pun1/CS — onto a corridor the fruit already runs, and then teaching the fruit when and where to switch them on. Five of our seven platform modules transfer to any sequenced genome unchanged; only these four are pepper-specific, and only they have to be carried across.

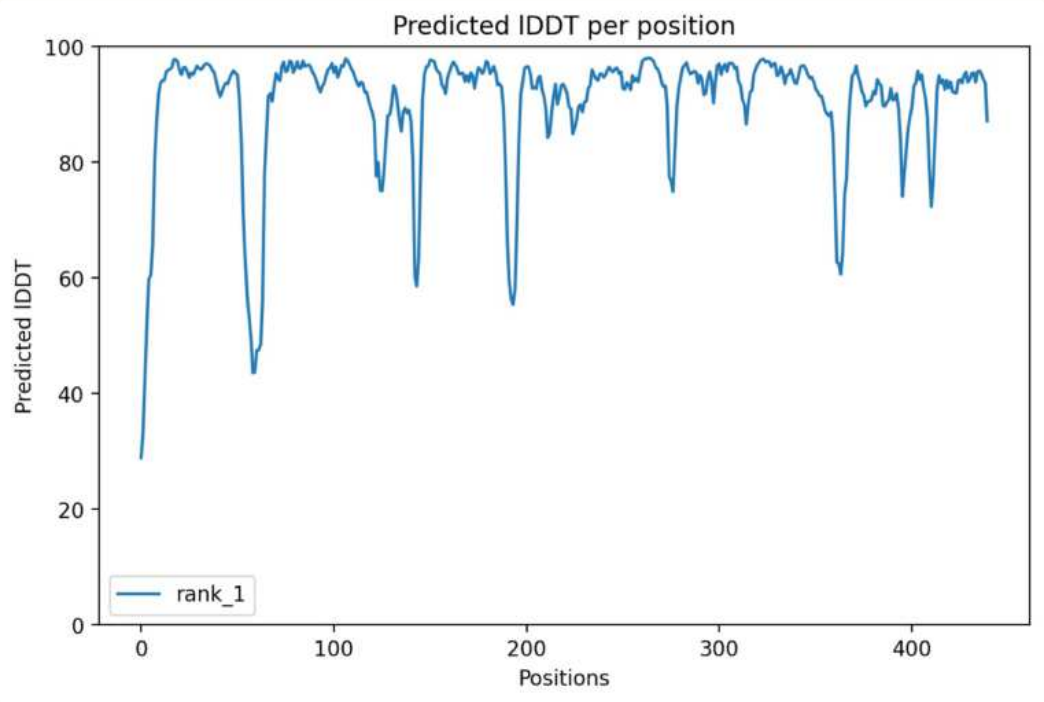

FIG. 11 — the Pun1/AT3 acyltransferase, coloured by model confidence; the substrate cleft that binds the branched acyl-CoA and the vanillylamine, and that we port whole into the nightshade, resolves at high confidence.

The length and saturation of that aliphatic tail are not a detail — they are the *shape* of the burn. An unsaturated 8-methyl-6-nonenoyl chain gives capsaicin proper, the bright, forward, fast-arriving heat. Its saturated partner gives dihydrocapsaicin, a slower, deeper, longer-dwelling warmth. KAS and the cell's desaturase balance set the ratio of the two, and so the module we install is also a dial: a tomato can be tuned toward the quick lift of capsaicin or, as in the eggplant, toward the dense, absorbent, dihydrocapsaicin-leaning heat that suits flesh built like a sponge.

How the tongue reads it — TRPV1, exactly

Capsaicin binds TRPV1, a non-selective cation channel standing in the membrane of the trigeminal nerve endings of the tongue — the same channel gated by temperatures above 43 °C and by protons. It is a homotetramer: four identical subunits around a central pore, and within each subunit a vanilloid pocket formed by the S3 and S4 helices and the S4–S5 linker, where a single tyrosine, Y511, anchors the vanillyl head of the molecule. Capsaicin does not warm the tissue. It binds that pocket and lowers the channel's thermal threshold, so a mouth at body temperature is suddenly above the channel's set-point and the nerve reports a fire that is not there. The sensation is real; the heat is a fiction the receptor has been talked into.

Because the channel has four pockets, its response to concentration is cooperative — the binding of one molecule makes the next bind a little more readily — and the dose–response curve is steep, with a

Hill coefficient near 2. Writing the fraction of maximal response as a function of concentration C against the half-activating concentration EC_{50}:

```
response = C² / (EC₅₀² + C²)
```

At $C = EC_{50}$ the two terms are equal and the response is exactly one half — that is what EC_{50} *means*, the concentration that half-opens the channel. The steepness is the part worth feeling. Hold EC_{50} fixed and walk the concentration across it:

C ÷ EC_{50}	TRPV1 RESPONSE	READING
0.25	5.9%	a whisper
0.58	25.2%	barely warm
0.93	46.4%	just under the target
1.00	50.0%	the target — half-activation
1.07	53.4%	just over
1.42	66.9%	clearly hot
2.00	80.0%	a double dose reads only 60% louder

Two things fall out of that column. First, near the target a small change in concentration makes a large change in sensation — the curve is nearly vertical through its middle, which is exactly the region a designed food has to live in. Second, far above the target the curve flattens: doubling the capsaicin past EC_{50} does not double the heat, it adds a fifth, because the channels are already mostly open. The ladder of superhots climbs into that flattening, which is why each successive flagship costs more capsaicin for less additional loudness — and why, for a tomato, the interesting, legible, teachable range sits *below* the knee, where a modest, well-held concentration does honest, proportional work.

Putting a number on the loudness

The Iyengar Flavor Index turns that half-activation into a single score. The axis of heat, IFI-H, is anchored so that **1,000** marks the half-activation of the native wild-type reference, and every candidate is read against it by inverting the EC_{50} and scaling by a thousand:

```
IFI-H(candidate) = 1,000 × EC₅₀(reference) ÷ EC₅₀(candidate)
```

The formula runs in both directions, and the backward direction is how we set a target. Decide that the lead Volcanic Tomato should carry an IFI-H of 900 — a true background warmth, deliberately *below* the wild-type pepper anchor — and the arithmetic hands back a concentration:

```
EC₅₀(tomato) = 1,000 × EC₅₀(ref) ÷ 900 = 1.11 × EC₅₀(ref)
```

The tomato needs about eleven percent more of its extract than the reference to half-open the same channel: a little quieter than a wild pepper, a warmth that greets you rather than seizes you. Run the same arithmetic on Phoenix, the full-spectrum superhot we use as our calibrated yardstick, reading 4,200:

$$EC_{50}(\text{Phoenix}) = 1{,}000 \times EC_{50}(\text{ref}) \div 4{,}200 = 0.238 \times EC_{50}(\text{ref})$$

Under a quarter of the reference concentration half-opens the channel — the plain meaning of a number that large. Laid along the proof-of-concept ladder, the index and the Scoville scale climb together:

CULTIVAR · SKU	PREDICTED SHU	IFI-H	$EC_{50} \div EC_{50}(REF)$	SHU PER IFI POINT
Ember · SS-0087	866 K	1,570	0.637	552
Cinder · SS-0088	1.05 M	2,290	0.437	459
Kindle · SS-0089	1.25 M	2,000	0.500	625
Blaze · SS-0091	1.72 M	2,630	0.380	654
Wildfire · SS-0092	1.96 M	2,930	0.341	669
Scorch · SS-0093	2.21 M	3,220	0.311	686
Firestorm · SS-0094	2.47 M	3,490	0.287	708
Solara · SS-0095	2.72 M	3,980	0.251	683
Phoenix · SS-0096	2.95 M	4,200	0.238	700

(SHU values are computational predictions from our flux model. IFI-H scores are model-based. Flare, SS-0090, sits at 1.48 M between the third and fourth rows; its IFI-H is not yet fixed, so we leave it off the arithmetic rather than guess.)

Read the last column across and the two scales quietly disagree. Cinder spends 459 SHU for each point of IFI-H; Firestorm spends 708 — more than half again as much — and the figure wanders, not drifts, up the ladder. The reason is exact and it is the whole case for the Index. Scoville counts capsaicinoids by mass and weights them all alike, but the receptor does not. The minor capsaicinoids a pod also makes — nordihydrocapsaicin, homocapsaicin, homodihydrocapsaicin — carry real Scoville mass yet speak to TRPV1 more softly than capsaicin itself, and the food matrix around them, its fats and sugars, can blunt or sharpen their delivery. So a cultivar whose KAS and desaturase settings tilt its output toward those quieter species reads high in SHU per IFI point: loud on paper, more measured in the mouth. IFI-H is the number the tongue would give; SHU is the number the balance gives. When the two separate, the separation is information about the *shape* of the fire, and it is why the whole garden is scored against the receptor and not against a mass.

IFI PROFILE

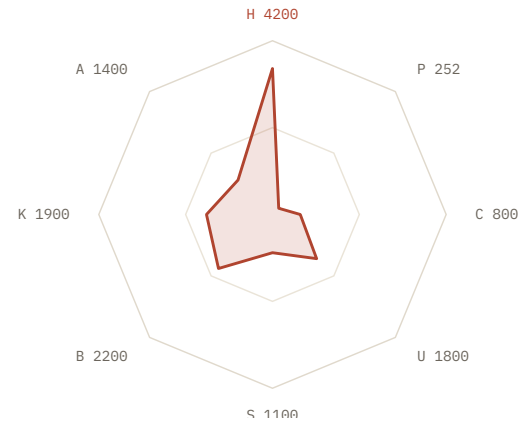

IFI-H	4,200
IFI-P	252
IFI-C	800
IFI-U	1,800
IFI-S	1,100
IFI-B	2,200
IFI-K	1,900
IFI-A	1,400

That profile is Phoenix — the pepper we built gene by gene, lifted in predicted output by two AI-designed substitutions in Pun1 (S39L, −2.1 kcal/mol; L345G, −1.8) that no breeder ever found by selecting phenotype and hoping. It is not a member of the garden; it is the ruler we hold the garden against. Every tomato, every eggplant, every potato in the pages ahead is measured on these same eight axes, with this same instrument.

Where the fire goes — localizing heat in a fruit that never made it

A pepper does not make its heat in the flesh you bite. It makes it in the placenta — the pale pith the seeds cling to — in specialized epidermal blister cells that secrete capsaicinoids into vacuoles and onto the tissue surface, sequestered away from the living cytoplasm. That sequestration is not incidental. Capsaicin is amphipathic and membrane-active; left loose in a cell's own membranes at the concentrations a superhot reaches, it is toxic to the cell that made it. The pepper solved its own autotoxicity by building a dedicated place to put the fire. A tomato has no such place.



Tomato in cross-section — the locular jelly and the thick pericarp wall, the two possible sinks for an installed heat; a fruit with no blister-cell apparatus of its own, which must be given an address for the molecule before it is given the molecule.

So the second half of the terminal module is not an enzyme at all — it is an address. We co-install localization and sequestration machinery that routes the finished capsaicinoid into a vacuolar or apoplastic compartment, and we drive the whole module from a promoter tied to the host's own ripening or maturation clock, so the heat appears late, in the fruit, and nowhere else. The consequences of getting that wrong are the failure modes we design against from the first construct. Route the molecule into living membranes and the plant poisons itself. Let the promoter leak into the leaves and the whole plant turns pungent, its photosynthesis and its ecology disturbed. Leave the pathway's flux unbalanced and intermediates pile up — free vanillin, free branched-chain acids, each with its own intrusive flavor — so the fruit tastes not of warmth but of the half-finished chemistry of warmth. And divert too much phenylalanine and too much valine into the new sink and you starve the host of what it needs to grow; the flux-balance model that places control upstream — at 4CL and HCT/CCoAOMT, not at the terminal Pun1 step — is the same model that keeps the tomato alive while it burns.

Each nightshade asks the address question differently, and the differences are the reason the queue has tiers. The eggplant's dense mesocarp drinks flavor like a sponge, so the fire is built *into* the flesh — present everywhere you bite, slow and deep. The tomato offers two sinks, the watery locular jelly and the firm pericarp wall, and the choice between them is a choice between a bright burst and a steady background. The potato is the hardest of the three, because a potato is not a fruit at all — it is a tuber, a swollen underground stem, a storage parenchyma with no ripening program to borrow, which is exactly why the Seismic Potato sits a tier out from the tomato rather than beside it.

Holding the number — a steep curve demands a tight hand

A declared heat is a promise, and a promise is only as good as the batch-to-batch variance behind it. Wild-type pepper is a notoriously unreliable narrator: its capsaicin content swings with a coefficient of variation near **42%** from harvest to harvest, sun to shade, stress to ease. Our engineered constructs hold capsaicinoid output to a CV of **≤7%**. On a spec sheet that is a tighter number. On the tongue, through the steep TRPV1 curve, it is the difference between a product and a lottery.

Lay the two variances across the dose–response and watch what the receptor does with them. A ±7% swing in concentration walks you from 0.93 to 1.07 of EC_{50} — the two rows that bracket the target — and the response moves only from 46.4% to 53.4%, a band of about seven points around half-activation. The same curve, fed the wild-type's ±42%, walks from 0.58 to 1.42 of EC_{50}, and the response swings from 25.2% to 66.9% — from *barely warm* to *clearly hot*, on the identical construct, from one jar to the next. A tomato promised at IFI-H 900 would arrive anywhere from a timid 522 to a startling 1,278. That is not a seasoning; it is a hazard, and it is precisely the kind of ambush this whole book exists to abolish. The steepness that makes designed heat expressive is the same steepness that makes it unforgiving, and so the ≤7% is not a quality boast tacked onto the science. It is load-bearing. The curve demands it.

A word on the units, because the scale is easy to lose. Scoville and concentration convert by roughly sixteen heat units per part-per-million of capsaicinoid, and a part-per-million of capsaicin — molar mass about 305 g/mol — is on the order of a few micromolar in water. The receptor itself half-activates in the sub-micromolar-to-micromolar range. These are vanishing quantities: the whole art is that a molecule present at parts per million, held to seven percent, routed to one compartment of one tissue at one point in ripening, is enough to make an entire fruit speak a sensation it was never born with.

The lead of the garden — the Volcanic Tomato

Lead the queue is the Volcanic Tomato, and its concept series reads like a weather report from inside a mountain: Ember, Magma, Pyroclast, Caldera, Krakatoa — a laddered line climbing from a warm background glow to a genuinely formidable fruit, each named for a stage of eruption. The appeal is not novelty for its own sake. A tomato that carries its own designed heat is a cook's dream, because it fuses two of the hardest things to balance — the deep glutamate umami a real tomato can be engineered to restore, and a fire tuned to sit inside that savory depth rather than on top of it. Heat and umami are built in the same fruit, by the same ripening program, so they arrive fused the way a habanero's fire and fragrance arrive fused. You cannot buy that integration in a jar. It has to be grown.

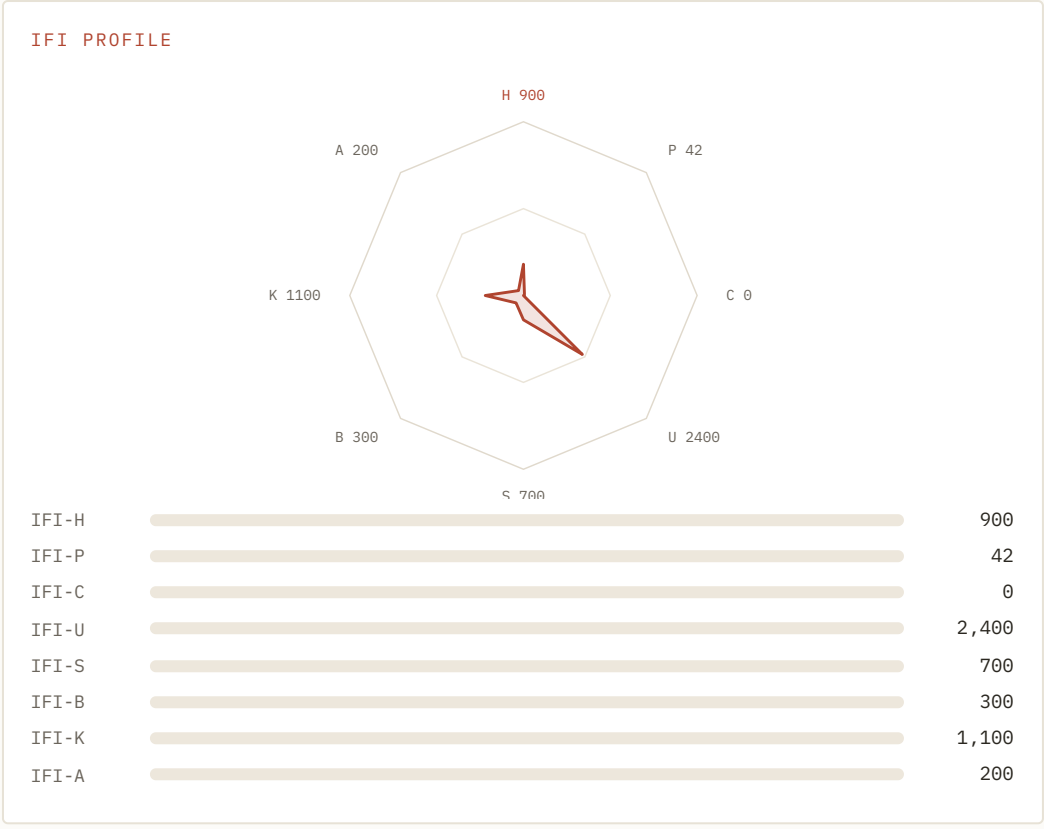

Read that profile and you can taste the ambition. A band-3 heat at 900 — gentle, just below the wild-pepper anchor — set *inside* a restored umami of 2,400 and a real kokumi floor of 1,100, with a touch of sweetness at 700 and the cool axis held at zero because nothing here should read as mint. This is not a hot tomato. It is a savory, warm, complete one: the sauce and the spice grown together on a single vine, heat composed against another axis within the organism, arrived already balanced — which is the one thing no sauce, assembled after the fact on a plate, can do.

Alongside it, the first two tiers hold the rest of the volcanic garden, each a different answer to the same question. Dark Matter Eggplant — Smolder, Nightshade, Eclipse, Obsidian, Event Horizon — builds its slow dense heat into the one vegetable whose flesh soaks up flavor like a sponge, so the fire is *inside* the eggplant, not painted on. Seismic Potato takes the most beloved blank canvas in the kitchen and gives it a warmth from within. Acid Tomatillo marries capsaicin to the tomatillo's native, bright, green tartness — a salsa verde that comes pre-composed, heat and acid grown into one husked fruit.



Potato — Seismic in concept: a warmth built into the most familiar staple on Earth, the blank canvas finally given a voice of its own, and the hardest of the near nightshades because a tuber has no ripening clock to borrow.

We owe the reader the honest state of all four. The tomato is an established chassis — more than two thousand constructs designed and built into it already, for sweetness and umami and their combinations. The potato is an established host too, though for bitterness, not heat. But *carrying capsaicin* is a module neither has shipped; the Volcanic Tomato and the Seismic Potato are near-term concepts of the first and second tiers, not products on a shelf. The eggplant and the tomatillo are a step less certain still — concept organisms, not yet among the forty-two hosts we have established. We will not blur that line. What is real today is the method and the chassis; what is on the roadmap is the heat inside these particular cousins, eighteen to thirty-six months of work away.

The further tiers, and a frontier that arrives labeled

The tiers past the nightshades are more distant and more wondrous. The sweet-heat strawberry earns a moment even here, because it is the clearest picture of where the fire is going once it leaves the pepper's family behind. Its concept line — Warm Kiss, Sugar Burn, Candied Fang, Devil's Dessert — climbs from a berry with a warm finish to one that genuinely bites, and the engineering trick is not heat but *timing*. The capsaicin is localized and its release tuned so the sweetness lands first and the warmth blooms a half-second behind it — a swicy time-course grown into a single fruit rather than assembled on a plate. A strawberry sweet going in and warm coming out is not a stunt; it is a demonstration that heat can be sequenced against another axis *inside the organism*, designed from the first taste to the last. It is a cross-family build of real difficulty, years out, and we say so plainly. Past it lie the grains and the

brassicas of the synthetic frontier — Storm Corn, Dynasty Rice, Wildfire Broccoli — where the genetic distance is greatest and the gene budgets run to a dozen, the builds stretching toward six years.

The method, though, is the same at every tier: confirm the native corridor, install and localize the four terminal modules, balance the flux so the host lives, hold the variance to seven percent. The only thing separating a designed tomato from a frontier grain is how much bespoke tuning the distance demands. We will not pretend the far tiers are near. We will say that the queue is real, ordered, and worked one cousin at a time.

A frontier of edible novelty like this one lives inside a real regulatory landscape, and we treat that as part of the design and not a hurdle to clear afterward. A tomato engineered to carry capsaicin is, to a regulator, a novel food and in most jurisdictions a genetically engineered one — which means compositional analysis to show it is exactly what we say and nothing we did not intend, allergenicity assessment of the new proteins, and clear labeling so a shopper knows precisely what they are buying. We take the labeling obligation most seriously of all, because a fruit that looks like an ordinary tomato and carries a real, built-in warmth is the exact surprise our whole ethos forbids. The heat must be declared as plainly as it was designed. None of this is a brake on the work; it is the same discipline that runs through everything our family does — that heat should be a pleasure offered, never an assault sprung — spoken in the language of food law. The volcanic garden will arrive labeled, characterized, and honest, or it will not arrive.

That is the fire that travels, arrived at its destination: not a weapon dropped into a vegetable, but a designed warmth grown into the flesh of a tomato that is also savory, also sweet, also complete — the sauce and the spice raised together on one vine, and read, to the last decimal, on the same eight axes as everything else we grow. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.



PART IV OF V

The Kitchen

how designed heat behaves in a pan · a real recipe at every band · the responsibility of the edge

Chapter 12 · Cooking With a Number — *how designed heat behaves in a pan.*

A number is only ever in service of a meal, and so the longest part of this book is the kitchen. But before the recipes come the physics, because capsaicin behaves in a pan according to rules that have nothing to do with how it behaves on a spreadsheet, and a cook who knows those rules can place heat exactly where the Index says to, while a cook who does not will watch a carefully dosed dish come out wrong. What follows is everything we had to learn at our own stove to make a designed cultivar keep the promise its profile makes — the molecule, the receptor it speaks to, the shape of the dose-response curve, and the two kinds of arithmetic that turn a single drop into dinner. Hold these, and heat stops being a gamble you make with a pinch and a prayer, and becomes something you compose the way you compose a chord.

The molecule, and why oil is its only honest solvent

Everything begins with the shape of one molecule. Capsaicin is 8-methyl-*N*-vanillyl-6-nonenamide: a small polar vanillyl head, an amide bond that acts as a hinge, and a long, branched, hydrophobic tail. That tail is the whole story in the pan. It gives capsaicin an oil-to-water partition coefficient — its log *P* — of roughly three to four, which is a chemist's way of saying the molecule would rather sit in fat than in water by a factor of a thousand to ten thousand. This is the single most consequential fact in the spicy kitchen, and it is why a swallow of water does nothing for a burning mouth: water cannot dissolve the fire, it only pushes the oil around. Milk works because its casein — an amphiphilic phosphoprotein that assembles into micelles with hydrophobic interiors — physically wraps the capsaicin and lifts it off the receptor, while its own fat gives the molecule somewhere to go.



A cross-section of the fruit, cut clean. The capsaicinoids are not in the flesh and not in the seeds but in the pale placental tissue — the ribs that hold the seeds — where the Pun1/AT3 acyltransferase assembles them. This is the tissue we weigh.

In the pan, that partition coefficient issues one instruction: heat must be *bloomed in fat*. Add a dried chili to a watery broth and much of its capsaicin stays locked in the placental tissue — the pale ribs inside the pod where the enzyme builds it — because water is the one solvent it refuses. Toast that same chili in oil first, or melt it into butter, and the fire releases into a medium that will carry it evenly through the whole dish. Fat is the solvent of heat, and good fat is also the carrier of the fat-soluble nourishment a plant holds, which is why we bloom in oil we would be glad to eat on its own. There is a second-order consequence worth knowing: because capsaicin lives in the fat phase, a thin soup with an oil slick on top delivers its heat unevenly — the burn rides the droplets and hits in ambushes — until something emulsifies it. A little mustard, the lecithin in a cooked allium, a starch that gelatinizes and thickens: any of these pulls the fat into suspension and spreads the fire smooth. Bloom it in fat, then bind the fat into the dish. That is the mechanical heart of every recipe we cook.

Where the fire actually sits — the dose-response curve

Capsaicin does not warm the tissue. It binds TRPV1, a non-selective cation channel on the trigeminal nerve endings of the tongue, slotting into the vanilloid pocket between the channel's S3 and S4 helices. The same channel opens above about 43 °C and in the presence of protons; capsaicin simply lowers its thermal threshold, so the body reports a fire that is not there and lets calcium and sodium flood a nerve that was built to warn you about a hot pan. The sensation is real; the heat is a rumor the molecule starts.

What matters for cooking is that a receptor does not answer a stimulus in proportion — it answers on a curve. Fractional receptor occupancy follows the Hill relation, $\theta(C) = C^n / (EC_{50}^n + C^n)$, where C is capsaicinoid concentration, EC_{50} is the concentration that drives the channel to half of its full response, and n is the Hill coefficient — near 2 for capsaicin at TRPV1, because the channel binds cooperatively. Two features of that curve govern everything at the stove. First, at $C = EC_{50}$ the occupancy is exactly one half, by definition — and that half-activation point is the anchor the entire Index is built on. Second, with $n \approx 2$ the curve is steep: the concentration needed to climb from one-tenth to nine-tenths of full response is a ratio of 9^(2/n) = 9, a mere ninefold window. Put numbers to it. Measuring concentration in units of EC_{50} and setting $n = 2$, occupancy $\theta = C^2 / (1 + C^2)$ climbs from 0.20 at half the EC_{50}, to 0.50 at the EC_{50} exactly, to 0.80 at twice it — so merely doubling the dose across the midpoint carries the burn from one-fifth of full to four-fifths. Stretch the span from a third of the EC_{50} to three times it and θ runs the whole 0.10-to-0.90, the ninefold window itself. Most of the perceptual rise from *barely warm* to *genuinely hot* happens across less than one order of magnitude in concentration. That is the cliff every cook feels and few can name — the reason "a little more chili" so often becomes "too much" in one shake. It is also the reason a cultivar held to a known, reproducible potency is worth so much: when the working window is ninefold, the difference between a weighed cultivar and a guessed one is the difference between landing on the curve and falling off it.

Reading the Index at the stove — from IFI-H to EC50 and back

The Index gives every heat a single number, IFI-H, defined exactly as

> IFI-H = 1,000 × EC_{50}(reference) ÷ EC_{50}(candidate).

Read it backwards and it becomes a working tool. A cultivar's EC_{50}, expressed relative to the reference, is simply 1,000 ÷ IFI-H. Phoenix carries IFI-H 4,200, so its EC_{50} is 1,000 ⁄ 4,200 = 0.238 — it drives TRPV1 to half-activation at a little under a quarter of the reference concentration. Ember at IFI-H 1,570 sits at EC_{50} = 1,000 ⁄ 1,570 = 0.637. Divide one by the other, 0.637 ⁄ 0.238 = 2.68, and you have a fact you can cook with: to reach the same half-occupancy, you need 2.68 times the capsaicinoid concentration of Ember that you would of Phoenix The Index is loudness, and loudness is the inverse of how much you must add. The whole ladder, converted this way, reads as a table of how little of each cultivar it takes to speak at the same volume.

CULTIVAR	SKU	PREDICTED SHU	IFI-H	EC_{50} (REL. TO REF)	AXES
Ember	SS-0087	866K	1,570	0.637	2/8
Kindle	SS-0089	1.25M	2,000	0.500	3/8
Cinder	SS-0088	1.05M	2,290	0.437	3/8
Blaze	SS-0091	1.72M	2,630	0.380	4/8
Wildfire	SS-0092	1.96M	2,930	0.341	5/8
Scorch	SS-0093	2.21M	3,220	0.311	5/8
Firestorm	SS-0094	2.47M	3,490	0.287	6/8
Solara	SS-0095	2.72M	3,980	0.251	6/8
Phoenix	SS-0096	2.95M	4,200	0.238	8/8

Notice that IFI-H and raw Scovilles track closely but not perfectly — Kindle out-measures Cinder in SHU yet reads quieter on the Index. That gap is not an error; it is information. IFI-H weights perceived loudness, and loudness depends on the ratio of capsaicin to dihydrocapsaicin, which the length and saturation of the acyl chain set at the enzyme. Capsaicin arrives fast and sharp at the front of the mouth; dihydrocapsaicin arrives slower and sits deeper and longer. Two pods of identical total capsaicinoid mass can therefore score differently on the axis because they draw their burn with a different hand. When we design a cultivar, we are tuning that ratio as deliberately as the total — which is why the apex of the ladder registers not only the highest heat but a full eight-axis chord:

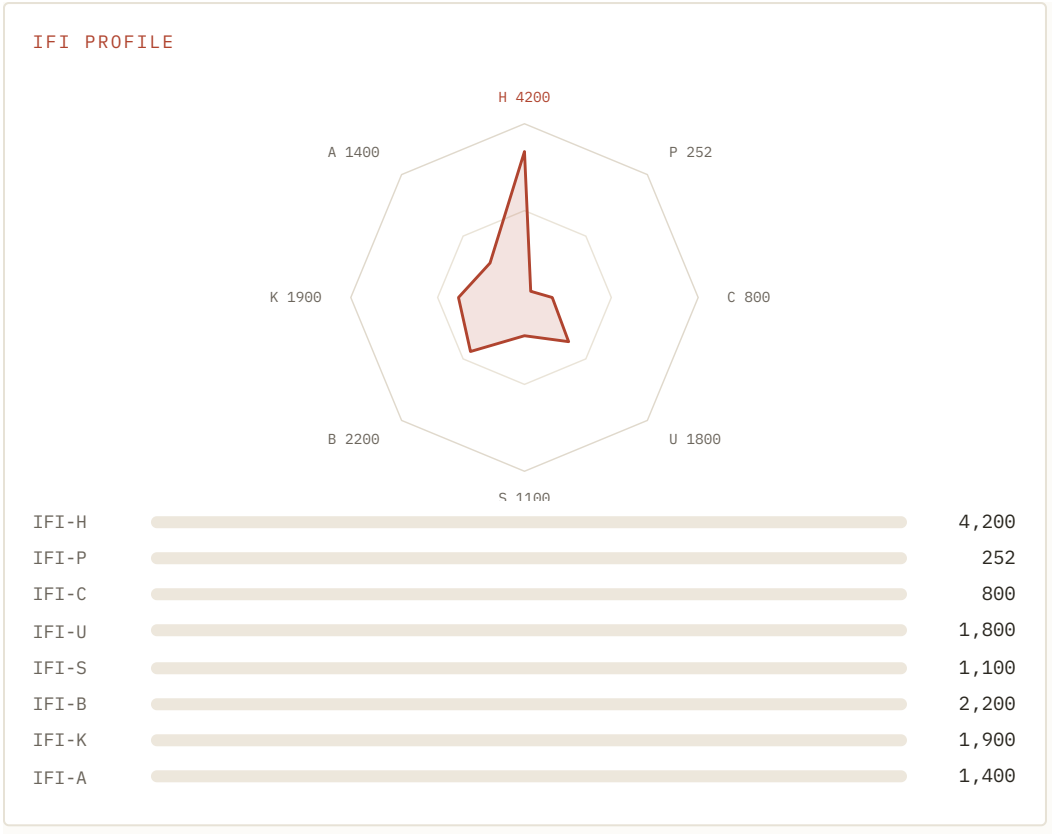

On a scorecard normalized to a per-axis maximum of 5,000, IFI-H 4,200 sits at 84 percent of the heat axis's ceiling — the loudest plant we grow, and still short of the theoretical top of the scale. Read the other seven lines and the same fruit speaks at every receptor at once: a pungency of 252 on the mustard-and-wasabi channel, a kokumi of 1,900 in the round savory register, an astringency of 1,400 in the tannin dimension. Each axis is calibrated against its own defined reference compound, so scores across axes each mark activation on that channel's own scale, not a single common yardstick. It is the only cultivar in the catalog that registers on all eight axes, and when it enters a dish it brings not merely fire but a whole built chord of taste — one more reason its half-gram does so much work in a pot, and one more reason it is weighed, never guessed.

Two currencies — why we dilute in Scovilles and read in Index points

Here is the subtlety that separates a cook who hits a target from one who hopes. Capsaicinoid concentration is linear and additive: combine ingredients and the Scoville value of the mixture is a clean mass-weighted average, $\mathrm{SHU(dish)} \approx \Sigma(m_i \cdot \mathrm{SHU_i}) / \Sigma m_i$. Index points are *not* additive this way, because IFI-H is a perceptual score and perception saturates — the Hill curve flattens as the channel fills. So the map from Scovilles to Index points is compressive, and it is compressive by an enormous factor. The table below reads the same ladder as two currencies and asks how many Scovilles it costs to buy one Index point in different regions of the scale.

REGION OF THE LADDER	PREDICTED SHU SPAN	IFI-H SPAN	SHU PER IFI-H POINT
Base of the scale (plain pot → a drop of Phoenix in a kilo)	0 → 2,935	0 → ~90	~33
Top of the scale (Ember → Phoenix)	866,000 → 2,937,948	1,570 → 4,200	~788

Near the bottom, roughly thirty-three Scovilles buy an Index point; near the top it takes nearly eight hundred — a compression of about twenty-four to one across the useful range, and far steeper still at

the very foot of the scale, where a shishito near IFI-H 40 and a comfortable kitchen chili are separated by only a few hundred Scovilles. The practical lesson is a two-step. *Dilute in Scovilles, convert to Index points.* Do the mixing math in the linear currency, where a mass-weighted average is honest, then read the result back onto the ladder to know what you will taste. Skip the first step and reason in Index points directly and you will overshoot wildly, because you will have treated a saturating perceptual scale as if it added like mass. It does not.

One worked example makes the discipline concrete. Say you want a finishing oil at a confident IFI-H 700, which the ladder places near 200,000 Scovilles, and you build it from a cultivar at IFI-H 1,400, roughly 700,000 Scovilles, bloomed into neutral oil. Dilution is linear in Scovilles, so the chili's mass fraction is simply the ratio of the two Scoville values: 200,000 ⁄ 700,000 = 0.286. Nearly a third of the whole mass would have to be solid chili — far too much for anything pourable — which tells you at once to either start from a milder cultivar or accept a thick paste you will dose by the drop. Run the same arithmetic from a cultivar near 50,000 Scovilles and the ratio exceeds one: you cannot *reach* 200,000 by blooming it at any sane fraction, which tells you the opposite — reach up a band. Before a single chili is toasted, the arithmetic, worked in Scovilles and read back onto the ladder, has already told you which shelf to reach for and roughly how much to weigh.

A drop becomes dinner

The gap between the two currencies is widest at the extreme, and it is the whole reason the top of the ladder is dosed not in spoons but in drops. Our model predicts Phoenix at approximately 2,937,948 Scovilles. Stir a single gram of it into a kilogram of base and the mixture's predicted Scoville value follows the mass-weighted average, (1 × 2,937,948) ⁄ (1,000 + 1) = 2,935 Scovilles — a gentle, sociable warmth near IFI-H 90, nowhere near the raw terror of the pod. By the standard conversion, in which pure capsaicin registers about 16,000,000 Scovilles and one part per million of capsaicinoids is worth roughly sixteen, that finished kilo carries only 2,935 ⁄ 16 ≈ 183 ppm of capsaicinoids — a fifth of a gram spread through a whole pot. A single gram of the fiercest plant we have ever grown, released into dinner for a family, lands as hospitality.

Seen from the pod's side, the same convention is sobering: our model's prediction of 2,937,948 SHU implies Phoenix is designed to carry capsaicinoids at something near eighteen percent of its tasted mass — a load no unmodified pepper approaches, where even the fiercest wild pods sit nearer one or two percent. That is what it means to have designed a plant gene by gene past a ceiling nature never reached, and it is precisely why a cultivar like this is weighed to the milligram and cooked by the drop. Potency this far up the scale is not a seasoning you shake; it is a reagent you dilute.



Phoenix (SS-0096) on the stage in the controlled chamber, where it is weighed to the milligram. At a predicted heat of approximately 2,937,948 Scovilles, it is cooked with a pipette, not a spoon — one gram carries the heat of a thousand grams of a kitchen chili.

The same arithmetic lays out a whole dosing ladder for a single apex cultivar, and it is worth seeing how linear it stays once you are working in the right currency. Each line below is a gram-count of Phoenix carried into one kilogram of otherwise heatless base, with the mixture's predicted Scovilles and its capsaicinoid load in parts per million. Double the chili and you double the heat — exactly, because Scovilles add.

- **0.1 g / kg** → 294 SHU → ~18 ppm: a whisper, a background warmth under other flavors.
- **0.5 g / kg** → 1,468 SHU → ~92 ppm: the floor of noticeable heat, an easy table.
- **1.0 g / kg** → 2,935 SHU → ~183 ppm: the drop-in-a-kilo above, near IFI-H 90.
- **2.0 g / kg** → 5,864 SHU → ~366 ppm: a confident everyday warmth.
- **5.0 g / kg** → 14,617 SHU → ~914 ppm: a bright, declarative heat.

That is the quiet power of cooking to a number. Because Phoenix's potency is anchored to the receptor and held batch to batch to a coefficient of variation of seven percent or less — against the forty-two percent of unselected wild-type Capsicum — the line above is reproducible. You can weigh to it. No chili-by-eye recipe in history could promise that the pot you cook will land where the last one did.

Heat stays — thermal stability and its honest edge

Capsaicin does not cook off. Its decomposition point sits above 240 °C, hotter than any braise, most roasts, and ordinary frying, so where many aromatic compounds volatilize and fade over a long cook, capsaicin simply *stays*. A chili simmered three hours is very nearly as hot at the end as at the start; the fire has only redistributed through the dish. This thermal stubbornness is what makes the Index's dosing math trustworthy at the stove — the Scovilles you put in are, to a good approximation, the Scovilles you serve, so a mass-weighted prediction made before cooking survives the cooking.

The honesty of the number ends at two edges worth naming. The first is scorch. Local temperatures at a dry pan's surface or in open flame can run well past that 240 °C ceiling, and a chili toasted until it blackens does lose capsaicin and turns acrid — so we toast to

fragrance and flexibility, never to char. *Toast, don't scorch.* The second edge is aroma. Capsaicin is stable, but the volatile terpenes, pyrazines, and fruity esters that give a cultivar its perfume are not; three hours of simmering keeps all of the heat and loses much of the top-note scent. A braise built on an aromatic cultivar therefore keeps its fire but goes quiet in the nose unless some of that chili is added late. The stability of capsaicin is a gift to the dosing math and a reminder that heat and aroma keep different clocks.

Building the curve in time

Because the heat survives but the perfume does not, and because a dish has a heat *curve* rather than a single level, you build that curve by choosing *when* the fire enters. Chili bloomed into the base oil at the start gives a deep, foundational, evenly distributed warmth — it has had time to dissolve fully into the fat and spread through every spoonful. Chili added at the end gives a bright, sharp burn that sits forward on the palate, its aroma intact and its heat not yet folded into the background. Most dishes we love use both: a base layer for architecture and a finishing layer for lift. This is why a single IFI-H target is so often hit with two cultivars doing two different jobs — one chosen for the depth of its dihydrocapsaicin-leaning burn in the foundation, another for the fast capsaicin brightness of a late addition. You are not merely setting a volume; you are drawing the shape of the arrival.

The counterweights, measured

Heat can be tamed, and the counterweights are well enough understood to be used with intent rather than panic. There are four, and each works by a different mechanism. **Fat and dairy** strip capsaicin off the receptor: casein sequesters it in micelles and fat dissolves it away from the nerve — the yogurt beside the curry, the cream in the sauce, the coconut milk, the cheese. **Sweetness** occupies and reshapes the perception, recruiting a different axis so the burn shares the stage rather than owning it — sugar, honey, ripe fruit. **Acid** brightens and cuts, and because TRPV1 is itself gated by protons, a squeeze of lime or a splash of vinegar changes the whole context the burn is read in. **Starch** buffers physically, giving the fire a bland surface to spread across — rice, bread, potato. A dish that lands too hot is not ruined; it is one counterweight away from balanced, and knowing which one tells you whether to reach for the yogurt or the lime. This is no accident of our kitchen; it is why the great spicy cuisines pair fire with exactly these things — the raita, the lime over the larb, the rice under everything, a culinary wisdom that turns out to be receptor pharmacology worked out by taste over centuries.



Garlic — not hot itself, but the kitchen's great mediator of heat: its kokumi length gives a burn something to hold onto, and its fat-friendly aromatics bloom in exactly the oil capsaicin needs to dissolve into.

Composing a pot — the mass-weighted blend and its error budget

Put the pieces together and a whole dish becomes a calculation you can trust. Picture a pot for eight, built from three cultivars doing three jobs: a weighed base bloomed into the sofrito, a top-note lift stirred in at the end, and a body of mild chili folded through for floor and flesh. The finished Scoville value is just the mass-weighted sum of what each contributes, and the table makes it plain.

COMPONENT	JOB	MASS	PREDICTED SHU	SHU·G	SHARE OF HEAT
Ember (SS-0087), bloomed in the base oil	foundation	6 g	866,000	5,196,000	73.0%
Phoenix (SS-0096), stirred in at the end	top-note lift	0.5 g	2,937,948	1,468,974	20.6%
Charred poblano, folded through	body + floor	300 g	1,500	450,000	6.3%
Beans, tomato, aromatics, stock	carrier	1,200 g	~0	~0	~0%
Finished pot		**1,506.5 g**		**7,114,974**	**100%**

Divide the summed SHU·grams by the total mass — 7,114,974 ⁄ 1,506.5 — and the pot lands at **4,723 Scovilles**, about 295 ppm of capsaicinoids, a warm and sociable heat in the low reaches of the ladder, the sort of pot our table actually eats. Three cultivars, three points on the scale, one chord: six grams of a flagship for the deep root, half a gram of the apex for color and lift, three hundred grams of a gentle chili for the body — and the arithmetic tells you the note the whole thing will sound before you taste it. Climb a band by adding base-cultivar grams; the heat rises linearly with them.

The number carries an error budget too, and it is better than any single ingredient's. A target held to a coefficient of variation of seven percent has a standard deviation of 0.07 of its value; at a dish aimed near IFI-H 450 that is about 31 points, so ninety-five percent of batches fall within roughly 387 to 513 — a window you taste as *the same dish*. Wild-type Capsicum, at 42 percent CV, would scatter the

same target from about 72 to 828, three bands wide and uncookable to a spec. And blending helps: when three heat sources are batched independently and contribute in similar measure, their random errors partly cancel, and the blend's CV falls toward $7\% / \sqrt{3} \approx 4.0\%$. Composing from several controlled cultivars is not only a way to draw a more interesting heat; it is a way to tighten reproducibility below what any one of them offers alone — provided the errors really are independent, which brings us to the edge of what the number can promise.

The honest limits of a number

Cooking to a number is a freedom, not a guarantee, and it is worth being exact about where the exactness stops. The variance arithmetic above assumes independent batches; three cultivars drawn from the same seed lot, dried in the same run, and stored in the same jar share their errors instead of cancelling them, and the $7\% / \sqrt{3}$ improvement quietly evaporates. The Scoville conversion we leaned on — sixteen per part per million — is a convention tuned to the capsaicin-plus-dihydrocapsaicin blend of ordinary pods; a cultivar whose ratio we have deliberately shifted will read a little off it, which is exactly why the Index measures loudness at the receptor rather than trusting a mass-to-SHU constant. The matrix matters, too: sugar, fat, acid, and salt all bend perceived heat, so the Scovilles in the pot and the fire on the tongue agree in the lab and drift in the mouth. And palates adapt within a single meal — the third bite reads cooler than the first as TRPV1 desensitizes — so no number can fix the experience of heat, only its chemical cause. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

What the number *does* promise is this: that two cooks, continents apart, weighing the same cultivar to the same mass, will put the same fire into the same pot — and that is something no chili-by-eye recipe could ever offer. So we cook, now, up the ladder — the Glow, the Burn, the Fire, and the responsible edge of the Reckoning — vegetarian-first as our table has always been, with seafood and meat written in the same reverent hand. A real dish at every band, each one scored, each one yours to reproduce exactly. Build in Scovilles, read in Index points, and serve the same bowl every time.

Chapter 13 · Recipes · The Glow — *heat as hospitality.*

Most of the cooking a family actually does lives low on the heat axis — in the first three bands, where a pepper is not a dare but a welcome. This is the register we love best and trust most: heat as a warmth you can set in front of anyone, present and kind and in service of everything around it. The dishes gathered here are built for exactly that. They are vegetarian, they are generous, and they are tuned so that *mild* is a promise kept, not a gamble taken.

The glow is not the absence of heat. It is heat held precisely where it does its gentlest work — warming the other flavors into themselves, loosening the shoulders of a meal, arriving on the tongue and then clearing. To hold it there on purpose, the same small warmth batch after batch, we had to understand molecule by molecule what the mouth is doing when it reports a heat that no thermometer in the room can find. The reverence came first, at our table, years ago. The mathematics came to keep the reverence honest.



Backlit greenhouse at golden hour, a single fruit hanging in the light — the low bands are grown with the same care as the extreme tier, because a warmth you can trust is harder to hold steady than a fire.

The receptor that reports a fire that is not there

Capsaicin binds TRPV1, a non-selective cation channel standing in the membrane of the trigeminal nerve endings that lace the surface of the tongue. It is a polymodal sentinel: the same channel opens to temperatures above 43 °C, to protons when the tissue turns acidic, and to the vanilloid molecules a pepper makes. It is, in other words, a genuine heat detector that capsaicin has learned to fool.

The deception is elegant and worth stating exactly, because the whole art of the low bands rests on it. Capsaicin does not warm the tissue. It slips into an intracellular pocket formed between the S3–S4 helices and the S4–S5 linker of the channel — the vanilloid pocket, anchored at a tyrosine the structural biologists number Y511 — and by binding there it lowers the channel's thermal threshold. A gate that should have waited for 43 °C now opens at the 37 °C the mouth already is. Calcium and sodium pour inward, the nerve depolarizes, and the trigeminal line carries a message of burning to the brain though not one cell has been harmed. The fire is real as a signal and absent as an injury. That gap — felt heat without thermal damage — is the entire reason heat can be a pleasure rather than a wound.

Two further facts of the channel govern how a gentle dish behaves. First, the response is graded, not switch-like: how many TRPV1 channels a dose opens, and how far, scales smoothly with concentration, which is what lets a number mean something. Second, sustained calcium influx drives TRPV1 into desensitization — a calcium- and calcineurin-dependent fade that is the molecular form of "building tolerance" within a meal, and the reason a strong burn eventually dulls. At the doses of the lower bands we are nowhere near that cliff. A band-0 or band-1 dish opens a modest fraction of the available channels, holds them in a warm, steady murmur, and lets go cleanly — activation without the heavy desensitization that a harsh dose forces. The glow is, physiologically, a light and reversible touch on a receptor built to shout.

Where the warmth is made — a sentence of enzymes

A pepper builds capsaicin from two rivers that meet at a single enzyme. The first river is the phenylpropanoid branch that runs in every vascular plant — PAL → C4H → 4CL → HCT → COMT — the same chain that makes lignin and the scent of so much of the plant world, routed here toward vanillin and then, by the aminotransferase pAMT, toward vanillylamine. The second river is a short branched-chain fatty acid, assembled on the fatty-acid side by BCAT and KAS into an 8-methyl-branched acyl group and activated as its CoA thioester. The two rivers meet at Pun1/CS — capsaicin synthase, the AT3 acyltransferase — which condenses the vanillylamine head onto the branched acyl tail. That single amide bond is capsaicin.

The character of the burn is written in the tail. When the branched acid carries its double bond, the product is capsaicin proper, and the heat it gives is forward and quick — it blooms at the front of the mouth and clears. When the tail is saturated, the product is dihydrocapsaicin, and its heat is slower, deeper, set further back in the throat and longer to leave. Nature mixes them, and the ratio is the *shape* of the fire. This is not a metaphor we reach for; it is the acyl chemistry doing exactly what the exemplar record says it does — chain length and saturation setting the balance between the two capsaicinoids, and with it the arc of what you feel.

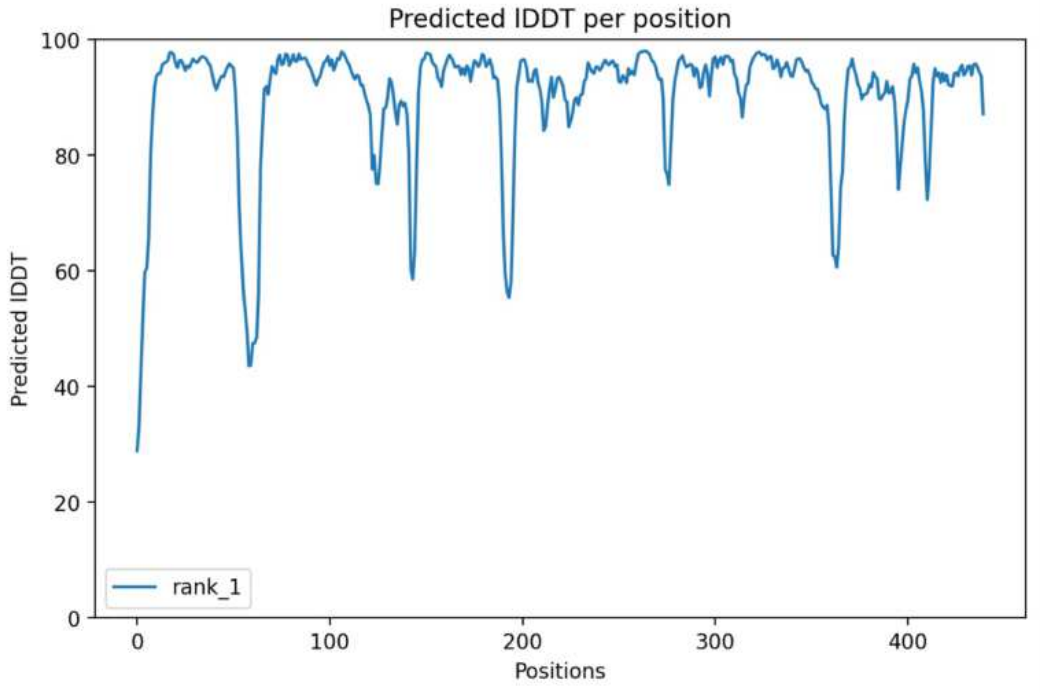

The Pun1/AT3 capsaicin-synthase model, colored by predicted confidence — the active-site cleft where vanillylamine and the branched acyl tail are joined into capsaicin; our flux model locates the heat ceiling upstream at 4CL and HCT, with Pun1 as the final synthase step.

Our flux model found that the real control points sit upstream, at 4CL and HCT/CCoAOMT in the phenylpropanoid trunk — the enzymes that govern how much substrate flows toward vanillylamine. Metabolic Control Analysis across five thousand Monte Carlo draws assigned Pun1 a flux control coefficient near zero; 4CL and HCT together account for 93–99% of pathway control. Pun1/CS is the final capsaicin synthase: it joins the two tributaries, it labels the molecule, but the heat ceiling is set before it arrives. That upstream architecture is what our extreme-tier designs address, and it is what gives the platform its range. Pun1 also carries a double variant — S39L+L345G, two AI-designed substitutions scored by a protein language model and evaluated against our flux model — that refines the terminal step; together with the upstream pathway design, our model predicts peak output of approximately 2.95 million Scoville Heat Units at a batch-to-batch coefficient of variation designed to ≤7%, against the wild-type plant's 42% variation. The low bands are that same upstream mastery pointed the other way: not a weak plant stumbling into mildness, but the same pathway held deliberately at a low, steady setpoint, and the capsaicinoid ratio biased toward the quick-clearing capsaicin form that a hospitable dish wants. Mildness held to ±7% is, if anything, the harder trick — a fire turned down is easy, a *warmth that behaves* every single time is engineering.

One number for loudness — the index at the gentle foot

Reverence can describe the glow; only a number can reproduce it. The Iyengar Flavor Index gives heat its number by characterizing not how much capsaicin a fruit holds but how loudly it speaks to the receptor:

> $\text{IFI-H} = 1{,}000 \times \text{EC}_{50}(\text{reference}) \div \text{EC}_{50}(\text{candidate})$

EC_{50} is the concentration at which a stimulus drives TRPV1 to half of its full response. A quieter pepper needs more of itself to reach that half-way point, so its EC_{50} is high and its score is low; a louder one reaches half-activation at a whisper of concentration, so its EC_{50} is small and its score climbs. The index is anchored so that 1,000 is the

half-activation of the native wild-type reference — the candidate that meets the reference exactly scores 1,000, and every other reading is that candidate's loudness read against the same fixed benchmark.

Invert the formula and it tells you, for any score, the relative concentration a pepper needs to reach half-max: EC_{50}(candidate) ÷ EC_{50}(reference) = 1,000 ÷ IFI-H. The lower bands, worked out, look like this.

PEPPER	BAND	IFI-H	EC_{50} VS REFERENCE (= 1000 ÷ IFI-H)	LOUDNESS VS REFERENCE (= IFI-H ÷ 1000)	CONVENTIONAL SHU*
Shishito	0	40	25.0×	0.04×	~50–200
Fresno	1	120	8.33×	0.12×	~2,500–10,000
Serrano	2	300	3.33×	0.30×	~10,000–23,000
reference anchor	—	1,000	1.00×	1.00×	–
Phoenix · SS-0096	—	4,200	0.238×	4.20×	2,937,948

**Conventional SHU ranges shown for reference peppers are external culinary estimates. For Phoenix (SS-0096), the SHU shown is a computational prediction from our model, not a certified measurement; our primary design specification uses the IFI-H score.*

Read the shishito line slowly, because it is the whole philosophy of this chapter in one row. At IFI-H 40, a shishito must be present at twenty-five times the reference concentration before it brings TRPV1 even to half-activation; put the other way, it is one twenty-fifth as loud as the anchor. That is why a shishito is a warmth and not a burn — the receptor barely registers it until there is a great deal of pepper in the bite. The serrano, at 300, is three and a third times quieter than the reference and a clean, bright 0.30× loud: enough to draw a green line through a dish, never enough to shout over it. And at the far end of the same axis, Phoenix reaches half-activation at under a quarter of the reference concentration — 4.2× as loud as the anchor, the top of a scale whose gentle foot we are cooking on here.

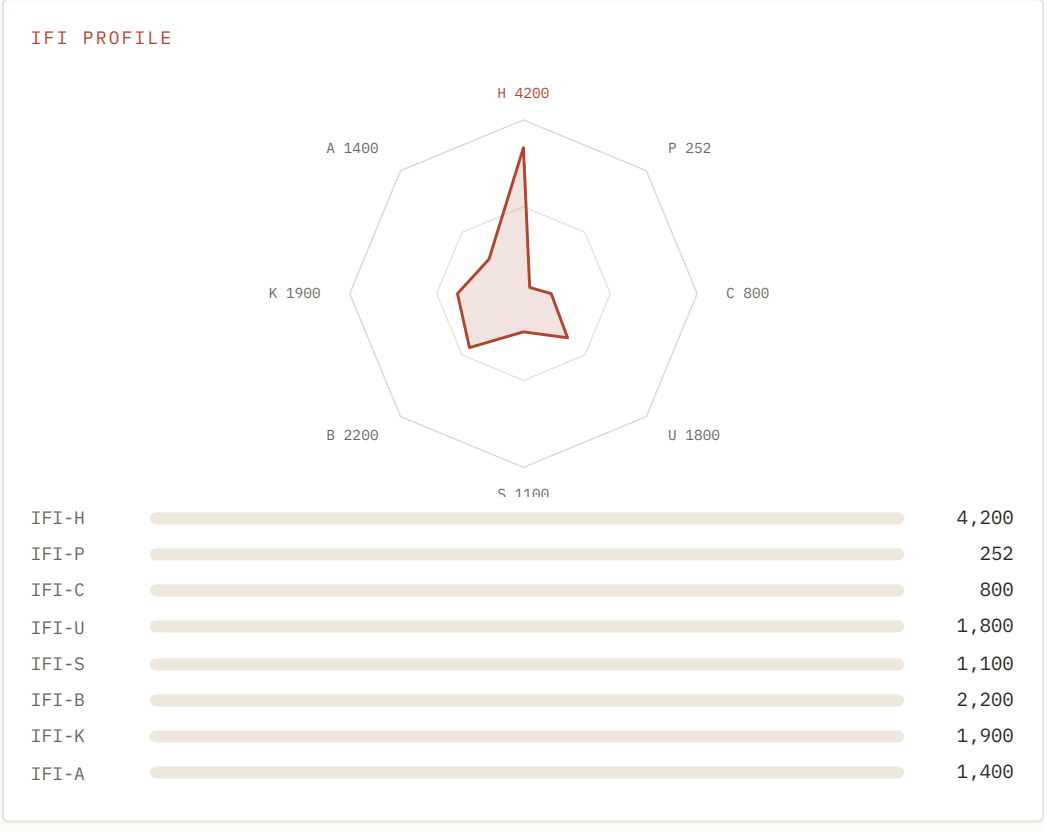

We print Phoenix's full eight-axis profile not to invite it to this table — it would set the whole thing alight — but because it fixes the far end of the ruler. The dishes below live three to four doublings down from that heat line, by design and by arithmetic.

It helps to ground the index in the molecule once, so the scores are not abstract. At capsaicin's own scale, the compound half-activates human TRPV1 near 0.7 micromolar. Converted to a mass concentration — 0.7×10^{-6} mol/L × 305.4 g/mol — that is about 2.1×10^{-4} grams per litre, roughly 0.21 milligrams of capsaicin in a litre of water, already enough to open half the channels it finds. The low bands keep the eaten concentration a long, deliberate way beneath that threshold.

From sugar water to half-activation — reading the ladder

The index did not invent the idea of measuring heat by how far you must dilute it; it sharpened a very old one. In 1912 Wilbur Scoville measured a pepper by steeping it and asking a panel how many fold he had to dilute the extract in sweetened water before the heat vanished from their tongues — the Scoville Heat Unit is, at its root, a dilution-to-threshold. Phoenix, our highest-heat cultivar at a model-predicted 2.95 million SHU, is a design you could dilute nearly three million times and still, at the very edge, detect. The IFI simply replaces the panel's vanishing-point with a defined receptor endpoint: not "where can a taster no longer feel it" but "where does the channel reach half of full response." One is a threshold, the other a half-max; both are inverse measures of potency, which is why they march together up the ladder — and why, where they disagree, the disagreement is itself information.

CULTIVAR · SKU	SHU*	IFI-H	SHU ÷ IFI-H
Ember · SS-0087	866,000	1,570	552
Cinder · SS-0088	1,050,000	2,290	459
Blaze · SS-0091	1,720,000	2,630	654
Wildfire · SS-0092	1,960,000	2,930	669
Scorch · SS-0093	2,210,000	3,220	686
Firestorm · SS-0094	2,470,000	3,490	708
Phoenix · SS-0096	2,937,948	4,200	700

**SHU values for Scoville Splice cultivars are computational predictions from our flux model, not certified measurements.*

If SHU and IFI-H measured the same thing, the ratio in the last column would be a constant. It is not — it wanders from 459 to 708, a spread of better than one and a half fold. That wander is not noise; it is the capsaicinoid *shape* we met at the Pun1 node, made visible. SHU weighs the total mass of capsaicinoids a fruit holds; IFI-H weighs how loudly that particular mixture speaks to TRPV1. Two peppers with the same milligrams of capsaicinoid but different capsaicin-to-dihydrocapsaicin ratios carry the same SHU and a different IFI-H, because the receptor does not read mass — it reads the molecules it is tuned to, each at its own potency. Cinder, low at 459, holds more of its heat in the slower, mass-heavy form; the top of the ladder, near 700, speaks more loudly per milligram. A single

number for heat was never going to be enough, which is exactly why there are two, and why keeping both honest is a form of respect for the person who has to cook with them.

The shape of the dose — why mild stays mild

The deepest reason the low bands can be trusted lives in the shape of the dose–response curve itself. TRPV1 answers a dose the way most receptors do, along a sigmoid described by the Hill equation,

$$\text{response} = [A]^n \div (EC_{50}^n + [A]^n),$$

where [A] is the delivered concentration and n, the Hill slope, runs near 2 for capsaicin at TRPV1 — a cooperative, moderately steep rise. Write the dose in units of the stimulus's own EC_{50} and the curve reads cleanly:

DOSE [A] ÷ EC_{50}	FRACTION OF TRPV1 OPEN
0.2×	0.038 (≈ 4%)
0.3×	0.083 (≈ 8%)
0.5×	0.200 (20%)
1.0×	0.500 (50%)
2.0×	0.800 (80%)

The lower bands are engineered to land a bite on the shallow foot of that curve — around a third of EC_{50}, where perhaps eight percent of the channels answer and the warmth is a murmur. That placement is not only about gentleness; it is about *reliability*, and here the ≤7% CV earns its keep twice over. The sensitivity of felt heat to a wobble in dose is, to first order, $\Delta\text{response} \approx n \cdot \text{response} \cdot (1 - \text{response}) \cdot (\Delta A/A)$. On the foot, with response ≈ 0.083 and our batch CV of 0.07, that is

$$\Delta\text{response} \approx 2 \times 0.083 \times 0.917 \times 0.07 \approx 0.011,$$

so a dish aimed at eight percent activation drifts only between about 0.072 and 0.094 — a swing the tongue cannot find. Hand the same recipe the wild-type plant's 42% variation and the arithmetic turns cruel: 2 × 0.083 × 0.917 × 0.42 ≈ 0.064, so the identical dish ranges from roughly 0.019 to 0.147 — from *is there any heat in this* to *oh, that's got a kick* — on nothing but which fruit the cook happened to pick up. "Mild means mild, every time" is not a slogan. It is the distance between a 7% and a 42% coefficient of variation, read through a Hill slope of 2 on the gentle foot of the curve. The steepness also explains why the brave bands feel like jumps: near EC_{50}, where response is 0.5, the same 7% wobble moves the felt heat by 2 × 0.25 × 0.07 ≈ 0.035 against a large baseline — a change you notice. The foot is forgiving; the shoulder is not. We cook the foot.



A pepper in cross-section — the pale placental ribs, not the seeds, are where capsaicin is made and held; reading that anatomy is the first act of dose control.

Five dishes at the foot of the curve

Here, then, is the glow made dinner — five dishes across three bands, each one placed on that shallow foot on purpose, and each handling its pepper in a different way to put it there. Not one asks anything of the eater but pleasure.

Blistered Shishito and White Beans with Warm Garlic Oil — *band 0, IFI-H 60.* The one-in-ten surprise pepper, gentled into a weeknight supper.

```
Serves 4 · 20 minutes
200 g shishito peppers (band-0 cultivar, IFI-H 40), whole
2 × 400 g tins white beans, drained
6 tbsp good olive oil
4 cloves garlic, thinly sliced
1 lemon — zest and juice
2 tbsp chopped parsley
Flaky salt
```

1. Warm the olive oil and sliced garlic together in a wide pan over low heat until the garlic is pale gold and the oil is fragrant, about 4 minutes. Lift out half the garlic and reserve as a crisp garnish.
2. Raise the heat to high. Add the whole shishitos and blister hard, tossing, until blackened in spots and collapsing, 4–5 minutes.
3. Add the white beans and a good pinch of salt; warm through, crushing a few against the pan to thicken.
4. Off the heat, fold in the lemon zest, half the juice, and the parsley. Taste, adjust salt and acid.
5. Spill onto a warm platter, scatter the reserved crisp garlic and flaky salt, and finish with the remaining lemon.

IFI PROFILE

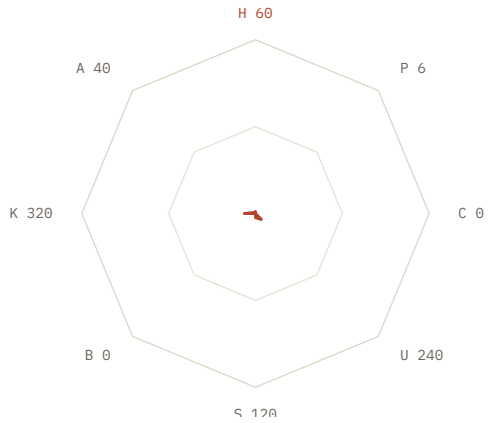

IFI-H	60
IFI-P	6
IFI-C	0
IFI-U	240
IFI-S	120
IFI-B	0
IFI-K	320
IFI-A	40

Charred-Corn Esquites with Fresno and Lime — *band 1, IFI-H 180.* The cup of warm, smoky, tangy corn that is a small act of hospitality in a bowl.

Serves 4 · 25 minutes
4 ears corn, kernels cut off (or 500 g frozen)
2 Fresno chilies (band-1 cultivar, IFI-H 120), finely diced
3 tbsp mayonnaise or thick yogurt
40 g crumbled feta or cotija
2 limes — juice
1 tsp smoked paprika
Small bunch cilantro, chopped
Salt

1. Heat a dry heavy skillet until smoking. Add the corn in one layer and leave it undisturbed to char, then toss, repeating until deeply browned in patches, 6–8 minutes. Patience here is the whole dish.
2. Scrape the hot corn into a bowl. While warm, fold in the diced Fresno, mayonnaise or yogurt, half the cheese, the juice of one lime, and the smoked paprika.
3. Taste and balance with the second lime and salt — it should be tangy, smoky, and gently warm.
4. Top with the rest of the cheese and the cilantro. Eat warm, with a spoon.

IFI PROFILE

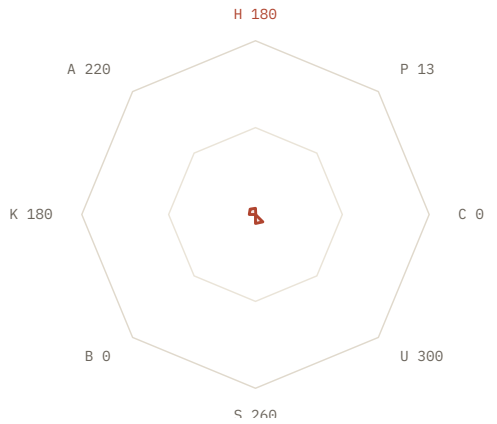

IFI-H	180
IFI-P	13
IFI-C	0
IFI-U	300
IFI-S	260
IFI-B	0
IFI-K	180
IFI-A	220



Charred corn, the sweet and smoky backbone of the esquites — the lower bands of heat are mostly in the business of making other flavors taste more like themselves.

Serrano-Herb Rice with Toasted Cumin — *band 2, IFI-H 300.* A pot of rice that is itself a dish, the serrano's clean bright burn threaded through every grain.

Serves 4 · 30 minutes
300 g basmati rice, rinsed until clear
2 serrano chilies (band-2 cultivar, IFI-H 300), slit lengthwise
1 tbsp ghee or neutral oil
1 tsp cumin seeds
3 cardamom pods, cracked
500 ml water or light stock
Large handful cilantro and mint, chopped
Salt

1. Soak the rinsed rice 15 minutes; drain. Warm the ghee in a heavy pot and toast the cumin and cardamom until fragrant and popping, about 1 minute.
2. Add the slit serranos and let them sizzle 30 seconds — slitting rather than dicing gives an even, moderate heat without hot pockets.
3. Add the drained rice; stir to coat in the spiced fat. Pour in the water and a good pinch of salt, bring to a boil, cover, and drop to the lowest heat for 12 minutes. Rest, covered, 10 minutes more.
4. Fluff with a fork, folding in the herbs. Remove the serranos or leave them for the brave.

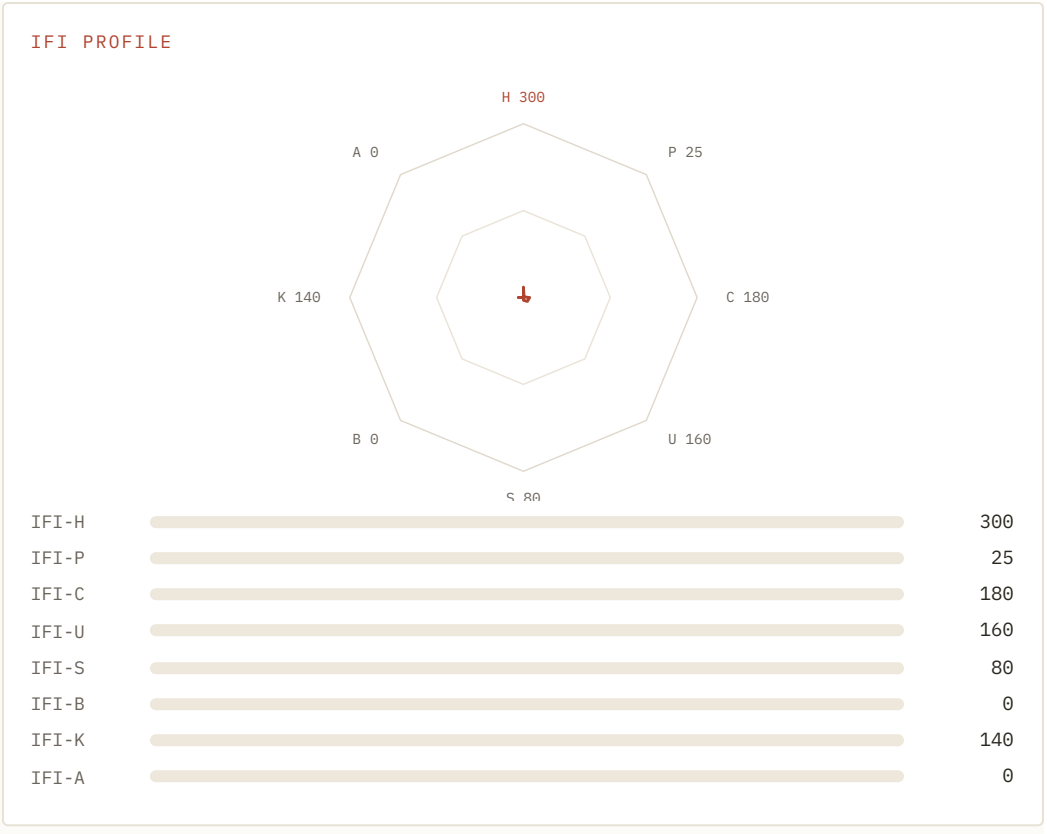

Charred Tomatillo and Serrano Salsa Verde — *band 2, IFI-H 280.* A raw-edged green salsa — smoke, acid, and a clean serrano line — built in a quarter of an hour.

Makes ~500 ml · 15 minutes
500 g tomatillos, husked and halved
2 serrano chilies (band-2 cultivar, IFI-H 300)
½ white onion
2 cloves garlic, unpeeled
Large handful cilantro
1 lime — juice
Salt

1. Heat a dry skillet or broiler until fierce. Char the tomatillos, serranos, onion, and unpeeled garlic hard, turning, until blackened in patches and softened, 6–8 minutes. The char is the dish; do not be timid.
2. Slip the garlic from its skins. Tip everything into a blender with the cilantro and a good pinch of salt.
3. Pulse to a coarse purée — a salsa verde wants texture, not a smoothie. Loosen with a splash of water if it needs it.
4. Finish with lime and salt to taste. It should be bright, smoky, and warm, the serrano a clean green line drawn through the tartness.

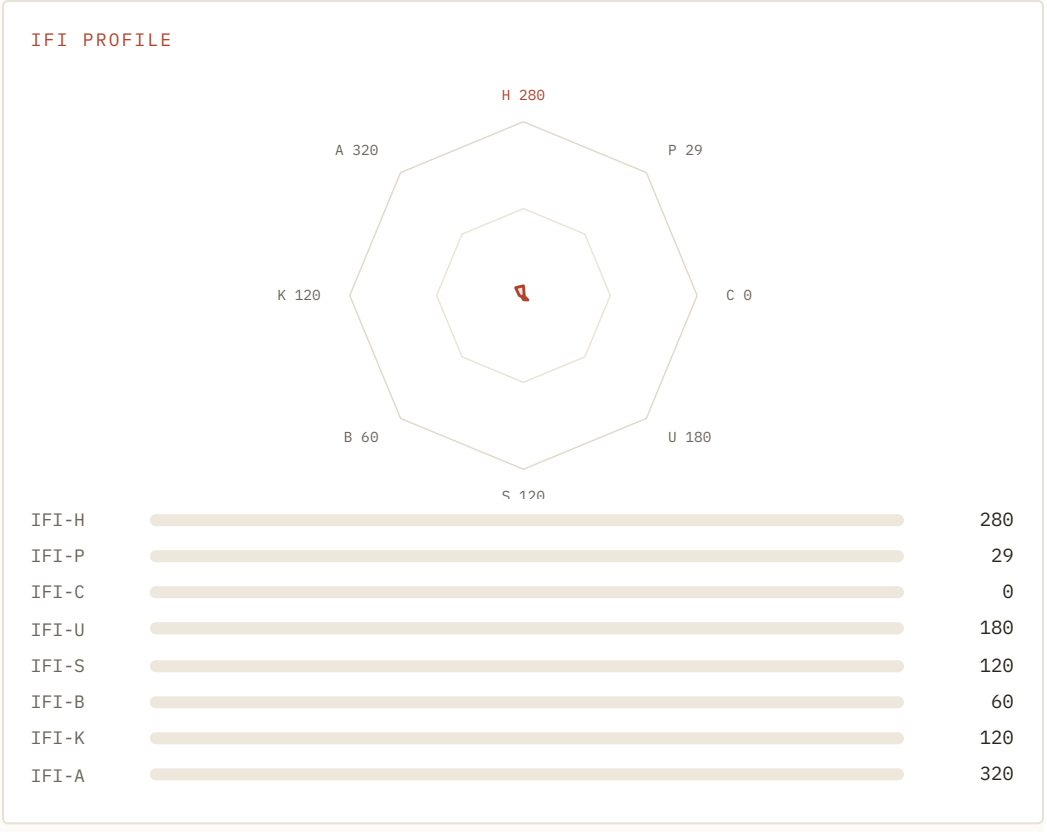

Quick-Pickled Fresno Rings — *band 1, IFI-H 140.* The condiment that lives in the fridge door — bright, sweet-sour, gently warm, good on almost everything.

Makes 1 jar · 10 minutes + 1 hour
4 Fresno chilies (band-1 cultivar, IFI-H 120), sliced into rings
120 ml rice or cider vinegar
120 ml water
1 tbsp sugar
1 tsp salt
1 clove garlic, smashed
½ tsp coriander seeds

1. Pack the Fresno rings and garlic into a clean jar with the coriander seeds.
2. Warm the vinegar, water, sugar, and salt until just dissolved — do not boil hard, or you blow off the brightness.
3. Pour the hot brine over the rings to cover. Cool, then refrigerate at least an hour; they keep two weeks and improve for the first few days.
4. The acid does more than flavor: a quick pickle rounds the Fresno's raw green edge and lets its gentle heat settle into something mellow and sweet-sharp — band 1 made friendly.

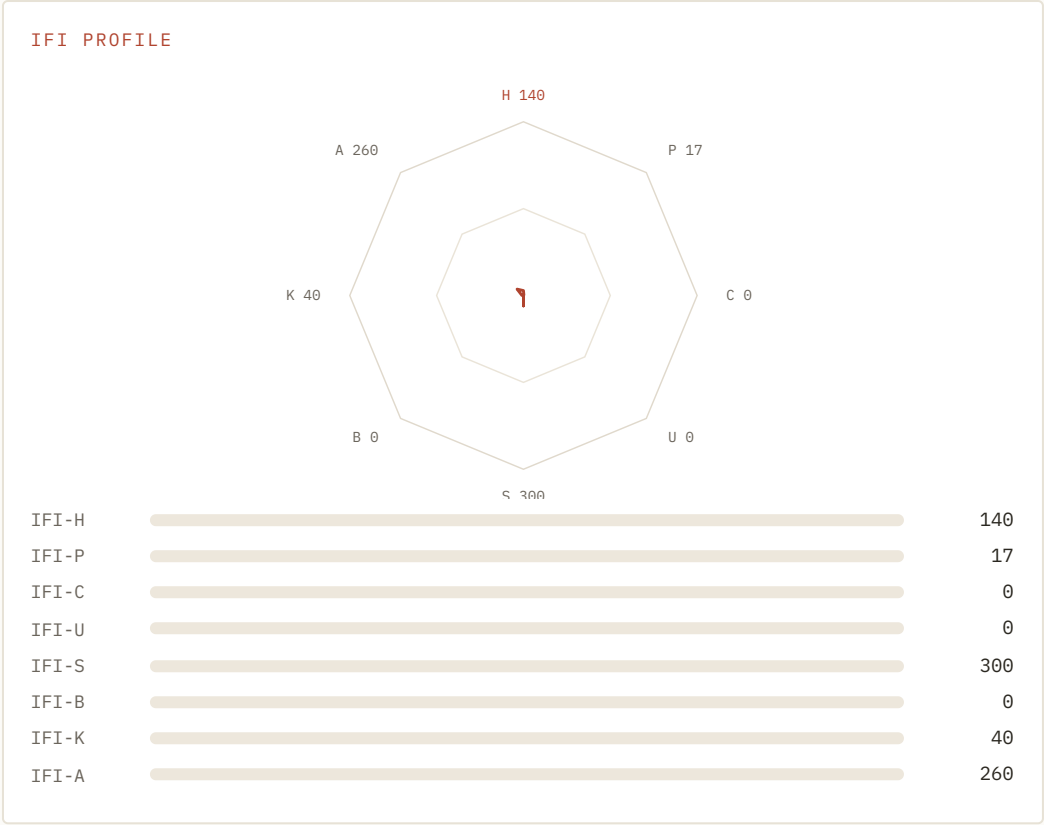

Cut is dose — the chemistry of handling

Notice that every one of these dishes ends with a different IFI-H than the cultivar it was built from, and that the difference is the cook's doing. The shishito cultivar is IFI-H 40 and its dish reads 60; the serrano cultivar is 300 and its slit-rice dish reads 300 exactly; the same serrano, charred and blended into half a kilo of tomatillo, lands at 280. The cultivar sets the loudness per gram; the kitchen sets how many grams speak, and how freely. We can write that plainly. The heat a dish actually delivers is the cultivar's intrinsic loudness

turned by two dials — a loading factor *L* (how much pepper per portion) and a release factor *φ* (how much of its capsaicin the handling frees to reach the tongue):

```
IFI-H(dish) ≈ IFI-H(cultivar) × L × φ.
```

Read against the recipes, the realized product L × φ falls out of the canonical numbers themselves:

DISH	CULTIVAR · IFI-H	HANDLING	IFI-H(DISH)	REALIZED L × Φ
Blistered Shishito & White Beans	shishito · 40	whole, blistered	60	1.50
Charred-Corn Esquites	Fresno · 120	finely diced, raw-folded	180	1.50
Quick-Pickled Fresno Rings	Fresno · 120	sliced, steeped in brine	140	1.17
Serrano-Herb Rice	serrano · 300	slit lengthwise	300	1.00
Charred Tomatillo Salsa Verde	serrano · 300	charred, blended, diluted	280	0.93

The two dishes that land at exactly 1.50 reach it by opposite routes, and that is the lesson. The shishito dish loads heavily — 200 grams of pepper across four plates, about 50 grams a portion, close to a twelfth of the plated weight — but blistered *whole*, its capsaicin stays locked in the placental ribs with little surface freed, so a high *L* meets a low *φ*. The esquites load far less pepper but dice it fine and fold it raw into warm corn, freeing almost all of it — a low *L* meets a high *φ*. Same destination, different arithmetic. The slit serrano in the rice is the control case made visible: slitting perfumes the pot without concentrating the fire into any one grain, holding *L* × *φ* at 1.00 so the dish reads its cultivar exactly. And the salsa, blended for maximum release but poured into 500 grams of tomatillo, is dilution doing the work — *φ* up, *L* down hard, net 0.93.

Underneath the dials is one simple molecular fact: capsaicin is intensely lipophilic, with a partition coefficient (log *P*) near 3.6. It dissolves in fat and oil and barely at all in water. That single property explains the whole handling table. The warm garlic oil of the shishito dish is a solvent — it pulls capsaicin off the blistered skins and distributes it evenly through the beans, smoothing delivery into a gentle wash instead of a few hot bites. The brine of the pickle is the opposite: vinegar and water extract almost no capsaicin, so the heat stays put inside the Fresno rings while the acid mellows how sharply the tongue reads what little escapes — which is why a pickled band-1 pepper eats softer than a raw one. And charring, for all its drama, does not burn the heat away: capsaicin is thermally stable well past normal cooking temperatures. The char ruptures cells and redistributes capsaicin; it does not destroy it. You cannot cook a pepper's heat off. You can only move it, dissolve it, or dilute it — which is the honest, useful truth behind "cut is dose."

One last dial costs nothing and is almost always overlooked: temperature at the moment of eating. TRPV1 is polymodal, remember — it integrates heat and capsaicin on the same gate. A dish served above the channel's 43 °C threshold arrives with warmth already co-

activating the receptor, so its capsaicin reads hotter; the same dish cold sits below threshold, and the capsaicin must carry the whole signal alone, reading milder. This is why the esquites are served warm, not scalding, and why the Fresno pickle — cold, acidic, lightly sweetened — eats gentler than its capsaicin alone would predict. Serving temperature is a heat knob you turn with the oven door, not the knife.

What a single number cannot promise

A number earns trust by being honest about its edges, so here are ours. The ≤7% CV is a promise about the *cultivar* — the fruit from the plant — not about the knife. A careless dice that drops the capsaicin-rich placental rib into one corner of the bowl will hand one guest a bite far above the dish's average, a spatial variance the batch statistic cannot touch. That is precisely why we slit the serrano for the rice rather than chop it: uniform delivery is the cook's half of the bargain. It is also worth retiring an old myth while we are here — the seeds are not the hot part. Capsaicin is made and stored in the pale placental tissue; seeds are hot only by the company they keep.

Three more limits belong on the table plainly. First, TRPV1 desensitizes within a meal, so the third spoonful of a band-1 dish genuinely reads milder than the first; the IFI-H is a per-stimulus measure of loudness, not a running tally of a tongue that is adapting as it eats. Second, people differ — TRPV1 polymorphisms and long habituation shift real sensitivity, and the index is anchored to the receptor's half-maximal response, a population-level biophysical constant, not a guarantee about any one person's palate. Third, the matrix talks back: fat and sugar blunt perceived heat, acid reshapes it, warmth amplifies it, and none of that changes a single milligram of capsaicin. A good number does not pretend these away. It gives you a trustworthy, reproducible starting point — the loudness of the fruit — and leaves the cook, honestly, in charge of the rest.

The glow, kept

Five dishes, three bands, and the same quiet intention running under all of them: that heat can be a kindness, a warmth you give the body on purpose rather than a test you set it. We grow these peppers the way we grow the fierce ones — to hold the most they can, cleanly, and to behave exactly as the knife predicts — because a thing you set in front of the people you love should be dependable before it is impressive. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory, and the precision described here is the precision our model holds and that good seed carries forward. That is the glow we keep: present, gentle, and calibrated so closely that *mild* is simply true. Warm the oil, slit the pepper, fold it into something sweet and green, and set it on the table. The receptor will report its small, borrowed fire, no one will be hurt by it, and the whole meal will taste more like itself.

Chapter 14 · Recipes · The Burn — *where the meal wakes up.*

Band 3 is where heat stops being seasoning and becomes a voice at the table. In the first two bands the warmth is a background hum — a note you notice only if you go looking for it, a warmth under the dish rather than in it. Here it steps forward and takes a seat. The four dishes in this chapter are built around that voice: the fruit-forward habanero and the deep scotch bonnet doing real architectural work, their heat not an accent laid over the plate but one of the load-bearing flavors holding it up. Three of the four are vegetarian, as most of our cooking is; one is a seafood dish, written with exactly the same care, because a prawn charred over fire deserves the same reverence as a bean — not an apology.

A band, though, is a promise. For "band 3" to mean the same warmth in your kitchen as in ours, the phrase has to rest on a chain of real and measurable things: a single receptor, two closely related molecules, a number anchored to that receptor, and a spoonful of fat that is doing more physics than it looks like. Cook these dishes well and you never have to think about any of it. But the reason they work — the reason the number on each card can be trusted — is worth holding in the light first, because once you can read the burn you can design with it.



Mango — the stone-fruit sweetness that is the habanero's oldest and truest partner, sugar reshaping the burn's time-course into something round and long.

The receptor these dishes are written for

Every one of these dishes is, at bottom, a message addressed to a single protein. Capsaicin — the molecule that carries almost all of a chili's fire — binds TRPV1, a non-selective cation channel standing in the membranes of the trigeminal nerve endings that wire the surface of the tongue. TRPV1 is, in its day job, a thermometer: it opens when the tissue climbs above roughly 43 °C, and it opens when the surroundings turn acidic. Capsaicin hijacks that thermometer. It does not warm the tissue by a single degree; it slips into a pocket on the channel's inner face, between two of its membrane-spanning

helices, and lowers the temperature at which the channel decides it is hot. The threshold that used to sit at 43 °C slides down below body temperature, the gate swings open at a heat that is not there, calcium and sodium pour into the nerve ending, and the neuron fires the one signal it knows how to send from that place: *burning.* The fire is real as an experience and fictional as a fact. You are not being harmed. You are being told you are, by a receptor that has been talked into lying.

Three things about that receptor shape everything a cook does with it. First, TRPV1 is built from four identical subunits arranged around the central pore — a tetramer — and a channel that works in four parts tends to respond cooperatively, all-or-nothing-ish rather than gently graded, which will matter the moment we try to dose it. Second, the molecule that opens it is amphipathic: a polar vanillyl head on one end, a long greasy tail on the other, a molecule that is half at home in water and fully at home in fat — which is why a later section is spent entirely on a spoonful of oil. Third, the receptor tires. Sustained calcium flux through TRPV1 drives the channel into a desensitized state; the burn you feel at the third bite is quieter than the burn at the first, not because the dish changed but because the receptor did. A good band-3 dish is built to ride that arc, not to fight it.



A pepper, halved. The fire is not in the seeds but in the pale placental ribs that cradle them — the tissue where the Pun1 enzyme writes capsaicin, molecule by molecule, into the living fruit.

It matters, too, where in the fruit the fire lives. It is not in the seeds — the seeds are only bystanders that happen to sit against the hot part. The capsaicin is made and stored in the pale placental ribs, the septum that runs down the inside of the pod, where the Pun1/AT3 acyltransferase assembles it molecule by molecule. That is why every recipe here says *seeded:* you are not pulling out the heat by discarding seeds, you are pulling out the pith that clings to them. And it is why, in the greens below, a whole unbroken scotch bonnet can perfume a pot for forty minutes without setting it alight — intact, the placenta keeps its capsaicin to itself, releasing aroma and only a whisper of heat until the skin is finally broken.

Two molecules, two timelines

We say "capsaicin" as shorthand, but a pepper's heat is really a duet. The two lead voices are capsaicin and dihydrocapsaicin, and together they carry the overwhelming bulk of any chili's fire. They are near-twins: the same vanillyl head, the same branched nine-carbon tail, built on the same bench by the same enzyme. The pepper makes both by condensing vanillylamine — itself produced from vanillin by the pAMT enzyme — with a small branched-chain fatty acid supplied by the KAS and BCAT machinery, and welding the two together at Pun1/AT3. The only difference between the twins is one double bond in that fatty-acid tail: capsaicin's tail carries it, dihydrocapsaicin's does not. A single point of unsaturation, and nothing else.

That one bond sets the *shape of the burn* — its behavior in time, which is the dimension a single heat number can never show you.

	CAPSAICIN	DIHYDROCAPSAICIN
Acyl tail	branched C9, one double bond (unsaturated)	branched C9, fully saturated
Assembly	Pun1/AT3 · vanillylamine + the unsaturated acyl-CoA	Pun1/AT3 · vanillylamine + the saturated acyl-CoA
Onset	fast, sharp	slower to arrive
Where it lands	the front and middle of the mouth	the back of the throat
In time	a bright, quick spike	a long, low, lingering tail
The dish built on it	the prawns' charred-chili butter	the scotch-bonnet pepper-pot greens

Read the table as a clock, not a thermometer. Capsaicin arrives early and leaves; it is the flash at the tip and middle of the tongue, the heat you register in the first second and that the sweetness of a mango can round off almost as fast as it comes. Dihydrocapsaicin is the slow one — it builds a beat behind, settles at the back of the throat, and stays, which is exactly the heat you want braised into greens that cook for the better part of an hour. The two recipes that lean hardest on this are opposite by design: the prawns ask for a capsaicin-forward chili bloomed quickly into butter so the fire hits bright and clean the instant you bite, and the pepper-pot greens ask for the bonnet's deep dihydrocapsaicin tail to melt down into the pot and surface late, long after the spoon has left your mouth.

This is also where the engineering earns its place in a cookbook. In a wild pepper the ratio of the two twins — and therefore the shape of the burn — wanders from fruit to fruit with the weather, the soil, the day it was picked. The designed cultivars hold it. The two AI-chosen edits that make up the double-variant Pun1 — **S39L** and **L345G** — also help it run consistently batch after batch, so the duet comes out in the same balance every time. You are not just dialing a level; you are fixing a *shape.* A band-3 habanero in these pages is a pepper whose fire has a known rise and a known fall, which is the only reason a recipe can promise you a feeling and not merely a quantity.

Dosing a pot to a band — the arithmetic

Heat on the Index is not a measure of how much capsaicin a pepper holds. It is a measure of how *loudly* that capsaicin speaks to the receptor. The whole Index is built on one relation:

$$IFI(candidate) = 1{,}000 \times EC_{50}(reference) \div EC_{50}(candidate)$$

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response — the receptor's own midpoint. The axis is pinned so that **1,000** is the loudness of the native wild-type reference: its EC_{50} is the yardstick, and everything else is read against it. Invert the ratio so that a *lower* EC_{50} — less material needed to reach half-max — reads as a *higher* score, scale by a thousand, and loudness becomes a number you can put on a recipe card.

Work it the other way and the number tells you something physical. Take the mango simmer's target of **IFI-H 700**. Rearranging the formula, EC_{50}(dish) = 1,000 ÷ 700 × EC_{50}(reference) = 1.43 × EC_{50}(reference): the finished dish needs about 43% *more* material than the reference to drive TRPV1 to its half-point, so it is roughly 0.7× as loud — present, forward, but not overwhelming. The prawns' band-2 chili at **IFI-H 300** needs 3.3× the material (1,000 ÷ 300), making it about a third as loud. The scotch bonnet cultivar at **IFI-H 900** is a quarter louder than the band-3 habanero at **720** — 900 ÷ 720 = 1.25 — which is precisely why the bonnet goes into the pot whole and cautious while the habanero can be minced straight in. The relationship is clean because it is anchored to a receptor and not to a sensation.

Here is the ladder we cook by. The band edges are our kitchen's own shorthand — a coarse grouping of a continuous axis, not a law of nature — but wild-type anchor peppers carry measured reference values; the designed cultivar entries are our model's predictions.

BAND	IFI-H (OUR SHORTHAND)	WHAT THE HEAT DOES	A PEPPER THAT LIVES THERE
1	~40–200	a warmth you could miss	shishito (~40)
2	~200–600	heat you clearly notice	the prawns' chili (300)
3	~600–1,000	heat as a main voice	band-3 habanero (720); scotch bonnet (900)
4	~1,000–2,000	heat as the subject	—
5	~2,000–4,200	the edge — not kitchen seasoning	the superhot flagships, up to Phoenix (predicted IFI-H 4,200)

That last rung is not a place you cook a weeknight dinner. The flagship at the ceiling of the whole axis carries a profile no pepper reaches by accident:

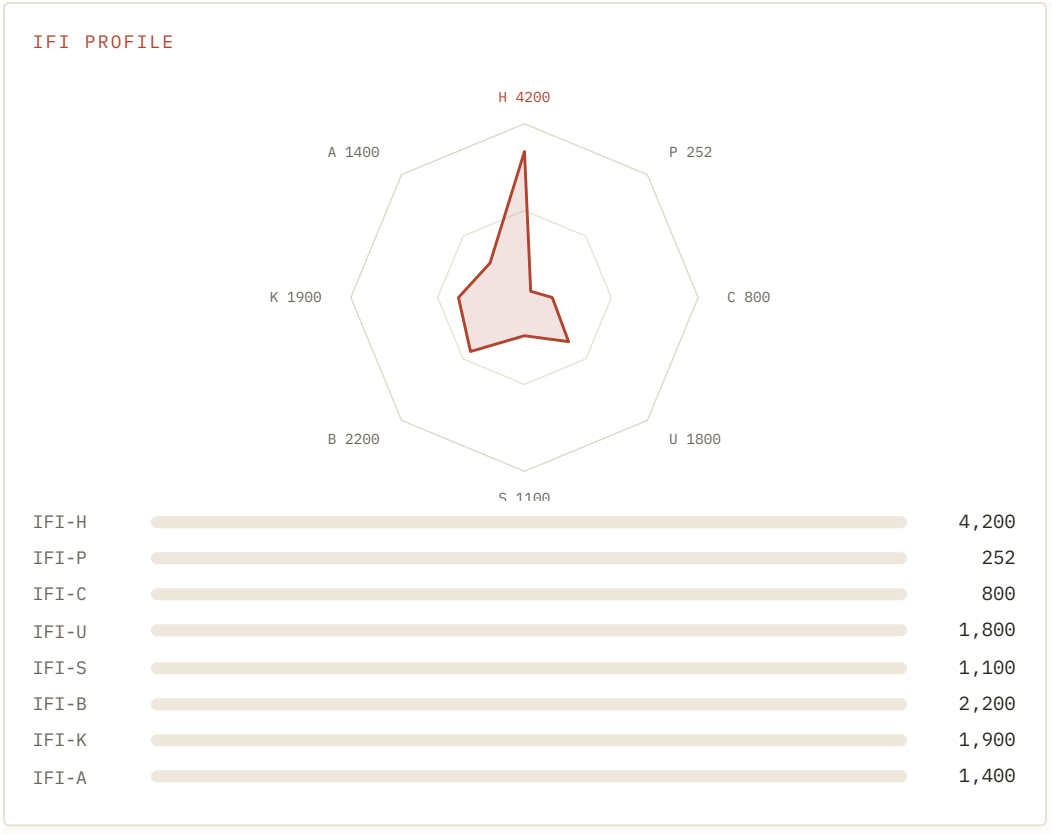

Nearly six times the loudness of the band-3 habanero you will actually cook with (4,200 ÷ 720 ≈ 5.8). We put it here only so you can see the scale you are standing *well* below. Band 3 is a table where you can still taste the mango.

Dosing, then, is dilution you can predict. Over the gentle lower reach of the receptor's response, perceived heat climbs almost in step with concentration — double the capsaicin in the bowl and you roughly double the signal. So hold the pot constant and vary only the pepper: halve the habanero in the mango simmer and the plated heat falls by about half, dropping the dish out of band 3's 600–1,000 and down into band 2; add a third pepper and you climb toward band 4. The number on the card is a *dose,* not a wish.

That near-proportionality is a low-dose courtesy, and it fails as you climb — which is the single most useful thing to understand about cooking with real heat. Because TRPV1 is that four-part cooperative channel, its response follows a Hill curve, close to

$$R(C) = C^2 \div (EC_{50}^2 + C^2)$$

with a Hill coefficient near 2. Walk the dose across the receptor's midpoint and watch the steepness. At $C = EC_{50}$, response is 0.5 — the definition of the half-point. Drop to half that dose ($C = 0.5 \cdot EC_{50}$) and response falls to 0.25 ÷ 1.25 = **0.2**; raise to double ($C = 2 \cdot EC_{50}$) and it climbs to 4 ÷ 5 = **0.8**. A mere fourfold swing in dose — a cook's careless pinch — moves the felt heat from a fifth of full to four-fifths of full. That is the cliff in the middle of the curve, and it is why a dish is dosed to sit on the *shoulder* of its band and never on the edge. It is also why, far up the ladder, each added unit of fire buys less and less: near saturation the curve flattens, and you must pour in disproportionate capsaicin to move the perceived heat at all.

You can see that flattening in the flagships themselves, by asking how much raw Scoville heat it costs to buy one point of receptor loudness:

CULTIVAR · SKU	PREDICTED SHU	IFI-H	SHU PER IFI-H POINT
Ember · SS-0087	866,000	1,570	≈ 552
Cinder · SS-0088	1,050,000	2,290	≈ 459
Kindle · SS-0089	1,250,000	2,000	≈ 625
Blaze · SS-0091	1,720,000	2,630	≈ 654
Firestorm · SS-0094	2,470,000	3,490	≈ 708
Phoenix · SS-0096	~2,950,000	4,200	≈ 700

All SHU and IFI-H values for Scoville Splice cultivars are computational predictions from our model.

The cost per point is not constant — it runs from about 459 to about 708 — and it is not even monotonic. Look at the second and third rows: Kindle is predicted to hold *more* raw Scoville heat than Cinder (1.25M against 1.05M) and yet reads *lower* on the Index (2,000 against 2,290). Two scales, two verdicts, in flat contradiction. This is not an error; it is the whole point of measuring at the receptor. Scoville counts total capsaicinoid by dilution — it cannot tell the capsaicin twin from the dihydrocapsaicin twin, though the two speak to TRPV1 with different force. The Index listens to the receptor, so a pepper whose fire leans toward the louder twin can out-score a pepper that holds more total heat. You cannot convert one number into the other with a single multiplier, which is exactly why a recipe that wants to promise a *feeling* dials in IFI-H and not SHU.

Why the fat is never optional

Notice that every one of these band-3 dishes is built around a fat — coconut milk in the mango simmer and the greens, butter around the prawns, the oil of a chili crisp soaking into eggplant. That is not a flourish and not a way to put the fire out. It is the thermodynamics of delivery, and it is doing real work.



Coconut, opened. Its fat is not a way to cool the heat but a way to carry it — the vehicle in which a fat-loving fire spreads evenly across a whole bowl instead of ambushing a single bite.

Capsaicin is strongly fat-loving — its oil-to-water partition coefficient sits near 1,000-to-1 (a logP close to 3), meaning that given the choice it will leave water for fat a thousand times over. In a watery dish with nowhere greasy to dissolve, capsaicin does not spread; it huddles on whatever few oil droplets it can find and clings to the pepper solids, so one spoonful of a thin broth can be mild and the next a landmine. Bloom the chili into fat first and the capsaicin

dissolves fully and evenly into a continuous oil phase that then disperses through the whole pot. Every spoonful carries the same dose. The fat has not lowered the heat; it has made the heat *fair.* This is the real reason the recipes all begin by letting the minced chili bloom a minute in the oil before anything else goes in — that minute is where the burn becomes even.

Fat and protein also reach the burn after it has started. Capsaicin stuck to TRPV1 is held there by the same greasy tail that let it in; wash the mouth with water and the water slides past it. But the casein in coconut milk and dairy is itself amphipathic — a molecule with greasy regions that love capsaicin and water-facing regions that keep it in suspension — so it behaves like a gentle detergent, surrounding the capsaicin and lifting it off the receptor. That is the mechanism under the old kitchen wisdom that milk beats water on a chili burn, and it is why a coconut braise feels round and carried rather than sharp: the fire is continually being picked up and set back down rather than pinned in place.

And the heat itself survives the stove. Capsaicin melts around 62 °C and does not break down until temperatures far past anything a pot ever reaches, so cooking never destroys a pepper's fire — it only *moves* it. Blooming relocates capsaicin into the fat; a long braise draws the slow dihydrocapsaicin tail out of the placenta and into the liquid; a whole unbroken bonnet, as in the greens, holds its payload in reserve until you choose to burst it. Heat is a way of placing the fire, not of removing it. A cook who understands that never over-seeds in a panic and never forgets the oil.

Four dishes, dosed to the band

What follows is the band made edible. Each dish carries its full eight-axis profile, so you can read not only how hot it is but what else it is doing — the sweetness the mango lends, the umami the miso builds, the kokumi richness the coconut carries underneath.

Habanero-Mango Simmer over Coconut Rice — *band 3, IFI-H 700.* The classic marriage: tropical fruit and fruit-forward fire, each making the other more itself. The habanero's heat is capsaicin-forward and quick, so the mango's sugar can meet it in the same instant, reshaping a sharp spike into something round and long.

> Serves 4 · 40 minutes
> 2 habanero chilies (band-3 cultivar, IFI-H 720), seeded and minced
> 2 ripe mangoes — one diced, one puréed
> 1 onion, finely diced
> 3 cloves garlic, minced
> 1 thumb ginger, grated
> 400 ml coconut milk
> 300 g jasmine rice
> 1 tbsp coconut oil
> 1 lime
> Salt

1. Cook the rice in 450 ml water with half the coconut milk and a pinch of salt; keep warm.
2. In coconut oil, soften the onion, garlic, and ginger until golden, 6–8 minutes. Add the minced habanero and let it bloom in the fat 1 minute — this is where the heat enters the oil and becomes even.

3. Add the mango purée and the rest of the coconut milk; simmer 15 minutes until glossy and thickened. Fold in the diced mango for the last 2 minutes so it keeps its shape.
4. Finish with lime juice and salt. Spoon over the coconut rice.

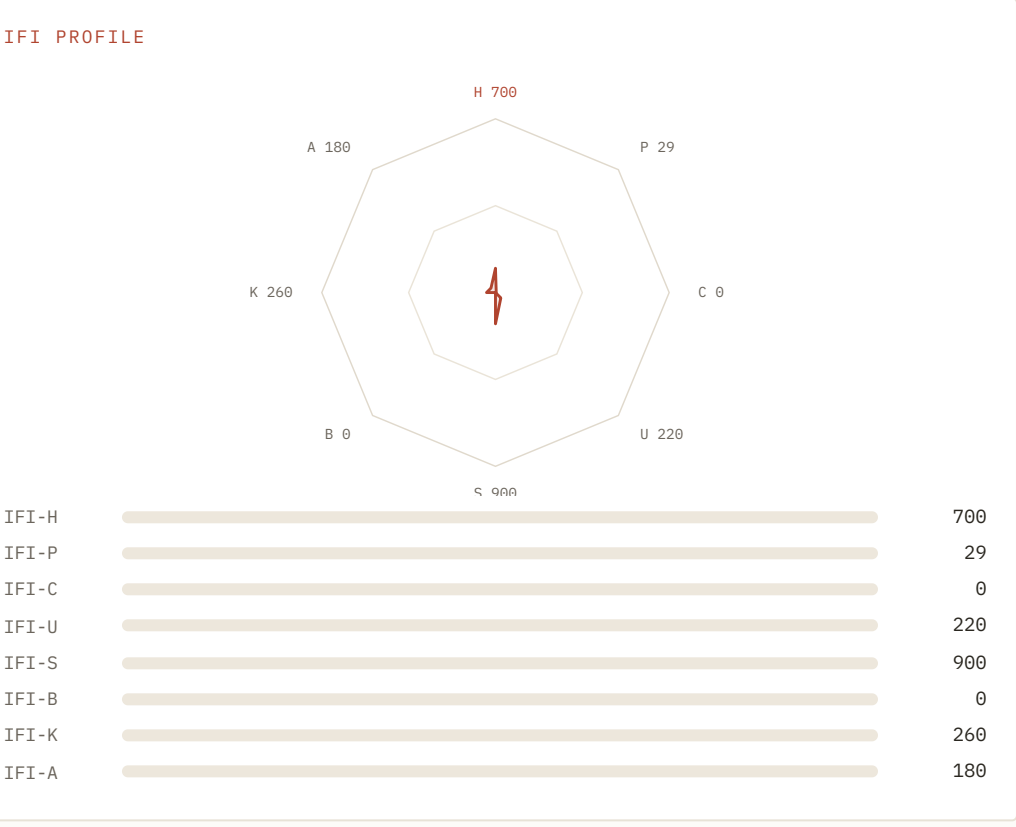

Scotch Bonnet Pepper-Pot Greens — *band 3, IFI-H 850.* A deep, slow, Caribbean-leaning braise where the dihydrocapsaicin tail of the bonnet cooks down into the greens for the better part of an hour, surfacing late and low at the back of the throat.

Serves 4–6 · 1 hour
1 scotch bonnet (band-3 cultivar, IFI-H 900), left whole, plus 1 seeded and minced
800 g mixed hearty greens (callaloo, collard, chard), chopped
1 onion, diced
4 cloves garlic
2 sprigs thyme
1 tbsp coconut or neutral oil
400 ml vegetable stock
200 ml coconut milk
Salt, black pepper

1. In oil, soften the onion, garlic, and thyme until fragrant. Add the minced scotch bonnet and bloom 1 minute.
2. Add the greens in batches, wilting each before the next. Pour in the stock and nestle the whole scotch bonnet into the pot — whole, it perfumes and warms without releasing its full fire.
3. Cover and braise low for 40 minutes, until silky. Remove the whole chili (or burst it in for a real climb).
4. Stir in the coconut milk, simmer 5 minutes to marry, and season. Serve over rice or with flatbread.

IFI PROFILE

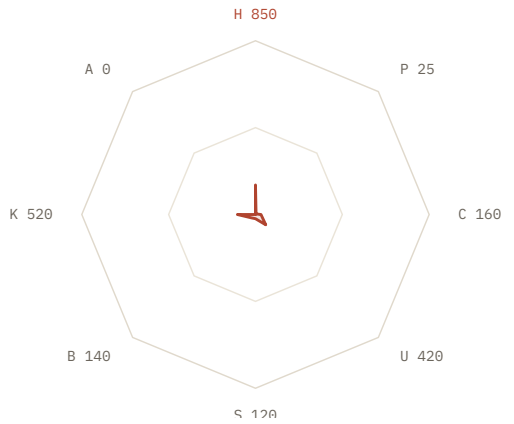

IFI-H		850
IFI-P		25
IFI-C		160
IFI-U		420
IFI-S		120
IFI-B		140
IFI-K		520
IFI-A		0

Grilled Prawns with Charred-Chili Butter — *band 3, IFI-H 600.* Seafood over live fire, written with the same reverence as any dish in this book. The chili here is capsaicin-forward — bright and fast — bloomed into butter so the heat dissolves into the fat and coats every prawn with the same even dose rather than ambushing one and sparing the next.

Serves 4 · 25 minutes
16 large prawns, shell-on, deveined
2 band-2 chilies (IFI-H 300), finely chopped
80 g butter
3 cloves garlic, grated
1 lemon
Small bunch parsley, chopped
Flaky salt

1. Melt the butter gently with the garlic and chopped chili; let it bloom over low heat 3 minutes until the butter is stained red-gold and fragrant. Keep warm.
2. Heat a grill or heavy pan to very hot. Toss the prawns in a spoonful of the chili butter and a pinch of salt.
3. Grill hard, 2 minutes a side, until just opaque and lightly charred — no more, or they toughen.
4. Pile into a warm bowl, pour over the remaining chili butter, and finish with lemon, parsley, and flaky salt. Eat with your hands, and eat it hot — TRPV1 reads warmth and capsaicin on the same channel, so a dish served steaming sits a touch higher than the same dish gone cold.

IFI PROFILE

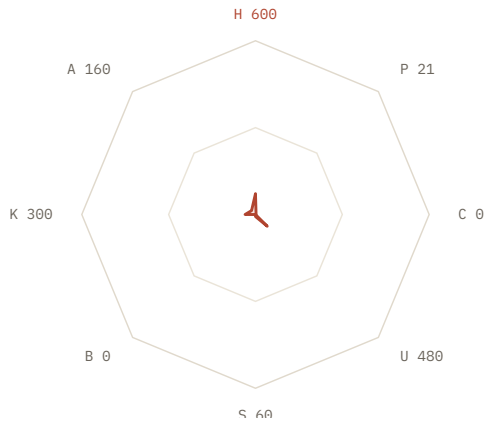

IFI-H	600
IFI-P	21
IFI-C	0
IFI-U	480
IFI-S	60
IFI-B	0
IFI-K	300
IFI-A	160

Miso-Glazed Eggplant with Chili Crisp — *band 3, IFI-H 650.* Umami and fire grown toward each other — the savory depth of miso carrying a real band-3 heat into soft, lacquered flesh, with the chili crisp's oil as the vehicle a fat-loving capsaicin is perfectly at home in.

Serves 4 · 35 minutes
2 large eggplants, halved and scored
3 tbsp white or red miso
1 tbsp mirin
1 tbsp maple syrup or honey
2 tbsp chili crisp (built on a band-3 cultivar, IFI-H 720)
1 tbsp toasted sesame oil
2 spring onions, sliced
1 tsp toasted sesame seeds

1. Score the cut faces of the eggplant in a diamond; brush with sesame oil and roast at 220 °C, cut-side down, 15 minutes, then turn them over.
2. Whisk the miso, mirin, and maple into a glaze. Brush it thickly over the eggplant and roast 10 minutes more, until lacquered and collapsing.
3. Spoon the chili crisp over the hot eggplant so its oil — and its capsaicin, fat-soluble and perfectly at home here — soaks down into the soft flesh.
4. Scatter with spring onion and sesame. The miso's umami grounds the fire; the two read as one deep, warm, savory thing rather than a hot thing and a salty thing side by side.

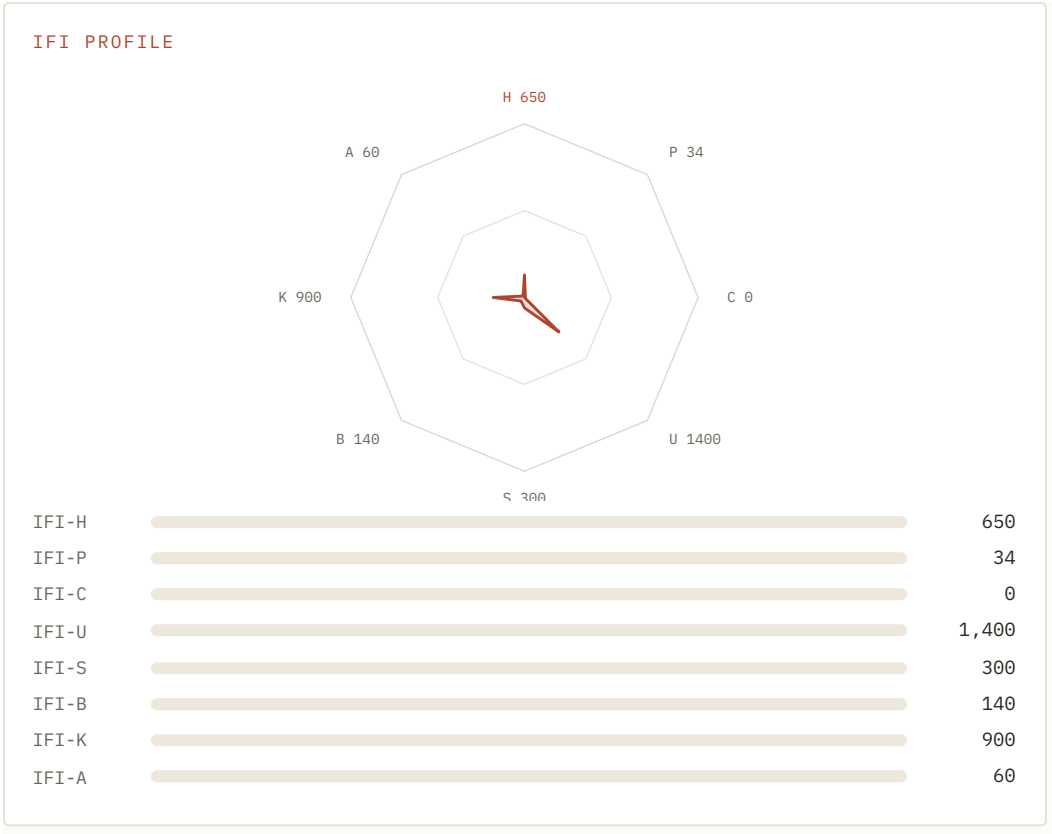

Why these peppers want fruit and ferment

Look again at the pairings and you will see they are not arbitrary. The habanero and the mango belong together at the molecular level, not just on the plate: the habanero's own aroma is built from stone-fruit and citrus esters — the volatiles of its skin genuinely smell of apricot and orange — so the mango does not cover the pepper's heat, it *finishes the pepper's own perfume,* the fruit on the plate completing the fruit in the fire. The scotch bonnet brings those same tropical aromatics to the greens. The miso meets the chili crisp's warmth with a savory depth the tongue reads through an entirely separate receptor, so the two land as one flavor rather than two. This is why band-3 *Capsicum chinense* belongs *with* fruit and ferment rather than against them. You are not taming the heat so much as completing a flavor the pepper only half-stated on its own — and because you are building toward a known number, the completion is reproducible. Cook these peppers with their aromatic kin and the fire stops being a separate event and becomes the hot center of a single, whole, fragrant thing.

Reading the fire by number — why the band holds

Return to the promise we started with: that "band 3" means the same warmth in your kitchen as in ours. A recipe can only make that promise if "a habanero" is a dependable unit — if the pepper you hold weighs the same fire as the pepper we held. This is the quiet place where the engineering pays off in a cookbook, and it is pure arithmetic.

A wild pepper's heat varies batch to batch with a coefficient of variation around **42%** — scatter so wide that one pod can be nearly twice another. The designed cultivars are built to hold a batch CV of **≤7%** (a design target), and the double-variant Pun1 — **S39L** and **L345G** — is engineered to run consistently, so far fewer inactive copies come off the line and the output stops wandering. Now watch what a pot does to that number. A dish is not one pepper; it is the *average* of several, and the scatter of an average of N independent units shrinks by the square root of N.

PEPPERS IN THE POT	ENGINEERED (UNIT CV ≤7%)	WILD-TYPE (UNIT CV ~42%)
1	≤7.0%	~42%
2	≤4.9%	~29.7%
4	≤3.5%	~21%

Two engineered habaneros in the mango simmer bring the dish's heat to a CV under 5% — a dish that lands on band 3 essentially every time you cook it. Two *wild* peppers leave you near 30%: a "band 3" recipe that could plate out anywhere from the top of band 2 to the bottom of band 4, which is to say a recipe that cannot honestly promise a band at all. The precision is not a boast. It is the thing that lets a number on a card survive the trip into your hands. Dose to the shoulder of the band, cook with a pepper whose fire holds to a few percent, and the warmth you taste is the warmth we designed.

What the number does not promise

A single honest number is a gift, and it is also a boundary. IFI-H is the *peak loudness* of a dish at the receptor — and that is all it is. It is worth knowing precisely what it leaves out, because every one of these is a way a well-measured dish can still surprise you.

- **The receptor is yours, not ours.** TRPV1 is coded by a gene that varies from person to person; some diners carry channels a little quicker to open than others. The plate's IFI-H is fixed, but the fire *felt* across a table is not. Dose to the dish, not to the single most heat-shy person in the room — and keep a coconut-rich side for them.
- **The receptor tires as you eat.** Calcium flowing through TRPV1 desensitizes it, so the fourth forkful reads cooler than the first within a single meal. The number describes the food, not the running state of the mouth meeting it.
- **A scalar cannot show the shape.** Two dishes at the same IFI-H can feel different in time — one a capsaicin spike, one a dihydrocapsaicin tail — because the single number reports the peak and says nothing about the rise and fall. The two-molecule table is the part the scalar drops.
- **Technique is part of the dose.** Skip the bloom and the same peppers deliver the same *average* heat as uneven spikes rather than an even glow; the mean is honest, the distribution is not. How you cook a chili is as much a lever as how much you use.
- **Serve it as written.** Because TRPV1 is heat-gated, a dish eaten steaming reads a little hotter than the identical dish gone cold — real for the prawns, worth remembering for anything you let sit.
- **The failure modes are predictable.** Forget the fat and a perfectly dosed dish *stabs* instead of glowing — same number, worse delivery. Over-seed in fear and you gut the fire you designed for. Burst the whole bonnet into the greens and you jump a band in a single stir. None of these are mysteries; all of them are the receptor and the chemistry behaving exactly as they must.

The meal is awake now. These are dishes with a spine — the heat is a voice at the table, not a whisper under it — and three of the four feed you from the plant alone, grown toward the fullest flavor and the fullest nourishment a fruit can carry. Every one is dosed to a

number you can trust, built on a receptor we can read, carried by a fat that makes the fire fair, and held steady by a pepper whose warmth does not wander. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. That is the whole art of band 3: not to make a dish hot, but to make a dish *reliably, beautifully warm* — the same for your table as for ours. Up one more rung.

Chapter 15 · Recipes · The Fire — *the edge of comfort.*

Band 4 is the ghost-and-scorpion register, IFI-H 1,100 to 2,300, and it is the edge of where dinner lives. Below it, heat is seasoning; above it, heat is a stunt. Here, for the span of a few hundred index points, fire is still food — formidable, a little frightening, but something a person who genuinely loves heat can sit down to and finish with their eyes bright instead of streaming. We cook in this register with care and with arithmetic, because the margin is narrow and the reward, when you hold it, is one of the great pleasures of the table: a warmth that fills the mouth and keeps arriving, bite after bite, without tipping into punishment.

Four dishes carry the band. Three are plant-based, which is where our family has always cooked and where we are most at home. One is a slow-smoked brisket, written with full respect and without apology, because barbecue is its own reverent craft — a discipline of smoke, time, and patience that any honest book about fire owes a seat at the table, whether or not it is the way we eat. All four obey the same law, and most of this chapter earns the right to state it plainly: at band 4 you do not guess. You weigh, you convert, and you build the relief into the dish before the fire reaches the plate.



Ember — SS-0087, the opening flagship, here reduced to its highest and best kitchen use: a measured chili oil, its 1,570 of heat tamed into a condiment you can dose by the spoon.

The register, and what the tongue is doing in it

To cook band 4 well you have to know, exactly, what the fire is. It is not warmth. Nothing in these dishes raises the temperature of your mouth by a single degree. What happens is a beautiful deception at the molecular scale, and understanding it is the difference between dosing with respect and dosing with bravado.

Capsaicin binds TRPV1, a non-selective cation channel standing in the membrane of the trigeminal and sensory nerve endings that web the surface of the tongue and lips. In its ordinary life TRPV1 is a heat sensor: it opens when the tissue around it climbs past roughly 43 °C,

and when it opens it lets sodium and calcium ions flood inward, depolarizing the nerve and sending the brain a single, unambiguous message — *this is hot, pull away*. Capsaicin slips into a vanilloid pocket inside the channel's transmembrane bundle and does something quieter and stranger than heating anything: it lowers the temperature at which the channel opens. A receptor that should have stayed shut at mouth temperature is now held ajar at 37 °C. The nerve fires. The brain, reading a channel it has trusted for a lifetime, reports a fire that is not there — and then, because the calcium pouring in triggers the release of calcitonin gene-related peptide and substance P, the body answers the phantom fire for real: the flush, the sweat, the watering eyes, the quickened pulse. The burn is imaginary. The response is entirely sincere.

Two things follow from this mechanism, and both matter in the kitchen. First, the calcium that carries the signal also, over seconds to minutes, desensitizes the channel — a calcium-dependent fatigue that is why the third bite of a well-built dish lands softer than the first, and why relief, once you give it, compounds. Second, because the receptor is a *protein with a binding pocket*, the fire can be carried off by anything that out-competes the tongue for the capsaicin molecule. That is the whole secret of the counterweight we build into every dish here — the dairy, the oil, and the fat that do the carrying.



FIG. 15-1 — Solara (SS-0095) in cross-section. The pale placental tissue along the inner wall, not the seed, is the factory: this is where the Pun1/AT3 acyltransferase completes the final step of capsaicin synthesis, labeling the product that gives the heat of a band-4 rub its character.

Where does the molecule come from? Not the seeds, as the folklore has it, but the pale placental tissue running down the inner wall of the pod — the blanched ribs you can see in any cut pepper. There, vanillylamine condenses with an 8-methyl branched-chain fatty acid at the Pun1/AT3 acyltransferase to form capsaicin. Pun1 is the terminal synthase — it labels the final product — while our flux model predicts that pathway control sits upstream, in the phenylpropanoid steps at 4CL and HCT/CCoAOMT, which set how much precursor flows through. The length and saturation of the acyl chain set the balance between capsaicin and its close cousin dihydrocapsaicin, and that ratio is not a footnote — it is the *shape* of the burn. A capsaicin-forward pod reads fast and high and bright, a

flash at the tip of the tongue; a dihydrocapsaicin-forward pod reads slower, deeper, and longer in the throat. Two pods can carry the same raw pungency and feel like different animals in the mouth — exactly the kind of thing the index exists to capture and the old Scoville scale cannot. Band 4, on that index, is the window from IFI-H 1,100 to 2,300 — the register where a cultivar's heat is enormous but a cook, armed with a scale and the arithmetic, can still place a dish anywhere inside it on purpose.

The index, derived

Everything in this chapter resolves to a single equation, so let us write it out and actually use it:

$$\text{IFI-H} = 1{,}000 \times \text{EC}_{50}(\text{reference}) \div \text{EC}_{50}(\text{candidate})$$

EC_{50} is the concentration of a stimulus that drives its receptor to half of full response. It is an inverse measure of potency: the *less* of something you need to reach half-activation, the smaller its EC_{50}, and the *louder* it speaks. The index inverts that inversion — it divides the reference's EC_{50} by the candidate's — so that a louder compound earns a *higher* score, which is how a number should behave if it is going to mean what your tongue means. The reference is a native wild-type pepper, and it anchors the scale: substitute it into its own formula and IFI-H = 1,000 × EC_{50}(ref) ÷ EC_{50}(ref) = 1,000, exactly, by construction.

From there every score is a potency multiple hiding in plain sight. Rearrange the equation:

$$\text{EC}_{50}(\text{candidate}) \div \text{EC}_{50}(\text{reference}) = 1{,}000 \div \text{IFI-H}$$

Our opening flagship, Ember (SS-0087), carries IFI-H 1,570. So its EC_{50} is 1,000 ÷ 1,570 = 0.637 of the reference's — it reaches half-activation at 63.7 % of the concentration, which is to say it is 1.57× as potent, loud for loud. A band-4 ghost-type at IFI-H 1,400 needs 1,000 ÷ 1,400 = 0.714 of the reference concentration: 1.40× as potent. The score, divided by a thousand, simply *is* the potency multiple. This is the quiet elegance of anchoring to a receptor rather than to a dilution: the units are honest, and they compose.

Why EC_{50} at all, and not the total amount of capsaicin a pod holds? Because a receptor does not respond in proportion to how much agonist is present — it responds along a saturating curve. The relationship is the Hill equation, which we can state cleanly as a fraction of the channel's maximum response:

$$\text{response} = [A]^n \div (\text{EC}_{50}{}^n + [A]^n)$$

where [A] is the agonist concentration and *n* is the Hill coefficient (for TRPV1, reported near 2, reflecting more than one capsaicin bound per channel). At [A] = EC_{50} the response is exactly one-half, which is the definition doing its work. The shape of that curve — gentle at the bottom, steep through the middle, flattening toward a ceiling it never quite reaches — is the single most important fact about cooking with fire, and we will spend the next section inside its consequences. The index samples that curve at its most stable,

most reproducible point, the half-maximal inflection, which is why it is a measurement you can trust batch to batch rather than a number that wanders with every change in assay.

This is also the place to retire a convenient lie, which is that the index is just the old Scoville Heat Unit scale with the decimal moved. It is not, and the difference is culinary, not pedantic. SHU measures *content* — the mass of capsaicinoids in a pod, historically by how many times you must dilute an extract before a panel can no longer detect the burn. IFI-H measures *loudness at the receptor*. The two correlate, because more capsaicinoid does, broadly, mean more signal — but they are not interchangeable, and the gap between them is the capsaicin-to-dihydrocapsaicin shape we met in the placenta. You can convert content to a mass fraction with the standard anchor that pure capsaicin registers about 16,000,000 SHU, so that

```
capsaicinoid fraction (by mass) = SHU ÷ 16,000,000, and 1 ppm
capsaicinoid ≈ 16 SHU.
```

Ember's predicted 866,000 SHU works out to 866,000 ÷ 16,000,000 = 5.41 % capsaicinoid by mass, or 54.1 mg per gram of dried pod. The superhot we grind into the brisket rub, Solara (SS-0095), predicted at 2.72M SHU, is 17.0 % — 170 mg per gram. Those are the numbers that will do the dosing. But watch what happens when you lay content and loudness side by side across the flagship ladder:

CULTIVAR	SKU	SHU	CAPSAICINOID	MG · G⁻¹	IFI-H	× REF	SHU / IFI-PT
Ember	SS-0087	866K	5.41%	54.1	1,570	1.57×	552
Cinder	SS-0088	1.05M	6.56%	65.6	2,290	2.29×	459
Kindle	SS-0089	1.25M	7.81%	78.1	2,000	2.00×	625
Blaze	SS-0091	1.72M	10.75%	107.5	2,630	2.63×	654
Firestorm	SS-0094	2.47M	15.44%	154.4	3,490	3.49×	708
Solara	SS-0095	2.72M	17.00%	170.0	3,980	3.98×	683
Phoenix	SS-0096	2.95M	18.44%	184.4	4,200	4.20×	702

SHU and IFI-H values are model predictions for all Scoville Splice cultivars.

Read the last column. If IFI-H were only a rescaled Scoville, the SHU-per-index-point ratio would be a constant; instead it climbs the ladder from 552 toward 700, and — far more tellingly — it *breaks order* in the middle. Cinder holds 1.05M SHU, Kindle holds 1.25M; Cinder has the *lower* content. Yet Cinder scores IFI-H 2,290 against Kindle's 2,000 — the *higher* loudness. Per Scoville unit, Cinder spends only 459 SHU to buy an index point where Kindle spends 625. The pod with less capsaicinoid speaks louder at TRPV1, because its capsaicinoid blend is weighted toward the molecule the receptor hears best. That inversion is invisible to SHU and obvious to the index, and it is the reason we dose band 4 by IFI-H and not by Scoville: the eater feels loudness, not content, and the index is the only one of the two numbers that measures what the eater feels.

Why band 4 is where we put down the spoon and pick up the scale

Here is the uncomfortable truth the Hill curve forces on us. Capsaicin is a staggeringly potent agonist: it activates TRPV1 with an EC_{50} reported in the sub-micromolar range, commonly near 0.7 µM. Capsaicin's molar mass is about 305 g/mol, so 0.7 µM works out to 0.7×10^{-6} mol/L × 305 g/mol ≈ 0.21 mg/L — roughly **0.21 ppm** in water. That is the half-activation concentration. Now consider what a mouthful of any real band-4 dish actually delivers. The brisket crust carries its capsaicinoid at around 850 ppm — more than *four thousand times* its EC_{50}. On a Hill curve with $n \approx 2$, a concentration four thousand times above EC_{50} drives the response to 0.99999 of maximum. The single channel is pinned wide open. So is it pinned by a neat pod at five-percent capsaicinoid. At culinary concentrations, every genuinely hot food saturates the individual receptor.

If the single channel is already maxed out, what is left for a cook to control? The answer is the thing that still scales: *dose*. Above saturation, intensity is no longer about how far one channel opens — it is about how many thousands of nerve fibers you recruit across the field of the tongue, how long they keep firing, how far the sensation spreads and how slowly it fades. Perceived intensity in this supra-threshold regime follows a psychophysical power law in total delivered capsaicinoid, not the single-channel Hill curve — which is precisely why a dish's heat tracks the *milligrams you put in*, and why a gram or two of a 17-percent superhot genuinely moves the whole experience. This is the mechanistic license for the one rule that governs band 4: **weigh, do not pinch.** A pinch is an uncertainty of perhaps ±0.3 g. At Solara's 17 % capsaicinoid, ±0.3 g is ±51 mg of capsaicinoid landing on a single plate — enough, on its own, to carry a serving clean out of the band it was built for. A scale closes that door. A hand never can.



FIG. 15-2 — the band-4 handoff: a weighed, gloved quantity of superhot leaving the bench. At these heats the dose is a measured mass, because a "pinch" is not a unit — it is a hazard.

So we dose by a capsaicinoid budget, and the arithmetic is short. Take the chili oil. Fifteen grams of dried Ember at 5.41 % carries 15 × 0.0541 = 0.81 g of capsaicinoid, dispersed through 250 mL of oil. Dosed by the half-teaspoon — 2.5 mL, one part in a hundred of the batch — a single serving delivers 0.81 g × (2.5 ÷ 250) = **8.1 mg** of capsaicinoid. That is a real, countable quantity you control with a spoon, which is the entire design of a condiment: the fire is loaded,

measured, and released in increments the eater can read. Now the brisket, which disperses rather than concentrates. Three tablespoons of ground Solara is close to 20 g; at 17 % that is 20 × 0.17 = 3.4 g of capsaicinoid. Spread across a 4 kg brisket, the capsaicinoid sits at 3.4 g ÷ 4,000 g ≈ 8.5 × 10^{-4} by mass — **850 ppm** in the rubbed, crust-bearing fraction of each slice, which is the number we used above to prove the receptor saturates. A band-5 cultivar, dispersed across ten hours and four kilograms, arrives at the tongue as a measured band-4 heat. The superhot was never the danger; the *dose* was, and the dose is a thing you can divide.

The reason this arithmetic is trustworthy at all — the reason a weighed 20 g means the heat you expect and not some other dish — is reproducibility, a manufacturing number before it is a kitchen one. Our cultivars are held to a batch-to-batch capsaicinoid CV of 7 % or less, against the 42 % that a wild landrace will swing across a season. Watch what that spec buys you at the plate. Target a dish at IFI-H 1,400, dead center of the band. A 7 % coefficient of variation puts the one-sigma spread at 0.07 × 1,400 = 98 index points; two sigma, covering about 95 % of batches, is ±196, a window of 1,204 to 1,596. Every bit of that sits comfortably inside band 4. Now run the same target on wild-type material at 42 %: one sigma is 588, two sigma is ±1,176, a window from 224 all the way to 2,576 — a dish that might land as a faint warmth or as an inedible band-5 assault, batch to unpredictable batch.

MATERIAL	BATCH CV	Σ AT TARGET 1,400	±2Σ WINDOW	STAYS IN BAND 4 (1,100–2,300)?
Engineered cultivar	≤ 7%	98	1,204 – 1,596	yes
Wild-type landrace	42%	588	224 – 2,576	no

The table is the case for the whole enterprise, made small enough to taste. You cannot cook precisely with an imprecise ingredient, and at band 4 imprecision is not a quirk, it is the line between a thrill and a trap. One honest footnote, because the math has a sharp edge worth respecting: the dal below targets IFI-H 1,200, only a hundred points above the band floor at 1,100. Even at a tidy 7 % CV, that places one sigma at 84 points, so the 1,100 floor sits just 1.19 sigma away — roughly one batch in eight will dip a touch into band 3. That is a soft, forgivable failure (a slightly gentler dal, never a dangerous one), but it teaches a real lesson in method: with finite variation you aim for the *middle* of a band, not its edge, and you reserve the edges for dishes where a little drift costs nothing.

Ghost-Pepper Dal with Burnt Butter — *band 4, IFI-H 1,200.*

A lentil dal built to carry real fire with the dairy and fat that tame it, so the heat is huge but held. The yogurt here is not a garnish and not a mercy served on the side; it is a molecular counterweight stirred into the architecture of the dish. Casein, the dominant protein in dairy, folds into micelles studded with hydrophobic pockets, and capsaicin — a fat-loving molecule, logP near 3.6, which is the same reason water does nothing but spread the fire — partitions eagerly into

those pockets and off the receptor. Milk out-competes your tongue for the agonist. That is the mechanism folklore gropes at when it reaches for a glass of milk instead of water, and here we simply build it into the bowl.

Serves 4–6 · 50 minutes
1 ghost-type chili (band-4 cultivar, IFI-H 1,400), seeded and minced — handle gloved
300 g red lentils, rinsed
1 onion, diced
4 cloves garlic, minced
1 thumb ginger, grated
2 tomatoes, chopped
1 tsp turmeric
1 tsp cumin
80 g butter or ghee
200 g thick yogurt, to serve
Salt

1. Simmer the lentils with turmeric and 900 ml water until collapsing, 25 minutes. Season with salt.
2. In a separate pan, bloom the cumin in half the butter, then soften the onion, garlic, and ginger until deeply golden. Add the tomatoes and cook to a jam.
3. Add the minced ghost chili to the masala and bloom 1 minute — gloves on, face back, this is band 4. The fat is doing double duty here: it extracts the capsaicinoids off the pod and into the dish evenly, so the fire disperses rather than hiding in hot pockets. Stir the masala into the lentils and simmer 10 minutes.
4. Brown the remaining butter until nutty and just-burnt; pour it foaming over the dal. Serve with a generous spoon of cold yogurt — the casein is not optional here, it is the counterweight that makes the heat a pleasure.

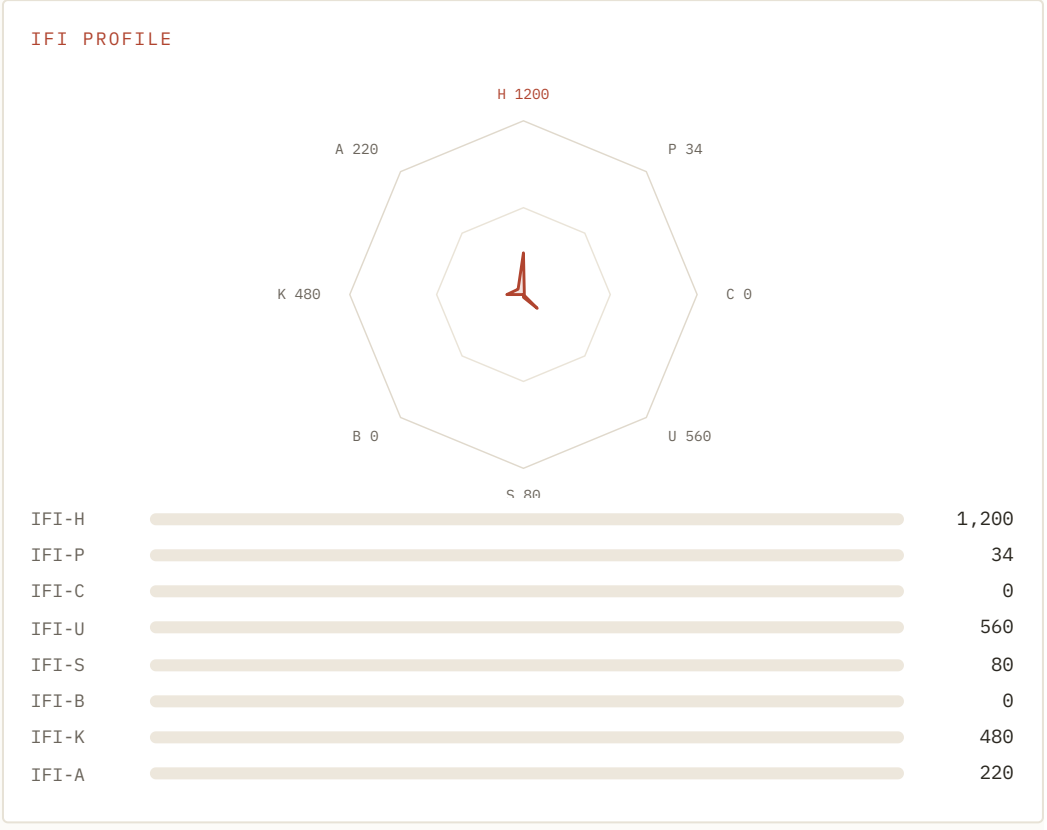

The profile reads the way the bowl eats: the heat at 1,200, but met by a kokumi of 480 from the long-cooked alliums and a umami of 560 from the tomato and lentil, the richness and savor rising to the same height as the fire so the whole thing lands round instead of sharp.

Ember Chili Oil — *band 4, IFI-H 1,570.*

The flagship cultivar at its best and safest — rendered into a measured oil you dose by the half-teaspoon. This recipe leans on a second property of capsaicin we have not yet used: it is not only fat-soluble but heat-stable, holding its structure through the temperatures of normal cooking, so the fire moves cleanly into hot oil and stays there. Oil is the ideal vehicle because capsaicin would rather be in fat than anywhere else; it migrates into the oil and distributes evenly, and once it is there you are no longer dosing a pod but a *solution* — 8.1 mg of capsaicinoid per half-teaspoon, a fire you can meter by the spoon.

```
Makes ~250 ml · 30 minutes + steeping
15 g dried Ember (SS-0087, IFI-H 1,570), stemmed — gloved
250 ml neutral oil
3 cloves garlic, smashed
2 star anise
1 tbsp Sichuan peppercorns (for tingle)
1 tbsp sesame seeds
1 tsp salt
```

1. Warm the oil with the garlic, star anise, and Sichuan peppercorns over low heat to 120 °C; steep 15 minutes to infuse, then lift out the aromatics.
2. Place the dried chili (whole or coarsely crushed for more heat) and the sesame seeds and salt in a heatproof jar.
3. Reheat the oil to 160 °C and pour it over the chili — it should sizzle. Capsaicin is fat-soluble and heat-stable, so this is where the fire moves cleanly into the oil.
4. Cool, cover, and rest overnight. Dose by the half-teaspoon over noodles, eggs, or rice — a little is a lot, by design.

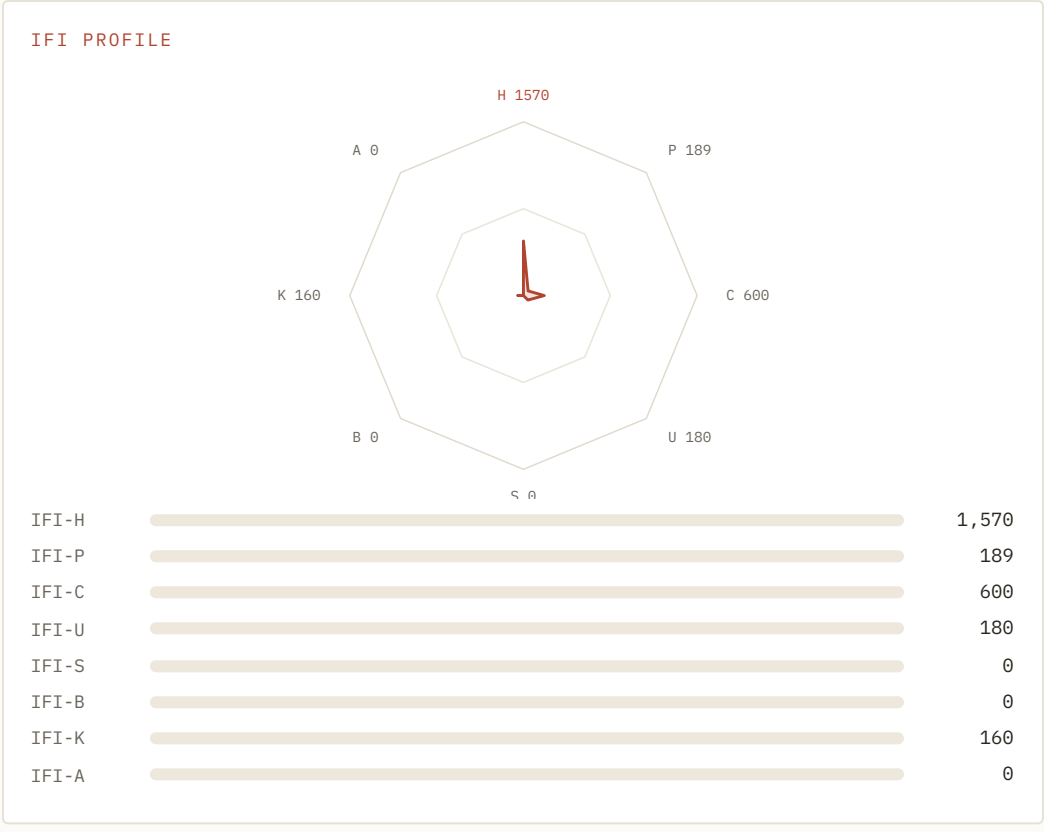

Note the pungency at 189 and the cool-burn at 600: those are the Sichuan peppercorn's hydroxy-alpha-sanshool working a different channel entirely, a tingling, vibratory sharpness that is not TRPV1 at all. The oil is a two-note chord — the steady fire of the index and the fluttering tingle beside it — which is why it reads as *mála* on the tongue and not as a single flat heat.

Smoked Brisket with an Extinction Rub — *band 4, IFI-H 1,400.*

Low-and-slow barbecue, the superhot cultivar ground into a dry rub so its heat is distributed across hours of smoke and fat. We write this one as students of a craft that is not ours to claim — the pitmaster's patience is a reverence we recognize, the same attention to the living material that we bring to a greenhouse at dawn. The science is dispersion, and we did the arithmetic for it above: three tablespoons of a predicted-2.72M-SHU cultivar, roughly 3.4 g of capsaicinoid, spread so thin across four kilograms that a band-5 pepper arrives as a measured band-4 heat. The fat is the counterweight, the way the dal's yogurt was — rendered intramuscular fat is a hydrophobic reservoir that holds and slows the capsaicin exactly as casein's pockets did — and the ten hours of smoke lay a wall of aromatic phenols over the top that reads as savor and depth rather than fire.

> Serves 8 · 10+ hours
> 1 whole beef brisket, ~4 kg
> 3 tbsp ground Solara chili (SS-0095) — a measured amount, gloved
> 3 tbsp smoked paprika
> 3 tbsp brown sugar
> 2 tbsp coarse salt
> 2 tbsp black pepper
> 1 tbsp garlic powder

1. Combine the rub ingredients. The band-5 cultivar is present in small proportion by design — across a 4 kg brisket, three tablespoons disperse to a formidable but eatable band-4 heat. Grind the superhot gloved and at arm's length; a fine superhot dust is a respiratory irritant long before it is a seasoning, and a mask is not fussiness here, it is sense.
2. Coat the brisket all over; rest uncovered in the fridge overnight so the rub sets and the salt penetrates.
3. Smoke at 110 °C over oak until the internal temperature reaches 90 °C and the meat probes like butter, 10–12 hours, wrapping in butcher paper at the stall around 70 °C.
4. Rest wrapped at least 1 hour. Slice against the grain. The fat is the counterweight — it carries and softens the fire the way the dal's yogurt did.

IFI PROFILE

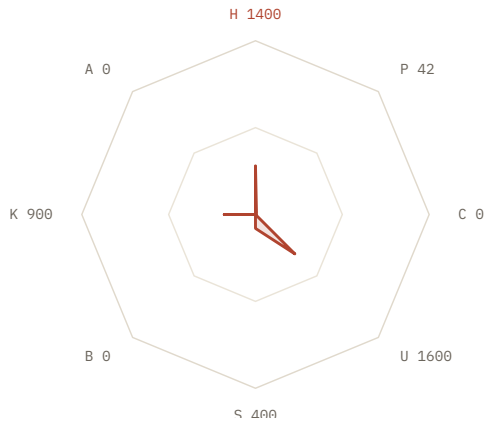

IFI-H	1,400
IFI-P	42
IFI-C	0
IFI-U	1,600
IFI-S	400
IFI-B	0
IFI-K	900
IFI-A	0

The umami here runs to 1,600 — higher than the heat — which is the entire point of the long smoke: glutamate and inosinate concentrating over hours until the savor stands taller than the fire and *holds it up*. A band-4 heat sitting on a band-matched savor and a band-matched kokumi is a heat you can eat a plateful of. A band-4 heat sitting on nothing is a dare.

Ghost-Pepper Peanut Noodles — *band 4, IFI-H 1,300.*

A cold bowl where the peanut's fat and the lime's acid are built, from the first step, to hold a formidable fire. Two mechanisms are at work and both are deliberate. The peanut butter is an emulsion of fat and protein, and it carries capsaicin the way the oil and the yogurt did — holding the molecule in suspension, releasing it slowly, keeping it off the receptor long enough to turn a spike into a glow. And the bowl is served *cold* on purpose. TRPV1 is a thermometer as much as a chemoreceptor; it sums the chemical signal and the thermal one. A hot bowl adds its own warmth to the capsaicin's and reads hotter; a cold bowl subtracts, letting you taste the fire's true shape without the temperature of the food inflating it. Chilling the noodles is not an aesthetic choice. It is a way of reading the heat cleanly.

```
Serves 4 · 25 minutes
1 ghost-type chili (band-4 cultivar, IFI-H 1,400), seeded and minced — gloved
400 g wheat or rice noodles
4 tbsp peanut butter
2 tbsp soy sauce
1 tbsp rice vinegar
1 tbsp maple syrup
2 cloves garlic, grated
1 thumb ginger, grated
2 limes — juice
Crushed roasted peanuts and cilantro, to finish
```

1. Cook and drain the noodles; rinse under cold water and set aside — a cold noodle reads the heat cleaner than a hot one.

2. Whisk the peanut butter, soy, vinegar, maple, garlic, ginger, and the juice of one lime into a loose sauce, thinning with warm water to a pourable cream. The peanut fat is the carrier and the counterweight at once.
3. Stir the minced ghost chili into the sauce a little at a time, tasting toward the target — gloved hands, and trust the arithmetic over bravado.
4. Toss the cold noodles through the sauce. Finish with the second lime, crushed peanuts, and cilantro. The peanut holds the fire in suspension; the lime keeps it bright.

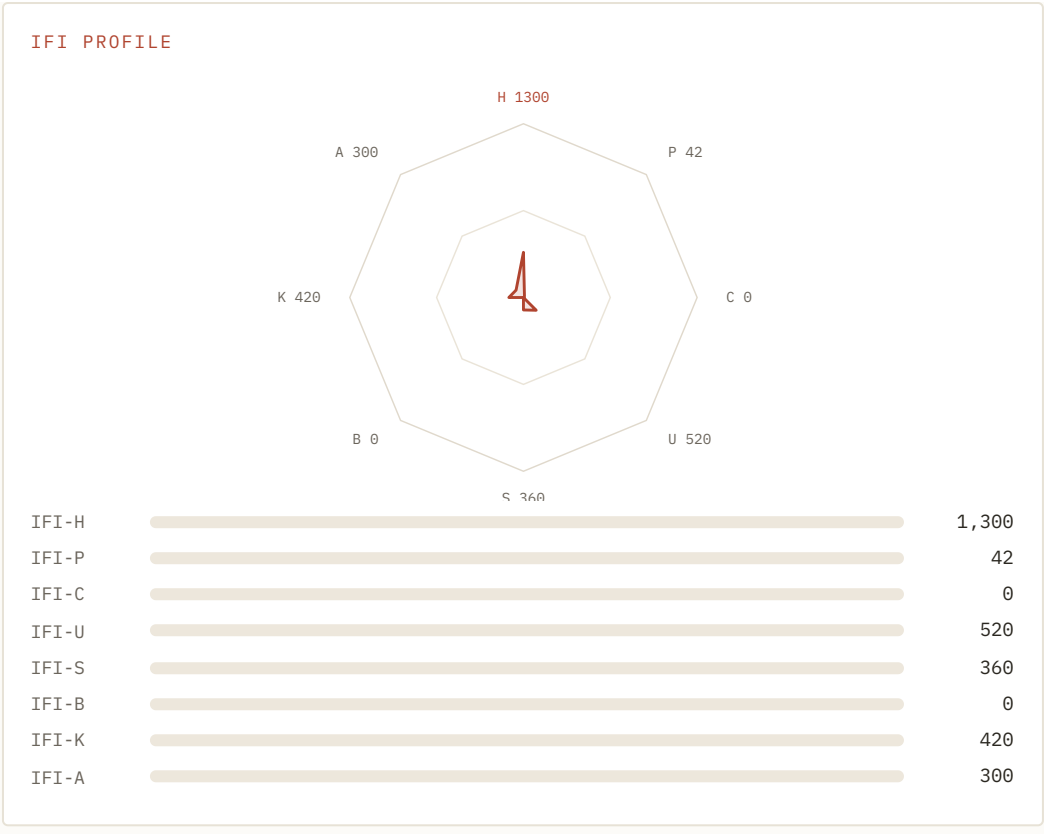

The counterweight is not a garnish

Band 4 is also where weighing replaces guessing, and we mean that literally. At these heats the difference between a dish a chili-lover relishes and one nobody can finish is a gram or two, far finer than a spoon can judge, so we dose these cultivars by mass on a scale and convert through the capsaicinoid arithmetic laid out above rather than trusting a hand. The dried chili in the dal, the ground superhot in the rub, the pods in the oil — each is a weighed quantity, and the recipes give grams precisely because at band 4 a "pinch" is not a

unit, it is a hazard. Precision is not fussiness here; it is the thing that keeps a thrilling heat on the right side of the line between pleasure and punishment.

And every recipe in this chapter obeys one rule that only grows stricter as the ladder climbs toward the ten flagships at its summit: the fire is always met by its specific antidote in the same bite. We can now say *why* that works and not merely that it does. Capsaicin is a hydrophobic molecule bound reversibly in a receptor pocket, and anything with a better hydrophobic home — the casein micelles in the dal's yogurt, the rendered fat in the brisket, the emulsified oil in the peanut sauce — pulls the molecule off the tongue and sequesters it. The ghost-pepper dal is unthinkable without its cold yogurt; the chili oil is dosed by the half-teaspoon because it is built to be; the brisket's three tablespoons of band-5 cultivar are safe only because they are dispersed across four kilograms of fat-rich meat and ten hours of smoke. This is the band-4 discipline in miniature: you may cook with serious fire here, but you must build its relief into the architecture of the dish and never leave the eater to go searching for it. Heat at this level given without its counterweight is not hospitality; it is a trap, and we do not set traps for people we are feeding.

Two honest limits close the chapter, because a number is only as good as what it refuses to promise. First, the index measures the *stimulus*, not the eater. TRPV1 itself varies from person to person — in its genetics, in how thoroughly a lifetime of chilies has desensitized it — so IFI-H 1,400 is a population target, a reliably placed center of gravity, not a guarantee that lands identically on every tongue. A band is a design window, not a contract. Second, the delivered dose is sensitive to things past the weighing: capsaicinoids settle out of a resting chili oil, so the jar wants a stir before each pour or the last spoonful carries a different fire than the first; serving temperature shifts the thermal half of the signal, which is why the noodles go out cold; and fat changes how much capsaicin reaches the receptor at all, which is why the same milligrams read gentler in the brisket than bare. We weigh to remove the largest uncertainty. We respect the rest. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

This is the edge of comfort, and we cook here with gloves on and respect up. Every dish pairs enormous heat with its specific counterweight — yogurt, oil, fat — so the fire is a thrill and never an injury. That is the whole craft of band 4, held in a single discipline: measure the fire, and build its mercy into the same bite that delivers it.

Chapter 16 · Recipes · The Reckoning — *cooking with the ten.*

Band 5 is not a place you cook dinner. It is a place you make *condiments* — concentrated, measured, bottled fire that you then dose, in drops, into dishes that live at a far gentler band. This chapter is short on ingredients and long on arithmetic, because responsibility, at this altitude, is mostly arithmetic. The ten flagships are not food to be eaten neat; they are instruments of extraordinary concentration, and the only honest way to cook with them is to dilute them with respect until a whole dish's worth of fire fits on the tip of a spoon. Heat is a pleasure and never an assault — and nothing in this book carries that line the way these three recipes do, because here, more than anywhere, the distance between pleasure and assault is a number you can write down.

So we will write it down — what the fire is at the level of a single channel on a single nerve ending, the two measurements that let us read it, and how both carry all the way into a jar of honey, a bottle of sauce, and a dish of salt. The lesson underneath every recipe is one sentence: you dose the cultivar, you never serve it.

The ten, and why they end up in a bottle

"The ten" are the proof-of-concept flagships — ten cultivars of *Capsicum chinense* we grew gene by gene to climb a ladder nature stops short of. They begin at a pepper that merely means business and end at a plant that registers on all eight axes of taste at once. Read them as a ladder of increasing computational design scope and axis activation, because that is how they were built.

CULTIVAR	SKU	PREDICTED SHU	PREDICTED IFI-H	AXES
Ember	SS-0087	866K	1,570	2/8
Cinder	SS-0088	1.05M	2,290	3/8
Kindle	SS-0089	1.25M	2,000	3/8
Flare	SS-0090	1.48M	—	4/8
Blaze	SS-0091	1.72M	2,630	4/8
Wildfire	SS-0092	1.96M	2,930	5/8
Scorch	SS-0093	2.21M	3,220	5/8
Firestorm	SS-0094	2.47M	3,490	6/8
Solara	SS-0095	2.72M	3,980	6/8
Phoenix	SS-0096	2.95M	4,200	**8/8**

Not one of these is a thing to put on a plate as itself. A single Phoenix pod is designed to carry roughly 2.94 million Scovilles — our model predicts approximately 2,937,948 SHU — in a few grams of flesh. That is not dinner; that is a reagent. The whole point of designing a fire hotter than any in nature was never to hand someone the pod. It was to have a source concentrated enough that a drop of it, diluted far, could be given as exactly the warmth the eater asked for and no more.

Phoenix earns a second look before we cook, because it is the only one of the ten that brings more than heat to a condiment. At 8/8 it is, by design, a full-spectrum pepper, and even a trace of it deposits faint signatures of the other seven axes into whatever it touches.

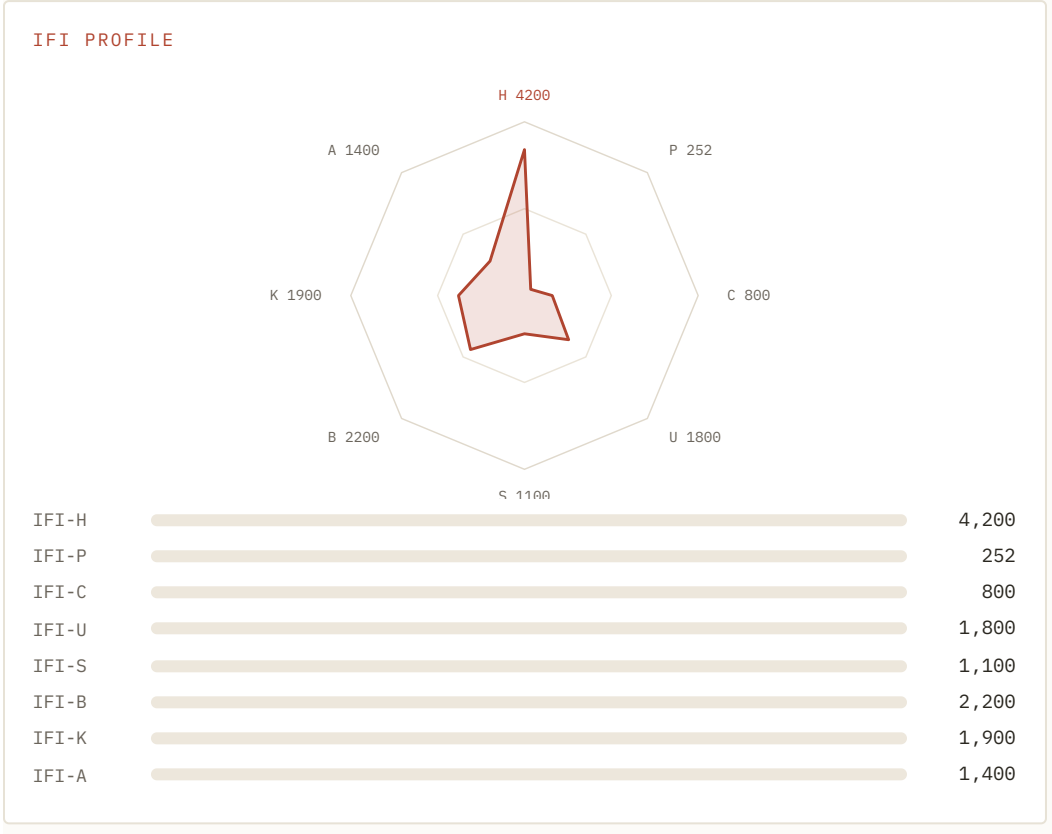

Hold that profile in mind. When we dilute Phoenix five-hundred-fold into a jar of honey, the heat drops to a whisper — but the pungency, the faint astringency, the ghost of its sweetness do not all vanish at the same rate, and that asymmetry is why its honey tastes like more than sweetened fire. We will see the receipt for that at the end.

What the fire actually is

There is no heat in a chili. There is only a molecule that persuades the body a heat is there.

Capsaicin binds TRPV1, a non-selective cation channel standing in the membrane of the trigeminal nerve endings that line the tongue and mouth. TRPV1 is, in its day job, a thermometer: it opens when the tissue around it climbs above about 43 °C, and it opens when the surroundings turn acidic, flooded with protons. Either way, when it opens it lets sodium and calcium ions pour through, the nerve depolarizes, and the brain receives a message it has only ever learned to read one way — *this is burning.* Capsaicin's trick is that it does not raise the temperature at all. It is lipophilic, it slips into the membrane from the fatty side, and it settles into a vanilloid-binding pocket cradled against the channel's voltage-sensing domain, where its vanillyl headgroup hydrogen-bonds to a conserved tyrosine and threonine and its long fatty tail threads into a lipid-facing groove. Bound there, it lowers the channel's thermal threshold — from 43 °C down toward body temperature. The channel now reads the mouth's own 37 °C as if it were scalding. The fire is real to the nerve and absent from the world. That gap is the entire sensation, and it is why a chili keeps burning long after it is swallowed: the molecule stays bound, and the body stays fooled.

Two features of that mechanism matter at the stove. The first is **desensitization.** The same calcium that carries the alarm also, once inside the cell, tells the channel to stand down: calcium-dependent pathways and the depletion of a membrane lipid called PIP_2 cause

TRPV1 to close and resist reopening. This is why the second bite of a hot dish lands softer than the first, why a cook tasting a sauce all afternoon slowly stops trusting their own tongue, and why these recipes insist you judge heat by the scale and the index, never by a palate that is quietly going numb.

The second is **the shape of the burn,** which is chemistry, not metaphor. A pepper's pungency is never one molecule but a blend of capsaicinoids, dominated by two: capsaicin itself and dihydrocapsaicin, together around nine-tenths of the total. They differ by a single double bond in that fatty tail, and they do not burn alike. Capsaicin arrives fast and forward — the tip and middle of the tongue, a bright sharp front. Dihydrocapsaicin is the slow one, blooming later and lower, in the back of the throat, and lingering after capsaicin has faded. Shift the ratio and you shift the *arc* of the heat — its onset, its peak, its long decline — independent of the total Scovilles. Designing a condiment is partly choosing that arc.

That ratio is set deep in the plant, at the end of a pathway we re-engineered. The backbone is the phenylpropanoid branch — PAL → C4H → 4CL → HCT → COMT — that runs in every vascular plant; onto it the pepper bolts four terminal enzymes of its own: pAMT, Pun1/CS, KAS, and BCAT. Vanillylamine, delivered by pAMT, condenses with a short branched-chain fatty acid, built by BCAT and KAS from the branched amino acids, at the Pun1/CS acyltransferase — the capsaicin synthase. Whether that acyl chain arrives unsaturated or saturated decides capsaicin versus dihydrocapsaicin, and with it the shape of the burn.

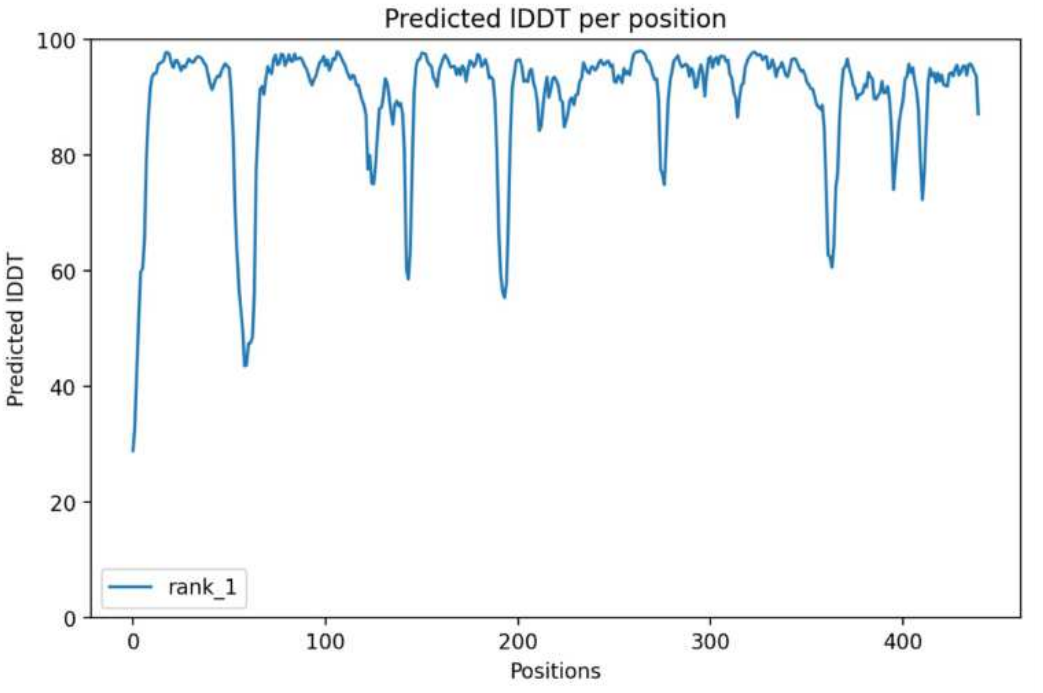

FIG. 16 — the Pun1/CS acyltransferase model, coloured by pLDDT confidence; the active-site cleft where vanillylamine meets its acyl partner is the high-confidence core, and the seat of both edits.

Every prior breeding program chased heat by selecting the hottest-tasting plant and hoping the genes followed. We inverted it: a protein language model scored every single-residue substitution in Pun1 for its predicted effect on catalytic efficiency and fold stability, while our genome-scale flux model traced where pathway control actually lives — upstream in the phenylpropanoid trunk, at the 4CL and HCT/CCoAOMT steps. Pun1 is the final capsaicin synthase; it places the last bond in the molecule, but our flux model predicts control sits

upstream, not at that terminal step. Two substitutions no breeder had ever stumbled into emerged from the protein modelling, and written in with CRISPR they each do a distinct job:

EDIT	WHAT IT DOES	ΔΔG (KCAL/MOL)
S39L	tightens substrate binding at the N-terminus	−2.1
L345G	relieves steric strain in the active-site cleft	−1.8

The number a cook should actually care about is not the peak Scoville those edits are designed to unlock. It is the **consistency.** Wild-type superhots vary batch to batch by around 42% — a coefficient of variation so wide that two pods off the same plant can belong to different heat bands. Our ten are designed to hold a batch-to-batch CV at or below **7%.** A recipe is a promise that the same inputs give the same result; at 42% variation that promise is a lottery, and at 7% it is a recipe. Everything that follows — every dose weighed to a tenth of a gram, every finished Scoville figure we dare to print — rests on that 7% design target. The molecular engineering is not an abstraction here. It is the reason the honey is the same honey twice.

Reading the loudness: the index, worked

We read these peppers two ways on purpose, because they answer two different questions.

The **Scoville Heat Unit** answers *how much.* It is, at heart, a dilution count — proportional to the concentration of capsaicinoids a material holds. Pure capsaicin sits near 16,000,000 SHU; by the standard conversion that is 16 SHU for every part-per-million of capsaicinoid, so a material's capsaicinoid content in ppm is simply its SHU divided by sixteen. SHU is *linear in content.* Double the capsaicin, double the Scovilles. That linearity is exactly what makes it the right currency for dilution: blend and dilute in SHU, and the arithmetic is honest grade-school proportion.

The **Iyengar Flavor Index** answers *how loud.* It is anchored to the receptor, not the bottle:

$$\mathrm{IFI}(c) = 1{,}000 \times \mathrm{EC}_{50}(\text{reference}) \div \mathrm{EC}_{50}(\text{candidate})$$

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response — a low EC_{50} means the channel yields to a faint dose, a loud compound; a high EC_{50} means it takes a flood, a quiet one. Invert it against a fixed reference, scale by a thousand, and loudness becomes a score. Each axis has its own stated reference constant — IFI = 1,000 is defined by convention on that axis, the way Scoville's original scale was defined by his dilution-to-threshold protocol. Read the formula and the arithmetic falls out: a cultivar's EC_{50}, as a fraction of the reference's, is just 1,000 ÷ IFI, and its loudness relative to the reference is IFI ÷ 1,000.

CULTIVAR	PREDICTED IFI-H	EC_{50} (× REFERENCE)	LOUDNESS (× REFERENCE)
Ember · SS-0087	1,570	0.637	1.57×
Blaze · SS-0091	2,630	0.380	2.63×
Firestorm · SS-0094	3,490	0.287	3.49×
Solara · SS-0095	3,980	0.251	3.98×
Phoenix · SS-0096	4,200	0.238	4.20×

So Phoenix, at predicted IFI-H 4,200, is designed to reach half-activation of TRPV1 at 0.238 times the concentration the reference needs. That is 4.2× louder. That single ratio is the clean, worked meaning of the top of the ladder.

Now the part that governs every recipe in this chapter — the reason we keep both numbers and never collapse them into one. **SHU and IFI do not dilute the same way,** because content and loudness are not the same thing. Content is linear; loudness compresses. Watch what happens to Phoenix on its way into a jar of honey. In Scovilles — in content — the pod falls from 2,937,948 to about 5,900, a dilution of very nearly **498-fold.** In index — in loudness — it falls from IFI-H 4,200 to about 160, a drop of only **26-fold.** We cut the capsaicin five hundred times over, and the receptor only grows twenty-six times quieter.

That is not a paradox; it is the tongue. The relationship between content and perceived loudness is compressive, close to a power law — across this pair the loudness scales roughly as content to the 0.53 power, a square-root-ish squashing that is the same law that lets your eye read both a candle and noon. The finishing salt tells the same story in its own numbers: content falls 170-fold (2.72M to about 16,000 SHU) while loudness falls only about 20-fold (IFI-H 3,980 to about 200), an exponent near 0.58. The exact exponent is not a universal constant — it shifts with the capsaicinoid blend and the matrix the fire ends up in, which is precisely why loudness must be *read* on the finished thing and cannot be *computed* from content. And that gives us the working rule for the whole chapter:

Compute the SHU. Read the IFI. Content dilutes by arithmetic we can do on paper; loudness is a receptor reading we take on the finished condiment. One is the plan; the other is the verdict.

The arithmetic of a condiment

Everything a superhot condiment is rests on one equation, and it is gentle. Put a mass m_c of a cultivar rated S_c into a base, for a total mass m_{total}; because Scovilles are linear in content, the finished material rates

$$\text{finished SHU} = S_c \times \left(m_c \div m_{total} \right)$$

the cultivar's heat scaled by the fraction of the jar that is cultivar. That is the entire engine. Run it forward for the hot honey: 0.6 g of Phoenix at 2,937,948 predicted SHU into ~300 g of honey gives 2,937,948 × (0.6 ÷ 300) ≈ **5,900 SHU** in the jar. The teaspoon of vinegar and the pinch of salt add a little mass and nudge the figure a hair lower; at these ratios they are a rounding error.

But a cook rarely wants to *discover* the heat after the fact — a cook wants to *aim* at it. So invert the equation. To land a base of mass m_{total} at a chosen target S_{target}, weigh out

$$m_c = m_{total} \times \left(S_{target} \div S_c \right)$$

Design the honey this way and the recipe writes itself: a 300 g jar at a soft band-1 target near 5,900 SHU, from Phoenix at 2,937,948, needs m_c = 300 × (5,900 ÷ 2,937,948) = 0.603 g — six-tenths of a

gram, exactly the dose the recipe calls for. The recipe is not a tradition handed down; it is this division, done once.

CONDIMENT	CULTIVAR (PREDICTED SHU)	DOSE	BASE	MASS FRACTION	FINISHED SHU	FINISHED IFI-H	BAND
Hot honey	Phoenix · 2.94M	0.6 g	300 g honey	0.0020	~5,900	~160	1
Fermented sauce	Firestorm · 2.47M	120 g	~500 g blend	0.24	~560,000	~1,100	4
Finishing salt	Solara · SS-0095 · 2.72M	1.2 g	200 g salt	0.0060	~16,000	~200	1

The honey and the salt are pure single-source dilutions, and their finished Scovilles fall straight out of the equation above. The sauce is the honest exception: it is a *blend,* not a dilution — a small charge of band-5 Firestorm carried mostly by band-2 red chilies, then softened again by ferment and vinegar — so its ~560,000 SHU is an approximation across several contributors rather than a clean one-line division, and its loudness, as always, is read on the finished bottle.

It is worth translating these numbers out of the abstract and into milligrams, because milligrams are where the restraint becomes visible. Phoenix's dried flesh, at a predicted 2,937,948 SHU, is designed to carry about 18% capsaicinoids by weight — 2,937,948 ÷ 16 ≈ 183,600 ppm. The 0.6 g we weigh therefore carries about 110 mg of predicted capsaicinoid, and we spread those 110 mg through the whole 300 g jar. A generous 5 g drizzle of the finished honey — a real spoonful over cheese — delivers 110 mg × (5 ÷ 300) ≈ **1.8 mg** of capsaicinoid. (Read it the other way for the check: 5,900 SHU ÷ 16 ≈ 369 ppm, times 5 g, is the same 1.8 mg — the two routes agree, which is how you know the arithmetic is sound.) Less than two milligrams. Our highest-designed cultivar, with a predicted heat exceeding any known pepper, arriving at the table as a milligram and a half of warmth spread across a shared plate. A drop spread far enough is a kindness and not a dare, and now we can say exactly how far, and exactly how much.

Weighing, handling, and the error in a tenth of a gram

Band 5 is the one place in this book where carelessness can genuinely hurt you, so we work the way we actually work. The flagship pods are handled gloved, in good ventilation, with eye protection, and never near the face. Airborne capsaicin from grinding or blending a superhot will set a whole room coughing — the same lipophilic molecule that lingers on the tongue transfers from skin to eye with miserable ease, and once it is on the cornea there is no rinse that helps quickly, because water cannot lift an oil-soluble molecule off tissue. That last fact is the chemistry behind the kitchen's oldest remedy, too: milk and yogurt quiet a burn where water fails because the casein protein in dairy behaves as a surfactant, its hydrophobic regions sequestering the capsaicin and carrying it off the receptor. We keep dairy nearby when we work, not as superstition but as applied partition chemistry.



Solara handed off gloved under lab light — at a near-record pungency, every transfer of the dried pods is weighed and bare skin never touches them.

Weigh, do not estimate. And here the mathematics turns strict in a way worth making explicit, because at band-5 doses the limiting error is not the pepper — it is the scale. Suppose you weigh the honey's 0.6 g of Phoenix on a jeweler's scale that reads to a tenth of a gram. Its resolution is ±0.05 g, which on 0.6 g is a relative uncertainty of 8.3% — larger, on its own, than the cultivar's entire 7% batch variation. Combine them in quadrature, the way independent errors combine, $\sqrt{(8.3^2 + 7^2)} \approx 10.9\%$, and your 5,900-SHU honey is really 5,900 ± 640 SHU. Switch to a scale that reads to a hundredth of a gram and the weighing uncertainty collapses to ±0.005 g on 0.6 g, 0.83%, and the combined error falls to $\sqrt{(0.83^2 + 7^2)} \approx 7.0\%$ — now cultivar-limited, 5,900 ± 415 SHU, which is as good as the design allows.

DOSE	SCALE READS TO	WEIGHING ERROR	CULTIVAR CV	COMBINED	FINISHED SHU ± 1Σ
0.6 g	0.1 g	8.3%	7%	10.9%	5,900 ± 640
0.6 g	0.01 g	0.83%	7%	7.0%	5,900 ± 415
1.2 g	0.1 g	4.2%	7%	8.2%	16,000 ± 1,310

Two lessons live in that table. First, the finishing salt's larger 1.2 g dose is inherently easier to hit than the honey's 0.6 g, because the same ±0.05 g scale error is a smaller slice of a bigger number — when a recipe lets you, weigh more and dilute more, rather than weighing a vanishing dose precisely. Second, and this is the whole argument of the ten made concrete: had we been working a wild-type superhot at 42% CV, the honey would read 5,900 ± roughly 2,500 SHU from the plant alone — a spread wide enough to cross from a gentle warmth into a different band entirely, the same recipe landing sweet one week and harsh the next. The ≤7% CV is not a spec-sheet boast. It is the reason a band-5 condiment can be a recipe at all.

One-Drop Phoenix Hot Honey

Finished near IFI-H 160, from a cultivar designed for 4,200. A teaching recipe before it is anything else: watch the dilution turn our highest-designed cultivar into a soft, sweet glow.



Phoenix (SS-0096) in cross-section — the full-spectrum pepper, eight axes in one pod, destined not for a plate but for a jar where a tenth of a gram of it becomes a drizzle you can give with intention.

```
Makes ~300 ml · 20 minutes + infusing
0.6 g dried Phoenix (SS-0096, predicted IFI-H 4,200) — weighed on a
jeweler's scale, gloved
300 g mild honey
1 tbsp cider vinegar
Pinch of salt
```

1. Weigh the Phoenix precisely — this is a recipe where grams matter and eyeballing is forbidden. Gloves, ventilation, no touching the face.
2. Warm the honey gently to 60 °C, no hotter, or it loses its character. Stir in the crushed chili, vinegar, and salt.
3. Infuse 20 minutes at that low warmth, stirring, then strain for a clear honey or leave the flecks for a fiercer one.
4. The arithmetic, which is the whole point: 0.6 g of Phoenix, designed to carry about 2.94 million Scovilles, dispersed through 300 g of honey yields roughly **5,900 SHU** in the jar — a gentle band-1 warmth near **IFI-H 160,** a soft drizzle for cheese, for pizza, for anything fried.

There is a reason to leave the flecks, and it is physical. Capsaicin is lipophilic and barely soluble in a sugar-and-water matrix like honey, so it does not dissolve evenly through the jar — it stays bound to those suspended particles of flesh. The heat is therefore *placed* rather than uniform: a drizzle that catches a fleck reads hotter than one that misses it, and the warmth arrives slow, released as the particle lingers. Strain them out and the honey turns even and mild; leave them and you keep the slow, uneven, living quality of the fruit. And because the honey itself is warm chemistry too, remember that TRPV1 is heat-gated: drizzle this over a hot slice and the dish's real warmth and the capsaicin's phantom one add together on the channel, so the same honey reads hotter on pizza than on cold cheese.

Here is the finished receipt — and the lesson Phoenix's eight axes promised earlier, now paid out:

IFI PROFILE

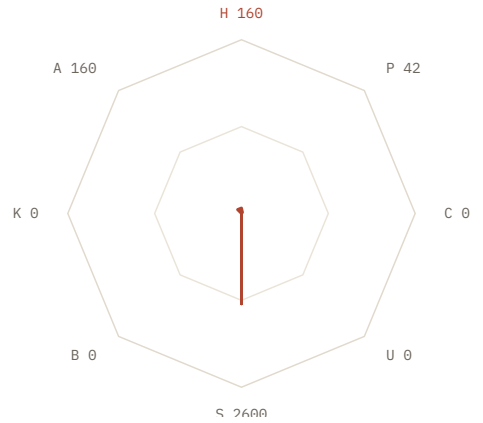

IFI-H	160
IFI-P	42
IFI-C	0
IFI-U	0
IFI-S	2,600
IFI-B	0
IFI-K	0
IFI-A	160

The heat is a whisper, 160. But notice what *else* survived the five-hundred-fold dilution and what did not. The sweetness reads loud — 2,600 — because that is the honey's own sugar speaking on TAS1R2·R3, lifted, not diluted. The pungency holds a faint 42 and the astringency a faint 160, the ghosts of Phoenix's own 252 and 1,400 on those axes plus the cider vinegar's edge, while cool, umami, bitter, and kokumi have all fallen below the threshold of report. The axes do not fade at one rate, because loudness is not linear — and that uneven fade is precisely why this honey tastes like a honey with a secret rather than like sweetened fire. Eight axes went into the jar; three came out loud enough to taste. That is design, not accident.

Firestorm Fermented Hot Sauce

Bottled near IFI-H 1,100, from a cultivar designed for 3,490. A proper lacto-fermented sauce, dosed into a bottle by the dash.



Firestorm in cross-section — six axes and a predicted 2.47 million Scovilles, charged in small into a ferment where lactic culture rounds the raw edge of the fire.

```
Makes ~500 ml · 20 minutes + 2 weeks ferment
120 g Firestorm pods (SS-0094, predicted IFI-H 3,490) — gloved — plus 300 g band-2 red chilies
20 g salt (2% of total vegetable weight)
4 cloves garlic
1 carrot, for sweetness and body
Filtered water to cover
120 ml vinegar, after fermenting
```

1. Weigh all the vegetables; weigh the salt at 2% of that total. Blitz the chilies, garlic, and carrot coarsely with the salt.
2. Pack into a clean jar, submerge under a weight, and ferment at room temperature two weeks, burping daily.
3. Blend the ferment smooth with the vinegar; the acid stabilizes and brightens. Bottle.
4. A small charge of band-5 cultivar carried by mostly band-2 chilies, then softened by ferment and vinegar, lands the bottle near **~560,000 SHU** and **IFI-H 1,100** — a serious but usable sauce, dosed by the dash, never the glug.

The fermentation is doing real chemical work, not just keeping time. The 2% salt selects for *Lactobacillus* and against spoilage organisms; as the culture consumes the sugars in the carrot and chilies it produces lactic acid, dropping the pH and preserving the mash, while the enzymatic breakdown of cell walls rounds the harsh, green, raw-capsaicin edge into something deeper and more rounded. The capsaicin content barely changes — fermentation does not meaningfully destroy capsaicinoids — but the *context* around the fire softens, which is why a fermented superhot sauce reads gentler and more complex than a fresh blend of the same Scovilles. This is the one recipe here built by blending rather than diluting, and its finished profile shows the company the heat now keeps:

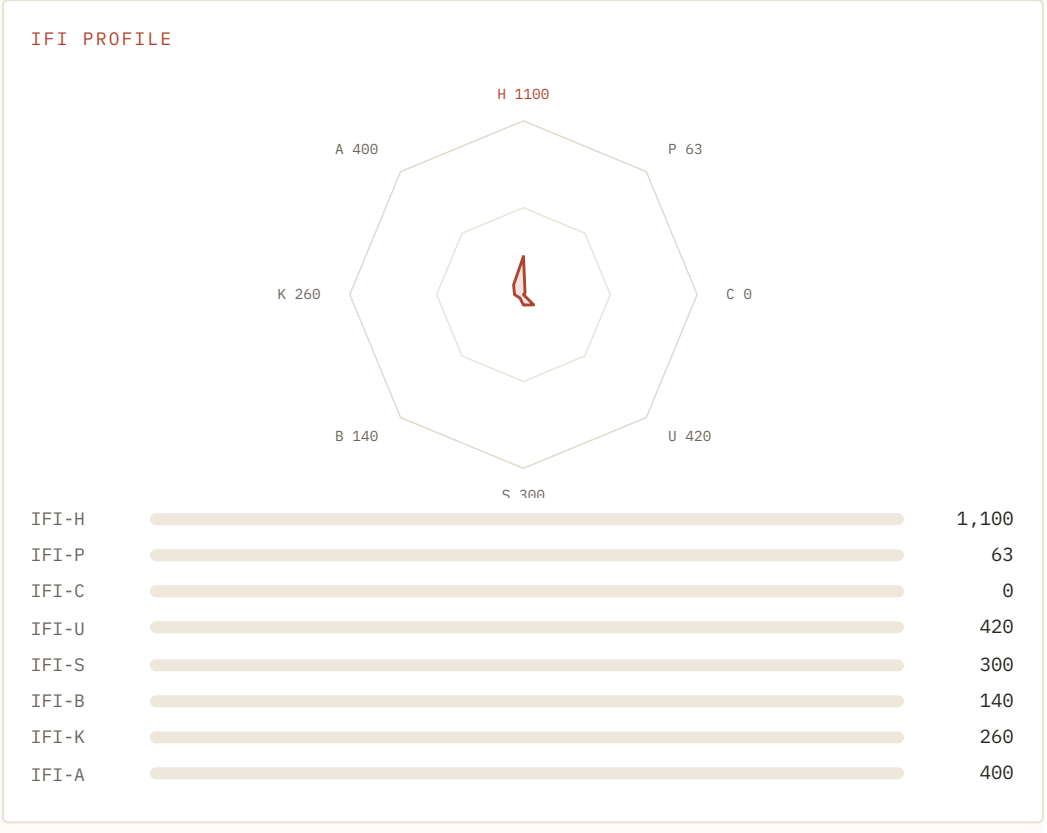

Band 4 on heat, yes — but read across the row. The ferment has raised a genuine umami (420) and a kokumi roundness (260) that no fresh purée would carry; the garlic lends its pungency and depth,

the carrot a quiet sweetness, the two weeks of culture a savory continuity that makes the sauce feel finished rather than merely hot. A dash of this is a composed thing, not a weapon.

Solara Finishing Salt

A near-record design at its gentlest possible use. A flaky salt you pinch over a finished plate, where a whisper of the second-highest-designed plant we grow becomes a single placed spark.

```
Makes ~200 g · 10 minutes + drying
1.2 g dried Solara (SS-0095, predicted IFI-H 3,980) — weighed, gloved
200 g flaky sea salt
Zest of 1 lemon (optional)
```

1. Grind the weighed chili to a fine powder in a dedicated mill — gloved, masked, well away from the face, because powdered superhots go airborne instantly.
2. Toss the powder thoroughly through the flaky salt until evenly colored, adding lemon zest if you want a brightness alongside the warmth.
3. If using zest, dry the mixture on a tray at the oven's lowest setting for 20 minutes; store airtight, labeled, out of reach of anyone who might mistake it for ordinary salt.
4. The arithmetic and its lesson: 1.2 g of a pepper designed for 2.72 million Scovilles through 200 g of salt makes a seasoning near **16,000 SHU** — about as hot as a mild hot sauce, if you ever ate it by the spoon, which you never would. A *pinch* over a finished plate delivers a clean, placed band-1 spark near **IFI-H 200,** exactly where you set it down.

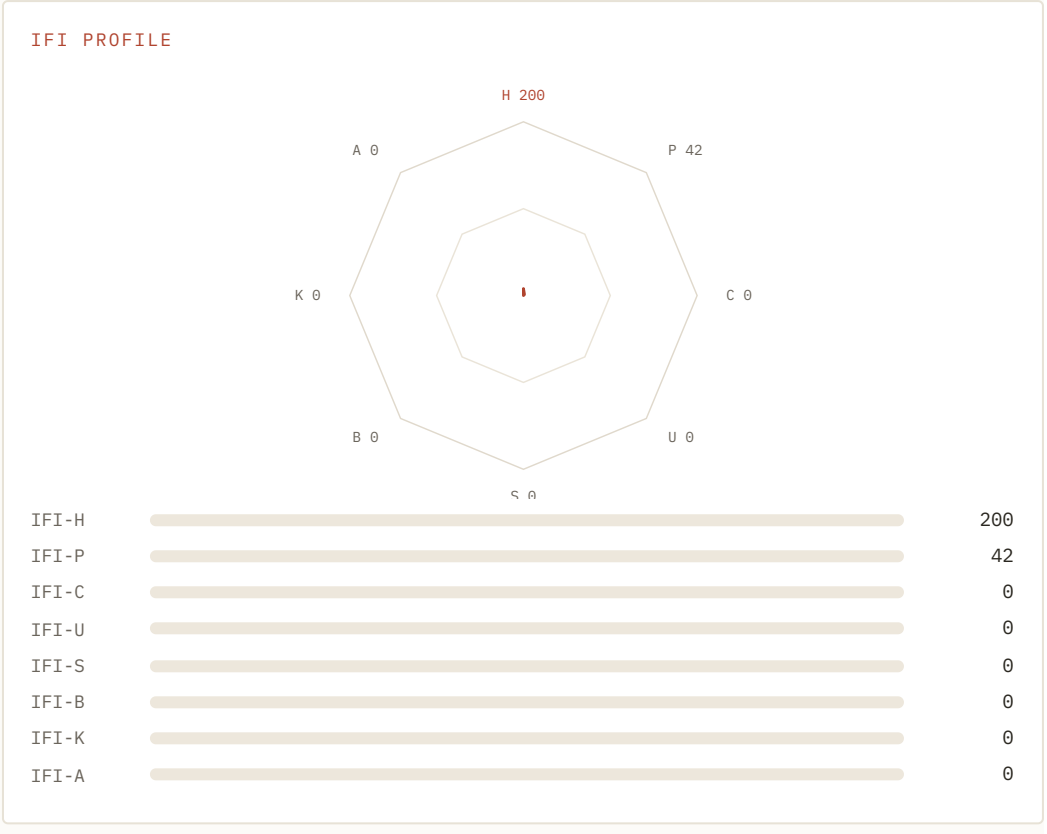

The salt is doing something the honey and the sauce cannot: delivering heat *dry,* as a surface dusting rather than a dispersion. There is no fat to dissolve the capsaicin evenly and no liquid to suspend it — the crystals simply carry the powder to the exact square centimeter of the plate where your fingers release the pinch, and the fire stays there. That is the whole virtue of the form. A liquid

sauce spreads its heat across a dish; this salt *places* it, a spark set down with the precision of a fingertip. Dose the pinch, never the pod — and note that the finished profile is heat and a breath of pungency and nothing else, because unlike Phoenix, Solara is a 6/8 pepper that leads overwhelmingly with the burn, and salt contributes no axis of its own but saltiness, which the Index does not score.

The discipline underneath

Three condiments, and a single discipline beneath all three: you dose the cultivar, you never serve it. Everything technical in this chapter is in service of that one sentence. We compute in Scovilles because content is linear and dilution is therefore honest arithmetic. We read the Index because loudness is what the tongue actually feels, and loudness compresses, so it must be read and not assumed. We weigh to a hundredth of a gram because at these doses the scale, not the pepper, is the error. We hold the cultivar's variation at or below 7% because a recipe is a promise, and a 42% promise is no promise at all. And we designed a fire hotter than any in nature not so that anyone could be hurt by it across a table, but so that its heat could be concentrated enough to be *bottled, diluted, and given* — a milligram and a half in a drizzle of honey, a dash of a composed sauce, a placed spark of salt — as exactly the pleasure the eater asked for and no more.

The failure modes are all failures of that discipline, and they are worth naming plainly. Over-dose, and the arithmetic you skipped turns a kindness into an assault. Under-disperse — a fleck left too concentrated, a powder tossed unevenly — and a single bite carries what a whole jar was meant to spread. Grind a superhot carelessly and the fire reaches you through the air before it ever reaches the plate. Trust a numbed palate over the scale and you will chase a heat your own tongue can no longer read. Every one of these is the same mistake: treating the pod as the serving. The pod is never the serving. The serving is what remains after the arithmetic has done its work.

These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. That is what the number was always for. The Reckoning is not the climax of the kitchen. The restraint is.



PART V OF V

The Craft of Consistency

heat is never eaten alone · the same warmth, every time · the fire we chose, on purpose

Chapter 17 · Pairing the Burn — *heat is never eaten alone.*

No one has ever eaten pure capsaicin for pleasure. We have held it in the lab — a fine white crystalline powder rated at sixteen million Scoville Heat Units, so pure it is almost an abstraction — and it is not food. It is a signal with nothing around it, a single channel shouting into an empty room. Heat, the kind you love, is always a *composition*: it arrives braided into sweetness, into coolness, into the sharp nasal lift of pungency and the round savory floor of umami, into a long slow tail of richness that keeps the whole sensation warm after the fire itself has gone. TRPV1 — the capsaicin channel threaded through the trigeminal nerve endings of the tongue, the same channel gated by heat above forty-three degrees and by acid — never fires into a silent mouth. It fires into a mouth already lit by its sibling receptors, and the way those signals meet, whether they reinforce or cancel or delay or lengthen one another, is the real grammar of the spicy kitchen. The most important work our family does is often not making heat louder. It is placing heat inside the other axes so that it sings instead of shouts.

That placement is not a matter of taste in the loose, unaccountable sense. It is arithmetic. Every axis on the Iyengar Flavor Index is a single number pinned to a single receptor, and when two numbers describe two receptors firing at once, there is a mathematics that predicts how the mouth will add them up — whether it will hear a chord or a collision. This chapter is that mathematics, worked out first at the receptor and then on the tongue. Heat is the loudest instrument in the orchestra; an orchestra it remains, and the scores are written in the language of channels interfering with one another on purpose.

The arithmetic of two receptors

Begin with the ruler, because the whole composition rests on it. Each axis is anchored to the half-maximal activation of a receptor through one formula:

> $\mathrm{IFI}(c) = 1{,}000 \times \mathrm{EC}_{50}(\mathrm{ref}) \div \mathrm{EC}_{50}(c)$

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response. A score is not how *much* of a compound a food holds; it is how *loudly* that compound speaks at the channel. The formula inverts cleanly, and the inversion is the thing worth memorizing. Rearranged, $\mathrm{EC}_{50}(c) = 1{,}000 \times \mathrm{EC}_{50}(\mathrm{ref}) \div \mathrm{IFI}$, so the ratio of a candidate's half-activation concentration to the reference's is simply $\mathrm{EC}_{50}(c)/\mathrm{EC}_{50}(\mathrm{ref}) = 1{,}000 \div \mathrm{IFI}$. Loudness is the inverse of the dose a stimulus needs to half-open its channel. When a sweet-heat construct carries a sweetness of 1,600, its sugar-free sweetener reaches half-activation of the TAS1R2·R3 receptor at $1{,}000 \div 1{,}600 = 0.625$ of the reference concentration — it speaks at under two-thirds the dose, which means it is louder than the anchor by precisely that factor. A heat of 800 reaches half-activation at

1,000 ÷ 800 = 1.25 times the reference dose: a true, substantial burn that is nonetheless, at the receptor, a measured step quieter than the reference fire. Two numbers, two inverted doses, and already the balance of a dish is visible on paper.

A single axis, taken alone, does not rise forever. Capsaicin on TRPV1 obeys a Hill relationship, the standard dose-response form,

$$R(C) = R_{max} \times C^{n} \div (EC_{50}^{n} + C^{n})$$

in which R is the response, C the concentration, and n the Hill coefficient — the cooperativity, greater than one for TRPV1, which sets how steeply the curve climbs through its half-activation point. At C equal to EC_{50} exactly half the channels are open; below it the response falls away sharply, and above it the curve bends over toward saturation, where adding more molecule buys almost no more signal. This flattening is the first reason composition exists at all, and it is visible in our own predicted range when our computational SHU estimate and the receptor measure are laid side by side.

Scoville Heat Units are a concentration measure: how many times a capsaicinoid extract must be diluted before the burn disappears, anchored so that pure capsaicin reads at sixteen million. IFI-H is a potency measure: the inverse EC_{50} at the channel. The two do not track perfectly, and where they diverge the whole argument for pairing is written in plain numbers.

CULTIVAR · SKU	PREDICTED SHU	IFI-H	EC_{50} RATIO (1000 ÷ IFI-H)	IFI-H PER 10^6 PREDICTED SHU
Ember · SS-0087	866,000	1,570	0.637	1,813
Cinder · SS-0088	1,050,000	2,290	0.437	2,181
Kindle · SS-0089	1,250,000	2,000	0.500	1,600
Phoenix · SS-0096	2,950,000	4,200	0.238	1,424

Read the third and fourth rows against each other and the lesson is stark. Our model predicts Kindle carries roughly nineteen percent more capsaicinoid than Cinder — a predicted 1.25 million SHU against 1.05 million — and yet it is *quieter* at the receptor, 2,000 against 2,290, a thirteen-percent lower IFI-H. More mass, less loudness. The reason is profile, not quantity: the balance between capsaicin and dihydrocapsaicin, set by the length and saturation of the acyl chain the plant builds its capsaicinoids from, shifts the effective EC_{50} and the shape of the burn. The final column, IFI-H earned per million predicted SHU, makes the saturation explicit — it peaks at Cinder's 2,181 and falls steadily to Phoenix's 1,424, so that at the very top of the design range each additional unit of concentration returns less and less receptor response. The Hill curve is bending over. You cannot make heat endlessly more interesting by cranking the concentration; past a point the channel

simply runs out of room. What you *can* do, without limit, is compose — set the burn against another receptor and let the mouth hear something new.

There are four grammars by which two axes combine, and every pairing in the kitchen is one of them. Two stimuli at *different* receptors that nonetheless overlap in central processing combine roughly **additively** in magnitude. Two stimuli at *neighboring* trigeminal channels that are cross-wired in the brainstem combine **subtractively**, one partly cancelling the other. A fast receptor set against a slow one combines **sequentially** in time, the quick signal arriving in front of the slow one. And an axis that alters not the peak of another but its rate of decay combines **durationally**, lengthening the tail while leaving the height untouched. Hold those four verbs — add, subtract, sequence, extend — and the rest of this chapter is their worked examples.

PARTNER AXIS	RECEPTOR	COMBINATION GRAMMAR	WHAT IT DOES TO THE BURN	EXAMPLE
Sweetness (S)	TAS1R2·R3	sequential	catches the burn, blooms it late and soft	`hot honey, mango-habanero`
Cool-Burn (C)	TRPM8	subtractive	lowers the perceived peak, lays frost over fire	`mint over curry`
Pungency (P)	TRPA1	additive top-note	a sharp nasal lift above the burn	`mustard, wasabi, ginger`
Umami (U)	mGluR1	floor-raising	grounds the fire in savory depth	`tomato heat, miso chili`
Kokumi (K)	CaSR	durational	lengthens and thickens the tail, no peak change	`garlic, aged ferment`
Astringency (A)	salivary PRPs	structural	a dry scaffold that holds the burn upright	`chili · dark chocolate`
Zero-Bitter (B)	TAS2R38	counter-tension	keeps sweet-heat from cloying	`chili-coffee, cocoa mole`

Swicy — heat and sweetness, in sequence



Strawberry — the frontier of sweet heat, where sugar is engineered to meet fire in a single fruit, and the TAS1R2·R3 sweetness receptor reshapes the TRPV1 burn into something that blooms late and lingers soft.

The dominant pairing of the modern kitchen is heat against sweetness, and its magic is entirely a matter of *kinetics* — of which receptor is fast and which is slow. Sweetness is read at TAS1R2·R3, a class-C G-protein-coupled receptor sitting at the surface of the taste cell in watery contact with the food. Its ligand binds, washes, and releases quickly: onset within about a second, a peak by two or three, and a clean fade over ten to fifteen as saliva carries the sugar away. Capsaicin could hardly be more different. It is lipophilic — it does not dissolve in saliva so much as partition into the lipid of the membrane to reach TRPV1 from within — and that partition is slow to load and slower to clear. The burn is late to arrive, five to ten seconds behind the first taste, climbs toward its peak over fifteen to thirty, and then releases over minutes as the molecule desorbs. Put the two on one axis of time and they do not fight; they *take turns*. The sugar occupies the front of the perception; the heat arrives a half-beat behind it, into a mouth already made generous, and blooms rounder and fades softer for having been caught. Expressed as representative time-intensity traces normalized to a hundred: at roughly one second sweetness reads near 90 and heat barely 5; by three seconds sweetness peaks at 100 and heat has woken to 20; by eight seconds sweetness has eased to 70 as heat rises through 60; by twenty seconds sweetness is down to 35 while heat tops out at 100; and by forty-five seconds sweetness is a trace at 10 and heat, still warm, holds near 75. The two signals interleave rather than cancel, and the result is a fire that feels given rather than inflicted.

This is heat composed rather than cranked — a swicy profile exactly as we design it, fire and sugar in deliberate balance:

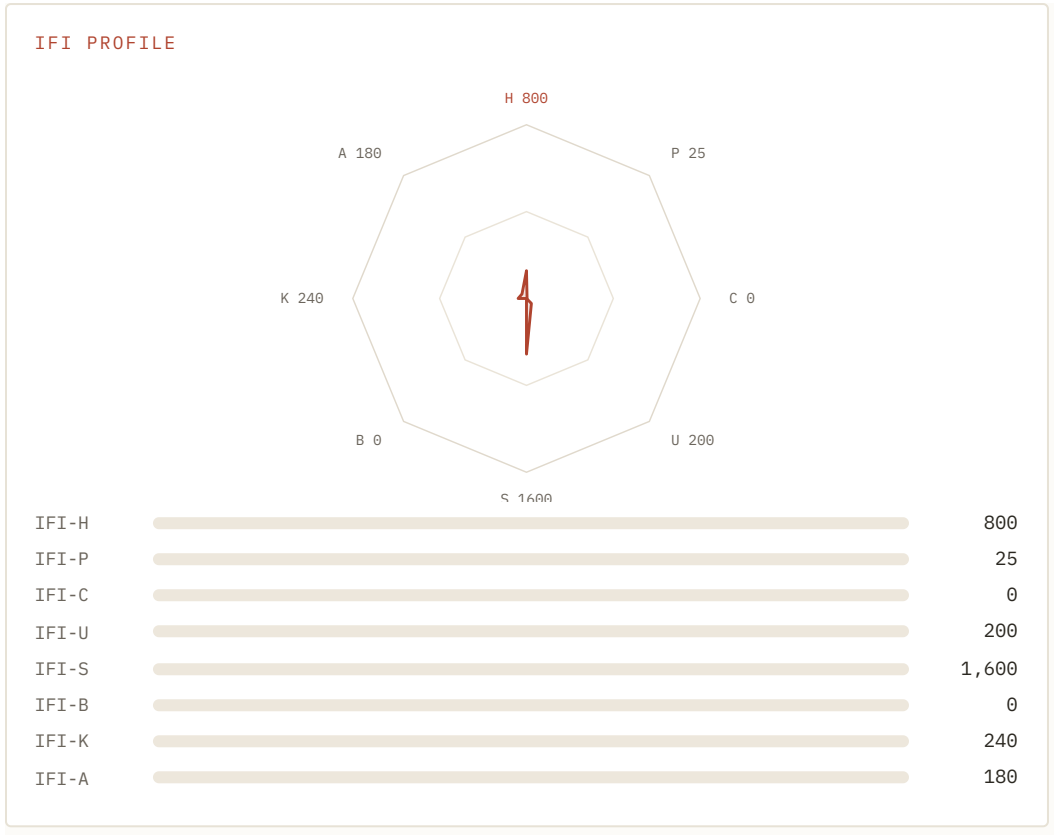

Run each number back through the inverted formula and the design becomes a table of doses — a reading of precisely how loud each receptor has been set.

AXIS	IFI	EC_{50} RATIO (1000 ÷ IFI)	ROLE IN THE COMPOSITION
Sweetness · S	1,600	0.625	loudest; the front that catches the burn
Heat · H	800	1.250	a true band-3 fire, placed behind the sugar
Kokumi · K	240	4.167	a faint lengthening of the tail
Umami · U	200	5.000	a savory floor under both
Astringency · A	180	5.556	a whisper of structure
Pungency · P	25	40.000	a nasal lift at the very top
Cool-Burn · C	0	—	silent
Zero-Bitter · B	0	—	silent

Read it and you can taste the architecture. The sweetener is the most potent stimulus in the construct, half-activating its receptor at 0.625 of the reference dose; the heat is real and substantial but set a step behind, half-activating at 1.25 times reference; and everything else is accent, an order of magnitude quieter, present only to round the edges. This is not a hot thing made tolerable by sugar. It is heat *orchestrated* — placed inside the other axes so the whole is more pleasurable than any part. There is, though, a window, and it is the chief failure mode of the style: push sweetness far enough above heat and the pairing stops *sequencing* the burn and starts *burying* it. Beyond roughly a two-and-a-half to one ratio of S to H the sugar no longer catches the fire; it smothers it, and the swicy collapses into candy with a memory of warmth. Composition lives inside that ratio

band, and the band is why the dual-axis libraries run to more than eight thousand constructs — each one a different point inside the window, held there on purpose.

The menthol subtraction — heat and cool-burn

Where sweetness sequences the burn, cool-burn *subtracts* it. TRPM8, the receptor for cold and for menthol, sits in the same trigeminal system as TRPV1, often in the very same nerve fibers, and the two are cross-wired in the brainstem such that a cooling signal is partly deducted from a hot one before the brain ever reports a number. Menthol does not touch capsaicin's chemistry; it does not cool the tissue any more than capsaicin heats it. It sends a competing message along a neighboring channel, and the central integration — the subtraction — does the rest. The simplest honest model of the perceived burn is linear in the cooling stimulus:

```
H′ = H − k_C × C
```

where H is the raw heat, C the cool-burn intensity, and k_C a cross-channel suppression coefficient fit from trained-panel data, typically landing somewhere near one-half. Take a dish built at H = 800 and lay a measured mint accent of C = 400 across it, with k_C = 0.5: the perceived heat falls to 800 − 0.5 × 400 = 600, a one-quarter reduction, while the dish gains a whole fresh dimension it did not have before. A cooling note over a fiery curry, a breath of menthol through a spicy drink — each buys relief without dulling a single thread of flavor, because it is not removing the capsaicin, only quieting the report of it.

Push the two together harder, in near-equal measure, and something stranger than relief appears: the mouth is told fire and frost at once, a contradiction it cannot resolve and keeps re-checking, re-lighting both signals in the attempt. That irresolvable simultaneity — genuinely hot and genuinely cold in the same instant — is the thrill at the center of the whole Cool-Burn axis, and it is only reachable by composition. The failure mode is its mirror image and worth naming: when k_C × C exceeds H, the subtraction runs past zero and the burn vanishes entirely, leaving a hollow, mentholated blank — the toothpaste note, cold and empty where the fire should be. The cool must be set *beneath* the heat it modulates, never over the top of it, or the dish loses the very thing it was built around.

Mála — three receptors, one sensation English cannot name



Ginger in cross-section — the sharp, volatile top-note of pungency, read not at the capsaicin channel but at TRPA1, the receptor that lifts a mála composition above its burn.

The most sophisticated heat composition in world cooking is the Sichuan *mála* — numbing-heat — and it is the clearest proof that heat is a composition problem, because no single receptor can make it. It is a chord of three. Capsaicin fires TRPV1 for the burn. Hydroxy-sanshool, the Sichuan-peppercorn molecule, fires a mix of touch-and-tingle channels and brushes the cool-burn system for the famous electric numbness — a vibration the mouth reads as a frequency rather than a taste. And a sharp pungency fires TRPA1, the mustard-and-wasabi channel, for a nasal top-note that lifts above both. Three receptors set in deliberate proportion produce a sensation that has no name in English because English never had a single-channel word for it. Quantified as a triad: the burn sits at a band heat near IFI-H 1,200, half-activating TRPV1 at 1,000 ÷ 1,200 = 0.83 of reference and arriving on capsaicin's slow lipophilic clock; the tingle rides the touch channels on a fast, flickering onset; and the pungency, at an IFI-P near 189 — half-activating TRPA1 at 5.29 times reference — opens almost instantly and clears almost as fast, a bright spike at the very front of the bite.

The numbness is not incidental, and here the composition does its most beautiful structural work. By partly deadening the trigeminal field, the sanshool raises the mouth's *tolerance* for the capsaicin riding alongside it — the numb tissue reports the same fire as less punishing — so a properly built mála oil can carry more heat than a straight chili oil and still read, start to finish, as pleasure. One axis makes room for another. That is composition as engineering, not garnish, and it is exactly what the triple-axis library, some 2,800 constructs, exists to build: multi-receptor experiences designed on purpose rather than stumbled into. A Sichuan chili oil assembled from the Index sets three dials at once, and the whole is tuned where no single dial could reach:

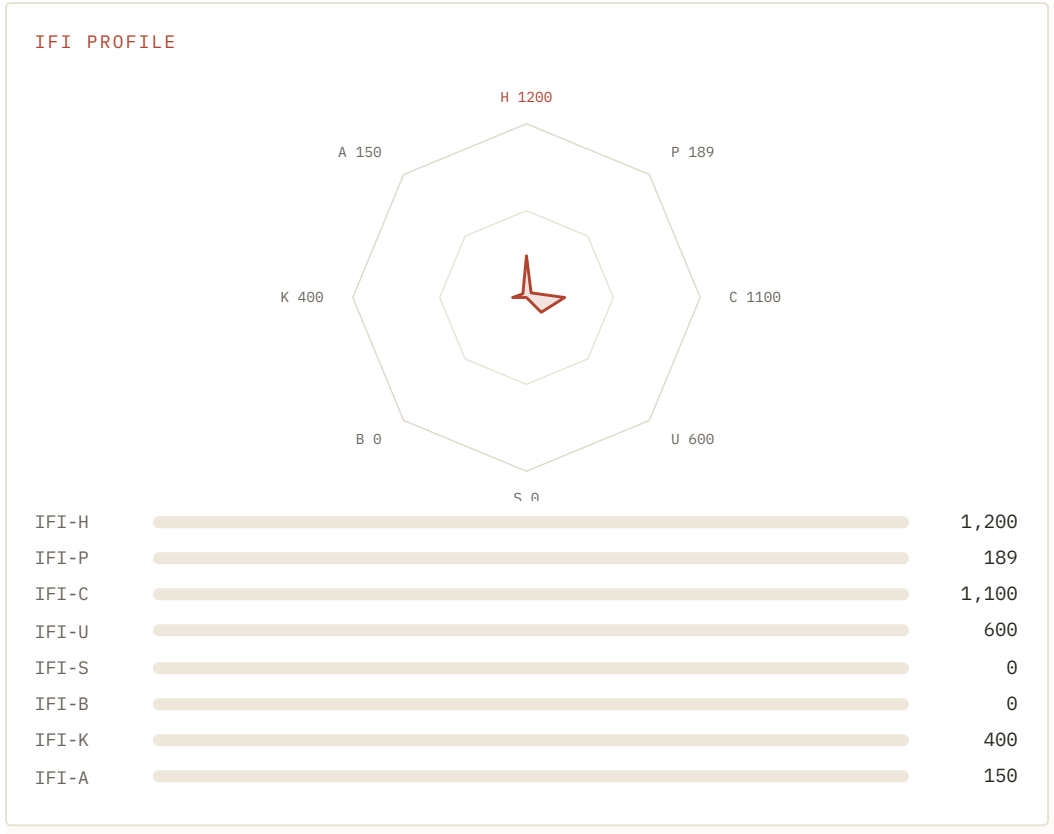

A band-3 heat for the burn, a strong pungency for the top-note, a strong cool-burn carrying the numb, and a savory floor of umami and kokumi from the oil and its ferments to hold the three bright signals over something warm and grounded — no sweetness, no bitterness, nothing to soften what is meant to stay electric. This is a plant composition in the fullest sense: everything in it is grown, pressed, and fermented, and the depth beneath the fire is built from the oil and the living culture, not borrowed from anything an animal made. The mouth receives a sensation designed to sit precisely where it sits, and the body receives a thing made with intention.

The axes that change time, not loudness — kokumi and umami

Two axes do something the others cannot: they alter the burn without ever touching its peak. Kokumi, read at the calcium-sensing receptor CaSR, adds no intensity at all. What it changes is *duration*. Model the decay of a burn as a simple exponential, $H(t) = H_{peak} \times \exp(-t/\tau)$, where τ is the time constant — the measure of how slowly the fire lets go. The total sensation, the full dose of experience the mouth receives, is the area under that curve, and the integral is clean: $AUC = H_{peak} \times \tau$. Kokumi lengthens τ while leaving H_{peak} exactly where it was. Lift τ from twelve seconds to twenty — a realistic potentiation from a well-built kokumi construct — and the area rises by $20 \div 12 = 1.67$, a sixty-seven-percent increase in total heat-experience with not one additional unit of peak. The half-life of the burn, $\tau \times \ln 2$, stretches from about 8.3 seconds to 13.9; the fire does not climb higher, it simply refuses to leave. This is why garlic, aged ferments, a long-cooked foundation make a spicy dish feel *deeper* without making it feel hotter — they are editing the tail, not the height.

Umami, read at mGluR1, works a third way again: it raises the floor the heat stands on. It grounds the fire in savory depth, so the burn reads not as a spike over emptiness but as the top note of something substantial — the tomato heat, the miso chili, the broth that makes a fierce bowl feel nourishing rather than merely aggressive. The edge

case for both axes is the same and instructive: an extended tail or a raised floor needs *something worth extending*. Kokumi laid under a thin, poorly made heat produces only a long chemical lingering, a tail that outstays its welcome because there was never enough there to lengthen. Duration and depth amplify what exists; they cannot invent it. The peak still has to be designed right first.

The grown-up pairings — astringency and bitterness



Cocoa — the astringency and controlled bitterness that give heat a dry, serious frame; tannins strip the saliva of its lubrication and the burn stands upright inside the structure they leave behind.

Two quieter pairings pull heat toward sophistication rather than fun, and both add structure where sweetness adds pleasure. Astringency is the strangest of all the axes because it is not, strictly, a taste at all — it has no receptor of its own. The tannins and larger polyphenols of tea, red wine, and dark chocolate bind the proline-rich proteins that lubricate saliva, cross-linking and precipitating them out of solution; the mouth, stripped of its slip, reports the resulting mechanical friction as a dry, puckering draw. IFI-A is therefore anchored not to a channel's half-activation but to the concentration at which that protein precipitation reaches half-maximal turbidity — the "receptor" is a precipitation assay, and the sensation is read by mechanoreceptors feeling the loss of lubrication rather than by any taste cell. What this does to heat is architectural. The dry tannic scaffold gives the burn a frame to stand inside, which is why a chili-laced dark chocolate or a peppery red-wine reduction feels grown-up where hot honey feels playful: the astringency holds the fire upright, structured, serious. Bitterness, read at the TAS2R family of receptors, does something adjacent from the other direction — a controlled counter-tension that keeps a sweet-heat from cloying and gives a composition a spine. A measured bitter is the logic of a chili-coffee rub, of cocoa and cayenne in a mole: not the loud, lovable pairings, but where heat goes when it wants to be taken seriously, composed against the two axes that add architecture rather than thrill, so that the fire reads as structure instead of a dare.

The cost of composition — error, limits, and honesty

Composition is only possible because the dials hold still between batches, and that is a quieter triumph than any single flavor. A designed construct lands its axes at a batch-to-batch coefficient of variation at or under seven percent, against the forty-two percent of a wild-type plant left to its own genetics. That difference is not cosmetic; it is the precondition for the entire discipline, and the

arithmetic shows why. A composition is defined not by any one axis but by the *ratio* between axes — the swicy *is* its S-to-H balance, the mála *is* its three-way proportion. For two independent axes, the uncertainty of their ratio propagates as the root-sum-square of their individual coefficients: $CV(H/S) \approx \sqrt{(CV_H^2 + CV_S^2)}$. At our seven-percent control that is $\sqrt{(0.07^2 + 0.07^2)} = 0.099$, a touch under ten percent — a balance that holds its character batch after batch. Run the same calculation on wild-type's forty-two percent and the ratio's variation is $\sqrt{(0.42^2 + 0.42^2)} = 0.594$, very nearly sixty percent. At sixty percent the swicy would lurch from "too hot" to "too sweet" every other batch; the composition would not exist as a repeatable thing at all. You cannot design a chord on instruments that retune themselves by sixty percent overnight. Low CV is not a quality boast — it is the thing that makes pairing possible, and it sets the design frontier too: to hold a composition's ratio to within five percent, each axis must be controlled to about 3.5 percent, since $5 \div \sqrt{2} \approx 3.54$. Tighter chords demand tighter dials.

Honesty requires naming what the model does not capture. The IFI sets the loudness of each receptor, but it reads each axis at a single dose and does not, by itself, predict the full nonlinearity of central integration — the brain does not always add and subtract as cleanly as the four grammars suggest, especially at high intensity, where capsaicin's own calcium-dependent desensitization of TRPV1 blunts the report of *subsequent* heat and can quietly flatten a composition built to climb. The numbers set the dial; a trained human panel still has to confirm the music, and the two together are the method — neither alone. And the mouth doing the tasting is not a constant. The receptors themselves vary from person to person: TAS2R38 runs in haplotypes that make some tasters far more sensitive to bitterness than others, and TRPV1 carries polymorphisms that shift an individual's baseline burn, so a single composition is read a little differently by every person who eats it. We design to the center of that distribution and we say so plainly. The Index is a ruler, not a guarantee; what it guarantees is that when we move a dial, it moves by the amount we intended, in the direction we intended, batch after batch — and on that reliability every pairing in this chapter is built. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Heat is never eaten alone, and the best heat is never designed alone either. It is sequenced behind sugar, subtracted beneath cool, lifted on pungency, grounded in savory, lengthened by richness, framed by tannin — placed, each time, inside the other axes so the whole is more than the fire. The burn is only ever the loudest voice at a table already full. Our work is to set the table well.

Chapter 18 · The Same Warmth, Every Time — *CV ≤7%, and why it is a kindness.*

We end where we began — with reliability — because it is the quiet foundation under every number in this book, and the clearest statement of what our family actually values. A heat you cannot reproduce is not a flavor; it is a wager. For ten thousand years spicy food was precisely that wager: the pod that ran hot, the curry that ambushed a guest, the "mild" that was anything but. The single most important thing our platform does is end the wager — not by making heat larger, but by making it *trustworthy*. This closing chapter is about the one number that sits beneath all the others, unglamorous and load-bearing, the number that decides whether every other number in these pages can be believed: the coefficient of variation, the measure of how stubbornly a thing refuses to hold still.

It is worth saying plainly, before the mathematics, what we are really defending. We are a family that has eaten with intention its whole life, and we think of a warmth arriving at a table the way you would think of anything you set in front of someone you love: it should be exactly what you promised, no more and no less. Consistency is not the dull part of the work that comes after the glamorous part. It is the part that makes the glamorous part *true*.

The arithmetic of a coefficient of variation

Begin with the figure that has shadowed heat since the first cultivated chili. Capsaicin output in a wild pepper varies — pod to pod, on a single plant, in a single field of a single cultivar — by more than **forty-two percent**. That number is a *coefficient of variation*: the standard deviation of a batch divided by its mean, written as a percentage,

```
CV = σ ⁄ μ × 100%.
```

The CV is the honest way to speak about spread, because it is dimensionless. A 42% CV says the same thing whether you are counting Scoville units, milligrams of capsaicinoid per gram, or a score on our Index — it is the width of the scatter expressed as a fraction of the thing itself. Against that wild 42%, across every engineered line we grow, we hold the coefficient of variation to **seven percent or below**: a better than sixfold tightening, 42⁄7 = 6.0, and the single most consequential ratio our family has ever computed.

Variance is invisible until it ruins something, so make it visible. Take a hot sauce formulated to a target of **IFI-H 800** — a confident, mid-ladder heat. The standard deviation of a batch is simply the CV times the mean, σ = CV × μ:

- Built from wild-type peppers at CV 42%: σ = 0.42 × 800 = **336**.
- Built from an engineered line at CV 7%: σ = 0.07 × 800 = **56**.

Those two sigmas are two different worlds, and the ranges they imply are what a customer actually tastes. For a roughly symmetric batch distribution, about 68% of bottles fall within ±1σ of the target and about 95% within ±2σ:

SOURCE	CV	Σ AT TARGET 800	±1Σ (≈68% OF BOTTLES)	±2Σ (≈95% OF BOTTLES)
Wild-type field	42%	336	464 – 1,136	128 – 1,472
Engineered line	7%	56	744 – 856	688 – 912

Read the wild-type row and feel the problem. One bottle in three lands outside 464–1,136 — a spread from a gentle warmth to nearly double the promised heat — and one bottle in twenty is hotter than 1,472 or so faint it barely registers. The eater never knows which bottle they are opening. Read the engineered row and the problem dissolves: ninety-five bottles in a hundred fall inside 688–912, a band so narrow that two of them set side by side cannot be told apart. The sixfold tightening of variance is not a laboratory nicety. It is the thing that turns a designed heat into a *shippable* one.



The controlled chamber where a cultivar is held to its number — fixed light, fixed temperature, fixed humidity, so that the only variable left is the one we engineered out. Constant conditions are half of a constant pepper; the stabilized enzyme is the other half.

What the tongue can actually tell apart

The deepest reason the coefficient of variation matters is not statistical but perceptual, and here the mathematics turns tender. Human senses do not read absolute differences; they read *ratios*. This is Weber's law, nearly two centuries old and still the most reliable rule in the study of sensation: the smallest difference a sense can detect, the just-noticeable difference or JND, is a roughly constant fraction of the stimulus already present,

> ΔI ⁄ I = k (the Weber fraction).

For the burn of capsaicin the perception literature puts that fraction high — somewhere in the neighborhood of fifteen to twenty-five percent, considerably coarser than the ear's or the eye's. Heat is a loud, slow, protective sense, built to warn rather than to measure, and it discriminates bluntly. Take the middle of that range, roughly twenty percent, as a working threshold: two heats must differ by about a fifth before a person can reliably tell them apart.

Now lay the two CVs against that threshold, and the word *kindness* in this chapter's title stops being sentiment and becomes arithmetic.

- An engineered line at CV 7%: ±1σ is 7%, ±2σ is 14%. Ninety-five percent of all batches sit within 14% of the target — *below* the roughly 20% a tongue can notice. The variation is not merely small. It is, to the person eating, **imperceptible**. The heat you asked for is the heat you get, not because the numbers are pretty on a chromatogram but because the difference has fallen beneath the floor of human sensation.
- A wild-type field at CV 42%: ±1σ alone is 42%, already more than twice the JND. Nearly every pod differs *noticeably* from its neighbor. The ambush is not bad luck. It is built into the plant.

Push the engineered line to its own extreme and the conclusion still holds. At CV 7% the three-sigma tail — the single batch in roughly a thousand that strays furthest — reaches ±21%, which only *now* grazes the lower edge of what a tongue can register, and that rarest outlier is precisely the one the release gate exists to catch before it ever leaves us. The wild type has no such floor: its *typical* pod already lives where the engineered line's one-in-a-thousand worst case merely visits once and is turned away.

This is the whole case in one line: we have not just made heat steadier, we have made its wobble smaller than the eater can feel. A heat that varies beneath perception is a heat that keeps a promise to the body, every single time, whether or not anyone is counting.

From receptor to number: the EC50 beneath the score

To hold a number steady you must first know exactly what the number measures, so follow the heat from molecule to score. Capsaicin binds TRPV1, a non-selective cation channel on the trigeminal nerve endings of the tongue — the same channel opened by temperatures above 43 °C and by acid. The molecule warms nothing; it lowers the channel's thermal threshold, so the nerve reports a fire that is not there, and calcium and sodium pour through the open pore. Upstream, in the pepper's placental tissue, the enzyme Pun1 (the AT3 acyltransferase) condenses vanillylamine with an 8-methyl branched-chain fatty acyl group to form capsaicin; the length and saturation of that acyl chain set the ratio of capsaicin to dihydrocapsaicin, and with it the *shape* of the burn — where it blooms, how long it stays.

How loudly that molecule speaks is captured by its dose–response curve, which for a cooperative channel like TRPV1 follows a Hill relation:

$$R / R_{max} = [A]^h / (EC_{50}^h + [A]^h).$$

Here R is the receptor current, [A] the agonist concentration, and h the Hill coefficient — about 2 for capsaicin at TRPV1, the signature of a channel that opens cooperatively. Set [A] = EC_{50} and the fraction collapses to exactly ½: **EC_{50} is the concentration that drives the receptor to half of its full response.** A potent extract reaches half-max at a low concentration; a faint one needs far more. That single number is the anchor of the entire Index:

$$IFI(candidate) = 1{,}000 \times EC_{50}(reference) / EC_{50}(candidate).$$

Invert the EC_{50} against a fixed reference and scale by a thousand, and loudness becomes a score. Fix the reference where half-maximal TRPV1 current arrives at EC_{50}(ref) = 50.0 mg·L^{-1} of lyophilized fruit equivalent — the concentration that *defines* IFI-H 1,000. Then every cultivar's EC_{50} falls out of its score as EC_{50}(candidate) = 50,000 ⁄ IFI mg·L^{-1}. Our model predicts **Phoenix**'s design would reach half-max TRPV1 current at 11.9 mg·L^{-1} — the number follows directly from its computed IFI-H score — and the Index reads

$$IFI\text{-}H = 1{,}000 \times 50.0 / 11.9 = \mathbf{4{,}202 \approx 4{,}200},$$

our predicted flagship heat — a plant we designed gene by gene, targeting a level no existing pepper approaches. The full ladder, read the same way, is a table of potencies — each EC_{50} expressed as a multiple of the reference, so the arithmetic needs no absolute anchor at all:

CULTIVAR	SKU	SHU	IFI-H	EC_{50} (× REFERENCE)	CAPSAICINOIDS (PPM)
Ember	SS-0087	866,000	1,570	0.637	54,100
Blaze	SS-0091	1,720,000	2,630	0.380	107,500
Firestorm	SS-0094	2,470,000	3,490	0.287	154,400
Solara	SS-0095	2,720,000	3,980	0.251	170,000
Phoenix	SS-0096	2,937,948	4,200	0.238	183,600

SHU, IFI scores, EC_{50} multiples, and capsaicinoid concentrations for engineered cultivars are computational predictions from our flux and receptor models, not certified measurements.

The EC_{50} column is just 1,000 ⁄ IFI-H: Phoenix's design reaches half-max at 0.238 times the reference concentration, meaning it is about 4.2 times as potent. And Phoenix carries that potency across not one axis but all eight at once — the first cultivar we built to register on the whole Index:

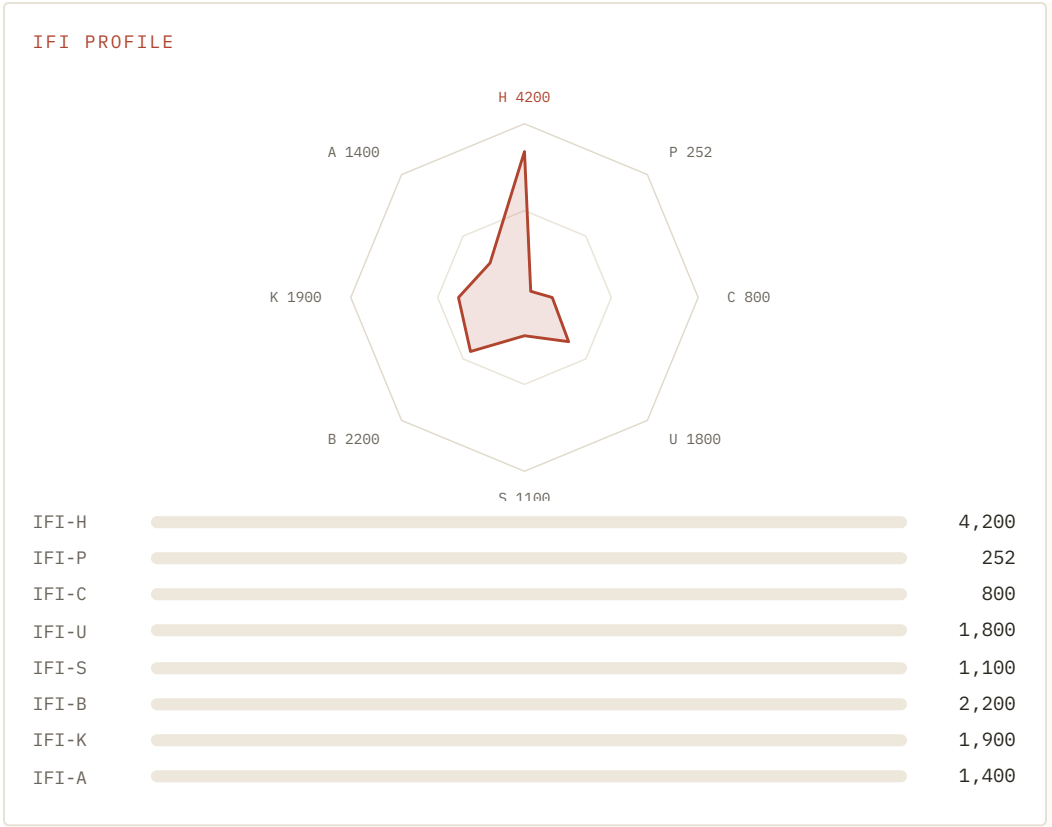

One honest subtlety deserves its moment, because it is the kind of thing a careful reader should catch us on. IFI and SHU do **not** convert by a single constant. From Ember to Phoenix the Scoville figure multiplies by 2,937,948 ⁄ 866,000 = 3.39, while the Index score multiplies by only 4,200 ⁄ 1,570 = 2.68. The Scoville scale is linear in capsaicinoid mass; the Index is anchored to the receptor, and as TRPV1 approaches saturation each additional milligram buys less additional *loudness*. SHU keeps climbing in a straight line while the Index gently compresses — which is exactly what you want from a number meant to track what a person feels rather than what a balance weighs.

Why a steadier enzyme makes a steadier pepper

The predicted record heat and the reliable heat are not two achievements. They are one piece of engineering seen from two sides, and the hinge is a protein. Every breeding program before ours chased heat by selecting the phenotype and praying the genotype followed. We inverted it: a protein language model scored every single-residue substitution in Pun1 by its predicted effect on catalytic efficiency and fold stability. Our flux model finds that metabolic control over the capsaicinoid rate sits upstream in the phenylpropanoid pathway — at 4CL and HCT — not at the terminal Pun1 step; Pun1 is the final synthase that assembles the product, and its stability governs *consistency* rather than output ceiling. Two substitutions emerged that no breeder had ever found, and when written in with CRISPR they do two things at once.

EDIT	ROLE	ΔΔG (KCAL·MOL⁻¹)	WHAT IT BUYS
S39L	N-terminal binding	−2.1	tighter substrate binding → higher turnover, more output
L345G	active-site cleft	−1.8	steric relief → faster catalysis *and* a stabler fold

A third substitution, C175S (ΔΔG −0.9, disulfide-misfold prevention), was identified during structural modelling as a potential stabiliser; it is retained as a structural annotation in our records but was not

included in the final shipped design.

The output and the steadiness come from different members of that pair, which is why one design delivers both. The binding edit, S39L, lifts how much capsaicin each enzyme makes. The folding edit, L345G, governs how *reliably* the enzyme stays in its active shape as a greenhouse breathes through its daily swing of temperature and light. That second effect is the one that tames variance, and it is pure thermodynamics. The ratio of folded (active) to misfolded (dead) enzyme follows a Boltzmann relation,

```
K_fold = exp(−ΔG_fold ⁄ RT),
```

where RT ≈ 0.59 kcal·mol^{-1} at 25 °C. The stabilising edit contributes ΔΔG = −1.8 kcal·mol^{-1}, which shifts that equilibrium by a factor of exp(1.8 ⁄ 0.59) = exp(3.05) ≈ **21** — more than twentyfold, toward the folded, working state. (Across both substitutions together, the double variant delivers a net ΔΔG of −3.9 kcal·mol^{-1}, a shift of roughly seven hundred-fold toward the native enzyme.) A protein held that far toward its native fold does not merely make *more*; it makes the *same* amount when the afternoon runs three degrees warm, because the fraction of enzyme in service barely moves. A stable enzyme in a flux-balanced pathway produces the same output batch after batch — and that, at the molecular root, is where a 42% wild-type scatter becomes a 7% engineered one.

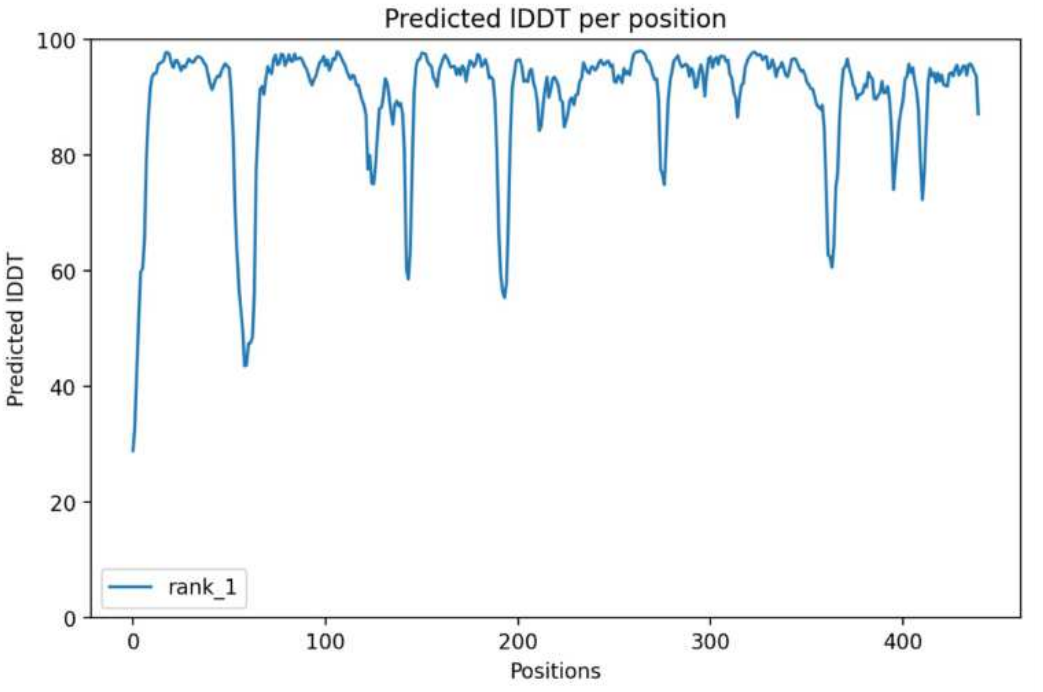

FIG. 18 — the Pun1/AT3 model, coloured by pLDDT confidence. The binding domain and active-site cleft resolve at high confidence; the two designed substitutions sit where the structure is most certain, which is why a prediction on a screen became a steadiness in the field.

How the error travels

A 7% coefficient of variation at the enzyme is only worth something if it survives the journey to the bottle, so follow the error as it propagates. Because the Index score is anchored to an inverse EC_{50}, and EC_{50} runs inversely with the capsaicinoid concentration C that the pepper delivers, IFI-H scales directly with C. For a direct proportion the relative error carries through one-to-one: CV(IFI) = CV(C). A 7% variation in what the plant produces becomes a 7%

variation in the delivered heat — no amplification, no hidden blow-up. The number on the label wobbles exactly as little as the pod does.

Formulation can do better still, through a piece of arithmetic that works quietly in our favor. Blend n independent batches of equal mean and the variances add while the mean does not, so the standard deviation of the blend falls as σ / √n and the coefficient of variation with it:

```
CV(blend) = CV(single) ⁄ √n.
```

Two independent lots: 7% / √2 = 4.9%. Four: 7% / 2 = 3.5%. A sauce drawn from several harvests is *steadier* than any one of them — a gift of scale that costs nothing. The honesty is in the word *independent*: the √n benefit is real only when the batches vary for unrelated reasons. A shared bad seed lot or a single heat wave correlates them, the covariance terms stop cancelling, and in the fully-correlated limit the blend is no steadier than one batch. We earn the averaging by deliberately decorrelating our lots — staggered plantings, independent chambers — rather than assuming it.

And here is a mercy that falls out of the receptor anchoring for free: when a sauce is built from a single cultivar and simply diluted to a target, the whole eight-axis profile holds its *shape*, because every compound in the pod rides down together. Dilute Phoenix's predicted profile until the heat lands at 800 — the factor is 800 / 4,200 = 0.190 — and each axis is just its predicted value times that same factor:

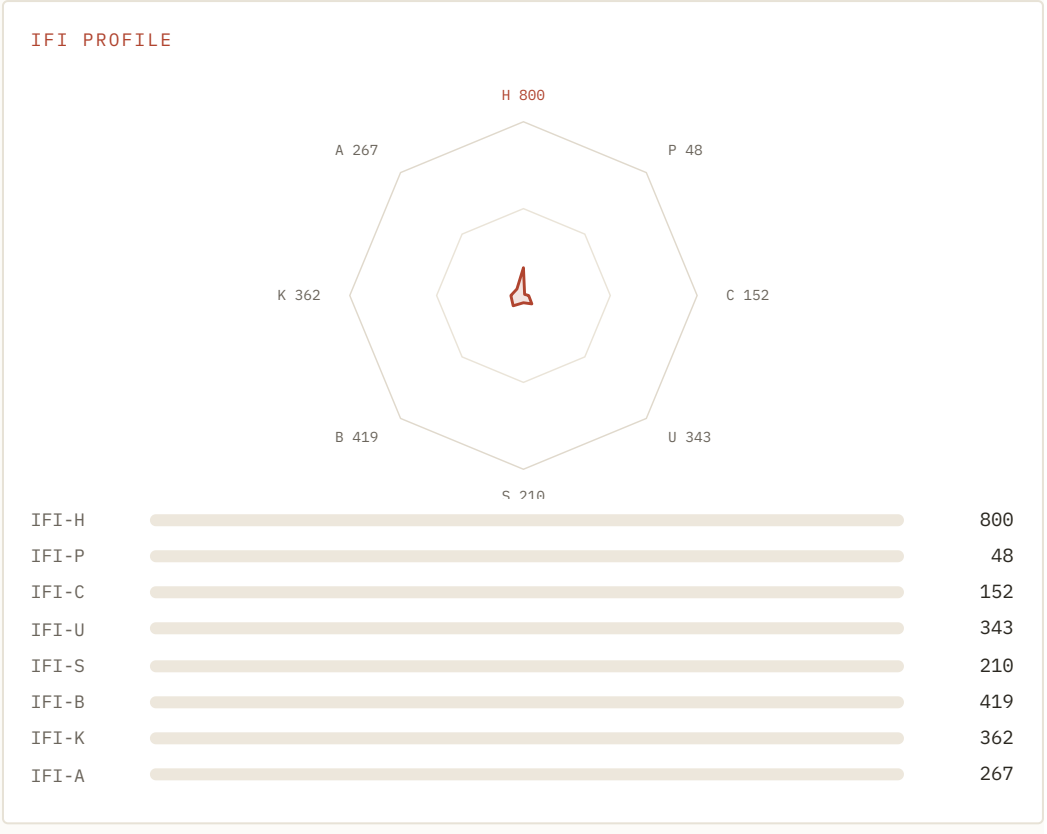

Every one of those figures is Phoenix's canonical predicted value multiplied by 0.190; nothing is invented, only diluted. The point is quiet but real: the 7% steadiness rides on *all eight axes equally*, so the whole flavor is as trustworthy as its heat — the richness and the pungency and the faint cool arrive as reliably as the fire.

The gate every batch must pass

None of this reaches a table unpoliced, and the policing is the least romantic, most important work we do. Each axis of the design loop has a measurement beneath it, which is what makes it a loop and not a wish: a profile is chosen as a target vector across the eight axes; a cultivar is selected because its predicted IFI-H matches the heat the profile asks for; the formulation is built by the weighted-average arithmetic of blending; and then every production batch is *verified*.

Verification runs on two instruments that must agree. First, HPLC resolves the capsaicinoids and reports their total concentration in parts per million, which converts to the familiar Scoville figure by the heat coefficient of capsaicin, SHU = 16 × C(ppm). It is worth doing one conversion by hand to feel the scale: our model predicts approximately 2,937,948 SHU for Phoenix, implying 2,937,948 ⁄ 16 = 183,622 ppm of capsaicinoids — about **18.4 grams per hundred grams of dry fruit**, a capsaicinoid density no unengineered pepper approaches, set against Ember's comparatively modest predicted 5.4%. The coefficient of 16 hides a small truth worth naming: it is itself a weighted blend of capsaicin's heat value, near 16.1, and dihydrocapsaicin's, near 15.2, and the balance between those two molecules is exactly what the acyl-chain length at Pun1 decides. By fixing the chain we fix the ratio, and by fixing the ratio we fix the conversion — so our Scoville figure and our Index score stay locked to one another batch to batch, rather than drifting apart as an uncontrolled capsaicin-to-dihydrocapsaicin mix would let them. Second, a sample of that same extract is run across the TRPV1 receptor assay as a spot-check, so that the molecule count and the receptor response confirm one another rather than either standing alone. A batch is released only if it lands inside its window; a batch that drifts is not shipped and softened with an apology, it is held, traced to its cause, and corrected — because the whole promise collapses the first time a bottle lies about its number. The window is not arbitrary. For the IFI-H 800 sauce we set release at the ±2σ band of a well-behaved line, 688 to 912: a lot assaying 905 ships, a lot assaying 918 is held — not because 918 would scald anyone, but because a bottle that reads 800 and delivers 918 has broken the single promise the number exists to keep.

One last piece of rigor keeps us honest about the headline itself. The 7% we observe is not all pepper; part of it is the ruler. Observed variance is the sum of the real production variance and the assay's own,

```
σ²(observed) = σ²(production) + σ²(assay),
```

so if the HPLC method carries a CV near 3%, the plant's true biological CV is $\sqrt{(7^2 - 3^2)} = \sqrt{40} \approx$ **6.3%**. We report the conservative *observed* ≤7% as the shippable guarantee precisely because it already folds in the noise of our own instruments — the living plant is a little steadier than the number we promise, never less.

The honest limits

A standard that asks the world to trust it owes the world its edges, so here are ours, plainly. The coefficient of variation is a ratio, and like any ratio it grows unstable as its denominator shrinks: for an axis a food barely expresses — a near-zero heat — the CV balloons toward meaninglessness, and there the honest report is the absolute spread, not the percentage. Our √n averaging across blended lots is real only for *independent* batches; a shared environmental shock or a single compromised seed line correlates them and erases the benefit, which is why we decorrelate deliberately rather than assume. The ±1σ and ±2σ bands in these pages assume a roughly symmetric distribution; capsaicinoid content is in truth mildly right-skewed, better described as log-normal, so the true upper tail runs a touch longer than a clean Gaussian would suggest, and we quote the conservative side. And the celebrated 42% of the wild type is itself a bundle — genetic scatter, environmental swing, and the plain fact that a pod near the crown of a plant differs from one at its foot — which is exactly why no amount of careful *selection* ever tamed it, and why it took re-engineering the enzyme to do so. We do not claim a variance of zero. We claim a variance we can measure, defend, and hold beneath the threshold of your own senses, and we publish the method so a skeptic can check us.

The same warmth, every time



The greenhouse at dusk, where the work comes to rest. Everything in this book — the receptor, the pathway, the enzyme steadied by two substitutions no one had found — exists so that a warmth asked for at a table arrives exactly as asked, and nourishes the person it reaches.

That seven percent is the quiet under everything romantic we have said about fire, because an Index score is only ever as good as the steadiness of the thing it names. A precise number on a wildly variable pepper is a precise reading of a moving target, useless to a cook; only a heat that holds still can honestly wear a number at all. Precision of measurement and precision of production are two halves of one promise — that when we tell you a cultivar is IFI-H 720, it is 720, not on average, not usually, but in the pod in your hand. We put our twenty-one peer-reviewed papers into the open, archived and licensed for anyone to read and build on, and filed our three patents in the daylight, because a company that asks the world

to trust a number owes the world its working: the EC_{50} assays, the reference anchoring, the scoring arithmetic are all published, so that the pod in the hand can settle every argument.

And underneath the reliability, the thing we have meant since the first page, said once and then let go. Heat, to us, is not a weapon and not a dare. It is one of the great pleasures a plant can offer a body, and our family has spent a decade learning to design it so that it can be *given* with intention — the exact warmth, meant for the exact person, grown in a plant tuned to carry the most it can and to nourish the one who eats it. That is the whole difference between a heat that happens to you and a heat that is *for* you. Every number in these pages — every receptor and pathway and mutation, every EC_{50} and sigma and release gate — exists in service of one small, enormous thing: a warmth you asked for, arriving exactly as asked, at a table that is glad you came. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. We are a family that reveres what it grows, and this is how we grow fire — cleanly, precisely, and always on the edible side of the line. The same warmth, every time.